The Complexity of the Dark Matter Sheet

Jens Oliver Stücker

Munich 2019

The Complexity of the Dark Matter Sheet

Jens Oliver Stücker

Dissertation
am Max-Planck-Institut für Astrophysik
der Ludwig–Maximilians–Universität
München

vorgelegt von
Jens Oliver Stücker
aus Hamm

München, den 18. April 2019

Erstgutachter: Simon D.M. White
Zweitgutachter: Volker Springel
Tag der mündlichen Prüfung: 24. Juni 2019

Contents

List of Figures

List of Tables

Zusammenfassung

Eine der wichtigsten offenen Fragen der modernen Physik ist: “Woraus besteht die Dunkle Materie?”. Astrophysikalische Ansätze, um mehr über die Natur der Dunklen Materie zu erfahren, erfordern eine zuverlässige Modellierung des *Fluids der Dunklen Materie* in kosmologischen Simulationen. In dieser Arbeit entwickeln wir neue mathematische Werkzeuge und numerische Simulationstechniken, um die Entwicklung der Dunklen Materie und ihrer Phasenraumverteilung zu verstehen und vorherzusagen.

Die Dunkle Materie hat eine unbekannte, aber kleine primordiale Geschwindigkeitsdispersion. Daher nimmt dieses Fluid im sechsdimensionalen Phasenraum effektiv eine dreidimensionale Lagrangesche Untermannigfaltigkeit ein - auch bekannt als das *dark matter sheet*. Dies hat viele Auswirkungen auf “feinkörnige” Merkmale im Dichtefeld, die Bildung von Strukturen und numerische Simulationsmöglichkeiten. Wir entwickeln neue Ansätze, um verschiedene großflächige Strukturtypen (Single-Stream-Regionen, Pfannkuchen, Filamente und Halos) aus der Perspektive des Phasenraums zu identifizieren und zu verstehen. Weiterhin entwickeln wir einen “Excursion-Set”-Ansatz für Single-Stream-Regionen, mit dem wir die Frage: *“‘Was ist die Mediandichte des Universums?”* beantworten und untersuchen, ob Single-Stream-Regionen einen zusammenhängenden Perkolationskomplex oder verschiedene isolierte Regionen bilden.

Weiterhin stellen wir einen “Sheet + Release”-Simulationsansatz vor, um zuverlässige kosmologische Simulationen von warmer dunkler Materie zu ermöglichen. Dies kombiniert ein “sheet”-basiertes Phasenraum-Interpolationsschema (Hahn & Angulo, 2016), das in Regionen mit niedriger Dichte (wie Single-Stream-Regionen, Pfannkuchen und Filamenten) verwendet wird, mit einem N-Körper Simulationsansatz in komplexen Regionen mit hoher Dichte (wie Halos). Dadurch werden Probleme in reinen “sheet”-Schemen (, die innerhalb von Halos zu komplex werden,) und N-Körpersimulationen (, die dazu neigen, in Regionen mit niedriger Dichte zu fragmentieren,) überwunden. Weiterhin stellen wir ein neuartiges Schema zur Berechnung von Kräften durch einen rekursiven Baum von Würfeln vor, das es ermöglicht, das “Sheet + Release” Schema bei hohen Kraftauflösungen zu verwenden. Dadurch präsentieren wir das erste Schema, das Simulationen warmer dunkler Materie ermöglicht, die von Regionen mit niedriger Dichte bis hin zu den sehr dichten und komplexen Zentren von Halos zuverlässig sind - und dies während zusätzlich viel feinkörnige Phasenrauminformation ausgerechnet wird. Als ersten Testfall wenden wir dieses Simulationsschema auf eine “Zoom-in”-Simulation eines kleinen Halos an (in einem Universum mit warmer dunkler Materie). Die Simulation enthält keine unphysikalischen Strukturen und die radiale Dichtestruktur des Halos konvergiert gut. Das Dichteprofil scheint konsitent mit einem NFW-Profil zu sein und unterscheidet sich nicht wesentlich von dem mit einer N-Körpersimulation berechneten.

In einem letzten Teil dieser Arbeit diskutieren wir die Möglichkeit der Rekonstruktion von warmen und heißen Phasenraumverteilungen lokal um kalte “Sheets” im Phasenraum. Dies könnte es ermöglichen, in Simulationen warmer dunkle Materie das Ausschmieren des Dichtefeldes aufgrund der primordialen Geschwindigkeitsdispersion explizit zu modellieren. Darüber hinaus könnte es verwendet werden, um den kosmischen Neutrino-Hintergrund mit einer relativ geringen Anzahl von erforderlichen Auflösungselementen zu simulieren.

Abstract

One of the most important open questions of modern physics is: *What is dark matter?* Astrophysical approaches to learn more about the nature of dark matter rely upon the numerical modelling of the *dark matter fluid* in cosmological simulations. In this thesis we develop new mathematical tools and numerical simulation techniques to understand and predict the evolution of the dark matter fluid and its phase space distribution.

Dark matter has an unknown, but small primordial velocity dispersion. Therefore the dark matter fluid effectively occupies a three dimensional Lagrangian submanifold in the six dimensional phase space - also know as the *dark matter sheet*. This has many implications for fine-grained features in the density field, the formation of structures and numerical simulation possibilities. We develop new tools to identify and understand different large-scale structures types (single-stream regions, pancakes, filaments and haloes) from a phase space perspective. Further, we develop an excursion set approach for single-stream regions which is used to answer the question *"what is the median density of the universe?"* and to investigate whether single-stream regions form a connected percolating complex or distinct isolated regions.

Further we introduce a "sheet + release" simulation approach to make reliable cosmological warm dark matter simulations possible. This combines a sheet-based phase space interpolation scheme (Hahn & Angulo, 2016) which is used in low density regions (like single-stream regions, pancakes and filaments) with a combined N-body + geodesic deviation equation simulation approach in complex high-density regions (like haloes). Thereby this overcomes problems in pure sheet-schemes (which become untraceably complex inside haloes) and N-body simulations (which tend to fragment in low-density regions). Further, we introduce a novel scheme for calculating forces from an oct-tree of cubes, which makes it possible to use the sheet + release scheme at high force resolution. Thereby we present the first scheme which makes possible warm dark matter simulations that are reliable from low-density regions up to the very dense and complex centers of haloes – and this while following a lot of fine-grained phase space information. As a first test case we apply this simulation scheme in a "zoom-in" simulation of a small warm dark matter halo. The simulation shows no artificial structures and the radial density structure of the halo converges well. The density profile seems to be consistent with an NFW-profile and does not differ significantly from an N-body simulation (which shows artificial fragments in the density field).

In a final part of this thesis we discuss the possibility of following warm and hot phase space distributions by reconstructing the phase space locally around cold sheets in phase space. This could allow to simulate warm dark matter while explicitly modelling the thermal smoothing due to the primordial velocity dispersion. Further, it might be used to simulate the cosmic neutrino background with a relatively small number of required resolution elements.

Chapter 1

Introduction

1.1 Motivation

The project of modern physics has lead us to understand and predict the behaviour of our universe to an extraordinary accuracy - from scales much smaller than nano-meters up to incredibly large scales of billions of light years. Despite this success, there are still open questions which are left to answer. One of the most intriguing questions is: *What is dark matter?*

A large variety of astrophysical observations show that something is missing if we try to interpret the effects of gravity on length scales of our galaxy or larger when only considering visible forms of matter. There are in principle two different approaches to this problem. The one class of solutions proposes that the general relativistic description of gravity is not correct on those large scales - therefore postulating a modified theory of gravity. The other solution is to postulate another unknown elementary particle which is commonly referred to as *dark matter* - an additional mass component which does not interact electromagnetically. While this debate is not completely over, a general consensus has formed that *dark matter* is the best explanation of the extra gravity. This is so, because dark matter can explain the clearest cosmological probe that we have - the cosmic microwave background - to an extraordinary accuracy while at the same time showing great consistency with late time observations of the universe. Therefore it gives a consistent quantitative picture of the evolution of the universe from very early stages, over the formation of structures up to the universe that we observe today. We will for the rest of this work assume the existence of dark matter as given. However, there has been no clear (non-gravitational) detection of dark matter yet, and it is one of the biggest open puzzles of modern physics what dark matter is.

The quest for the nature of dark matter can be pursued in two different ways.

(1) The particle physics approach pursues dark matter by detecting it by its interactions with other particles or by the products of its self annihilation. There is a large variety of postulated particles which could be dark matter and accordingly there is a large variety of experiments which try to detect different versions of dark matter. There has not been

any clear detection of dark matter yet, but experiments are continuously decreasing the parameter space of potential particles.

(2) The other approach is to learn about the nature of dark matter from its impact upon cosmological probes. For example, in the early years of the theory of dark matter, neutrinos were a natural candidate for dark matter. Neutrinos are a known elementary particle which in principle has the necessary properties. That is, they have a mass, they do not interact electromagnetically and they are produced in large quantities in the early universe. However, it was figured out quickly that neutrinos have a too large thermal velocity dispersion to be the main component of dark matter. In a neutrino dominated universe hardly any structures like our galaxy could form. Therefore neutrinos were excluded as (the main component of) dark matter (White et al., 1983). The cosmological approach to learn more about dark matter can still be applied nowadays. Different dark matter candidates can have properties which lead to differences in the observable universe. Such are its self-interaction, its quantum nature and its warmth. We will discuss this in more detail later.

It is for both approaches of fundamental importance that we are able to predict the evolution and the final state of our universe reliably – as a function of the underlying dark matter model. Cosmological simulations start from the very well probed early state of our universe and simulate the evolution of its ingredients until today. In the philosophy of approach (2) the simulated final state (and intermediate stages) can then be compared with observations of our actual universe to see what dark matter models can be valid. In the philosophy of approach (1) it is important for direct and indirect detection experiments to understand the local distribution of dark matter in our solar system and galaxy. The expected signal and, for example, its annual behavior can depend significantly on the details of this distribution.

In this thesis we make several contributions to the theoretical and numerical modelling of the dark matter content of our universe. These are based on a fine-grained view of the dark matter distribution in phase space which we will motivate below. However, we will explain it in much more detail throughout the second chapter of this thesis.

1.1.1 The fine-grained dark matter distribution

In Figure 1.1 we show the dark matter density field from a simulation of our universe visualized in two different ways. In the top panel we show a parallel-projection over a large volume with a depth of 40 Mpc/h. In a strongly simplifying manner we could say that this is a bit how our sky would look like if we could directly observe dark matter. Most of the mass clusters strongly in roundish objects - also known as halos. These halos assemble along some larger scale patterns known as the cosmic web. So far we cannot observe dark matter directly, but instead we observe the light from galaxies which form in the centers of halos. This “coarse-grained view” is the common perspective that is taken in the astrophysical community to understand the behaviour of the dark matter fluid.

However, we can get a different picture of what is happening if we make a thin slice through the dark matter density field, as can be seen in the bottom panel of Figure 1.1.

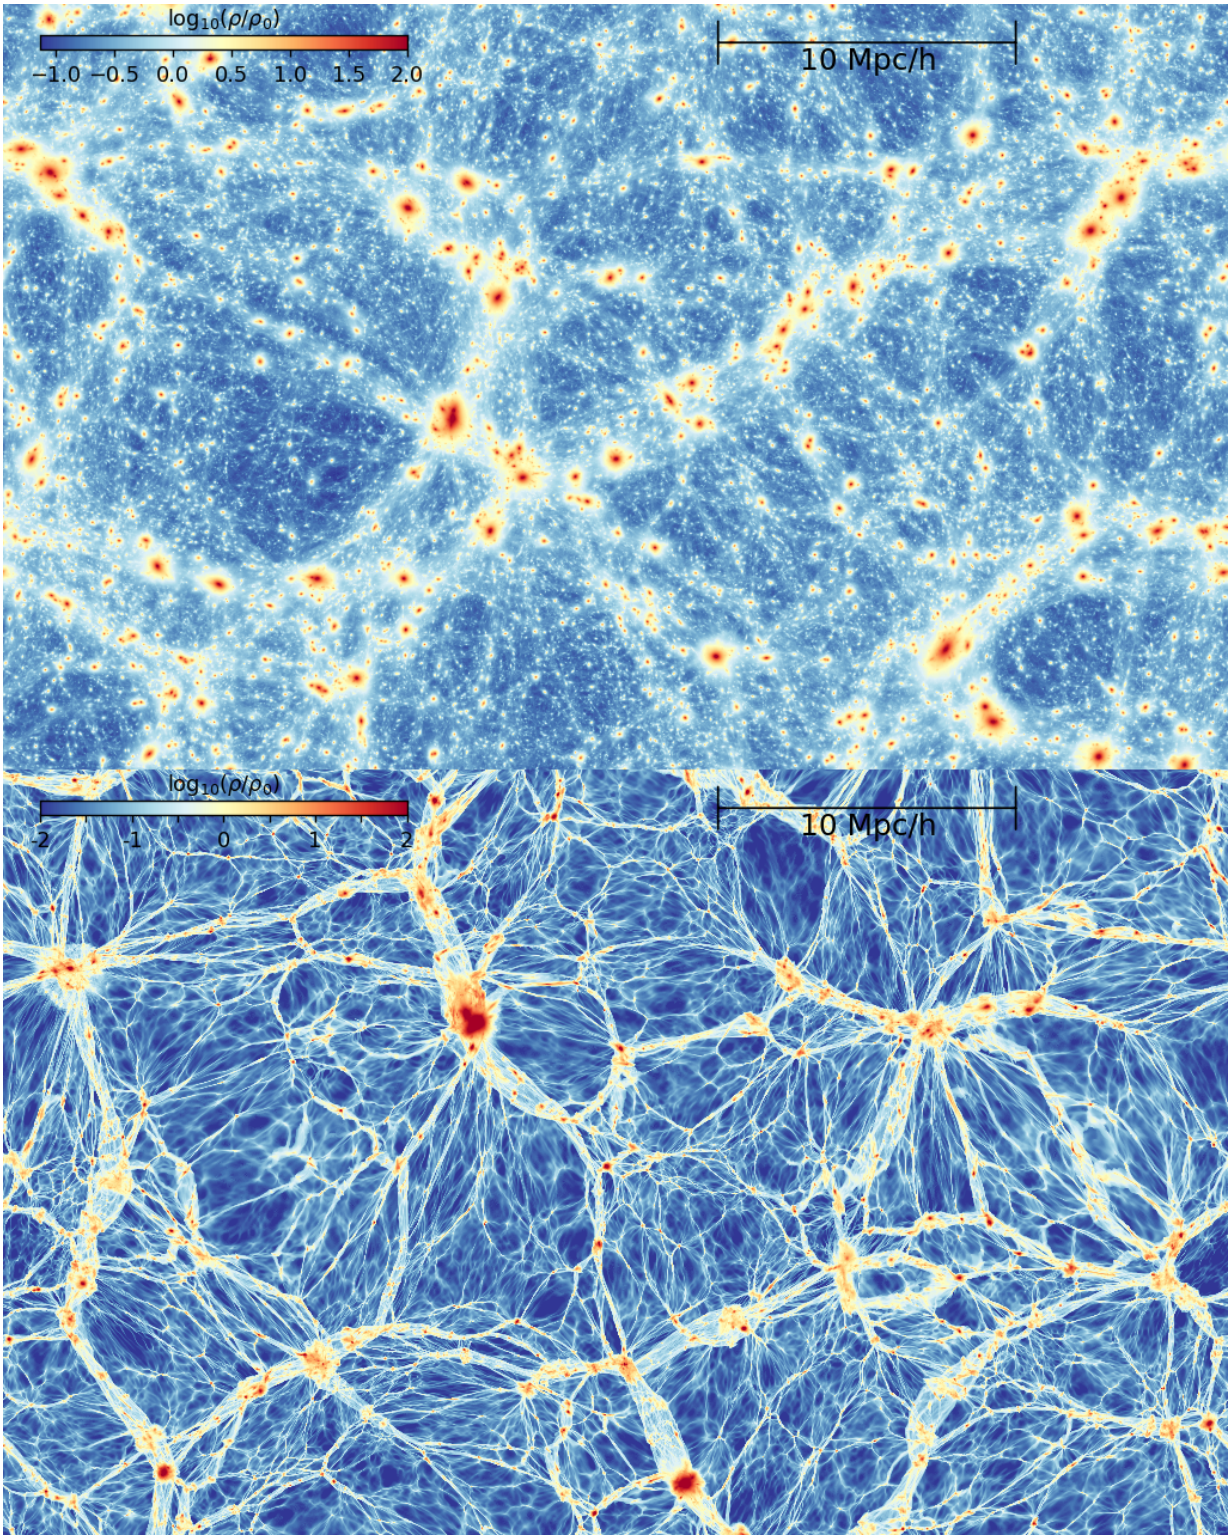

Figure 1.1: The simulated dark matter distribution of our universe visualized in two different ways. Top: showing a projection over a depth of 40 Mpc/h. One can identify halos as dominant structures in our universe (in terms of mass) and identify some large scale patterns (the cosmic web) on which they assemble. Bottom: a thin slice through the dark matter density field (width 20 kpc/h). A variety of fine-grained features become visible. Halos only make up a tiny fraction of the volume of the universe.

A thin slice enables us to see a lot of the fine-grained features of the dark matter fluid which get averaged out when making a thick volume projection. Among these fine-grained features are for example the diffuse single-stream regions which fill most of the volume, caustics which manifest as sharp edges in the density field, and a variety of collapsed structures of many different morphologies. I will refer to this way of interpreting the dark matter fluid as the "fine-grained view".

In most contexts the coarse-grained view is good enough to explain the behaviour of dark matter and possible observables. That is so, because we usually measure dark matter in a very coarse-grained sense where it is averaged over large volumes - convolved with the long-ranged gravitational interaction. However, in some contexts the fine-grained view can be of major importance. Observationally this can be important when asking the question for dark matter self-annihilation signals. Self-annihilation scales as the square of the local density and can potentially be enhanced by local fine-grained features of the dark matter distribution. This has for example been investigated by Vogelsberger & White (2011). While this is not a central topic of this thesis, we will still explain it in a bit more detail in section 2.5.

Further, the fine-grained view is of importance for many applications of theoretical and numerical nature. The fine-grained view can help to understand how the dark matter fluid behaves and how we can understand the kind of structures that we find in our universe. Further it can be used to come up with new simulation schemes. For example there are some doubts on the functionality of the standard-simulation technique of cosmological structure formation (Wang & White, 2007). N-body simulations tend to create artificial small scale structures when applied to initial conditions which do not contain any small-scale perturbations (like warm dark matter simulations). While this does not appear to be a problem for most cosmological simulations (which focus on a cold dark matter paradigm), it severely affects the prediction power of simulations which explore dark matter models with a small-scale cut-off like warm dark matter. New simulation techniques which consider the fine-grained structure of the dark matter fluid can help out in these cases and produce non-fragmenting dark matter simulations (Hahn & Angulo, 2016; Sousbie & Colombi, 2016). Amongst other projects, we will extend upon these here, to make them work reliably everywhere.

In this thesis we will take the perspective of this "fine-grained view". It offers a full narrative of what is happening and includes many features which easily get forgotten when coarse-graining too much. The fine-grained view is based on the idea that the dark matter fluid only occupies a three-dimensional submanifold in the six-dimensional phase space. This is so, because its primordial velocity dispersion is very low. This sub-manifold is often referred to as the dark matter sheet in phase space and it explains a lot of the features that we see in the bottom panel of Figure 1.1. Thus the title of this thesis - "The complexity of the dark matter sheet".

Among many smaller topic specific questions, in this thesis we answer the questions: *What can we learn from the fine-grained view of dark matter? What is the median density of the universe? How to make reliable simulations of warm dark matter? How well can we trust the N-body method? Can sheet-techniques be used to simulate the cosmic neutrino*

background?

1.1.2 An overview of the topics in this thesis

The topics in this thesis are based on the fine-grained view of the dark matter sheet, but are quite diverse. We will give a short overview here.

In the remainder of chapter 1 we will explain the physical context of this work in more detail. This includes a short review of our cosmological standard model - the ΛCDM model. This review is of course not an extensive introduction into cosmology. Instead it discusses only the minimum set of cosmological observations and physical notions that are necessary to understand this work.

In chapter 2 we will give an overview over the "fine-grained view" of the dark matter sheet and discuss a set of mathematical and numerical tools that are based on these notions. This includes a discussion of the Vlasov Poisson system and the tools of linear theory, the Zel'dovich approximation, N-body simulations, sheet-based simulation techniques and the geodesic deviation equation. Further we will develop a few additional tools that can be used to classify and understand cosmological structures. This chapter is quite technical, and not all introduced techniques are required for understanding the rest of this thesis. Therefore I actively encourage the reader to skip parts in this chapter if they appear too technical. The reader can revisit them later when the motivation becomes clearer from the application context of the later chapters.

Chapter 3 is about an interesting side project of my PhD which we have published in Stücker et al. (2018). It is about the question: *what is the median density of the universe?* While the mean density of our universe is very well known, its typical (or median) density is hardly known at all, since its value depends on dynamical processes on too many different length scales. We develop an excursion set formalism to calculate the typical density of our universe for the first time. Further, we try to understand properties of the diffuse low-density single-stream regions of our universe - for example whether they percolate.

In chapter 4 we present some work which is soon to be published. It is about the question: *How can we make reliable simulations of warm dark matter?* Therefore we extend upon the already existing Hahn & Angulo (2016) scheme and solve their problem of too much complexity in high-density regions by switching to an N-body approach in such critical regions. Further we develop a new scheme for calculating forces in cosmological simulations, by using an oct-tree of cubes. This allows us to use a high force resolution in combination with the sheet-based simulation approach. The combination of these schemes – which we refer to as "sheet + release" method – makes possible for the first time warm dark matter simulations that reliably resolve everything from low-density regions up to the very dense and complex centers of halos.

In chapter 5 we present one of the first such simulations. We make a zoom-in simulation of a small halo with mass $5 \cdot 10^9 M_\odot/h$ in a 3.5 keV warm dark matter cosmology. We simulate this halo with both – the N-body and the novel "sheet + release" simulation approach. We test whether they give consistent results for the structure of the halo or whether the structure is influenced by the artificial fragments in the N-body simulation.

Further, we have a short look at the phase space structure of this halo – though we note that some of these phase space quantities are not reliable yet.

In chapter 6 we discuss the possibility of simulating warm and hot distribution functions through usage of one or more Lagrangian submanifolds (sheets). We develop a formalism that can be used to reconstruct a warm phase space distribution around a cold sheet. As a proof of concept we apply it for a two-dimensional sheet to estimate the smoothing-effect of the primordial velocity distribution. We notice that this could be applied as a post-processing step to any cosmological simulations that follow a cold sheet. Further it could be used in simulations of warm dark matter to explicitly model the effect of the warm velocity distribution. As another application it could be used to reconstruct the high-velocity part of hot phase space distributions by using multiple sheets and reconstructing the phase space distribution locally around each of those. Thereby this might open up new possibilities to simulate the cosmic neutrino background.

1.2 The ΛCDM model

In this section we will briefly review the key observations of our universe that are important to understand the physical context of this thesis.

1.2.1 The expanding universe

It was discovered in the 1920's that the universe itself seems to be expanding. While there had already been some suspicions and minor evidence before, in Hubble (1929) Edwin Hubble put together for the first time compelling observational evidence for this. The basic observation is that other galaxies appear to be moving away from us (removing individual proper motions). The farther they are away, the faster the distance between us and them increases. This is summarized by the velocity-distance relation

$$v(d) = H_0 \cdot d \tag{1.1}$$

where d is the distance to an object, v the apparent velocity at which it is moving away, c the speed of light and H_0 the Hubble constant (today). This is also known as the Hubble-Lemaître law. The soon established interpretation of this law was that the universe itself is expanding (everywhere). Therefore the distance between objects at different locations is increasing over time. Assuming homogeneity and isotropy (on large scales) the expansion of the universe can easily be derived as a result of general relativity. It is described by the Friedman equations

$$H^2 = \left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho - \frac{kc^2}{a^2} + \frac{\Lambda c^2}{3} \tag{1.2}$$

$$\dot{H} + H^2 = \frac{\ddot{a}}{a} = -\frac{4\pi G}{3c^2}\left(\rho c^2 + 3p\right) + \frac{\Lambda c^3}{3} \tag{1.3}$$

where H is the time-dependent Hubble parameter, $a = 1/(1+z)$ is the scale factor (which describes the relative size of the universe), ρ is the total density, p is the total pressure, k is a parameter describing the curvature of the universe and Λ is the cosmological constant. It is not necessary to discuss the details of the Friedman equations here. As a simple explanation they relate the (time-dependent) expansion rate of the universe (described by H) to a budget of its contents. The contents of the universe are for the most part baryons, dark matter, radiation, neutrinos and (depending on the interpretation) dark energy which all enter the equations by their density and pressure budget. The cosmological constant Λ is a parameter which remains free in Einstein's theory of general relativity. It can either be interpreted as a property of gravity itself or alternatively as an additional ingredient of the universe – that is dark energy or vacuum energy – with an equation of state of the form $p = -\rho$ which leads to behaviour equivalent to a cosmological constant.

Nowadays the expansion rate of the universe and its history have been measured to a high accuracy. It seems that the universe is flat $k = 0$, we have a non-zero cosmological constant $\Lambda > 0$ and it has become clear that in today's energy budget the effect of the cosmological constant dominates the expansions rate – leading to an accelerated expansion of the universe.

1.2.2 The cosmic microwave background

That the relative size of the universe is increasing over time ($H > 0$), leads immediately to the recognition that it must have been much smaller in the past. However, the amount of mass in the universe is constant more or less. Therefore densities (and temperatures) must have been much higher at earlier times. If one extrapolates back far enough one finds that there must have been an epoch where the densities of baryons was high enough to maintain a local thermodynamic equilibrium with the radiation field. At this time baryons and photons were interacting frequently enough that they effectively formed one hot uniform fluid. However, densities diluted and the interactions decreased in frequency so that the universe effectively became transparent for photons. This happened when the universe was roughly a thousand times smaller than today ($z \sim 1100$) and we can directly observe the photons from this time. These are known as the *Cosmic Microwave Background* (CMB).

Quite surprisingly the CMB has roughly the same temperature of $T_{\mathrm{CMB}} \sim 2.7$ K in every direction. (The temperature scales like a^{-1} to earlier times.) This is surprising, because if one calculates the evolution of the scale factor $a(t)$ using only the known ingredients of the universe (matter, neutrinos, radiation, dark energy) one has to come to the conclusion that most regions of the observable universe were never in causal contact. Therefore there would be no mechanism which could establish something so close to a thermal equilibrium across the whole universe. So why should the universe have a uniform temperature then? This question can be answered by the theory of inflation which puts the universe into causal contact at very early times on scales which are much larger than the observable universe. For sake of brevity, we will not discuss this in more detail here, but instead take the initial state of the universe as given.

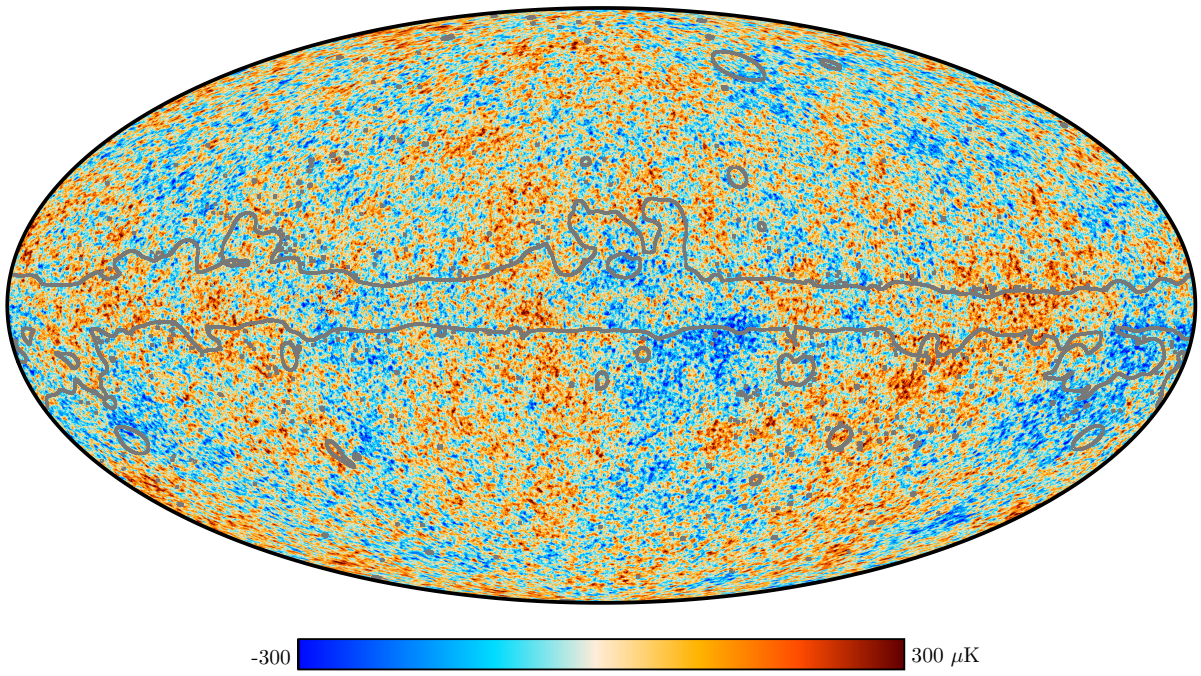

Figure 1.2: The temperature fluctuations of the Cosmic Microwave Background as observed by the Planck satellite. The CMB has an almost uniform temperature of 2.7 K with small temperature fluctuations on top of that which can be seen in this image. The grey lines delineate the areas of the CMB that have actually been reliably measured from those which have been masked out because they are contaminated by fore-ground structures. Full image credit goes to the Planck collaboration and the European Space Agency (Planck Collaboration et al., 2018).

While the CMB has roughly a constant temperature, there are small temperature fluctuations of the order $10^{-5} \cdot T_{\mathrm{CMB}} \sim 300\mu\mathrm{K}$. These have been measured to an extraordinary accuracy – by the Planck satellite as shown in Figure 1.2.

In the current best understanding of our universe – the ΛCDM model – these temperature fluctuations share a common origin with the same perturbations that have grown into the cosmic web, halos and galaxies nowadays. The physical picture here can briefly be described as follows. The universe starts at some very early time with some set of density perturbations which can be well described by a two parameter power spectrum $P(k) = A_s k^{n_s}$ with some amplitude A_s and a slope n_s. It is not entirely clear what is their source, but a very good explanation is given by the theory of inflation, where they are quantum fluctuations which are blown up to cosmic scales by a period of accelerated expansion. These density fluctuations are present in all species in the early universe, but evolve differently over time depending on the properties of each species. For baryons and photons these cannot grow freely, but instead get suppressed – by a different amount depending on how long they have been able to interact (which is scale dependent). At the time of the CMB these have relative fluctuations of the order 10^{-5}. Note that if baryons were the only matter component of the universe, these fluctuations would have grown by roughly a factor of $1/a_{\mathrm{CMB}} \sim 10^3$ till today – so they would still be of order 10^{-2} – thereby resembling an almost homogeneous universe (with no galaxies at all).

However, the CMB (and a lot of other observational evidence) suggests that there is an additional component of *dark matter* which co-evolves with the other components. While the baryonic perturbations are dampened down to the amplitude that we can observe in the CMB, the perturbations in the dark matter can grow relatively freely. At the point of decoupling they are several orders of magnitude larger than the baryonic perturbations. After the baryons stop interacting with the photon field they can fall into the potential wells that are provided by the dark matter. These over-densities then grow later into structures and form galaxies. (This process of structure formation is modelled in cosmological simulations and will be discussed in this thesis in great detail.)

The physics in the early universe can be described very accurately by perturbation theory. By assuming the additional component of dark matter (and a cosmological constant for the late time evolution) the statistics of the cosmic microwave background can be explained to an extraordinary accuracy. We show the temperature power spectrum from Planck Collaboration et al. (2018) in the top panel of Figure 1.3. The Figure also shows the best-fit according to the ΛCDM model. Further panels show different versions of power-spectra and cross-power spectra. We will not discuss all these observables here, but instead just want to point out that the spectrum of the CMB is fit to an extraordinary accuracy by the ΛCDM model. According to this our universe can be described accurately by a set of 6 parameters. We list those in table 1.1. Note that one can argue that there are only 5 really meaningful parameters. The parameter τ is not a fundamental property of our universe, but just a fit parameter which is left free, because the relevant physics of reionization is not constrained well enough. On the other hand one could argue that Ω_Λ should actually be a free parameter. The presented fit assumes that the universe is flat so that this parameter is given automatically as $\Omega_\Lambda = 1 - (\Omega_{\mathrm{dm}} + \Omega_b)$. Throughout this thesis

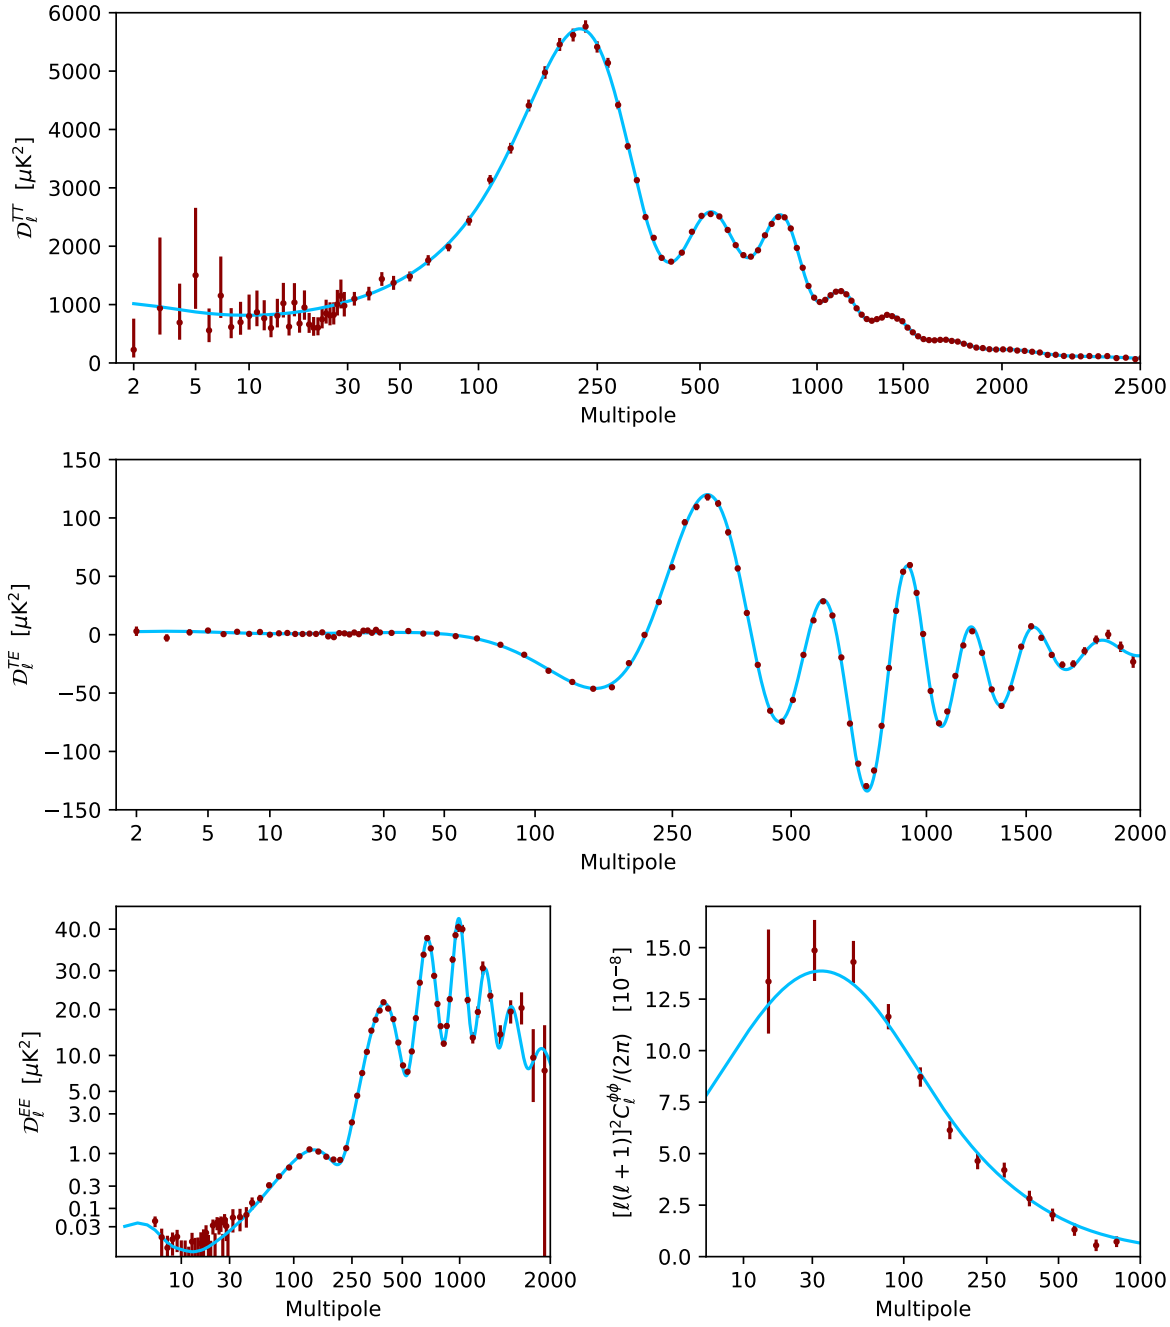

Figure 1.3: Top: the angular power spectrum of temperature fluctuations of the CMB as observed by the Planck satellite. The other panels show various other spectra and cross-spectra. (Second panel: cross spectrum of temperature and polarization, bottom left panel: the E-mode of polarization and bottom right panel: the lensing potential.) These spectra are described by the ΛCDM model (blue line, best fit) to an extraordinary accuracy. The CMB is the most clear cosmological probe for our standard model of cosmology. Full image credit goes to the Planck collaboration and the European Space Agency (Planck Collaboration et al., 2018).

Parameter	Value	Describes
Ω_{dm}	0.257	dark Matter density
Ω_b	0.048	baryon density
H_0	67.9 km/s/Mpc	Hubble constant
A_s	$2.14 \cdot 10^{-9}$	amplitude of the power spectrum
n_s	0.968	slope of the power spectrum
τ	0.05	optical depth to reionization
Ω_Λ	0.695	cosmological constant / dark energy density

Table 1.1: The main parameters of the ΛCDM model. The ΛCDM model is described by a set of 6 parameters which describe fundamental properties of our universe (except for τ). The parameter Ω_Λ is listed separately, because it is given by the other parameters if the universe is assumed to be flat $\Omega_\Lambda = 1 - (\Omega_{\mathrm{dm}} + \Omega_b)$. These are the parameters from Planck Collaboration et al. (2016) which we use in our modelling in this thesis. However, slightly updated values are available in Planck Collaboration et al. (2018).

we assume the ΛCDM model and the parameters from table 1.1 as given properties of the universe.

There are a large set of other cosmological probes from later times which confirm this picture, but arguably the CMB is the clearest cosmological probe: our universe seems to have an additional component of dark matter which has no interactions with light, but can mostly be traced by its gravitational influence upon the CMB, the formation of structures and the expansion of the universe. It is one of the most important quests of modern physics to find out: *What is dark matter?* There are a large number of well motivated candidates for dark matter, but so far none of them has been detected.

Note that the CMB is historically not the first evidence of dark matter. This was rather the missing mass in galaxy clusters (Zwicky, 1933) and the rotation curves of galaxies (Rubin & Ford, 1970). However, from today's perspective, the CMB is the evidence which is clearest to interpret and suffers least from prediction inaccuracies. To my knowledge there is no serious way of explaining the CMB without dark matter or at least something which behaves equivalently to a dark matter component. It is a bit of a pity that alternative theories of gravity which try to provide an alternative to dark matter always focus on giving an alternative explanation for late time probes (like galaxy rotation curves), but hardly ever try to explain the CMB which is arguably the cleanest test laboratory of our universe.

1.3 The nature of dark matter

It is not yet clear what dark matter is. From theoretical particle-physics there exists a whole zoo of well motivated candidates. Such are, for example, weakly interacting massive particles (WIMPs), axions and sterile neutrinos. Each of these candidates has different properties and is being searched for by different direct (and/or indirect) detection experiments. However, it is not only possible to learn more about the nature of dark matter by

detecting it directly, but also by observing its impact on structures in our universe.

There are a few such properties of dark matter which can have an impact on the formation of structures in the universe. These are its self-interaction, its quantum nature and its warmth.

Most dark matter candidates have some degree of self-interaction. However, this self-interaction must be much smaller than the one of baryons. An upper boundary of the self-interaction cross-section can for example be constrained from observations of the Bullet cluster (e.g. Robertson et al., 2017). The Bullet cluster shows two galaxy clusters which have passed through each other. The gas shows collisions, but most of their matter content shows a collisionless behaviour. So far dark matter seems to be consistent with a perfectly collisionless particle. However, if dark matter has a significant self-interaction, and the observational evidence and the numerical predictions of the relevant scales become better, this might change in the future.

Further dark matter probably has a quantum mechanical description. While we mostly do not consider quantum effects relevant on scales which are much larger than a couple of atoms in the case of baryons, we have no knowledge of the scales on which these become relevant for dark matter. For example, if dark matter would be an ultra light axion with mass of order $m \sim 10^{-22}$eV it would have a Compton wavelength $h/(mc)$ of 0.4 pc (which is the lower-most limit to its position uncertainty). Its position uncertainty could then have relevant effects on observable scales by creating large-scale quantum phenomena. Such scenarios have been proposed as possible solutions to the cusp-core problem[1]. Whether this is a likely scenario or not - the quantum nature is part of the parameter space of possible dark matter candidates that can be constrained from large-scale structure observations.

Finally, dark matter has a warmth. This refers not to the temperature, but rather to the primordial velocity dispersion of dark matter.

1.3.1 The warmth of dark matter

The effect of the warmth of dark matter is very significant for all dark matter models. The primordial velocity dispersion smooths out the formation of structures on small scales. For a "cold" dark matter (CDM) candidate like a WIMP this means that there cannot be any structures with masses much lower than earth-masses, for so called "warm" dark matter models (WDM) this means that no structures can be formed below for example 10^7 solar masses (depending on the model) and for neutrinos which are referred to as "hot" dark matter (HDM) there could not even be formed anything like our Milky Way. As an example we show a comparison between the simulated density fields of dark matter models of different warmth in Figure 1.4. Note: both of the shown dark matter models are already excluded by observations, but the point here is purely qualitative.

[1]The cusp-core problem is an inconsistency between the observed central density structures of dwarf galaxies (which appear cored) and the predicted ones (which are cuspy - that means steeper). However, this problem could originate from uncertainties in the predictions for such small haloes (due to baryonic processes) or uncertainties in the interpretation of the observations.

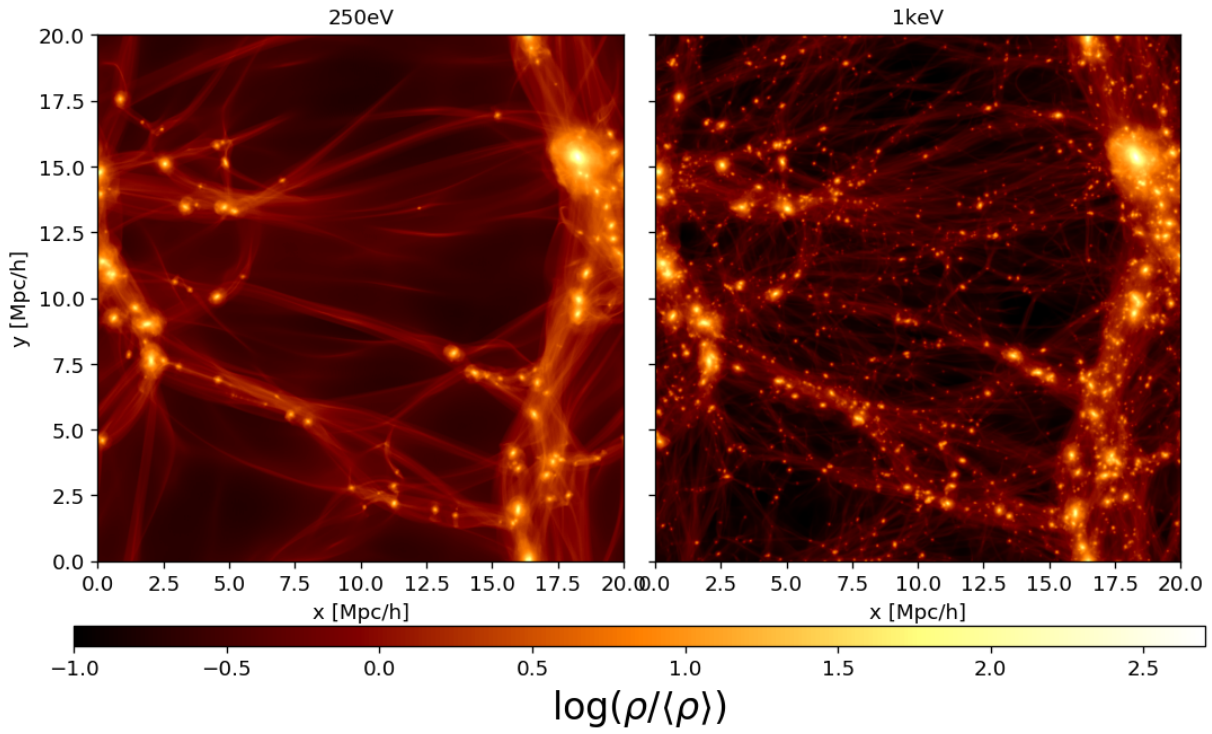

Figure 1.4: Projected density fields (width = 20Mpc/h) for simulations of dark matter with two different warmths. The left case uses a dark matter model with very high velocity dispersion which smoothes out structures on small scales. The right panel shows a colder (but still warm) dark matter model. The number of small mass objects is strongly related to the warmth of dark matter. For example, neutrinos have been excluded as the main component of dark matter, because they are too hot to produce small scale structures.

While neutrinos have been excluded as the dark matter that we are searching for (White et al., 1983), cold dark matter and warm dark matter models are still viable. When it is spoken of the Λ-cold dark matter model, it is typically meant that so far everything seems to be consistent with the limit of a vanishing velocity dispersion. However, there is a large parameter space of warm dark matter models which is not excluded.

Since, we have major parts of this thesis dedicated to the numerical modelling of warm dark matter, we want to point out one possible way to measure the warmth of dark matter. This is strongly connected to the question *what is the mass of the smallest dark matter halos?* One promising way of searching for the smallest dark matter halos is to detect perturbations in strongly lensed gravitational systems. As an illustration of this method we show in Figure 1.5 the system SDSS J120602.09+514229.5 from Vegetti et al. (2010). A blue background galaxy is lensed through a system which is located somewhere in-between the observer and the actual position of the galaxy. This creates several images of that galaxy and a strongly distorted arc which appears like a part of an Einstein ring. The effect of gravitational lensing is caused by the mass distribution in between which effectively focuses the light of the background galaxy to the observer. As such it is very sensitive to the mass distribution of the "lens". In this system this effect is particular pronounced around the galaxy G4 which creates a strong distortion to the lensing arc. When the system and the lensing effect is modeled properly this can be used to estimate the mass of G4.

While this effect is caused here by an anyways visible galaxy, it can also be caused by otherwise invisible halos. If one observes a large enough number of different lens systems, one would expect to measure some number of such perturbations due to small mass halos. If such systems exist, a lot of warm dark matter models can be ruled out. If they don't exist, that can be explained by the warmth of dark matter. Of course this requires a lot of detailed modelling and it is a quite active field of research. However, it is of big importance for this modelling that accurate predictions are available of the (sub-)halo mass function of warm dark matter universes. So far it was a big problem to determine these numerically, since N-body simulations tend to create artificial halos in warm dark matter cosmologies. Our new scheme for warm dark matter simulations can solve this problem and make such measurements possible. By the time I write this thesis, we have not made measurements of the (sub-)halo mass function yet, but this will be a future project.

Note that there are also other cosmological probes which can be used to constrain the warmth of dark matter. Very famous here are the Lyman-α-forest which probes the early universe where the difference between warm dark matter models and cold dark matter models is more significant. Further the so called "small scale crisis of dark matter" is an often stated reason to consider alternative dark matter models. However, we do not want to give a comprehensive overview here. Instead we picked the lensing case to point out why it is important that we model warm dark matter correctly in simulations.

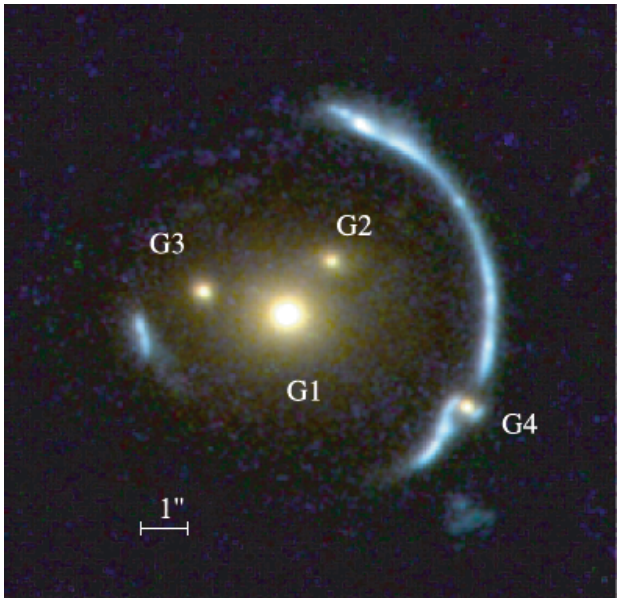

Figure 1.5: The system SDSS J120602.09+514229.5 which shows a strong gravitational lens. A blue background galaxy is lensed by a foreground system (including the galaxies G1-G4) to several images and an Einstein arc. The arc has an additional perturbation from the galaxy G4. Such a perturbation allows for an accurate mass determination of the object G4. In principle this method can be used to find (or not find) very small halos - maybe down to $10^5 M_\odot$ - thereby constraining dark matter. Image credit goes to Vegetti et al. (2010), Lin et al. (2009) and NASA/HST. This image is from the publication Vegetti et al. (2010) who created it from Hubble Space Telescope data. The system was originally discovered by Lin et al. (2009).

Chapter 2

The dark matter sheet in phase space

As already mentioned in the introduction, the "fine-grained view" of the dark matter sheet in phase space gives rise to a whole new set of simulation and analysis techniques. In this chapter we will give an introduction to these and further discuss some additional tools that we have developed in this context.

Therefore we will explain the basics of numerical structure formation with a special emphasis on the fine-grained phase space distribution of dark matter in sections 2.1 to 2.4. A summarizing illustration of this can be found in Figure 2.1. The physics and the necessary mathematical descriptions will be given in greater detail in the corresponding sections: section 2.1 describes the state of the early universe, section 2.2 the physical evolution equations, section 2.3 the N-body method and section 2.4 sheet-based simulation schemes.

Further we will describe in section 2.5 the geodesic deviation equation which provides another tool to follow the evolution of the fine-grained phase space distribution. Finally, in section 2.6 we will discuss several features in the dark matter distribution which arise from the dark matter sheet. Further, we develop some new tools here to identify regions with different phase-space properties.

Many of the concepts in this chapter appear again in later chapters. The general idea is that they are introduced here in a conceptual and mathematical way. Later they are used in a quantitative application context. This chapter contains a lot of mathematical details. They serve the purpose of a mathematical completeness, but they are often not necessary to understand the general ideas and concepts. We actively encourage the reader to skip sections which appear too technical and to possibly revisit them later if a bigger interest in the mathematical details should arise.

Though I have written everything in this chapter in my own words, some sections are mentally largely inspired by the introduction of Hahn & Angulo (2016).

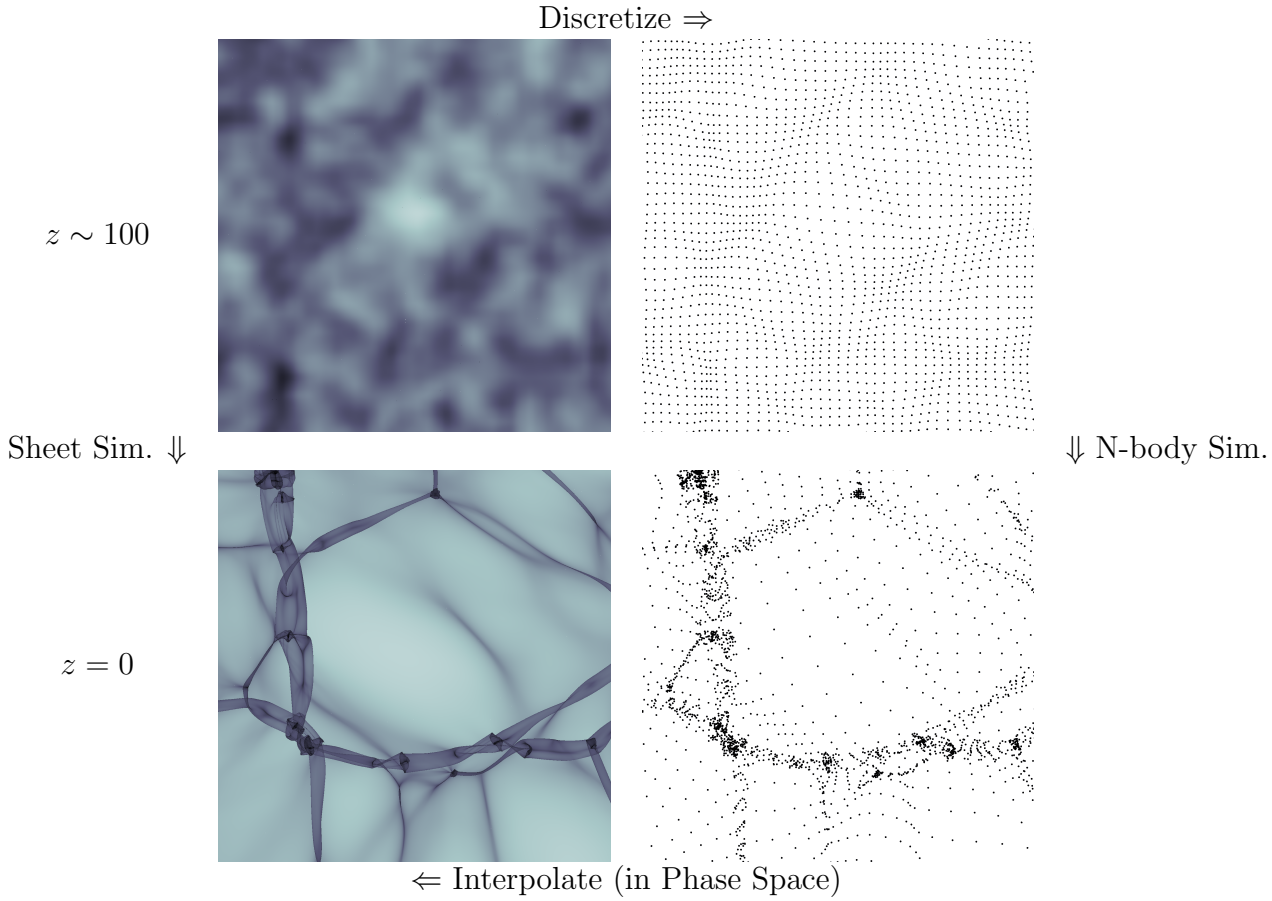

Figure 2.1: Illustration of the density distribution in the early universe and its connection to the late time distribution through N-body or Sheet simulation techniques. The continuous density field in the early universe (top left) is known mathematically very precisely (through the CMB). It is the goal of cosmological simulations to predict the late-time density field (and other observables), starting from this early stage. N-body simulations discretize the continuum to a set of particles (top right) and then evolve these by their Newtonian gravitational interactions to a final state (bottom right). This state is typically interpreted as a Monte-Carlo realization of the actual continuous system. However, one can ask the question: how to obtain an estimate of the actual continuous distribution? As we will describe in this chapter, this can be done by interpolation between particles in phase space (Shandarin et al., 2012; Abel et al., 2012) – leading to the continuum density estimate at the bottom left. Sheet-based dark matter simulation techniques (from top left to bottom left) try to directly follow the evolution of the continuous system to infer the final state. These Figures have been created with the "sheet2d"-code that I have worked on in collaboration with Oliver Hahn.

2.1 The early universe

In the ΛCDM model the state of our universe around redshift $z \sim 100$ is known to an extraordinary accuracy. That is so, since in the early universe $z \gtrsim 100$ the relative density contrast

$$\delta = \frac{\rho(\mathbf{x}) - \rho_0}{\rho_0} \tag{2.1}$$

is very small $\delta \ll 1$. Therefore it is well approximated by linear theory which expands around a small value of δ. Boltzmann codes like CAMB (Lewis & Bridle, 2002) or CLASS (Blas et al., 2011) solve the co-evolution of the distributions of all relevant species in the early universe. These have been used in combination with the observations of the CMB to constrain our cosmological parameters to an extraordinary accuracy (see section 1.2.2). Further, they give a very clear picture of how our universe looks in the linear regime. However, they cannot describe the evolution in the nonlinear regime $z \ll 100$ since the approximations do not hold if $\delta \gg 0$.

It is the goal of cosmological simulations to understand and predict the evolution of the different constituents of our universe after the point where linear theory becomes invalid. The predictions of linear theory can be used to set the initial conditions in the linear regime and those are then evolved by a non-linear set of equations until today.

We briefly review in this section how our universe looks at a redshift of $z \sim 100$ and therefore how the initial conditions for our cosmological simulations are defined. Therefore we explain in section 2.1.1 what is the warmth of dark matter – referring to its velocity distribution in the early universe – in section 2.1.2 we explain how linear theory is used to predict the density field in the early universe $z \sim 100$ and in section 2.1.3 how the Zel'dovich approximation can be used to estimate positions and velocities as initial conditions for cosmological simulations.

2.1.1 The warmth of dark matter

All of the physically well motivated dark matter candidates already have a thermal (or non-thermal) velocity dispersion in the early (almost homogeneous) universe $z \gtrsim 100$. We speak of *warm dark matter* (WDM) if the velocity dispersion might be large enough to influence structure formation on observable scales whereas we speak of *cold dark matter* (CDM) if the velocity dispersion has a negligible amplitude.

The exact shape of the velocity distribution depends on the properties of dark matter. However, for many considerations the shape of the distribution does not matter, but only its width $\sigma_v^2 = \langle v^2 \rangle$. It is therefore convenient to treat the warmth of dark matter as a one parameter model. It would make sense to simply quote the value of σ_v when speaking of the warmth of dark matter. However, it is a convention in the literature to assume the velocity distribution of a thermal relic particle and to quote its mass m_X instead of its velocity dispersion σ_v. We will follow this convention. In the following we will derive the relation between the velocity distribution and the mass of a thermal relic.

A thermal relic

A thermal relic is a particle which is produced in thermal creation/annihilation equilibrium in the early universe. If it decouples while being relativistic and has negligible self-interactions after decoupling, its phase space distribution is quite simple. It becomes non relativistic long after decoupling. The prime example for a thermal relic are neutrinos. However, neutrinos cannot be the massive missing dark matter component in our universe, since their temperature T_ν (which is well known) is too high to achieve the correct density parameter Ω_{dm} (constrained from cosmology) and a small enough velocity dispersion (limited by structure formation requirements) at the same time. A thermal relic can be seen as a generalization of the neutrino case where the temperature T_X is kept as a free parameter (that is equivalent to leaving its mass m_X as a free parameter) so that a much smaller velocity dispersion can be achieved when fixing the density parameter $\Omega_x = \Omega_{\mathrm{dm}}$ to the dark matter value.

If we assume that a thermal relic is a fermion which is created in relativistic equilibrium, its phase space distribution function is given by

$$f(\mathbf{x}, \mathbf{p}) = \frac{1}{h^3} \frac{g_i}{\exp\left(\frac{E-\mu}{kT}\right) + 1} \tag{2.2}$$

$$= \frac{1}{h^3} \frac{g_i}{\exp\left(\frac{pc}{kT}\right) + 1} \tag{2.3}$$

where $\mathbf{p}$ is the momentum, p its absolute value, T is the temperature, k the Boltzmann constant, h the Planck constant, c the speed of light, g_i the degeneracy factor of the particle (presumably 2 because of spin degeneracy) and μ the chemical potential. We have used here that the chemical potential is zero because of the annihilation/creation equilibrium and the energy can be well approximated by $E \approx pc$ since we are talking about an ultra-relativistic particle. Further we have assumed that the universe is homogeneous $f(\mathbf{x}, \mathbf{p}) = f(\mathbf{p})$ at the time of interest which is a good approximation at early times $z \gg 100$.

The number density of particles n_X is given as an integral over the phase space distribution

$$n_X = \int_{\mathbb{R}^3} f(p)\, \mathrm{d}^3\mathbf{p} \tag{2.4}$$

$$= \frac{g_i}{h^3} \int_0^\infty \frac{4\pi p^2}{\exp\left(\frac{pc}{kT}\right) + 1}\, \mathrm{d}p \tag{2.5}$$

$$= \frac{g_i}{2\pi^2} \left(\frac{kT}{\hbar c}\right)^3 \int_0^\infty \frac{u^2}{\exp(u) + 1}\, \mathrm{d}u \tag{2.6}$$

$$= \frac{3}{4} \frac{g_i}{\pi^2} \left(\frac{kT}{\hbar c}\right)^3 \zeta(3) \tag{2.7}$$

where we have used $\hbar = \frac{h}{2\pi}$ and the Riemann zeta function ($\zeta(3) \approx 1.20$) which solves the

integral

$$\int_0^\infty \frac{u^n}{\exp(u)+1}\,\mathrm{d}u = n!\,\zeta(n+1)(1-2^{-n}) \tag{2.8}$$

This is the number density of a thermal relic when it is relativistic and in creation/annihilation equilibrium. After it decouples and becomes non-relativistic there is no mechanism which maintains the thermal equilibrium. However, coincidentally the number density of particles decreases with the scale factor as $n_X \propto a^{-3}$ and the momentum redshifts like $p \propto a^{-1}$. Therefore one can write in the relativistic and in the non-relativistic regime the distribution function as a relativistic Fermi-Dirac distribution with a modified temperature $T = T_0(\frac{a}{a_0})^{-1}$. Note that it is a bit misleading to speak of a temperature in the non-relativistic regime since the particles are not really in thermal equilibrium. If the particles were able to maintain a kinetic thermal equilibrium (by collisions in the non-relativistic regime) they would approach a Maxwell-Boltzmann distribution instead.

The number density of the thermal relic relates to its density parameter Ω_X as

$$\begin{aligned}
\Omega_X(a) &= \frac{\rho_X(a)}{\rho_{tot}(a)} && (2.9)\\
&= \frac{8\pi G}{3H^2(a)} m_X n_X(a) && (2.10)\\
&= \frac{8\pi G}{3H^2(a)} m_X \frac{3}{4}\frac{g_i}{\pi^2}\left(\frac{kT_X(a)}{\hbar c}\right)^3 \zeta(3) && (2.11)\\
&= \frac{1}{h^2}\left(\frac{g_X}{2}\right)\left(\frac{m_X c^2}{94\,\mathrm{eV}}\right)\left(\frac{T_X}{T_\nu}\right)^3 && (2.12)
\end{aligned}$$

where $h = H_0/(100\,\mathrm{km\,s^{-1}\,Mpc^{-1}})$ and $T_\nu \approx 1.95\,\mathrm{K}$ is the temperature of the neutrino background. The last step used the fact that neutrinos are a thermal relic to normalize to the case of neutrinos.

Note that the causal logic goes in the order: when given a temperature T_X and a mass m_X of the thermal relic particle that leads to some density parameter Ω_X. However, when assuming that our thermal relic particle is dark matter (in the sense of the massive missing component in our universe) we actually know the density parameter $\Omega_X = \Omega_{\mathrm{dm}}$ today from observations whereas we know nothing about its mass m_X and at most an upper limit for the temperature T_X (that upper limit is the neutrino temperature). So if we are choosing Ω_X from observations and m_X as a free parameter those two parameters automatically define T_X and therefore the velocity distribution of the thermal relic.

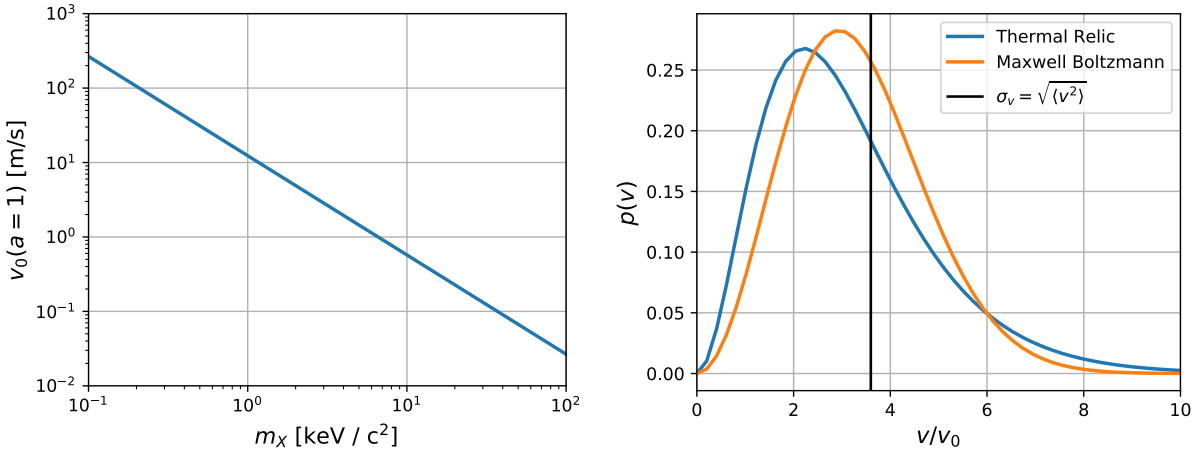

Figure 2.2: Left: The characteristic velocity v_0 as a function of the thermal relic mass for the Planck 2016 cosmology. Right: Velocity distribution function of a thermal relic in comparison to a Maxwell Boltzmann distribution with the same velocity dispersion $\sigma_v = \sqrt{\langle v^2 \rangle} = 3.60 v_0$. The thermal relic distribution function peaks at smaller velocities, but has a much wider tail than the Maxwell Boltzmann distribution.

Its velocity distribution is then given by

$$p(v) = \frac{2}{3\zeta(3)} \frac{v^2}{\exp\left(\frac{v}{v_0(a)}\right) + 1} \tag{2.13}$$

$$v_0(a) := \frac{kT_X}{a\hbar m_X c} \tag{2.14}$$

$$= 12.6\,\mathrm{m\,s^{-1}} \left(\frac{\Omega_X}{0.3}\right)^{1/3} \left(\frac{h}{0.65}\right)^{2/3} \left(\frac{1.5}{g_X}\right)^{1/3} \left(\frac{\mathrm{keV}}{m_X}\right)^{4/3} a^{-1} \tag{2.15}$$

where we used that the thermal relic is non-relativistic today $p = m_X v$. v_0 relates to the velocity dispersion $\sigma_v = \sqrt{\langle v^2 \rangle}$ as $\sigma_v = 3.60 v_0$. Here we normalized to $g_X = 1.5$ (case of a gravitino) for correspondence with Bode et al. (2001). In the left panel of Figure 2.2 we show v_0 as a function of thermal relic mass m_X assuming that thermal relic is dark matter ($\Omega_X = \Omega_{\mathrm{dm}}$) and the Planck Collaboration et al. (2016) cosmology as listed in Table 1.1. For example for a thermal relic with $m_X = 3\,\mathrm{keV}$ the thermal velocity dispersion would be $\sigma_v(a) \approx a^{-1} \cdot 10\,\mathrm{m\,s^{-1}}$. In the right panel of Figure 2.2 the velocity distribution is shown. It is slightly different from a Maxwell Boltzmann distribution with the same velocity dispersion. It peaks at lower velocities, but has a longer high velocity tail.

The velocity dispersion of dark matter has a relatively simple effect in the early universe: It smooths out density perturbations on scales that are smaller than the free-streaming scale - that is the distance particles with a typical thermal velocity would travel in total due to that thermal velocity. These effects are well described by linear theory models.

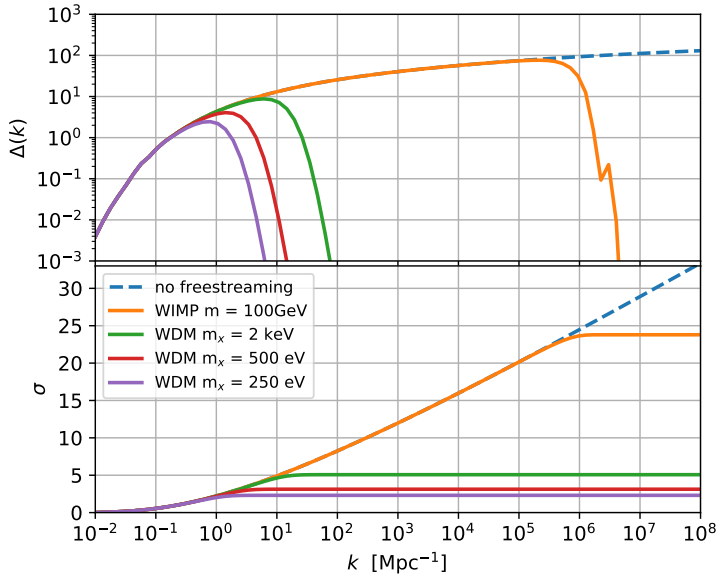

Figure 2.3: Top: the dimensionless linear power spectrum for different different dark matter models. Bottom: the σ values obtained by integrating the power spectra up to a scale k according to (2.18). Dark matter models with a higher velocity dispersion have a cut-off on larger scales and a smaller value of σ.

2.1.2 Linear theory

In the linear regime the statistics of the density field of our universe can be summarized by

$$\langle \delta(\mathbf{k})\delta(\mathbf{k}')\rangle = (2\pi)^3 P(k)\delta_D(\mathbf{k}-\mathbf{k}') \tag{2.16}$$

where $\delta(\mathbf{k})$ is the Fourier transformation of the (over-)density field, δ_D denotes the 3d Dirac delta function, $P(k)$ is the power spectrum and $\langle ...\rangle$ denotes an expectation value. Note that this definition already assumes that there are no correlations between different modes $\mathbf{k} \neq \mathbf{k}'$ in the linear regime and further that our universe is isotropic $P(\mathbf{k}) = P(k)$. Note that for example if there is an ubiquitous tidal field it would lead to an angle-dependent power spectrum that depends on the direction of the $\mathbf{k}$-vector (Barreira & Schmidt, 2017; Schmidt et al., 2018).

The power spectrum is predicted by linear theory and depends on the cosmological parameters and the considered dark matter model. Further it is slightly different for dark matter and baryons, but this is a higher order effect that we can neglect here. In Figure 2.3 we show the dimensionless power spectrum $\Delta(k)$ and the total standard deviation $\sigma(k)$

which are given by

$$\Delta(k) = \frac{k^3}{2\pi^2} P(k) \tag{2.17}$$

$$\sigma(k) = \left(\int_{-\infty}^{\ln(k)} \Delta(k') d\ln k' \right)^{1/2} \tag{2.18}$$

for an example set of dark matter models. Although the power spectra can accurately be calculated with Boltzmann codes, we use the approximate description formulas from Eisenstein & Hu (1999), Bode et al. (2001) and Green et al. (2005) here (as will be described more clearly in section 3.3.5). These are still quite accurate and this avoids dealing with problems with Boltzmann codes when trying to evaluate very small scales $k \gg 10^3 \mathrm{Mpc}/h$ which we will use in chapter 3.

The power spectrum is given at redshift 0 (where it is actually not valid anymore). However, in linear theory the power spectrum can be rescaled to earlier times by the linear growth factor D:

$$\Delta(k, a) = D(a)\Delta(k, a = 1) \tag{2.19}$$

$$D_*(a) = \frac{5\Omega_m}{2} \frac{H(a)}{H_0} \int_0^a \left(a \frac{H(a)}{H_0} \right)^{-3} da \tag{2.20}$$

$$D(a) = \frac{D_*(a)}{D_*(a = 1)} \tag{2.21}$$

where D_* is normalized so that it is $D_* \sim a$ during matter domination and D is normalized so that it is $D = 1$ at $a = 1$.

We can think of our universe as a randomly drawn realization from all the possible universes that fulfil the same power spectrum. When we make a cosmological simulation, we simulate a different random realization than our universe which fulfils the same statistics (as far as we know). As a side product of the project that we will discuss in chapter 3 we have written a small python code which can generate real space initial conditions for dark matter simulations. We will shortly discuss here how one can generate a real space density field $\delta(\mathbf{x})$ given a power spectrum. Note that an additional assumption that is made during that step is that the density field has Gaussian statistics, i.e.

$$p(\delta(\mathbf{k})) = \frac{1}{\sqrt{2\pi P(k)}} \exp\left(-\frac{1}{2P(k)} \right) \tag{2.22}$$

which is the case for the primordial density field as far as we know. However, trying to measure possible non-Gaussianities in the primordial density field is an active research topic.

To get a real space realization on a grid of a periodic density field which follows a power spectrum $P(\mathbf{x})$ one can simply do the following

1. Set up a scalar valued grid in Fourier space where each grid point is sampled from a Gaussian normal distribution with center 0 and standard deviation $\sigma = 1$.

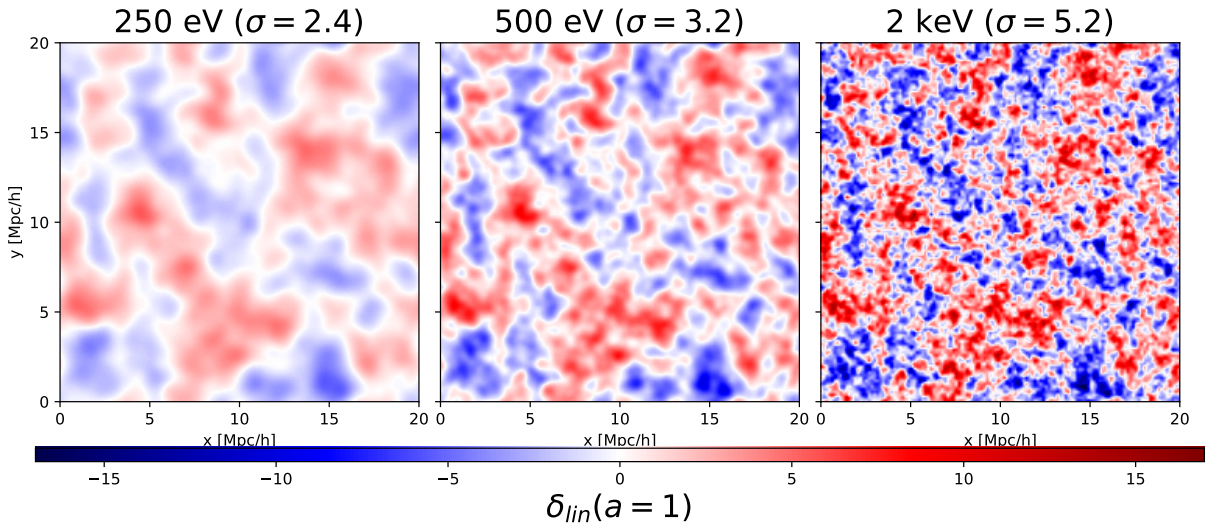

Figure 2.4: Thin slices through the linear theory density fields of different warm dark matter universes. This is approximately how our universe could look around $z \sim 100$ (though the amplitude of δ is normalized to $z = 0$ here). We additionally denoted the standard deviation of the density fields as given by $\sigma(k \to \infty)$ according to 2.18.

2. Multiply each grid point depending on its frequency $\mathbf{k}$ by $\sqrt{P(|\mathbf{k}|)}$. Now we have $\delta(\mathbf{k})$.

3. Apply an inverse Fourier transformation. Now we have $\delta(\mathbf{x})$.

4. (Depending on the definition of the used Fourier transformation additional constant normalization factors might have to be applied.)

With this simple procedure one can already obtain an image of our universe at $z \sim 100$. Such an image is shown in Figure 2.4 for different dark matter models. (The amplitude is scaled to $a = 1$ however.) Depending on the warmth of dark matter, small scale density perturbations are smoothed out on different scales. Also it becomes quite visible that warmer dark matter models have a lower amount of total power which we quantified by the value of σ.

The reader might wonder why we do not plot the case of cold dark matter here. We do not plot it, because the cold dark matter case is not very well defined at any finite resolution. It depends strongly on the resolution of the used grid, since the power spectrum is always cut off implicitly at the Nyquist frequency. We instead choose in this thesis to always cut-off the power spectrum explicitly by using a thermal cut-off. This way we can easier disentangle true resolution effects from effects that depend on the power-spectrum.

2.1.3 The Zel'dovich approximation

While the density fields in Figure 2.4 give the right impression of how our universe looks around $z \sim 100$, they do not even get close to how it looks at $z \sim 0$. For example the density perturbations from linear theory can easily get much smaller than $\delta \ll -1$ which corresponds to negative densities. Clearly they must be wrong at $z \sim 0$. A more realistic picture can be obtained is given by the Zel'dovich approximation (Zel'dovich, 1970).

$$\mathbf{x}(a) = \mathbf{q} + D(a)\mathbf{s}(\mathbf{q}) \tag{2.23}$$

The Zel'dovich approximation estimates at which location $\mathbf{x}$ a particle would be at time a if it started at the Lagrangian coordinates $\mathbf{q}$ (which is its hypothetical location for $a \to 0$). The displacement field $\mathbf{s}$ relates to the linear density field as

$$\mathbf{s}(\mathbf{q}) = -\boldsymbol{\nabla}_q \phi_s(\mathbf{q}) \tag{2.24}$$

$$\nabla_q^2 \phi_s = \delta_{\mathrm{lin}} \tag{2.25}$$

where we defined the displacement field potential ϕ_s. To create initial conditions for a cosmological simulation (a set of position $\mathbf{x}$ and velocity variables $\mathbf{v}$) one can now proceed as following:

1. Obtain a Fourier space density field $\delta(\mathbf{k})$ as already described in section 2.1.2.
2. Obtain the displacement field in Fourier space $s(\mathbf{k})$ by multiplying with $-i\mathbf{k}/(ik)^2$ which is equivalent to the real space operator $-\boldsymbol{\nabla}_q/\nabla_q^2$.
3. Transform back to real space to get the displacement field $\mathbf{s}(\mathbf{q})$.
4. (Multiply by correct normalization depending on Fourier convention.)

Positions $\mathbf{x}$ can then be obtained like in (2.23) and velocities $\mathbf{v}$ are given by

$$\mathbf{v} = \dot{\mathbf{x}} a^2 \tag{2.26}$$

$$= f H(a) D \mathbf{s} \tag{2.27}$$

where the first line defines the velocity variable that we will use by default in this thesis unless stated otherwise. It is a standard choice for simulations, because it makes the equations of motion quite simple, as we will discuss in section 2.2. $f = \mathrm{d}\ln D/\mathrm{d}\ln a$ is the growth rate which can be approximated by $f \sim \Omega_m^{5/9}(a)$ (Bouchet et al., 1995) and is close to unity during matter domination.

In Figure 2.5 we show a slice through the density field that is obtained by applying the Zel'dovich approximation for an $m_X = 500\,\mathrm{eV}$ thermal relic in comparison to the one predicted by linear theory. For early times they agree quite well. For intermediate times $D\sigma \lesssim 1$ the Zel'dovich approximation shows a way more dynamic situation which is relatively close to what happens in actual simulations. However, at later times $D\sigma \gg 1$, it also fails catastrophically, since deaccelerating forces are missing. We will come back to

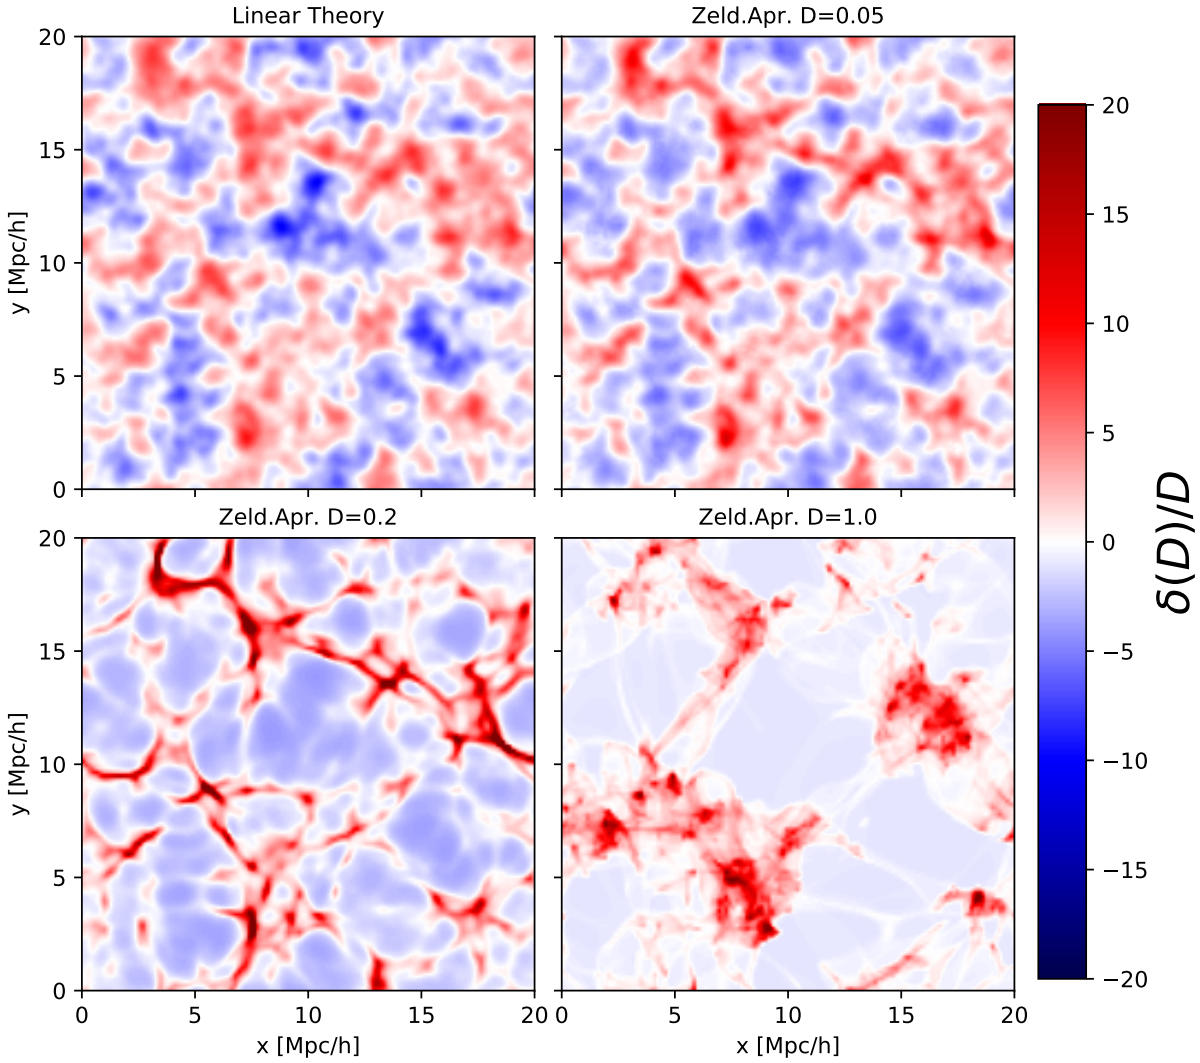

Figure 2.5: Thin slice through a 3D density field ($m_X = 500\,\mathrm{eV}$, $\sigma = 3.2$) for using linear theory (top left) and the Zel'dovich approximation at different growth factors as stated in the headlines of each plot. The color space is normalized as $\delta(D)/D$ so that linear theory is independent of D whereas the Zel'dovich approximation is strongly dependent on D in the non-linear regime. At early times $D\sigma \ll 1$ the Zel'dovich approximation and linear theory agree well. At intermediate times $D\sigma \lesssim 1$ it creates a physically meaningful picture of collapsing structures in our universe. At late times $D\sigma \gg 1$ it creates unphysical over-extending shells, since deaccelerating forces are missing.

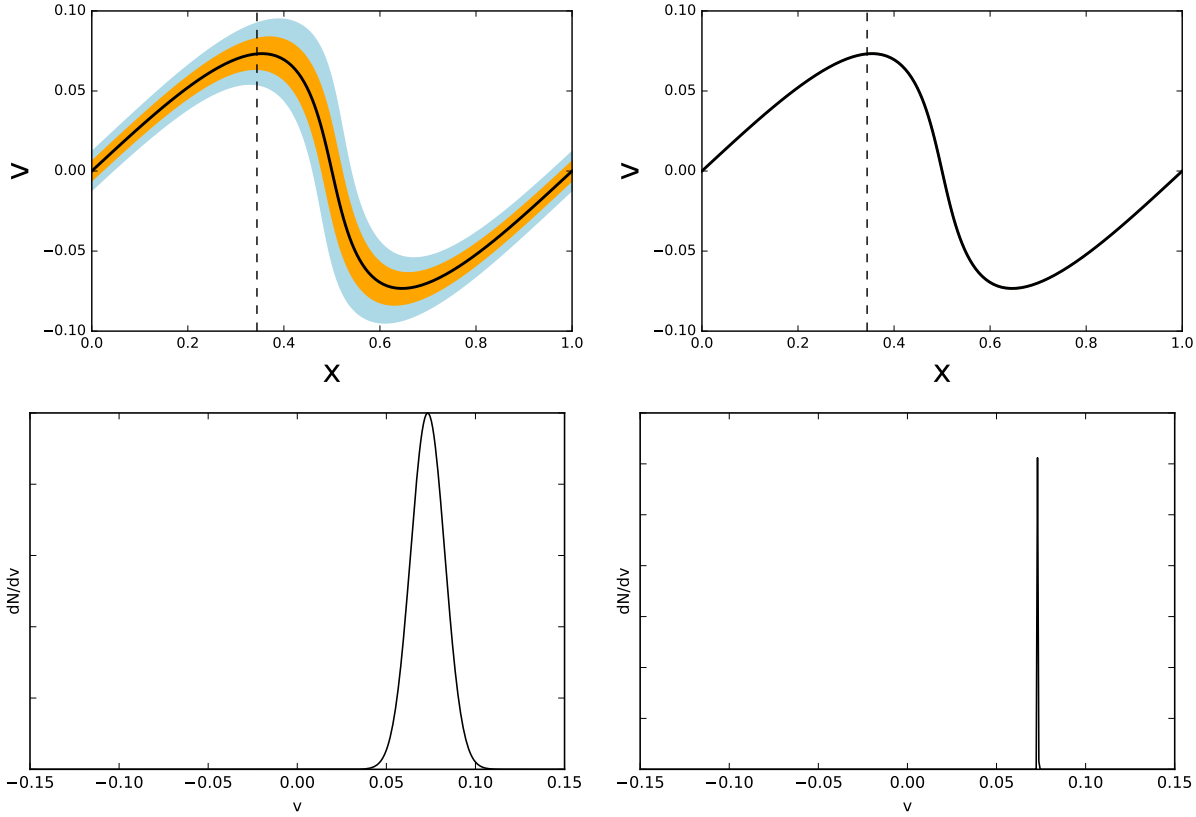

Figure 2.6: An illustration of the phase space distribution of dark matter in the early universe. Top: phase space distribution as a contour plot with the 1σ and 2σ contours marked. Bottom: the velocity distribution at the x location that is marked in the upper plots. Left: if dark matter had hot initial conditions (for example like neutrinos). The amplitude of the velocity dispersion is high enough that it is important dynamically and it cannot be neglected in simulations. Right: cold initial conditions σ_v is much smaller than the bulk velocities. It can effectively be neglected as we will see later.

this topic in chapter 3 where we will use the Zel'dovich approximation as a toy model to make analytic estimates about the density distribution of our universe.

The Zel'dovich approximation already describes a lot of the fine-grained features that exist in the dark matter distribution qualitatively. We will describe them later in the context of actual simulations. However, for an interesting treatment in the context of catastrophe theory, we would like to refer the reader to Hidding et al. (2014).

2.1.4 The phase space distribution

The way we stated linear theory and the Zel'dovich approximation, they only give information about the moments of the phase space distribution (before shell crossing)

$$\rho(\mathbf{x}) = m_X \int_0^\infty f(\mathbf{x}, \mathbf{v}) \mathrm{d}^3\mathbf{v} \tag{2.28}$$

$$\mathbf{v}_{\mathrm{bulk}}(\mathbf{x}) = \frac{\int_0^\infty \mathbf{v} f(\mathbf{x}, \mathbf{v}) \mathrm{d}^3\mathbf{v}}{\int_0^\infty f(\mathbf{x}, \mathbf{v}) \mathrm{d}^3\mathbf{v}} \tag{2.29}$$

but do not actually define the full phase space distribution of dark matter. However, on a more fundamental level, dark matter must be treated as a collisionless fluid in phase space (as we shall discuss in more detail in section 2.2). Therefore we have to answer the question here: how does the full phase space distribution of dark matter look around the time a simulation would have to start?

In Figure 2.6 we illustrate two different possible cases of phase space distributions that we could have in the early universe. The left side shows a distribution where the bulk velocities are of comparable amplitude to the velocity dispersion $\sigma_v \sim v_{\mathrm{bulk}}(\mathbf{x})$. In cases like this we would speak of a hot distribution function. The right side shows a cold distribution function where $\sigma_v \ll v_{\mathrm{bulk}}(\mathbf{x})$. The question is: which case do we have for the dark matter fluid?

We estimate from the variance of the displacement field

$$\sqrt{\langle s^2 \rangle} = \sqrt{\int_{-\infty}^{\infty} k^{-2} \Delta(k') d \ln k'} \approx 10\,\mathrm{Mpc} \tag{2.30}$$

$$\sqrt{\langle v_{\mathrm{bulk}}^2 \rangle} = f H_0 \sqrt{\langle s^2 \rangle} \approx 675\,\mathrm{km\,s^{-1}} \tag{2.31}$$

that typical bulk velocities today are of order $10^3 \mathrm{km\,s^{-1}}$, whereas the warmest warm dark matter models ($m_x \sim 3\,\mathrm{keV}$) have $\sigma_v \sim 10\,\mathrm{m\,s^{-1}}$ today which is 5 orders of magnitude smaller. Comparing these two values with each other is not exactly fair, because the bulk velocities grow over time, whereas the thermal velocities decrease over time (if peculiar velocities $u = v(a)/a$ are used as velocity variable). Therefore the ratio of these two has been smaller at earlier times. However, at times where simulations would start typically, bulk velocities are already much bigger than thermal velocities $\sigma_v \ll v_{\mathrm{bulk}}$ so the situation looks already similar to the right panel of Figure 2.6.

That said, it is a good approximation to treat dark matter dynamically as a cold fluid from the point where one would start a simulation. That means even for warm dark matter we can set σ_v to zero at that time. The warmth of dark matter is then only incorporated by modifying the initial power spectrum already shown in Figure 2.3 and 2.4. We will discuss possibilities of how to incorporate the full phase space distribution in simulations later in chapter 6. However, for most of this thesis we will make the approximation $\sigma_v \sim 0$ in the initial conditions of the simulation.

In the limit of cold initial conditions $\sigma_v \to 0$ the dark matter fluid occupies a thin three dimensional sub-manifold in phase space. This is also referred to as the *dark matter sheet* in phase space (Arnold et al., 1982; Shandarin & Zeldovich, 1989; White & Vogelsberger, 2009; Abel et al., 2012; Shandarin et al., 2012; Hahn & Angulo, 2016; Sousbie & Colombi, 2016). It is subject of the next sections to understand how to evolve the dark matter sheet in the non-linear regime.

2.2 Evolution equations

We have described in the last section, how the early universe looks like, and thereby what kind of initial conditions are suitable for cosmological simulations. In this section we want to review shortly the evolution equations of the dark matter fluid. Therefore we will first derive the equations of motion for a single particle that is surrounded by some mass distribution $\rho(\mathbf{x})$. Afterwards we will discuss the collisionless Boltzmann equation, which is the most fundamental description of the collisionless dark matter fluid. Combined with the Poisson equation this forms the Vlasov-Poisson system. The underlying assumption is that dark matter can be treated as a collisionless fluid.

It is quite easy to see that dark matter can be treated as a fluid. By this we mean that it is well described by some continuous phase space density $f(\mathbf{x}, \mathbf{v})$. It consists of a gigantic number of particles when compared to the cosmological scales that we will speak of in this thesis. For example if it is made out of a $m \sim 100\,\mathrm{GeV}$ particle (e.g. a WIMP), the number density is approximately $\rho_0/m \sim 0.03\,\mathrm{m}^{-3}$. That corresponds to $\sim 10^{48}$ particles in a cubic parsec volume. Therefore at cosmological length scales (kilo-parsecs to Giga-parsecs) dark matter can be considered as a continuous fluid.

The treatment of dark matter as collisionless is less obvious. It depends on the nature of dark matter whether it has self-interactions and how strong these self-interactions are. However, for most considered dark matter models self-interactions are very weak. Further, the self-interactions have been constrained by astrophysical observations like the Bullet cluster to be quite small (Robertson et al., 2017). We want to point out here that measuring the degree of self-interaction from astrophysical observations might be a way to learn more about its nature. However, so far it is consistent with a completely collisionless picture. For simplicity we will assume dark matter as collision-less for the remainder of this thesis.

Therefore dark matter can be treated as a *self-gravitating collisionless fluid*.

2.2.1 The equations of motion

The correct way of deriving the equations of motion in an expanding universe is to derive a general relativistic description of its metric and then set up the geodesic equation. However, this is not very intuitive and makes it quite hard to see how Hamiltonian phase space dynamics generalize into the comoving frame. Assuming that all involved velocities are non-relativistic, locally the equations of motion have to revert to a Newtonian description. To get an intuitive picture of what changes when viewing the world from the comoving

frame, we start from the Newtonian equations of motion and transform them to comoving coordinates. For this, we take the Friedman equations as given and assume a universe which contains only matter ($\Lambda = 0$). We will discuss afterwards what has to change if a cosmological constant is included.

In the Newtonian frame the equations of motion of a single particle are given by

$$\ddot{\mathbf{r}} = -\boldsymbol{\nabla}_r \phi_r \tag{2.32}$$

$$\nabla_r^2 \phi = 4\pi G \rho_r \tag{2.33}$$

where r are its physical coordinates, ϕ_r is the Newtonian potential, ρ_r is the mass density, $\boldsymbol{\nabla}_r$ denotes a gradient operator with respect to physical coordinates and dotted variables denote time derivatives.

We can choose to transform the coordinates $\mathbf{r}$ into any other frame and reexpress the equations of motion in that frame. For an expanding universe a convenient choice is

$$\mathbf{x} := \frac{\mathbf{r}}{a} \tag{2.34}$$

$$\mathbf{v} := \dot{\mathbf{x}} a^2 = \dot{\mathbf{r}} a - \mathbf{r}\dot{a} \tag{2.35}$$

$$\boldsymbol{\nabla}_x = a \boldsymbol{\nabla}_r \tag{2.36}$$

where $\mathbf{x}$ are comoving coordinates and $\mathbf{v}$ is a velocity variable (where peculiar velocities are given by $\mathbf{v}/a$). With the new coordinates we can reexpress the equations of motions as

$$\dot{\mathbf{x}} = \frac{\mathbf{v}}{a^2} \tag{2.37}$$

$$\dot{\mathbf{v}} = \ddot{\mathbf{r}} a - \mathbf{r}\ddot{a} \tag{2.38}$$

$$=: -\frac{\boldsymbol{\nabla}_x \phi_c}{a} \tag{2.39}$$

$$\nabla_x^2 \phi = a^2 \nabla_r \left(-\ddot{\mathbf{r}} a + \mathbf{r}\ddot{a}\right) \tag{2.40}$$

$$= a^3 \left(4\pi G \rho_r + 3\frac{\ddot{a}}{a}\right) \tag{2.41}$$

where we have defined the peculiar potential ϕ. The choices for $\mathbf{x}, \mathbf{v}$ and ϕ seem somewhat arbitrary up to this point. However, if the function $a(t)$ is chosen as the expansion factor of the universe we can use the second Friedmann equation (1.3) to eliminate $\ddot{a}/a$ from the new Poisson equation:

$$\nabla_x^2 \phi = a^3 \left(4\pi G \rho_r - 4\pi G \rho_{bg,m}\right) \tag{2.42}$$

$$= 4\pi G a^3 \rho_m \frac{(\rho_r - \rho_m)}{\rho_m} \tag{2.43}$$

$$= 4\pi G \rho_0 \delta \tag{2.44}$$

where we have defined the relative overdensity δ and used that matter dilutes as $\rho_m = a^{-3}\rho_0$ where ρ_0 is the mean matter density of the universe at $a = 1$. We summarize the equations of motion and the Poisson equation:

$$\dot{\mathbf{x}} := \frac{\mathbf{v}}{a^2} \tag{2.45}$$

$$\dot{\mathbf{v}} := -\frac{\boldsymbol{\nabla}_x \phi}{a} \tag{2.46}$$

$$\nabla_x^2 \phi = 4\pi G \rho_0 \delta = \frac{3}{2}\Omega_0 H_0^2 \delta \tag{2.47}$$

(cf. Peebles, 1980) where we also gave an alternative version for the Poisson equation by re-expressing the mean matter density as $\rho_0 = \rho_{\rm crit}\Omega_0 = 3H_0^2\Omega_0/(8\pi G)$. The equations of motion only appear that simple because of the particular choices of $\mathbf{x}$, $\mathbf{v}$ and ϕ. This coordinate choice ensures that $\mathbf{x}$ is constant as a function of time if $\mathbf{v}$ is zero, $\mathbf{v}$ is constant if the mass distribution is homogeneous $\delta = 0$ and ϕ is constant if δ is constant. In other coordinate choices (e.g. when using the peculiar velocity $\mathbf{u} = \mathbf{v}/a$ as velocity variable) additional time and/or coordinate dependencies can arise.

A cosmological constant in the Newtonian Frame

In our derivation we assumed $\Lambda = 0$ and neglected radiation ($p = 0$). However, the final equations of motion are still valid in a universe with a cosmological constant. What changes when including a cosmological constant are the Newtonian equations of motion that we started from. For a non-zero cosmological constant they must read

$$\ddot{\mathbf{r}} = -\boldsymbol{\nabla}_r \phi_r + \frac{\Lambda c^2}{3}\mathbf{r} \tag{2.48}$$

if we do not absorb the cosmological constant into the definition of the potential. In the Newtonian frame the cosmological constant acts as an additional distance dependent force. In the case of a negligible gravitational force $\boldsymbol{\nabla}\phi_r = 0$ the physical distance between two test particles would grow exponentially over time:

$$r(t) = r_0 \exp\left(\sqrt{\frac{\Lambda c^2}{3}}t\right) \tag{2.49}$$

However in the case of strong gravitational forces $\|\boldsymbol{\nabla}_r \phi_r\| \gg \Lambda c^2 r$ the effect of the cosmological constant can generally be neglected. For example let us consider the effect of the additional force due to the cosmological constant in the case of a test-particle in an (otherwise) Keplerian potential:

$$\ddot{r} = -\frac{MG}{r^2} + \frac{\Lambda c^2}{3}\mathbf{r} \tag{2.50}$$

This system has an effective attractive central potential (assuming Λ is small). Therefore the test particle does not spiral outwards. Instead it has a slightly modified bound orbit. Interestingly in the case of the Keplerian potential orbits do not close anymore, but instead

have a precession. However that precession is far too small to be measured. For example in the case of Mercury it is more than 10 orders of magnitude smaller than the measurement uncertainty (Adkins et al., 2007). Post-Newtonian corrections are much larger than the effect of the cosmological constant in our solar system. However, it is not completely unthinkable that in a distant future we could also measure the cosmological constant by its effect on closed orbits.

2.2.2 The Vlasov-Poisson system

Dark matter is collisionless and as a fluid it can be well described by a continuous phase space density $f(\mathbf{x}, \mathbf{v})$. In the comoving frame that we are using here this phase space density is conserved along trajectories. This is described by the collisionless Boltzmann equation

$$\frac{\mathrm{d}f(\mathbf{x}, \mathbf{v}, t)}{\mathrm{d}t} = \frac{\partial f}{\partial t} + \frac{\partial f}{\partial \mathbf{x}} \cdot \frac{\mathrm{d}\mathbf{x}}{\mathrm{d}t} + \frac{\partial f}{\partial \mathbf{v}} \cdot \frac{\mathrm{d}\mathbf{v}}{\mathrm{d}t} = 0 \tag{2.51}$$

The combination with the Poisson equation (2.47) and using the equations of motion from (2.45) - (2.46) leads to the Vlasov Poisson system

$$0 = \frac{\partial f}{\partial t} + \boldsymbol{\nabla}_x f \cdot \frac{\mathbf{v}}{a^2} - \boldsymbol{\nabla}_v f \cdot \frac{\boldsymbol{\nabla}_x \phi}{a} \tag{2.52}$$

$$\rho(\mathbf{x}, t) = m_x \int f(\mathbf{x}, \mathbf{v}, t) \mathrm{d}^3 v \tag{2.53}$$

$$\nabla_x^2 \phi = 4\pi G(\rho - \rho_0) \tag{2.54}$$

Note that the choice of position and velocity variables ensures that the associated phase space density is conserved $\frac{\mathrm{d}f}{\mathrm{d}t} = 0$. If e.g. the peculiar velocity was chosen as velocity variable this would not be the case. Interestingly more or less the same system appears in the description of collisionless plasmas, since the electrodynamical potential also follows the Poisson equation.

2.2.3 Impact of baryons

All cosmological gravity-only simulations (without self-interactions) try to solve the evolution of this system (in some approximate manner). If baryonic interactions are included, the evolution of the collisionless dark matter part of the system still stays equivalent, only the density receiving an additional time dependent component

$$\rho(\mathbf{x}, t) = \int f_{\mathrm{dm}}(\mathbf{x}, \mathbf{v}, t) \mathrm{d}^3 v + \rho_b(\mathbf{x}, t) \tag{2.55}$$

where the baryon density $\rho_b(\mathbf{x}, t)$ evolves according to a different set equations, since it is collisional. In most cases it is assumed that the baryons are in local thermodynamic

equilibrium and therefore can be approximated by hydrodynamic equations. These evolve projected distribution functions that are only a function of space $\mathbf{x}$, but not velocity (like f). However, we will not further worry about the influence of baryons in this thesis, but instead focus on the gravity-only or dark matter only case $\rho_b = 0$. We just want to note here that all qualitative statements in this thesis about the phase space distribution of dark matter still hold in the case that baryons were included. That is so, because dark matter still follows the collisionless Boltzmann equation and the Baryons have only a very indirect impact by modifying the source term in the Poisson equation. However, baryons can modify quantitative results.

2.3 The N-body method

The Vlasov Poisson system of equations are the most fundamental evolution equations to describe a continuous dark matter fluid. A straight-forward approach for solving these would be to set up a Eulerian 6-dimensional phase space grid and solve the system of equations on this grid. However, the dynamical range of the scales that are involved in the evolution of the dark matter fluid, the fact that the initial distribution function is extremely cold, and the fact that the number of required resolution elements scales with the 6th power of spatial/velocity resolution scale, would make such an enterprise relatively hopeless (Yoshikawa et al., 2013, for example).

Instead, the N-body method has established as the method of choice for modeling the evolution of the cold collisionless dark matter fluid. The N-body method tries to sample the phase space distribution f in a Monte-Carlo like fashion. In cosmological simulations typically the cold limit is assumed $\sigma_v \to 0$ so that effectively only the three dimensional sub-space that is occupied by the dark matter sheet has to be sampled by particles. Each of those particles has then a unique location on this Lagrangian submanifold $\mathbf{q}_i$ and has time dependent Eulerian phase space coordinates $\mathbf{x}_i$ and $\mathbf{v}_i$. The particles are evolved as a Hamiltonian system using (softened) point mass interactions to reach a final set of position and velocity coordinates. It is then often assumed at the final state that the N particles can be treated like a Monte-Carlo realization of the true continuous phase space distribution. We illustrate the N-body method schematically in Figure 2.7.

In this section we will briefly describe the Hamiltonian system which is followed in N-body simulations and then point out some of the major achievements of this method. Further we will discuss the issue of artificial fragmentation and discuss the question how far the Monte-Carlo view is appropriate.

2.3.1 The N-body system

From the continuous initial fields of linear theory a discrete approximate realization of N particles is created. These can be for example created by sampling the Zel'dovich

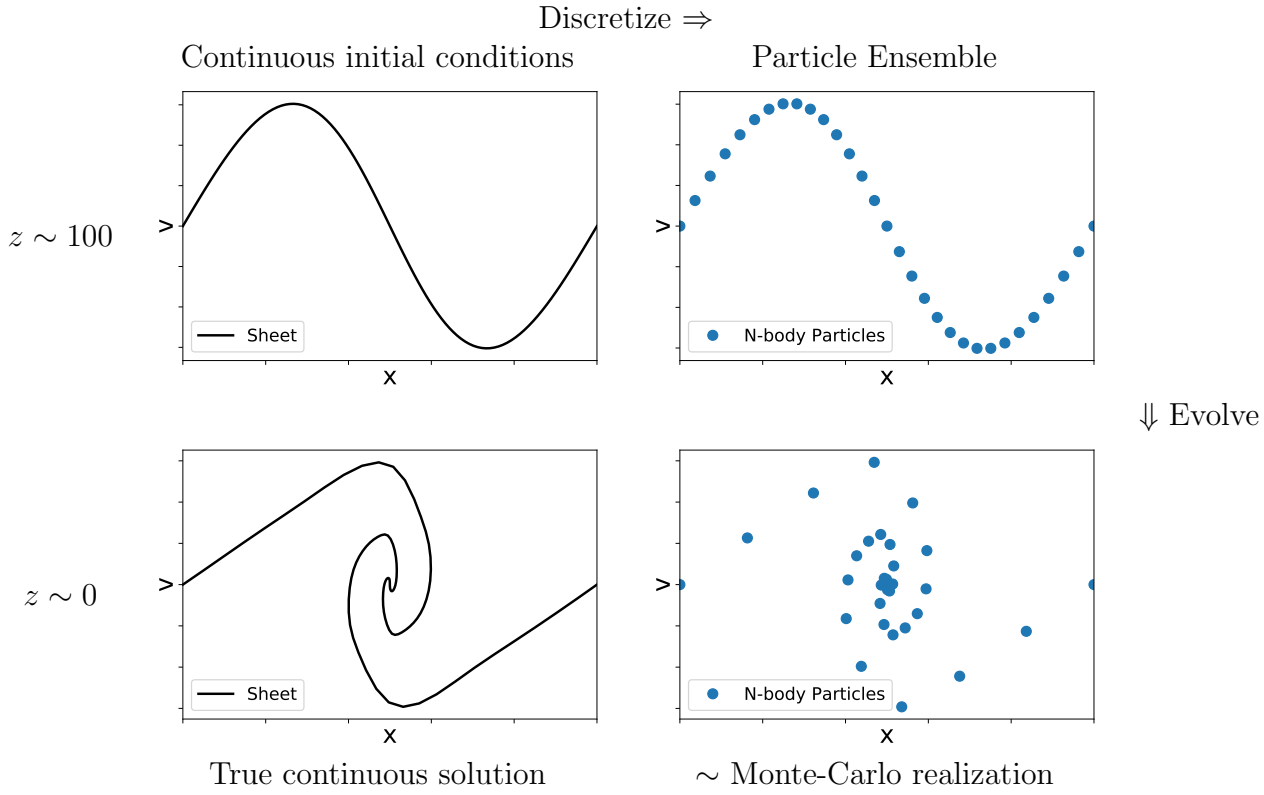

Figure 2.7: Illustration of the N-body method in phase space. From linear theory we have given a very accurate continuous representation of the initial conditions around $z \sim 100$ (top left). These initial conditions are discretized to a set of particles. This can be done for example by using the Eulerian $\mathbf{x}_i, \mathbf{v}_i$ coordinates of a Lagrangian $\mathbf{q}_i$ grid (top right). The particle ensemble is then evolved as a Hamiltonian system using (softened) point-mass interactions to reach a final set of Eulerian coordinates (bottom right). These coordinates are typically interpreted as a Monte-Carlo realization of the true continuous solution of the Vlasov-Poisson system (bottom left).

approximation on a Lagrangian grid

$$\mathbf{q}_{ijk} = \Delta q \begin{pmatrix} i \\ j \\ k \end{pmatrix} \tag{2.56}$$

$$\mathbf{x}_{ijk} = x_{\mathrm{za}}(\mathbf{q}_{ijk}) \tag{2.57}$$

$$\mathbf{v}_{ijk} = v_{\mathrm{za}}(\mathbf{q}_{ijk}) \tag{2.58}$$

$$m_{ijk} = \rho_0 \Delta q^3 \tag{2.59}$$

These particles then follow the Newtonian equations of motion like given in (2.45) and

(2.46). The potential is inferred from a density field which is approximated as

$$\rho(\mathbf{x}) = \sum m_i \overline{\delta}(\mathbf{x} - \mathbf{x_i}) \tag{2.60}$$

where $\overline{\delta}$ is normally not chosen as an actual Dirac δ-function, but instead some softened version of this, e.g. in Springel (2005a) as

$$\overline{\delta}(\mathbf{r}) = \frac{1}{h^3} W\left(\frac{|\mathbf{r}|}{h}\right) \tag{2.61}$$

$$W(u) = \frac{8}{\pi} \begin{cases} 1 - 6u^2 + 6u^3, & u \leq \frac{1}{2} \\ 2(1-u)^3, & \frac{1}{2} < u \leq 1 \\ 0, & u > 1 \end{cases} \tag{2.62}$$

where $h = 2.8\epsilon$ is the radius of the softening kernel and ϵ is the Plummer equivalent softening parameter. With these definitions it is also possible to write the system as a Hamiltonian

$$\sum_i \frac{m_i \mathbf{v}_i^2}{2a^2(t)} + \frac{1}{2} \sum_{i,j} \frac{m_i m_j \psi(\mathbf{x}_i - \mathbf{x}_j)}{a(t)} \tag{2.63}$$

where $\psi(\mathbf{x})$ is the pair potential which incorporates the softening and possibly periodic boundary conditions.

The density estimate as in (2.60) (which we will refer to as N-body density estimate in the future) is a very crude approximation to the actual density field. Typically the softening kernel is chosen much smaller than the mean particle separation $h \ll \Delta q$. As a consequence of this the N-body density is zero in the majority of the volume. As we shall see later, the actual density cannot be zero anywhere if dark matter is a continuous fluid (see chapter 3). However, the density estimate does not need to be accurate to get a reasonable evolution of positions and velocities. For those the force field is more relevant. The force field is the density field convolved with a long range kernel which drops as $1/r^2$ and is not very sensitive to the immediate (noisy) surrounding of each particle.

2.3.2 Visualizations

The N-body method has been extremely successful at predicting the large scale structure of the universe as well as smaller scale properties of the dark matter density field. As such it was found that the density profiles of dark matter haloes can be described by the NFW-profile (Navarro et al., 1996)

$$\rho(r) = \frac{\rho_c}{\frac{r}{r_s}\left(1 + \frac{r}{r_s}\right)^2} \tag{2.64}$$

This two-parameter profile can be seen as an elementary building block of the non-linear dark matter density field. There have been a large variety of predictions and successes which

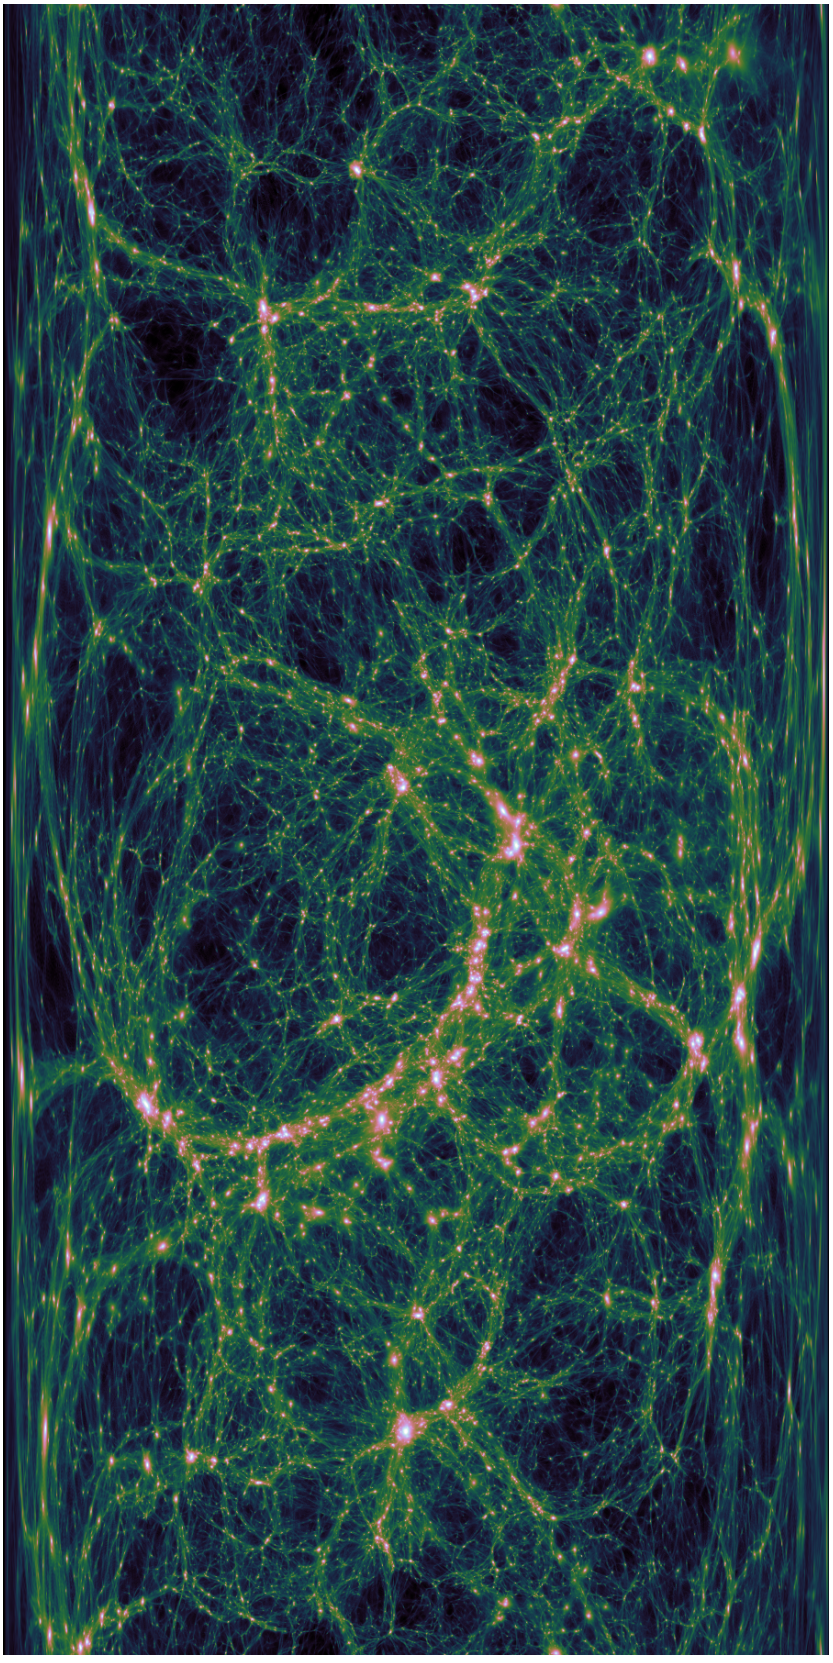

Figure 2.8: Equirectangular projection of the non-linear dark matter density field. This is a single frame from a 360° "virtual reality" movie that can be found under `https://youtu.be/uMBvgCYiUiI`. The non-linear density field contains a large variety of different structures on a large variety of different length scales. Most prominently seen can be here haloes and filaments – aligning in a cosmic web.

were achieved on the back of N-body simulations (Frenk & White, 2012, for an overview). Further they have provided us with intuitive and beautiful images and visualizations of the dark matter density field. Most famous are here probably the images and movies of the Millennium simulation (Springel et al., 2005a).

Since I have some of my own visualization material, I will use this here to provide the reader with an intuitive impression of how the dark matter density field looks in the nonlinear regime. In Figure 2.8 we show an equirectangular projection of the non-linear dark matter density field from a $L = 155\,\mathrm{Mpc}/h$ cold dark matter N-body simulation with 512^3 particles. This image is a frame from a 360° virtual reality (vr) - video which has been created in collaboration with Raul Angulo and can be found under `https://youtu.be/uMBvgCYiUiI`. If opened on youtube, it is possible to change the direction of the view (with the mouse on a computer or by turning around with a phone). With a google card board device[1] one can even watch this movie as a full vr movie. I can highly recommend trying this. I also provide a non-vr version under `https://youtu.be/5r8iY4_7FRI`. The videos use some of the interpolation techniques that will be described in section 2.4 to get a high quality density estimate.

The non-linear density field shows a variety of different structures - such as halos, filaments and pancakes - which assemble on a spider-web like pattern - also known as the cosmic web.

2.3.3 Artificial fragments

While cold dark matter simulations are mostly believed to give reliable results, simulations of warm dark matter create artificial numerical fragments during the early nonlinear stages of evolution. In Figure 2.9 we show examples of these. Wang & White (2007) have shown that these small haloes which align like "beads in a string" inside the filaments are numerical artefacts of the discreteness of N-body simulations. They are not expected in the true continuum solution of the Vlasov Poisson system. As a simple numerical experiment Wang & White (2007) set up the collapse of the N-body realization of a perfectly homogeneous filament which, however, also produced such numerical fragments.

It is important to point out here again that N-body simulations of warm dark matter and cold dark matter only differ in their choice of initial conditions. CDM simulations use a power-spectrum that creates physical halos on all mass scales whereas WDM simulations have power spectra that do not create physical halos below some finite mass scale. It is therefore not entirely clear whether CDM are completely unaffected by this type of problem, or whether the problem is simply not so visible because it is overshadowed by actually physically forming halos.

The reason for the artificial fragmentation is not entirely clear. A possible argumentation (largely from Oliver Hahn, personal conversation) goes like this: When the N-body method is applied to a hot three-dimensional system like a halo, it profits from ergodicity. In a time averaged sense the distribution function of the N-body system gives an accurate

[1] which can be obtained for less than 20 euros if one has a compatible phone

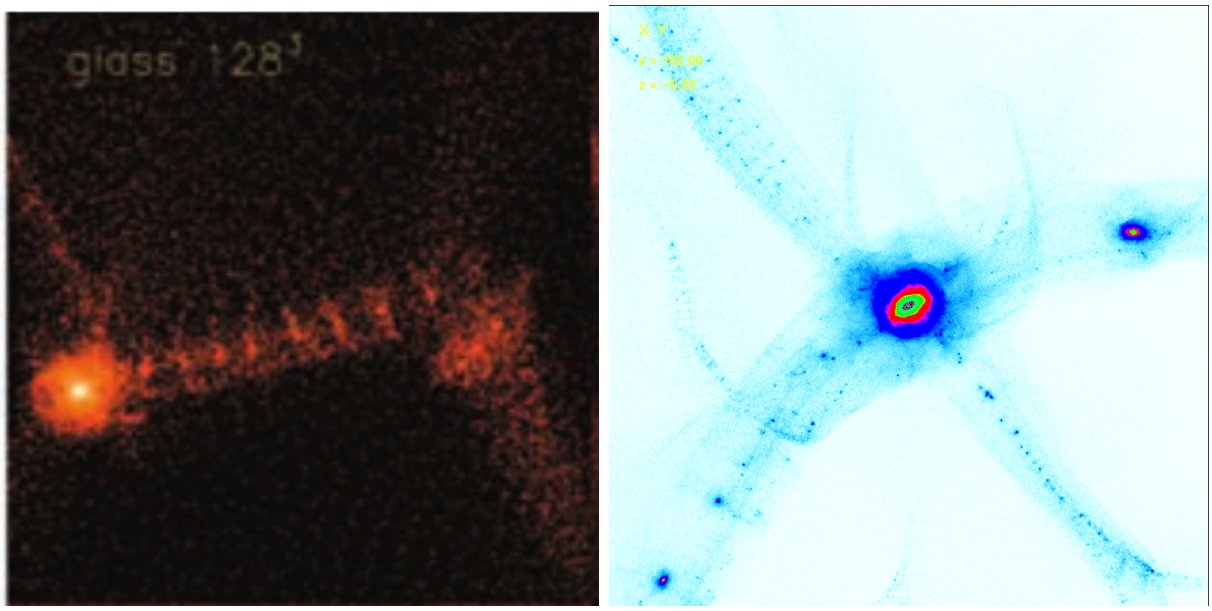

Figure 2.9: Artificial fragmentation in N-body simulations of warm dark matter. Left: Figure 2 from Wang & White (2007) showing regularly spaced fragments in a filament. Right: (figure from my master thesis Stücker (2015)) the filaments around a zoom-in halo are fragmenting into small lumps. In these warm dark matter simulations there should be no such small structures.

representation of the underlying distribution function, even if at any particular moment in time the representation is not quite perfect. However, this is only the case for systems which are hot in every dimension, i.e. which have a non-zero velocity dispersion in every direction. During the early stages of anisotropic collapse, however, systems are formed which only have a velocity dispersion support in one dimension (pancakes) or along two dimensions (filaments). The remaining dimensions remain cold in these cases. In these dimensions the realization-dependent noise does not cancel out in a time-averaged sense. Instead any imperfect perturbation can grow over time and then form for example collapsed artificial halos.

The discreteness effects in warm dark matter simulations, put a question mark on the validity of the N-body scheme in general. In fragmented warm dark matter simulations the Monte-Carlo view is certainly not correct. That is the final N-body distribution can not be seen as a random realization of the true distribution function. Instead it contains additional features that remain after coarse graining. However, it is not clear how far the final distribution deviates from a Poisson sampling of the true distribution. Possibly it is good enough to identify problematic regions (e.g. artificial haloes) and ignore them in the analysis (Lovell et al., 2014) while the Monte-Carlo view might still apply in other regions - like e.g. the centers of massive haloes. It is important to understand this problem quantitatively to test the validity of the N-body scheme and its predictions. Therefore we will discuss in the next sections alternative simulation approaches that do not suffer from discreteness effects. They can be used to benchmark the N-body method.

2.4 Schemes that trace the dark matter sheet

The artificial fragmentation of N-body schemes can be overcome by employing a different density estimate. One which does not suffer that much from shot noise, but instead comes much closer to the continuum limit. Such a density estimate can be inferred by reconsidering the phase space structure of the dark matter sheet.

This section is largely based on the ideas in Shandarin et al. (2012), Abel et al. (2012), Hahn & Angulo (2016) and Sousbie & Colombi (2016). However, I will use my own words and figures for an explanation.

2.4.1 Parameterizing the dark matter sheet

As already discussed, to a good approximation dark matter occupies a three dimensional subspace in six dimensional phase space. This space is parameterized by the Lagrangian coordinate $\mathbf{q}$:

$$\mathbf{q} \rightarrow (\mathbf{x}, \mathbf{v}) \tag{2.65}$$

We can describe the whole dark matter fluid by the two (time-dependent) functions $\mathbf{x}(\mathbf{q})$ and $\mathbf{v}(\mathbf{q})$. We illustrate this in Figure 2.10. It is sufficient to trace the functions $\mathbf{x}(\mathbf{q})$ and $\mathbf{v}(\mathbf{q})$ to reconstruct the final state in phase space. This is a convenient parameterization of

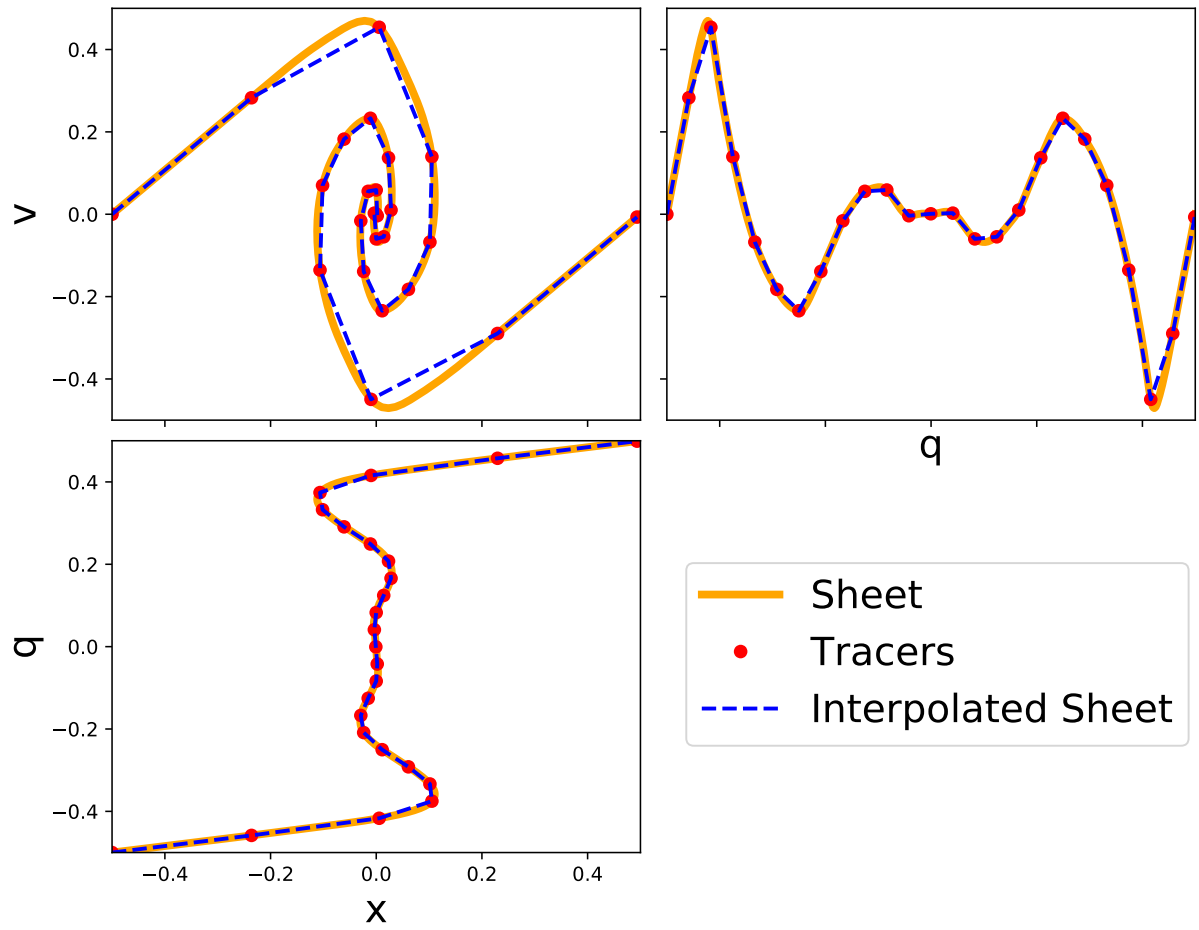

Figure 2.10: Illustration of the parameterization of the dark matter sheet. The phase space sheet (top left) can be parameterized by two functions $\mathbf{x}(\mathbf{q})$ (bottom left with flipped axes) and $\mathbf{v}(\mathbf{q})$ (top right). These functions and their first derivatives are smooth and single-valued. Therefore they can be reconstructed by interpolation techniques from a finite set of tracer particles.

the system, since those functions are single-valued and have some smoothness properties. Saying they are single-valued is equivalent to saying each particle can only have one position and velocity and has a unique Lagrangian coordinate.

We can think of particles $(\mathbf{x}_i, \mathbf{v}_i)$ in simulations as sampling points of these functions

$$\mathbf{x}_i = \mathbf{x}(\mathbf{q}_i) \tag{2.66}$$
$$\mathbf{v}_i = \mathbf{v}(\mathbf{q}_i) \tag{2.67}$$

Turning it the other way round we can reconstruct those functions from particles by interpolation schemes. This is possible if the Lagrangian (= initial) coordinates $\mathbf{q}_i$ of each particle are known. In simulations these are typically encoded into the id of each particle. Therefore it is already possible to reconstruct the dark matter sheet by interpolation from the outputs of typical N-body simulations. We illustrate this in Figure 2.10 by the dashed line which is reconstructed by (here linear) interpolation of the red particles.

Further we show in Figure 2.11 the density field that can be inferred from a two dimensional particle distribution when using the phase space interpolation. With the phase space interpolation it is possible to infer a density estimate that does not show discreteness effects. If a simulation is evolved with the N-body density estimate (left panel of that Figure) the discreteness effects in the density field will evolve into artificial structures. However they are not actually contained in the density estimate of the continuum limit. The interpolated density estimate (right panel) gets very close to the continuum limit (we will discuss more quantitatively in chapter 4). Therefore simulations which use this density estimate do not fragment into artificial haloes.

Since the first mentioning of the possibility of phase space interpolation (Shandarin et al., 2012; Abel et al., 2012) these techniques have been used in different contexts. The one application area is to use them as a pure post-processing step of N-body simulations to infer additional information about the phase space structure of the dark matter sheet (e.g. Shandarin & Medvedev, 2014; Ramachandra & Shandarin, 2017). We will describe some of such applications in section 2.6. The other application area is to use them to infer the (close to) continuum density estimate at run-time of a simulation and thereby evolve a system that should ideally be closer to the true Vlasov-Poisson solution. In the remainder of this section we will briefly describe what lies at the heart of such approaches. That is (a) how to estimate the density field from the function $\mathbf{x}(\mathbf{q})$ and (b) what techniques can be used to reconstruct $\mathbf{x}(\mathbf{q})$ and $\mathbf{v}(\mathbf{q})$ given a finite set of tracer points $(\mathbf{q}_i, \mathbf{x}_i, \mathbf{v}_i)$. Further we will briefly discuss refinement techniques which add additional tracing points in the course of a simulation to improve the interpolation accuracy in critical regions.

2.4.2 Density estimates

The exact density can be evaluated from a phase space distribution as

$$\rho(\mathbf{x}) = m_x \int_{\mathbb{R}^3} f(\mathbf{x}, \mathbf{v}) d^3v \tag{2.68}$$

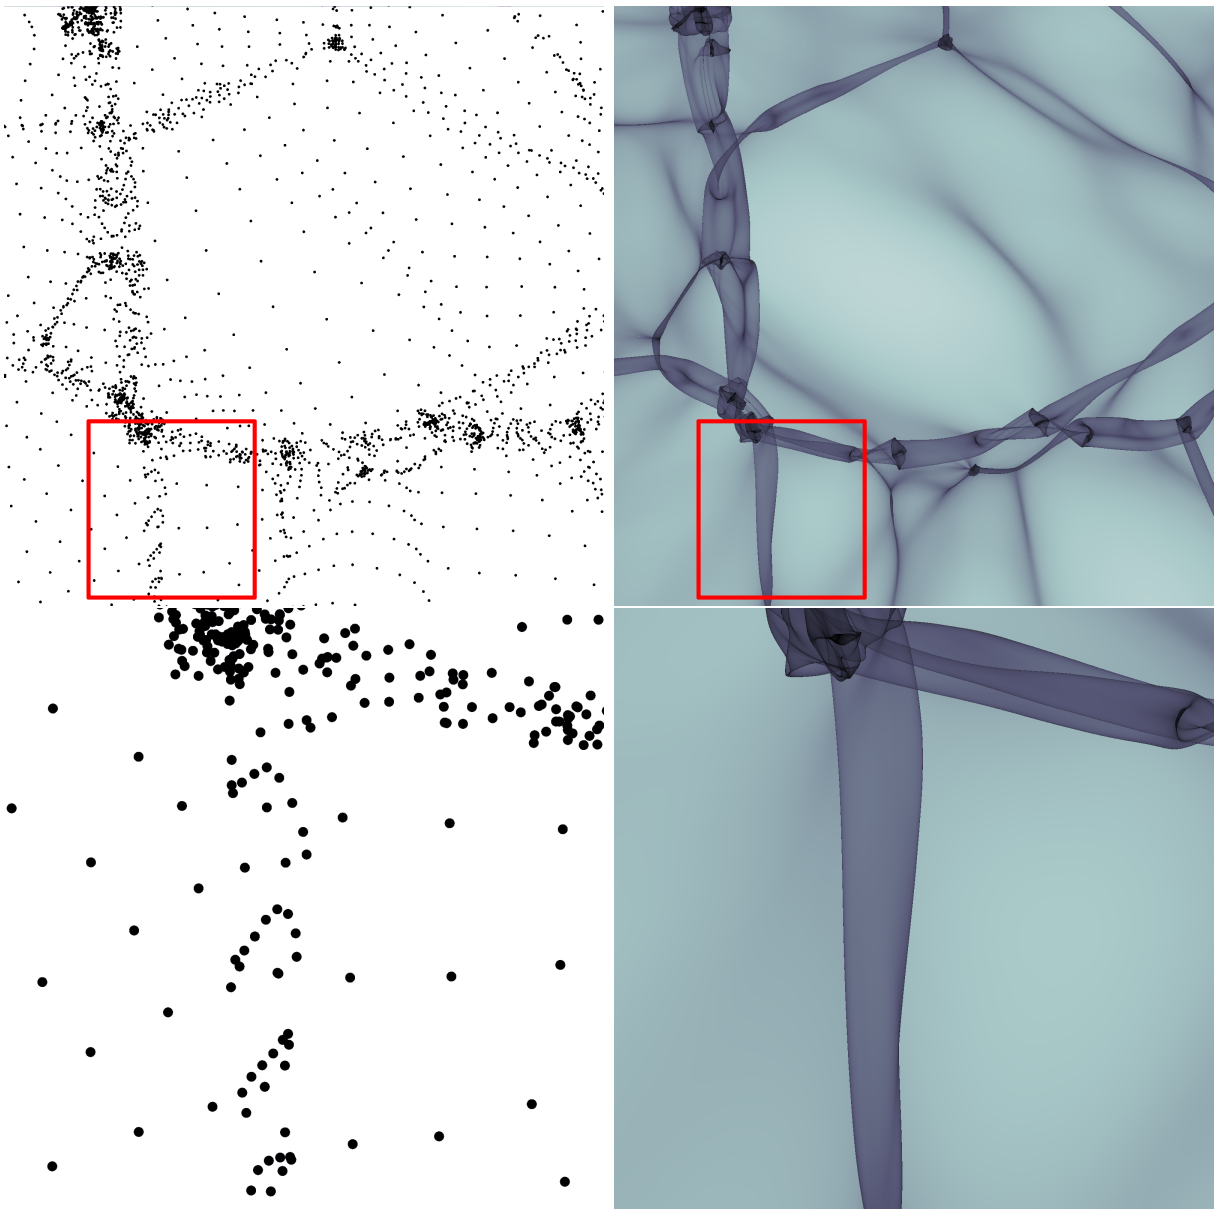

Figure 2.11: N-body density estimate (left) versus a density estimate inferred by interpolation of the dark matter sheet in phase space (right). The top row shows a wider region whereas the bottom row shows a zoom into a filamentary region. The N-body density estimate shows regular lumps which will grow into fragmented artificial halos when evolved as N-body simulation. However, the continuum density estimate shows no such artefacts. These Figures have been created with the "sheet2d"-code that I have worked on in collaboration with Oliver Hahn.

This corresponds to a projection along the velocity dimensions. Due to the discrete δ-function-like nature of the phase space distribution this integral collapses into a simple summation over all streams that share the same location:

$$\rho(\mathbf{x}) = \sum_{\text{streams}(\mathbf{x})} \rho_s(\mathbf{q}) \tag{2.69}$$

$$\text{streams}(\mathbf{x}) := \left\{\text{all } \mathbf{q} \in R^3 \text{ with } \mathbf{x}(\mathbf{q}) = \mathbf{x}\right\} \tag{2.70}$$

$$\rho_\text{s}(\mathbf{q}) = \left|\det\left(\frac{\partial \mathbf{x}}{\partial \mathbf{q}}\right)\right|^{-1} \tag{2.71}$$

We illustrate how this "true" density estimate works in Figure 2.12. The upper left panel shows the phase space distribution and the bottom left panel shows the real space density $\rho(\mathbf{x})$ which is the projection along the velocity axis. Note that the density becomes infinite at the places of caustics where a stream has $\det\left(\frac{\partial \mathbf{x}}{\partial \mathbf{q}}\right) = 0$. We marked such caustics by dashed lines. Note that these singularities are probably not problematic for cosmological simulations since they contain an infinitesimal amount of mass. In physical reality their densities are only limited by the small (but unknown) primordial velocity dispersion of dark matter. Another important effect is that the number of streams changes discretely at those caustics. This leads to a step function like behaviour of the density field. This sharp edges and the steps can also be seen in the 2d density field of the right panels of Figure 2.11. The density can easily jump by a factor larger than three at such edges.

The right panels of Figure 2.12 show how the situation looks in Lagrangian space. The upper right panel simply shows the function $\mathbf{x}(\mathbf{q})$ which is continuous and perfectly well defined as well as its derivative. Caustics appear where the derivative is zero, therefore the stream density ρ_s (shown in the bottom right panel) going to infinity. We also show the Eulerian space density as a function of Lagrangian coordinate $\rho(\mathbf{x}(\mathbf{q}))$. Looking at this quantity in Lagrangian space can be quite confusing. Just note that the two densities agree exactly in regions with one stream, whereas they are different for particles which are in multi-stream regions. In regions with many streams the Eulerian density $\rho(\mathbf{x}(\mathbf{q}))$ is typically much higher than individual stream densities $\rho_\text{s}(\mathbf{q})$. However, close to caustics the contribution from one stream can dominate the local density $\rho_s \sim \rho$ even in regions with a large number of streams. This has fact has lead Vogelsberger et al. (2008); Vogelsberger & White (2011) to an investigation of the dark matter self-annihilation signal that could potentially be created from these caustics.

Determining the exact density at a location $\mathbf{x}$ requires to invert the Lagrangian mapping $\mathbf{q} \to \mathbf{x}$ – that is finding all Lagrangian coordinates that map to the same location $\mathbf{x}$. That is a complicated mathematical problem. However, this is not necessary nor is it desirable to do in the context of cosmological simulations. In what exact form a density estimate is desirable depends on the discretization that is used to solve Poisson's equation. So far all implementations of sheet schemes have used a mesh-based Poisson solver. That is they map the mass-distribution onto a mesh and then use Fourier-techniques to solve Poisson's equation. Therefore it is not really of interest to estimate the density in a single point, but

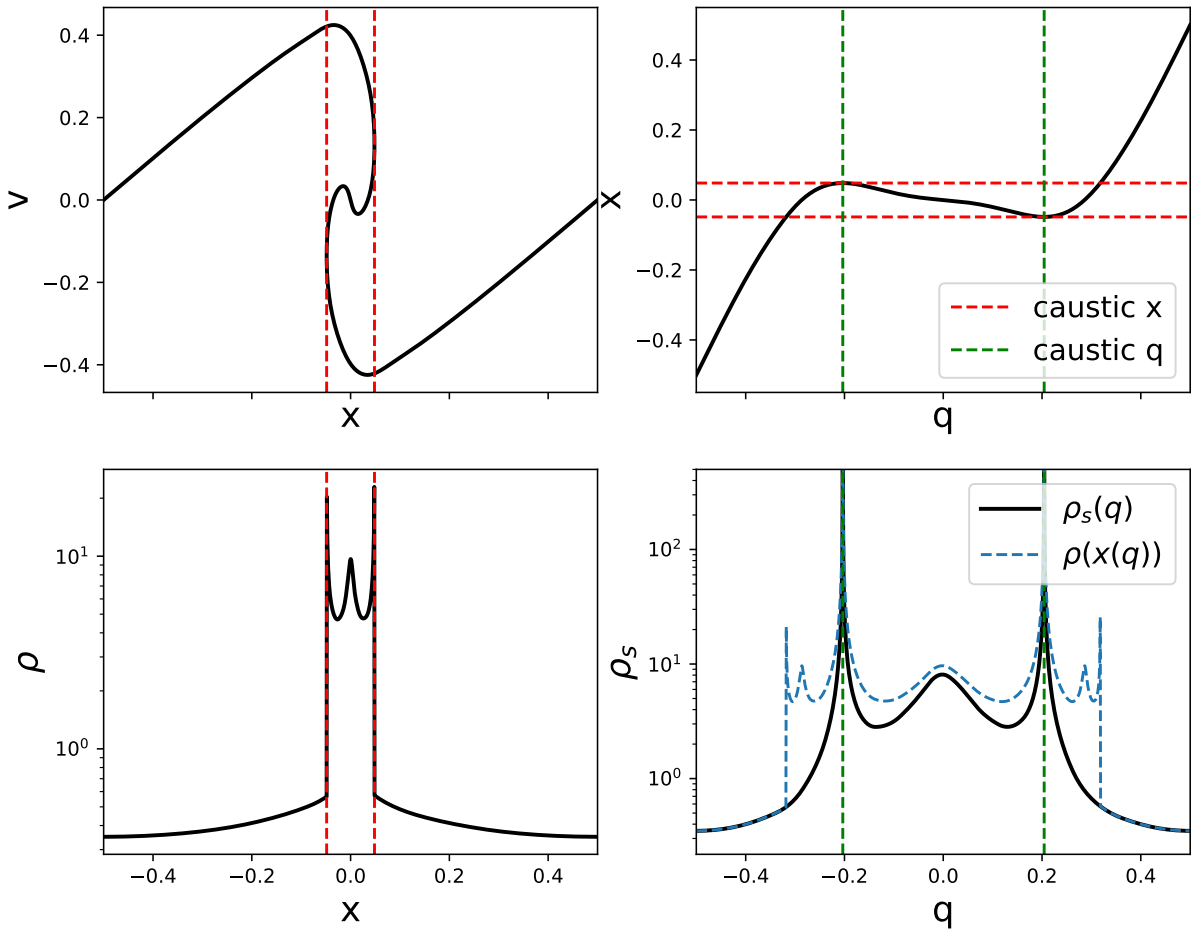

Figure 2.12: Illustration of the continuum density field in a two dimensional phase space. Top left: phase space distribution. Bottom left: the density field which is simply the projection of the phase space densities along the velocity coordinate. Note that infinite densities appear at caustics which we have marked as dashed lines. Further at caustics the density changes in a step function behaviour, because the number of streams changes (here) from one to three. Top right: the function $x(q)$ with the x and q coordinates of the caustics marked. Bottom right: the stream density ρ_s as a function of the Lagrangian coordinate $\mathbf{q}$. Unlike its Eulerian counterpart this function does not have discontinuities (but still singularities), because $\frac{\partial x}{\partial q}$ has none. Also you can find as a blue dashed line the Eulerian space density here which is only identical in single-stream regions.

rather something like the mean density of a cell, e.g.

$$\rho_{\text{cell}} = \frac{1}{V_{\text{cell}}} \int_{\mathbf{x}\in \text{ cell}} \rho(\mathbf{x}) \mathrm{d}^3 x \tag{2.72}$$

(The details of this can vary depending on whether a kernel is used in the assignment.) Also it is not really of interest to determine this for a single cell, but instead for the whole space at once. Therefore it is practical to map the whole Lagrangian space into Eulerian space and determine the mass contribution of each Lagrangian element to each cell it intersects. There are two approaches which have become established here.

The first approach, which is elaborated in Powell & Abel (2015) and used by Sousbie & Colombi (2016), is to construct from the map $\mathbf{x}(\mathbf{q})$ a set of simplices (tetrahedra) and then determine the mass in each cell by their intersections with the cells and an additional interpolation of the internal density distribution. This approach has the advantage that it is quite accurate and always produces a continuous density estimate (i.e. it is impossible to ever get an empty cell). The disadvantage is, that there is an additional discretization step into tetrahedra that can compromise the quality of the final density field if a too small number of tetrahaeda is chosen.

Another approach is to create a large number of pseudo-particles that are then deposited onto the mesh using e.g. cloud-in-cell binning. The main advantage of this approach is that it is very simple to implement and does not require additional consideration of what tetrahedra resolution is required in each Lagrangian region. The disadvantage is that it can create quite noisy density fields when too small a number of pseudo-particles (defined by a deposit level) is chosen. In this thesis and the modified HA16 scheme we always use this approach unless otherwise stated. We always make sure that enough particles are deposited so that the density field is essentially continuous and noise-free.

2.4.3 Interpolation schemes

There is a large variety of different interpolation techniques available that allow reconstruction of the mapping $\mathbf{q} \to \mathbf{x}$. Here we will only discuss interpolation techniques that assume that the initial particle distribution is sampled on a grid in Lagrangian space (or adaptive grid in the next section). However, also for other cases – for example if the Lagrangian positions are chosen from glass-like initial conditions – interpolation techniques can be used to reconstruct this mapping. These then would have to be meshless interpolation techniques – examples of such would be Voronoi-interpolation or radial basis functions.

In this thesis we use the scheme of Hahn & Angulo (2016) and make several additions to the code. Therefore we will focus on their interpolation scheme here which is based on tri-quadratic splines. In this scheme one groups particles in Lagrangian space into 3x3x3 patches (Lagrangian elements) to define a three dimensional interpolation of quadratic order. Particles on the boundary of each element also appear in neighbouring elements. If one requires that each point in Lagrangian space can only be contained by one patch, this leads to a total number of patches which is one patch per 2x2x2 particles. In Figure

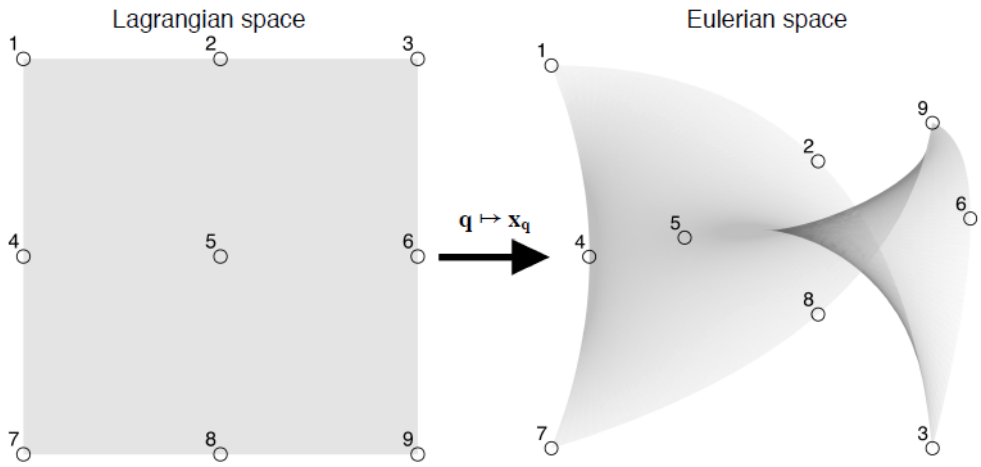

Figure 2.13: A figure from Hahn & Angulo (2016). Illustration of the mapping $\mathbf{q} \to \mathbf{x}$ defined from a single bi-quadratic element. The element is defined by 3x3 particles.

2.13, which is from Hahn & Angulo (2016), a bi-quadratic element is illustrated for the two-dimensional case.

When I started working on my project there was already a full working implementation of tri-quadratic interpolation in the code from Hahn & Angulo (2016). However, I found that their code could be simplified and speeded up by using a more generalized description of the interpolation than they had previously. I will briefly describe how to set up a d -dimensional interpolation of a function f of order n from $(n+1)^d$ given points on a grid. Tri-quadratic interpolation would be a specialization of this for $d = 3$ and $n = 2$. The mapping $\mathbf{q} \to \mathbf{x}$ is then given by three such functions, i.e. one needs to repeat the process for $f = x_1(\mathbf{q})$, $f = x_2(\mathbf{q})$ and $f = x_3(\mathbf{q})$.

General n-th order interpolation

I will start by writing the n-th order interpolation for a *three* dimensional (Lagrangian) space here, and then later point out where it has to be generalized to get to arbitrary dimensions d. We are given $(n+1)^3$ function values $f_{ijk} = f(\mathbf{q}_{ijk})$ where the sampling points $\mathbf{q}_{ijk}$ are located on a grid

$$q_{ijk} = \begin{pmatrix} i/n \\ j/n \\ k/n \end{pmatrix} \tag{2.73}$$

where n is the order of the interpolation and each of the indices goes from 0 up to n. We normalized the input space of the interpolation to $\mathbf{q} \in [0,1]^3$ (without loss of generality).

We can generally write the n-th order 3d polynomial that is to be used for the interpo-

lation as

$$f(\mathbf{q}) = \sum_{i,j,k=0}^{n} A_{ijk} q_1^i q_2^j q_3^k \tag{2.74}$$

$$=: A_{ijk} q_1^i q_2^j q_3^k \tag{2.75}$$

where the sum goes over three dimensions. As indicated in the second line, we will skip writing summations in the remaining equations in this section and instead use as convention that every index that appears twice in any term is summed over. It is the goal to determine the interpolation coefficients A_{ijk} by using the given function values f_{ijk}. By evaluating (2.75) on the grid from equation (2.73) we can write those function values as

$$f_{abc} = A_{ijk} \left(\frac{a}{n}\right)^i \left(\frac{b}{n}\right)^j \left(\frac{c}{n}\right)^k \tag{2.76}$$

where $0^0 = 1$ has to be used wherever that appears. The question here is how to invert the equation so that we get A_{ijk}. One way of viewing this is the following: we have an $(n+1)\times(n+1)\times(n+1)$ tensor on the left side, and a product of the $(n+1)\times(n+1)\times(n+1)$ tensor A and another $((n+1)\times(n+1)\times(n+1))^2$ tensor on the right side. The idea is to somehow multiply by the inverse of that second tensor. It is however, quite confusing to think of the inverse of a $(n+1)^6$ dimensional tensor. We can simplify this process by pointing out that on the right side of (2.76) three times the same sub-tensor appears:

$$f_{abc} = A_{ijk} B_{ai} B_{bj} B_{ck} \tag{2.77}$$

$$B_{lm} = \left(\frac{l}{n}\right)^m \tag{2.78}$$

Therefore we can simply multiply the whole tensor equation by the inverses of $\mathbf{B}$ which we will call $\mathbf{M}$:

$$\mathbf{M} := \mathbf{B}^{-1} \tag{2.79}$$

$$f_{abc} M_{ap} M_{bq} M_{cr} = A_{ijk} B_{ai} B_{bj} B_{ck} M_{ap} M_{bq} M_{cr} \tag{2.80}$$

$$= A_{ijk} \delta_{ip} \delta_{jq} \delta_{kr} \tag{2.81}$$

where δ denotes a unit matrix here. Then after renaming some indices we have

$$A_{ijk} = f_{abc} M_{ai} M_{bj} M_{ck} \tag{2.82}$$

The generalization of this to d dimensions is straight-forward. In that case we have

$$f(\mathbf{q}) = A_{i_1 \ldots i_d} q_1^{i_1} \ldots q_d^{i_d} \tag{2.83}$$

$$A_{i_1 \ldots i_d} = f_{a_1 \ldots a_d} M_{a_1 i_1} \ldots M_{a_d i_d} \tag{2.84}$$

The matrix $\mathbf{M}$ only depends on the order n of the interpolation, but not on the dimension d. This is an expression of the fact that a d-dimensional interpolation operation can be seen as a chain of d one-dimensional interpolation operations where the one-dimensional interpolation is described by M. For future reference we list the matrices $\mathbf{M}$ up to order 4 in Table 2.1. Note that we usually only use the tri-quadratic case $d = 3, n = 2$ throughout this thesis. However, the code also has the cases $n = 1$ and $n = 4$ available. The original version of the HA16 code used a different way of calculating the interpolation which hard-coded most of the operations (which was however at some point generated with mathematica or so). Using the generalization I could replace of order 5000 lines of code by just a handful of lines - which decreased the compile time by a minute or so.

Evaluating the polynomial (2.83) at arbitrary locations requires of order $(n+1)^3$ operations (for $d = 3$). However, in the HA16 scheme it is often needed to evaluate the polynomial on a grid. In that case there is a trick one can use when evaluating the polynomial. The straight forward way of evaluating on a grid of m^3 locations would be to evaluate equation (2.83) m^3 times which would require of order $(n+1)^3 \cdot m^3$ operations. However, in the grid case one can also rewrite the operation as

$$f_{ijk} = A_{abc} q^a_{1,i} q^b_{2,j} q^c_{3,k} \tag{2.85}$$

$$= ((A_{abc} q^a_{1,i}) q^b_{2,j}) q^c_{3,k} \tag{2.86}$$

where the first line indicates a way of evaluation where it is summed over a, b and c for each i, j, k. However, one can also restructure the summation like in the second line, denoting an $(n+1)^3$ tensor which is multiplied by three different $(n+1, m)$ tensors. This requires of order $m^3(n+1)$ operations in the case that $m \gg n$. Using this more efficient summation speeded up evaluation by far more than a factor of two (at the negligible cost of slightly higher memory requirement during evaluation).

Some notes on interpolation

While the tri-quadratic interpolation is certainly an improvement over a tri-linear interpolation and also over tesselation techniques, it is still not quite optimal. Actually a nicer way of defining an interpolation would be to use an interpolation technique which ensures continuity of the function and its derivatives. In the tri-quadratic case the derivatives can be discontinuous at the boundaries of each element. Better would be a tri-cubic interpolation which is defined on elements of 2^3 particles, but constrains derivatives from neighbouring elements and smoothness requirements. However, we did not decide to make any such changes, since the details of the interpolation seem to matter relatively little when facing the problem of the enormous complexity of the dark matter sheet as we shall discuss in chapter 4.

Trigonometric Interpolation

Another interesting possibility as a choice for the interpolation scheme is trigonometric interpolation. Trigonometric interpolation refers to an interpolation which writes a function

n	$B_{lm} = \left(\frac{l}{n}\right)^m$	$M = B^{-1}$
1	$\begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 \\ -1 & 1 \end{pmatrix}$
2	$\begin{pmatrix} 1 & 0 & 0 \\ 1 & 1/2 & 1/4 \\ 1 & 1 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 \\ -3 & 4 & 1 \\ 2 & -4 & 2 \end{pmatrix}$
3	$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 1 & 1/3 & 1/9 & 1/27 \\ 1 & 2/3 & 4/9 & 8/27 \\ 1 & 1 & 1 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 & 0 \\ -5.5 & 9. & -4.5 & 1. \\ 9 & 22.5 & 18. & -4.5 \\ -4.5 & 13.5 & -13.5 & 4.5 \end{pmatrix}$
4	$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 1 & 1/4 & 1/16 & 1/64 & 1/256 \\ 1 & 1/2 & 1/4 & 1/8 & 1/16 \\ 1 & 3/4 & 9/16 & 27/64 & 81/256 \\ 1 & 1 & 1 & 1 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ -25/3 & 16 & -12 & 16/3 & -1 \\ 70/3 & -208/3 & 76 & -112/3 & 22/3 \\ -80/3 & 96 & -128 & 224/3 & -16 \\ 32/3 & -128/3 & 64 & -128/3 & 32/3 \end{pmatrix}$

Table 2.1: Matrices that are used to set up a general interpolation. n (left column) is the order of the interpolation and $M = B^{-1}$ (right column) is the matrix that is used to compute the coefficients as in (2.84). B is given in the central column where $B_{lm} = \left(\frac{l}{n}\right)^m$.

as a sum of sinus and cosinus functions. Something along the lines

$$f(\mathbf{q}) = \sum_{klm} A_{klm} \cos(kq_1 + lq_2 + mq_3) + B_{klm} \sin(kq_1 + lq_2 + mq_3) \tag{2.87}$$

Note that this is more or less exactly the basis that the discrete Fourier transformation assumes for a function. Therefore we can use the Fourier transformation as a short-cut to evaluate the interpolation. Note that this basis assumes that the function is periodic. This is not the case for $\mathbf{x}(\mathbf{q})$, but it is naturally the case for the displacement field $\mathbf{s}(\mathbf{q}) = \mathbf{x}(\mathbf{q}) - \mathbf{q}$ of most cosmological simulations.

Let's assume we have a regular Lagrangian grid of N^3 function values on which we have given a periodic function f. If we do a Fourier transformation we obtain something like the coefficients A_{klm} and B_{klm}. Given those we could in principle evaluate the interpolation at arbitrary locations - each evaluation would be a sum over N^3 sine and cosine pairs. However, if we were to evaluate the interpolation on a grid instead, we can actually use the inverse Fourier transformation to obtain a large number of evaluated locations very efficiently. I have developed the following scheme to create a large number of pseudo-particles to infer a density estimate (by binning them). I have implemented this in a small, but highly efficient mpi code which uses the fftw3 library.

We compute the displacement field

$$\mathbf{s}(\mathbf{q}) = \mathbf{x}(\mathbf{q}) - \mathbf{q} \tag{2.88}$$

which is a periodic function in most cosmological simulations due to the periodic boundary conditions. Note that in practice one has to undo the periodic wrapping of the position

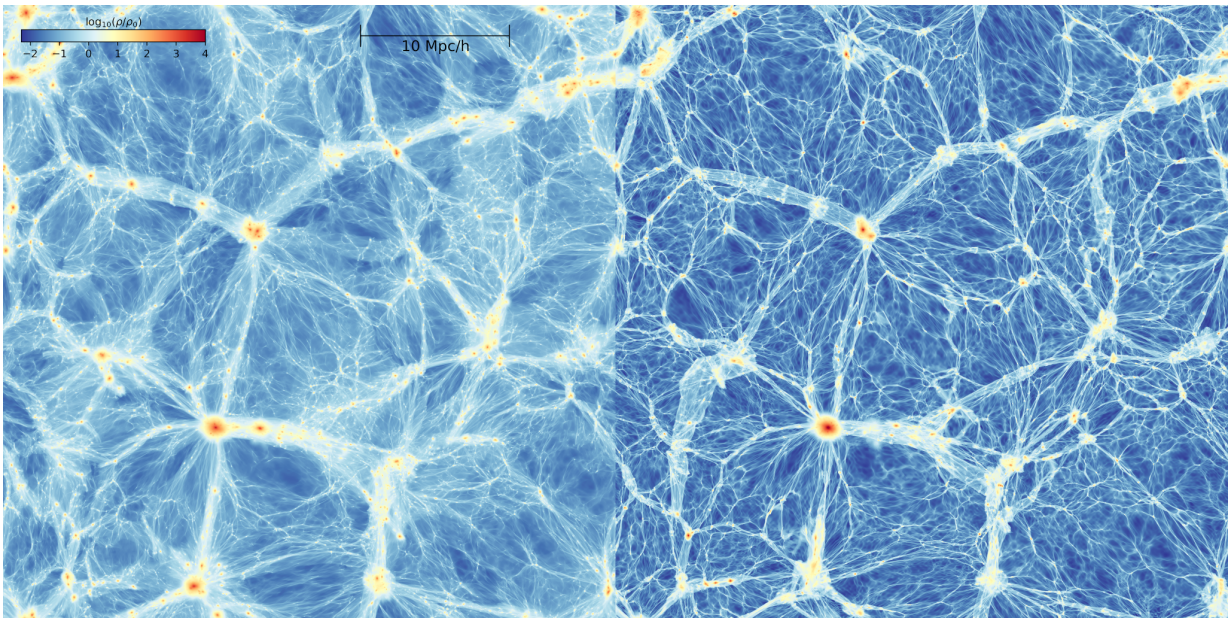

Figure 2.14: Density field that we obtained from trigonometric resampling of a CDM simulation. Left: a projection of width $2\,\mathrm{Mpc}/h$, right: a razor thin slice of $20\,\mathrm{kpc}/h$. By using interpolation techniques a lot of the fine-grained features of the density field become visible. As such one can see the sharp edges in the density field also known as caustics. We discuss these phenomena in more detail in section 2.6.

function e.g. by wrapping $\mathbf{s}$ so that all components are in the range $-L/2 \ll s_i \ll L/2$ where L is the size of the box.

As a periodic function the displacement field can be re-sampled to a higher resolution grid by trigonometric interpolation in the following way. One transforms the grid $\mathbf{s}$ into Fourier space $\mathbf{s}_k$, zero pads it to a higher resolution grid $\mathbf{s}_{k,\mathrm{hr}}$, so that all the low frequency components are kept the same and originally unresolved higher frequency components are set to zero, and then transforms the higher resolution Fourier grid back to real-space to obtain the re-sampled displacement field $\mathbf{s}_{\mathrm{hr}}$. We call this method *trigonometric re-sampling* here. The re-sampled displacement field can then be binned (adding $\mathbf{q}$ to obtain positions) onto a mesh to obtain a density field.

While this is already an efficient and simple algorithm, it is still a bit impractical in terms of memory consumption, since all resampled particles have to be present in memory at the same time. Therefore as a more memory efficient version of the same scheme, instead of zero padding the Fourier transformed low resolution grid $\mathbf{s}_k$, we shift it by a small displacement $\Delta\mathbf{q}$ in Fourier space

$$\mathbf{s}_{k,\mathrm{shift}} = \exp(i\mathbf{k}\cdot\Delta\mathbf{q})\mathbf{s}_k \tag{2.89}$$

transform it back to real space $\mathbf{s}_{\mathrm{shift}}$. Then to recover the position field we have to add the shifted Lagrangian coordinate

$$\mathbf{x}_{\mathrm{shift}} = \mathbf{s}_{\mathrm{shift}} + (\mathbf{q} - \Delta\mathbf{q}) \tag{2.90}$$

Now one does this procedure repeatedly for several different shifts and each time simply bins the positions $\mathbf{x}_{\mathrm{shift}}$ and discards them afterwards. For example to resample from a 2-dimensional grid with N^2 grid cells to a higher resolution grid of $(2N)^2$ grid cells, one applies the 4 shifts

$$\Delta\mathbf{q}_{1,2,3,4} = \begin{pmatrix} \pm\Delta g/4 \\ \pm\Delta g/4 \end{pmatrix} \tag{2.91}$$

where $\Delta g = L/N$ is the gridspacing, L is the boxsize and N the number of grid points per dimension.

One might be worried that the Fourier interpolation might create artefacts in the density field due to Gibb's phenomena. To benchmark the scheme we test it in a two dimensional setup. From a cosmological simulation we select a 2d slice in Lagrangian space (that means all particles that are in one plane in the initial conditions) and project it into Eulerian space with the trigonometric resampling. Further we do the same with a bicubic interpolation as a reference case. Note that - unlike tricubic interpolation - bicubic interpolation is availble as a simple `python` function in `scipy.interpolate.RectBivariateSpline`. We show a zoom into a part of the result in Figure 2.15. Against justified worries, we cannot see any artefacts due to a possible Gibb's phenomenon. It seems that the trigonometric interpolation performs equally well to a cubic interpolation here.

In Figure 2.14 we show the density field that we have obtained from an $N = 1024^3, L = 40\,\mathrm{Mpc}/h$ CDM simulation. With the interpolation scheme we are even able to create a

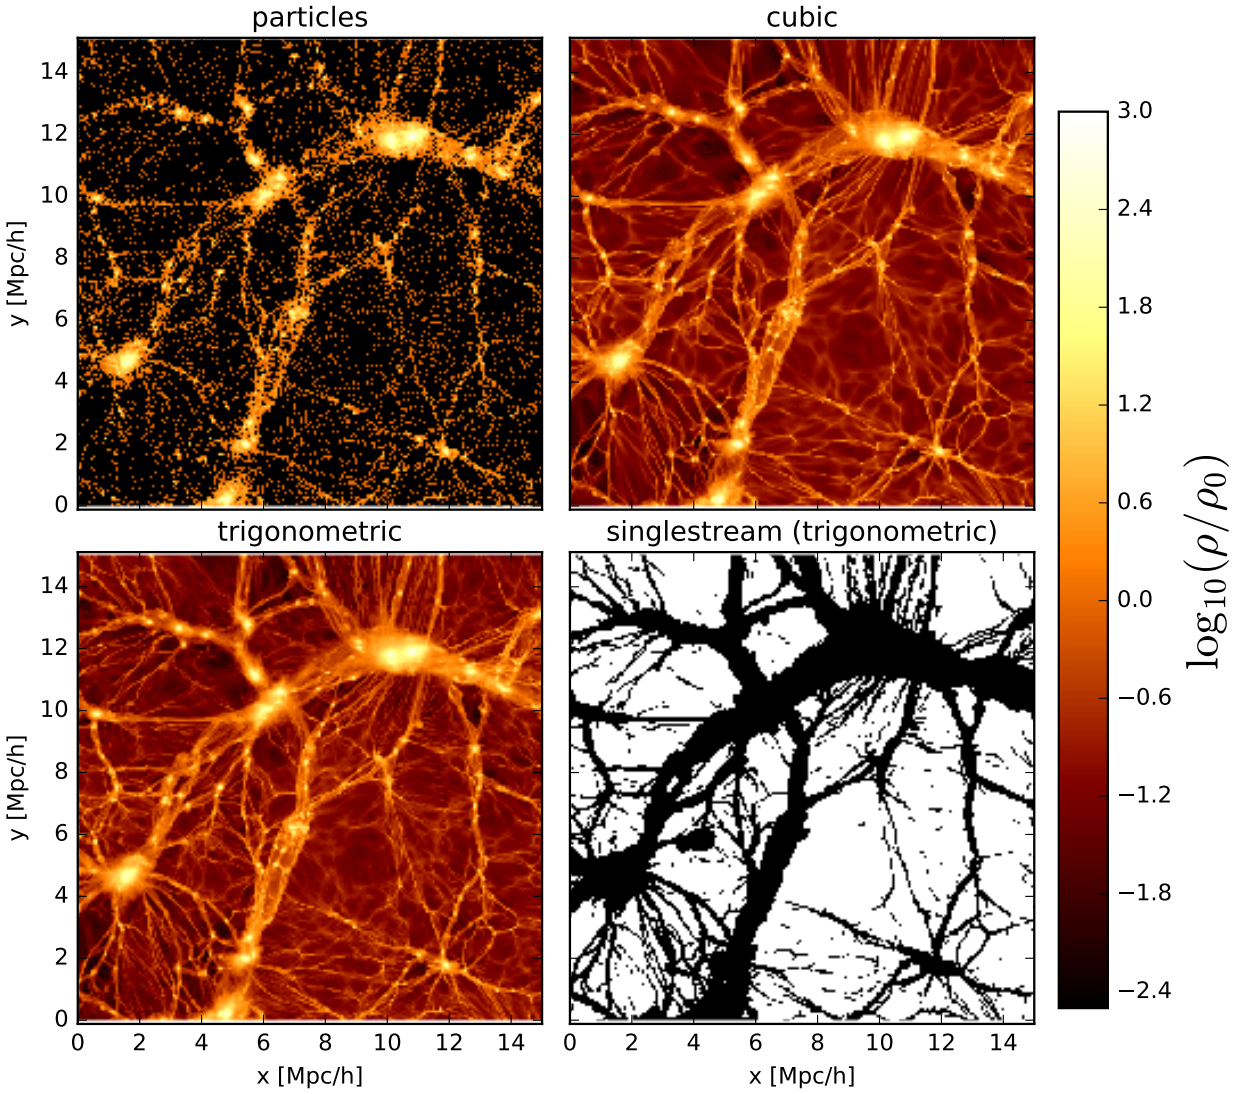

Figure 2.15: Benchmark of the trigonometric interpolation by projecting a Lagrangian 2D slice from a simulation into Eulerian space. Top Left: density field inferred from binning the particles. Top Right: resampling using third order interpolation. Bottom Left: resampling using trigonometric interpolation. Bottom Right: single-stream regions (white) as inferred by using the trigonometric interpolation. We explain in section 2.6.3 how the single-stream classification was inferred and in more detail what this notion means. The cubic and trigonometric interpolation agree very well. The identified single-stream regions match with the ones that one would determine by eye from the density field.

razor thin slice with a width of 20 kpc. This allows us to see many of the fine-grained features that typically remain hidden when doing thick projections of simulations. Most prominently these are the sharp edges in the density field, also known as caustics. We will discuss these phenomena in more detail in section 2.6. These Figures were created by resampling $32768^3 \sim 10^{13}$ pseudo-particles to a 2048^3 grid. This number is quite redundant to infer this density field, but with the mpi/fftw code it was possible in a relatively short amount of time.

2.4.4 Adaptive refinement

The complexity of the dark matter sheet grows over time. For example, in one dimension the phase space spiral (as for example in Figure 2.10) winds up over time. Therefore, if a constant number of interpolating particles is used, the interpolation works initially, but breaks down in later stages, since the functions $\mathbf{x}(\mathbf{q})$ and $\mathbf{v}(\mathbf{q})$ vary faster than can traced from the finite set of particles. We will analyze this in more detail in chapter 4.

To deal with this problem Hahn & Angulo (2016) and Sousbie & Colombi (2016) have developed adaptive refinement schemes.

In these schemes a refinement criterion detects Lagrangian elements which are about to become too complex to be traced. For example Hahn & Angulo (2016) use estimates of force errors and consistency requirements between interpolation orders as refinement criteria whereas Sousbie & Colombi (2016) use Poincaré invariants. When an element is flagged for refinement, additional particles are created from the interpolated functions $\mathbf{x}(\mathbf{q})$ and $\mathbf{v}(\mathbf{q})$. For example in Hahn & Angulo (2016) from a 3x3x3 element with Lagrangian width Δq a set of 8 new 3x3x3 elements with half the size $\Delta q/2$ is created. Effectively this technique can be called adaptive mesh refinement in Lagrangian space.

These schemes offer a completely independent alternative to N-body schemes and in principle can be brought to a complete "fine-grained" convergence if refinement is allowed to arbitrary levels. However, Sousbie & Colombi (2016) find that the number of required resolution elements increases very dramatically with time. We will discuss this problem in more detail in chapter 4 where we propose a "release" of elements in critical regions as an alternative.

2.5 The geodesic deviation equation

Sheet-based schemes only work properly when the functions $\mathbf{x}(\mathbf{q})$ and $\mathbf{v}(\mathbf{q})$ are reconstructed well enough. In chapter 4 we will discuss more thoroughly when and where this is the case. However, if one is interested in properties of the fine-grained phase space distribution in a regime where an interpolation-reconstruction is not possible, the geodesic deviation equation offers a viable tool. With the geodesic deviation equation one can follow the evolution of the dark matter sheet in phase space in a statistical manner. The contents of this section are largely based on the concepts presented in White & Vogelsberger (2009)

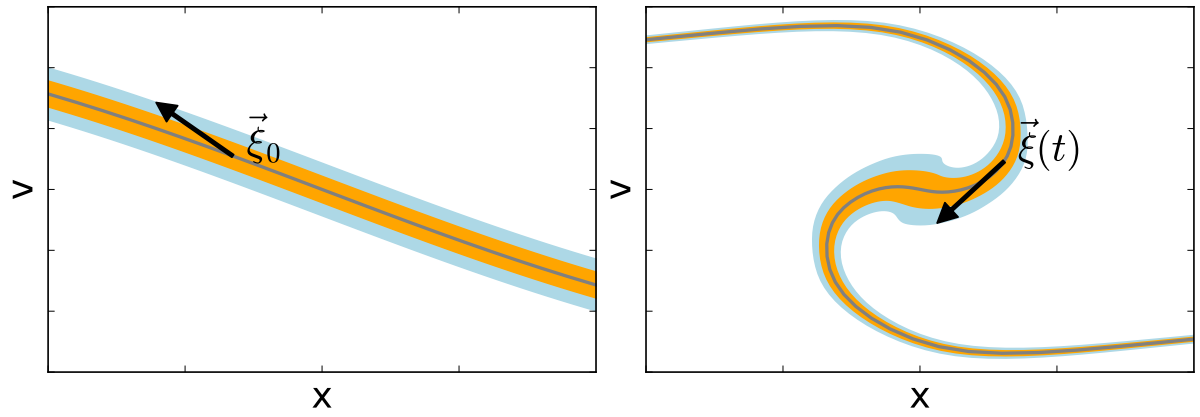

Figure 2.16: Example for the evolution of a small displacement in phase space with time. On the left side the displacement in some initial phase space and on the right at a later time. One can imagine the small displacement $\boldsymbol{\xi}$ always pointing from a particle to a close-by neighbouring particle.

and Vogelsberger & White (2011). However, I will explain these with my own words and figures.

The geodesic deviation equation (GDE) answers the question: how does a small displacement in phase space

$$\boldsymbol{\xi}(t) := \begin{pmatrix} \Delta\mathbf{x} \\ \Delta\mathbf{v} \end{pmatrix} \tag{2.92}$$

evolve with time? We illustrate an example evolution of such a displacement in Figure (2.16).

2.5.1 The distortion tensor

If the displacement is taken to be infinitesimally small it can be parameterized by the distortion tensor

$$\boldsymbol{\xi}(t) = \mathbf{D}(t)\boldsymbol{\xi}_0 \tag{2.93}$$

$$\mathbf{D} = \begin{pmatrix} \frac{\partial\mathbf{x}}{\partial\mathbf{q}} & \frac{\partial\mathbf{x}}{\partial\mathbf{p}} \\ \frac{\partial\mathbf{v}}{\partial\mathbf{q}} & \frac{\partial\mathbf{v}}{\partial\mathbf{p}} \end{pmatrix} \tag{2.94}$$

$$=: \begin{pmatrix} \mathbf{D}_{\mathrm{xq}} & \mathbf{D}_{\mathrm{xp}} \\ \mathbf{D}_{\mathrm{vq}} & \mathbf{D}_{\mathrm{vp}} \end{pmatrix} \tag{2.95}$$

where we have labelled the initial phase space coordinates as $\mathbf{q}$ and $\mathbf{p}$, the final coordinates as $\mathbf{x}$ and $\mathbf{v}$ and labeled all 3×3 sub tensors of the distortion tensor. Note that we labelled the spatial initial coordinate $\mathbf{q}$ in the same way we labelled the Lagrangian coordinate in previous sections. These two are identical if one defines the initial frame at time $a \to 0$. That is, if one sets $\mathbf{D}(a \to 0) = \mathbb{1}$. It is not necessary to make specifically this choice, but

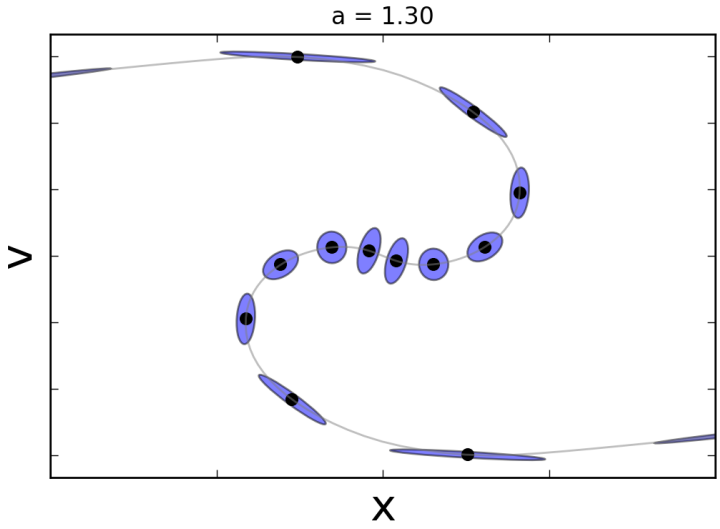

Figure 2.17: An illustration of the phase space volume elements that can be traced by the distortion tensor. It is possible to reconstruct the space tangential and the space orthogonal to the dark matter sheet from the distortion tensor. The phase space distribution can be reconstructed locally around each particle.

this one turns out to be very convenient, since the $\mathbf{D}_{\mathrm{xq}}$ and the $\mathbf{D}_{\mathrm{vq}}$ part of the distortion tensor correspond exactly to the derivatives of the dark matter sheet. For simplicity, but without loss of generality, we will always use this frame in this thesis (unless stated otherwise).

Since the displacement $\boldsymbol{\xi}_0$ can be taken in any direction, the distortion tensor effectively describes an infinitesimal volume element around each particle. If we think of the initial volume element as a small sphere, the distortion tensor describes how it gets stretched or squeezed and rotated or mirrored into some ellipsoidal shape at the final time. We illustrate such phase space ellipsoids in Figure 2.17. Note that the deformations tangential to the sheet are always given by $\mathbf{D}_{\mathrm{xq}}$ and $\mathbf{D}_{\mathrm{vq}}$. Intuitively one might think that the deformations "orthogonal" are then always given by $\mathbf{D}_{\mathrm{xp}}$ and $\mathbf{D}_{\mathrm{vp}}$. While this is correct in some sense, it is also wrong in some sense, since the concept "orthogonal" is not very well defined in phase space. I will discuss this topic later in chapter 6 where I will also discuss how to make a plot like Figure 2.17. Just note for now, that that Figure does not directly plot $\mathbf{D}_{\mathrm{xp}}$ and $\mathbf{D}_{\mathrm{vp}}$ in the orthogonal direction.

For most of this thesis we will be interested in the parts of the distortion tensor which describe displacements along the sheet $\mathbf{D}_{\mathrm{xq}}$ and $\mathbf{D}_{\mathrm{vq}}$. As already discussed, the primordial velocity dispersion of dark matter can be mostly neglected inside dark matter simulations. However, it becomes relevant when asking questions like e.g. what is the width of caustics.

2.5.2 Evolution equations

Recall from earlier (equations (2.45) and (2.46)) that the equations of motion of a test particle in the comoving frame read

$$\dot{\mathbf{x}} = \mathbf{v}a^{-2} \tag{2.96}$$
$$\dot{\mathbf{v}} = -\boldsymbol{\nabla}\phi(\mathbf{x})a^{-1} \tag{2.97}$$

We can derive the geodesic deviation equation by differentiating these with respect to the Lagrangian coordinate $\mathbf{q}$:

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{\partial \mathbf{x}}{\partial \mathbf{q}}\right) = \frac{\partial \dot{\mathbf{x}}}{\partial \mathbf{q}} \tag{2.98}$$
$$= \frac{\partial \mathbf{v}}{\partial \mathbf{q}}a^{-2} \tag{2.99}$$
$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{\partial \mathbf{v}}{\partial \mathbf{q}}\right) = \frac{\partial \dot{\mathbf{v}}}{\partial \mathbf{q}} \tag{2.100}$$
$$= -a^{-1}(\boldsymbol{\nabla}_x \otimes \boldsymbol{\nabla}_q)\phi \tag{2.101}$$
$$= -a^{-1}(\boldsymbol{\nabla}_x \otimes \boldsymbol{\nabla}_x)\phi\frac{\partial \mathbf{x}}{\partial \mathbf{q}} \tag{2.102}$$

where $\otimes$ denotes an outer product (the result being a 3×3 tensor here). To avoid confusions: in index form this can be written as $(\boldsymbol{\nabla}_x \otimes \boldsymbol{\nabla}_x)_{ij} = \frac{\partial}{\partial x_i}\frac{\partial}{\partial x_j}$. Defining the tidal tensor as $\mathbf{T} = -(\boldsymbol{\nabla}_x \otimes \boldsymbol{\nabla}_x)\phi$ and doing an analogue derivation for the $\mathbf{D}_{\mathrm{xp}}$ and $\mathbf{D}_{\mathrm{vp}}$ tensors, we can write the evolution equations as

$$\dot{\mathbf{D}}_{\mathrm{xq}} = a^{-2}\mathbf{D}_{\mathrm{vq}} \qquad \dot{\mathbf{D}}_{\mathrm{xp}} = a^{-2}\mathbf{D}_{\mathrm{vp}} \tag{2.103}$$
$$\dot{\mathbf{D}}_{\mathrm{vq}} = a^{-1}\mathbf{T}\mathbf{D}_{\mathrm{xq}} \qquad \dot{\mathbf{D}}_{\mathrm{vp}} = a^{-1}\mathbf{T}\mathbf{D}_{\mathrm{xp}} \tag{2.104}$$

also known as the geodesic deviation equation (Vogelsberger & White, 2011).

The GDE can be integrated in cosmological simulations along with the other equations of motion for each particle. Thereby one follows the tangent space of the dark matter sheet around each particle. Therefore a lot of the properties of the phase space sheet can be reconstructed in a statistical manner. Based on Mark Vogelsberger's implementation of the GDE from Vogelsberger & White (2011) I have added the GDE scheme into the code of Hahn & Angulo (2016).

Note that $\mathbf{D}_{\mathrm{xq}}$ and $\mathbf{D}_{\mathrm{vq}}$ form a completely independent system from $\mathbf{D}_{\mathrm{xp}}$ and $\mathbf{D}_{\mathrm{vp}}$. Therefore it is possible to only evolve the space tangential to the sheet if one is only interested in it. In our simulations we still follow the whole 6×6 distortion tensor, but for some applications half of the memory associated with the distortion tensor could be saved in principle. Note that the evolution equations of the left column and the right column of the distortion tensor are exactly identical if one switches the "q" in the index by a "p". Usually their only difference comes from the choice of initial conditions of the distortion

tensor - that is $\mathbf{D}_{\mathrm{xq}}(a \to 0) = \mathbf{D}_{\mathrm{vq}}(a \to 0) = \mathbb{1}$ and $\mathbf{D}_{\mathrm{vq}}(a \to 0) = \mathbf{D}_{\mathrm{xp}}(a \to 0) = 0$. The structure of the equations leads to some relations between the components of the distortion tensor.

2.5.3 Symplectic constraints

The distortion tensor is a symplectic matrix. A symplectic 6x6 matrix D has the property

$$D^T \cdot \begin{pmatrix} 0 & \mathbb{1} \\ -\mathbb{1} & 0 \end{pmatrix} \cdot D = \begin{pmatrix} 0 & \mathbb{1} \\ -\mathbb{1} & 0 \end{pmatrix} \tag{2.105}$$

This is equivalent to the following conditions on the sub-tensors

1. $\mathbf{D}_{\mathrm{xq}}^T \mathbf{D}_{\mathrm{vq}}$ is symmetric
2. $\mathbf{D}_{\mathrm{xp}}^T \mathbf{D}_{\mathrm{vp}}$ is symmetric
3. $\mathbf{D}_{\mathrm{xq}}^T \mathbf{D}_{\mathrm{vp}} - \mathbf{D}_{\mathrm{vq}}^T \mathbf{D}_{\mathrm{xp}} = \mathbb{1}$

It is a bit hard to grasp what is going on here. First let us count the number of constraints that are set by this. Requiring that a 3×3 tensor is symmetric effectively constrains 3 components of that tensor while six components are left free. Therefore the first two conditions constrain 3 components each. The last condition is an equality which constrains 3×3 components. Therefore in total these conditions constrain 15 degrees of freedom of the distortion tensor. Since it is a 6×6 tensor with 36 components this leaves in the end 21 actual degrees of freedom.

Let us shortly discuss, how one can intuitively understand this. Typically when we talk about the phase space, we are used to state that the phase space density is conserved. In one dimension this is equivalent to something like the number of particles contained within a phase space area element $\sim \Delta x \Delta v$ is constant. We then often generalize this to 3 dimensions saying that the number of particles in a 6d phase space volume element $\sim \Delta x^3 \Delta v^3$ is conserved. While this is certainly the case, this generalization neglects a large number of other quantities that are additionally conserved when increasing the number of dimensions. A different generalization of this conservation law is given by Poincaré invariants. However, these are not very intuitive by themselves. Let me simplify a bit and say that these more or less define a sense in which the number of particles that are contained on any kind of phase-space *area element* is conserved. In the infinitesimal case this manifests in the following way: every possible pairing of coordinate axes in the initial phase space is associated with a conservation law. For example, if we trace full 6d displacements in the initial phase space, there are $(6 \cdot 5)/2 = 15$ possible combinations of axes - therefore 15 constraints on the distortion tensor. As a different example, if one would only be interested in initial displacements in position space $\mathbf{q}$, but not initial velocity space $\mathbf{p}$, then we have $(3 \cdot 2)/2 = 3$ different combinations of axes (in the initial space), leading to 3 constraints on the combination of $\mathbf{D}_{\mathrm{xq}}$ and $\mathbf{D}_{\mathrm{vq}}$. This condition is expressed in the first condition from above. Interestingly that means that the two functions $\mathbf{x}(\mathbf{q})$

and $\mathbf{v}(\mathbf{q})$ are not completely independent, but have some joint symmetry requirements in their derivatives. For example, Sousbie & Colombi (2016) have used these as a refinement criterion in their code (in the form of Poincaré invariants).

Lack of symmetry in numerical force-fields

Based on the implementation of Vogelsberger & White (2011) I have added the GDE scheme into the code of Hahn & Angulo (2016). In that process I came across a problem which makes it hard to ensure full phase space consistency in the continuum limit. I will briefly discuss it here. Most readers might want to skip this subsection though, because it goes deeply into technical implementation details.

Cosmological simulation codes like for example Gadget2 (Springel, 2005a) typically use a mesh to solve Poisson's equation for the long-range components of the force. Typically this is done in the following way: (1) Density values ρ_{ijk} are measured on a mesh (infered from the particle distribution). (2) The density field is convolved with the gravitational potential kernel $1/r$ (by using Fourier techniques) to obtain the potential field ϕ_{ijk}. (3) The potential field is derived (three times) by finite-differences to obtain the force field $\mathbf{F}_{ijk}$. (4) The force grid is interpolated for each particle p to obtain the force that acts on it $\mathbf{F}(\mathbf{x}_p)$.

This way of calculating the force field is straight forward, quite efficient and for most applications accurate enough. However, I find that in an infinitesimal sense Schwartz's theorem is violated

$$\frac{\partial F_i}{\partial x_j} \neq \frac{\partial F_j}{\partial x_i} \tag{2.106}$$

Just to be clear here: it still holds to a very good approximation $\frac{\partial F_i}{\partial x_j} \approx \frac{\partial F_j}{\partial x_i}$ and also the violation decreases when decreasing the size of the grid cells. However, in the limit of infinite particles at a finite mesh resolution the system will not be exactly consistent with Hamiltonian phase space dynamics, because the force-field cannot exactly be derived from a potential. Let us first discuss why $\frac{\partial F_i}{\partial x_j} \neq \frac{\partial F_j}{\partial x_i}$ and then afterwards discuss in more detail when this might be an issue.

In step (3) from above the force field is calculated on the grid using finite differencing methods. Let us give an example in two dimensions:

$$F_{x,i,j} = -\frac{\phi_{i,j+1} - \phi_{i,j-1}}{2\Delta x} \tag{2.107}$$

$$F_{y,i,j} = -\frac{\phi_{i+1,j} - \phi_{i-1,j}}{2\Delta y} \tag{2.108}$$

Also other finite-differencing schemes are possible, but these do not change our argumentation here. When the forces for individual particles are calculated these get assigned by some interpolation scheme, for example

$$F_x(x,y) = \sum_{i,j} F_{x,i,j} \cdot A_{i,j}(x,y) \tag{2.109}$$

where $A_{i,j}(x,y)$ are kernels which define the interpolation and the sum goes over different grid points. Now if one derives the force field (in the infinitesimal sense) we have something like

$$\frac{\partial}{\partial y}F_x(x,y) = \sum_{i,j} F_{x,i,j} \cdot \frac{\partial}{\partial y}A_{i,j}(x,y) \tag{2.110}$$

$$\frac{\partial}{\partial x}F_y(x,y) = \sum_{i,j} F_{y,i,j} \cdot \frac{\partial}{\partial x}A_{i,j}(x,y) \tag{2.111}$$

that is, one needs to derive the interpolation kernel to get the infinitesimal derivative. By infinitesimal derivative I mean the difference in force that an actual particle at e.g. $x + \mathrm{d}x$ would experience. By construction we are violating Schwartz's theorem here $\frac{\partial}{\partial x}F_y \neq \frac{\partial}{\partial y}F_x$. I think it is not possible to come up with any sensible choice of kernel or finite differences scheme where $\frac{\partial}{\partial x}F_y = \frac{\partial}{\partial y}F_x$. The conceptual issue here is the following: We are deriving the potential twice. The first time we derive it on a grid with a finite difference operation and the second time we derive it in the infinitesimal sense by deriving the interpolating function.

Probably this is not very relevant in typical cosmological simulations, because it still holds $\frac{\partial}{\partial x}F_y \approx \frac{\partial}{\partial y}F_x$ in a very good approximate sense. However, it causes for example some issues when defining the associated tidal tensor. There are now two possible approaches to define it

1. Derive the potential twice by finite differences on the grid and then interpolate on the resulting tidal-tensor grid.

2. Derive the force field in the infinitesimal sense by deriving the interpolating polynomial.

Approach (1) seems the most straight forward approach to use and delivers an exactly symmetric tidal tensor. However, the issue here is that the tidal field is not consistent with the force-field in the infinitesimal sense. Therefore, if we consider the derivatives of the dark matter sheet, they will not be exactly the same if we reconstruct them from finite differencing from a converged particle distribution $N \to \infty$ or by using the evolved GDE.

Approach (2) has the advantage that this consistency is exactly given i.e. the GDE distortion tensor is exactly the continuum limit of the finite differences distortion tensor. However, the drawback is that the tidal tensor is non-symmetric and therefore the symplectic constraints are broken up.

Beyond the problems with the tidal tensor, I think that there might generally be a problem if fine-grained phase space quantities are considered in the continuum limit. Probably the Poincaré invariants are not exactly conserved for simulations with this Poisson solver. For example, that can compromise the refinement strategy that Sousbie & Colombi (2016) use in their code (i.e. leading to more refinement than necessary).

I could not find a simple and good solution to this issue, and probably it can be completely ignored in all cases where one does not bother with fine-grained phase space

properties in the continuum limit. However, I think the ideal way of obtaining the force and tidal-field from a potential field on a grid would be to define a continuous interpolated potential $\phi(\mathbf{x})$ from the discretely given values on the grid, and then derive this interpolated function in the infinitesimal sense to get $\mathbf{F}(\mathbf{x})$ and $\mathbf{T}(\mathbf{x})$. However, it is quite tough to come up with a version of such an interpolation scheme which meets sufficient smoothness requirements on the derivatives so that the tidal field does not end up being some kind of step function.

During this thesis we will by default make the choice of using approach (1) from above to obtain the tidal tensor (which is also used in the implementation from Vogelsberger & White (2011)). However, in cases where we explicitly test the convergence of the continuum particle distribution versus the infinitesimal distortion tensor I use approach (2) to ensure that those are consistent.

2.5.4 Correspondence with interpolation schemes

The distortion tensors $\mathbf{D}_{\mathrm{Xq}}$ and $\mathbf{D}_{\mathrm{Vq}}$ can now be obtained in two ways. (1) By reconstructing the functions $\mathbf{x}(\mathbf{q})$ and $\mathbf{v}(\mathbf{q})$ and measuring the derivative – which typically leads to some finite-differencing scheme. (2) By integrating the GDE. If the interpolated functions are consistent with the continuum limit, then both approaches should give the same answer. In Figure 2.18 we show the stream densities $\rho_s/\rho_0 = \left|\det\left(\mathbf{D}_{\mathrm{Xq}}\right)\right|^{-1}$ from both different approaches in Lagrangian space. This is from an $m_x = 250eV$ simulation that uses the HA16 scheme without refinement. Note that the HA16 scheme also improves the quality of the GDE integration a lot (in comparison to an N-body case like in Vogelsberger & White (2011)), because the higher quality density estimate leads to a more reliable estimate of the tidal tensor.

The stream densities agree very well in lower density regions and even in many shell-crossed regions. However, there are also regions where stream densities vary quite quickly with the Lagrangian coordinate and they disagree between the GDE and finite-difference scheme. That is so, because the dark matter sheet is too complex here to be reconstructed by the interpolation approach. These regions are predominantly haloes. We will discuss this topic in much more quantitative detail in chapter 4.

2.6 Features of the dark matter sheet

The consideration of the fine-grained phase space nature of the dark matter fluid leads to an interesting set of features which give additional insight into what is happening beneath the coarse-grained density and velocity fields which are typically investigated in cosmological simulations. In this section I will describe a set of those which I consider quite enlightening. These make it possible to look at cosmic structures from a quite different angle than just a pure classification by densities.

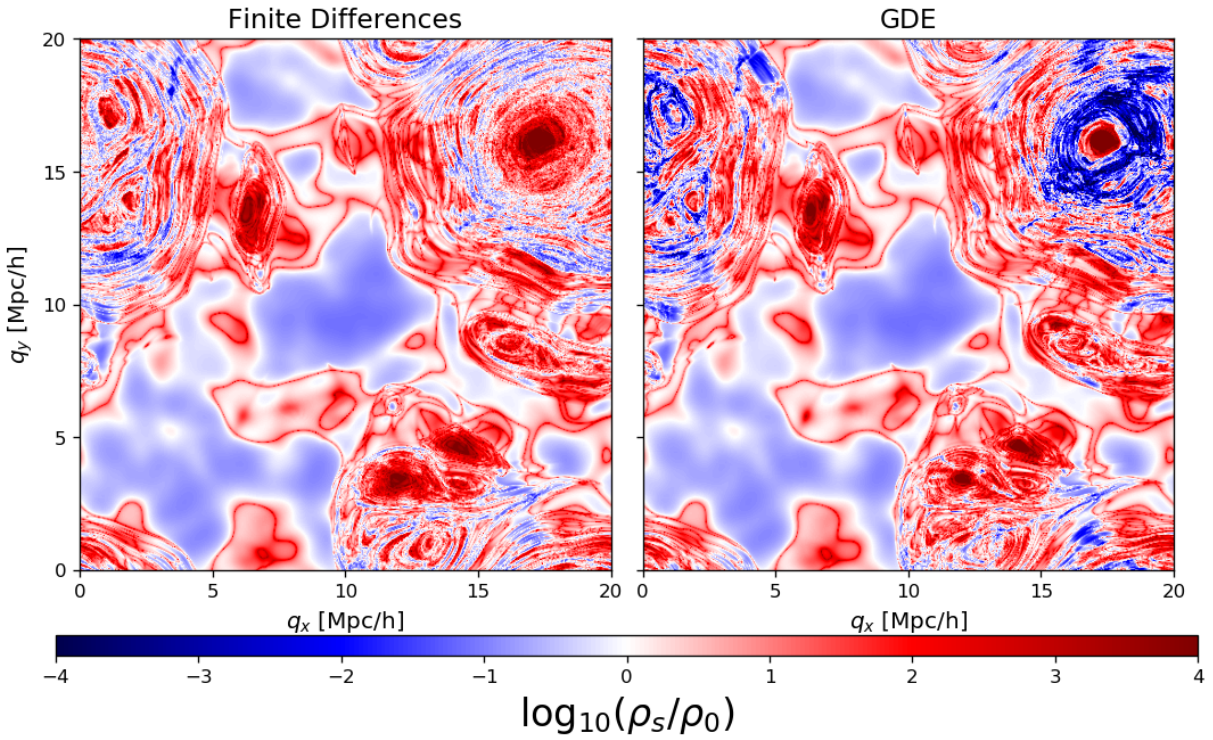

Figure 2.18: Stream densities in Lagrangian space. Left: as reconstructed from a finite-differencing scheme inferred from the interpolated dark matter sheet. Right: the infinitesimal derivatives as inferred from the geodesic deviation equation. Both schemes agree very well in regions where the stream densities vary slowly as a function of the Lagrangian coordinate. These are regions where the interpolation is converged. However, in regions where this is not the case both schemes disagree. As we shall see later, these are predominantly halos.

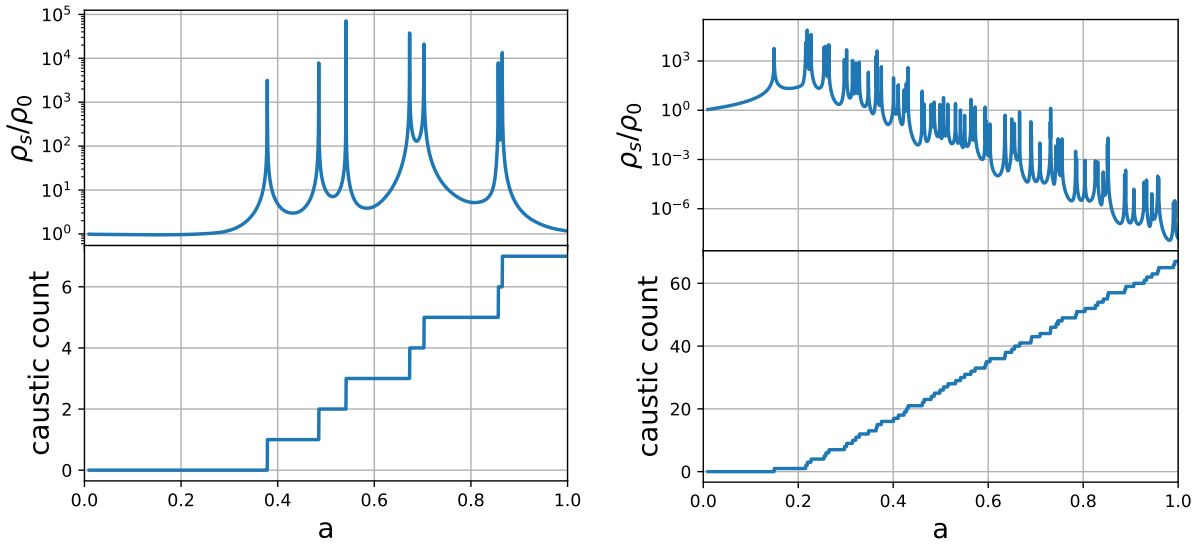

Figure 2.19: (Top) Stream density evolution and (bottom) number of caustics a particle has gone through for two different particles (left and right) from one of our simulations (from chapter 5). When a particle goes through a caustic its associated stream-density goes (formally) to infinity. The right case shows a particle which is in a halo. The left could maybe be a particle inside a filament or pancake.

2.6.1 Caustics

As already mentioned in early sections, the density field has sharp edges where a dark matter stream has a zero determinant of the distortion tensor $\det\left(\mathbf{D}_{\mathrm{xq}}\right) = 0$. These places are called caustics and the local density can get very large (depending on the velocity dispersion) and goes formally to infinity in the limit of a cold distribution. Vogelsberger & White (2011) have investigated the dark matter self-annihilation signal from such caustics with the help of the geodesic deviation equation. The GDE allows to determine the stream-densities of the infinitesimal (phase space) regions around each particle. We show in Figure 2.19 how the evolution of the stream density looks for two different example particles. A particle can go through caustics many times through its evolution. Each time the determinant of the distortion tensor flips its sign. Therefore in simulations one can count the number of caustics a particle has gone through by counting the number of times the distortion tensor has flipped its sign. The caustic count is also shown in the bottom panels of Figure 2.19.

The density within each caustic is only limited by the primordial velocity dispersion σ_v of dark matter (which is very small). This has already been investigated in more detail by White & Vogelsberger (2009); Vogelsberger & White (2011).

The width of caustics is extremely small, and proportional to the primordial velocity dispersion. In a small weekend project I tried to figure out: what is the typical width of a caustic? For example, are the largest caustics as big as our galaxy, as big as our solar

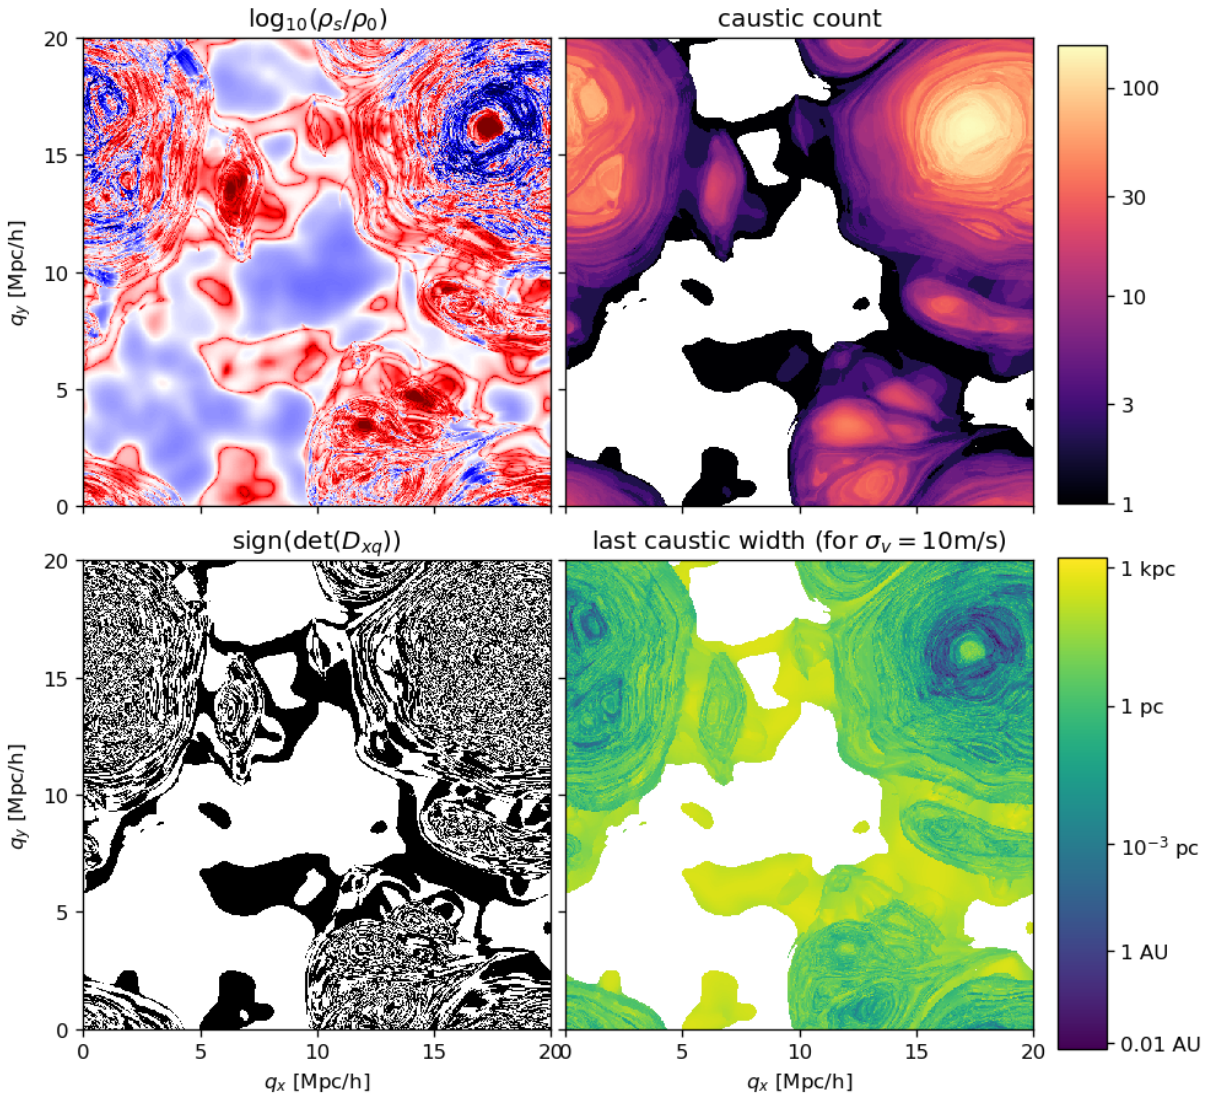

Figure 2.20: Caustic related quantities in Lagrangian space. Top left stream-densitiy (color is normalized as in Figure 2.18) - caustics are visible as strong enhancements in the stream-density. Top right: number of caustics a particle has gone through (white corresponds to 0). In low-density regions the counter increases quite slowly one by one whereas it increases rapidly inside of halos. Bottom left: sign of the distortion tensor. This shows the boundaries of individual streams. Inside of halos the sign seems almost uncorrelated between neighbouring particles. Right: the width of the last caustic for a velocity dispersion $\sigma_v = 10$m/s of a typical WDM case. For CDM models this value has to be multiplied by a factor of order 10^{-3}. The largest caustics get as big as kilo-parsecs. The smallest caustics are much smaller than our solar system ($\sim$ AU).

system or even much smaller than that? The width of a caustic is given by

$$\Delta x = \frac{\sigma_v}{|\mathbf{D}_{\mathrm{vq}} \cdot \mathbf{e_c}|} \tag{2.112}$$

where $\mathbf{e}_c$ is a unit vector which points in the direction of the caustic and $\mathbf{D}_{\mathrm{vq}}$ is evaluated at the time where $\mathbf{D}_{\mathrm{xq}}$ is degenerate. I will derive this relation in chapter 6 where I also show how to define $\mathbf{e}_c$. I made a simulation where I saved the width of the last caustic that each particle has gone through:

In Figure 2.20 we show several caustic related quantities for a Lagrangian slice from a 250eV warm dark matter simulation. The top left is the GDE stream density, the top right the caustic count, the bottom left the sign of the distortion tensor $\mathbf{D}_{\mathrm{xq}}$ and the bottom right shows the width of the last caustic that a particle has gone through for the case of $\sigma_v = 10\,\mathrm{m\,s^{-1}}$. This can easily be rescaled to other dark matter models, through multiplication with the ratio of their velocity dispersions. For example by a factor of order 10^{-3} for CDM models. The smallest caustics are much smaller than our solar system whereas the largest caustics can be of order kilo-parsecs (for WDM models) or parsecs (for CDM models). These estimates should even go down if a higher force-resolution and a colder power-spectrum are used.

2.6.2 Stream multiplicity

Another interesting feature of the dark matter sheet is the stream multiplicity. The stream multiplicity $n_s(\mathbf{x})$ is the number of different streams that are present at the same location. That means that one can find particles from n_s different original locations at $\mathbf{x}$. Therefore the velocity distribution function at $\mathbf{x}$ can be seen as a sum of n_s delta functions with different stream velocities v_s. We illustrate this in Figure 2.21. Note that the delta functions are actually something like a Gaussian with a very small width centered at the velocity of each stream. However, since the primordial velocity dispersion is so small, they appear like δ-functions. In Figure 2.21 I actually used a very large primordial velocity dispersion so that at least some of the Gaussians become visible.

The stream multiplicity n_s can be any uneven number - i.e. 1, 3, 5, and so on. One can argue that it should have even numbers at caustics, but in terms of volume these are a zero set anyways. Recently there have been a few investigations trying to define cosmic structures using the stream multiplicity (Ramachandra & Shandarin, 2015, 2017). However, it is not obvious in how far something fundamental would change when going e.g. from a 9 to an 11 stream region. Now of course one can try to smooth this field and define effective thresholds and everything, but in the end one is just following a much harder to understand and much worse converged version of the density field. My personal opinion is therefore that the stream multiplicity in general is not so suitable for identifying structures of different kinds. However, it is still an interesting quantity which can be used to understand the local complexity of the dark matter sheet. However, we will not use it in this thesis, because it is hardly possible to get a converged version of this field inside of high density regions.

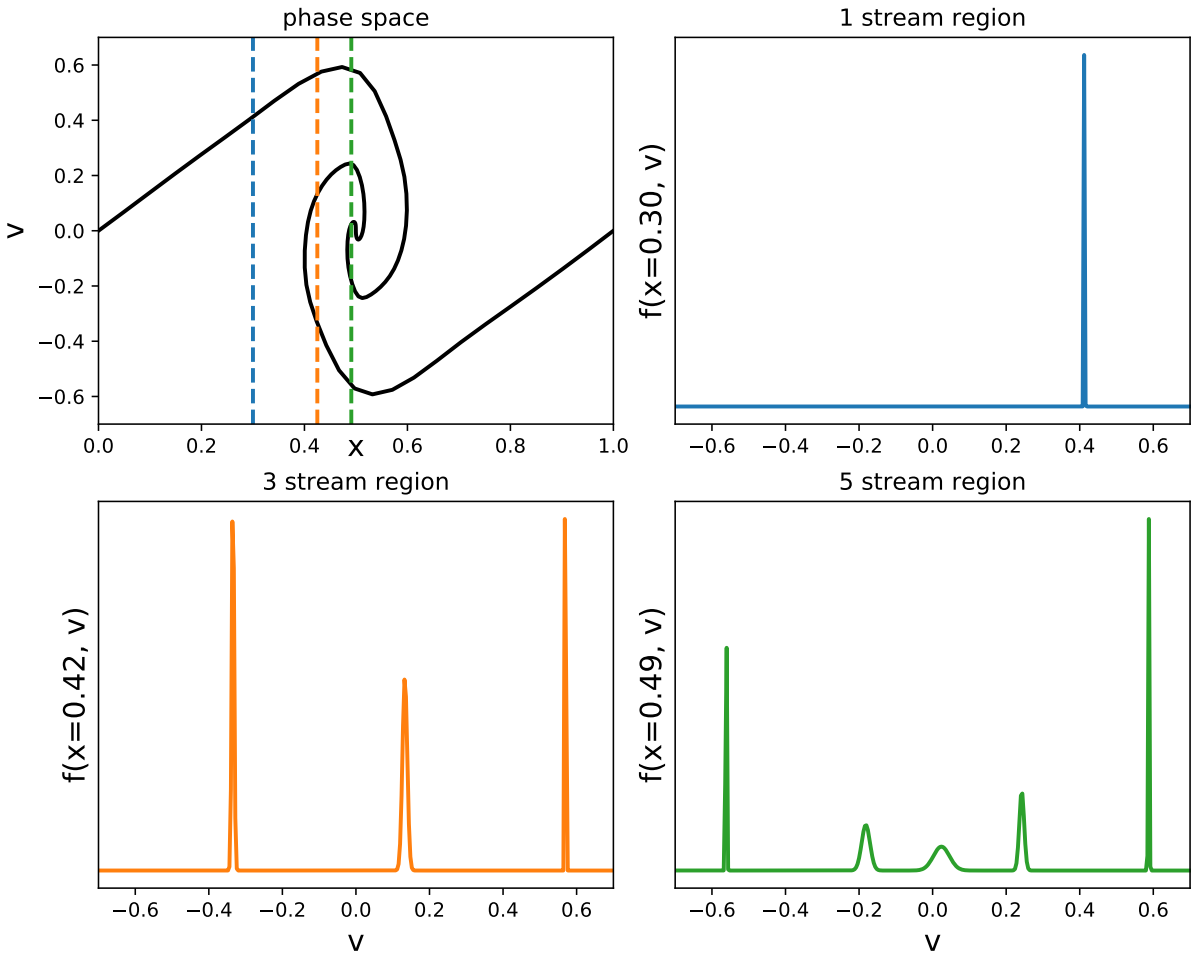

Figure 2.21: Illustration of the stream multiplicity. The top left shows the phase space distribution with different x = const. slices marked as dashed lines. The other three panels show the velocity distribution along those slices. To make the widths of the distributions visible, I have used a very large primordial velocity dispersion for the velocity distributions. In praxis these would be like delta functions in all cases. In a single-stream region the velocity distribution appears as a single delta function, in a 3 stream region as a sum of three delta functions and in a 5 stream-region as a sum of 5 delta functions and so on...

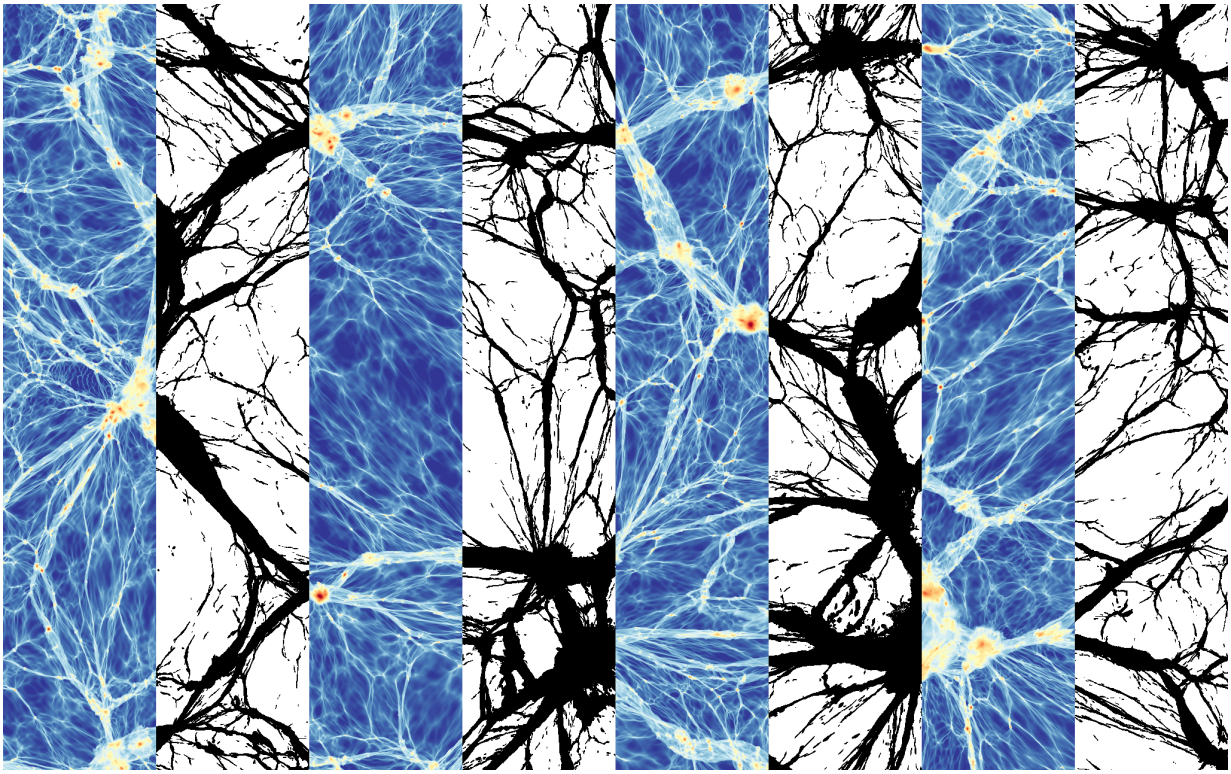

Figure 2.22: A segmentation of space into single-stream regions and multi-stream regions. Plotted are alternating slices of the logarithmic density field (blue-red) and the segmentation into single-stream regions (white) and multi-stream regions (black). This is a razor thin slice (depth = 20 kpc) with a width of 40 Mpc and a height of 25 Mpc. The transition from single-stream to multi-stream regions is visible as sharp edges in the density field (caustics) and is defined very accurately.

While it is not so obvious that there is a simple relation between stream multiplicity and structures in general, there is one threshold which is undoubtedly of qualitative and topological importance. This is the difference between single-stream regions which have $n_s = 1$ and multi-stream regions which have $n_s > 1$.

2.6.3 Single-stream regions

Single-stream regions are qualitatively very different from multi-stream regions. In the mathematical sense single-stream regions are topologically special, because the mapping from Lagrangian space $\mathbf{q}$ to Eulerian space $\mathbf{x}$ is injective here. In single-stream regions every Lagrangian coordinate gets mapped to a different unique Eulerian coordinate $\mathbf{x}$. Therefore the mapping is in principle invertible inside single-stream regions.

In a physical sense, single-stream regions are the most meaningful notion of what we could consider as diffuse uncollapsed regions whereas multi-stream regions include all kinds of collapsed structures - e.g. pancakes, filaments and halos. One can even argue that single-

stream regions might be used as a more meaningful definition of the notion "void". Usually voids are defined as strongly under-dense regions. Such definitions usually require the selection of a relative arbitrary quantitative threshold. However, the single-stream notion is very precise and *a-priori* meaningful.

We show in Figure 2.22 a slice through a density field combined with a segmentation into single-stream regions as inferred from our trigonometric resampling scheme. To determine the single-stream regions we additionally resample the real space distortion tensor $\mathbf{D}_{\mathrm{xq}}$. This can be done by deriving the displacement field in Fourier space:

$$D_{\mathrm{xq},lm} = \mathbb{1} + \mathrm{IFFT}(ik_m s_{k,l}) \tag{2.113}$$

where $\mathbf{s}_k$ is the displacement field in Fourier space and IFFT denotes an inverse Fourier transformation. Now every multi-stream region contains at least one stream with negative determinant of the distortion tensor $\det(D_{xq}) < 0$. One can imagine this by thinking of following along the phase space spiral as for example in Figure 2.21. If one ever has to go left ($\frac{\partial x}{\partial q} < 0$) one knows that one has to be in a multi-stream region. In single-stream regions on the other hand there exists only one single stream with $\frac{\partial x}{\partial q} > 0$. This is generalized to higher dimensions by using the sign of the determinant of the distortion tensor. In our resampling scheme we determine in every bin in Eulerian space the minimum-sign of the distortion-tensor determinant of all particles that have been binned to that bin. That is, if any particle inside a bin has a negative sign, we classify the bin as multi-stream region.

An interesting question is whether these single-stream regions percolate or not. So if one started in a single-stream region, would it be possible to travel through the whole universe without ever entering a multi-stream region? We have dedicated a whole section to this in section 3.6.

2.6.4 Morphology and rotation

Just from looking at projections and three-dimensional visualizations of density fields, one can already suspect that it is not enough to classify our universe only into over-densities and under-densities, halos and voids or into multi-stream regions and single-stream regions. Each of these bimodal classifications is missing out on some of the varieties of structures in the universe. One can roughly distinguish between four different stages of collapse

1. Single-stream regions or voids: These correspond to diffuse three-dimensional regions which have not collapsed at all.

2. Pancakes or planar sheet-like structures: These are structures which have collapsed along one dimension. Therefore the phase space distribution has heated up in that dimension, but the other two dimensions remain cold and dynamically not activated.

3. Filaments: These have collapsed along two dimensions, but still maintain one axis along which they remain cold and dynamically inactive. In the two-dimensional dynamically activated space rotations become possible.

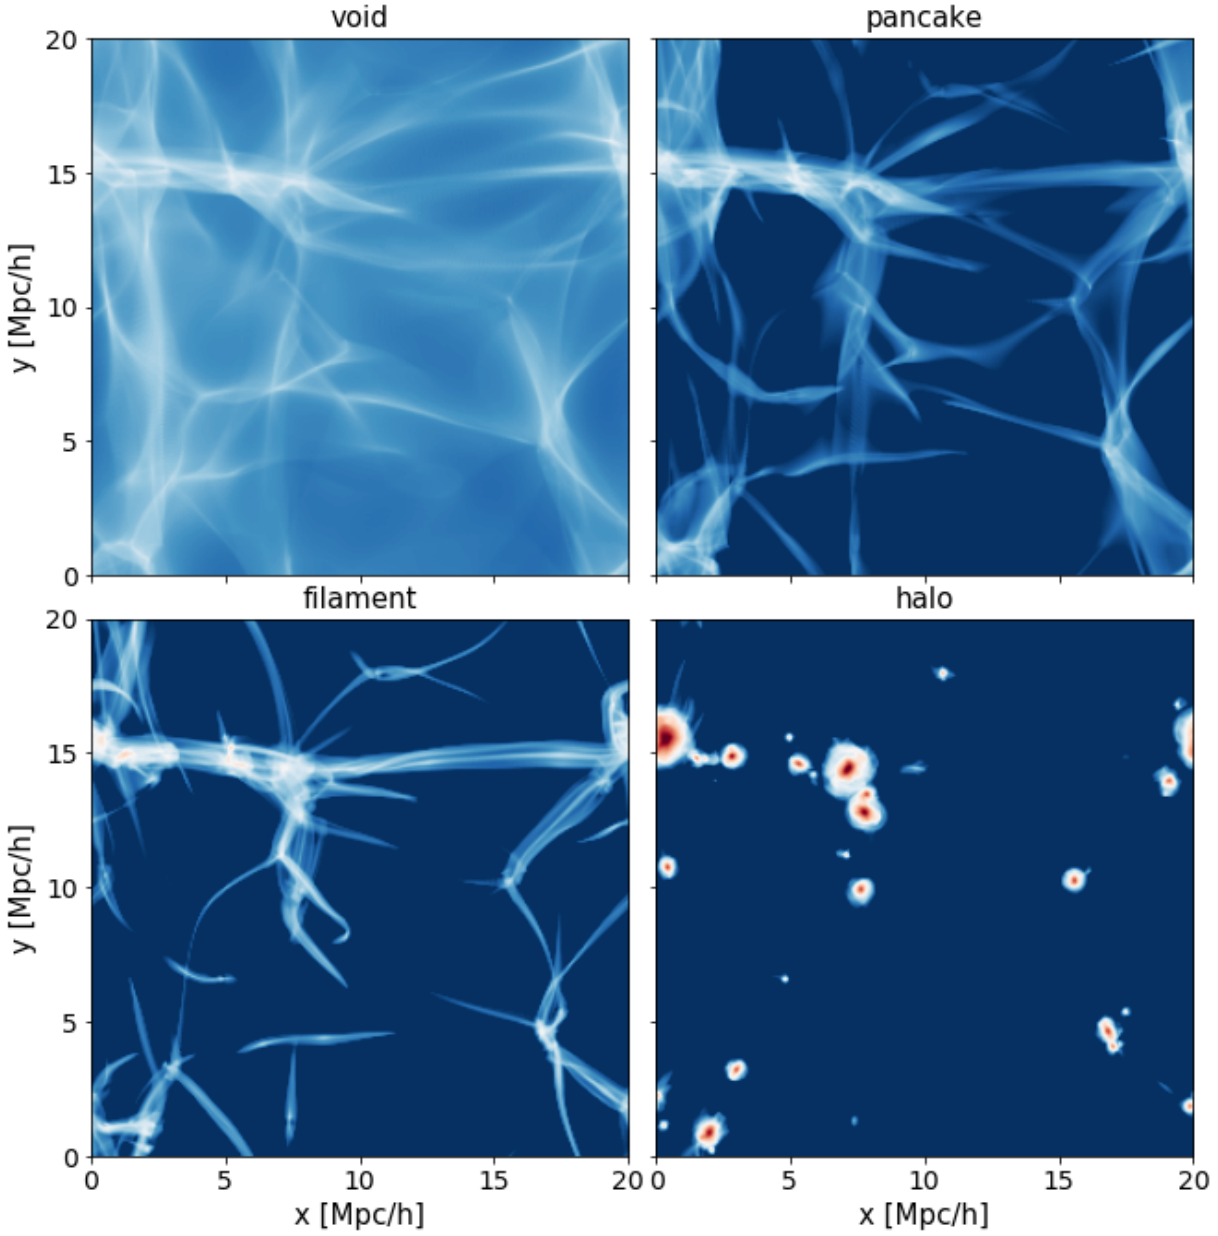

Figure 2.23: Density projections of a 20 Mpc/h cosmological warm dark matter box selecting different subgroups of the particles according to their structure class. The structure classification is done by counting the number of axes of the Lagrangian volume elements that are aligned with initial orientation: All axes aligned corresponds to a void, one axis misaligned (or flipped) to a pancake, two axis misaligned to a filament and all axes rotated arbitrarily in comparison to their initial aligned corresponds to a halo.

4. Halos: Objects which have collapsed along all three-dimensions into a point-like (or rather small spherical) object. These objects are dynamically active in all three dimensions.

We show the result of a morphological classification that we develop here in Figure 2.23. However, we will describe how we did this later in the text.

While one can try to distinguish regions of such different morphologies by quantitative investigations of density fields, such an approach would be not very insightful in terms of understanding what is going on, but more of an experimental *a-posteriori* nature. However, the qualitative explanation of these morphologies already contains possible ways to define those regions in more meaningful ways. As such one could define structures by the number of axes that have collapsed in each region. This has actually been done and we would like to refer the reader to Falck et al. (2012) for further interest. However, it is actually not very straight-forward how to mathematically define the number of axes an object has collapsed along. It depends on the (quite arbitrary) choice of axes when done like in Falck et al. (2012). Further there is quite some ambiguity between collapse and rotation in a three-dimensional world.

Instead we try here to disentangle structures by counting the number of dimensions which are dynamically fully activated. We can determine these by thinking of a small volume element around each particle. This volume element is described by the distortion tensor. If no dimensions of that volume element are dynamically active, then the volume element just gets slowly stretched or compressed along its different dimensions. However, all its major axes are aligned with its initial orientation. If one dimension is (fully) activated, then the volume element can flip its orientation along that axis (i.e. be mirrored) whereas the other two dimensions remain aligned with their initial orientation. If a second axis is activated the volume element can rotate and/or flip along two axes. However, the last axis still remains aligned with its initial orientation. If all three dimensions are active, the volume element can rotate or flip arbitrarily.

That said we can determine how many dimensions are dynamically active or "collapsed" by comparing the initial orientation of each Lagrangian volume element with its final orientation in Eulerian space. The degree of rotation between these will be the basis of our classifictaion scheme. We will describe here how to quantitatively describe this in a meaningful way.

Singular value decomposition

To disentangle stretching from rotations we use the singular-value decomposition of the distortion tensor. Any matrix can be decomposed in the form

$$\mathbf{D}_{\mathrm{xq}} = \mathbf{U}\mathbf{S}\mathbf{V}^T \tag{2.114}$$

where $\mathbf{U}$ and $\mathbf{V}$ are orthogonal matrices and $\mathbf{S}$ is a diagonal matrix where the diagonal elements are called the singular values s_i. If the Matrix $\mathbf{D}_{\mathrm{xq}}$ is symmetric, the singular-value decomposition becomes equivalent to the eigenvalue decomposition: then the singular values are the absolute values of the eigenvalues and $\mathbf{U} = \mathbf{V}$. However the singular

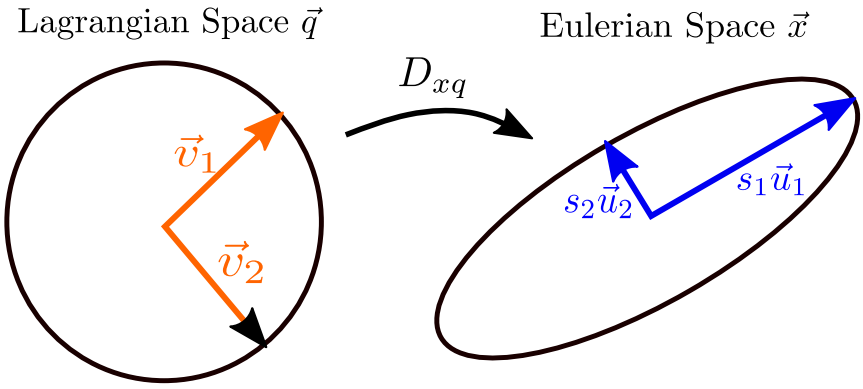

Figure 2.24: Illustration of the components of the singular value decomposition $D_{xq} = USV^T$. The Matrix D_{xq} maps a sphere from Lagrangian space to an ellipsoid in Eulerian space. The column vectors of U give the orientations of the major axes in Eulerian space, the singular values give the relative sizes of the axes, and the column vectors of V give the orientations in Lagrangian space.

value decomposition also has a simple geometric interpretation in the case of general non-symmetric matrices (like the distortion tensor). We illustrate this in Figure 2.24. The distortion tensor $\mathbf{D}_{\mathrm{xq}}$ maps a unit sphere in Lagrangian space to a distorted ellipsoid in Eulerian space. The column vectors $\mathbf{u}_i$ of the matrix $\mathbf{U}$ give the orientations of the major axes of the ellipsoid in Eulerian space. The singular values s_i quantify the stretching along the major axes. The column vectors $\mathbf{v}_i$ of $\mathbf{V}$ give the orientations of the major axes in Lagrangian space. So the general vector $a\mathbf{v}_1 + b\mathbf{v}_2$ gets mapped by $\mathbf{D}_{\mathrm{xq}}$ as

$$\mathbf{D}_{\mathrm{xq}} \cdot (a\mathbf{v}_1 + b\mathbf{v}_2) = as_1\mathbf{u}_1 + bs_2\mathbf{u}_2 \tag{2.115}$$

To quantify rotations we define three angles α_i from the singular value decomposition

$$\alpha_i = \arccos(\mathbf{v_i} \cdot \mathbf{u_i}) \tag{2.116}$$

which are the relative angles between the orientation of the major axes in Lagrangian space and in Eulerian space. Note that this choice of angles is independent of the coordinate system (unlike most possible angle definitions). In the example from Figure 2.24 the angle α_1 would be relatively small whereas the angle α_2 would be close to 180°. The angles α_i and the singular values s_i form together a set of six variables which contain frame independent meaningful information about the distortion tensor. The remaining three components only contain information about the absolute orientation and are of no real interest because of the isotropy of the universe.

We show these angles in Lagrangian space in the second row of Figure 2.25. From these angles we could try to directly classify particles into structures. However, note that during any particular time during the orbit of a particle its volume-element can be aligned with its initial orientation just by chance. Therefore an additional step is required which abstracts a bit from the current particular configuration of a particle's volume-element. We have come up of two possible ways of doing that.

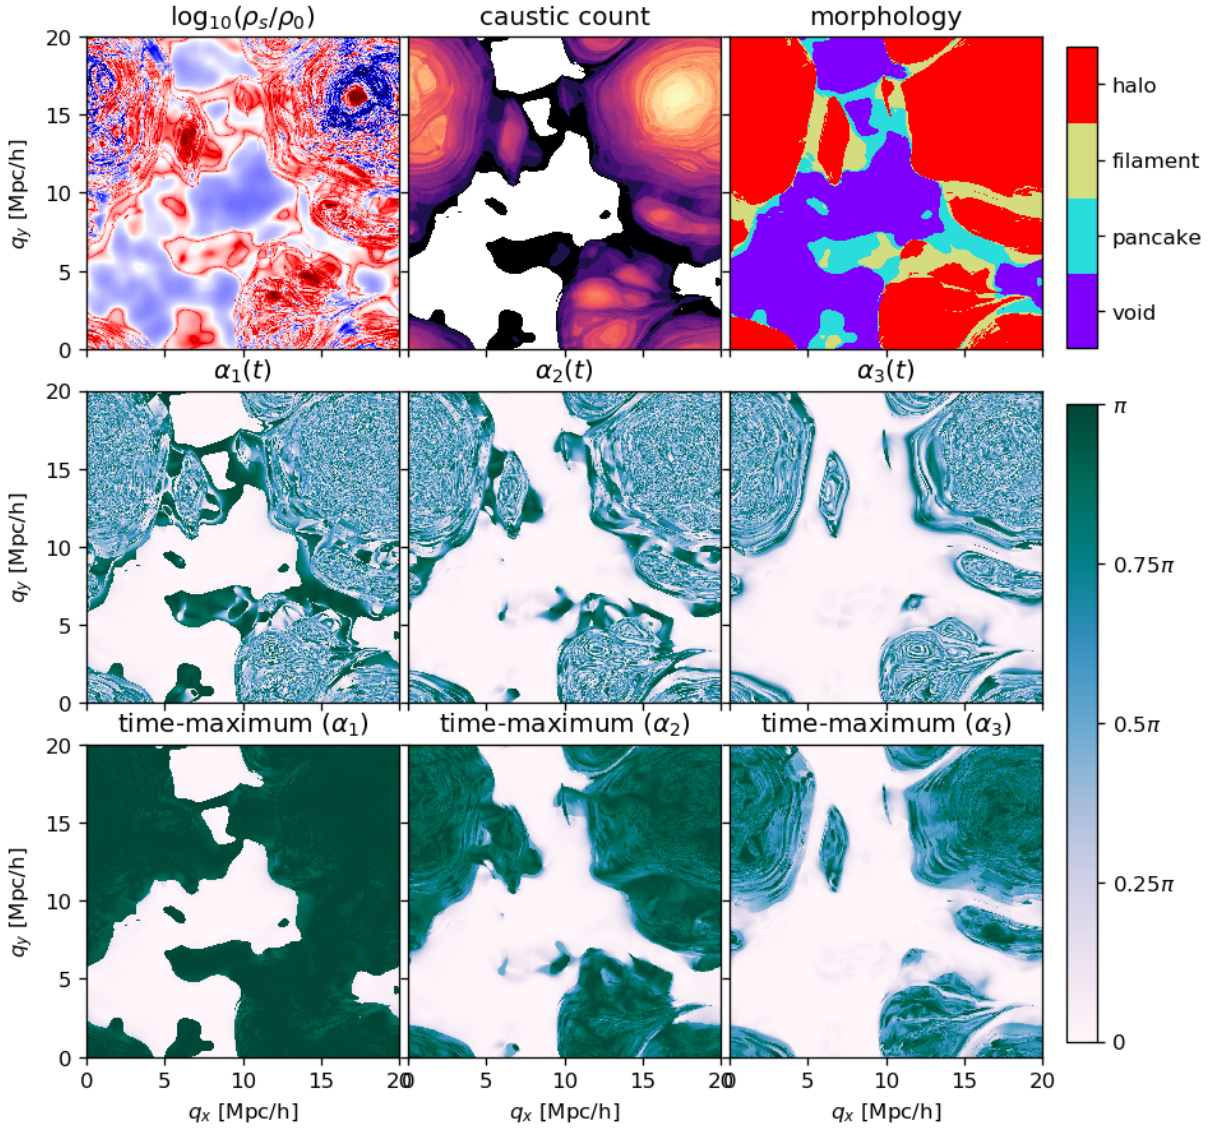

Figure 2.25: A slice in Lagrangian space through a cosmological warm dark matter simulation with m_X =250eV. First two panels: the stream density and the caustic count with the same color normalization as in 2.20. Third panel: morphology classification by the angles as described in the text. Second row: Angles of the distortion tensor at a single time ($a = 1$). Third row: Time-maximum of these angles over the whole history of each particle. It is striking that the angles are activated in clearly distinct Lagrangian regions. Therefore the morphology classification appears to be very robust. However, it is necessary to take the maximum value of the angles to avoid misclassification for cases where axes align by chance. Note also that the pancake/void boundary intersects more or less exactly with the caustic count 0 to 1 boundary.

Taking the time-maximum

Instead of using the angles α_i for classification, we can use the time-maximum of these angles for the classification. We trace the angles α_i in our simulations and determine the time-maxima of them for every particle. We show the difference between the particular state of the angles at some time t and the time-maximum of these in Figure 2.25. Note that the angles appear to be active or inactive in very distinct Lagrangian regions with quite sharp boundaries. That said the angles that we defined indeed point out qualitatively distinct regions.

We classify structures by the number of angles for which the maximum exceeds $\pi/4$: 0 corresponding to a void, 1 to a pancake, 2 to a filament and 3 to a halo. Note that the threshold of $\pi/4$ is relatively arbitrary, but the classification is quite robust against the detailed choice of this angle. We show the classification result in Lagrangian space in the top right panel of Figure 2.25. We have already shown the result of this classification in Eulerian space in Figure 2.23. Arguably the classification selects regions in the same way one would intuitively classify them. However, we want to point out that since our classification is based on the dynamical behavior of particles and not pure Eulerian space properties, there can coexist particles at the same location which are assigned to different morphological structures.

Taking the stream-maximum

While the previous approach of using the time-maximum of the distortion tensor angles for structure classification seems to work quite well, it has two drawbacks for practical applications. (1) It does not lead to unique classification in Eulerian space - for example at the same location there might co-exist void-particles and pancake-particles. (2) It requires the additional effort of tracing the distortion tensor (either from finite differences or by the GDE) during the simulation and doing a singular value decomposition at each time-step. This makes it impossible to use it as a pure post-processing step on already existing dark matter simulations. These two drawbacks are not relevant for the applications we use this classifications for in this thesis (in chapter 4). However, it would be nice to also have a version of this scheme which can be applied as a distinct Eulerian-space classification in a pure post-processing step.

Such a scheme could be obtained by taking the stream-maximum of the angles. With stream-maximum I mean that the class of a given Eulerian location $\mathbf{x}$ is determined by the maximum of the angles $\alpha_i(\mathbf{q_s})$ over all streams $\mathbf{q_s}$ that are present at location $\mathbf{x}$ - i.e. $\mathbf{x}(\mathbf{q_s}) = \mathbf{x}$.

This definition allows to use a resampling approach to determine the morphology e.g. on a grid. I have implemented this into the mpi trigonometric resampling code that was described earlier. Here the maximum of the angles over all pseudo-particles that get deposited into the same bin is determined. This maximum should become equivalent to the stream-maximum in the limit of arbitrary many pseudo-particles. In Figure 2.26 we show the result of this classification. Here I have chosen (quite arbitrarily) a threshold of $\pi/2$.

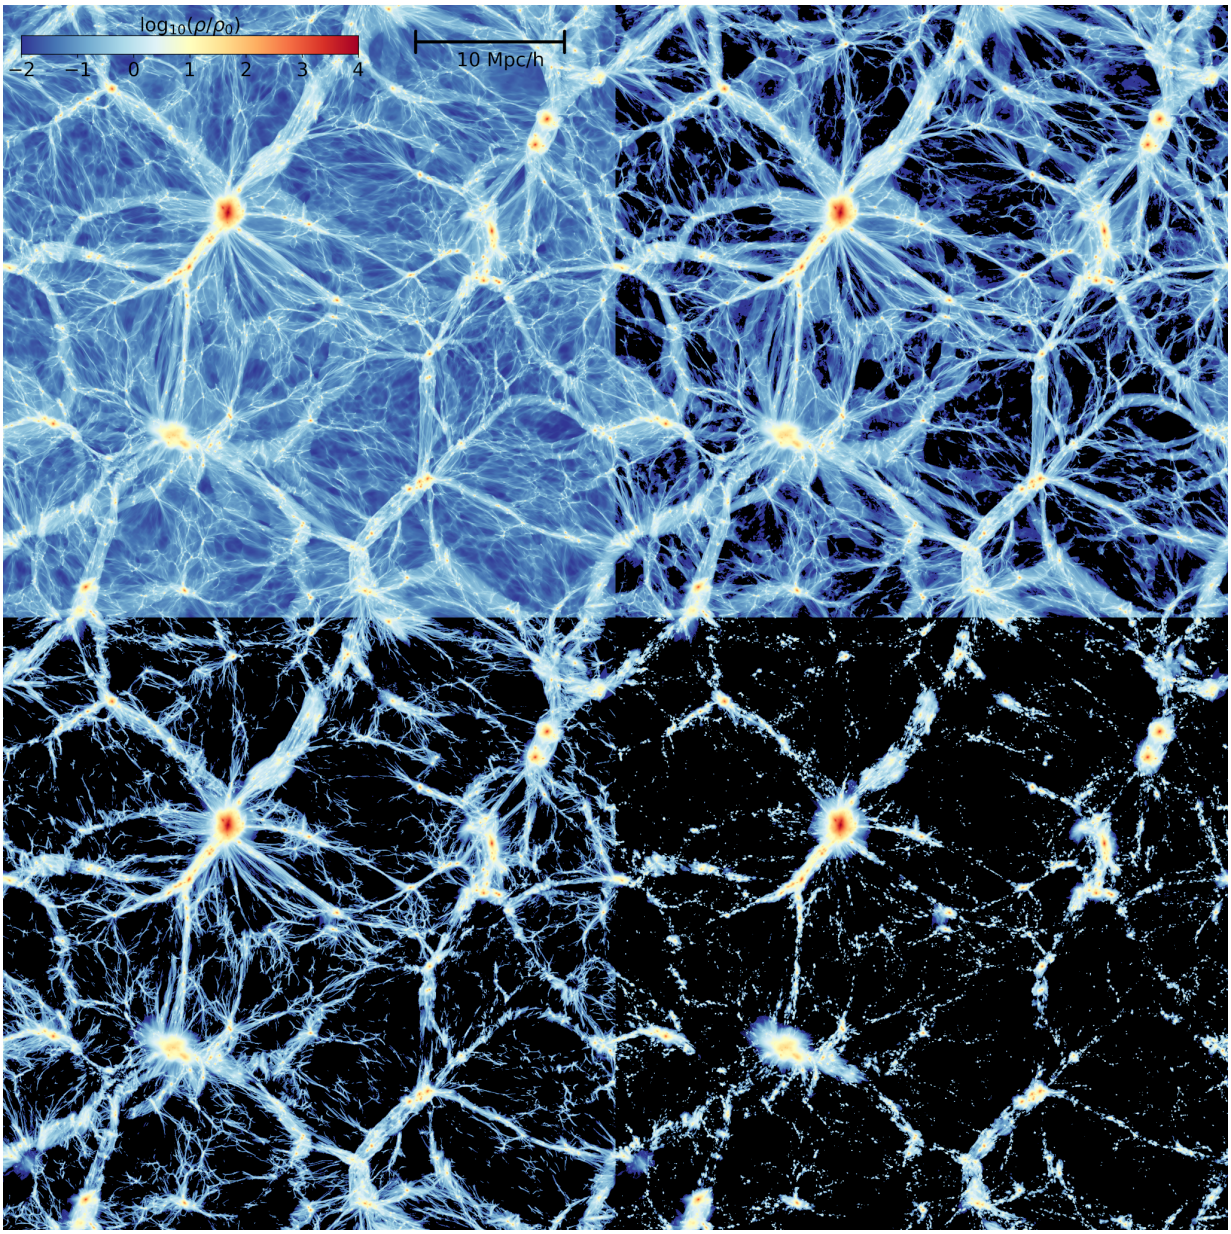

Figure 2.26: Projections of the density field (depth = 0.8 Mpc/h) of a CDM simulation selected on different morphological types. Top left: all mass, top right: only mass which is in pancakes + filaments + halos, bottom left: mass which is in filament or halos, bottom right: only mass in halos. Arguably the classification selects the structures close to the way one would intuitively classify them. We provide other versions of this Figure in the appendix A.1 and A.2

Further in this Figure I show always all mass which has morphology class $\geq$ a given class. So the first plot shows all mass, the second plot pancakes, filaments and haloes but not voids, the third plot shows filaments and halos and the last plot only halos. In the appendix I also provide a Figure where the mass is selected by exactly one morphology class (A.1) and one where the mass is selected $\leq$ a given class (A.2). I recommend having a look at those figures, too.

2.7 Summary

In this chapter we have explained the notion of the dark matter sheet in phase space. In this context we discussed N-body simulations and sheet based simulation schemes, the geodesic deviation equation and different features and structure classification schemes that arise from this fine-grained view. In the next chapters these things will be used in different application areas. In chapter 3 we will focus on the investigation of single-stream regions, develop an (almost) analytical excursion-set model to predict their properties and investigate whether they percolate. In chapter 4 we will try to solve the fragmentation of N-body simulations by combining sheet-based simulation techniques with an N-body + GDE approach in critical regions. In chapter 5 we will investigate the density and phase space structure of an isolated small warm dark matter halo with the newly developed scheme. In chapter 6 we will explore the possibility of reconstructing a warm phase space distribution from a single cold sheet and discuss how this might be applied in the context of warm dark matter simulations or for simulations of the cosmic neutrino background.

Chapter 3

The median density of the Universe

This chapter is about the project "The median density of the Universe" which we published in Stücker et al. (2018). It started with Simon White asking me and Philipp Busch: What is the median density of the universe? This started of a rich discussion and evolved into a whole project out of curiosity. Easily one can see that the median density must be much smaller than the mean density, because most mass of the mass in the universe assembles in high density regions like halos – thereby evacuating "typical" regions. Further, one can easily see that the median density of the universe cannot be zero. That is so, because in the beginning the whole universe is filled by the continuously connected dark matter sheet. Phase space densities are conserved along trajectories and trajectories cannot intersect. It is not possible to "tear apart" this continuous sheet (but only to stretch and fold it).

Beyond these two qualitative statements we have come up with a quantitative scheme to estimate the median density of the universe. The following text will be about this.

Note that this is almost one to one from the publication in Stücker et al. (2018). However, in some places I added some extra material that can simplify the understanding process. Further I removed parts which are redundant with the explanations in chapter 2 and replaced them by shorter summaries. The details of the publication are listed in Table

title:	The median density of the Universe
authors:	Jens Stücker, Philipp Busch and Simon D.M. White
publication date:	07/2018
journal:	Monthly Notices of the Royal Astronomical Society
volume:	Volume 477, Issue 3, p.3230-3246
misc:	has been awarded with the Rudolf Kippenhahn award for the best student publication (submitted) in 2017 at MPA (jointly with Aniket Agrawal)

Table 3.1: Details of the publication that this chapter is based on.

3.1

3.1 Abstract

Despite the fact that the mean matter density of the universe has been measured to an accuracy of a few percent within the standard ΛCDM paradigm, its median density is not known even to order of magnitude. Typical points lie in low-density regions and are not part of a collapsed structure of any scale. Locally, the dark matter distribution is then simply a stretched version of that in the early universe. In this single-stream regime, the distribution of unsmoothed density is sensitive to the initial power spectrum on all scales, in particular on very small scales, and hence to the nature of the dark matter. It cannot be estimated reliably using conventional cosmological simulations because of the enormous dynamic range involved, but a suitable excursion set procedure can be used instead. For the Planck cosmological parameters, a 100 GeV WIMP, corresponding to a free-streaming mass $\sim 10^{-6} M_\odot$, results in a median density of $\sim 4 \times 10^{-3}$ in units of the mean density, whereas a 10 μeV axion with free-streaming mass $\sim 10^{-12} M_\odot$ gives $\sim 3 \times 10^{-3}$, and Warm Dark Matter with a (thermal relic) mass of 1 keV gives $\sim 8 \times 10^{-2}$. In CDM (but not in WDM) universes, single-stream regions are predicted to be topologically isolated by the excursion set formalism. A test by direct N-Body simulations seems to confirm this prediction, although it is still subject to finite size and resolution effects. Unfortunately, it is unlikely that any of these properties is observable and so suitable for constraining the properties of dark matter.

3.2 Introduction

Let us start this chapter with a simple thought experiment: Imagine we were able to reliably measure the mass density in small volume elements, let us say cubes with a side length of a kilometre. And let us further assume we would be able to do this measurement anywhere in the universe. Now if we would do this measurement for a large number of randomly placed cubes, what would the distribution of their densities look like? Or rephrasing: *What is the one-point density distribution of the Universe at very high resolution?* What is its median density? What is its shape? What determines the behaviour of its high- and low-density tails? And if we knew the distribution, could we learn something about dark matter?

Although such a measurement is not possible today nor will it be possible anywhere in the near future, trying to answer this question from purely theoretical arguments turns out to give considerable insight. It helps paint a simple picture of what is happening in the majority of the volume - from the largest to the smallest scales. We here propose an excursion set formalism which provides such a qualitative picture and further enables us to estimate the unsmoothed density distribution of the universe.

In the current best fitting model of cosmic structure formation, the main gravitating component is dark matter. In the early universe it is distributed almost homogeneously

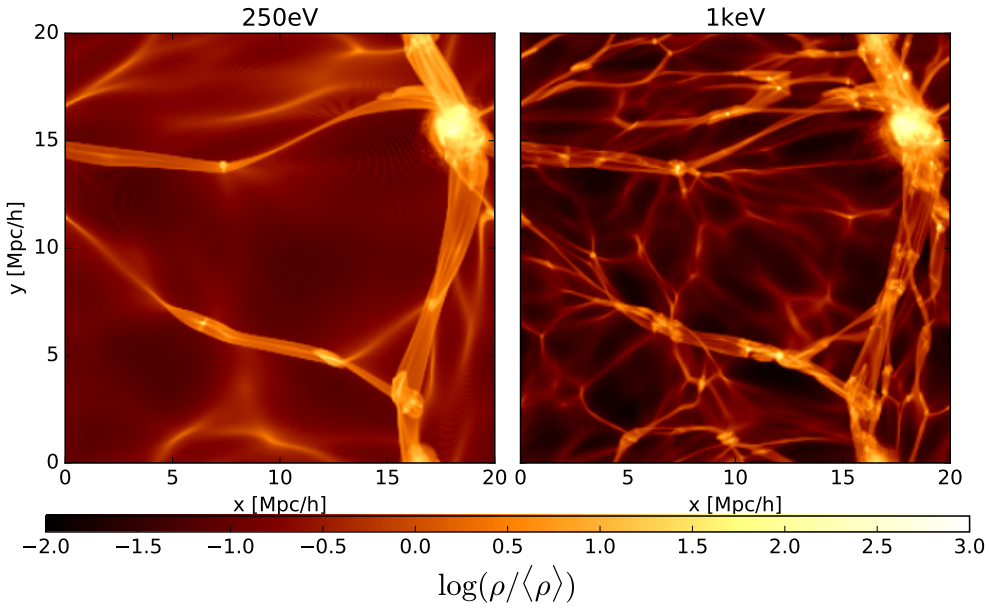

Figure 3.1: A razor thin slice through two WDM simulations with different thermal cut-offs in the power spectra corresponding to a 250eV thermal relic (left) and a 1keV thermal relic (right). In the un-collapsed single-stream regions the density distribution depends strongly on the dark matter model and its Lagrangian smoothing scale.

with only small perturbations from the mean density. At this time the distribution of density perturbations is expected to be given by a simple Gaussian distribution.

However as the universe expands, the perturbations grow - early on in a simple linear way, and thereafter in a non-linear and partially chaotic manner. While the linear regime can be well described by analytic methods, the investigation of the dark matter distribution in the non-linear regime usually requires N-Body simulations which explicitly follow the evolution of a large set of tracer particles in a three dimensional cosmological volume.

The power spectrum of density perturbations is extremely flat for cold dark matter cosmologies in the sense that density perturbations on all scales from hundreds of mega-parsecs down to a thermal smoothing scale of e.g. parsecs (for WIMPs) are relevant to determine what happens to the unsmoothed density field in the non-linear regime. To encompass nearly homogeneous scales, a cosmological simulation to follow the unsmoothed density field of WIMP-like dark matter would need to resolve about 8 orders of magnitude in spatial scale, requiring of order 10^{24} resolution elements which is still far from what is possible.

This problem is usually tackled by smoothing the initial conditions of dark matter simulations on a relatively large length scale - either explicitly by introducing a cut-off scale into the power spectrum as in warm dark matter simulations, or implicitly by the Nyquist frequency of the mesh that samples the initial density field. The conclusions that can be made from these simulations are then limited to features that do not depend on initial perturbations that are smaller than this Lagrangian smoothing scale. As most

observations involve a relatively large smoothing anyway, the smoothing in Lagrangian space is usually of little importance for the comparison with observations. However, in our thought experiment we are asking explicitly for the unsmoothed density field.

The unsmoothed density field depends strongly on the small scale cut-off of the dark matter power spectrum (which is equivalent to a Lagrangian smoothing scale). To illustrate this we show in Figure 3.1 a thin slice through two warm dark matter simulations with different free-streaming scales. The smaller the smoothing scale, the more diffuse material fragments into small scale structures, and the lower the typical density of the universe becomes.

While the Lagrangian smoothing scale is incorporated explicitly in these two simulations, it is also present implicitly in all classical cold dark matter simulations. If the resolution of a cold dark matter simulation is changed, the maximum spatial frequency of the imposed initial perturbations shifts. That leads to additional small scale structure which strongly modifies the density distribution (Yang et al., 2015). As an example of this we show the volume-weighted density distribution of the Millennium Simulation (Springel et al., 2005b) in comparison to the much higher resolution Millennium Simulation II (Boylan-Kolchin et al., 2009) in Figure 3.2. Here the density field is approximated by attributing the volume of each cell in a Voronoi tessellation of the particle distribution to the particle at its centre and using this to provide a density estimate. Additionally we show the 50-, 90- and 99-percentiles of the distributions. It is evident that these density distributions are far from converged. While the relatively good convergence in the high density tails of these distributions was previously discussed by Pandey et al. (2013), here we show that there is a factor of 2 between the median densities. A similar factor lies between the peaks in the density distributions and the minimum particle densities ($(\rho_{\min}/\rho_0)_{MSI} = 1.1 \cdot 10^{-2}$ and $(\rho_{\min}/\rho_0)_{MSII} = 5.6 \cdot 10^{-3}$).

It is worth noting that while most of the mass is part of collapsed structures, most of the volume is part of single-stream regions. Therefore the density distribution of the universe is mostly given by the density distribution of single-stream regions. Only the high density tail will be affected by collapsed structures.

As a reminder of what was already explained in 2.6.3: single-stream regions are regions which have not undergone any kind of collapse. That means no Lagrangian patch has yet passed any caustic. Therefore only a single dark matter stream is present and the density is given locally by the stream-density of that one stream. This makes single-stream regions mathematically much simpler than multi-stream regions. In Figure 3.3 the notion of single-stream regions is illustrated. The concept of single-stream regions is similar to the idea of voids. The term void is, however, often used to refer to the largest underdensities in the universe after smoothing on Mpc scale or larger (e.g. van de Weygaert & Platen (2011)). Such voids actually contain many collapsed objects of smaller scale. In contrast our subject of interest here is to describe the regions of the universe which contain no collapsed object of any scale, motivating our definition of single-stream regions. As already discussed in chapter 2, tracing the detailed structure of the dark matter phase sheet has recently become possible (Shandarin et al., 2012; Abel et al., 2012; Hahn & Angulo, 2016; Sousbie & Colombi, 2016) allowing the stream multiplicity to be measured

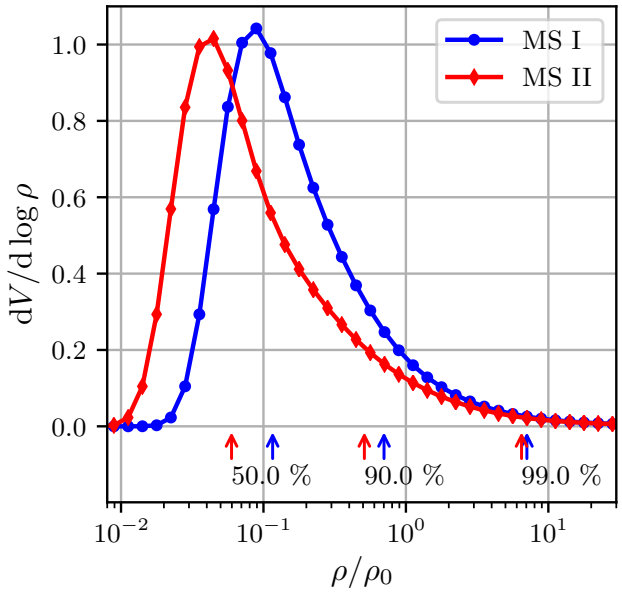

Figure 3.2: Comparison between the volume-weighted density distribution of the particles in the Millennium I (MS I) and Millennium II (MS II) simulations, using the volume of each particle's Voronoi cell to calculate its density. The density distribution in classical cold dark matter simulations is still far from from converged with the resolved small-scale power. Image credit goes to Philipp Busch.

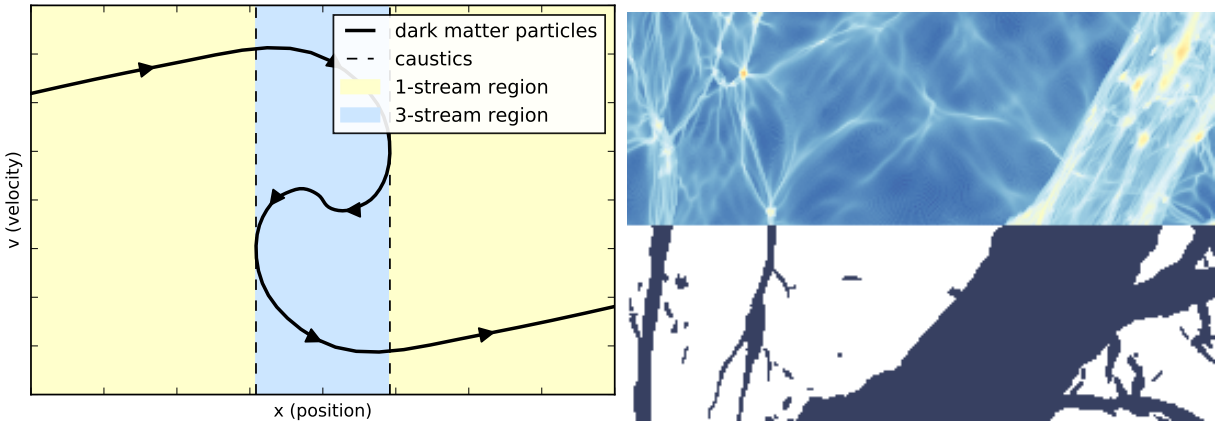

Figure 3.3: Summarizing illustration of what is meant by single-stream regions - left: in phase space and right: (top) the density field and (bottom) a segmentation into single-stream (white) and multi-stream region (black). In single-stream regions only one single dark matter stream is present whereas in multi-stream regions more than one stream is present. The transition is marked by sharp edges in the density field (caustics). Single-stream regions can be described by mathematically simpler models than multi-stream regions. For a more detailed explanation please refer back to section 2.6.3. (This Figure was not part of the original publication.)

in simulations, and giving an interesting new perspective on structures in the universe (Ramachandra & Shandarin, 2017).

We propose an excursion set formalism that allows prediction of the density distribution of single-stream regions. The formalism defines a collapse criterion which detects the first caustic crossing of a particle which occurs as it first becomes part of a two dimensional sheet-like structure (a pancake). It then checks whether this criterion is fulfilled by the smoothed linear density field at the point corresponding to a particular particle for *any* Lagrangian smoothing scale. If it is not fulfilled on any scale, the particle is assumed to be part of a single-stream region. In that case we expect the local density to be well described by a simple model such as the Zel'dovich approximation or the triaxial collapse model which we introduce here. We use this excursion set formalism to predict the density distribution of single-stream regions, and the total amount of mass expected within single-stream regions.

Another interesting question that can be answered within this context is whether single-stream regions form distinct regions enclosed by collapsed structures (i.e. by multi-stream structures), or whether they form one connected infinitely large percolating region. It is well known that galaxies are arranged in a percolating web-like structure (Zeldovich et al., 1982) and Shandarin et al. (2006) showed that also voids percolate if defined as regions below a certain smoothed density threshold.

The percolation of single-stream regions has already been investigated by Falck & Neyrinck (2015) and Ramachandra & Shandarin (2017) who find that the single-stream regions in their simulations percolate. In contrast we show in section 3.6 that our excursion set formalism predicts that single-stream regions do not percolate in cold dark matter universes. The regime where single-stream regions stop percolating lies beyond the resolution limit of the simulations of Falck & Neyrinck (2015). We attempt to test this regime with an N-Body simulation. When we test for percolation in Eulerian space, we find that the sizes of individual single-stream regions depend significantly on resolution parameters. However, a percolation test in Lagrangian space, which we consider more robust against numerical artefacts, shows no percolation. We thus infer that single-stream regions do not percolate in the continuum limit of cold dark matter.

Finally we note that in this chapter we only consider the idealised problem of a ΛCDM universe without baryons. The inclusion of realistic galaxy formation physics and a realistic baryon fraction would significantly alter our conclusions, since pressure effects after reionisation would remove most of the baryons from low mass haloes and distribute them smoothly throughout low-density regions. This would in turn inhibit the late-time growth of the low-mass haloes themselves.

3.3 An excursion set formalism for single-stream regions

In this section we introduce an excursion set formalism that can be used to predict the density distribution of single-stream regions. We briefly review classic excursion set formalisms, then introduce two alternative models for the collapse barrier, and summarize the mathematical background needed for the six dimensional random walk of the deformation tensor. Finally, we note that the predictions of the excursion set models depend only on the variance of the unsmoothed linear dark matter density field.

3.3.1 Excursion set formalisms

Excursion Set formalisms have long been used as simplified models of structure formation in the non-linear regime. Probably the most prominent one is the extended Press-Schechter formalism, hereafter EPS, by Bond et al. (1991) which has been used to predict conditional and unconditional halo mass functions and halo clustering bias. The unconditional mass function turns out to be very similar to that originally obtained by Press & Schechter (1974).

The EPS formalism is based on a simple assumption derived from the spherical collapse model: any uniform spherical perturbation which has a linear-theory density contrast larger than $\delta_c = 1.68$ is assumed to have collapsed. Given a realisation of a linear cosmological density field all particles for which $\delta > \delta_c$ for some smoothing scale are assumed to be part of a halo in a fully non-linear calculation. The mass of the halo is assumed to correspond to the largest smoothing scale for which the collapse criterion is satisfied.

In the simplest version of the EPS formalism the largest scale is identified by smoothing with a top-hat filter in Fourier space. One starts with an infinitely large smoothing length scale $R_s \propto k_s^{-1}$ and decreases it smoothly. The density contrast at a single point then effectively makes an uncorrelated random walk. If the density contrast first crosses the barrier δ_c the particle is assumed to be part of a halo with Lagrangian size corresponding to the smoothing scale R_s at first crossing. In the EPS model the distribution of R_s at first crossing thus determines the halo mass function.

Despite its simplicity the excursion set formalism has been shown to give reasonably good predictions for the halo mass function (e.g. Zentner, 2007). It has been improved by introducing more sophisticated barriers which consider deviations from spherical symmetry (Sheth et al., 2001). Arguably a major reason why it works so well is the way large scale modes interact with small scale modes in the non-linear regime. Large modes strongly influence what is happening on small scales, while small scales barely influence any large scale structure. Therefore if something can be determined to collapse when the density field is smoothed on a large scale, it will almost certainly also collapse in the unsmoothed density field, since the smaller scale perturbations do not influence the large scale structure. In Eulerian space this can be seen in Figure 3.1 where the dominant large structures remain almost unchanged by decreasing the Lagrangian smoothing scale. Additionally we

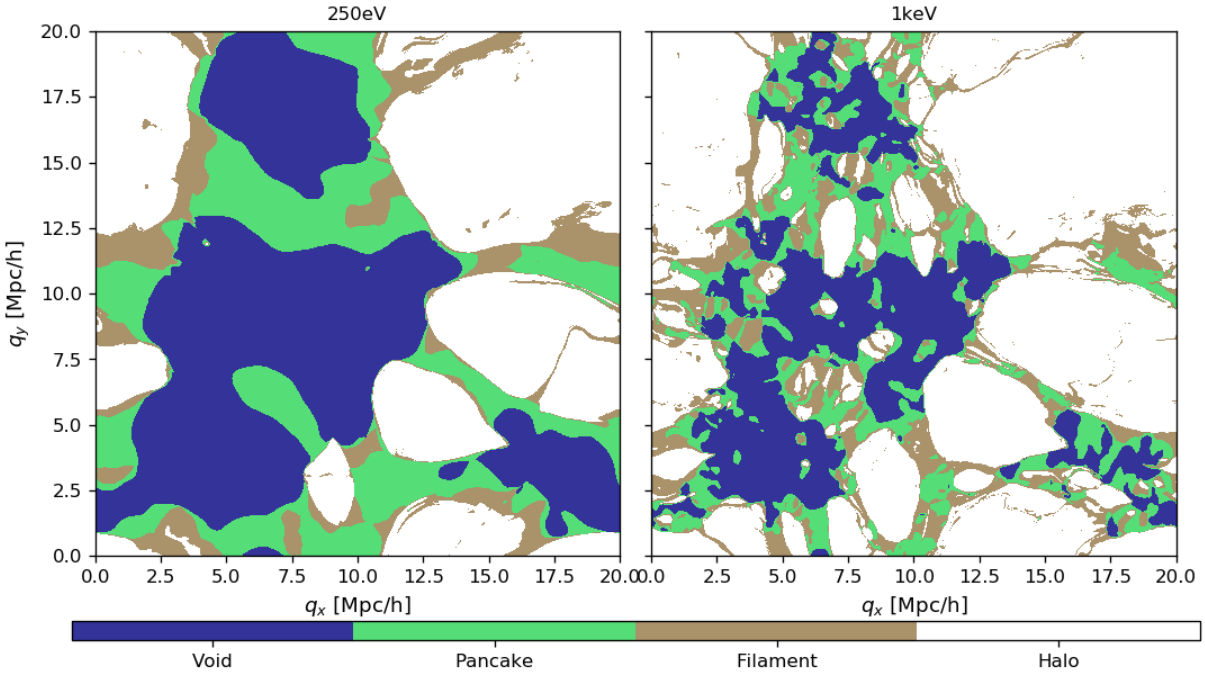

Figure 3.4: Morphology classification as described in section 2.6.4 for two different warm dark matter simulations in Lagrangian space. The left panel shows a warmer case ($m_x = 250$eV) with a larger Lagrangian smoothing scale ($\rightarrow \sigma = 2.4$) and the right panel shows a colder case ($m_x = 1$keV) with a smaller Lagrangian smoothing scale ($\rightarrow \sigma = 4.1$). Decreasing the Lagrangian smoothing scale leads to additional small scale structures forming. Typically this allows particles only to move up on the morphology scale (void $\rightarrow$ pancake $\rightarrow$ filament $\rightarrow$ halo). This is also reflected in excursion set formalisms. In the EPS formalism particles can only change from "non-halo" to halo when increasing σ (by decreasing the Lagrangian smoothing scale). In the formalism that we propose here, particles can only change from single-stream ($\approx$ void in the picture) to multi-stream (pancake/filament/halo in the picture) when increasing σ. (This Figure was not part of the original publication.)

demonstrate this in Lagrangian space in Figure 3.4.

While the EPS formalism seems to describe the formation of haloes reasonably well, it does not try to explain the behaviour of the material outside of haloes. This material can either be part of string-like filaments, planar sheet-like "pancakes" or diffuse three dimensional single-stream regions. Single-stream regions are regions which have not yet collapsed on any scale.

We propose an excursion set formalism here that tries to predict properties of these single-stream regions. It considers a particle to be part of a single-stream region if it does not fulfil a collapse criterion on any length scale. In this case the collapse criterion does not mark the point where a particle becomes part of a halo, but it marks the point where it goes through its first caustic - which normally happens in a pancake or filament. We consider two models for the collapse criterion here (1) the Zel'dovich approximation and (2) a triaxial collapse model. Further we assume the stream density of particles to be given by these simple models if they are part of a single-stream region. This allows us to evaluate different statistics of single-stream regions.

Note that while the EPS formalism only requires following the random walk of the density, our formalism requires following the three eigenvalues of the deformation tensor. The idea of following the eigenvalues of the deformation tensor in the excursion set formalism has already been explored in previous work in the context of halo formation and its relation to the cosmic web(Chiueh & Lee, 2001; Sandvik et al., 2007). Here we will use it to learn more about single-stream regions.

3.3.2 The Zel'dovich approximation

As a first idea to model single-stream regions we consider the Zel'dovich approximation. Recall from section 2.1.3 that that the Zel'dovich approximation relates the comoving Lagrangian coordinates $\mathbf{q}$ at an initial time ($a = 0$) to the comoving Eulerian coordinates $\mathbf{x}$ at a later time

$$\mathbf{x}(a) = \mathbf{q} + D(a)\mathbf{s}(\mathbf{q}) \tag{3.1}$$

The comoving stream-densities at any scale factor a can then be evaluated as

$$\frac{\rho_s(a)}{\rho_0} = \left|\det\left(\frac{\partial \mathbf{x}}{\partial \mathbf{q}}\right)\right|^{-1} \tag{3.2}$$

$$= \left|\det\left(\mathbb{1} - D(a)\frac{\partial \mathbf{s}}{\partial \mathbf{q}}\right)\right|^{-1} \tag{3.3}$$

$$= |(1 - D(a)\lambda_1)(1 - D(a)\lambda_2)(1 - D(a)\lambda_3)|^{-1} \tag{3.4}$$

where $\lambda_1 \geq \lambda_2 \geq \lambda_3$ are the eigenvalues of the deformation tensor

$$d_{ij} = \frac{\partial s_i}{\partial q_j} \quad . \tag{3.5}$$

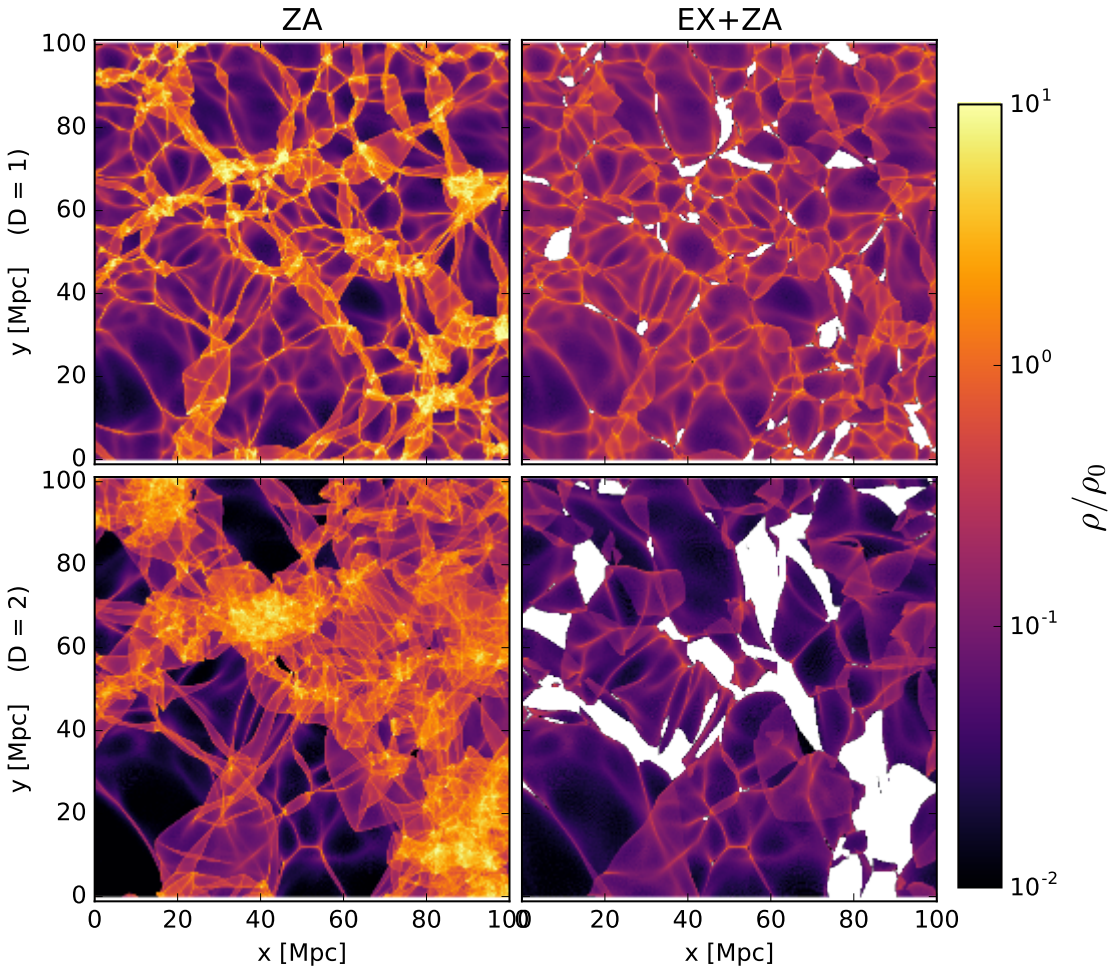

Figure 3.5: Left column: Evolution under the Zel'dovich approximation for a two dimensional density field using a power spectrum with normalization $\sigma = 2.36$ (as defined in (3.17)) at a growth factor of $D = 1$ (top) and $D = 2$ (bottom). Right: the same original density field, but cutting out particles that were classified as belonging to a multi-stream region with our excursion set formalism. The linearly evolved two point correlation function in this two dimensional test problem corresponds to that of a $m_X = 250$eV WDM cosmology where dark matter has a free-streaming scale of about 5 Mpc.

Within the Zel'dovich approximation a particle passes its first caustic when $D(a)\lambda_1 = 1$. The particle then becomes part of a structure which is collapsed in one dimension, but remains extended in the two other dimensions - commonly referred to as a pancake or wall-like structure. We use this as our first collapse criterion.

In Figure 3.5 we show the excursion set + Zel'dovich approximation (from now on EX+ZA) approach in practice for a two dimensional density field. Even though a large part of the density field has undergone shell-crossing, and is lacking any decelerating forces within the ZA, the ZA still appears to give a reasonable qualitative picture of regions that have been classified as single-stream regions within the EX+ZA formalism.

3.3.3 Triaxial collapse model

While the Zel'dovich approximation gives a good qualitative description for the behaviour of single-stream regions, it fails quantitatively in the non-linear regime as we shall see in section 3.4. We have therefore developed a triaxial model for the evolution of infinitesimal volume elements:

$$\dot{x}_i = a^{-2} p_i \tag{3.6}$$

$$\dot{p}_i = -\frac{4\pi G}{3}\rho_{\mathrm{bg}} a^{-1} x_i \left(\delta + \alpha(t)(3\lambda_i - \delta_0)\right) \tag{3.7}$$

$$\delta = \frac{1}{x_1 x_2 x_3} - 1 \tag{3.8}$$

for $i = 1, 2, 3$ where the x_i represent the individual Lagrangian to Eulerian expansion factors of the three principal axes of a volume element, and the p_i are the related momentum variables. λ_i are the eigenvalues of the deformation tensor, and $\delta_0 = \lambda_1 + \lambda_2 + \lambda_3$ is the linear density contrast at a scale factor $a = 1$. ρ_{bg} is the mean matter density of the universe today and δ is the relative over-density of the considered volume element. To not interrupt the flow of the text too much, we shifted the derivation of this model to the end of this chapter in section 3.8. It describes the general evolution of a single-stream volume element under the influence of external tidal forces. The time dependent factor $\alpha(t)$ parametrizes how the external tidal field grows with time. In linear theory $\alpha(t) = D(t)$. While this is certainly correct in early stages of evolution, it leads to strongly over-estimated tidal forces in the non-linear regime (see section 3.8.5). To limit the external tidal field in the non-linear regime we instead use

$$\alpha(t) = \frac{D(t)}{1 + |\delta_0| D(t)} \quad . \tag{3.9}$$

Note that other choices for $\alpha(t)$ are possible, and they lead to similar results. It is mostly important that linear theory is recovered for early stages and that the external tidal field becomes sub-dominant in the strongly non-linear regime.

This triaxial model is conceptually very similar to the ellipsoidal collapse model in Bond & Myers (1996). Both models try to follow the evolution of a volume element that is subject

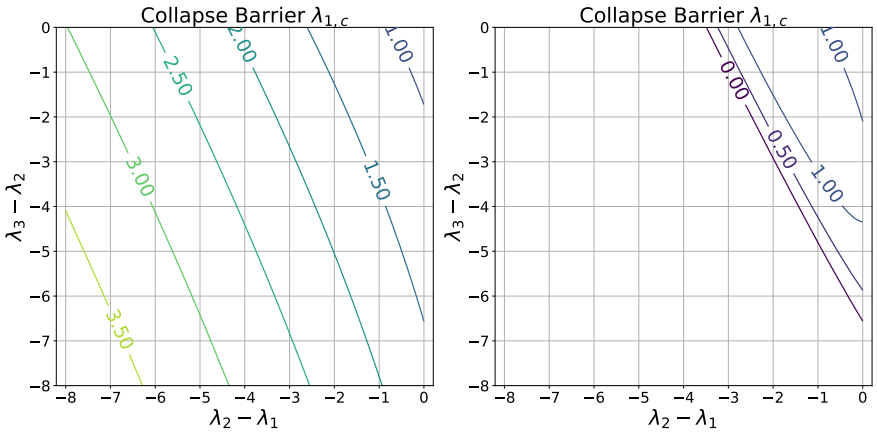

Figure 3.6: The collapse barrier for the biggest eigenvalue λ_1 of the deformation tensor as a function of the two smaller eigenvalues. If $\lambda_1 > \lambda_{1,c}$ a particle is predicted to be part of a collapsed structure by the triaxial collapse model. Left: with fading tidal field as in (3.9). If the two larger axes get stretched by the tidal field $\lambda_3 \leq \lambda_2 \ll 0$ the threshold $\lambda_{1,c}$ becomes larger and the collapse thereby more unlikely. Right: Triaxial collapse with the tidal field from linear theory ($\alpha(t) = D(t)$). The tidal field becomes too large and can lead to premature collapse in cases with $\lambda_3 \leq \lambda_2 \ll 0$. Note that the spherical collapse barrier is correctly reproduced for $\lambda_3 = \lambda_2 = \lambda_{1,c} \approx 1.68/3$.

to a tidal field that is described by the eigenvalues of the deformation tensor. However, the BM96 model assumes that the described perturbation has an ellipsoidal shape. This assumption can be dropped when speaking of an infinitesimal volume element which results in the simpler differential equations in (3.6) - (3.8).

To be able to evaluate the density prediction of the triaxial collapse model as a function of the eigenvalues of the deformation tensor $\rho_s(\lambda_1, \lambda_2, \lambda_3, a = 1)$ we integrate the equations of motion for a large set of parameters, store them in an interpolation table and interpolate it to the requested values later. Further we determine the collapse barrier which we define as the point where the smallest axis becomes $x_1 = 0.1$ and parametrize it as a threshold for the largest eigenvalue λ_1 depending on the two smallest eigenvalues of the deformation tensor. In Figure 3.6 we show this collapse barrier $\lambda_{1,c}(\lambda_2, \lambda_3)$ for the model with fading field (as defined in equation (3.9)) and for the tidal field from linear theory ($\alpha(t) = D(t)$). The model that uses the pure tidal field from linear theory can apparently lead to premature collapse for cases with strongly negative eigenvalues $\lambda_3 < \lambda_2 \ll 0$. For the excursion set formalism we therefore only consider the triaxial collapse model with fading tidal field from equation (3.9).

3.3.4 The six-dimensional random walk of the deformation tensor

Following Chiueh & Lee (2001) for a cosmological density field with rms density fluctuation σ a random realization of the deformation tensor can be generated by drawing six independent random variables $\{y_1, ...y_6\}$ from a normal distribution with dispersion σ and using the transformation

$$d_{11} = -\frac{1}{3}\left(y_1 + \frac{3}{\sqrt{15}}y_2 + \frac{1}{\sqrt{5}}y_3\right) \tag{3.10}$$

$$d_{22} = -\frac{1}{3}\left(y_1 - \frac{2}{\sqrt{5}}y_3\right) \tag{3.11}$$

$$d_{33} = -\frac{1}{3}\left(y_1 - \frac{3}{\sqrt{15}}y_2 + \frac{1}{\sqrt{5}}y_3\right) \tag{3.12}$$

$$d_{12} = d_{21} = \frac{1}{\sqrt{15}}y_4 \tag{3.13}$$

$$d_{23} = d_{32} = \frac{1}{\sqrt{15}}y_5 \tag{3.14}$$

$$d_{13} = d_{31} = \frac{1}{\sqrt{15}}y_6 \tag{3.15}$$

A random walk within the deformation tensor can then be constructed by choosing n intervals between $\sigma_0 = 0$ and $\sigma_n = \sigma_{max}$ and subsequently evaluating

$$d^{(k)} = d^{(k-1)} + \Delta d_k(\Delta\sigma_k^2) \tag{3.16}$$

where one starts with $d_{ij}^{(0)} = 0$ and draws each step a random Δd_k as explained above with a dispersion $\Delta\sigma_k^2 = \sigma_k^2 - \sigma_{k-1}^2$. Note that the only cosmology dependent parameter which enters this random walk is the standard deviation of the final unsmoothed density field

$$\sigma_{max}^2 = \frac{1}{2\pi^2}\int_0^\infty P(k)k^2dk \quad . \tag{3.17}$$

The value σ_{max} is unknown and depends on the free streaming cut-off of the considered dark matter model. Therefore, a measurement of the (unsmoothed) dark matter density within single-stream regions would directly constrain σ_{max}, allowing conclusions about the nature of the dark matter particle. Unfortunately this is unlikely ever to be possible.

In each step of the random walk we diagonalize the deformation tensor to obtain its three eigenvalues and test whether a collapse criterion is fulfilled $\lambda_1 \geq \lambda_{1,c}(\lambda_2, \lambda_3)$. For each random walk trajectory we save whether it has ever been outside of the barrier, and we also store the values of the eigenvalues λ_1,λ_2 and λ_3 at the final step σ_n. At the end of a random walk we assume for all single-stream particles (which have never crossed the barrier) that their density is given by

$$\rho_s = \rho_m(\lambda_1, \lambda_2, \lambda_3) \tag{3.18}$$

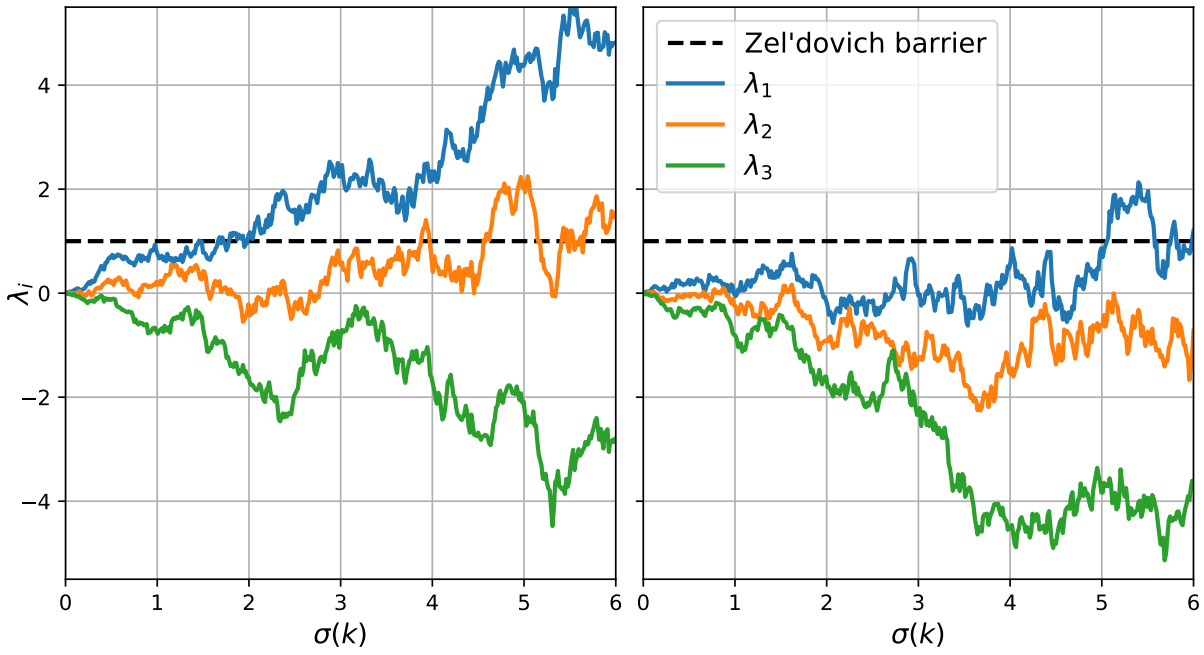

Figure 3.7: Two examples for randomwalks of the deformation tensor. Shown are the ordered eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3$ and the collapse barrier that is predicted by the Zel'dovich approximation. Left: a typical case. The collapse barrier is first crossed around $\sigma \sim 1.8$ - therefore this particle is only part of single-stream regions for dark matter models with $\sigma < 1.8$. Right: a rarer case where the barrier is crossed much later. This particle is considered to be part of a single-stream region for dark matter models with $\sigma < 5$. (This Figure was not part of the original publication.)

where ρ_m is

$$\rho_{za}(\lambda_1, \lambda_2, \lambda_3) = \frac{\rho_0}{(1-\lambda_1)(1-\lambda_2)(1-\lambda_3)} \tag{3.19}$$

for the Zel'dovich approximation and

$$\rho_{tc}(\lambda_1, \lambda_2, \lambda_3) = \frac{\rho_0}{x_1 x_2 x_3} \tag{3.20}$$

for the triaxial collapse model where the x_i have been evaluated via numerical integration.

We illustrate two examples of random walks of the deformation tensor together with the simple collapse barrier of the Zel'dovich approximation in Figure 3.7.

3.3.5 The thermal cutoff

The predictions of the excursion set formalism are determined entirely by the rms density fluctuation σ of the unsmoothed density field. The value of σ is not known and depends both on the physics of inflation and the particle physics properties of the dark matter. We will give estimates of σ for different dark matter models here.

Since the dark matter power spectrum has a slope close to -3 on small scales, the integral in (3.17) is nearly logarithmically divergent for a generic cold dark matter model that ignores the effects of free streaming (e.g. the dashed line in Figure 3.8). However, most dark matter models exhibit a cut-off of the power spectrum on small scales, since small scale pertubations are smoothed out either by the effect of thermal free streaming or by quantum effects. The value of σ will depend strongly on the scale of the cutoff. Here we consider cutoff models for warm dark matter (WDM), WIMP-based cold dark matter and axion-based cold dark matter. For the warm dark matter models we use the cold dark matter power spectrum parametrization of Eisenstein & Hu (1999) and apply the warm dark matter thermal cutoff from Bode et al. (2001). For the WIMP models we use the Eisenstein & Hu (1999) spectrum for $k \leq 10^2 h\text{Mpc}$ and use the small scale parametrization of Green et al. (2005) for $k > 10^2 h\text{Mpc}$. We normalize the Eisenstein & Hu (1999) spectrum to the cosmological value of σ_8 and choose the normalization of the Green et al. (2005) to match $P_{e,h}(k = 10^2 h\text{Mpc}) = P_g(k = 10^2 h\text{Mpc})$. For the axion model we use the same cut-off parametrization as in the WDM models, with the effectively rescaled mass relation from Marsh (2016) equation (118). We present the dimensionless power spectra

$$\Delta(k) = P(k)\frac{k^3}{2\pi^2} \tag{3.21}$$

and the values of

$$\sigma(k) = \left(\int_0^{\ln(k)} \Delta(k') d\ln k'\right)^{1/2} \tag{3.22}$$

for four different dark matter models in Figure 3.8. This leads us to the total rms density fluctuations σ which are listed in Table 3.2. Note that these models are just intended to give a rough impression of the range of possible values for σ which is quite large given the weak current constraints on the non-gravitational properties of dark matter.

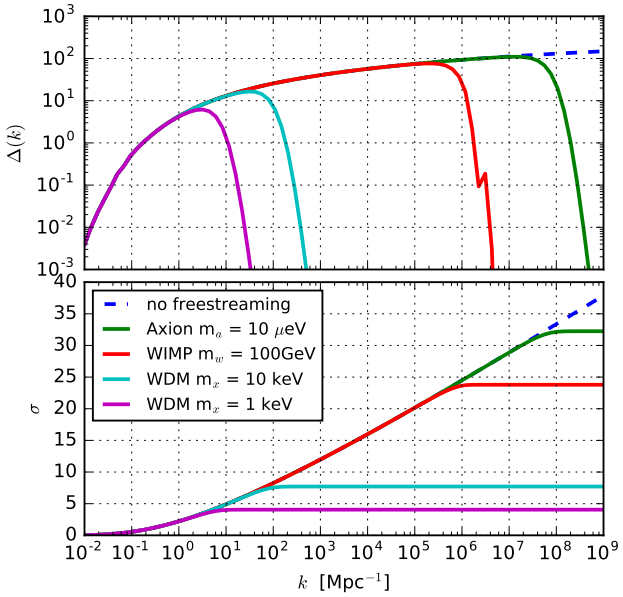

Figure 3.8: Top: the dimensionless linear power spectrum (top) for different different dark matter models. Bottom: the σ values obtained by integrating the power spectra up to a scale k according to (3.22).

3.4 Test on simulations

To evaluate quantitatively whether the excursion set formalism produces reasonable results, we compare it with a 1keV warm dark matter simulation with 512^3 particles in a box with a side-length of $L = 20$Mpc. The simulation is carried out using a modified version of Gadget 3 (Springel, 2005b) using the HA16 scheme (Hahn & Angulo, 2016) without refinement. The HA16 uses phase-space interpolation techniques to obtain a high quality density estimate that allows running fragmentation-free warm dark matter simulations. Additionally we track the distortion tensor $\frac{d\mathbf{x}}{d\mathbf{q}}$ with the Geodesic Deviation Equation (Vogelsberger & White, 2011) and determine the point in time when a particle goes through its first caustic where $\det(\mathrm{d}\mathbf{x}/\mathrm{d}\mathbf{q}) = 0$. Further we use the distortion tensor to get accurate estimates of the stream densities of individual particles $\rho_s = 1/|\det \frac{\partial \mathbf{x}}{\partial \mathbf{q}}|$.

When creating the initial conditions for the simulation we also evaluate the excursion set predictions for the corresponding density field by evaluating the deformation tensor of every particle when using sharp k-space filters with different length scales.

To test whether the Zel'dovich approximation and the Triaxial Collapse model give reasonable predictions for the densities of single-stream particles we show in Figure 3.9 the measured versus the predicted stream densities of a bootstrapped sample of 1000 particles which have not gone through a caustic in the simulation, and are therefore likely to belong to a single-stream region. The Zel'dovich approximation underpredicts the stream densities of single-stream particles systematically. Note that Sahni & Shandarin (1996) have shown that higher order Lagrangian perturbation theories even worsen the accuracies at such late

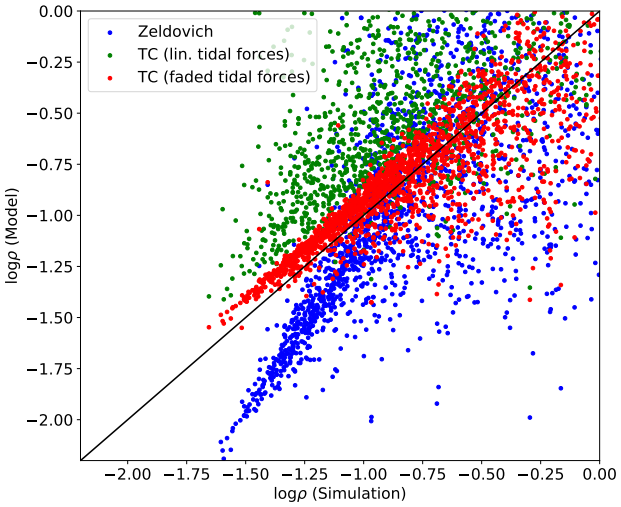

Figure 3.9: Predicted versus measured stream densities for particles that have not crossed a caustic at $a = 1$ for different models. In blue: the Zel'dovich approximation which shows a tight correlation, but systematically underpredicts the stream densities. In green: the triaxial collapse model with the tidal field from linear theory. It has a much larger scatter and seems to over-predict the densities. In red: the triaxial collapse model with fading tidal field as in (3.9). This results in a tight correlation with the right slope, but a small offset (over-prediction) in the densities.

times as we are investigating here. The triaxial collapse model with the tidal field from linear theory seems to over-predict stream densities, and has a much larger scatter than the pure Zel'dovich approximation. This is likely due to the tidal field from linear theory being unreasonably large in the non-linear regime. However, the model with fading tidal field performs significantly better - showing a tight correlation with the right slope, but a slight over-prediction of the densities. There is in principle room for improvement by considering more sophisticated descriptions of the tidal field, but we expect this to be good enough for a first estimate of the density distribution.

Note, while spherical evolution models have been shown to describe the evolution of densities in under-dense regions well (Neyrinck, 2013), these are not appropriate for our purposes here, since they do not determine the point of first caustic crossing. However, for the interested reader we want to mention that we find that in comparison to the triaxial collapse model, spherical collapse performs similarly well on the lowest densities, but slightly under-predicts densities in the intermediate range.

As the second ingredient to the excursion set formalism, the performance of the collapse criterion has to be tested in practice. Therefore we show the in Figure 3.10 the classification of particles into single-stream (colored) and multi-stream (grey) in a slice through Lagrangian space. Further we color code the predicted stream densities of single-stream particles. The EX+ZA approach apparently generally underpredicts the stream densities of single-stream particles and assigns too little material to single-stream regions. This means

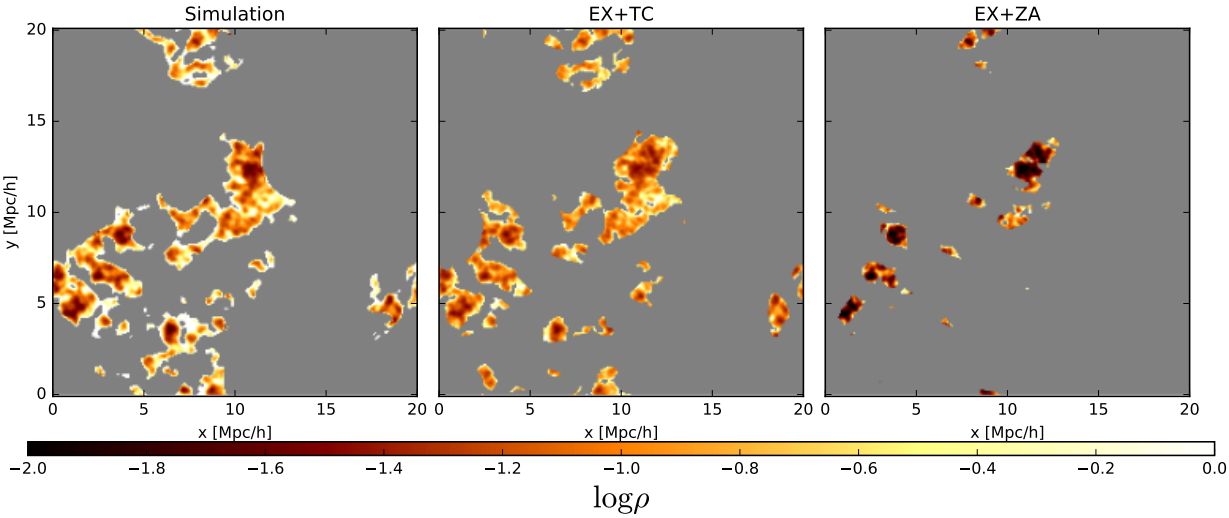

Figure 3.10: A slice through Lagrangian space showing the stream densities of particles that are considered part of a single-stream region by different approaches. Left: Stream density measured in the simulation for particles that have not yet crossed a caustic. Center: Prediction from the EX+TC approach. Right: Prediction from the EX+ZA approach. The EX+TC agrees fairly well with the simulation. The EX+ZA approach underpredicts the amount of single-stream material and underpredicts the stream densities.

that the collapse barrier of the Zel'dovich approximation $\lambda_{1,c} = 1$ is too tight in most cases. However, the agreement between simulation and EX+TC is fairly good. This shows that the collapse barrier for single-stream particles should indeed depend on the the two other eigenvalues as in Figure 3.6 left. If e.g. two axes of a volume element are expanding rapidly (λ_2, $\lambda_3 \ll 0$), the tidal field needs to act much more strongly than in the Zel'dovich case, to bring the smallest axis to collapse.

As a final benchmark of the scheme, we compare in Figure 3.11 the volume weighted density distribution that is predicted by the excursion set formalisms with the distribution that is actually measured in the simulation. The simulation densities are here measured on a mesh in Eulerian space. The high density tail in the simulation is due to multi-stream regions, and it is not expected to be captured by the models. Apparently the EX+ZA formalism strongly underpredicts the densities whereas the EX+TC formalism tends to slightly over-predict densities. Our predictions are not of high accuracy, but they are good enough to provide a reasonable picture and a first estimate of the density distribution. In principle the quality of the triaxial collapse predictions could be improved by finding a more precise description of the tidal field in the non-linear regime. For simplicity however, we stick to our EX+TC model here.

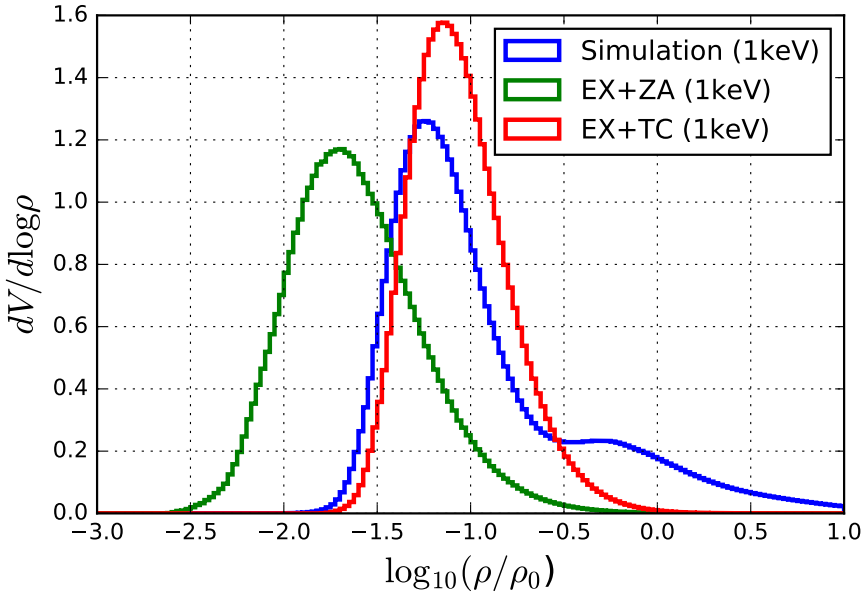

Figure 3.11: The normalized (volume-weighted) density distribution predictions for single-stream regions by the EX+ZA and EX+TC formalisms in comparison to the measured distribution in a simulation. The high density tail ($\rho \gtrsim 10^{-0.5}\rho_0$) in the density distribution of the simulation mostly originates from multi-stream regions (pancakes, filaments, haloes). These are by construction not represented in the excursion set formalisms.

3.5 Predictions

In this section we will use the excursion set formalisms to predict the mass- and volume-fractions of single-stream regions and their density distribution. This distribution should approximately correspond to the volume-weighted density distribution of the universe, except for the high density tail which is dominated by multi-stream regions.

We create a sample of $4 \cdot 10^7$ random walks which we evaluate at

$$\sigma_k = \frac{k}{n}\sigma_{\max} \tag{3.23}$$

for $n = 3000$ steps and $\sigma_{\max} = 30$. We determine at each step the eigenvalues of the deformation tensor λ_{1-3} and test whether one of the Zel'dovich or the triaxial collapse criteria is fulfilled. Further we compute several statistics among the particles which are still classified as belonging to a single-stream region. We present these statistics here. σ is the only free parameter. Thus, a prediction for a dark matter model today can be read off at its value of σ_{Model} and a prediction for the model at an earlier time can be simply obtained by reading off at

$$\sigma(a) = \frac{D(a)}{D(a=1)}\sigma_{\text{Model}} \tag{3.24}$$

Table 3.2: Predictions which are obtained from the excursion set + Zel'dovich approximation (EX+ZA) approach and the excursion set + triaxial collapse (EX+TC) approach evaluated for different dark matter models.

Dark Matter Model	WDM (1keV)	WDM (10keV)	WIMP (100GeV)	Axion (10μeV)
Linear Theory rms density perturbation $\sigma_{\rm Model}$	4.1	7.7	23.8	32.2
Single-Stream Mass Fraction (EX+ZA)	$3.98 \cdot 10^{-2}$	$6.39 \cdot 10^{-3}$	$2.31 \cdot 10^{-4}$	$9.54 \cdot 10^{-5}$
Single-Stream Mass Fraction (EX+TC)	$1.11 \cdot 10^{-1}$	$3.27 \cdot 10^{-2}$	$3.09 \cdot 10^{-3}$	$1.58 \cdot 10^{-3}$
Median Density (EX+ZA)	$2.18 \cdot 10^{-2}$	$3.47 \cdot 10^{-3}$	$1.22 \cdot 10^{-4}$	$4.91 \cdot 10^{-5}$
Median Density (EX+TC)	$8.15 \cdot 10^{-2}$	$2.93 \cdot 10^{-2}$	$4.41 \cdot 10^{-3}$	$2.55 \cdot 10^{-3}$
Single-Stream Volume (EX+ZA)	1.11	1.15	1.21	1.22
Single-Stream Volume (EX+TC)	1.01	0.84	0.52	0.45

For σ-dependent plots we add the corresponding scale factor of the universe (assuming $\Omega_\Lambda = 0.7$ and $\Omega_m = 0.3$) for the WIMP model with $\sigma = 23.8$ as an alternative x-scale. We also list some of the predictions for three selected dark matter models in Table 3.2.

3.5.1 Mass- and volume- fractions

In Figure 3.12 we present the mass fraction of single-stream regions with respect to the total mass for the EX+ZA and the EX+TC models. Further we show the amount of material which is predicted not to be part of a halo for the EPS model with barrier $\delta_{sc} = 1.686$ and the EPS model with the moving barrier described in Sheth et al. (2001) with the parameterization that is used in Angulo & White (2010)

$$\delta_{ec}(\sigma) = \sqrt{q}\delta_{sc}\left(1 + \beta\left(\frac{\sigma^2}{q\delta_{sc}^2}\right)^\gamma\right) \tag{3.25}$$

with $q = 0.5$, $\beta = 0.55$, $\gamma = 0.5$.

Apparently only a small fraction of the material which is not in haloes ($\sim$ 10% for a WIMP) is actually in single-stream regions (0.3% for a WIMP). Most of it is part of filaments and pancakes. Although only a small fraction of the total mass is part of single-stream regions, most of the volume is occupied by them.

We present the volume fractions that are obtained from the excursion set formalisms in Figure 3.13. They are calculated as

$$\frac{V_{ss}}{V_{tot}} = \frac{1}{N_{tot}} \sum_{i \in \text{singlestream}} \frac{\rho_0}{\rho_i} \tag{3.26}$$

where N_{tot} is the total number of particles and the sum only goes over those particles which are considered to be part of single-stream regions. We do not expect these values to be precise, since the density estimates are still relatively crude and the classification operates in Lagrangian Space and some of the particles which are classified as single-stream will actually be in multi-stream regions in Eulerian space. It provides however a benchmark for the scheme, since the total volume of the single-stream regions should be of order

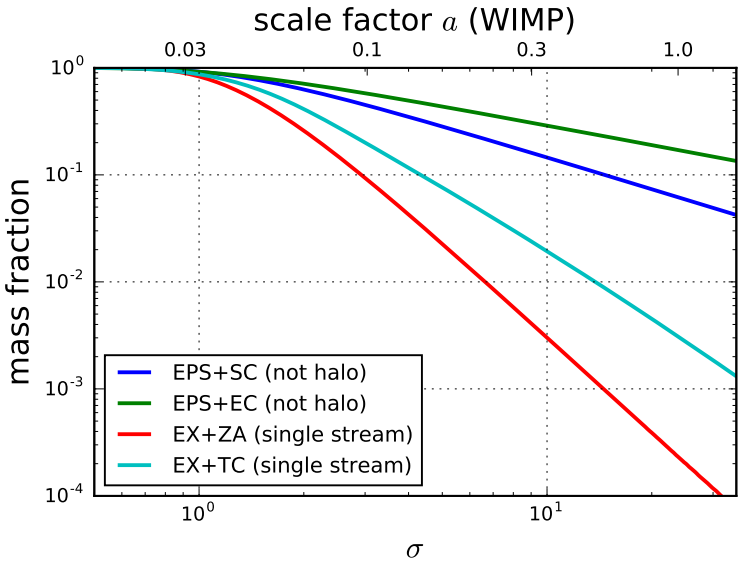

Figure 3.12: Evolution of the mass fraction of single-stream regions with respect to the total mass (red and cyan). Also indicated the amount of mass that is not part of the halo for the EPS with spherical collapse barrier (blue), and with the moving barrier described in (3.25) (green). At $\sigma = 0$ nothing is collapsed and everything is a single-stream region. With increasing density perturbations more and more mass collapses and gets removed from the singlestream regions.

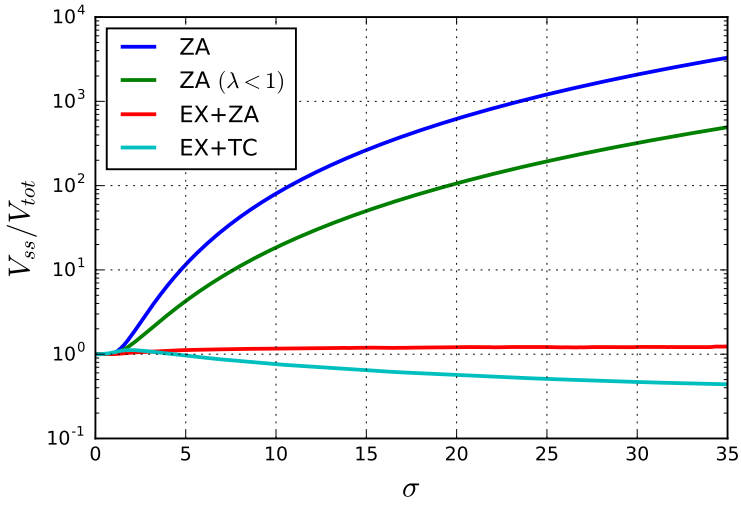

Figure 3.13: Evolution of the volume fraction of single-stream regions with respect to the total volume. For a pure Zel'dovich approximation the volume over counting increases dramatically for $\sigma \gg 1$ and even reaches a factor of 400 around $\sigma = 30$. With the excursion set formalisms the volume fractions stay of order unity. However, we do not expect the volume fractions to be accurate enough for a quantitative prediction of the single-stream volume. They can merely serve as a consistency check here.

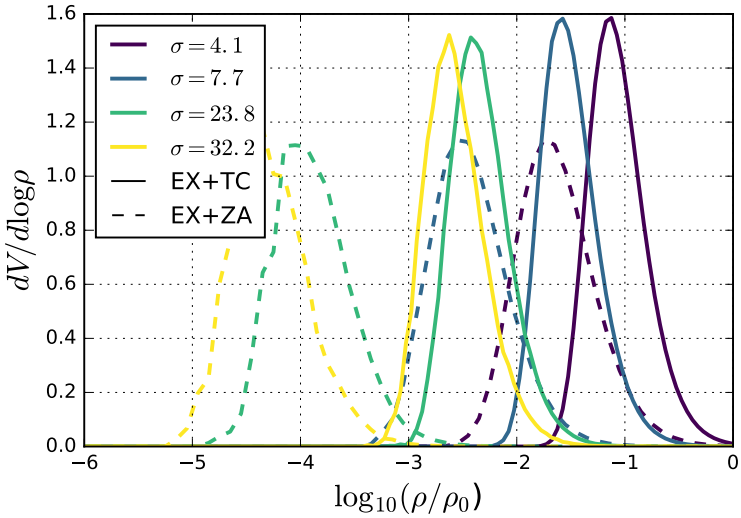

Figure 3.14: Volume weighted density histogram for the single-stream regions for four different dark matter models: WDM (1keV) in black, WDM (10keV) in blue, WIMP (100GeV) in green and an axion (10μeV) model in yellow. All distributions are normalized to 1. The dashed lines represent the EX+ZA prediction and the solid lines the EX+TC model prediction. We expect the true distributions to lie a bit to the left of the EX+TC predictions.

unity. This is indeed the case for the EX+ZA scheme which is at about $V_{ss} = 1.19V_{tot}$ at $\sigma = 23.8$. This is remarkable since the densities tend to be underestimated within this scheme (compare Figure 3.11) and too many particles are classified as collapsed (compare Figure 3.10). However these effects seem to balance out so that an Eulerian volume of order unity is achieved. To emphasize this we also show the total volume of the sheet which is obtained in a pure Zel'dovich approximation when running the sum of (3.26) over all particles and when only selecting particles that at the last step of the randomwalk have all eigenvalues with $\lambda < 1$. In both of these cases the volume is overestimated by several orders of magnitude. This is due to shells expanding after shell-crossing without any de-accelerating forces as we have already seen on the left side of Figure 3.5. The EX+TC scheme seems to underpredict the volume a bit which leads to $V_{ss} \sim 0.52V_{tot}$ at $\sigma = 23.8$. This is probably due to the densities being slightly overestimated as we have seen in Figure 3.11. However, the deviation is still well below a factor of two and therefore unproblematic for our current approximate estimates.

3.5.2 Density distribution

We present the volume weighted density distribution for the single-stream regions of the four considered dark matter models in Figure 3.14. For each model we present the EX+ZA and the EX+TC prediction. From Figure 3.11 we would expect that the EX+ZA scheme significantly underpredicts the densities and the EX+TC distribution slightly over predicts

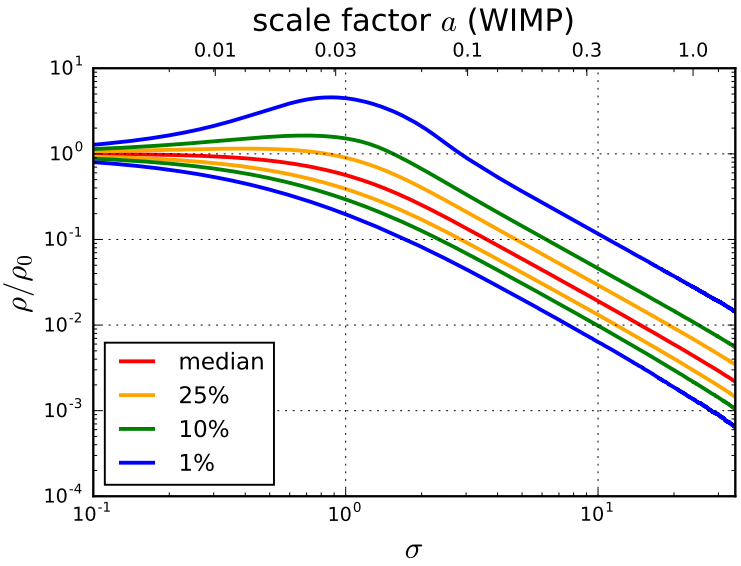

Figure 3.15: Quantiles of the volume weighted density distribution for the EX+TC model as a function of the linear theory rms density perturbation σ. The upper x-scale shows the corresponding scale factors for a ΛCDM cosmology with a 100 GeV WIMP as dark matter. While for $\sigma \lesssim 0.2$ the distribution is symmetric around the mean density $\rho = \rho_0$, for $\sigma \gg 1$ the density distribution is dominated by the under dense regions.

them. The true density distribution probably lies slightly to the left of the EX+TC distributions. Interestingly the shape of the distributions changes little as the variance of the density field increases; rather their center just shifts to lower densities. We list the median densities of the distributions in Table 3.2. The EX+TC scheme predicts for a WIMP of mass 100 GeV a median density of $\sim 4 \cdot 10^{-3}\rho_0$. Following our argumentation from above, if the dark matter is indeed such a WIMP, this should roughly be the *median density of the universe*.

Further we show the evolution of the quantiles of the distribution for the EX+TC scheme in Figure 3.15. Note how the the normal distribution around $\rho = \rho_0$ is reproduced in the linear regime for $\sigma \lesssim 0.2$. In this regime the median density is approximately the mean density ρ_0, since over- and underdense regions enter the distribution roughly with the same weight. For the WIMP model the volume density distribution starts deviating significantly from the linear case for $a \gtrsim 0.01$. In the strongly non-linear regime $\sigma \gg 1$ the median density roughly scales like $\rho \propto \sigma^{-3/2}$. At that point the volume density distribution is completely asymmetric, and dominated by underdense regions.

3.6 Percolation

Whereas above we focused on predicting the density distribution, and the volume- and mass-fractions of single-stream regions, we will try to find out more about their topology

in this section. A profound question is whether it actually makes sense to speak of individual single-stream regions, or whether almost all single-stream volume is connected in one universe filling complex. That is, *do single-stream regions percolate?*

3.6.1 Previous work

This question has already been investigated by Falck & Neyrinck (2015) and Ramachandra & Shandarin (2017) who found independently that the single-stream regions in their simulations do indeed percolate.

Falck & Neyrinck (2015) use simulations of resolution up to $L/N^{1/3} = 0.2h^{-1}\mathrm{Mpc}$ (boxsize over particles per dimension). As can be read off from Figure 3.8 this corresponds to a resolved rms density perturbation of approximately $\sigma(k_{\mathrm{Ny}} \sim 16h\mathrm{Mpc}^{-1}) \sim 5$. They use the `ORIGAMI` (Falck et al., 2012) method to identify single-stream regions in their simulations. Each particle is tested for shell-crossing along three orthogonal Lagrangian axes. A single-stream particle is then a particle with no shell-crossing along any axis.

Ramachandra & Shandarin (2017) use simulations with resolutions up to $L/N^{1/3} = 0.39h^{-1}\mathrm{Mpc}$. They identify single-stream regions by testing that the Lagrangian Submanifold has not gone through any foldings at each given location.

Despite their different methods, both Falck & Neyrinck (2015) and Ramachandra & Shandarin (2017) find that in all their tests the single-stream regions appear to percolate. Further Falck & Neyrinck (2015) find that the mass fraction of single-stream regions decreases when going to higher resolution, whereas the volume fraction does not decrease significantly. From this, they infer that single-stream regions may percolate even in the case of infinite resolution.

We argue here, however that the fact that the mass fraction of single-stream regions decreases significantly with increasing resolution (or equivalently: with increasing σ), makes it quite unlikely that single-stream regions percolate: In single-stream regions the mapping from Lagrangian space $\mathbf{q}$ to Eulerian space $\mathbf{x}$ is one to one (per definition). Therefore if single-stream regions form an infinitely large percolating structure in Eulerian space, they also form an infinitely large percolating structure in Lagrangian space. Now, at infinite resolution only a tiny fraction of the mass is part of single-stream regions (e.g. 0.3% for a WIMP). That means single-stream regions would have to be percolating in Lagrangian space at a (Lagrangian) volume fraction of $\sim 0.3\%$ (for a WIMP).

While we agree that single-stream regions are likely to percolate as long as they contain significant fractions of the total mass, we expect them to stop percolating when they reach a low enough mass fraction. As already shown in Figure 3.12 the mass fraction is only dependent on σ, and therefore we ask the question: *Do single-stream regions always percolate, and if not, at which value of σ do they stop percolating?*

3.6.2 Percolation with the excursion set formalism

We try to search the percolation threshold of single-stream regions with the EX+TC scheme here. As already shown in Figure 3.10, the EX+TC scheme performs fairly well in pre-

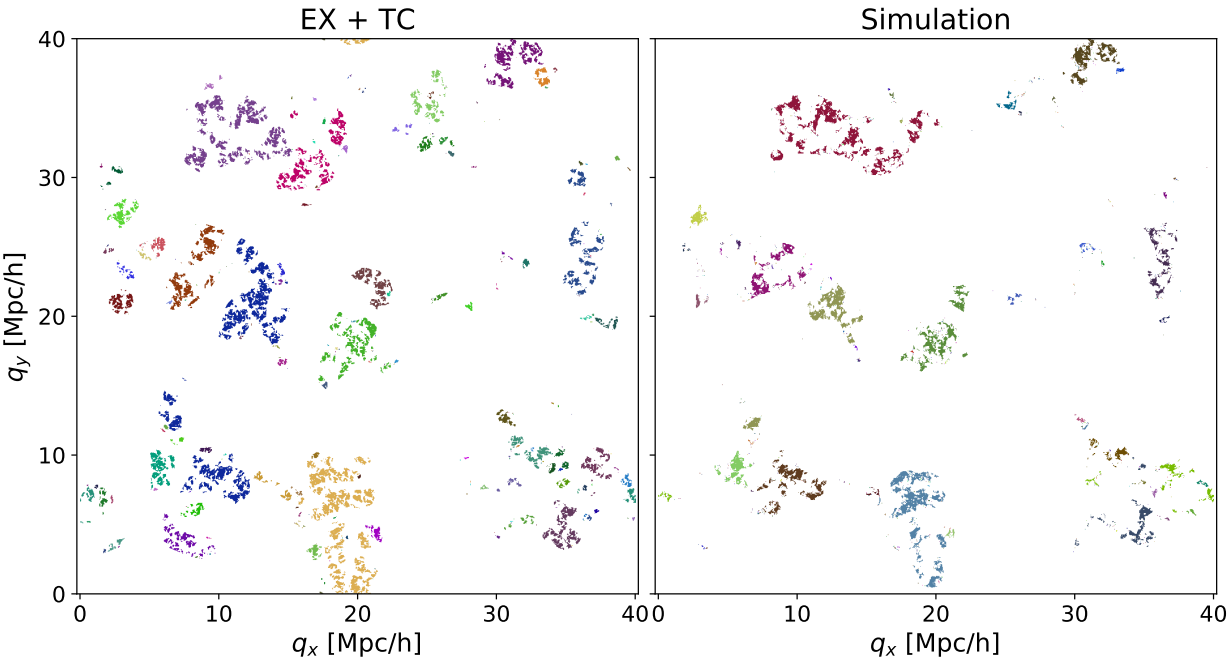

Figure 3.16: Lagrangian slice classifying particles into single-stream (colored) and multi-stream (white) for a cosmological box with $\sigma = 6.4$. Each distinct connected component is assigned a random color. Left: Classification according to the EX+TC excursion set prediction. Right: Classification inferred from a simulation. The particles are classified according to the nearest grid point in Eulerian Space as explained in section 3.6.3. In both cases the single-stream regions do not percolate in any dimension.

dicting which particles are going to become part of a single-stream region. Therefore we can use it to test for the percolation of single-stream regions in Lagrangian space to get an estimate of the percolation threshold. It is computationally far cheaper to evaluate than to run a whole simulation.

We set up a Lagrangian mesh of 1024^3 particles within a periodic box of a width of $L = 40\text{Mpc/h}$. We use the same power spectrum as shown in Figure 3.8 with a sharp k-space cutoff at $k_c = \frac{2\pi}{0.1Mpc} = 63\text{Mpc}^{-1}$ leading to $\sigma_0 = 6.4$. For each particle the EX+TC scheme is used to determine the growth factor at collapse D_{col}. In Figure 3.16 (left) we show a thin slice through Lagrangian space for particles which have $D_{\text{col}} > D(a = 1) = 1$, i.e. particles which are single-stream at $a = 1$.

We test for different growth factors D_{thresh} whether the boolean grid defined by $D_{\text{col}} > D_{\text{thresh}}$ percolates. To test for percolation, we determine all connected single-stream regions in Lagrangian space, linking the cells over the faces to their 6 nearest Langrangian neighbours. A region percolates if there is a path between any cell and any of its periodic replications that does not exit the connected component.

We find the single-stream regions to percolate only for $D_{\text{thresh}} < 0.77$ - corresponding to $D \cdot \sigma_0 = 4.9$ and a single-stream mass fraction of 8.0%. Although there might be a dependence of the percolation threshold on the precise shape of the power spectrum and also a residual realisation dependence (due to the small box size), we expect these influences to be small in comparison to the importance of the mass fraction which exclusively depends on the value of $D \cdot \sigma_0$.

So, assuming that these factors are small, we predict with the EX+TC scheme that single-stream regions do not percolate if

$$\sigma_{\text{DM}}(a) \gtrsim 5 \tag{3.27}$$

which corresponds for a cold dark matter simulation roughly to an inter-particle spacing of $L/N^{1/3} = 0.2h^{-1}\text{Mpc}$. It is an unfortunate coincidence that this value is so close to the resolution limit of Falck & Neyrinck (2015) and Ramachandra & Shandarin (2017). Possibly their single-stream regions are not far away from ceasing to percolate - at one or two resolution levels higher they might have found them to no longer percolate.

However, it is not clear whether the EX+TC prediction is accurate enough for this to be a reliable conclusion. Fortunately $\sigma \gtrsim 5$ is still in a regime which is accessible through classical dark matter simulations so that it can be tested numerically.

3.6.3 Percolation in a simulation

We attempt to test the prediction that single-stream regions do not percolate for $\sigma \gtrsim 5$ with a simulation. Therefore we use the same particle grid that has been used for the EX+TC calculation ($\sigma = 6.4$) to create initial conditions with the Zel'dovich approximation at a scale factor of $a = 0.01$. We run a classical N-Body simulation using the Planck ΛCDM cosmology $\Omega_m = 0.30$, $\Omega_\Lambda = 0.69$ and $h = 0.68$ (Planck Collaboration et al., 2016) up to a scale factor of $a = 1$. Then we use only the positions of the particles at the final time to

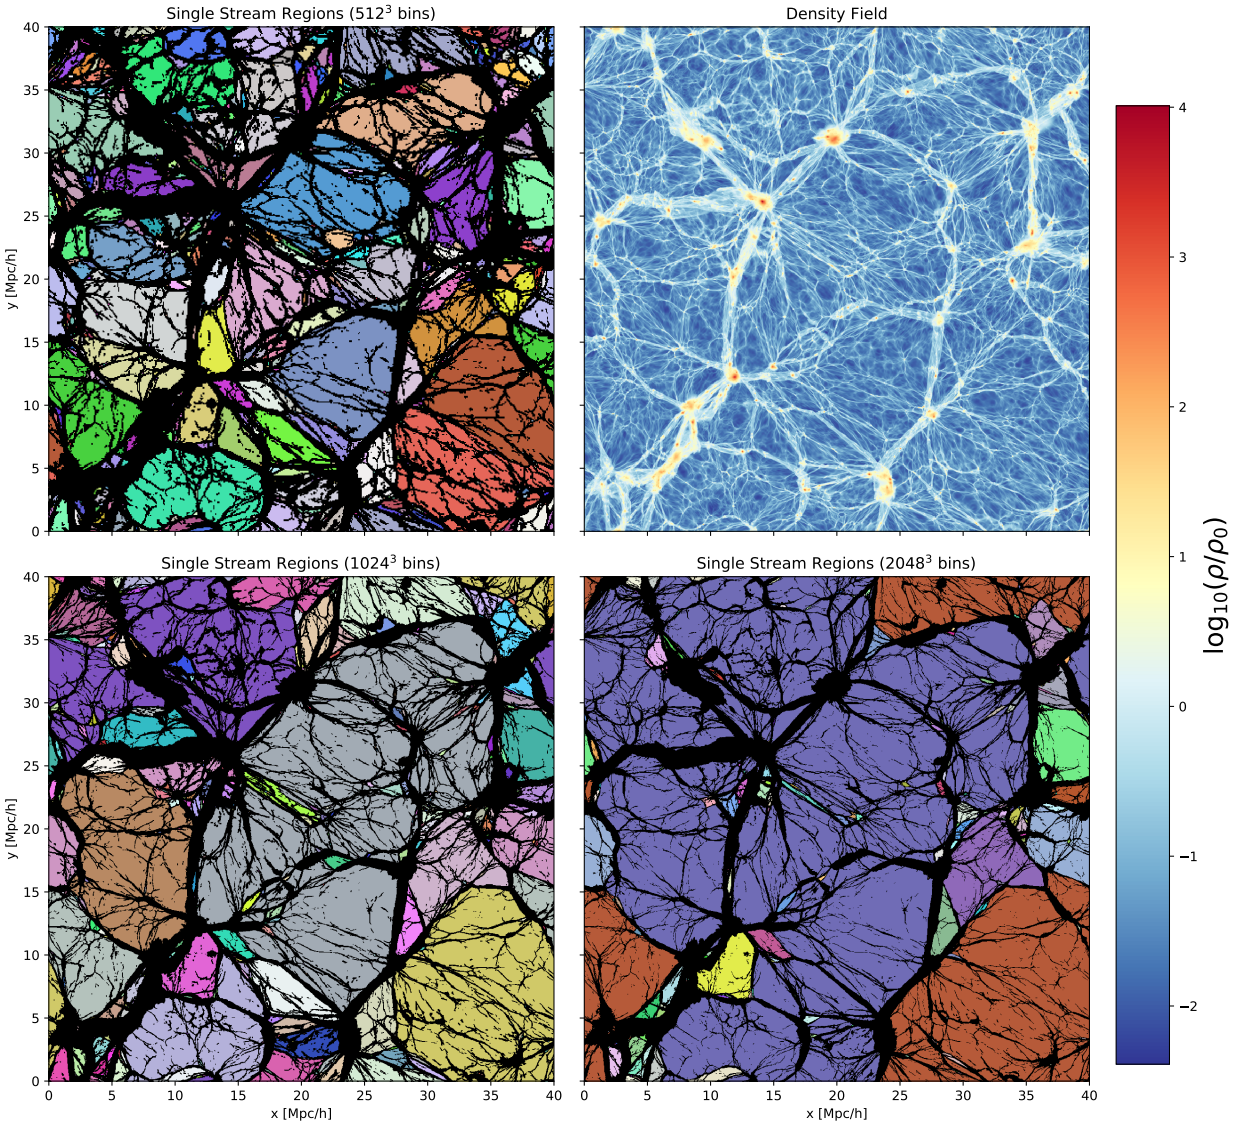

Figure 3.17: A thin slice (width 40 kpc) through a cosmological ΛCDM simulation. Top right: density field (logarithmic). Other panels: Classification into multi-stream regions (in black) and single-stream regions (color) for different different grid resolutions: top left 512^3 bins, bottom left 1024^3 and bottom right 2048^3. Each distinct single-stream region is assigned a random color. Many of the connected regions do not appear to be connected within this slice, but are connected through the third dimension. In the cases with 512^3 and 1024^3 bins the single-stream regions do note percolate. In the 2048^3 case there is one region (dark blue) which percolates in the y-dimension, but not in the x- and z-dimensions. Note that since this is a thin slice through Eulerian space, most structures that appear string-like are slices through pancakes. We provide movies that scroll through the z-coordinate of these slices under wwwmpa.mpa-garching.mpg.de/paper/singlestream2017/percolation.html.

calculate which regions of the space are single-stream regions and which are multi-stream regions.

Single-stream classification

We use the trigonometric resampling scheme as described in sections 2.4.3 and 2.6.3 to create from the 1024^3 original particles a much larger number (32768^3) of re-sampled particles which we bin onto a 512^3, a 1024^3 and a 2048^3 cubic mesh to infer a high quality density field. Additionally, we determine for the re-sampled particles the determinant of the real space distortion tensor

$$\mathbf{D}_{\mathrm{xq}} = \frac{\partial \mathbf{x}}{\partial \mathbf{q}} \tag{3.28}$$

where $\mathbf{x}$ are Eulerian and $\mathbf{q}$ are Lagrangian coordinates. We determine all bins which contain any re-sampled particles that have a negative determinant and classify them as multi-stream regions. Intuitively this can be understood for the one-dimensional case in Figure 3.3: if one follows along the sheet and has to go left $\frac{\partial x}{\partial q} < 0$ at any point one knows that one has to be in a multi-stream region. For a more detailed description please refer back to sections 2.4.3 and 2.6.3. Note that this way of determining single-stream regions is very robust and simple to implement. It can in principle be used as a pure post-processing step on any cosmological dark matter simulation which uses grid like initial conditions.

We show a slice through the determined density and single-stream fields in Figure 3.17. We color distinct single-stream regions in random colors.Every cell is linked with its 6 nearest neighbouring cells (no diagonals). Additionally we provide a set of movies and additional material at the following address:
wwwmpa.mpa-garching.mpg.de/paper/singlestream2017.
This includes, for example, movies which scroll through the z-coordinate over time. Visualizing the whole three dimensional volume this way helps to understand how regions are connected.

Percolation in Eulerian space

In the case with 512^3 bins we clearly find many individual single-stream regions that do not percolate. However, if we increase the number of bins and therefore decrease the bin size, many of the originally independent regions connect together to larger regions. This leads to much larger typical single-stream regions in the case with 1024^3 bins and even percolation in one dimension (as we will discuss below) in the case with 2048^3 bins. That means that single-stream regions with diameters of a few Mpc get linked together through gaps smaller than a binwidth of 39kpc.

This can also be seen in the volume histograms of single-stream regions which we show in Figure 3.18. In all three cases, as well as in the excursion set prediction, the number count of single-stream regions follows a power law with a slope $n(V) \propto V^{-0.5}$ up to a largest region. However, while there seems to be no preferred largest region in the 512^3

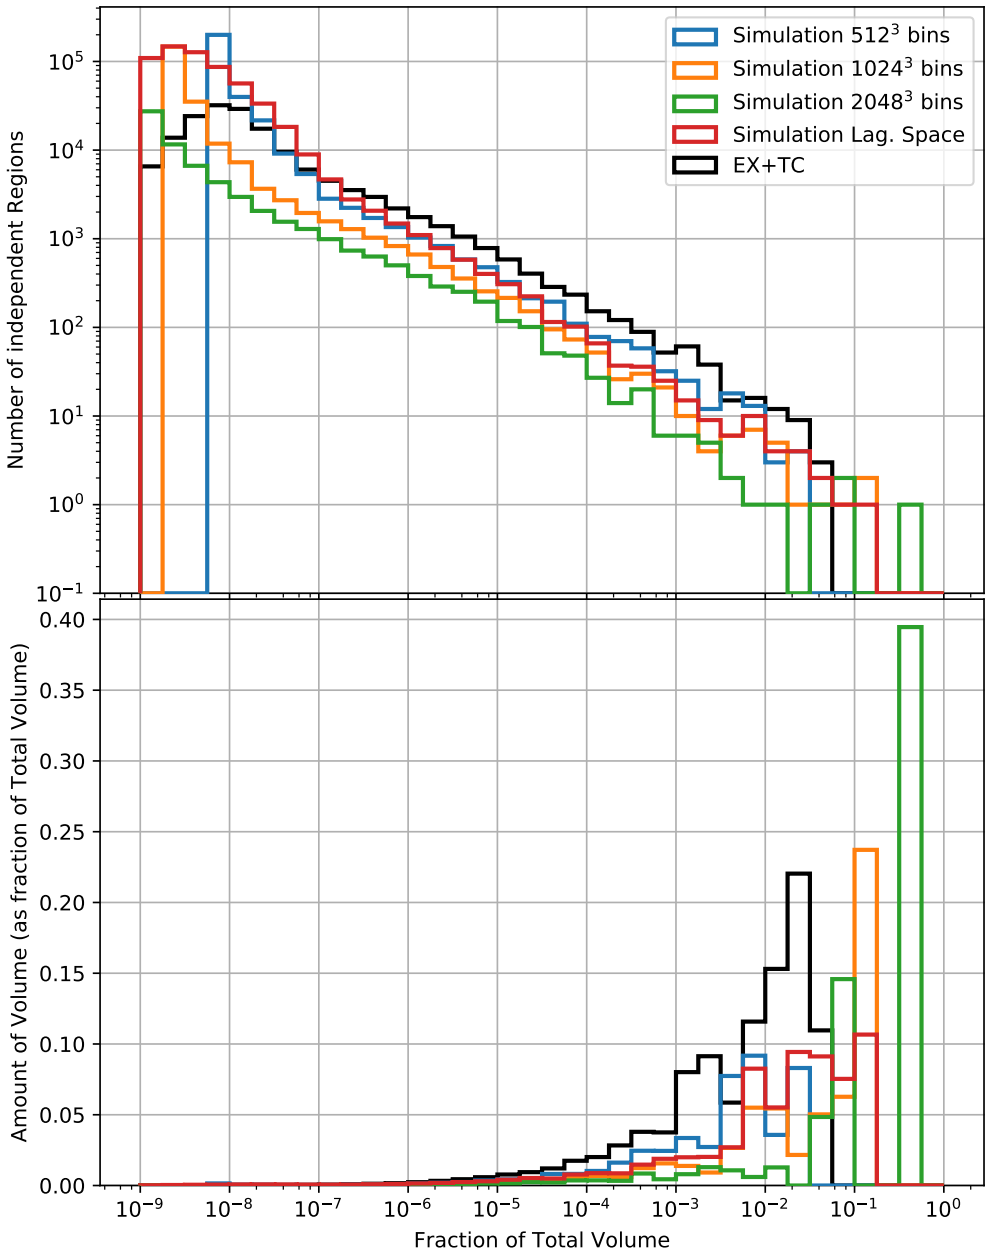

Figure 3.18: Top: Histogram of the Number of independent single-stream regions of a given size. The x-scale is given as fraction of the total volume $V_{\rm tot} = (40{\rm Mpc}/h)^3$. Bottom: Volume weighted histogram. Apparently in the case with 2048^3 bins there is a single-stream region which takes about 40% of the total volume and which percolates in one dimension. The areas under the curves correspond to the total single-stream volume and are listed in Table 3.3.

Table 3.3: Volume and mass-fractions of single-stream regions. Note that for the excursion set case (EX+TC) the numbers given here are for the Eulerian volumes corresponding to connected zero caustic crossing regions in Lagrangian space. These are expected to be larger than Eulerian single-stream regions and to sum to somewhat more than the total (Eulerian) volume.

bins	binwidth	total ss. vol.	total ss. mass	largest component vol.	second largest component vol.	percolation
512^3	78 kpc/h	46.5%	1.5 %	2.1 %	2.3 %	none
1024^3	39 kpc/h	59.6%	2.2 %	12.4 %	11.3 %	none
2048^3	20 kpc/h	67.4%	2.7 %	39.5 %	7.5 %	1D
Lag. Space		66.2%	2.8 %	10.7 %	7.5 %	none
EX+TC		102%	5.1 %	8.5 %	8.4 %	none

and the 1024^3 bins cases, the largest region in the 2048^3 bins case makes up 39.5% of the total volume whereas the second largest makes up only 7.5%. We list these values among other relevant mass- and volume-fractions in Table 3.3.

It is hard to conclude from this alone whether single-stream regions percolate or not. The largest component clearly stands out in the 2048^3 bin case. However, if it was percolating, it would be surprising that it still only makes up a bit more than half of the total single-stream volume of 67.4%.

Through a more sophisticated percolation test we find that the largest region in the 2048^3 bin case is percolating in the y-dimension, but not in the x- and z-dimension. That means the region together with its periodic replications would form an infinitely large string-like structure along the y-dimension, but would not be connected to its periodic images in the $x-$ and $z-$ directions. Percolation in a single dimension is practically impossible in an infinite universe, and the fact that we find it in this periodic box, shows that we are limited by finite size effects.

Resolution effects

It is somewhat surprising that the connectivity of the single-stream regions depends so strongly on the resolution of the mesh. The reason for this could be either (1) that there are tiny holes in the multi-stream regions that only get resolved at higher mesh resolutions or (2) that the likelihood of linking regions together through numerical artefacts increases strongly with the mesh resolution.

We find some evidence that (2) might be the case here. While we use exclusively a resampling with 32768^3 particles in the plots in this section, we tested also how the resampling resolution affects the results. Generally we find that the single-stream field is very well converged with the used resampling resolution. For example when switching from 16384^3 to 32768^3 resampled particles, the volume fraction of single-stream regions in the 1024^3 bins case only decreases by 0.3% from 59.9% to 59.6%. While this proves that there is only a small uncertainty in the classification of the cells, such a small difference can lead to major differences in the linked regions. For example the same resolution switch changed the volume of the largest connected region from 22.7% to 12.4% - apparently

by disconnecting the largest connected component into two independent subcomponents. Similar disconnections might be expected for the 2048^3 bins case when further increasing the resampling resolution.

We speculate that the dependence of the connectivity on such tiny details is probably due to the large surface area of the single-stream - multi-stream intersection and the small thickness of the multi-stream regions. The large surface area leads to a large number of boundary pixels that need to be classified correctly, while the small thickness of multi-stream regions leads to a larger chance of single-stream regions connecting through misclassified pixels. When increasing the number of bins at constant resampling resolution, the number of surface area pixels increases dramatically, and the chance of misclassifying individual pixels increases slightly - thereby increasing the chance of non-physical connections significantly.

Percolation in Lagrangian space

These resolution issues can in principle be solved by using even higher resampling resolutions so that the chances of misclassifications diminish. However, a cheaper alternative is to map the single-stream field into Lagrangian space, and link resolution elements there. In the continuum limit, linking the single-stream regions in Lagrangian space or in Eulerian space should be equivalent, since the mapping from Lagrangian to Eulerian space is one-to-one and continuous in single-stream regions. However, at finite resolution the chances of linking together unconnected regions due to misclassified cells is much smaller in Lagrangian space. This is because the disconnecting multi-stream regions are much larger in Lagrangian space, and the number of pixels that define the intersections is much smaller. This becomes obvious when comparing Figure 3.16 with Figure 3.17.

We classify each particle as single-stream or multi-stream according to the class of the nearest grid point from the 2048^3 bin Eulerian single-stream field. Then we determine the connected components by linking each particle with its 6 nearest neighbours in Lagrangian space[1]. We show a slice through the classification in Lagrangian space in the right panel of Figure 3.16. Further we determine the volumes of the connected components by weighting with $1/\rho$ where ρ is the Eulerian density at the particle positions. We provide the volume histogram of the single-stream regions as the red line in Figure 3.18.

We find that the single-stream regions do not percolate in Lagrangian space. We checked that this result is not affected by using a different linking scheme, e.g. by linking each cell with its 18 or its 26 nearest neighbours (including diagonal links). The volume distribution of single-stream regions appears to be relatively similar to the one in the 1024^3 bin case. This could possibly mean that resolution effects are best under control in that case.

However, while we think that the test for percolation in Lagrangian space should be more stable against resolution effects than the Eulerian one, we are still affected by finite size effects here. The largest single-stream region still makes up 10.7% of the total volume, and the volume histograms look quite noisy due to the low number statistics. In principle

[1]We provide movies showing the connected components in 3D at wwwmpa.mpa-garching.mpg.de/paper/singlestream2017/percolation.html.

a simulation with larger boxsize could help out. However, at the same time it appears to be important that the resolution is sufficiently high to capture the power spectrum far enough ($\sigma \gg 5$) and to resolve small features (20-40kpc) in the single-stream field. This would make a better test simulation computationally difficult.

3.6.4 Do single-stream regions percolate?

Unlike Falck & Neyrinck (2015) we do not find any case where single-stream regions clearly percolate. This may reflect differences in methodology, or may be because we test for percolation at a much higher variance of the density field $\sigma \sim 6.4$ where the mass fraction of single-stream regions is significantly smaller. While we find a case which gets close to full three dimensional percolation (2048^3 bins in Eulerian space), we suspect that resolution effects are enhancing connectivity here. A more robust test in Lagrangian space shows no sign of percolation. However, the single-stream regions that we find still occupy a major fraction of the box volume, and thus are still arguably quite close to percolation. A larger simulation could be made to test for finite size effects, but would be computationally expensive. Nevertheless, we suspect that single-stream regions do not percolate in the continuum limit of cold dark matter. The more small scale power is resolved, the less mass remains in the single-stream regions, the smaller the volume of Lagrangian space they occupy, and the less likely percolation becomes.

3.7 Conclusions

Simulating the unsmoothed dark matter distribution in the deeply non-linear regime is not yet possible for WIMP-like dark matter models, which have a high variance and require resolving many orders of magnitude in spatial scale. We propose here an excursion set formalism to get an estimate for the volume weighted density distribution. This allows us to estimate the median density of the universe $4 \cdot 10^{-3}\rho_0$ and the fraction of mass in single-stream regions $3 \cdot 10^{-3}$ if dark matter is a 100 GeV WIMP. The only parameter in the excursion set model is the rms relative density perturbation σ which is directly related to the properties of the specific dark matter model considered. Thus, in principle, a measurement of e.g. the *median density of the universe* would provide information about the nature of the dark matter particle. Unfortunately this is unlikely ever to be possible, since all measurements of dark matter are indirect and involve smoothing on a relatively large scale.

The proposed excursion set formalism gives a reasonable qualitative picture and an first estimate for the quantitative properties of single-stream regions. However, there is still a lot of room for improvement. A significant limitation comes from our heuristic assumptions about the tidal field in the non-linear regime. A more sophisticated examination of the statistics of the tidal field within simulations could lead to a superior model. Further it would be straightforward to go to a full tensor-valued evolution model (given by the Geodesic Deviation Equation). This would require dealing with changes in the orientation

of the tidal field.

Further the excursion set formalism predicts that single-stream regions do not percolate if the resolved linear rms density perturbation satisfies $\sigma \gtrsim 5$. We suspect that previous studies (Falck & Neyrinck, 2015) did not reach high enough resolution to reach the state where single-stream regions no longer percolate. We attempted to test this prediction with a smaller, but higher resolution simulation with $\sigma \sim 6.5$. This simulation shows many small distinct single-stream regions, but their distribution depends strongly on the choice of resolution parameters, even if the classification of individual cells into single and multi-stream regions seems to be converged reasonably well. We suspect that percolation detection in Eulerian space is susceptible to numerical artefacts, since the surface area of the single-stream regions is enormous whereas the thickness of the disconnecting multi-stream regions is relatively small. A percolation test in Lagrangian space, which should be numerically more robust, shows no signs of percolation.

It would be interesting to run a larger simulation with the same mass resolution to get a better statistical grasp of the volume distribution of single-stream regions, and to test this on a scale, where finite size effects do not matter anymore. Our expectation is that single-stream regions will not percolate in the continuum limit of cold dark matter. At $\sigma \sim 6.4$ we already observe a state where single-stream regions are typically disconnected. Cold dark matter has a much larger variance of the density field (e.g. $\sigma = 23.8$ for a 100GeV WIMP) and therefore a significantly smaller mass fraction in single-stream regions. As a result, these occupy a substantially smaller volume in Lagrangian space and hence seem much less likely to percolate than in the (already non-percolating) simulation we have been able to carry out.

Finally we repeat the warning from our introduction, that our quantitative results apply only to idealised ΛCDM universes without baryons. They would be significantly modified by the inclusion of a realistic baryon fraction and realistic galaxy formation physics.

Acknowledgements

The authors thank Mark Vogelsberger for the access to the geodesic deviation equation code. JS thanks Oliver Hahn and Raul Angulo for helpful discussions and for providing a motivating and inspiring atmosphere within the collaboration.

3.8 Paper appendix - Derivation of the triaxial collapse model

In single-stream regions the local density is given by the stream density. We want to develop a model here that describes the evolution of infinitesimal volume elements in such single-stream regions so that the stream densities can be predicted. Fortunately we have already discussed a scheme which describes such infinitesimal volume elements – that is the geodesic deviation equation (as described in section 2.5). In single-stream regions we

can assume the trace of the tidal tensor (= local density) to be given by the stream-density and assume that the trace-free part of the tidal tensor is described by linear theory (or something close to it as we shall discuss later). In linear theory the distortion tensor can be completely described by its three eigenvalues (because its symmetric) and therefore this description can be interpreted as a triaxial model. Note that this model naturally recovers the spherical collapse model, but the interpretation is different: it is the evolution of a small volume element in the absence of an external tidal field[2].

3.8.1 The geodesic deviation equation

Recall from section 2.5 that the geodesic deviation equation describes the distortions of a small volume element by the distortion tensor. If we are only interested in the parts of the distortion tensor $\mathbf{D}_{\mathrm{Xq}}$ and $\mathbf{D}_{\mathrm{Vq}}$ that describe the deformations of the dark matter sheet, the evolution equations read

$$\dot{\mathbf{D}}_{\mathrm{Xq}} = a^{-2}\mathbf{D}_{\mathrm{Vq}} \tag{3.29}$$

$$\dot{\mathbf{D}}_{\mathrm{Vq}} = a^{-1}\mathbf{T}\mathbf{D}_{\mathrm{Xq}} \tag{3.30}$$

where the initial conditions are $\mathbf{D}_{\mathrm{Xq}}(a \to 0) = \mathbb{1}$ and $\mathbf{D}_{\mathrm{Vq}}(a \to 0) = 0$. The stream density of the Lagrangian volume element around a particle can be evaluated as

$$\rho_s = \frac{\rho_{\mathrm{bg}}}{|\det \mathbf{D}_{\mathrm{Xq}}|} \tag{3.31}$$

3.8.2 Single-stream regions

In single-stream regions the local density ρ is exactly given by the stream density ρ_s of the particles that occupy that point in space. The trace of the tidal tensor therefore has to match

$$\sum_i T_{ii} = \sum_i \frac{\partial^2 \phi}{\partial x_i \partial x_i} \tag{3.32}$$

$$= -\Delta\phi \tag{3.33}$$

$$= -4\pi G \rho_{\mathrm{bg}} \delta \tag{3.34}$$

$$= -\frac{4\pi G}{3}\rho_{\mathrm{bg}} \left(\frac{1}{|\det D_{xq}|} - 1 \right) \tag{3.35}$$

We therefore separate the trace from the trace-free part T_{ext} of the tidal tensor

$$\mathbf{T} = \frac{4\pi G \rho_{\mathrm{bg}}}{3} \left(\frac{1}{|\det \mathbf{D}_{\mathrm{Xq}}|} - 1 \right) \mathbb{1} + \mathbf{T}_{\mathrm{ext}} \tag{3.36}$$

[2]Arguably this explains why the spherical collapse model works so well in the first place. Most points in the density field show nothing close to spherical symmetry. However, it seems reasonable that a good first order approximation to the evolution can be obtained by setting the external tidal field to zero.

Note that for a given point in space $\mathbf{T}_{\rm ext}$ is completely non-local and depends on the density distribution of the surroundings.

3.8.3 Non rotating tidal field

The tidal tensor is symmetric by definition and therefore generally has six components. Since it is symmetric, there is always a system where it is diagonal. In the non linear cosmological density field this system can change its orientation over time leading to all components of the distortion tensor becoming active. However, in linear theory it stays diagonal in the same system at all times. In this case of a tidal field that does not change its orientation we can simplify the Geodesic Deviation Equation, since in the system where the tidal tensor is diagonal also $\mathbf{D}_{\rm Xq}$ and $\mathbf{D}_{\rm Vq}$ are diagonal from the beginning, and the non diagonal components never become active. We can then simplify the GDE to six differential equations

$$\dot{x}_i = a^{-2} p_i \tag{3.37}$$

$$\dot{p}_i = a^{-1} x_i \left(-\frac{4\pi G}{3} \rho_{\rm bg} \delta + T_{ext,i} \right) \tag{3.38}$$

for $i = 1, 2, 3$ where we label by x_i the diagonal components of D_{xq} and by p_i the diagonal components of D_{vq}, and use the abbreviation

$$\delta = \frac{1}{|\det \mathbf{D}_{\rm Xq}|} - 1 = \frac{1}{x_1 x_2 x_3} - 1 \tag{3.39}$$

We call models of this form *triaxial collapse models*. They can differ in their assumptions about the external tidal field $T_{\rm ext}$. Note that in the absence of an external tidal field $T_{{\rm ext},i} = 0$ this naturally recovers the spherical collapse model.

3.8.4 Linear theory

In linear theory the tidal tensor takes the form

$$T_{ij} = -4\pi G \rho_{\rm bg} d_{ij} D(t) \tag{3.40}$$

where d_{ij} is the deformation tensor and D is the linear growth factor. d_{ij} is a symmetric tensor and we can transform into the system where it is diagonal with the eigenvalues $d_{ii} = \lambda_i$. The external tidal tensor is then given by

$$T_{{\rm ext},ii} = -4\pi G \rho_{\rm bg} D(t) (\lambda_i - \delta_0/3) \tag{3.41}$$

$$T_{{\rm ext},ij} = 0 \text{ for } i \neq j \tag{3.42}$$

where $\delta_0 = \lambda_1 + \lambda_2 + \lambda_3$. We then arrive at

$$\dot{x}_i = a^{-2} p_i \tag{3.43}$$

$$\dot{p}_i = -\frac{4\pi G}{3} \rho_{\rm bg} a^{-1} x_i \left(\delta + D(t)(3\lambda - \delta_0) \right) \tag{3.44}$$

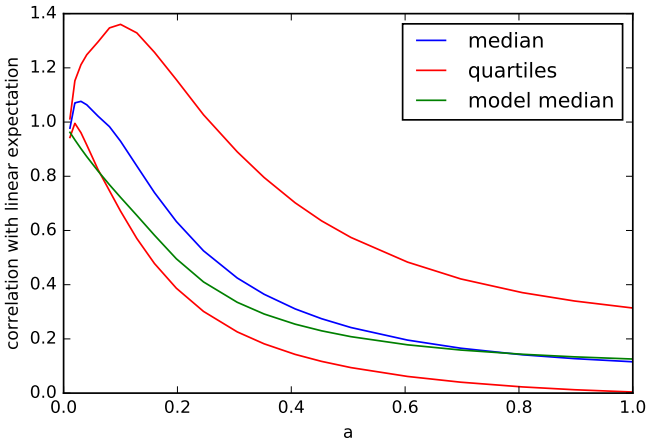

Figure 3.19: The degree of correspondence of the measured tidal field to the tidal field predicted from linear theory.

where only the external tidal field has been assumed to be given by linear theory but not the local density. Note that it is easy to recover the Zel'dovich approximation from this equations by approximating $\delta \approx D(t)\delta_0$, $x_i \sim 1$, leading to

$$\dot{x}_i = a^{-2} p_i \tag{3.45}$$

$$\dot{p}_i = -4\pi G \rho_{\rm bg} a^{-1} D(t) \lambda_i \tag{3.46}$$

which is solved by $x(t) = 1 + D(t)\lambda$.

3.8.5 Fading tidal field

While the triaxial collapse model using the external tidal field from linear theory gives a reasonable description for the evolution of volume elements in the linear regime and slightly non linear regime, it certainly cannot be extrapolated to strongly non-linear stages $D(t)\delta_0 \gg 1$. Generally the tidal field from linear theory becomes too large and completely dominates the evolution of volume elements. As can be seen in Figure 3.19 where we show as a similarity measure between the simulated tidal tensor and the one from linear theory the volume weighted median of the distribution of

$$\beta = \frac{\sum_{i,j} T_{\rm ext,sim,ij} T_{\rm ext,lin,ij}}{\sum_{i,j} T^2_{\rm ext,lin,ij}} \tag{3.47}$$

In realistic cases however, the non linear tidal field in single-stream regions is becoming smaller and also independent from its initial alignment - therefore becoming sub-dominant in comparison to the local self gravity. Therefore for the excursion set formalism we consider a triaxial collapse model where we multiply the linear external tidal field by a damping

factor to explicitly fade out the tidal forces in the non-linear regime. This leads to an acceleration equation of the form

$$\dot{p}_i = -\frac{4\pi G}{3}\rho_{\rm bg} a^{-1} x_i \left(\delta + \alpha(t)(3\lambda - \delta_0)\right) \tag{3.48}$$

where α has to be $D(t)$ in the linear regime and should be limited in the non linear regime. A simple parametrization which seemed to work well for us is

$$\alpha(t) = \frac{D(t)}{1 + |\delta_0| D(t)}. \tag{3.49}$$

For simplicity we stick to this model for now, but other choices are possible. In principle a more sophisticated model could be found by measuring the statistics of the external tidal field in the non linear regime and quantifying them adequately.

Chapter 4

Simulating the complexity of the dark matter sheet

This chapter is about the question *how to make warm dark matter simulations that are reliable everywhere?* I have worked on this question together with Oliver Hahn, Raul Angulo and Simon White for quite a while and we developed several extensions to the numerical scheme from Hahn & Angulo (2016). In this collaboration we have made a paper draft that we will submit soon under the (probable) title *Simulating the Complexity of the Dark Matter Sheet I: Numerical Algorithms*. I have written most of the text and created the figures, but some parts of the text have been edited or reworked by the co-authors. The concepts and ideas in this paper draft are the joint work of all authors. This chapter is mostly given by this paper draft. However, the draft also contains a few of the concepts which I have explained in chapter 2 of this thesis – most significantly the morphology classification scheme that is described in 2.6.4. I have replaced the description of these concepts here by short summaries to reduce redundancy. The details of the upcoming publication are listed in Table 4.1.

Very recently we have noticed some problems with the consistency of the estimation of stream densities with the geodesic deviation equation when the novel tree techniques (that are discussed in section 4.4.1) are used. This is the main reason why we have not published this paper yet, but we will try to fix this soon after I hand in this thesis. I will shortly describe the problem in section 4.4.5.

title:	Simulating the Complexity of the Dark Matter Sheet I: Numerical Algorithms
authors:	Jens Stücker, Oliver Hahn, Raul E. Angulo and Simon D.M. White
estimated submission date:	May or June 2019
journal:	Monthly Notices of the Royal Astronomical Society

Table 4.1: Details of the upcoming publication that this chapter is based on

4.1 Abstract

At early times dark matter has a thermal velocity dispersion of unknown amplitude which, for warm dark matter models, can influence the formation of nonlinear structure on observable scales. We propose a new scheme to simulate this process that combines two previous methods in a way that avoids the numerical artefacts which have so far prevented either from producing fully reliable results. At low densities and throughout most of the cosmological volume, we represent the dark matter phase-sheet directly using high-accuracy interpolation, thereby avoiding the artificial fragmentation which afflicts particle-based methods in this regime. Such phase-sheet methods are, however, unable to follow the rapidly increasing complexity of the denser regions of dark matter haloes, so for these we switch to an N-body scheme which uses the geodesic deviation equation to track phase-sheet properties local to each particle. In addition, we present a novel high-resolution force calculation scheme based on an oct-tree of cubic force resolution elements which is well suited to approximate the force-field of our combined sheet+particle distribution. Our hybrid simulation scheme enables the first reliable simulations of the internal structure of low-mass haloes in a warm dark matter cosmology.

4.2 Introduction

While the ΛCDM model explains most key observations of our universe remarkably well, the physical nature of its main ingredients – dark energy and dark matter – remains unknown. So far dark matter appears to be cold and collisionless on those scales which have been reliably probed. However, in all dark matter models which are well motivated from particle physics, dark matter has a residual thermal (or non-thermal) velocity dispersion at early times. In the case of weakly interacting massive particles (WIMPs) the thermal velocity dispersion is so small that it only affects structure formation on very small scales ($\sim$ Earth mass). However, for other dark matter candidates - for example, sterile neutrinos - the thermal velocity dispersion might be large enough to cause effects on observable (or soon observable) scales. In this case we speak of warm dark matter. If we were to measure such effects, we could constrain the nature of dark matter. It is therefore of fundamental interest to understand the implications of the "warmth" of dark matter, and to search for deviations from the perfectly cold scenario. It is the aim of this paper to improve upon existing methods for simulating warm dark matter (WDM), so that reliable predictions for warm dark matter cosmologies become possible.

Cosmological N-body simulations have been very successful at predicting the large-scale structure of the universe for cold dark matter scenarios (see e.g. Frenk & White, 2012, for a review). In N-body simulations, the (statistically) well known linear density field of the early universe is discretized to a finite number of macroscopic particles. These particles are then evolved under their self-gravity to infer an approximation to the late-time non-linear density field. The particles are interpreted as a Poisson-sampled realization of the continuous non-linear density field of the dark matter fluid.

However, in the case of warm dark matter cosmologies, N-body simulations give rise to numerical artefacts during the first phases of nonlinear structure formation. In practice the difference between cold and warm dark matter simulations lies merely in the choice of initial conditions. The initial density field is smoothed by the free streaming motion of the particles and therefore its power spectrum has a small scale cut-off, which is on larger scales for warmer dark matter. Since at later times the thermal velocity dispersion is relatively small when compared to the bulk velocities of the dark matter fluid, the thermal velocity dispersion can be neglected once the cut-off in power on small scales is established (Bode et al., 2001). Therefore in all cases one simulates a perfectly cold fluid, however either with (WDM) or without (CDM) an additional truncation scale in the initial perturbation spectrum. N-body simulations of warm dark matter form a large number of small haloes - most prominently found regularly spaced in filaments - aligning like beads on a string (Bode et al., 2001). Wang & White (2007) showed that these haloes are not of physical nature, but are merely numerical artefacts. They found such fragments even in the case of N-body simulations of the collapse of a perfectly homogeneous filament.

The fragmentation is a natural consequence of the anisotropic collapse with incomplete thermalisation in cosmology. This anisotropy of collapse means that, as structure forms, it collapses first to a one-dimensional sheet, or "pancake", followed by collapse to a filamentary strucuture, and only then a halo (e.g. Bond et al., 1996). In each case, the structures are supported by velocity dispersion only along the already collapsed directions, while the temperature is still effectively zero in the uncollapsed dimensions (cf. Buehlmann & Hahn, 2018), making them unstable to spurious collapse seeded by numerical noise. The underlying reason is of course that in a collisionless fluid no thermalisation (and therefore isotropisation of the temperature) takes place.

In recent years, a new set of simulation schemes has been designed which are unaffected by this artificial fragmentation (Hahn et al., 2013; Hahn & Angulo, 2016; Sousbie & Colombi, 2016). These employ a density estimate that is much closer to the continuum limit than that obtained from the particles in standard N-body simulations. This density estimate is obtained by interpolating between the positions of tracer particles in phase space. This is possible since in the limit of a cold distribution function, these tracers occupy only a three-dimensional (Lagrangian) submanifold of phase space, also known as the dark matter sheet (Arnold et al., 1982; Shandarin & Zeldovich, 1989; Shandarin et al., 2012; Abel et al., 2012). While this approach has successfully been used in Angulo et al. (2013) to measure the WDM halo mass function below the cut-off scale, there are still major limitations to the range of applications of the schemes. Inside haloes the dark matter sheet grows rapidly in complexity making it hard to reconstruct the sheet accurately (Vogelsberger & White, 2011; Sousbie & Colombi, 2016). Therefore schemes which do not refine the resolution of the interpolated mass elements give biased densities inside haloes, and schemes which use refinement (Hahn & Angulo, 2016; Sousbie & Colombi, 2016) quickly become unfeasibly complex. This is so since the detailed fine-grained evolution of the distribution function has to be followed at any time, so that new tracers need to be inserted in order to not lose information about the dynamics.

This is an important difference between the two approaches. In the N-body method,

one benefits from ergodicity, i.e. in a time-averaged sense one obtains an accurate representation of the underlying distribution function, even if at any moment in time, the particular realisation might not be perfect. This is also the underlying reason, why the N-body method has problems with anisotropic collapse from cold initial conditions: ergodicity has not been established in the uncollapsed subspace, where the mean field dynamics is now very noisy. This is circumvented by following the distribution function explicitly, which can be done as long as its structure is not yet too complex. In this case, there is no noise and the cold uncollapsed subspaces can be followed accurately. Ultimately however, rapid phase and chaotic mixing inside of haloes lead to close to ergodic dynamics, rendering it increasingly complex and ultimately impossible to follow the evolution of the sheet, but making N-body attractive, since it relies on exactly that assumption.

Finally, another short-coming of the previous implementations of the sheet method is that they have so far only worked at very low force-resolution. So far these have used only a single mesh for the force calculation which smooths the density field on scales much larger than what is necessary to resolve the centers of haloes. An accurate treatment requires an adaptive scheme for the force-calculation and the time-stepping.

In this paper, we propose solutions to these short-comings of sheet-based dark matter simulations. We employ a hybrid scheme which uses sheet-based simulation techniques wherever the interpolation is reliable, and switches to N-body based simulation techniques where the sheet becomes untraceable, but where we are reasonably confident that in a time-averaged sense the particles reproduce the correct mean field dynamics. We illustrate this in Figure 4.1 by a projection of the density field in and around a halo. We present - going from top left to bottom right - a low resolution $N = 128^3$ N-body simulation, a high resolution $N = 512^3$ N-body simulation as reference, a low resolution $N = 128^3$ sheet simulation (without using refinement techniques) and a low resolution $N = 128^3$ "sheet + release" simulation (without refinement) that switches to an N-body scheme when the sheet becomes too complex. The N-body case produces the correct halo shape, but fragments in low density regions. The pure sheet case captures the low density regions with stunning accuracy, but produces a deformed halo. However, the sheet + release case inherits the best of both worlds and seems authentic everywhere - thereby coming closest to the much higher resolution reference simulation at a much reduced number of degrees of freedom in the simulation, since instead of 512^3 only 128^3 freely moving tracers are used. It is the subject of section 4.3 to guide the reader to an understanding of how and how well the sheet + release scheme works.

In section 4.4 we further develop a new tree-based discretization of the force-field which is compatible with both N-body and sheet approaches. This makes it possible to use the sheet + release simulation approach all the way down to the small force resolutions scales that are needed to resolve the centers of haloes.

Thus, in this chapter, we present for the first time a full scheme which makes possible non-fragmenting and unbiased warm dark matter simulations with high force-resolution. In the subsequent chapter we will present its predictions for the case of one of the smallest haloes in a warm dark matter universe.

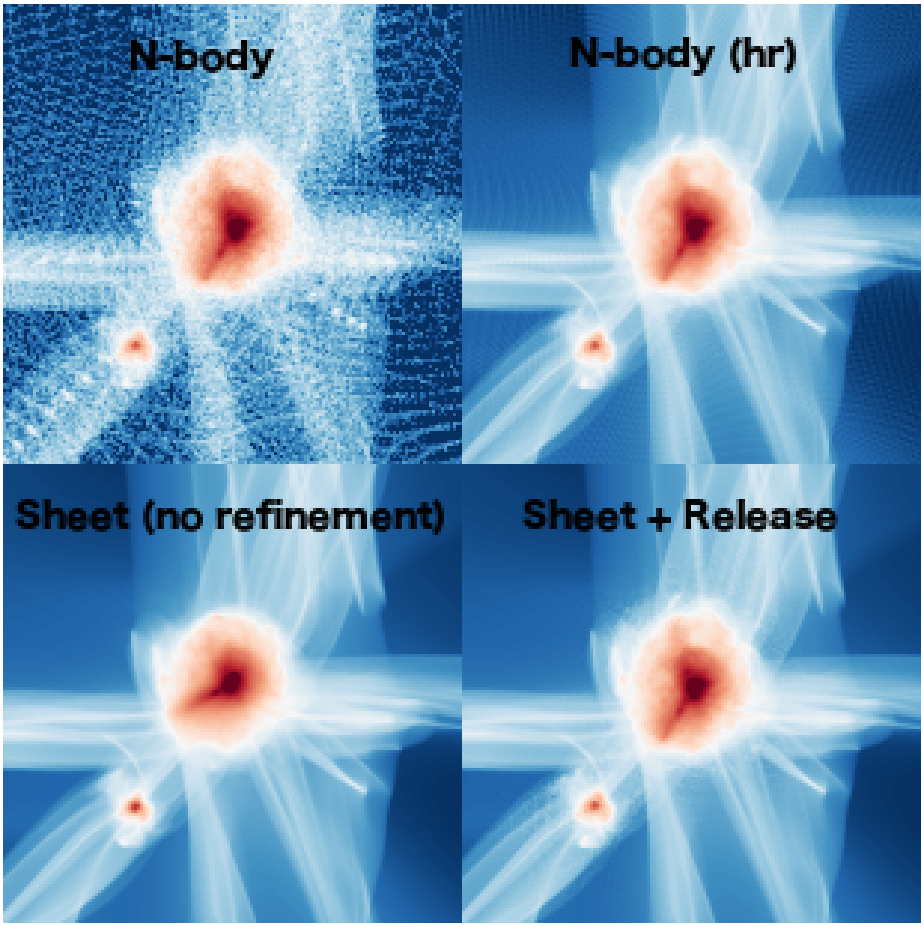

Figure 4.1: A projection of the density field in and around a halo for different simulation setups. Top left: N-body with $N = 128^3$ particles. Top right: N-body with $N = 512^3$ as a high-resolution reference case. Bottom left: sheet (no refinement) with $N = 128^3$. Bottom right: sheet + release (no refinement) with $N = 128^3$. "Sheet + release" means that most of the mass is traced by the sheet interpolation, but mass elements which were detected to become too complex are traced by an N-body approach instead. The low resolution N-body scheme appears to get the shape of the center of the halo correct, but fragments in the low density regions. The pure sheet scheme captures the low density regions very well, but creates a biased overly-round halo, since its phase space structure is too complex for reconstruction by interpolation. The sheet + release case inherits the best of both worlds and avoids the problems with fragmentation or biased halo structure. It comes closest to the much higher resolution reference case in the top right panel. Note that the $N = 512^3$ N-body case would also fragment if the force-resolution were increased significantly, while the sheet cases can also avoid fragmentation in that scenario. It is the subject of section 4.3 to elaborate the details of the sheet + release scheme.

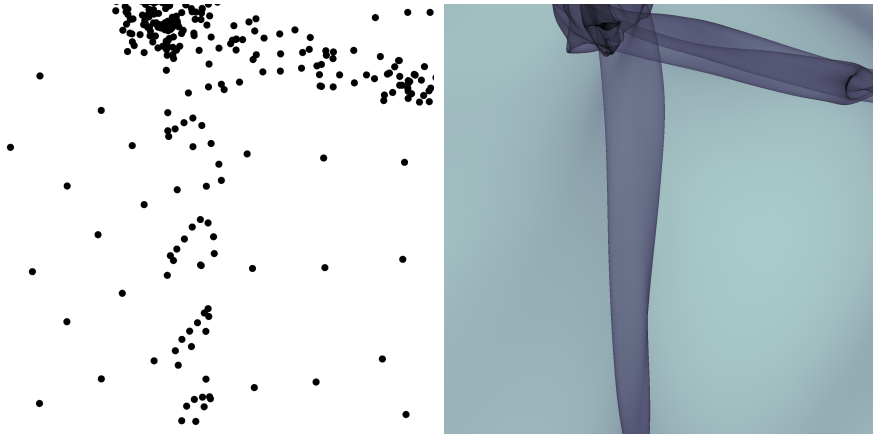

Figure 4.2: N-body density estimate versus a density estimate inferred by interpolation of the dark matter sheet in phase space. The N-body density estimate shows regular lumps which will grow into fragmented artificial haloes. However, the continuum density estimate shows no such artefacts.

4.3 A fragmentation-free and unbiased scheme for cosmological warm dark matter simulations

4.3.1 The dark matter sheet

The artificial fragmentation of N-body simulations (as for example in Figure 4.1 top left) appears to be a major shortcoming of the N-body method. The difference between simulations of warm dark matter (which tend to fragment) and cold dark matter (where artificial fragments are subdominant to real small-scale clumps) lies merely in the choice of the initial conditions. In principle simulations of warm dark matter should be simpler than cold dark matter ones, since there is much less structure and complexity. Solving the problem of fragmentation of N-body simulations is thus not only important for testing warm dark matter models, but also as a test, beyond the customary numerical convergence tests, of the validity of the N-body scheme as a whole in the cold dark matter case.

As already discussed in section 2.4, this fragmentation can be overcome by schemes which trace the dark matter sheet. These use a smoother density estimate in their Poisson solver which is inferred from the interpolated dark matter sheet. As a short summary, please have a look at Figure 4.2 (which we have already shown earlier in 2.11) and the top panel of Figure 4.3. We will briefly discuss here what are the limitations of sheet-schemes.

As can already be suspected from the top panel of Figure 4.3, the complexity of the dark matter sheet grows over time. While the sheet is very well represented by a linear interpolation between the particles in the initial conditions (top left panel), the representation by interpolating between the particles at later times is already considerably worse. The function that is to be captured by the interpolation has grown in complexity. A better

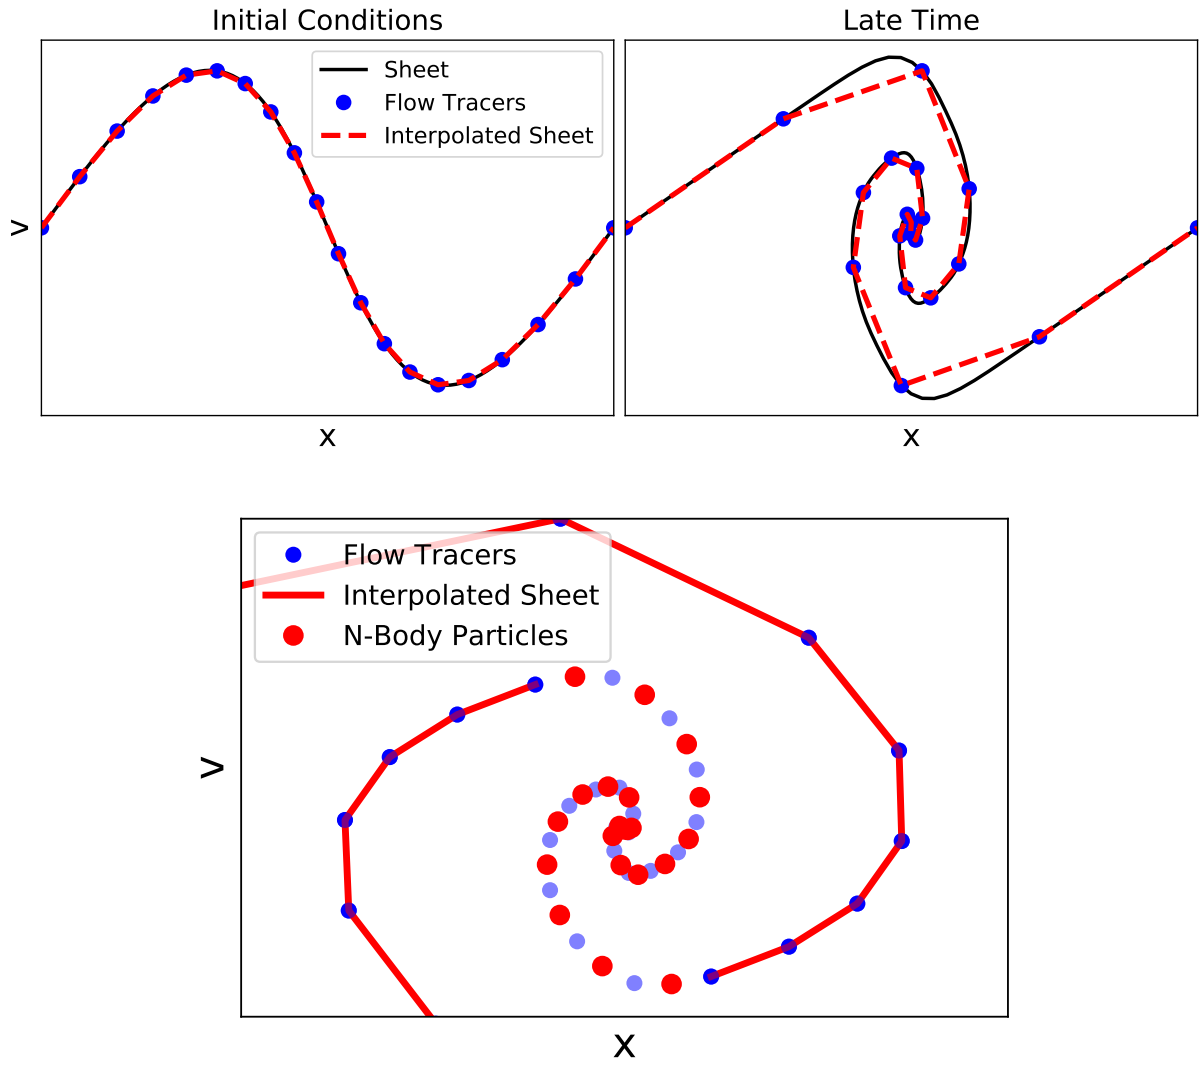

Figure 4.3: Top: Illustration of the dark matter sheet in phase space. Dark matter occupies a submanifold (a "sheet") in phase space. This submanifold can be reconstructed by interpolation from a finite number of tracers. Note: in this image we purposely use only linear interpolation to emphasize the difference between the true dark matter sheet and the one reconstructed from interpolation. When a finite number of tracers is used, the interpolation gets worse with time, because the complexity of the sheet grows with time. Bottom: Illustration of the release. Originally the mass (in red) in all Lagrangian volume elements was traced by the interpolated sheet. However, in the course of the simulation some mass elements have been flagged for release and are now represented by N-body particles instead.

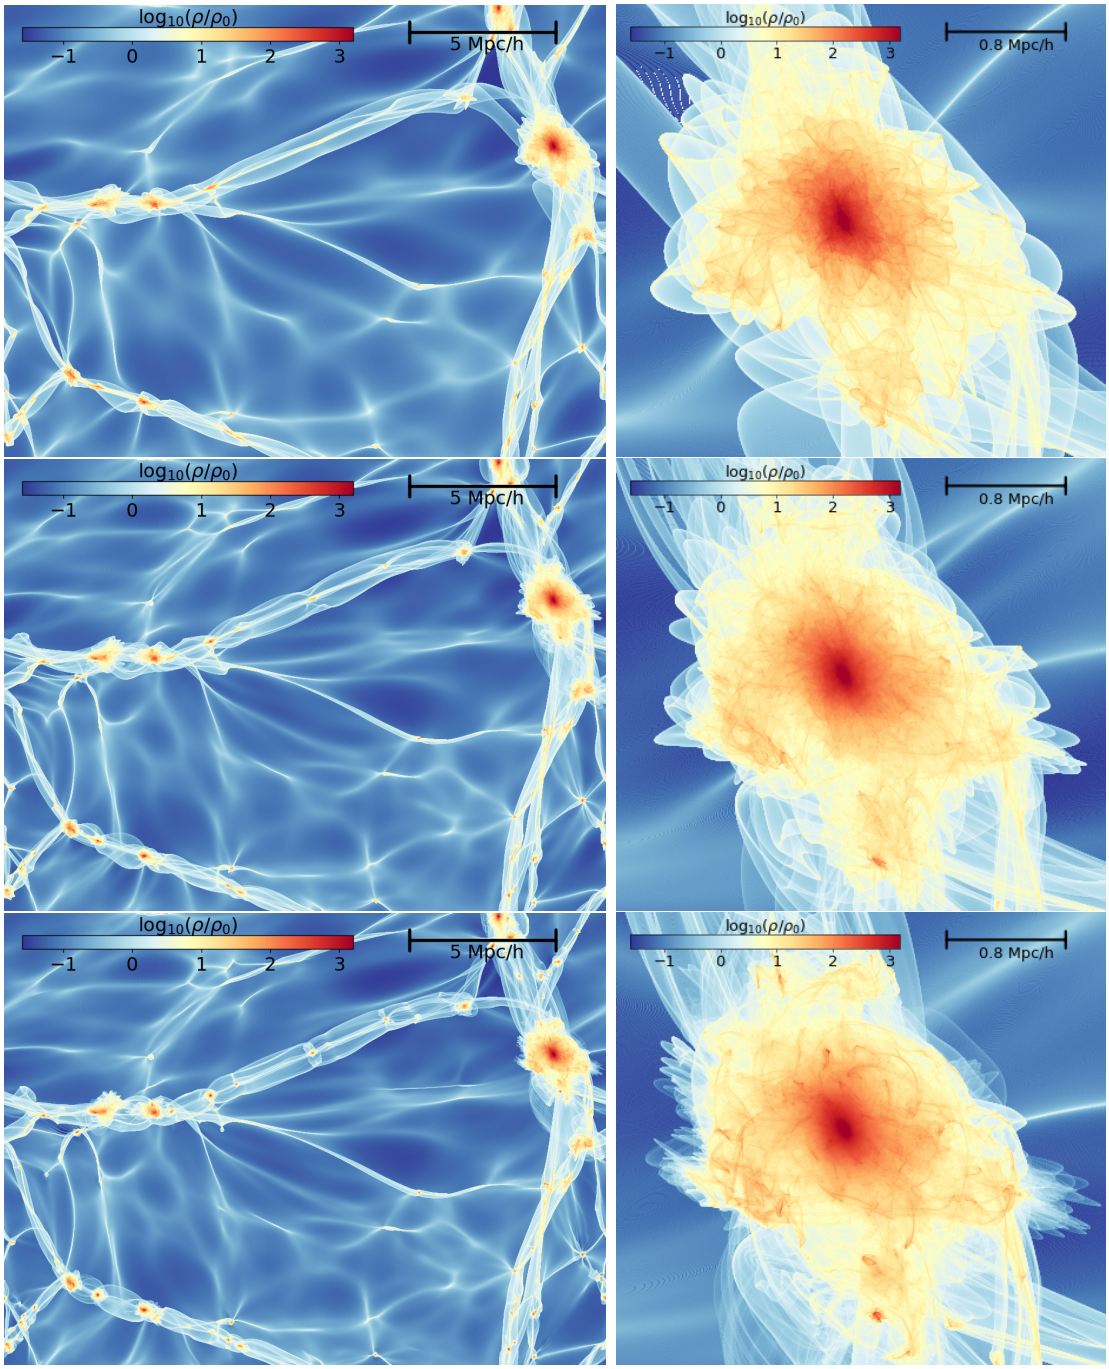

Figure 4.4: The projection of a pseudo-2d sheet into Eulerian space (see text for details) as reconstructed with varying number of tracer particles. Top with 64^2 particles, center with 128^2 particles and bottom with 512^2 particles. Left: showing a large region with various different structure types. Inside low-density regions (which are the majority of the volume) the interpolation seems to be well converged already at moderate resolution levels – thereby getting close to the continuum limit. Right: zoom onto a halo. The reconstructed density field in and around haloes still seems to be resolution dependent even when most other regions are already converged at those resolutions. (This figure is not part of the actual paper draft.)

representation of the dark matter sheet can then be obtained by using a higher interpolation order, or by an adaptive higher number of particles as has been explored by Hahn & Angulo (2016) and Sousbie & Colombi (2016).

Let us have a short look at how the quality of the inferred density field changes with the number of available tracers. For this we select from an $N = 512^3$ simulation of 1 keV warm dark matter (without refinement) a pseudo-2D sheet by selecting a slice with a fixed Lagrangian z coordinate $q_z = \text{const}$. In Figure 4.4 we show how its projection to Eulerian space changes with the resolution. We change the resolution that is used for reconstruction to 512^2, 128^2 and 64^2 particles by selecting every, every 4th and every 8th particle along each axis as support points of the interpolation. Inspection of Figure 4.4 already gives qualitatively the insight that we will test more quantitatively in the other sections of this paper: The density field is reconstructed very well in low-density regions like single-stream regions, pancakes and filaments. At finite resolution the sheet density estimate gets already extremely close to the continuum limit. However, in and around halos the density field still seems resolution dependent at relatively high resolution, because the dark matter sheet is quite complex here.

As already explained in section 2.4.4, the schemes of Hahn & Angulo (2016) and Sousbie & Colombi (2016) use adaptive refinement approaches to increase the Lagrangian resolution where the sheet has a higher complexity and keep the resolution low where convergence can be achieved with a low number of tracers. By a refinement criterion, these schemes try to make sure that the interpolated sheet is consistent with the continuum limit at all times (up to some error tolerance).

With these schemes it is, in principle, possible to make a simulation that exactly traces the dark matter sheet. However, this turns out to be impossible in practice, since the dark matter sheet grows in complexity very rapidly. Tracing it requires an extraordinary amount of computational resources. Sousbie & Colombi (2016) managed to carry out a $m_X = 250\,\text{eV}$ WDM simulation in a 28 Mpc box until $a = 0.31$ and found the number of simplexes required to scale with the twelfth power of time. Assuming that this scaling remains valid until $a = 1$, running that simulation until the present time would require roughly 10^6 times as much memory and probably also computational time. However, it is already an optimistic assumption that this scaling can be extrapolated so easily. So soon after their formation, their haloes had probably had no mergers yet and so had maintained a relatively simple phase space structure. It cannot be excluded that chaotic orbits with an exponentially growing complexity arise from merging haloes. Further the simulation described in Sousbie & Colombi (2016) has a relatively low force-resolution of $28\,\text{Mpc}/1024 \approx 27\,\text{kpc}$. Therefore the central structure of haloes is resolved poorly. The complexity of the sheet is expected to grow most rapidly in the centers of haloes. Vogelsberger & White (2011) find that the number of streams at a single point in the center of the halo of a Milky-Way-type dark matter halo in a cold dark matter universe might already get as high as 10^{16}. If that is true, a dark matter sheet plus refinement based simulation scheme for such a halo would require far more than 10^{16} resolution elements.

Even in that most optimistic case it seems unlikely that a simulation like that in Sousbie & Colombi (2016) can be run until the present day, $a = 1$. We will demonstrate in this

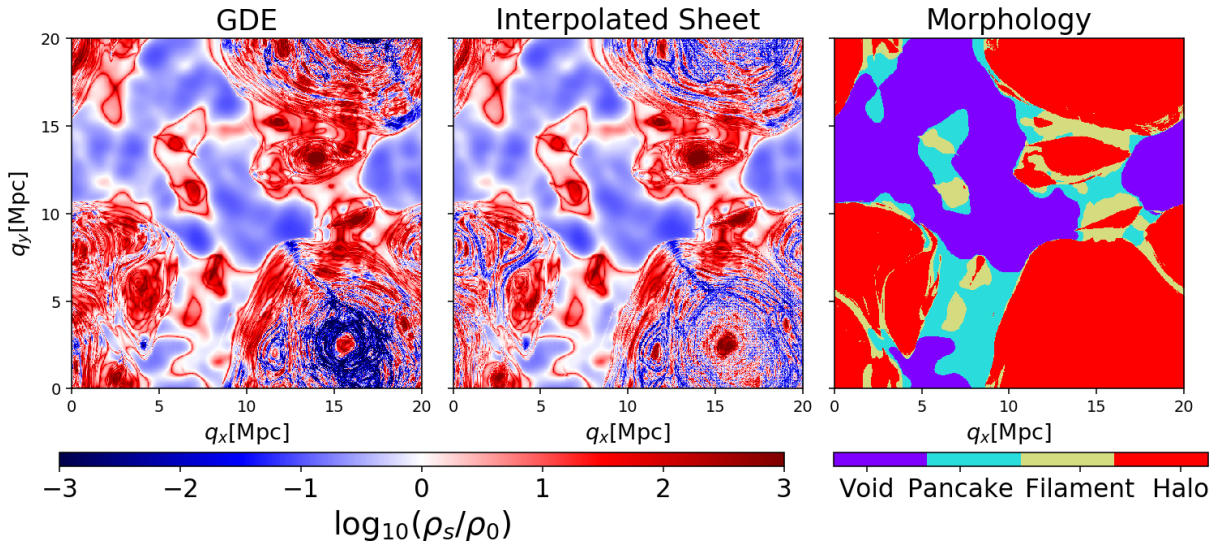

Figure 4.5: Stream-Densities on an infinitesimally thin plane in Lagrangian Space. Left: Exact stream densities from the GDE distortion tensor. Center: Finite differences approximation, representing the derivatives of the interpolated sheet. Right: Morphology classification as described in section 2.6.4. The stream densities agree extraordinarily well in regions where they vary slowly with the Lagrangian coordinates, but get into complete disagreement in regions where they vary rapidly - that is in haloes. The sheet is too complex here for accurate reconstruction.

chapter how to deal with this with affordable computational costs. We propose a simulation scheme with a "release" mechanism that uses a sheet-interpolation scheme where it is well converged, and switches to a particle based N-body approach in regions where the sheet becomes too complex. This allows us to perform the first warm dark matter simulations that do not fragment in low density regions while remaining accurate in the inner regions of haloes. We illustrate qualitatively in the bottom panel of Figure 4.3 how the release could look in the phase space of a one dimensional world. Note that in the case of a three dimensional simulation the complexity in the released region would be much higher - for example it could have $\sim 10^{16}$ foldings in the same region (Vogelsberger & White, 2011).

The next sections will explain how we identify regions where the interpolation scheme breaks down.

4.3.2 Quantifying complexity - the geodesic deviation equation

Recall that the geodesic deviation equation (GDE) makes it possible to follow the distortion tensor for each particle in a cosmological simulation (as explained in section 2.5). We have implemented the GDE into the HA16 scheme to follow the tangent space of the sheet

locally at the Lagrangian location of each particle

$$\mathbf{D}_{\mathrm{xq}} = \frac{\partial \mathbf{x}}{\partial \mathbf{q}} \qquad\qquad \mathbf{D}_{\mathrm{vq}} = \frac{\partial \mathbf{v}}{\partial \mathbf{q}} \tag{4.1}$$

The distortion tensor can be initialized by using finite differencing methods on the initial conditions, so that at 2nd order

$$\frac{\partial x_i}{\partial q_j} \approx \frac{x_i(\mathbf{q} + \Delta q \mathbf{e}_j) - x_i(\mathbf{q} - \Delta q \mathbf{e}_j)}{2\Delta q} \tag{4.2}$$

$$\frac{\partial v_i}{\partial q_j} \approx \frac{v_i(\mathbf{q} + \Delta q \mathbf{e}_j) - v_i(\mathbf{q} - \Delta q \mathbf{e}_j)}{2\Delta q} \tag{4.3}$$

where $\mathbf{e}_j$ is the unit vector along the j-th coordinate axis. If the simulation particles are initially located on a grid, the evaluation points of the finite differencing can be chosen so that the $\mathbf{D}_{\mathrm{xq}}$ and $\mathbf{D}_{\mathrm{vq}}$ distortion tensors can be approximated purely by taking differences of particle positions and velocities. We note that these tensors can also be explicitly calculated when generating cosmological initial conditions. We found however that finite differences are accurate enough, and can be more convient, since the initial condition files do not need to be modified.

In the initial conditions of a typical cosmological simulation this will always be a reasonably good approximation, since initial conditions are typically set at a time where the displacement field varies only moderately between Lagrangian neighbors. However, the finite-differencing scheme can also be used to obtain an approximation for the distortion tensor at any later time. If the Lagrangian map $\mathbf{x}(\mathbf{q})$ varies slowly with the Lagrangian coordinate this will be a good approximation, but if it varies rapidly, the approximation will break down. These are the cases where the dark matter sheet becomes too complex to be reconstructed from particle positions.

To illustrate this we show in Figure 4.5 a comparison of the stream densities

$$\rho_s = \frac{\rho_0}{|\det(\mathbf{D}_{\mathrm{xq}})|} \tag{4.4}$$

that can be obtained from the finite difference distortion tensor as in (4.2) and the infinitesimal distortion tensor that has been evaluated by the GDE. Additionally we show the result of the morphology classification based on the distortion tensor as already described in section 2.6.4. The Figure shows a slice through Lagrangian space, where each particle is plotted at its initial comoving location $\mathbf{q}$ (for $a \to 0$), but colored with the stream density it has at a later point in the simulation. In Lagrangian space the volume is proportional to the mass, therefore for example haloes appear as large regions in Lagrangian space whereas they make up a rather small part of the volume in Eulerian space. It can be seen that the GDE and the finite difference distortion tensor are in excellent agreement wherever the stream-density varies slowly - that is in single-stream regions, pancakes and filaments, as we shall see later. However, there are also regions (i.e. haloes) where the stream density

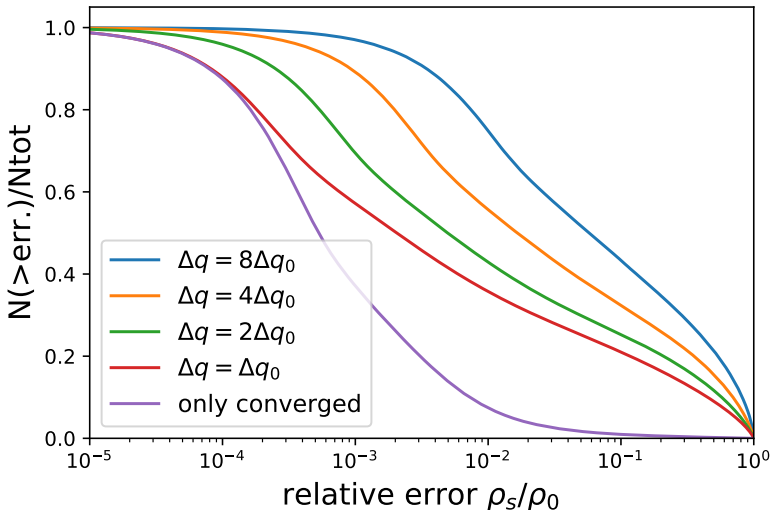

Figure 4.6: Relative error between the stream densities of the GDE distortion tensor and finite-difference approximations to the distortion tensor of the interpolated sheet with different spacings of the sampling points ($\Delta q_0 = 20\,\mathrm{Mpc}/512$). The finite difference distortion tensors converge quickly to the GDE distortion tensor for most of the particles. For roughly $30 - 40\%$ of the particles the convergence is rather slow. Particles which have a converged distortion tensor (at the 10% level) between different resolution levels (purple line) are in excellent agreement with the GDE distortion tensor, showing that the GDE distortion tensor is indeed the limit of the sheet derivative for infinite particle resolution $\Delta q \to 0$.

varies rapidly. The finite-differencing scheme breaks down here and becomes resolution-dependent, whereas the local distortion tensor of the GDE still remains valid. For infinite resolution the finite-differencing scheme should converge to the GDE result.

We demonstrate this more quantitatively in Figure 4.6 where we plot the relative difference between the GDE stream densities $\rho_{s,gde}$ and the stream densities inferred from finite differences $\rho_{s,fd}$

$$\epsilon = \left| \frac{\rho_{s,gde} - \rho_{s,fd}}{\rho_{s,gde} + \rho_{s,fd}} \right| \tag{4.5}$$

To test the convergence, we compare the finite differencing for different resolution levels, using all particles, every second particle (per dimension), every 4th and every 8th. It can be clearly seen that the finite difference stream densities converge to the GDE stream density. The GDE provides the limit for $\Delta q \to 0$ of the sheet distortion tensor – that is the infinitesimal derivative of the continuum dark matter sheet which is not affected by interpolation uncertainties. On the other hand the finite differencing provides the derivatives of the interpolated sheet, which are wrong where the interpolation is not converged.

To further emphasize this, we additionally select a subset of particles where the finite differencing is converged, which we define by their stream densities not changing by more than 10 % when only selecting every second or every fourth particle. For particles where the finite differencing has converged, the agreement with the GDE is remarkable. We conclude that the comparison of GDE and finite difference distortion tensor can be reliably used as a benchmark for the accuracy of sheet interpolation schemes.

Looking a bit closer at Figure 4.6, it seems puzzling that while for roughly 60% of the particles the stream densities converge quickly, for the other 40 % of the particles the stream densities converge rather slowly. We shall see that those regions where convergence is slow are mostly haloes, and that achieving true convergence here is almost hopeless. Any sheet interpolation scheme will either break down (without using refinement) or become too expensive to be followed (when using refinement) at some point.

To illustrate how important it is that the interpolation is converged, we show in Figure 4.7 power-spectra that have been inferred on sheet-based dark matter simulations (without any refinement) with different resolutions at two different points in time. While it seems that at $a = 0.3$ the power spectra are converged, showing that the sheet-interpolation is working reasonably well at that time, the situation is very different at $a = 1.0$: The complexity of the true dark matter sheet is too high to be captured by the interpolation scheme with a limited number of particles. Therefore the interpolation and subsequently the power spectrum is far from converged with the particle number. Typically the densities in the centers of haloes get strongly overestimated by a poorly interpolated reconstruction, as has been already demonstrated by Hahn & Angulo (2016).

4.3.3 Structure classification

As we have already mentioned in the introduction, the anisotropic nature of cosmological gravitational collapse together with the absence of thermalisation in collisionless dynamics

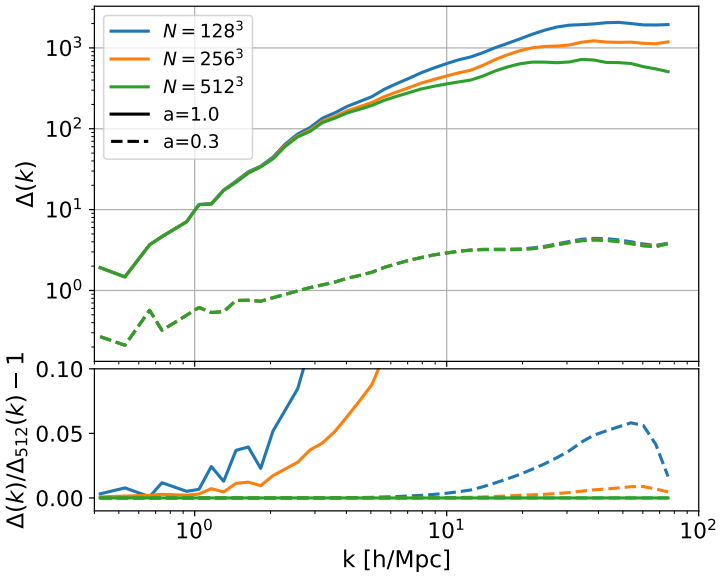

Figure 4.7: Comparison of dimensionless power-spectra (top) and residuals with respect to the highest resolution (bottom) at two different times. At an early time $a = 0.3$ the power spectra converge well, since the resolution is sufficient for reconstructing the sheet. At late times $a = 1.0$ the power spectra converge only slowly, the sheet is too complex for reconstruction with a fixed finite resolution.

(cf. Buehlmann & Hahn, 2018) makes anisotropically collapsed structure particularly vulnerable to particle noise. At the same time, the lower dimensional dynamics in those regions restricts the dynamics severely, so that ultimately it would be desirable to disentangle the unproblematic regions where dynamics is close to ergodic in all dimensions (haloes) from the problematic ones where this only true for dynamics in subspaces (filaments, sheets and voids). Therefore we developed a scheme to classify particles into void, pancake, filament or halo particles, purely by their dynamics. This scheme was already explained in section 2.6.4. (However, it is part of the paper draft that this chapter is based on.)

We show in Figure 4.8 the relative differences between GDE and finite difference stream densities for different structure types. Clearly particles in haloes have significantly larger errors on the stream densities than those in other structures. Simplifying a bit we can say that the sheet interpolation works well outside of haloes, but tends to break down inside of haloes. This is good news for the possibility of fragmentation-free warm dark matter simulations: It is well known that N-body simulations tend to fragment in filaments. Here small discrete errors in the density estimate can quickly evolve into artificial fragments, since the (coarse-grained) distribution is heated up in two dimensions, but still cold within the third dimension. Errors can easily couple in this situation and amplify. However, in the case of a halo, where the distribution is heated up in all dimensions, it is unlikely that small errors cause significant larger scale errors. On the other hand sheet schemes do not have the problem of fragmentation, because their density estimate is much more

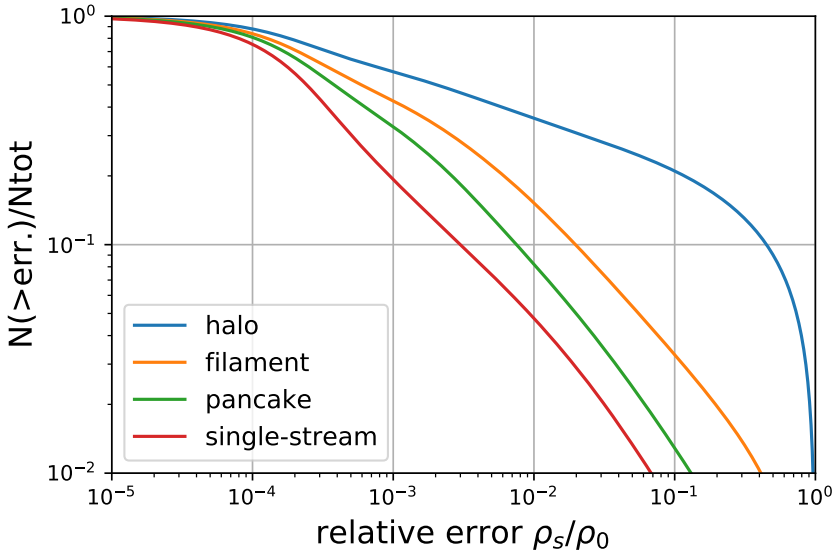

Figure 4.8: Cumulative histogram of the relative differences of the GDE and finite difference stream densities for different structure types. The dark matter sheet can be reconstructed well outside of haloes, but poorly within haloes.

accurate and less noisy than the N-body density estimate in low-density regions such as single-stream regions, pancakes and filaments. However, they become intractable inside of haloes because of the rapidly growing complexity.

It is an obvious next step to combine the benefits of the two schemes to make fragmentation-free unbiased warm dark-matter simulations. Therefore we develop a scheme that initially traces all matter by a sheet-interpolation scheme, then detects (Lagrangian) regions where the sheet becomes too complex to be followed accurately by interpolation and switches to an N-body approach for those. We call this switch to an N-body approach "the release".

4.3.4 The release

Initially we trace all mass in the simulation by a sheet interpolation scheme. That means we follow the dynamics by a set of mass-less particles which we call *flow tracers* which follow normal Newtonian equations of motion. However, forces are not estimated from the usual N-body interactions, but instead by assigning the mass of the interpolated sheet to a mesh (or another Eulerian discretization structure as will be discussed in section 4.4) and solving the Poisson equation from there. When we say that we release a Lagrangian volume element, we mean that from that point on its mass is no longer assigned by using the sheet interpolation, but simply by depositing independently traced N-body particles. We illustrate this in the bottom panel of Figure 4.3.

In our code we have two different methods available for choosing the positions and velocities of the released N-body particles. The first method simply evaluates the phase

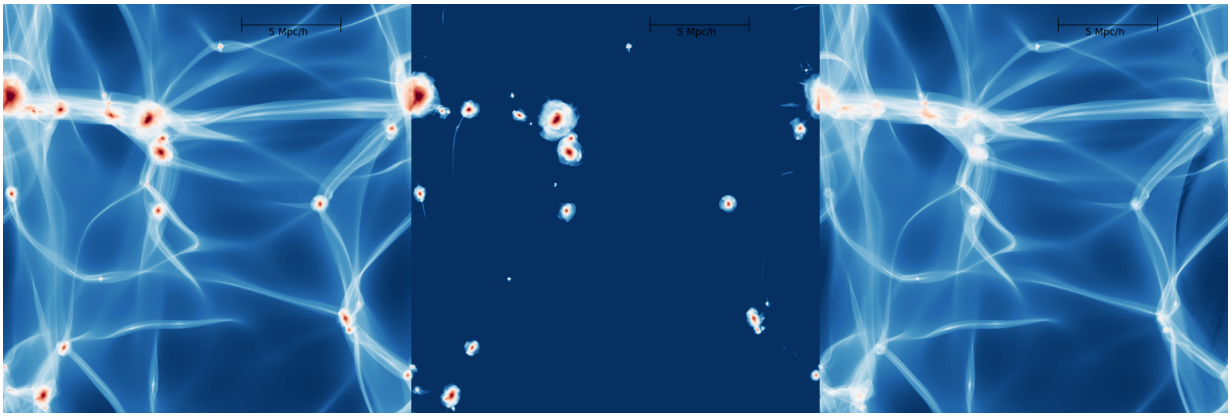

Figure 4.9: Density projection of a simulation using a hybrid Sheet/N-body scheme. Left: Projection of the total mass. Center: Projection of the mass that is traced by N-body particles. Right: Mass that is traced by the sheet-interpolation. This combination of schemes produces simulations with high quality density estimates in low-density regions which are simultaneously correct in haloes.

space coordinates of the interpolated dark matter sheet $(\mathbf{x}(\mathbf{q}), \mathbf{v}(\mathbf{q}))$ at a set of Lagrangian coordinates q_{ijk}. For example if a cubic Lagrangian volume element starting at $\mathbf{q}_0$ with Lagrangian side length Δq should be released into N^3 N-body particles, their parameters would be chosen the following way:

$$\mathbf{q}_{ijk} = \mathbf{q}_0 + \frac{\Delta q}{2(N+1)} \left(1+2i,\, 1+2j,\, 1+2k\right)^T \tag{4.6}$$

$$\mathbf{x}_{ijk} = \mathbf{x}(\mathbf{q}_{ijk}) \tag{4.7}$$

$$\mathbf{v}_{ijk} = \mathbf{v}(\mathbf{q}_{ijk}) \tag{4.8}$$

$$m_{ijk} = N^{-3} \rho_0 \Delta q^3 \tag{4.9}$$

where i, j and k are going from 0 to $N-1$. The nice thing about this release scheme is, that it is only necessary to trace N-body particles for the released mass, and it is in principle possible to control the mass resolution within released regions (mostly haloes as we shall see) independently. However, while this works fine in principle, it is relatively sensitive to making even a small error in the interpolation at the time of the release. If e.g. the velocity of the newly created N-body particles is off by just a percent it can already affect their future trajectories by a large amount. We found that this can lead to peculiar effects in some cases - for example some N-body particles which are expelled slightly beyond the splashback radius of a halo.

To make sure that there are no artefacts created from small inaccuracies at release, we implemented an additional release method. In this method we trace an additional set of mass-less particles at the locations $\mathbf{q}_{ijk}$ from the beginning of the simulation. We call these particles *silent particles* since they have no impact on the simulation prior to their release. However, when their Lagrangian volume element is released they are converted to N-body

particles with the mass m_{ijk} - thus creating new particles at exactly the correct location with exactly the correct velocity and no interpolation errors. We use this release method for the remainder of this paper.

Note that the release is defined per Lagrangian volume element. Therefore at the same location in Eulerian space released N-body particles and interpolated sheet-elements can coexist.

4.3.5 Release criterion

It is important to reliably identify the Lagrangian elements for which the sheet becomes too complex to be traced. We have developed two different criteria to flag elements for the release. The first criterion compares for each particle in a volume element the finite-difference distortion tensor $D_{xq,fd}$ with the GDE distortion tensor $D_{xq,gde}$. If their alignment $a(D_{xq,fd}, D_{xq,gde})$, which we define as

$$a(A, B) = \frac{\sum_{i,j} A_{ij} B_{ij}}{\sqrt{\sum_{ij} A_{ij}^2 \sum_{ij} B_{ij}^2}}, \tag{4.10}$$

becomes smaller than a threshold value a_{min} for any particle in a volume element, we flag that element for release. We show density projections of a simulation which uses this release criterion with $a_{min} = 0.99$ in Figure 4.9. The benefits of the release technique become clearly evident here. It allows us to get realistic haloes, while at the same time getting the accurate non-fragmenting sheet-density estimate in the low-density regions.

Additionally we defined an alternative criterion which simply releases all Lagrangian volume elements where at least one particle becomes part of a halo. To detect whether a particle becomes part of a halo we use the criterion based on the angles of the distortion tensor as described in section 2.6.4.

At first sight it might seem unnecessary to define an additional criterion here, since the other criterion seems to work well, and has a clearer quantitative justification. However, the benefits of this "halo-criterion" are (1) it can also be used without needing to integrate the GDE for all particles (which can be quite expensive) since also the finite differencing distortion tensor can be used for detecting haloes. (2) It can be hard to get the GDE and finite-difference distortion tensors into exact agreement in the continuum limit. This requires that the tidal tensor T_{ij} corresponds exactly to the derivatives of the forces in the sense that $F(\mathbf{x} + \Delta\mathbf{x}) = F(\mathbf{x}) + T\Delta\mathbf{x}$. At first sight this might seem relatively trivial to achieve, but we want to point out here that this is not so for a mesh based Poisson solver like that in Gadget2 which does not even exactly ensure that $\partial_i F_j = \partial_j F_i$.

While we find it to be possible to bring the two distortion tensors into exact agreement in the continuum limit when using a mesh based scheme, it seems relatively hard when using a tree as we will explain in section 4.4. Therefore we will use criterion (1) whenever possible, but fall back to the halo-criterion when necessary. However, we will show that the difference is negligible in most cases and the two criteria lead to the same results.

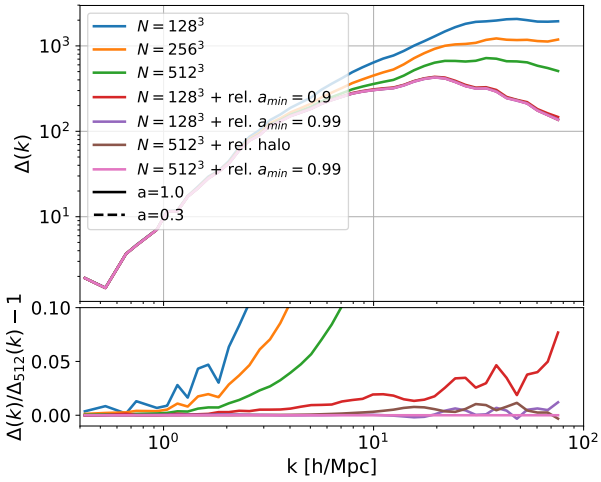

Figure 4.10: Convergence of the dimensionless power-spectra when using the release. In the top panel the power spectra and in the bottom panel the residuals with respect to the (pink) case with $N = 512^3$ and using $a_{\min} = 0.99$ for the release criterion. All simulations that use a release appear to be converged very well relatively independent of the release criterion or resolution. The only release case that disagrees above the percent level is the case with $N = 128^3$ and $a_{\min} = 0.9$ showing that the more conservative choice of $a_{\min} = 0.99$ is better.

4.3.6 Convergence of power spectra

To test the validity and the convergence of the release method, we run a series of sheet-simulations at particle resolutions of $N = 128^3$, $N = 256^3$ and $N = 512^3$ in a 20 Mpc/h box in an $m_X = 250\,\mathrm{eV}$ warm dark matter cosmology. The number of silent particles which are used for the release are two times as many per dimension in each case. For these simulations we do not use any refinement. We test the different release criteria for these methods; on the one hand the release with the halo criterion and on the other the release with the alignment criterion for different values of $a_{\min}$ as in equation (4.10). We show the power spectra of these runs in Figure 4.10 and also plot the residuals in comparison to the case which we consider most accurate - that is $N = 512^3$ with a release criterion of $a_{\min} = 0.99$ which is the same one plotted in Figure 4.9.

Clearly the pure sheet simulations (without refinement) have very biased power spectra on small scales and are far from converged. Note however, that in most of the volume their density estimate is exactly the same as in the simulations with release (compare right panel of Figure 4.9). The power spectrum is dominated by the density distribution in haloes. All simulations that use a release agree fairly well and their results seem to be relatively independent of the release criterion or the resolution. In all cases almost no elements outside of haloes are selected for release (since the dark matter sheet is relatively simple outside of haloes) and almost all mass inside of haloes is released. It seems that the distortion tensor based released criterion with $a_{\min} = 0.99$ is a good choice, but also that

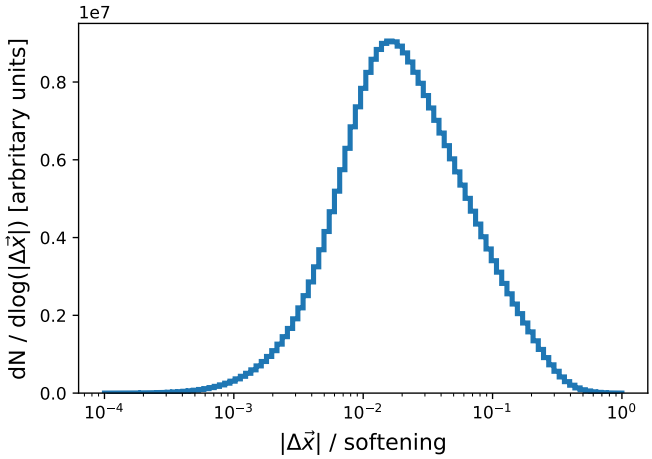

Figure 4.11: Histogram of the position error of the sheet interpolation at the point of release for the $N = 512^3$, $a_{\min} = 0.99$ simulation. This measures the error of the sheet interpolation at the point where it is largest - after this point it is considered unreliable and is no longer used for the corresponding elements. Even at its worst point, the interpolation error is still much smaller than a softening length (the grid spacing) in all cases. This shows that the release is happening early enough to make sure the interpolation is always reliable when used.

the release criterion based on the halo classification works well.

To get a quantitative idea of how large the interpolation error is in the worst case, we rerun the simulation with $N = 512^3$ and the $a_{\min} = 0.99$ release criterion. We save for every element that is released the difference in position between the silent particles (which are at the exactly correct position) and the corresponding positions predicted by the interpolated sheet (which can be slightly biased) at the time of release. This measures the absolute error of the interpolation at the point where it is largest and after which our code does not use it anymore for the corresponding mass element. We show those errors in Figure 4.11. Even these worst case errors are much smaller than a softening length (the grid spacing) in all cases. That means that the usage of the interpolation until that point cannot bias the force calculation significantly. The release criterion makes sure that mass elements are released early enough.

To highlight the improvement of the sheet simulations with release over pure sheet or pure N-body simulations we show in Figure 4.1 a projection of the density field around the highest mass halo for four different cases: an $N = 128^3$ N-body simulation (top left), a high resolution N-body simulation ($N = 512^3$) as reference (top right), a pure $N = 128^3$ sheet simulation (bottom left) and a $N = 128^3$ sheet + release ($a_{\min} = 0.99$) simulation (bottom right). In the pure sheet simulation the halo appears much rounder than it should. In the sheet + release simulation the halo has the same shape as in the high resolution N-body simulation but there is no artificial clumping in the surrounding structures as in the $N = 128^3$ N-body case. Clearly the sheet + release scheme inherits the best of each

of the methods. Note that the $N = 512^3$ N-body simulation is fragmentation-free, since here we are using a relatively low force resolution. However if the force resolution were higher than the initial separation between particles (as usually assumed) this case would also fragment.

We conclude that realistic, fragmentation-free warm dark matter simulations can be achieved when combining the benefits of sheet and N-body methods into a sheet + release scheme. In the following section we will address how to achieve higher force resolution in the sheet + release scheme.

4.4 Towards higher force-resolution

We have shown in section 4.3 how to make unbiased and fragmentation-free warm dark matter simulations with a fixed, relatively large force resolution and a global time-step. However, the force-resolution that can be achieved with a regular mesh is much below that necessary to get convergence in the centers of haloes. Therefore we develop here a new scheme to calculate gravitational forces in cosmological simulations.

N-body simulations typically discretize forces as interactions between point-particles. All additional components of the force-calculation like trees, particle-meshes or adaptive mesh-refinement are just means to speed up the force-calculation. While this approach of pair-wise interactions is relatively simple and works remarkably well, there are also some peculiarities that arise from this approach when compared to the true continuous physical system. One of those is that the "N-body density-estimate" is exactly zero in the large majority of the volume: In N-body simulations typically a softening length as small as 1/20 of the mean particle separation is chosen. The density estimate $\rho = \nabla^2\phi/4\pi G$ is zero outside of a particle's softening radius. Consequently that means that less than $(1/20)^3 \sim 10^{-4}$ of the volume has a non-zero density in such an N-body simulation. However, we know that the true density can be zero nowhere, since at every point in space at least one dark matter stream must be present. Depending on the nature of dark matter, the lowest density regions in the universe may still have densities of $\rho \sim 10^{-1} - 10^{-3}\rho_0$ (Stücker et al., 2018). Also the fragmentation of N-body simulations is a peculiar effect of the granularity of the N-body density estimate.

We here propose a different discretization of the density field and the force-calculations that accounts better for the continuous nature of the dark matter sheet. We discretize the density field by a space-filling oct-tree of cubes. The depth of the tree is chosen adaptively and depends (roughly) on the local density. The density of the cubes and their density gradients (assumed uniform) are calculated by sampling them with a larger number of pseudo-particles which are created from the interpolated sheet. The force that a single particle experiences is then the sum of interactions with cubes.

In the following sections we describe in more detail how the tree is constructed, how the densities and density gradients of the tree-nodes are computed and how the interaction with a cube can be computed. Subsequently we demonstrate for the case of a Hernquist sphere that this discretization of the force field is as good as an N-body representation with

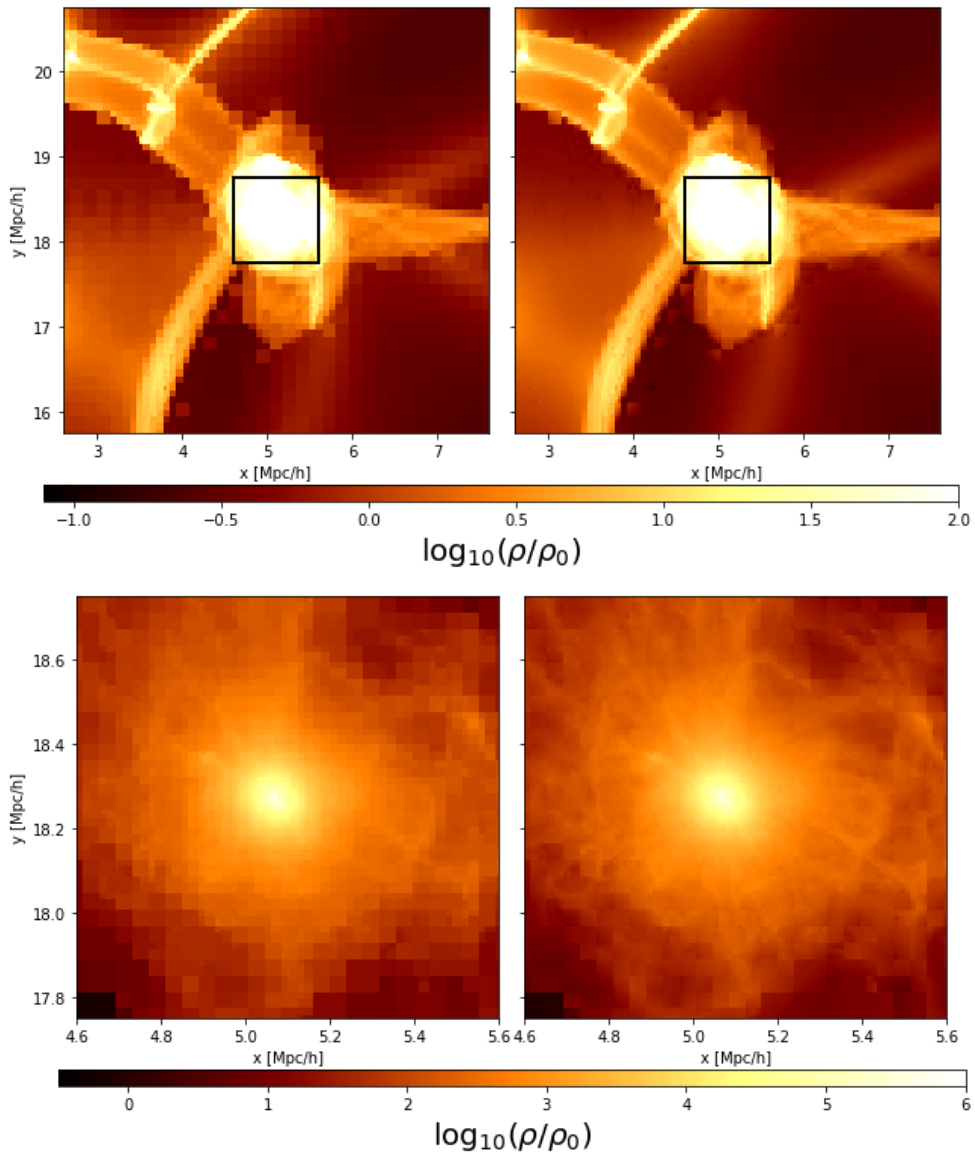

Figure 4.12: The density field that is represented by an oct-tree of cubes from one of our simulations. Left: using only homogeneous cubes, right: using cubes with constant density gradients. Top: 5 Mpc region - the colormap is clipped at a density of $100\rho_0$ to make the features in the low density regions visible. Bottom: 1 Mpc region showing the internal structure of the halo with an appropriate colormap range.

the same number of mass-resolution elements, but a much larger number of force-resolution elements.

4.4.1 A tree of cubes

We discretize the density field in our simulations as an oct-tree of cubes. An oct-tree is a recursive partition of a three dimensional volume into a set of cubes. One starts with a cube which represents the whole simulation box. It is split into eight equal size sub-cubes, each of them representing another oct-tree. Each attached oct-tree can either be a leaf or again be split into eight sub-cubes representing their own subtrees - and so on. It depends on the tree-building procedure which sub-trees are split.

To make minimal changes to the code, we employ the same tree building mechanism as the Gadget 2 code as presented in Springel et al. (2005b). That is, given a set of particles, the tree is split recursively until each leaf contains either one or zero particles or a minimum node size has been reached. That minimum node size is an additional parameter in the code and determines the limit of the force resolution. We will later refer to this parameter as the softening scale of the tree. In our simulations the tree structure is not built from all particles, but from a smaller set of distinctly defined ghost particles.

After the tree has been created it is just considered as a mass-less partition of the volume of the simulation box. Afterwards the mass is assigned to it in an additional independent step. Lagrangian regions that are still described by the sheet interpolation create a large number of mass-carrying pseudo-particles that are deposited into the tree whereas N-body particles are directly deposited into the tree. Each mass-carrying particle assigns its mass and its mass weighted position to the tree leaf it belongs to. Afterwards this information is propagated upwards in the tree so that at the end of the procedure each node has a well defined mass and center of mass. Typically we choose a much higher number of mass tracing particles than of tree building particles, so that the mass in each tree node is well sampled. For example we have 4^3 times more N-body particles than particles which define the tree structure, and the number of pseudo particles that are deposited from the sheet interpolation is again much higher than the number of N-body particles. Therefore each tree leaf is sampled by the order of 64 particles in completely released regions and by many more in regions where the interpolation is active.

Now we interpret this tree as a set of cubic volume elements which have a mean density given by their mass and volume, and which have a density gradient given by the center of mass. The density distribution within each cube is then approximated as

$$\rho(\mathbf{r}) = \rho_0 + \mathbf{g} \cdot \mathbf{r} \tag{4.11}$$

$$\rho_0 = M/L^3 \tag{4.12}$$

$$\mathbf{g} = \frac{12\rho_0}{L^2}\mathbf{r}_{\rm cm} \tag{4.13}$$

where M is the mass of the cube, L is the side-length, $\mathbf{r}_{\rm cm}$ is the offset of the center of mass from the center, and $\mathbf{r}$ is the offset from the geometric center at which the density is

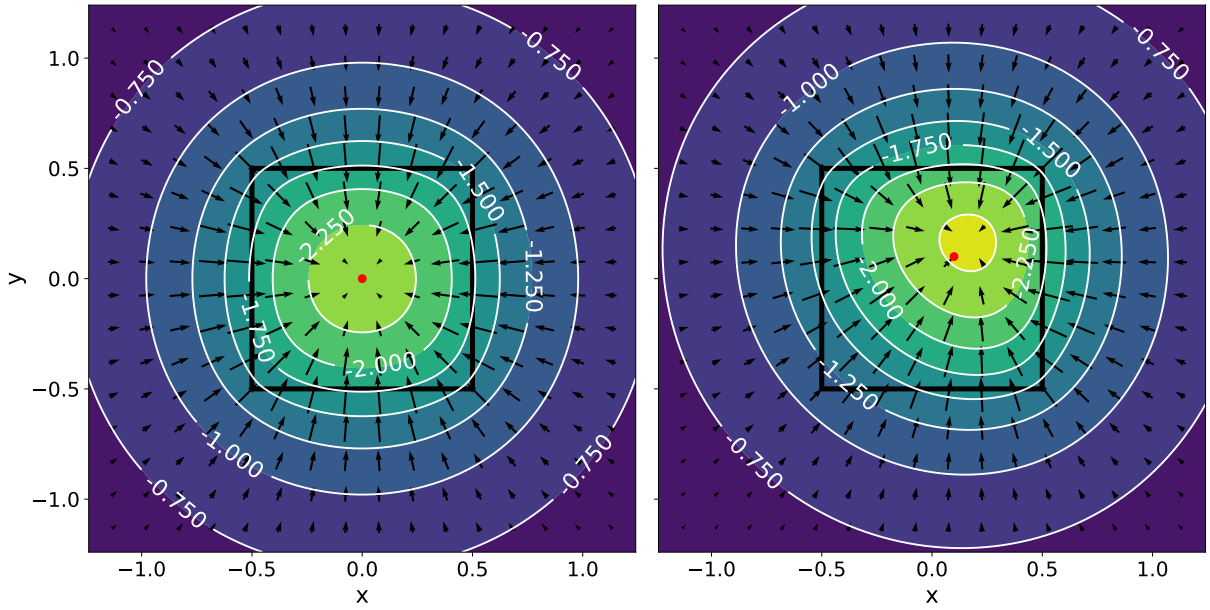

Figure 4.13: Potential (contours) and force field (arrows) of a homogeneous cube (left) and a cube with constant gradient (right). The equipotential-lines are closer to a sphere than the cubic mass distribution and farther away from the cube they become more and more spherical (centering around the center of mass - marked as a red dot).

to be evaluated. Note that the gradient is chosen so that the cube with gradient has the same center of mass as the node.

We show the density field that is represented by the oct-tree in one of our simulations in Figure 4.12. Clearly this seems to be a promising way of discretizing the density field.

4.4.2 The potential of a cube

While it is relatively simple to define the density field of this oct-tree, calculating its force field is mathematically relatively elaborate. To ease the reading flow we moved the full potential calculation to the end of this chapter 4.6.1 - 4.6.3. We just give a brief summary of the necessary steps here.

The potential of the cubic mass distribution can be calculated as the convolution of its density distribution with the Green's function of the potential

$$\phi(\mathbf{x}) = \rho \circ G_\phi \tag{4.14}$$

$$\rho(\mathbf{x}) = \begin{cases} \rho_0 + \mathbf{gx} & \text{if } -L/2 \leq x_i \leq L/2 \forall i \in \{1,2,3\} \\ 0 & \text{else} \end{cases} \tag{4.15}$$

$$G_\phi(\mathbf{r}) = -\frac{G}{\|\mathbf{r}\|} \tag{4.16}$$

where G is the gravitational constant. It turns out that the result of this convolution is analytical but gives a rather complicated expression. We present this in detail in section 4.6.1. In Figure 4.13 we plot the potential in the $z = 0$ plane for the example of a homogeneous cube ($\rho_0 = G = 1$, $\mathbf{g} = (0, 0, 0)^T$) and the case of a cube with constant density gradient ($\mathbf{g} = (1.2, 1.2, 0)^T$).

However this is not the kind of interaction that is actually summed over in Gadget. In Gadget the potential is split into a long-range part and a short-range part (cf. Hockney & Eastwood, 1981). The long-range part is a smoothed version of the potential (described by a convolution with a kernel f) and is calculated with Fourier methods on a mesh. The short-range part is then the remaining part of the potential

$$\phi_s = \rho \circ G_\phi \circ (1 - f) \tag{4.17}$$

This is the interaction that needs to be evaluated on the tree for maintaining consistency with the force-split. It turns out that for the Gaussian force-split that is used in Gadget, this convolution cannot be solved analytically for our mass distribution. Therefore we decided to use a different force-split kernel

$$f(\mathbf{r}) = \begin{cases} \frac{3(a-\|\mathbf{r}\|)}{a^4\pi} & \text{if } \|\mathbf{r}\| \leq a \\ 0 & \text{else} \end{cases} \tag{4.18}$$

for which the potential becomes analytic. As a drawback this piecewise defined kernel creates a large number of different cases. We only calculated the explicit expression for the most typical case that the whole cube is inside the range of the kernel. The full calculation can be found in section 4.6.2. For all other cases we use numerical approximations by sampling the mass distribution with point-masses or splitting the cube into eight sub-cubes.

Further the cost of the calculations can be drastically reduced by using multipole expansions. We give an overview of the approximations that we use in different cases in 4.6.3. We make sure that the errors in the forces and the tidal tensor due to a single cubic element stay well below 1% in all cases.

4.4.3 Force field of a Hernquist sphere

To test whether the calculation of the force field by the tree of cubes gives reasonable forces in complex three dimensional scenarios, we set up a Hernquist sphere (Hernquist, 1990). To do that we create a set of N particles that are a Poisson sampling of the 3D Hernquist density profile

$$\rho(\mathbf{r}) = \frac{M}{2\pi} \frac{a}{\|\mathbf{r}\|} \frac{1}{(\|\mathbf{r}\| + a)^3} \tag{4.19}$$

where a is the scale radius and M is the mass of the Hernquist sphere. We calculate the forces for that particle realization using the typical N-body softened monopole approach

and the cube-tree approach presented in this paper. The structure of the cube-tree is build on a set of N_c particles which are an independent set of Poisson sampled particles and typically $N_c \ll N$. The mass and gradients of the nodes are then determined by depositing the N particles into the tree. We show the results of the force computation and the difference with respect to the analytical force

$$F(r) = \frac{GM(r)}{r^2} \tag{4.20}$$

$$M(r) = M\frac{r^2}{(r+a)^2} \tag{4.21}$$

in Figure 4.14. To calculate the residuals we use only the radial component of the force:

$$F_{p,r} = \frac{\mathbf{r} \cdot \mathbf{F}_p}{\|\mathbf{r}\|} \tag{4.22}$$

$$R_p = \frac{F_{p,r} - F(r)}{F(r)} \tag{4.23}$$

where $\mathbf{F}_p$ is the full force vector of a particle and R_p is the relative error of the radial force as plotted in the lower panels of Figure 4.14.

It seems that the cube tree and the monopole interactions give very similar results when compared at the same number of mass-carrying particles N. They have a similar spread in the distribution which is mostly caused by shot noise and decreases when increasing the particle number N. The error around $r/a \sim 40$ that does not converge to zero is likely caused by inaccuracies in the Tree-PM force-split. It has a different (smoother) shape in the N-body monopole calculation, since that calculation uses a Gaussian kernel for the force-split where the short-range force is known exactly in all cases and does not need to be approximated. However, the error from the force-split is of the order 1% and therefore negligible. The force deviates at very large distances, because we use periodic boundary conditions for the force calculation. Further it deviates at very small scales because of the softening. In the cube-tree case the softening is defined as the minimal allowed node size.

While the cube-tree has a similar accuracy to the monopole-method in the case where the same number of mass-depositing particles is used, it has a significantly higher accuracy when compared at the same number of force resolution elements. In the cube tree this is of the order N_c and in the N-body case this is N. For example the cube-tree case with $N_c = 16^3$ and $N = 64^3$ has roughly the same accuracy as the monopole case with $N = 64^3$ whereas it has roughly 64 times fewer force resolution elements.

We conclude that the representation of the force field by a tree of cubes will be of similar accuracy to the monopole method inside of haloes where we assume all particles to be released and roughly a Poisson-sampling of the density distribution. However, outside of haloes the cube-tree can provide a much more accurate description of the force-field than the monopole tree, since we can use the sheet-interpolation to sample the tree with many more mass-carrying particles than would be possible in a pure monopole approach. Therefore the tree of cubes can be used to obtain at the same time a smooth and continuous

representation of the density- and force-field outside of haloes and a reasonably good one inside of haloes.

As another possible application it could be used to store an accurate representation of the force-field of an N-body simulation at much lower cost than saving all the individual particles. It would be enough to output the nodes of the tree down to some level. Forces could then be recalculated as a post-processing step as a sum over the node-interactions.

It is of further interest for us to check the accuracy of the calculated tidal field in the different approaches. In Figure 4.15 we plot the relative errors of the tidal tensor

$$R = \sqrt{\frac{\sum_{i,j}(T_{p,ij} - T_{ij}(r))^2}{\sum_{i,j} T_{ij}^2(r)}} \tag{4.24}$$

where T_p is the tidal tensor that has been evaluated and $T(r)$ is the analytical tidal tensor. It is striking that the errors easily can reach order unity depending on the setup. The tidal field is a quantity which is much harder to determine from the noisy particle distribution in a simulation. Unlike the force field, it depends crucially on the local density distribution.

The error in the tidal tensor is very large in the monopole interaction cases, and seems to be lower than that in almost all of the cube-tree cases. The noise in the cube-tree cases seems to be smaller, the larger the ratio of mass depositing particles N to force resolution elements N_c. For example the case with $N_c = 16^3$ and $N = 128^3$ performs better than the higher force-resolution case with $N_c = 32^3$ and $N = 128^3$. This case is very similar to $N_c = 16^3$, $N = 64^3$, so it seems to be N/N_c which matters. In the case with a higher N/N_c the moments of the lowest level nodes are determined more accurately and therefore the noise in their density estimate is smaller. Note that the density estimate constitutes the trace of the tidal tensor and its accuracy is of crucial importance for the accuracy of the tidal tensor.

With the cube-tree we can significantly reduce the errors in the calculations of the tidal tensor. However, the errors are still of worrying magnitude. It can be significantly reduced by increasing the number of deposited mass tracers. Therefore increasing N at a fixed number of force resolution elements N_c can be used to test the implications of the noise in the tidal tensor for the accuracy of the GDE integration.

4.4.4 Time evolution of a Hernquist sphere

In our scheme for the force calculation, the force resolution can vary in space and in time. That is similar to the case of an adaptive gravitational softening (Price & Monaghan, 2007; Iannuzzi & Dolag, 2011). As a consequence the time evolution will not be strictly energy conserving and not formally symplectic. To test whether this causes any major problems for the evolution of systems typical for cosmological simulations, we perform a time evolution of the Hernquist-sphere. For the cube-tree cases that means that both the (massless) tree-structure-defining particles and the mass-carrying particles are integrated along their orbits, and the structure of the tree can change with time. As an additional

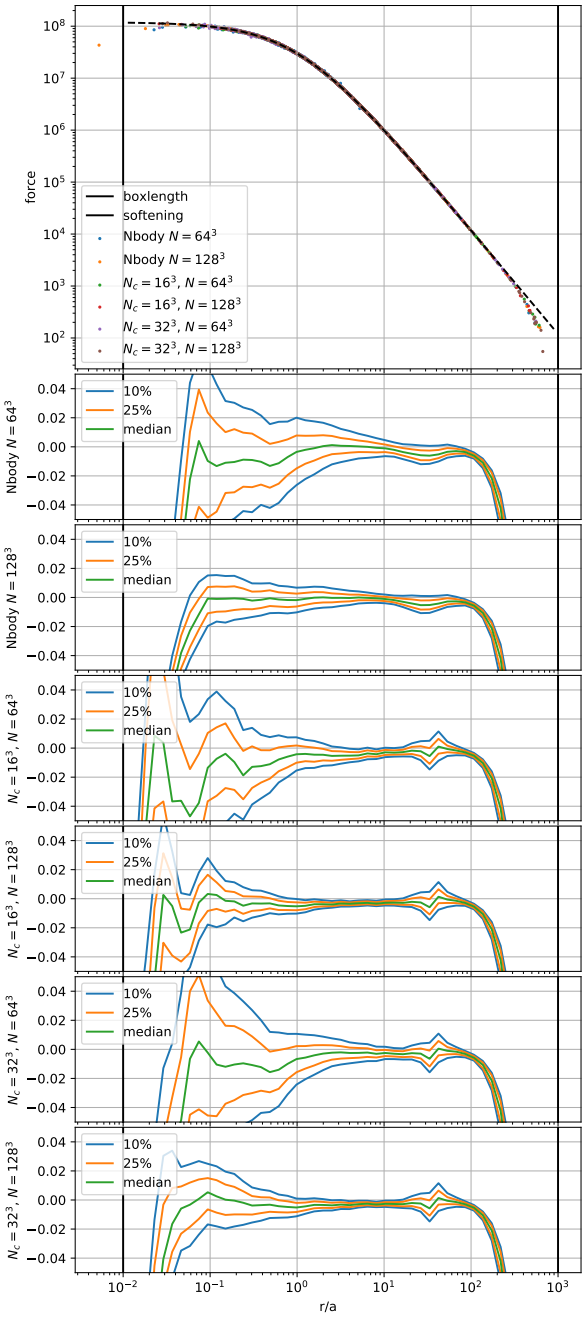

Figure 4.14: Absolute values of the force (top panel) and relative errors (other panels) for the numerically evaluated forces in a Hernquist sphere using different setups. Second and third panel: the quantiles of the force error distributions when using the traditional N-body approach for evaluating the forces with two different particle numbers $N = 64^3$ and $N = 128^3$. Bottom four panels: Using the cube tree with different number of tree building particles N_c and mass depositing particles N. In the cube tree cases the softening indicates the minimal allowed node length. The cube tree gives similar results as the N-body monopole force calculations at the same number of mass depositing particles. It gives significantly better results for the same number of force resolution elements.

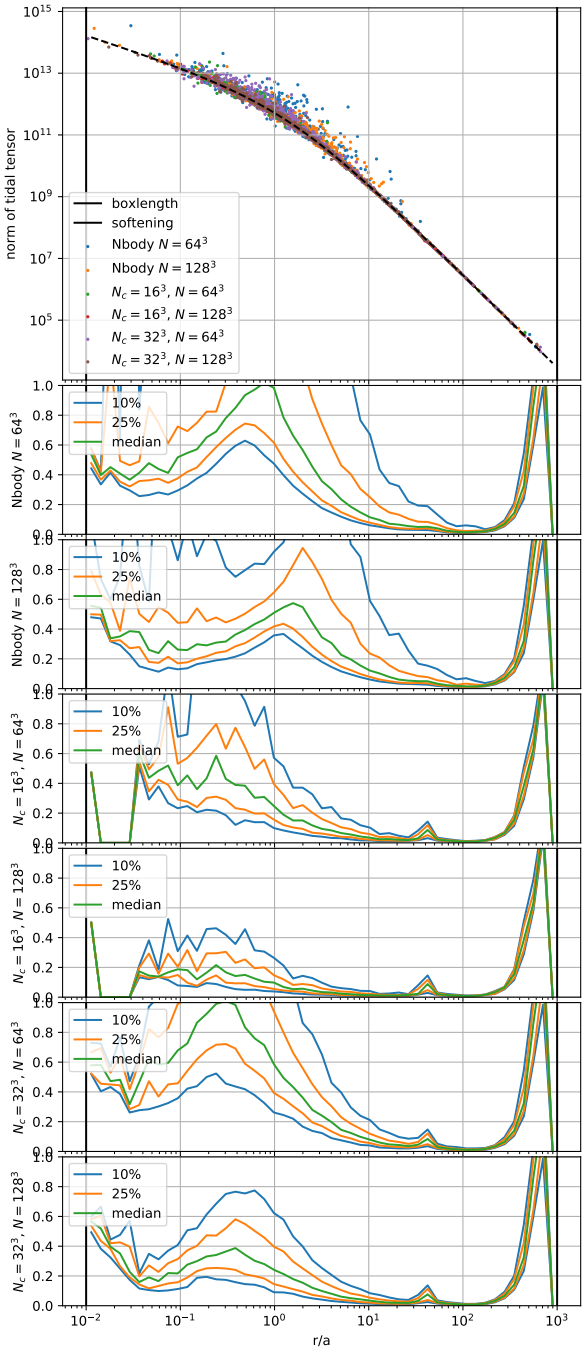

Figure 4.15: L2 norm of the tidal tensor (top panel) and relative errors (other panels) for the numerically evaluated tidal tensors in a Hernquist sphere using different setups. Second and third panel: the quantiles of the error distributions when using the traditional N-body approach for evaluating the tidal tensor with two different particle numbers $N = 64^3$ and $N = 128^3$. Bottom four panels: Using the cube tree with different number of tree building particles N_c and mass depositing particles N. In the cube tree cases the softening (the thick vertical black line) indicates the minimal allowed node length. The cube tree gives a better estimate of the tidal tensor than the N-body calculation in all setups. The estimates improve when increasing mass resolution or decreasing force resolution.

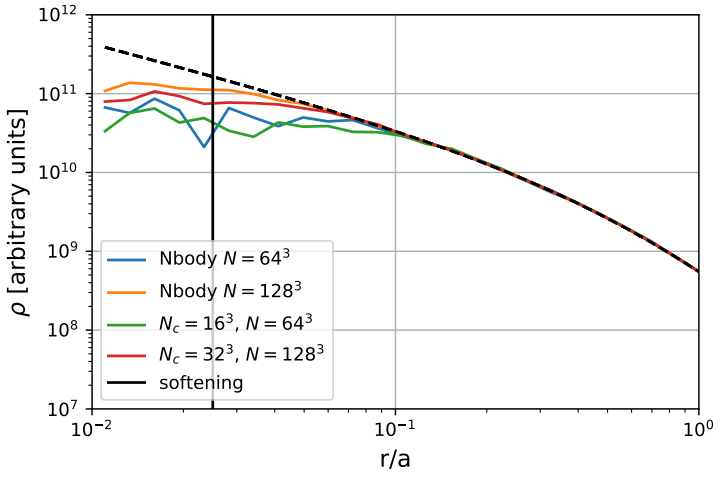

Figure 4.16: Density profiles of an evolved Hernquist-sphere. The profiles from the cube-tree method (green and red) show a similar degree of two-body relaxation to the ones from the N-body scheme (blue and orange) when compared at the same number of mass resolution elements N.

way of reducing aliasing effects due to the positioning of the oct-tree, we randomize its base positioning each time-step.

We evolve the Hernquist-sphere for a time span of $t = 110t_*$ where t_* is the natural time scale of the system

$$t_* = \sqrt{\frac{a^3}{MG}} \,. \tag{4.25}$$

In that time a particle on a circular orbit at a radius of $0.1a$ will have gone through about 90 periods. If there are any effects beyond the typical two-body relaxation, these should show up by that point.

In Figure 4.16 we show the density profile of the evolved Hernquist sphere for the different numerical setups and in Figure 4.17 we show the distributions of energy and angular momentum errors. It appears that N-body approach and cube-tree approach have almost equivalent accuracy when compared at the same mass resolution. The profiles diverge from the theoretical ones at similar radii and the energy- and angular-momentum errors are of similar amplitude. Note, however, that the cube-tree approach uses of order 64 times fewer force-resolution elements in those cases. The errors in the cube-tree scale just like the N-body ones with the mass-resolution and are therefore due to two-body effects. We do not find any major additional errors caused by the variable force resolution.

We conclude that the cube-tree performs as well as a pure N-body scheme when compared at the same mass resolution while requiring significantly less force-resolution elements. It is therefore well suited to be used in a hybrid N-body/sheet scheme, since it can capitalize on the high mass resolution that becomes available through the sheet, without compromising the consistency of the force calculation.

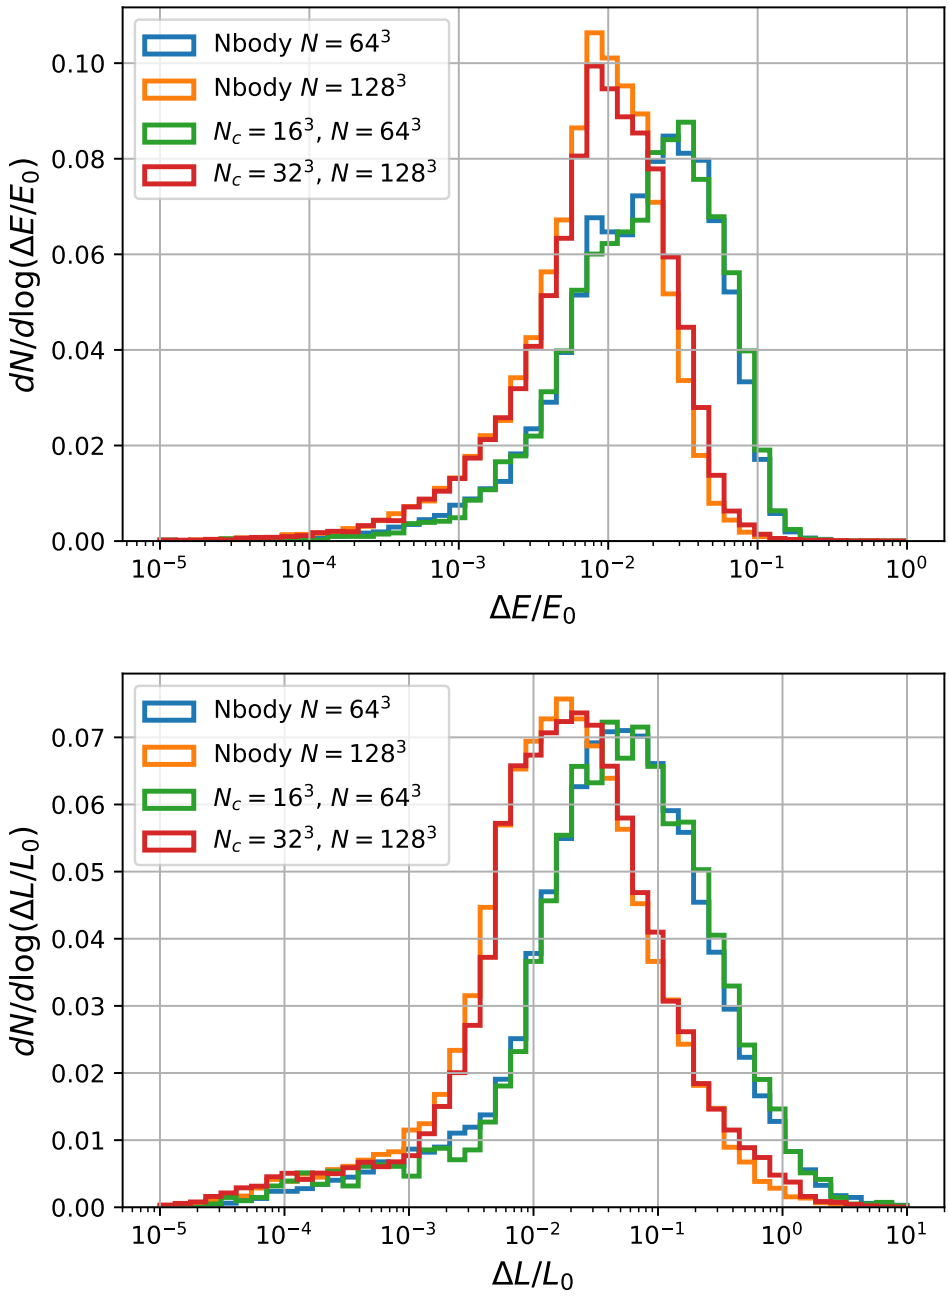

Figure 4.17: Histogram of the relative change in the energy (top) and angular momentum (bottom) of the particles in the evolved Hernquist-sphere. If there were no numerical errors these should be zero. The numerical errors of the N-body scheme (blue and orange) and of the cube-tree are almost of identical amplitude when compared at the same number of mass resolution elements N.

4.4.5 Time evolution of the distortion tensor in the Hernquist sphere

Additionally we test for the Hernquist-sphere whether the integration of the GDE gives reasonable results when the cube-tree is used. We have made these tests a few weeks before I had to hand in my PhD-thesis. Therefore this section is purely preliminary and will be different in the paper – depending on future tests and potential bug-fixes. However, it is important to point out this uncertainty before discussing possible results in chapter 5.

We integrate the GDE with initial conditions $\mathbf{D}_{\mathrm{Xq}}(t=0) = \mathbb{1}$ and $\mathbf{D}_{\mathrm{Vq}}(t=0) = 0$ in the same setup as described in the previous section. The tidal field is calculated as the infinitesimal derivative of the force field (up to some approximation errors). At the intermediate time $0.3t_f$ and at the final time t_f we calculate the median profiles of the stream densities and caustic counts (Figure 4.18) and the 10% and 90% quantiles (Figure 4.19). This is done by selecting all particles in thin spherical shells $|\mathbf{x}| \in (r, r+\Delta r)$ and determining the median (and quantiles) of the distribution in each shell. Additionally we provide a reference distribution which is obtained by integrating the orbits and the GDE for particles in the analytic Hernquist potential with very small time-steps.

The caustic counts seem to be robustly converged towards the correct solution, relatively independent of the method. However, this is not the case for the stream-densities. As can be seen in the left panels of Figure 4.18 and Figure 4.19, the cube-tree gives a very good estimate of the stream-density distribution at the early time $0.3t_f$ (outside of a softening affected radius $r > 0.1a$). However, at late times (the right panels of the respective Figures) the stream-densities are systematically underestimated. It is very puzzling, that the cube-tree seems to be converged so well with the resolution towards a wrong result (compare the purple and brown line in Figure 4.18). We have actually tested a lot of different parameters, such as softening, time-stepping and other combinations of force- and mass-resolution (that means changing N and N_c independently). In all these tests the result turned out to be the same, wrong curve for the stream-densities. In principle noise in the estimation of the tidal tensor can explain an underestimation of stream-densities. However, if shot noise was the problem, the error should get smaller with some of the parameters that we tested.

Resolving this issue will require additional investigations. At the time of handing in this thesis it has not yet been resolved. There are a lot of things that can be checked as consistency tests of the code:

- Does $\mathbf{F}(\mathbf{x}+\Delta\mathbf{x}) = \mathbf{F}(\mathbf{x}) + \mathbf{T}(\mathbf{x})\Delta\mathbf{x}$ hold on an infinitesimal level $\Delta x \to 0$?
- How well is the energy in the tangent space conserved ($\frac{\partial H}{\partial \mathbf{q}} = \frac{\partial H}{\partial \mathbf{x}} \cdot \mathbf{D}_{\mathrm{Xq}} + \frac{\partial H}{\partial \mathbf{v}} \cdot \mathbf{D}_{\mathrm{Vq}}$ where H is the Hamiltonian of the system.)
- If we set up a small Lagrangian patch, does its tangent space evolve consistently with the GDE?

We will leave the resolution of this issue open at this point. Note, however, that the GDE results in section 4.3 are not affected by this. These were well converged and quite robust. The problem that was pointed out here must be related to the tidal field of the cube-tree.

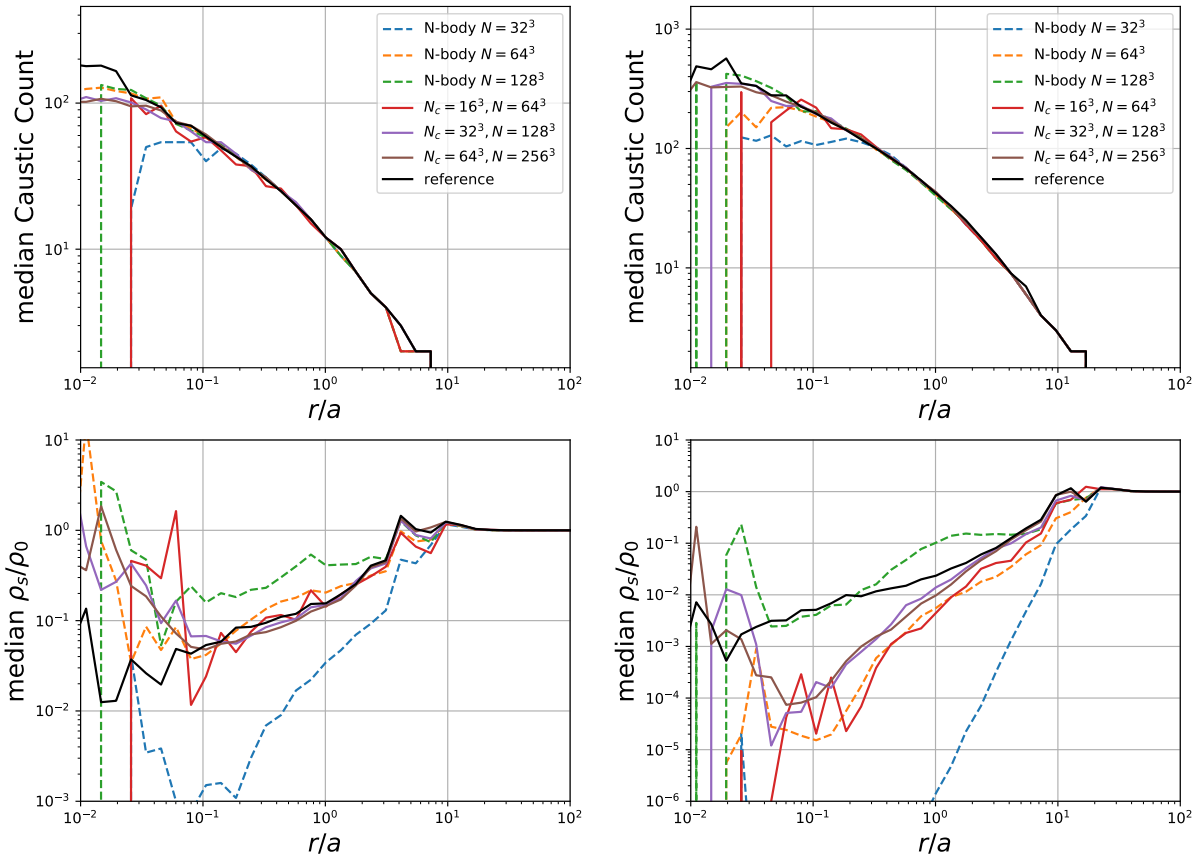

Figure 4.18: Median-profiles of the stream-densities (bottom) and caustic counts (top) in an evolved Hernquist sphere. The dashed lines show N-body simulations of varying resolutions and the solid lines simulations that use the cube-tree for the force- and tidal-field calculations. The black solid line is a reference case that was obtained by integrating the GDE in the analytic potential. Left: at an early time where particles in the central region have gone through ~ 30 orbits. The cube-tree seems to be well converged to the reference solution whereas the N-body simulation does not seem to reflect what is happening. (Note: the region $r < 0.1a$ is affected by the size of the softening $\epsilon = 0.025a$). Right: at a later time where typical particles in the center have made ~ 100 orbits. The cube-tree seems to converge to the wrong stream-densities. The caustic counts seem to converge well in all cases.

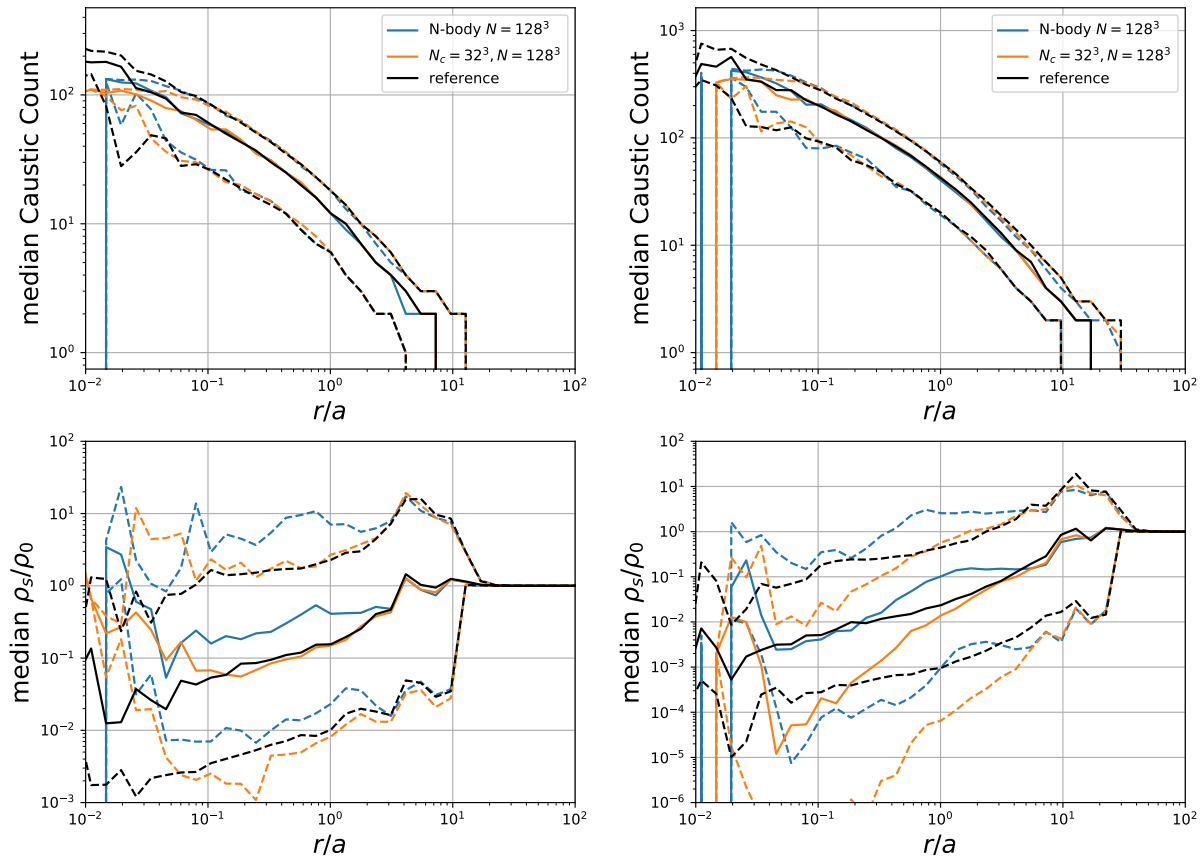

Figure 4.19: Quantile profiles of the stream-densities and caustic counts in an evolved Hernquist sphere - the median in solid and the 10%-and 90%-quantiles in dashed. Left: At an earlier time the stream density distribution that has been calculated by the combination of GDE and cube-tree (orange line) seems to reflect the analytic reference case quite well for $r > 0.1a$. Right: same plot, but at a later time - the distribution seems to turn away to too low stream densities.

4.5 Conclusions

Cosmological N-body simulations of warm dark matter suffer from artificial clumping in the sheets and filaments which precede formation of virialised haloes. Simulations based on the dark matter sheet are able to avoid this fragmentation by employing a density estimate that is less noisy and more accurate in low-density and anisotropically collapsed regions. However, they suffer from the intractable complexity of the dark matter sheet in strongly mixing regions like haloes. The two methods are thus optimal in different regions, and we have shown above that their respective strengths can be optimally used in a combined approach: a scheme that infers densities from the dark matter sheet interpolation wherever the interpolation is valid, and that switches to an N-body approach for mass elements that are too complex to be reconstructed by the interpolation. Through such a combined approach we obtain a fragmentation-free scheme for warm dark matter simulations that converges well inside and outside of haloes at affordable cost.

Further, we have developed a new scheme to calculate the forces in such simulations. N-body simulations typically use a multipole expansion of the interactions of point-like particles as the basis for the force calculations. Our scheme instead partitions space into an oct-tree and uses cubic nodes with uniform density gradients as the basic force resolution elements. In the regime where we can reconstruct the dark matter sheet, we can determine the masses and gradients of those nodes with very high accuracy. In the regime where the mass is traced by released N-body particles the cubic nodes still lead to an approximation of the force-field that compares favorably with the N-body approach: The accuracy is similar if compared at the same number of mass-resolution elements, but the accuracy is much higher if compared at the same number of force-resolution elements. Although this force calculation scheme is not exactly Hamiltonian, the energy and angular momentum errors are of the same amplitude as in the pure N-body case - even when integrated over hundreds of orbits. However, the integration of the GDE shows some numerical problems for the evaluation of stream densities (but not caustic counts) if integrated for long times, whereas it gives consistent results when integrated for shorter periods. We will try to resolve this issue before publishing this chapter.

With these new numerical methods, we are able to carry out reliable non-fragmenting warm dark matter simulations at high force-resolution. Additionally these simulations have a large degree of phase space information available. We will present such simulations in the next chapter.

Acknowledgements

We thank Mark Vogelsberger for providing us with the code for the geodesic deviation equation. Further we want to thank Volker Springel for writing the Gadget 3 code on top of which we implemented our changes. OH acknowledges funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 679145, project 'COSMO-SIMS').

4.6 Paper appendix - The potential of a cube

4.6.1 The total potential

The gravitational potential $\phi(\mathbf{x})$ of a mass distribution $\rho(\mathbf{x})$ can be obtained by convolving it with the Green's function $G_\phi(\mathbf{x})$ of the gravitational potential

$$\phi(\mathbf{x}) = \rho \circ G_\phi \tag{4.26}$$

$$= \int\int\int \rho(\mathbf{r}) G_\phi(\mathbf{x}-\mathbf{r}) d^3\mathbf{r} \tag{4.27}$$

$$G_\phi(\mathbf{r}) = -\frac{G}{\|\mathbf{r}\|} \tag{4.28}$$

where G is the gravitational constant. The mass distribution of a homogeneous cube is given by

$$\rho(\mathbf{x}) = \begin{cases} \rho_0 & \text{if } -L/2 \leq x_i \leq L/2 \forall i \in \{1,2,3\} \\ 0 & \text{else} \end{cases} \tag{4.29}$$

The potential of a homogeneous cube has already been derived in Macmillan (1958) and can be written as

$$\phi(\mathbf{x}) = -\int_{-L/2}^{L/2}\int_{-L/2}^{L/2}\int_{-L/2}^{L/2} \frac{G\rho_0}{\|\mathbf{r}-\mathbf{x}\|} d^3\mathbf{r} \tag{4.30}$$

$$= -G\rho_0 \int_{-L/2-x_1}^{L/2-x_1}\int_{-L/2-x_2}^{L/2-x_2}\int_{-L/2-x_3}^{L/2-x_3} \frac{1}{\|\mathbf{r}\|} dr_3 dr_2 dr_1 \tag{4.31}$$

$$= -G\rho_0 \left[\left[[F_{\text{hom}}(\mathbf{r})]_{r_1=-L/2-x_1}^{L/2-x_1} \right]_{r_2=-L/2-x_2}^{L/2-x_1} \right]_{r_3=-L/2-x_3}^{L/2-x_3} \tag{4.32}$$

$$F_{\text{hom}}(\mathbf{r}) = \sum_{i=1}^{3} \left(\frac{r_1 r_2 r_3}{r_i} \ln(r_i + \|\mathbf{r}\|) - \frac{r_i}{2} \arctan\left(\frac{r_1 r_2 r_3}{r_i^2 \|\mathbf{r}\|} \right) \right) \tag{4.33}$$

where we labeled the Cartesian 3d parent function of the integrand $1/r$ by $F_{\text{hom}}(\vec{r})$ and we used the notation from Chappell et al. (2013). Note that the type of integral shown in (4.31) is most easily evaluated in Cartesian coordinates. Transforming to spherical coordinates simplifies the integrand, but makes the integration boundaries very complicated, and is therefore not viable.

We show a slice through the $z = 0$ plane of the potential and the force-field of a homogeneous cube with $G = 1$, $\rho_0 = 1$ and $L = 1$ in the left panel of Figure 4.13. Close to the center of the cube forces get close zero. Close to the boundary the forces are largest and the equipotential lines deviate most from spherical ones. Then going farther away from the cube the equipotential lines approach spherical symmetry around the center-of mass.

The potential of a cube with constant gradient is given by

$$\phi(\mathbf{x}) = -G \int_{-L/2}^{L/2} \int_{-L/2}^{L/2} \int_{-L/2}^{L/2} \frac{\rho_0 + \vec{g} \cdot \mathbf{r}}{|\mathbf{r} - \mathbf{x}|} d^3\mathbf{r} \tag{4.34}$$

$$= -G \int_{-L/2-x_1}^{L/2-x_1} \int_{-L/2-x_2}^{L/2-x_2} \int_{-L/2-x_3}^{L/2-x_3} \left(\frac{\rho_0 - \mathbf{g} \cdot \mathbf{x}}{\|\mathbf{r}\|} + \frac{\mathbf{g} \cdot \mathbf{r}}{\|\mathbf{r}\|} \right) dr_3 dr_2 dr_1 \tag{4.35}$$

$$= -G \left[\left[\left[(\rho_0 - \mathbf{g} \cdot \mathbf{x}) F_{\text{hom}}(\mathbf{r}) + \sum g_i F_{\text{lin},i}(\mathbf{r}) \right]_{r_1=..} \right]_{r_2=..} \right]_{r_3=..} \tag{4.36}$$

We already know the first parent function and only need to calculate the second part which is given by

$$\begin{aligned} F_{\text{lin},i}(\mathbf{r}) &= \int\int\int \frac{r_i}{\sqrt{r_1^2 + r_2^2 + r_3^2}} dr_1 dr_2 dr_3 \\ &= \frac{1}{6} \left(2 \frac{r_1 r_2 r_3}{r_i} \|\mathbf{r}\| - 2 r_i^3 \arctan \left(\frac{r_1 r_2 r_3}{r_i^2 \|\mathbf{r}\|} \right) \right. \\ &\qquad \left. + r_j (3 r_i^2 + r_j^2) \ln(r_k + \|\mathbf{r}\|) + r_k (3 r_i^2 + r_j^2) \ln(r_j + \|\mathbf{r}\|) \right) \end{aligned} \tag{4.37}$$

where

$$j = i + 1 \mod 3 \tag{4.38}$$

$$k = i + 2 \mod 3 \tag{4.39}$$

We show the potential of a cube ($\rho_0 = L = G = 1$) with a density gradient $\mathbf{g} = (1.2, 1.2, 0.)$ in the right panel of Figure 4.13. In comparison to the homogeneous cube the center of mass and the deepest point in the potential shift in the direction of the density gradient. At large radii the equipotential lines approach sphericity around the center of mass.

4.6.2 The Tree-PM force-split

In Gadget forces are calculated using a mixture of a short-range tree summation and a long-range force calculation on a periodic particle mesh (PM). Therefore the potential is split

$$\phi(\vec{r}) = \phi_s(\mathbf{r}) + \phi_l(\mathbf{r}) \tag{4.40}$$

where the long-range potential ϕ_l is given by the true potential convolved with a smoothing kernel

$$\phi_l = (\rho \circ G_\phi) \circ f \tag{4.41}$$

The long-range potential is calculated on a periodic particle mesh. The particles are binned with a clouds-in-cell assignment onto a periodic mesh to get the real-space density field

ρ. The mesh cannot resolve structures in the density field which have a smaller size than a mesh-cell. However, if the smoothing kernel f is large enough, e.g. the size of a few mesh-cells, the contributions of these small-scale structures to the long-range potential is negligible and the long-range force can be calculated very accurately on the mesh. It is then easy to obtain the long-range potential, simply by Fourier-transforming the density field, multiplying it by all the convolution components and then transforming the obtained potential back to real space.

$$\phi_{l,k} = \rho_k G_{\phi,k} f_k \tag{4.42}$$

where we denoted 3d-Fourier-transformed functions by a small index k.

The short-range part of the potential cannot be represented accurately on the mesh and is instead calculated by a tree-walk in real-space. It is given by

$$\begin{aligned} \phi_s &= (\rho \circ G_\phi) \circ (1 - f) && (4.43)\\ &= \rho \circ (G_\phi \circ (1 - f)) && (4.44)\\ &=: \rho \circ G_{\phi,s} && (4.45)\\ & && (4.46) \end{aligned}$$

where we defined the Green's function of the short range potential $G_{\phi,s}$. For the simple case of a point mass the short-range potential is given by $G_{\phi,s}$. However, in the case of a cubic mass-distribution it is much more complicated to obtain the short-range force, since the integrand in (4.34) must be changed. The default choice of a force-split kernel in Gadget 2 is a Gaussian kernel

$$f_{G2}(\mathbf{r}) = \frac{1}{8\pi^3 r_s^3} \exp\left(-\frac{\|\mathbf{r}\|^2}{4r_s^2}\right) \tag{4.47}$$

$$f_{k,G2}(\mathbf{k}) = \exp\left(-\|\mathbf{k}\|^2 r_s^2\right) \tag{4.48}$$

$$G_{s,G2}(\mathbf{r}) = -\frac{G}{\|\mathbf{r}\|}\mathrm{erfc}\left(\frac{\|\mathbf{r}\|}{2r_s}\right) \tag{4.49}$$

However we find that this is not a practical choice in our case, since we cannot find an analytical solution to the convolution of (4.49) with the density field of a cube. Instead we choose a different force-split kernel

$$f(\mathbf{r}) = \begin{cases} \frac{3(a-\|\mathbf{r}\|)}{a^4\pi} & \text{if } \|\mathbf{r}\| \le a \\ 0 & \text{else} \end{cases} \tag{4.50}$$

$$f_k(\mathbf{k}) = \frac{12}{a^4\|\mathbf{k}\|^4}\left(2 - 2\cos(a\|\mathbf{k}\|) + a\|\mathbf{k}\|\sin(a\|\mathbf{k}\|)\right) \tag{4.51}$$

$$G_s(\mathbf{r}) = \begin{cases} -G\frac{(a-\|\mathbf{r}\|)^3(a+\|\mathbf{r}\|)}{a^4\|\mathbf{r}\|} & \text{if } \|\mathbf{r}\| \le a \\ 0 & if\|\mathbf{r}\| > a \end{cases} \tag{4.52}$$

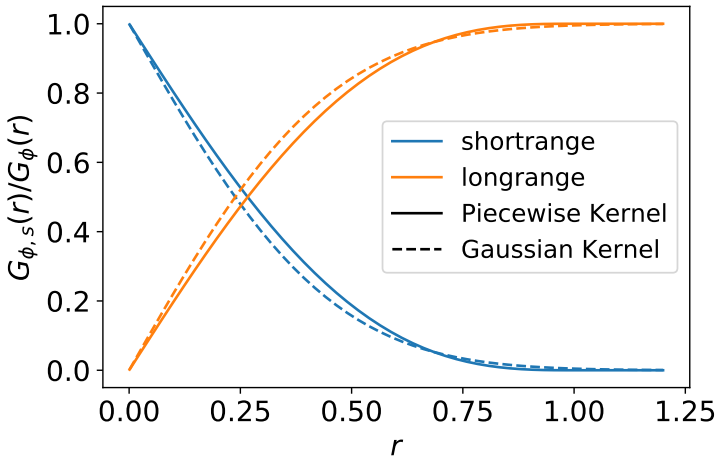

Figure 4.20: The short-range and long-range parts of the potential, defined as the fractional contribution to the potential at radius r for the two different cases of a piecewise force-split as in (4.50) with $a = 1$ and a Gaussian force-split as in (4.47) with $r_s = 0.25$. The two force-splits have a similar shape for $a \sim 4r_2$. However, the piecewise kernel has the advantage that the short-range potential becomes exactly zero at the finite radius a.

We plot the relative contributions of the long-range and short-range part of the potential in Figure 4.20. The Gaussian and the piecewise force-split have very similar shapes for $a \sim 4r_s$ and are both reasonable choices. However, the piecewise split has the advantage that the short-range contribution becomes exactly zero at the finite radius a. Therefore it is exactly correct to stop the summation of short-range forces beyond that radius, whereas the Gaussian short-range force never becomes exactly zero and can only be neglected approximately beyond some radius (typically $4.5 - 6r_s$).

To get the short-range potential of a cube with a gradient we have to solve the integral

$$\phi_s = \int_{-L/2-x_1}^{L/2-x_1} \int_{-L/2-x_2}^{L/2-x_2} \int_{-L/2-x_3}^{L/2-x_3} (\rho_0 - \mathbf{g}\cdot\mathbf{x} + \mathbf{g}\cdot\mathbf{r}) G_s(\vec{r}) d^3\mathbf{r} \tag{4.53}$$

Note that solving this integral in the general case is very complicated, because the interaction of the integral boundaries with the boundary of the kernel G_s introduces many different possible cases in the integral. However, in our simulations most of the interactions will be at short distance in comparison to the force-split scale $\|\mathbf{x}\| \ll a$ and $L \ll a$. Therefore we can calculate the analytical solution for these simpler cases, and use a numerical approximation for the other cases. For all cases where the farthest edge is still within the kernel radius a

$$\sum_i \left(|x_i| + \frac{L}{2}\right)^2 \leq a^2 \tag{4.54}$$

we can simplify

$$\phi_s = -G \int_{-L/2-x_1}^{L/2-x_1} \int_{-L/2-x_2}^{L/2-x_2} \int_{-L/2-x_3}^{L/2-x_3} (\rho_0 - \mathbf{g}\cdot\mathbf{x} + \mathbf{g}\cdot\mathbf{r}) \tag{4.55}$$

$$\frac{(a-\|\mathbf{r}\|)^3(a+\|\mathbf{r}\|)}{a^4\|\mathbf{r}\|} d^3\mathbf{r} \tag{4.56}$$

We define the 3d parent function of the short-range potential analogous to (4.36) and find

$$\begin{aligned} F_{\mathrm{hom},s}(\mathbf{r}) &= \int\int\int \frac{(a-\|\mathbf{r}\|)^3(a+\|\mathbf{r}\|)}{a^4|\|\mathbf{r}\||} dr_1 dr_2 dr_3 \\ &= \frac{r_1 r_2 r_3}{15}\left(30a^3 - 10a\|\mathbf{r}\|^2 + 2\|\mathbf{r}\|^3\right) \\ &\quad + \frac{1}{120a^4}\sum_i \left(r_i^2(60a^4 - 4r_i^4)\arctan\left(\frac{r_1 r_2 r_3}{\|\mathbf{r}\| r_i}\right)\right. \\ &\quad \left. + r_j r_k\left(9r_j^4 + 10r_j^2 r_k^2 + 9r_k^4 - 120r_k^4\right)\log(\|\mathbf{r}\| + z)\right) \end{aligned} \tag{4.57}$$

and

$$\begin{aligned} F_{\mathrm{lin},s,i}(\mathbf{r}) &= \int\int\int r_i \frac{(a-\|\mathbf{r}\|)^3(a+\|\mathbf{r}\|)}{a^4\|\mathbf{r}\|} dr_1 dr_2 dr_3 \\ &= \frac{1}{1680a^4}\left[yz\left(41r^5 + 52r^3 r_i^2 - 560ar^2 r_i^2 - 560a^4 r + 81r r_i^4 - 30r y^2 z^2 + 1680a^3 r_i^2\right.\right. \\ &\qquad \left. - 280a r_i^4\right) \\ &\qquad + \arctan\left(\frac{r_j r_k}{r r_i}\right)\left(560a^4 r_i^3 - 48 r_i^7\right) \\ &\quad + \log(r + r_j)\left(-840a^4 r_i^2 r_k + 105 r_i^6 r_k - 280a^4 r_k^3 + 105 r_i^4 r_k^3 + 63 r_i^2 r_k^5 + 15 r_k^7\right) \\ &\quad \left. + \log(r + r_k)\left(-840a^4 r_i^2 r_j + 105 r_i^6 r_j - 280a^4 r_j^3 + 105 r_i^4 r_j^3 + 63 r_i^2 r_j^5 + 15 r_j^7\right)\right] \end{aligned} \tag{4.58}$$

where the indices j and k change with the summation index i as in (4.38) and (4.39). In Figure 4.21 we show the short and- long-range potential and force of the same cube as the right panel of Figure 4.13 for a force-cut parameter of $a = 1.6$.

For cases which are not fully smoothed by the definition of 4.54 we approximate the potential numerically by a sum of point-masses:

$$\phi_s(\mathbf{x}) \approx \sum_{i=0}^{n-1}\sum_{j=0}^{n-1}\sum_{k=0}^{n-1} \rho(\mathbf{r}_{ijk})\Delta r^3 G_s(\mathbf{x} - \mathbf{r}_{ijk}) \tag{4.59}$$

$$\Delta r = L/n \tag{4.60}$$

$$\mathbf{r}_{ijk} = \begin{pmatrix} -\frac{L}{2} + (0.5+i)\Delta r \\ -\frac{L}{2} + (0.5+j)\Delta r \\ -\frac{L}{2} + (0.5+k)\Delta r \end{pmatrix} \tag{4.61}$$

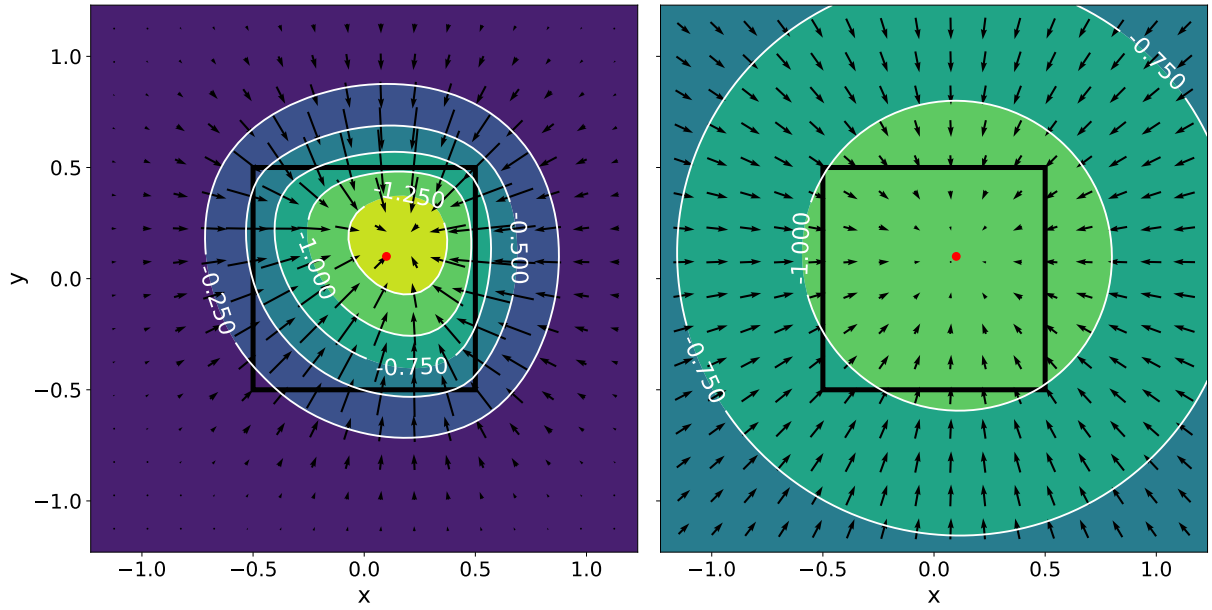

Figure 4.21: Potential (contours) and force field (arrows) of the same cube as in the right panel of Figure 4.13 with a piecewise force cut with $a = 1.6$. Left: The short-range component and right: the long-range component.

where n^3 is the number of point-masses per dimension. n controls the accuracy of the approximation and we adaptively choose n depending on the distance from the cube. We list the scenarios in which we use different point-mass approximations in the lower part of Table 4.2. In the left panel of Figure 4.22 we show the relative errors of the point-mass approximations

$$\epsilon = \frac{|\mathbf{F}_{s,\text{approx}} - \mathbf{F}_{s,\text{exact}}|}{|\mathbf{F}_{\text{tot,exact}}|} \tag{4.62}$$

where $\mathbf{F}_{s,\text{approx}}$ is the calculated force, $\mathbf{F}_{s,\text{exact}}$ is a much more accurate reference force (calculated by splitting the cube into 4^3 sub-cubes and calculating their forces) and $\mathbf{F}_{tot,\text{exact}}$ is the total exact force (without force-cut). Our choice of point mass approximations still maintains a relative accuracy better than 10^{-2} in the rare but challenging cases of $L \sim a$.

One special case which will almost never happen in a simulation, but which we still include for completeness, is if the whole force-split kernel resides within the cube. In that case the integrand of (4.53) is only non-zero on the spherical domain $\|\mathbf{x}\| \leq a$:

$$\phi_s = \int_0^a 4\pi r^2(\rho_0 - \mathbf{g}\cdot\mathbf{x})G_s(r)dr \tag{4.63}$$

$$= -\frac{4\pi a^2(\rho_0 - \mathbf{g}\cdot\mathbf{x})}{15} \tag{4.64}$$

where we dropped the term $\mathbf{g}\cdot\mathbf{r}$ from (4.53), because it must be zero since it is anti-symmetric with respect to $\vec{r}$.

Table 4.2: Full overview over different cases and the numerical approximations that are used. r_{max} is the distance to the farthest corner of the cube, r to its center and r_{min} to its closest boundary.

Case			Approximation
Fully Smoothed Cases			
$r_{\text{max}} \leq a$	and	$r \leq 1.5L$	analytic
$r_{\text{max}} \leq a$	and	$1.5 < r/L \leq \sqrt{5}$	multipole 4th order
$r_{\text{max}} \leq a$	and	$r/L > \sqrt{5}$	multipole 2nd order
Intersecting Cases			
$r_{\text{min}} > a$	and	$r_i < L/2\ \forall i$	analytic
$r_{\text{max}} > a$	and	$r/L < \sqrt{3}/2$	Split into 8 subcubes
$r_{\text{max}} > a$	and	$\sqrt{3}/2 \leq r/L < \sqrt{2}$	point masses $n = 4$
$r_{\text{max}} > a$	and	$\sqrt{2} \leq r/L < 2$	point masses $n = 3$
$r_{\text{max}} > a$	and	$2 \leq r/L < 4$	point masses $n = 2$
$r_{\text{max}} > a$	and	$r/L > \sqrt{16}$	point masses $n = 1$

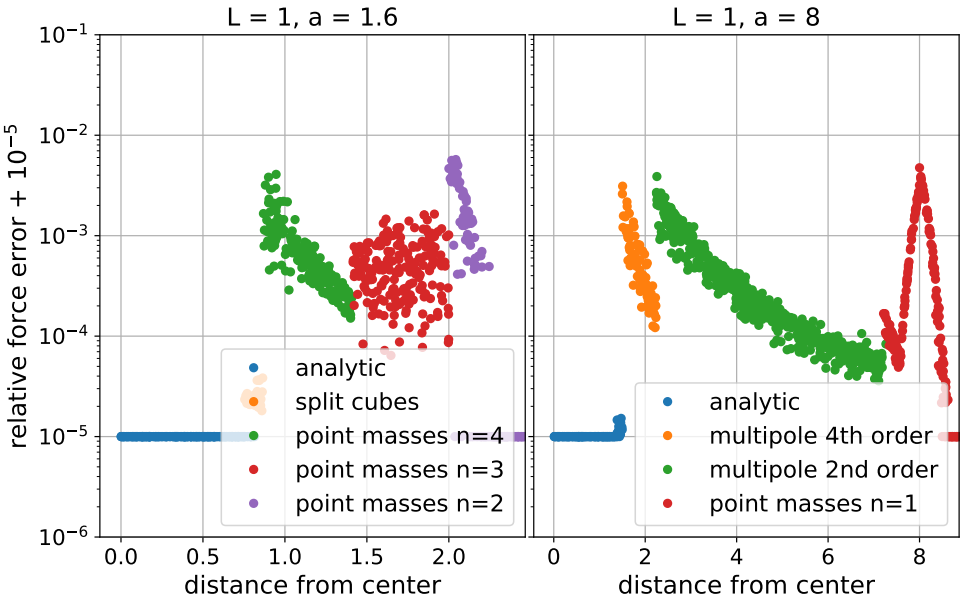

Figure 4.22: Relative force errors due to different numerical approximations as a function of distance from the center of a cube with $L = 1$, $G = 1$, $\mathbf{g} = (0.5, 0.3, 0.1)^T$. Errors are clipped at 10^{-5}. Left: with a force-split scale $a = 1.6$ which is similar to the size of the cube. This is a very unusual and challenging case, since our analytical solution is not valid in most of the range. It is however still well approximated (relative error smaller than 10^{-2}) with the point-mass approximations as explained in the text. Right: for a force-split scale $a = 8$ which is significantly larger than the cube. Not too close to the cube the forces can be very accurately described by multipole expansions which are numerically much cheaper than the analytic solution.

4.6.3 Multipole expansion of the far-field

Since the evaluation of the true parent function of the shortrange potential of the cube is very expensive, we use a multipole expansion to get much cheaper, but still accurate approximations for interactions at distances $\|\mathbf{x}\| > 1.5L$ from the cube. The multipole expansion is obtained from (4.53) by expanding the Green's function around a point $\mathbf{x}_0$:

$$\phi_s(\mathbf{x}) = \int\int\int \rho(\mathbf{r}) G_s(\mathbf{r}-\mathbf{x}) d^3\vec{r} \tag{4.65}$$

$$\approx G_s(\mathbf{x}-\mathbf{x_0}) \int\int\int \rho(\mathbf{r}) d^3\mathbf{r} \tag{4.66}$$

$$+ \sum_i \left(\partial_i G_s(\mathbf{x}-\mathbf{x_0}) \int\int\int \rho(\mathbf{r})(r_i - x_{0,i}) d^3\mathbf{r} \right) \tag{4.67}$$

$$+ \frac{1}{2} \sum_{i,j} \left(\partial_i\partial_j G_s(\mathbf{x}-\mathbf{x_0}) \int\int\int \rho(\mathbf{r})(r_i - x_{0,i})(r_j - x_{0,j}) d^3\mathbf{r} \right) \tag{4.68}$$

$$+ ... \tag{4.69}$$

In our case it is the most convenient to choose $\mathbf{x_0} = 0$ as the expansion point. Since the integration domain is symmetric around this point, many terms drop out. We find the expansion

$$\begin{aligned}
\phi_s(\vec{x}) \approx & -\frac{GL^3\rho_0(a-\|\mathbf{x}\|)^3(a+\|\mathbf{x}\|)}{a^4\|\mathbf{x}\|} \\
& -\frac{GL^5\left(a^4 - 4a\|\mathbf{x}\|^3 + 3\|\mathbf{x}\|^4\right)\sum_i g_i x_i}{12a^4\|\mathbf{x}\|^3} \\
& +\frac{GL^5\rho_0(\|\mathbf{x}\|-a)}{2a^4} \\
& -\frac{GL^7\left(\left(\|\mathbf{x}\|^4 - 5a^4\right)\sum_i g_i x_i^3 + \left(3a^4\|\mathbf{x}\|^2 + 7\|\mathbf{x}\|^6\right)\sum_i g_i x_i\right)}{240a^4\|\mathbf{x}\|^7} \\
& -\frac{GL^7\rho_0\left(\left(3\|\mathbf{x}\|^4 - 35a^4\right)\sum_i x_i^4 + 21a^4\|\mathbf{x}\|^4 - 17r^8\right)}{960a^4r^9}
\end{aligned} \tag{4.70}$$

where we have written one expansion order per line (from 0 to 4). Note that the dipole moment (order 1) does not vanish, since we expanded around the geometric center of the cube, not its center of mass. However, the other terms are simpler than in an expansion around the center of mass. In the case of a homogeneous cube $\mathbf{g} = 0$ in the absence of a force-cut $a \to \infty$ this simplifies to

$$\phi_{\text{hom}}(\mathbf{x}) \approx -\frac{G\rho_0 L^3}{|\mathbf{x}|}\left(1 + \frac{7L^4}{320|\mathbf{x}|^4} - \frac{7L^4\sum x_i^4}{192|\mathbf{x}|^8}\right) \tag{4.71}$$

which is in agreement with the multipole expansion in Hummer (1996). We use the multipole expansion up to fourth order if $1.5L < r < 2L$ and only up to second order if $r > 2L$.

We summarize these different cases together with the ones where the point mass approximations have to be used in Table 4.2. In the right panel of Figure 4.22 we show the errors of the multi pole expansion for a typical case $a = 8$, $L = 1$. In the ranges where we use the multipole expansions they have a relative accuracy better than 10^{-2}.

Chapter 5

The phase space structure of a warm dark matter halo

In this chapter we test the "sheet + release" scheme that we have developed in the last chapter in a zoomed simulation of one of the smallest haloes that would form in a warm dark matter universe. Such a small halo is a rather simple system. In a zoom simulation we can achieve quite high resolution in it and can therefore test the code in a relatively controlled system. Since we still need to fix the problem with the geodesic deviation equation integration that we have described in section 4.4.5, the contents in this chapter are preliminary and much shorter than they would be in a future publication. We keep the descriptions quite short and just give a basic overview of first results. A future publication about this topic will involve me, Raul Angulo, Oliver Hahn and Simon White as authors.

5.1 Halo selection

We use the MUSIC code (Hahn & Abel, 2011) to set up initial conditions for a "parent simulation" with $m_x = 3.5\text{keV}$, $N = 512^3$, $L = 20\,\text{Mpc}/h$. We run this with the simplest mode of the HA16 scheme (no refinement, no release, no tree) to get the evolved non-linear density field. This simulation is lacking small scale forces, but we only use it for selecting a halo for the zoomed simulation (which it is good enough for). In the left panel of Figure 5.1 we show a slice through the density field of this parent simulation and in the right panel we show a zoom onto the halo that we have selected for the zoom-in simulation. The halo was selected to be one of the lower mass objects that form in this cosmology. Further, for this zoom simulation we used a region that is expanding in all three dimensions. That makes it possible to use a relatively small extent for the high-resolution region, since the zoom region has no larger scale shell crossings. Therefore the selected halo is also a quite special one and not necessarily a "typical halo" (for whatever that means). However, for our first test simulations, and as a numerical convergence study this can be justified. It will be interesting to make a larger set of zoom-in simulations with different kinds of small halos, once we have solved the remaining numerical issues.

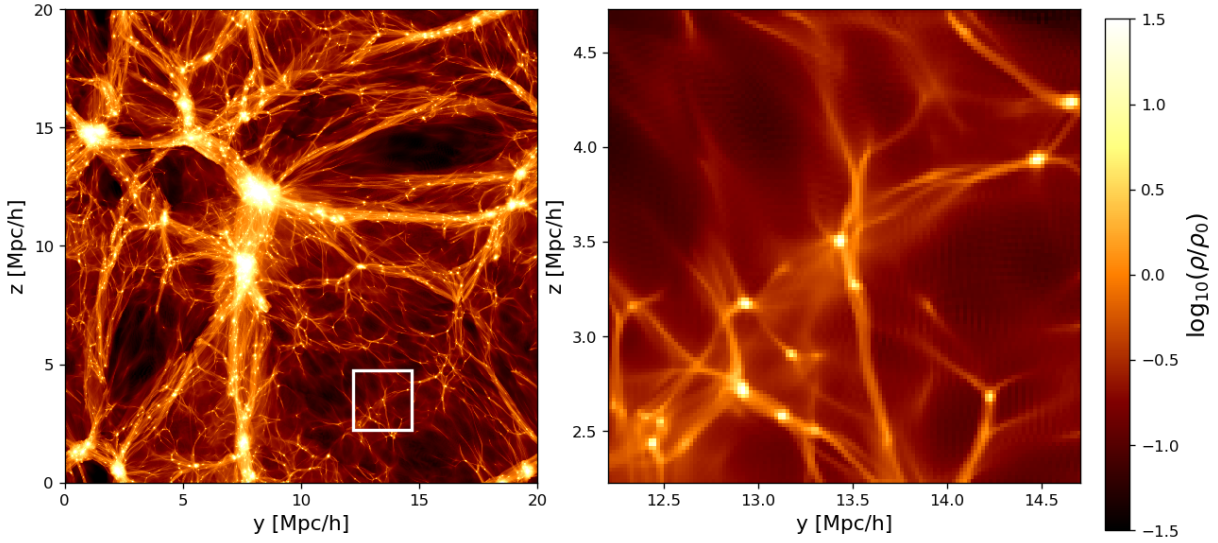

Figure 5.1: A slice through the parent simulation with a width of $\Delta x = 2.5\,\mathrm{Mpc}/h$. Left: the whole box and right: a zoom on the region around the halo that we have selected ($y = 13.5, z = 3.5$).

We trace particles of the selected halo back into the initial conditions and determine their Lagrangian center. Then we use MUSIC to set up a zoomed initial conditions that are focused on this center. The initial conditions are set up so that the low resolution region is completely represented by N-body particles, and only the high-resolution region uses the continuous Lagrangian sheet-elements. (I modified the MUSIC code to make this zoom set-up possible.) We show a picture of the zoomed initial conditions in the left panel of Figure 5.2.

We evolve these initial conditions with the sheet + release scheme and the cube-tree as described in the previous chapter. At the release each Lagrangian 3x3x3 element creates 4^3 N-body particles. Thereby for a flow tracer resolution of $N = 256^3$ the upper limit to the possible number of N-body particles is 512^3 (though the actual number is much lower). The refinement level of the cube-tree is defined on every second flow tracer per dimension, leading to $N_c = 128^3$ as we would have stated it in the previous chapter. However, we note that in the innermost regions of halos the cube-tree is almost always refined to the lowest level that is allowed by the force-resolution limit ϵ (which is the softening for the N-body simulations and the minimal allowed node length for the cube-tree cases).

On the right side of Figure 5.2 we show a slice through the evolved density field of the same Eulerian space region as in the left panel. As already mentioned, this region expands in all three dimensons. Figure 5.2 shows one of our first simulations which uses a Lagrangian size of $1.25\,\mathrm{Mpc}/h$, but the simulations that we present later even use a region with the half size of $0.625\,\mathrm{Mpc}/h$. We made sure that no low resolution particles get close into the virial radius of the halo. The density field in Figure 5.2 looks a bit different than the one in the right panel of Figure 5.1. That is mostly because the slice in Figure 5.2 has

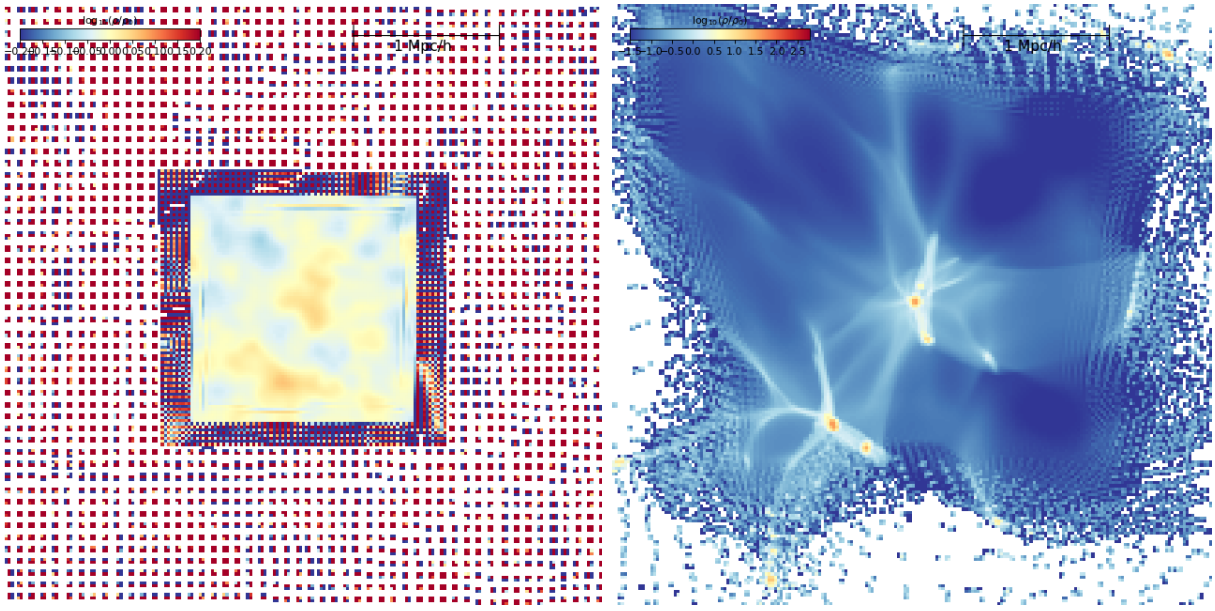

Figure 5.2: Zoomed initial conditions (left) and the same Eulerian space region in the evolved simulation (right). The width of the slice is $0.625\,\mathrm{Mpc}/h$ and the extend of the zoom region is $1.25\,\mathrm{Mpc}/h$ in Lagrangian space.

a much smaller width of $0.625\,\mathrm{Mpc}/h$ - thereby not containing many of the structures that can be seen in Figure 5.1.

Although we planned originally to simulate one of the smallest haloes that could potentially form in this WDM cosmology, the selected halo turned out to be a bit more massive than that. I.e. its virial mass is $M_{200} = 4.8 \cdot 10^9 M_\odot/h$ and its virial radius $r_{200} = 95\,\mathrm{kpc}/h$ - thereby the halo is only in the slightly suppressed regime of the halo mass function (the half mode mass of the 3.5keV thermal relic is $0.9 \cdot 10^9 M_\odot$ as calculated from the formula in Schneider et al. (2012)). However, we can still use this object for numerical convergence studies, and some qualitative remarks. We will later try to simulate other (potentially smaller) objects.

In Figure 5.3 we show a slice through the density field in and around the halo for N-body and sheet + release simulations. The sheet + release scheme does not fragment – even when using the cube-tree with a maximum force resolution of $\epsilon = 0.08\,\mathrm{kpc}/h$.

5.2 Density structure

We present the radial density structure of the halo at the final time $a = 1$ in Figure 5.4. This Figure contains different (equivalent) ways of showing the radial density structure: The top panel shows the density profile, the bottom left panel shows the mass $M(< r)$ that is contained within some radius r and the bottom right shows the circular rotational

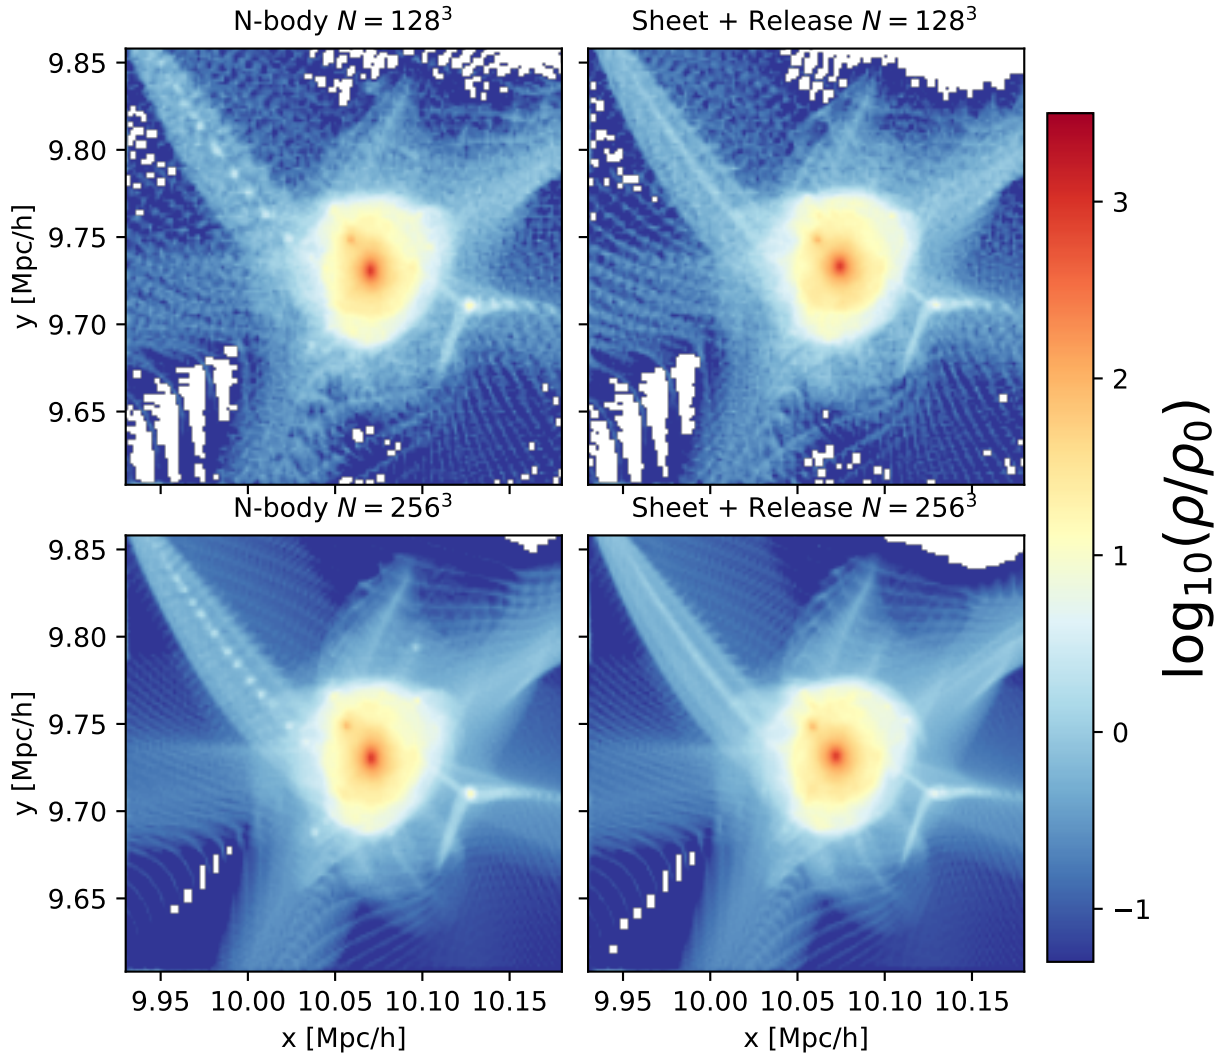

Figure 5.3: Density field in a slice with a thickness of $0.125\,\mathrm{Mpc}/h$ in and around the zoomed halo. Top: with $N = 128^3$ particles and bottom with $N = 256^3$, Left for an N-body simulation and right for the sheet + release scheme. The sheet + release scheme does not show artificial fragmentation - even though we use the very high force-resolution of $\epsilon = 0.08\,\mathrm{kpc}/h$ here. Note that for both schemes we plot the "N-body" density-estimate here (=cic binned particles) so that the two schemes can be compared fairly. However, the sheet + release scheme actually uses a density field with much higher quality in the simulations.

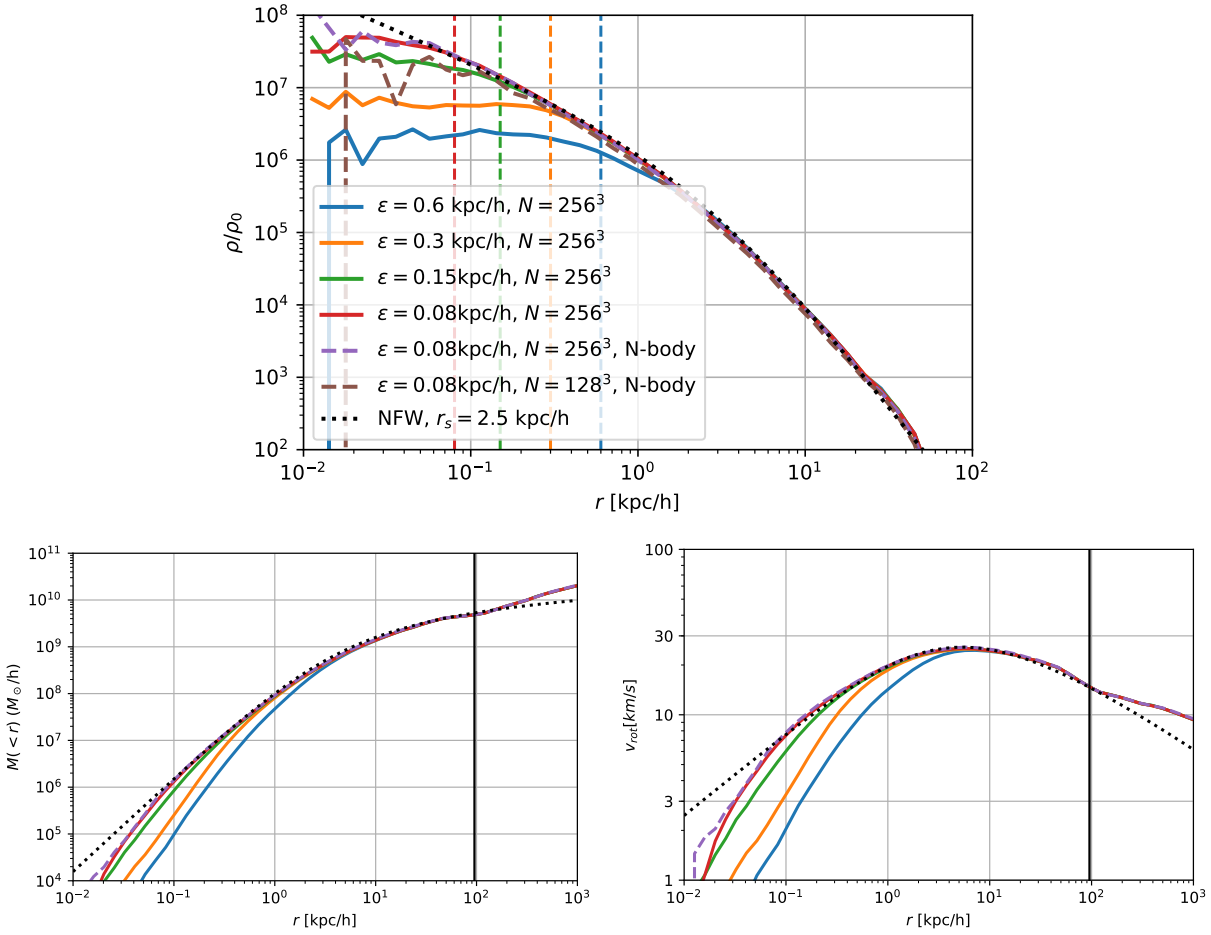

Figure 5.4: Density profile (top), enclosed mass (bottom left) and rotational velocity (bottom right) for the simulated warm dark matter halo. Vertical dashed lines in the upper plot mark the softening scales of the different simulations. The sheet + release and the N-body simulation produce consistent results and both agree with an NFW profile (in regions which are converged with the force resolution scale ϵ). It seems that the fragmentation that can be seen in Figure 5.3 has no significant impact on the structure of the halo.

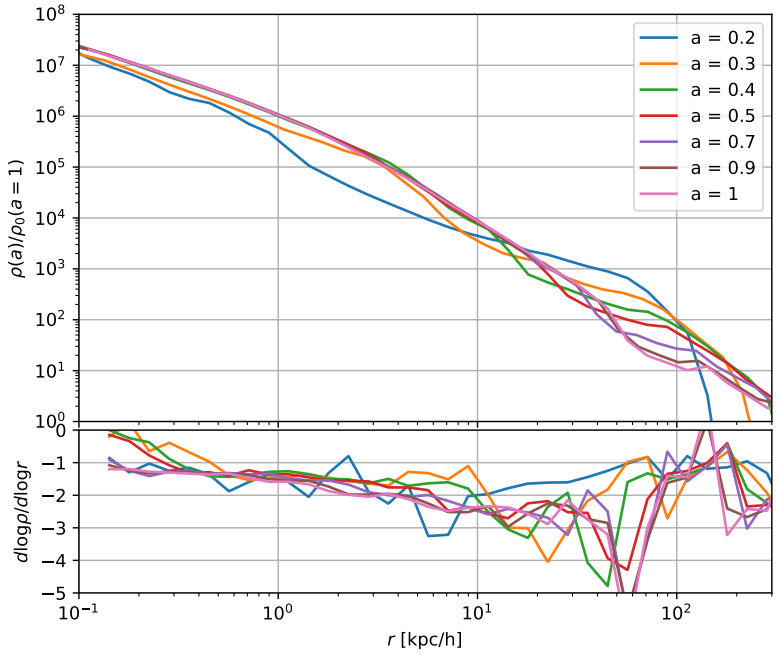

Figure 5.5: Profile evolution for the sheet + release simulation with $N = 256^3$ and $\epsilon = 0.08\,\mathrm{kpc}/h$. Plotted are physical density and physical radius. The profile mostly stops evolving after $a = 0.4$. Bottom: the slope of the profile. At $a = 1$ it approaches a central slope close to -1.

velocity v_{rot}:

$$M(r) = \int_0^r 4\pi r'^2 \rho(r')\,\mathrm{d}r' \tag{5.1}$$

$$v_{\mathrm{rot}} = \sqrt{\frac{GM(r)}{r}} \tag{5.2}$$

Further, each panel contains as dotted line an NFW profile (see equation (2.64)) that we have fitted by eye ($r_s = 2.5\,\mathrm{kpc}/h$, $\rho_c = 9 \cdot 10^5 \rho_0$).

The density profiles converge very well with the force resolution ϵ. We have also made sure everything is converged with time-stepping and other parameters of the schemes. The profiles of the sheet + release agree very well with the N-body simulations and both are consistent with an NFW-profile. It seems that the artificial fragmentation from Figure 5.3 does not alter the halo structure (due to, for example, the accretion of fragments). Instead the halo structure seems to be consistent between the N-body simulations and the sheet + release simulations. However, we should test this in the future in other setups, where the fragments are more dominant. In this case the fragments are relatively well behaved and cause hardly any significant additional substructure in the halo.

In Figure 5.5 we show the evolution of the density profile over time (for the sheet + release case with $\epsilon = 0.08\,\mathrm{kpc}/h$). The densities and radii are given in physical coordinates.

The density profile approaches a stable state around $a = 0.4$.

5.3 Phase-space structure

In Figure 5.6 we show the radial profiles of the median caustic count and the median stream density (determined in spherical shells around the halo). The stream densities seem to converge very well, but as discussed in section 4.4.5, this does not mean that they are correct. We therefore skip discussing the stream-density profile here and leave this discussion for a point in the future where we have resolved the issue that was described in section 4.4.5. The caustic count is a reliable quantity, however. Typical caustic counts reach order 400 in the center of the halo, therefore particles have gone through order 10^2 orbits.

In Figure 5.7 we show a 2d histogram of the stream-densities versus the caustic count. Interestingly at high cautic counts there seems to be a power-law with a very steep slope $\rho_s \propto n_{\rm caco}^{-15}$. So, if the caustic count grows roughly linearly with time, that means that the stream densities decrease roughly like the 15th power of the time. This is not too far away from the result of Sousbie & Colombi (2016) where the number of required simplices (to resolve the sheet accurately) scaled like the 12th power of time. However, as already mentioned, our results on the stream-densities are not reliable – at least for $n_{\rm caco} \gtrsim 30$. Stream-densities might be correct though for particles which have not gone through too many caustics yet $n_{\rm caco} \lesssim 30$, so the left half of the histogram might be correct. We note that there is an interesting trend for stream-densities to increase with the caustic count for low caustic counts $n_{\rm caco} \lesssim 5 - 10$ whereas at higher caustic counts the typical stream-densities only decrease. Possibly this is roughly number of caustics crossings that is needed until a typical particle has reached ergodicity – before this point the local volume element can shrink due to gravitational collapse, but as soon as the particle is on some typical orbit around the halo the volume element cannot shrink anymore (on average).

There are a lot of interesting phase space quantities that one could look at. These include beneath the caustic counts and stream densities also the evolution of the singular values of the distortion tensor (and also their associated directions). Probably there is always one singular value that shows a very strong stretching (because its Lagrangian axes is associated with displacements in energy levels) whereas the two other singular values are only mildly stretched (because they are associated with angular displacements). However, it does not make sense to investigate this in more detail while we cannot ensure that the results are reliable. Therefore we leave a more detailed discussion of the phase space structure for the time after we resolved the issue from section 4.4.5. Therefore we also skip a comparison with the results on the stream-densities of Vogelsberger & White (2011) for now.

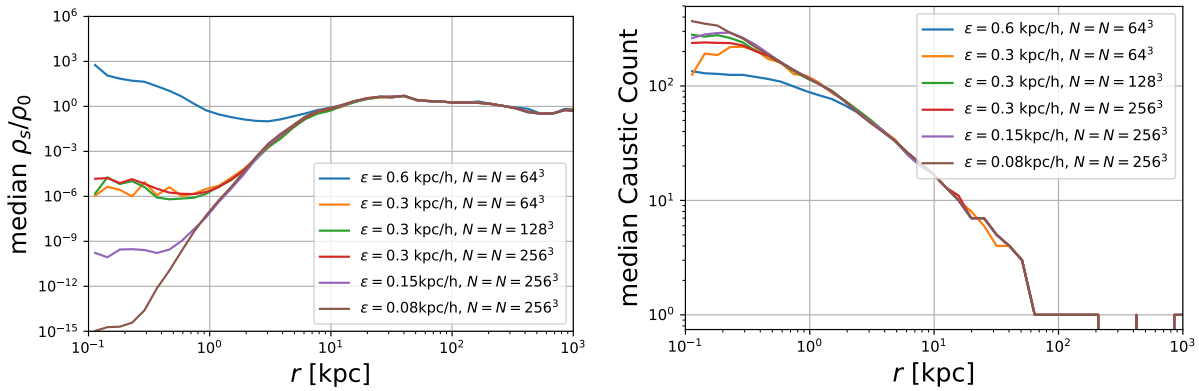

Figure 5.6: Median stream-density and caustic count profiles. The stream-density profiles converge well, but we think they are not reliable yet, as pointed out in section 4.4.5. Right: the caustic-count profile. In the center of the halo typical particles have gone through hundreds of caustics.

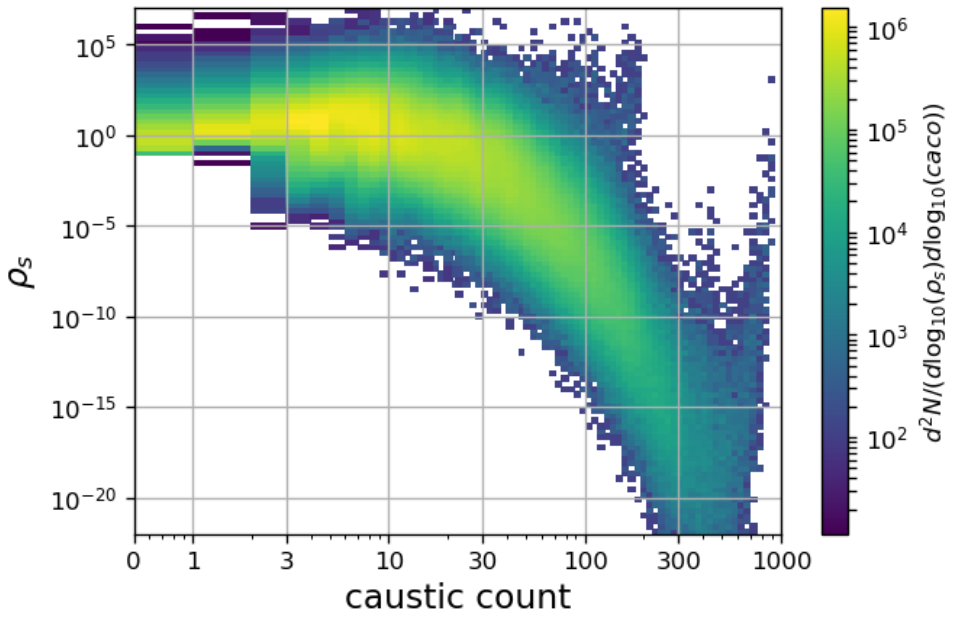

Figure 5.7: 2d-histogram of stream densities versus caustic count for the sheet + release simulation with $N = 256^3$ and $\epsilon = 0.08\,\mathrm{kpc}/h$. It seems that at high caustic counts there is a power law with a form $\rho_s \propto n_{\mathrm{caco}}^{-15}$. However, this result is not reliable yet, as discussed in section 4.4.5. Stream densities with lower caustic counts $N_{\mathrm{cac}} \lesssim 30$ are more likely to be correct. We selected for this histogram all particles within the virial radius of the halo.

5.4 Summary and outlook

We have presented some preliminary results on the structure and phase space structure of a small warm dark matter halo. In the future we will try to improve the reliability of the stream-densities. Then we can have a detailed discussion of the phase space structure of the simulated halo. Further, we will consider to make more simulations of different halos, also some which are in the even stronger suppressed regime of the halo mass function.

The N-body simulations of the warm dark matter halo produce artificial fragments whereas the sheet + release simulations do not (even when a very high force resolution is used). However, it seems that measurements of the halo structure lead to similar results in both schemes which are consistent with an NFW-profile. This is so, although the sheet + release scheme uses very different numerical techniques: The discrete elements that are used for solving Poisson's equation are completely different (cubes instead of particles), low-density regions are traced by the sheet-scheme (instead of an N-body scheme) and only inside halos the scheme reverts to an N-body approach. That this (almost) independent method reaches similar results to the N-body scheme is good news for the reliability of N-body simulations in warm dark matter cosmologies. However, we want to emphasize here, that we have only tested this for the density profiles of a halo – these involve a lot of coarse-graining. Other quantities which involve less coarse-graining might still be biased and have to be checked individually. We leave more detailed studies for the future.

Chapter 6

On the reconstruction of warm and hot phase space distributions from Lagrangian submanifolds

In the previous chapters we have always discussed the dark matter sheet in the limit of a vanishing primordial velocity dispersion $\sigma_v \to 0$. However, it is also possible to reconstruct a warm phase space distribution from a cold sheet if the velocity dispersion is small enough. This opens up possibilities to simulate six dimensional warm distributions while only actively tracing a single three dimensional sub-manifold in phase space. Further, hot distributions can potentially be simulated by tracing multiple sheets and treating the phase space distribution locally like a warm one. This might open up new simulation techniques for warm dark matter models and the cosmic neutrino fluid. In this chapter we will describe the necessary mathematical steps and discuss possible implementations.

This chapter is based on several ideas that I had at various points of my PhD, but I have never really put them together to a consistent picture. When writing down this thesis, I thought it would be nice to include some of these ideas. Therefore I started putting them together in the last days before handing in this thesis. I did not have the time to do a detailed literature research and therefore possible references could be missing. I have only cited publications here that I knew about when writing this text. Further, the ideas in this chapter are somewhat experimental and have not been tested under a multitude of different conditions, yet. Therefore, please read this chapter as a draft of possibilities, rather than a detailed scientific study.

6.1 The relation between initial and final phase space

By tracing the cold sheet $(\mathbf{x}(\mathbf{q}), \mathbf{v}(\mathbf{q}))^T$ in cosmological simulations, we are able to answer any questions that relate to displacements in this Lagrangian subspace. For example, we can estimate how a displacement in the Lagrangian spatial coordinate $\Delta\mathbf{q}$ relates to phase

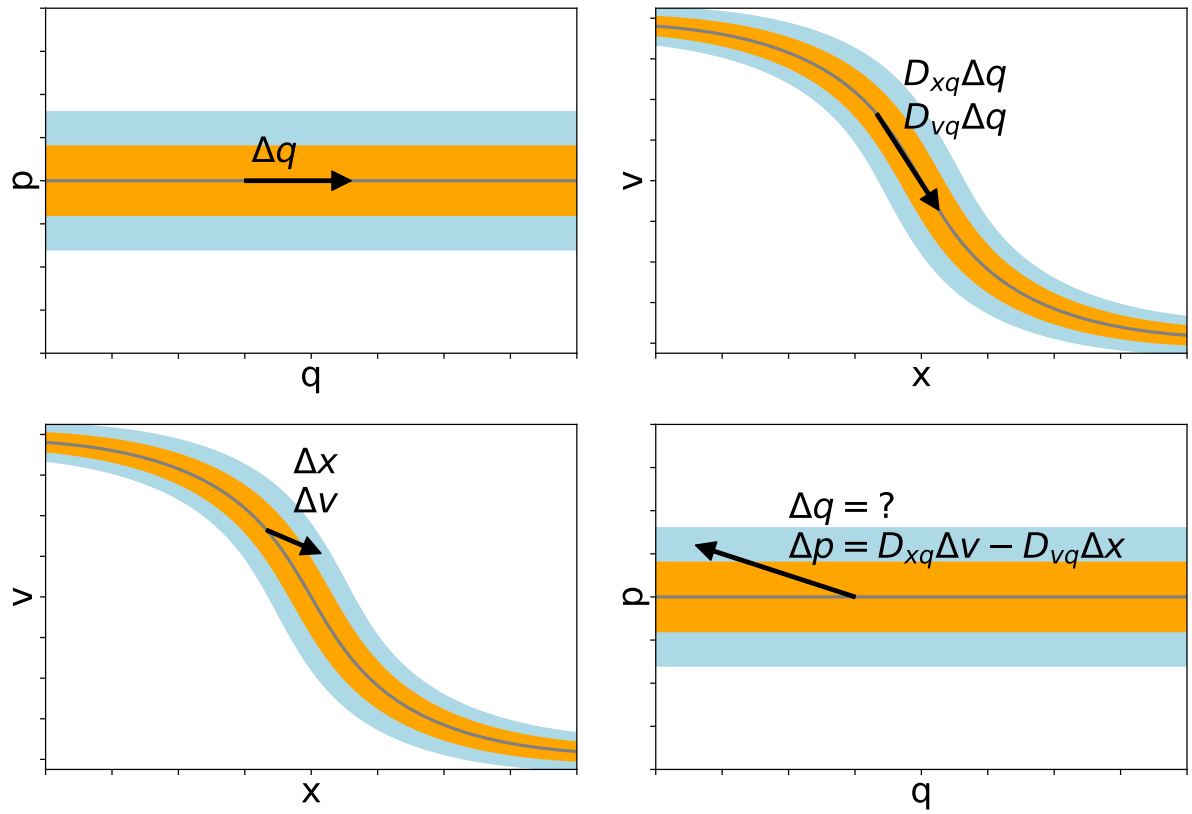

Figure 6.1: Illustration of the phase space information that we have when tracing a cold sheet in phase space. Top: We have information about how an initial displacement in the spatial Lagrangian coordinate $\Delta\mathbf{q}$ (left) evolves into a phase space displacement at later time (right) as given by the distortion tensor. However, due to the symplectic properties this also contains some vital information about the inverse map. Bottom: A phase space displacement at the final time (left) is mapped into the initial Lagrangian space (right). The available information allows to reconstruct the displacement in the initial velocity coordinate $\Delta\mathbf{p}$ whereas the information about the initial spatial displacement $\Delta\mathbf{q}$ is lost.

space displacements at a later time:

$$\begin{pmatrix}\Delta\mathbf{x}\\ \Delta\mathbf{v}\end{pmatrix} = \begin{pmatrix}\mathbf{D}_{\mathrm{xq}}\Delta\mathbf{q}\\ \mathbf{D}_{\mathrm{vq}}\Delta\mathbf{q}\end{pmatrix} \tag{6.1}$$

this is illustrated in the top panels of Figure 6.1.

So far, nothing new. However, it is also possible to ask the reverse question: If we think think of a small phase space displacement $(\Delta\mathbf{x}, \Delta\mathbf{v})^T$ in the final phase space, how does it map to a displacement in the initial phase space. This is illustrated in the bottom panel of 6.1. (We will explain later why this question is important.) To answer this question one can define the inverse distortion tensor

$$\mathbf{D^{-1}} = \begin{pmatrix}\frac{\partial\mathbf{q}}{\partial\mathbf{x}} & \frac{\partial\mathbf{q}}{\partial\mathbf{v}}\\ \frac{\partial\mathbf{p}}{\partial\mathbf{x}} & \frac{\partial\mathbf{p}}{\partial\mathbf{v}}\end{pmatrix} \tag{6.2}$$

$$=: \begin{pmatrix}\mathbf{D}_{\mathrm{qx}} & \mathbf{D}_{\mathrm{qv}}\\ \mathbf{D}_{\mathrm{px}} & \mathbf{D}_{\mathrm{pv}}\end{pmatrix} \tag{6.3}$$

where in the second line we named its sub tensors. Note that these are defined over partial derivatives in a way that for example $\mathbf{D}_{\mathrm{qx}}$ is not simply the inverse of $\mathbf{D}_{\mathrm{xq}}$. Actually one can calculate the components of $\mathbf{D}^{-1}$ by inverting the distortion tensor $\mathbf{D}$ (while using the symplectic constraints) and find the relations

$$\mathbf{D}_{\mathrm{qx}} = \mathbf{D}_{\mathrm{vp}} \qquad \mathbf{D}_{\mathrm{qv}} = -\mathbf{D}_{\mathrm{xp}} \tag{6.4}$$

$$\mathbf{D}_{\mathrm{px}} = -\mathbf{D}_{\mathrm{vq}} \qquad \mathbf{D}_{\mathrm{pv}} = \mathbf{D}_{\mathrm{xq}} \tag{6.5}$$

(compare for example White & Vogelsberger (2009)). Now if we map our late time phase space displacement $(\Delta\mathbf{x}, \Delta\mathbf{v})^T$ into the initial space, we find

$$\begin{pmatrix}\Delta\mathbf{q}\\ \Delta\mathbf{p}\end{pmatrix} = \begin{pmatrix}\mathbf{D}_{\mathrm{qx}}\Delta\mathbf{x} + \mathbf{D}_{\mathrm{qv}}\Delta\mathbf{v}\\ \mathbf{D}_{\mathrm{px}}\Delta\mathbf{x} + \mathbf{D}_{\mathrm{pv}}\Delta\mathbf{v}\end{pmatrix} \tag{6.6}$$

$$= \begin{pmatrix}\mathbf{D}_{\mathrm{vp}}\Delta\mathbf{x} - \mathbf{D}_{\mathrm{xp}}\Delta\mathbf{v}\\ -\mathbf{D}_{\mathrm{vq}}\Delta\mathbf{x} + \mathbf{D}_{\mathrm{xq}}\Delta\mathbf{v}\end{pmatrix} \tag{6.7}$$

$$= \begin{pmatrix}?\\ -\mathbf{D}_{\mathrm{vq}}\Delta\mathbf{x} + \mathbf{D}_{\mathrm{xq}}\Delta\mathbf{v}\end{pmatrix} \tag{6.8}$$

where we have emphasized in the last line that the $\Delta\mathbf{q}$ component cannot be calculated from our knowledge of the distortions of the sheet[1], but the $\Delta\mathbf{p}$ component can be. That means if we trace a cold sheet in phase space (instead of the full 6d distribution), we loose locally information about a persistent drift that happens due to displacements in the velocity variables. However, we still maintain some knowledge about the phase space extent of the sheet. We illustrate the inverse mapping in the bottom panels of Figure 6.1.

[1] Of course one could reconstruct it when the GDE is used, but we are assuming only a cold sheet here, as could be reconstructed from the positions and velocities of particles in typical cosmological simulations.

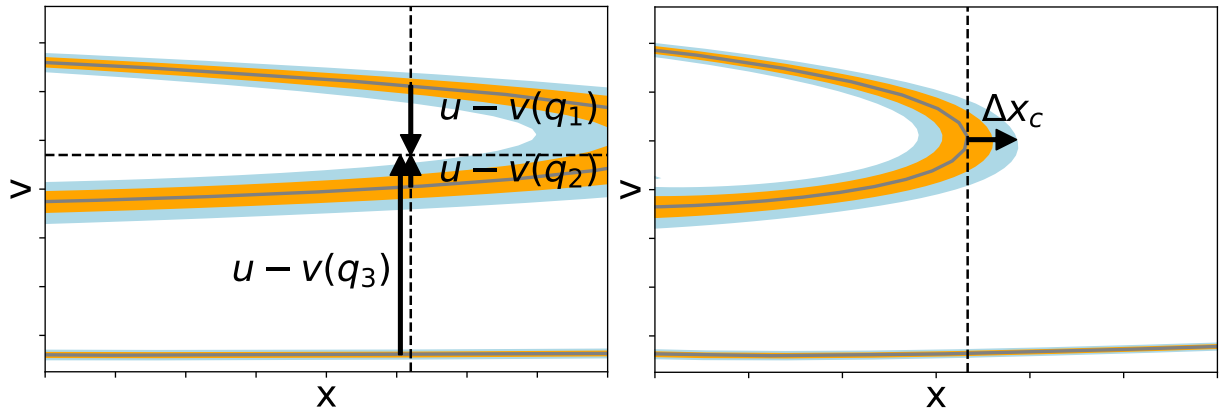

Figure 6.2: Left: illustration of the calculation of the velocity distribution function, the dashed lines delineate the coordinate x at which the distribution is measured and the velocity u at which it is evaluated. In this example three streams contribute, but at the selected velocity u only stream 2 will have a significant contribution. Right: Calculation of the width of a caustic. The dashed line shows the position of the caustic and Δx shows its extend which is inferred by mapping its orthogonal vector into the initial space where we can match it to width of the primordial velocity distribution.

While we have no information about the phase space drift when only considering a cold sheet, this information is actually not all that relevant. That is so, because the (hypothetical) initial space (for $a \to 0$) is homogeneous and its phase space distribution is given as a function that only depends on the velocity, i.e. $f_0(\mathbf{q}, \mathbf{p}) = f_0(\mathbf{p})$. We can still reconstruct the full phase space distribution without knowledge of $\Delta\mathbf{q}$. The value of the final phase space distribution f at a location that is close to a particle $(\mathbf{x}, \mathbf{v})^T$ on the central "cold" sheet (defined by $\mathbf{p} = 0$) is given by

$$\begin{aligned} f(\mathbf{x} + \Delta\mathbf{x}, \mathbf{v} + \Delta\mathbf{v}) &= f_0(\mathbf{q} + \Delta\mathbf{q}, \mathbf{p} + \Delta\mathbf{p}) \\ &= f_0(\Delta\mathbf{p}) \\ &= f_0(-\mathbf{D}_{\mathrm{vq}}\Delta\mathbf{x} + \mathbf{D}_{\mathrm{xq}}\Delta\mathbf{v}) \end{aligned} \tag{6.9}$$

Note that most of this was already realized in White & Vogelsberger (2009) where the implications for stream densities and the self-annihilation signal were discussed.

If we have a warm (but not hot) phase space distribution, we can approximate that the phase space distribution is zero for points far away from the sheet. For points close to the sheet we can estimate it from (6.9). Note that "close to the sheet" and "far away from the sheet" are actually concepts that are relatively hard to define. We will discuss this in more detail later in section 6.2. However, let us first consider two simpler applications of (6.9).

6.1.1 The velocity distribution

The velocity distribution function $f_{\mathbf{x}}(\mathbf{u})$ at a location $\mathbf{x}$ can be approximated (very well) by

$$f_{\mathbf{x}}(\mathbf{u}) = \sum_{q \in \text{ streams}(\mathbf{x})} f_0(\mathbf{D}_{\text{Xq}}(\mathbf{q}) \cdot (\mathbf{u} - \mathbf{v}(\mathbf{q})) \tag{6.10}$$

where we left the normalization arbitrary and labeled the initial velocity distribution function f_0. Note that we have used equation (6.9) with $\Delta x = 0$, because we are interested in the phase space distribution at a constant $\mathbf{x}$ as illustrated in the left panel of Figure 6.2. We have written this as a sum over all streams here instead of only selecting the closest stream to $(\mathbf{x}, \mathbf{u})^T$. This is, however, practically equivalent since the initial velocity distribution f_0 goes to zero very quickly and effectively only the closest stream contributes. For an example of the velocity distributions please refer back to Figure 2.21 where we have shown velocity distributions that were calculated in the way that is presented here.

6.1.2 The widths of caustics

Caustics are at locations where the spatial distortion tensor has a vanishing determinant $\det \mathbf{D}_{\text{Xq}} = 0$. The direction orthogonal to the caustic can be determined by making the singular value decomposition of (the degenerate) $\mathbf{D}_{\text{Xq}}$ and selecting the vectors $\mathbf{u}_c$ and $\mathbf{v}_c$ that are associated with the vanishing singular value $s_c = 0$ ($\mathbf{u}$ and $\mathbf{v}$ are as defined in section 2.6.4). $\mathbf{v}_c$ is then the Lagrangian direction associated with the caustic $\mathbf{D}_{\text{Xq}}\mathbf{v}_c = 0$ and $\mathbf{u}_c$ is the Eulerian axis perpendicular to the caustic. For clarity please refer back to section 2.6.4 where we explain the singular value decomposition in more detail. To not cause confusion with velocity coordinates we call $\mathbf{u}_c = \mathbf{e}_c$ in the following.

A Eulerian space vector $\mathbf{e}_c$ with length Δx (associated with the caustic) maps into the initial space as

$$\Delta \mathbf{p} = -\mathbf{D}_{\text{Vq}} \cdot (\Delta x \mathbf{e}_c) \tag{6.11}$$

where we have neglected the Lagrangian displacement Δq and used a zero velocity displacement $\Delta \mathbf{v} = 0$. This is illustrated in the right panel of Figure 6.2. As a reverse conclusion we can calculate the effective width of the caustic as

$$\Delta x_c = \frac{\sigma_v}{|\mathbf{D}_{\text{Vq}}\mathbf{e}_c|} \tag{6.12}$$

where σ_v is the typical amplitude of primordial thermal velocities. We have already shown how the widths of caustics behave in section 2.6.1 where we have used exactly these calculations.

6.2 Reconstructing sheets that are displaced in velocity space

While we have shown in the last two sections (6.1.1 and 6.1.2) reconstructions that are of pure velocity or pure spatial nature (i.e. either $\Delta x = 0$ or $\Delta v = 0$) the situation becomes less clear if one wants make reconstructions which involve displacements in position and velocity coordinates at the same time.

The question we try to answer in this section is: given a cold sheet e.g. $\mathbf{p} = 0$, how can we reconstruct a sheet that is slightly displaced in velocity space $\mathbf{p} = \mathbf{p}^*$? We will denote its position and velocity coordinates as $\mathbf{x}^*(\mathbf{q})$ and $\mathbf{v}^*(\mathbf{q})$. Let us first discuss this question in one dimension and then consider later how to generalize to higher dimensions.

If we take a phase space vector with an arbitrary direction in the final phase space $\alpha(\Delta x, \Delta v)^T$, we can generally find a value α so that this vector lies on the $p = p^*$ sheet:

$$p^* \stackrel{!}{=} \Delta p = \alpha(-D_{\mathrm{vq}}\Delta x + D_{\mathrm{xq}}\Delta v) \tag{6.13}$$

$$\alpha = \frac{p^*}{D_{\mathrm{xq}}\Delta v - D_{\mathrm{vq}}\Delta x} \tag{6.14}$$

that means we can choose the direction of this phase space vector relatively arbitrarily, but its length has to be re-normalized in such a way that it "hits" the p^*-sheet. We illustrate this in the left panel of Figure 6.3.

In principle we can reconstruct the p^*-sheet by choosing at each location $(x(q), v(q))$ in the central sheet a re-normalized displacement $\alpha(\Delta x(q), \Delta v(q))^T$ so that the whole p^*-sheet gets reconstructed. This is illustrated in the right panel of 6.3. Now, there is the question how to systematically choose the direction of these support vectors. It seems to be a good choice to support every point on the p^*-sheet by the closest point on the $p = 0$ sheet. Or said differently, we want to choose $(\Delta x, \Delta v)$ so that it is orthogonal to the $p = 0$ sheet. A vector parallel to the sheet is given by $(D_{\mathrm{xq}}, D_{\mathrm{vq}})^T$. Usually one can find an orthogonal vector by setting the scalar product to zero:

$$\begin{pmatrix} D_{\mathrm{xq}} \\ D_{\mathrm{vq}} \end{pmatrix} * \begin{pmatrix} \Delta x \\ \Delta v \end{pmatrix} = 0 \tag{6.15}$$

where $*$ denotes a scalar product. However, the scalar product is actually not well defined in phase space, since we are mixing position and velocity variables here. It requires the assumption of a metric m sothat

$$\begin{pmatrix} D_{\mathrm{xq}} \\ D_{\mathrm{vq}} \end{pmatrix} * \begin{pmatrix} \Delta x \\ \Delta v \end{pmatrix} = D_{\mathrm{xq}}\Delta x + m^2 D_{\mathrm{vq}}\Delta v \stackrel{!}{=} 0 \tag{6.16}$$

$$\begin{pmatrix} \Delta x \\ \Delta v \end{pmatrix} = \begin{pmatrix} -m^2 D_{\mathrm{vq}} \\ D_{\mathrm{xq}} \end{pmatrix} \tag{6.17}$$

where we have chosen in the last line one vector (of the infinitely many) that is orthogonal to the sheet. Now it is not at all clear what one should choose as value for the metric

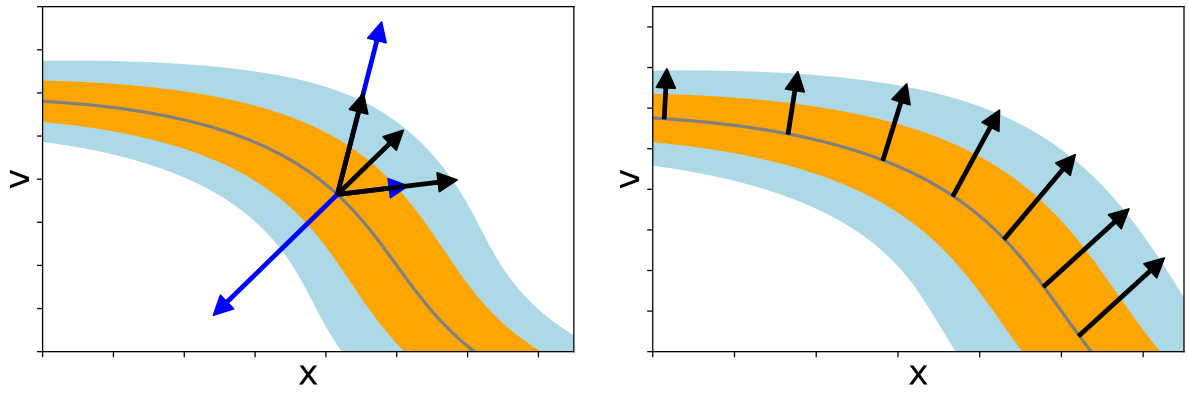

Figure 6.3: Left: Illustration of the vector re-normalization. We can choose a vector in an arbitrary phase space direction (blue vectors), but we can always re-normalize it (black vectors) so that it "hits" the p^*-sheet (here illustrated by the boundary of the upper blue contour). Right: reconstruction of the p^*-sheet by using a re-normalized vector at each location of the central sheet. The vectors are chosen to be orthogonal to the sheet. This required the assumption of a metric as discussed in the text.

m. It is also not clear whether one should even have a metric that depends on the spatial location (for example as a function of the tidal field). This will stay a bit of an open gap in the scheme that I am proposing here. However, it can in principle be tested how robust some reconstruction is against the choice of the metric. The smaller the velocity dispersion is, the less the choice of the metric matters, because the local volume is well described by displacements in any direction. The best motivated choice of a (global) metric that I could come up with is what we will call the *free-streaming metric*. It relates spatial and velocity coordinates by the question, how far can a particle with some velocity Δv travel in the absence of gravitational forces for $t \to \infty$? Using the equations of motion (2.45) and $v = \text{const.}$ (compare equation (2.46)) we find

$$\Delta x = \int_t^\infty \Delta v a^{-2}\,\mathrm{d}t \tag{6.18}$$

$$= \int_a^\infty \Delta v a^{-2} \left(\frac{\mathrm{d}a}{\mathrm{d}t}\right)^{-1} \mathrm{d}a \tag{6.19}$$

$$= \int_a^\infty \frac{\Delta v}{H a^3}\,\mathrm{d}a \tag{6.20}$$

$$\approx \int_a^\infty \frac{\Delta v}{H_0 a^{3/2}}\,\mathrm{d}a \tag{6.21}$$

$$= \frac{2\Delta v}{H_0 a^{1/2}} \tag{6.22}$$

under the approximation of an Einstein - de Sitter universe $H(a) = H_0 a^{-3/2}$. More detailed calculations can be done later. Let us just use as a metric $m = 2/(H_0 a^{1/2})$ for now. I found

this metric to work quite well in one-dimensional test cases. The contours of the warm phase space distributions throughout this thesis have mostly been created by usage of this metric.

Using $\Delta x = -m^2 D_{\mathrm{vq}}$ and $\Delta v = D_{\mathrm{xq}}$ for orthogonality and the free-streaming metric we can reconstruct the p^*-sheet as

$$\begin{pmatrix} x^*(q) \\ v^*(q) \end{pmatrix} = \begin{pmatrix} x(q) + \Delta x(q) \\ v(q) + \Delta v(q) \end{pmatrix} \tag{6.23}$$

$$\begin{pmatrix} \Delta x(q) \\ \Delta v(q) \end{pmatrix} = \frac{p^*}{D_{\mathrm{xq}^2} + m^2 D_{\mathrm{vq}^2}} \begin{pmatrix} -m^2 D_{\mathrm{vq}} \\ D_{\mathrm{xq}} \end{pmatrix} \tag{6.24}$$

$$m = \frac{2}{H_0 a^{1/2}} \tag{6.25}$$

This has been used to define the support vectors in the right panel of Figure 6.3. The reconstruction seems to work well in one dimension. So let us now try to figure out, how we can reconstruct a $\mathbf{p}^*$-sheet in a higher dimension. Actually, we can mostly keep the same parameterization. The main difference is, that many scalars get replaced by a sub-space that is described by tensors.

We can parameterize the space tangent to the sheet by a freely chooseable 3-vector $\boldsymbol{\beta}$ and the distortion tensors $\mathbf{D}_{\mathrm{Xq}}$ and $\mathbf{D}_{\mathrm{Vq}}$ and the space orthogonal to it by the free 3-vector $\boldsymbol{\alpha}$ and the two matrices $\mathbf{B}$ and $\mathbf{C}$ that we need to determine through the orthogonality condition:

$$0 = \begin{pmatrix} \mathbf{D}_{\mathrm{Xq}}\boldsymbol{\beta} \\ \mathbf{D}_{\mathrm{Vq}}\boldsymbol{\beta} \end{pmatrix} * \begin{pmatrix} \mathbf{B}\boldsymbol{\alpha} \\ \mathbf{C}\boldsymbol{\alpha} \end{pmatrix} \tag{6.26}$$

$$= (\mathbf{D}_{\mathrm{Xq}}\boldsymbol{\beta})^T \mathbf{B}\boldsymbol{\alpha} + m^2 (\mathbf{D}_{\mathrm{Vq}}\boldsymbol{\beta})^T \mathbf{C}\boldsymbol{\alpha} \tag{6.27}$$

$$= \boldsymbol{\beta}^T (\mathbf{D}_{\mathrm{Xq}}^T \mathbf{B} + m^2 \mathbf{D}_{\mathrm{Vq}}^T \mathbf{C})\boldsymbol{\alpha} \tag{6.28}$$

where we have assumed a scalar metric[2]. This is always solved by the choice $\mathbf{B} = -m^2 \mathbf{D}_{\mathrm{Vq}}^T$ and $\mathbf{C} = \mathbf{D}_{\mathrm{Xq}}^T$. Therefore we can describe the whole subspace that is orthogonal to the sheet by

$$\begin{pmatrix} \Delta\mathbf{x} \\ \Delta\mathbf{v} \end{pmatrix} = \begin{pmatrix} -m^2 \mathbf{D}_{\mathrm{Vq}}^T \cdot \boldsymbol{\alpha} \\ \mathbf{D}_{\mathrm{Xq}}^T \cdot \boldsymbol{\alpha} \end{pmatrix} \tag{6.29}$$

where $\boldsymbol{\alpha}$ is a 3-vector which can be chosen freely. The phase space vector in (6.29) maps into initial velocity space as

$$\Delta\mathbf{p} = (\mathbf{D}_{\mathrm{Xq}}\mathbf{D}_{\mathrm{Xq}}^T + m^2 \mathbf{D}_{\mathrm{Vq}}\mathbf{D}_{\mathrm{Vq}}^T)\boldsymbol{\alpha} \stackrel{!}{=} \mathbf{p}^* \tag{6.30}$$

[2]If a global metric is used (like it is here) then only the scalar metric makes sense. If a metric is used that is not global, also a tensor metric might make sense. However, the generalization to a tensor metric is straightforward and changes the argumentation nowhere.

So we can find the vector α that is associated with the $\mathbf{p}^*$-sheet by inverting this equation

$$\boldsymbol{\alpha} = (\mathbf{D}_{\mathrm{Xq}}\mathbf{D}_{\mathrm{Xq}}^T + m^2\mathbf{D}_{\mathrm{vq}}\mathbf{D}_{\mathrm{vq}}^T)^{-1}\mathbf{p}^* \tag{6.31}$$

This vector can then be used in equation (6.29) to get the phase space displacement that is associated with the $\mathbf{p}^*$-sheet. This way we obtain the position shift that is associated with the $\mathbf{p}^*$-sheet

$$\Delta\mathbf{x}^*(\mathbf{q}) = -m^2\left(\mathbf{D}_{\mathrm{vq}}^T(\mathbf{D}_{\mathrm{Xq}}\mathbf{D}_{\mathrm{Xq}}^T + m^2\mathbf{D}_{\mathrm{vq}}\mathbf{D}_{\mathrm{vq}}^T)^{-1}\right)\mathbf{p}^* \tag{6.32}$$

$$=: \mathbf{A}(\mathbf{q})\mathbf{p}^* \tag{6.33}$$

Note that we have defined something like a natural softening tensor $\mathbf{A}(\mathbf{q})$ here. If, for example, the distribution of $\mathbf{p}^*$ is taken be a spheroid with Gaussian kernel in initial velocity space, it will get mapped onto some ellipsoid with a Gaussian kernel in Eulerian space. The axes of that ellipsoid can be determined by a singular value decomposition of $\mathbf{A}$. Actually one could describe this behaviour by an anisotropic gravitational softening in N-body simulations. I have already developed an extension for Gadget 3 which can solve the Poisson equation for such ellipsoidal kernels in my master thesis (Stücker, 2015) - which I never published unfortunately. In my thesis I used it to follow Lagrangian volume elements to fight the fragmentation in WDM simulations. It would be straightforward to extend this behaviour to an anisotropic gravitational softening that incorporates this velocity smoothing as well.

6.2.1 Sampling a warm distribution from a cold sheet.

We use the math that we have developed in the last sections to sample a cosmological density field which includes a thermal smoothing due to the warmth of the primordial velocity distribution: This can be done, in principle, as a pure post-processing step in any cosmological simulation. We are able to reconstruct the thermal smoothing from a single cold sheet. As a proof of concept we show this here for the case of a two dimensional Zel'dovich approximation. For this setup this can be done relatively easily and quickly, since we only have to sample a four dimensional phase space (instead of a six-dimensional one). However, everything generalizes straightforwardly to six phase space dimensions and it can also be done with the cold sheet of an evolved particle distribution. We leave this open for possible future investigations.

We set up the 2d Zel'dovich approximation for a boxsize of $L = 10\,\mathrm{Mpc}/h$ and a power-spectrum that has the same two-point correlation function like an $m_x = 100\,\mathrm{eV}$ warm dark matter case. We use an Einstein - de Sitter universe and use $D = a = 1.$, but increase the amplitude of the power spectrum by a factor 1.7 so that there are already are some shell crossings. The full setup can be seen in Figure 6.5 (which we will discuss later). Note that this setup is very unphysical, but that does not matter for this proof of concept. Mathematically it has most of the properties that an evolved cosmological simulation has. From a (Lagrangian) grid of particles we reconstruct the sheet interpolating functions $\mathbf{x}(\mathbf{q})$

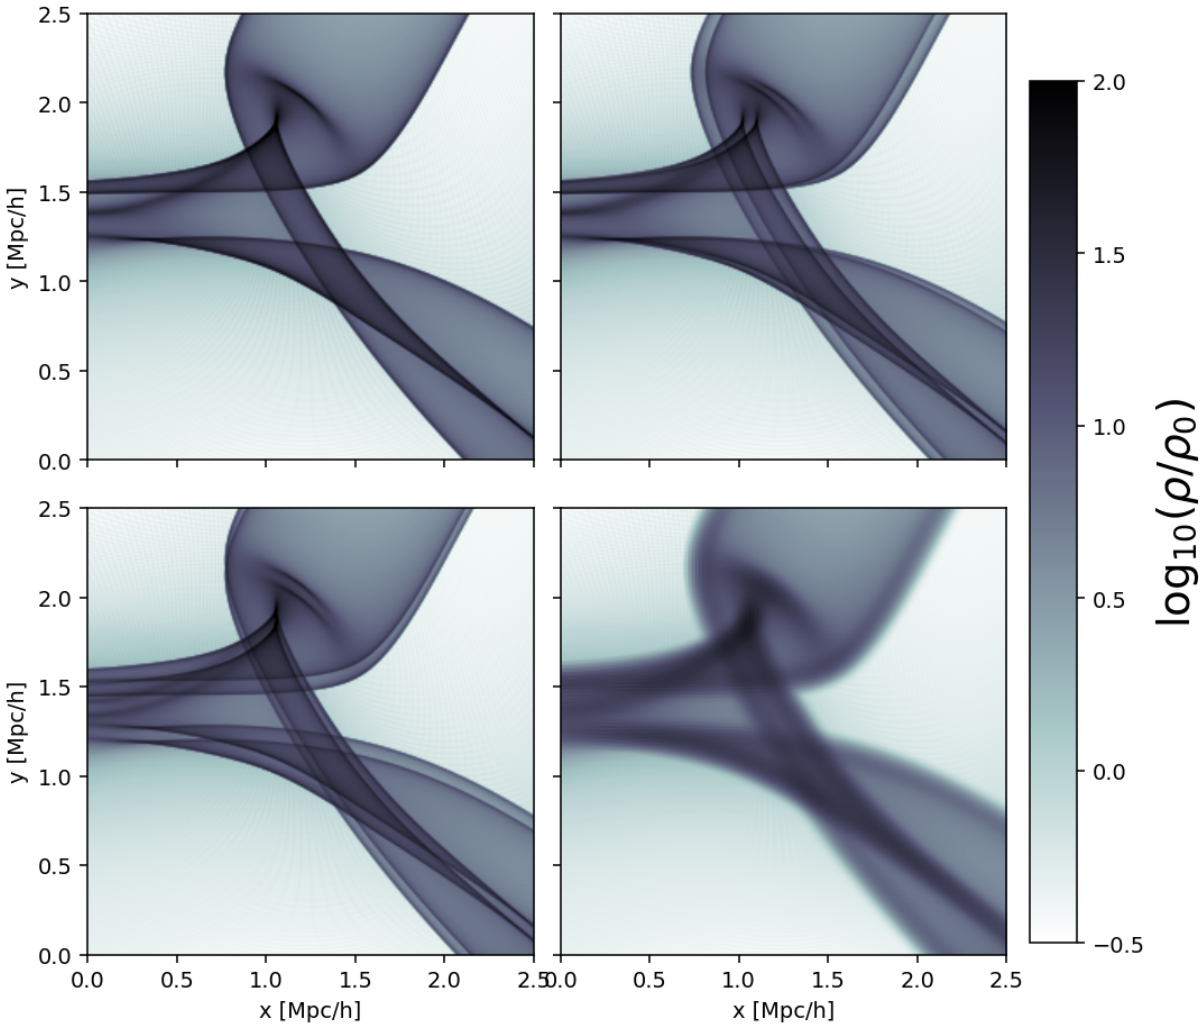

Figure 6.4: Density fields of a 2d Zel'dovich approximation (see text) for different sheets. Top left: the central "cold" sheet $\mathbf{p}^* = (0,0)^T$. Top right: the average of two sheets that are displaced in the x-direction of initial velocity space by $\mathbf{p}^* = (-1\,\mathrm{km\,s^{-1}}, 0)^T$ and $\mathbf{p}^* = (1\,\mathrm{km\,s^{-1}}, 0)^T$. Bottom left: analogue, but using displacements in the y-direction $\mathbf{p}^* = (0, -1\,\mathrm{km\,s^{-1}})^T$ and $\mathbf{p}^* = (0, 1\,\mathrm{km\,s^{-1}})^T$. With the procedures described in this section we are able to reconstruct these sheets from the cold sheet in the top left panel without any additional information. Bottom right: The density field of a warm distribution, as reconstructed by the sum of 13^2 different pseudo-sheets (see text). The density field of the warm distribution can be reconstructed from the dynamical information of a single cold sheet.

and $\mathbf{v}(\mathbf{q})$ and also their derivatives $\mathbf{D}_{\mathrm{Xq}}(\mathbf{q})$ and $\mathbf{D}_{\mathrm{Vq}}(\mathbf{q})$. In the top left panel of Figure 6.4 we show the density field that can be obtained from resampling this cold sheet.

Additionally we resample the density field for a set of different $\mathbf{p}^*$-sheets. To get an intuition of how these look and whether the reconstruction behaves properly we show the sum of two sheets that are displaced in initial x-velocity space ($\mathbf{p}^* = (-1\,\mathrm{km\,s^{-1}}, 0)^T$ and $\mathbf{p}^* = (1\,\mathrm{km\,s^{-1}}, 0)^T$) in the top right panel of Figure 6.5 and for two sheets that are displaced in y-velocity space in the bottom left panel. Note that the value of $1\,\mathrm{km\,s^{-1}}$ is very close to the typical thermal velocities of a 100 eV particle which we calculate as $\sqrt{\langle v^2 \rangle} \approx 980\,\mathrm{m\,s^{-1}}$ (compare Figure 2.2 and use $\sqrt{\langle v^2 \rangle} \approx 3.6 v_0$). Therefore the displacements of these sheets are comparable to the actual thermal displacements of typical particles in such a WDM cosmology. Note that far away from caustics the sheets can almost not be distinguished, but close to caustics the difference is quite dramatic. As a related effect the approximation of the stream-density as $\rho_0 |\det \mathbf{D}_{\mathrm{Xq}}|^{-1}$ works remarkably well in most of the volume, but breaks down in caustics (where the approximation goes to infinity, whereas the true density of the stream is limited by the thermal softening) - see White & Vogelsberger (2009) for more details.

We assume a 2d Gaussian velocity distribution for the initial velocities

$$f(\mathbf{p}) = \frac{1}{2\pi\sigma} \exp\left(\frac{-|\mathbf{p}|^2}{2\sigma^2}\right) \tag{6.34}$$

and use $\sigma = 1$km/s. Note that this does not correspond to the velocity distribution of the thermal relic, but we only make rough estimates here anyways. (Typical velocities should match within a factor of two with the 100eV relic.) We then calculate the density field of 13^2 sheets which range uniformly from -3σ to 3σ in each velocity coordinate. We determine as the total density

$$\rho_{\mathrm{tot}} = \sum_{p_x} \sum_{p_y} \rho_{\mathbf{p}} f(p_x, p_y) \Delta p^2 \tag{6.35}$$

where $\Delta p = 0.5\sigma$ is the spacing of the velocity grid. We display the result in the bottom right panel of Figure 6.4 and in Figure 6.5. Note that the fuzziness of these images is not due to resolution problems (the resolution is the same as in the other panels of Figure 6.4), but purely due to the thermal broadening. In this Zel'dovich case it appears like there is almost everywhere a relatively similar smoothing. That is so, since $\mathbf{D}_{\mathrm{Vq}}$ (which determines the width of caustics) has at least in order of magnitude similar values at the caustics. However, this is not expected to be the case in an actual cosmological setup. As was show in section 2.6.1, the thermal smoothing in caustics ranges easily over 9 orders of magnitude. However, possibly the first and largest caustics behave similarly to what can be seen in the bottom right panel of Figure 6.4.

Note that this way of re-sampling the velocity distribution by reconstructing pseudo-sheets is actually quite inefficient. It would require sampling a 6d phase space distribution in the case of a full cosmological setup. However, this is not at all necessary. Instead it would be faster (and also more accurate) to sample the 3d central cold sheet, but deposit

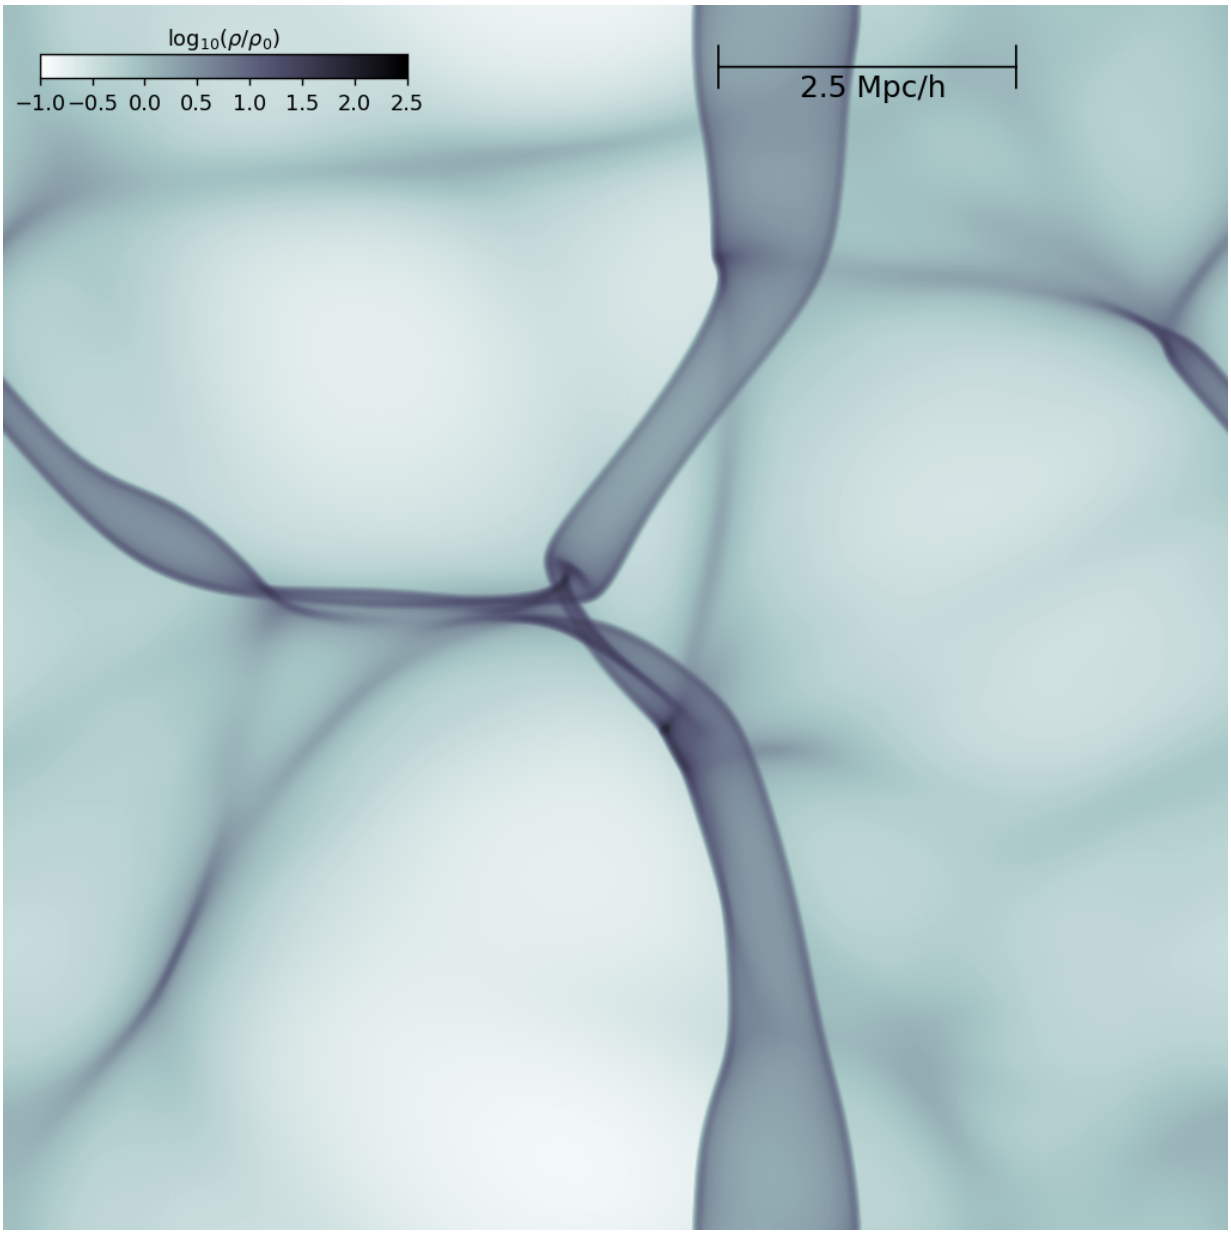

Figure 6.5: The density field of a warm sheet in the Zel'dovich approximation. Even for this relatively warm dark matter model of 100eV, the thermal smoothing does not alter large scale structures significantly, but only introduces a "fuzziness" into the density field. (The fuzziness is only due to the thermal broadening, but not due to resolution issues).

a Gaussian ellipsoid for each sampled pseudo-particle - instead of δ-function-like particles. That ellipsoid is given by a singular value decomposition of the matrix A from equation (6.33)

$$A = \mathbf{USV}^T \tag{6.36}$$

the singular values s_i are the semi-major axes and their directions are given by the column vectors $\mathbf{u}_i$.

Now this operation of depositing a large set of ellipsoids with different alignments is an operation which would require some time[3] to be implemented efficiently. However, if an efficient implementation is found, it will be straightforward to make simulations of warm dark matter which include a thermal smoothing.

We conclude that *the density field of a warm distribution function can be reconstructed from a single cold sheet*. This works in one dimension (as can be seen in many Figures throughout this thesis), but it also works in higher dimensional cases. This opens up possibilities for understanding the implications of warm distribution functions as a pure post-processing step of cold simulations. Further, it opens up the possibility of warm dark matter simulations which do not operate in the cold limit, but instead actively model the smoothing of the density field due to the thermal velocities. The additional requirement is a convolution with an ellipsoidal kernel that depends on the Lagrangian coordinate $\mathbf{q}$. However, The effect of thermal velocities is expected to be small, so it is not clear whether this would be worth the effort. In the next section, we discuss the possibility of applying the presented method in the context of the modelling of the cosmic neutrino background.

6.3 On the treatment of hot distribution functions (like neutrinos)

While it might be irrelevant to model the thermal smoothing in the case of warm dark matter models, it is certainly relevant in the case of the cosmic neutrino fluid (also known as cosmic neutrino background). The neutrino fluid is expected to be non-relativistic today, but it maintains a large enough velocity dispersion that typical neutrinos can escape the potential wells of smaller structures. As an example have a look at the left panel of Figure 6.6. This is how a neutrino fluid could look like in one dimension. This is from a one-dimensional simulation where we put all of the mass into the central cold sheet (as is a good approximation to the cosmological case where the cold dark matter sheet dominates the mass distribution).

We have a hot distribution function here and it seems quite impossible to reconstruct the whole distribution just from a single cold sheet (as was possible in warm cases). That is so, because the distribution function can not be approximated in a locally linear way, but has many higher order terms. In the cosmological case the central sheet will wind up to become extremely complex (as we have discussed in previous chapters already), whereas

[3] which I do not have left before handing in this thesis

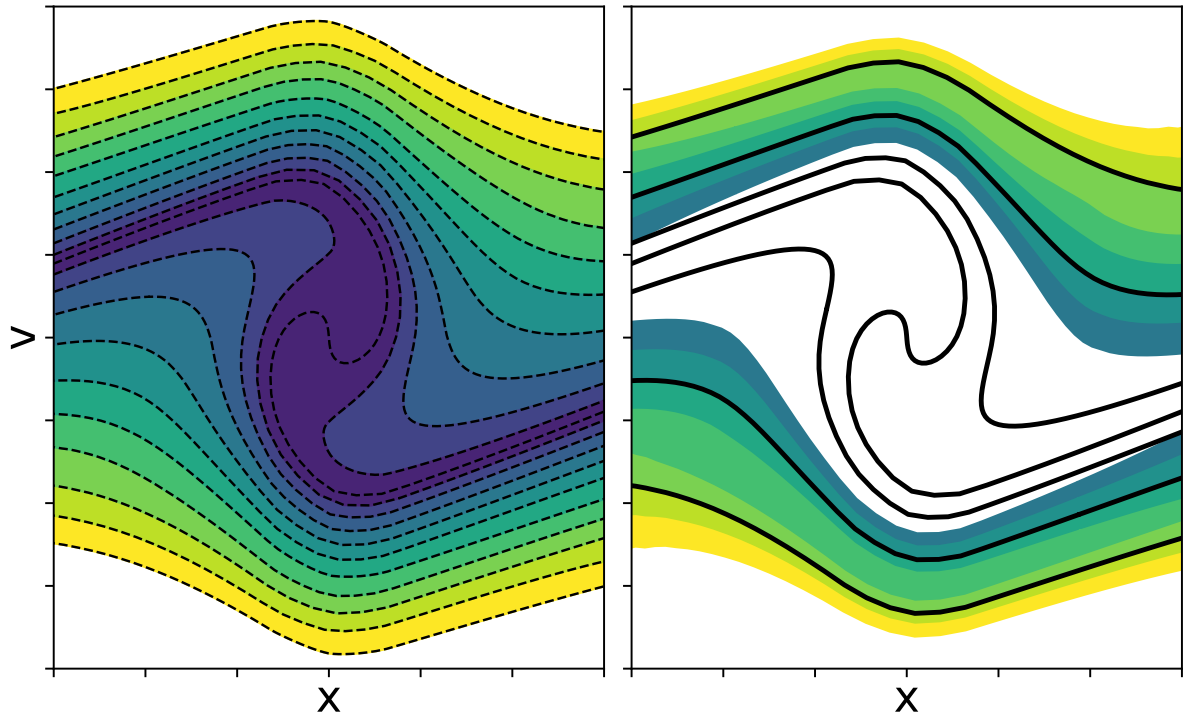

Figure 6.6: Left: A hot distribution in phase space (as for example expected for the cosmic neutrino background). The central parts of the distribution have complex geometries whereas higher velocity parts have relatively simple geometries. Right: This leads to the idea of approximating the higher-velocity parts of the distribution function by a few sheets and reconstructing the surrounding phase space from these. The black lines show some sheets that have been traced in the simulation and the contours show the pseudo-sheets that I have reconstructed from these - they can be compared with the "true" simulated sheets in the left panel. The reconstruction works well in the high-velocity regime, but it is close to impossible in the central regime. If there is little mass in the low-velocity parts of the distribution, it might be possible to skip a reconstruction of the phase space distribution here.

sheets further outside of the distribution stay much simpler. Because of the high velocity of particles in the outer regions they either do not get bound at all or they get bound, but go through a much smaller number of caustics than the central sheet. In the case of the cosmological neutrino fluid, maybe typical neutrino-sheets go through one or two caustics. Therefore they are much simpler than the central sheet which goes through hundreds of caustic crossings. So while the central sheet of the hot neutrino distribution has the same complexity as the cold dark matter sheet, typical sheets of it have much lower complexity.

This observation leads me to the suggestion to simulate the cosmic neutrino background by just a small set of sheets. The sheets with higher velocities should remain relatively simple, and it should be easy to reconstruct them with interpolation techniques. These are then used to reconstruct a "locally warm" phase space distribution around them. I show in panel Figure 6.6 how I have reconstructed from just four high velocity sheets (two for negative and positive velocities each) the major part of the high-velocity distribution. Note that the reconstruction is not perfect, but it can be improved by using more sheets (or possibly by somehow connecting sheets in the velocity dimension to interpolate between them). The colder part of the distribution is hard to reconstruct, and in the cosmological case it is practically impossible (as it is for the cold central sheet as discussed in previous chapters). However, if only a small fraction of particles resides in the low-velocity part of the distribution, it might be ok to approximate these parts by cold sheets or neglect them all together.

This scheme could then work in the following way

- At the beginning of the simulation construct a set of sheets which are displaced in velocity space
- Follow these sheets throughout the simulation and detect somehow which sheets can locally reconstruct their phase space neighbourhood
- Whenever a density field needs to be calculated: deposit the reliable sheets and their associated phase space neighbourhood as warm sheets (as described in the previous section). Deposit the unreliable sheets as cold sheets (or with an artificially reduced thermal broadening).

Of course such a scheme would require that a lot of practical details are worked out. However, I think if these practical details can be resolved reasonably well, it might be more accurate and efficient than existing techniques to simulate the cosmic neutrino background.

I apologize for missing literature references in this chapter. I have not had the time to do a more complete literature research, since I have only started writing this chapter and developing the schemes here in the very last days before handing in my thesis. Therefore I have only referenced publications that I know about. However, at the very least I have to point out here that the idea of using multiple sheets to approximate a hot phase space distribution has already been explored in Kates-Harbeck et al. (2016). However, these were kept as cold sheets in that case whereas I have proposed the usage of "warm sheets" here.

Chapter 7

Conclusions and outlook

In this thesis we have discussed the dark matter fluid from the phase space perspective. Based on this we have developed new mathematical tools and numerical simulation techniques.

The mathematical tools include an algorithm for detecting single-stream regions in evolved cosmological simulations, a structure classification scheme which is based on the distortions of the fine-grained dark matter sheet, and an excursion set formalism for single-stream regions. We have for example estimated that the typical density of our universe is only of order 0.4% of the mean density if dark matter is a 100GeV WIMP whereas it would be 8% of the mean density if dark matter were a 1keV thermal relic. Further, we found that single-stream regions form distinct regions which are separated by collapsed multi-stream regions.

As new numerical simulation techniques we have introduced the "sheet + release" simulation approach here. This simulation scheme uses an interpolation reconstruction of the dark matter sheet in phase space in Lagrangian regions where such a reconstruction is possible and reliable. It switches to an N-body simulation approach when the interpolation scheme becomes unreliable (which is mostly the case in haloes). Further, we developed a novel discretization of the force-field which is based on an oct-tree of cubes. By combining these approaches we present for the first time a new simulation scheme which makes possible warm dark matter simulations which are very accurate in low-density regions (like voids, pancakes and filaments) while at the same time being able to resolve the complex and dense centers of halos. Further, this scheme gives "perfect" phase space information in regions where the sheet can be reconstructed and still has statistical phase space information available in released regions thanks to the geodesic deviation equation. However, we notice that there are still some issues with the convergence of stream-densities which we want to address before publishing this.

To test the new simulation code, we have presented a zoom-in simulation of a small warm dark matter halo. While N-body simulations showed artificial fragments, the sheet + release simulations did not show such discreteness effects – even at very high force-resolution. The density profile of the halo converged very well and we found that it is consistent with an NFW-profile. Further, it agrees well with the N-body case, although this

one showed artificial fragmentation. The preliminary conclusion is that N-body simulations can still give reliable results when investigating quantities that involve a lot of coarse-graining – like the density-profile. However, this has to be investigated in more detail in a larger set of simulations which also include more significant artificial fragments.

Finally, we have discussed, how warm distribution functions can be reconstructed in phase space from a single cold sheet. That means, for example, that one can estimate the thermal broadening of the density field in a pure post-processing step of typical warm (or cold) dark matter simulations. Further, one could use this to make warm dark matter simulations that explicitly model the effect of the primordial velocity dispersion as an additional smoothing. This smoothing is anisotropic and is a function of the Lagrangian coordinates. As another application one could simulate the high-velocity part of the cosmic neutrino background by a set of sheets which are used to locally reconstruct their phase space neighbourhood. While a lot of practical details would still need to be addressed for such schemes to work efficiently, we have shown here that they are possible in principle.

The schemes that have been developed in this thesis can be applied in many different contexts in the future. For example, we can make a detailled investigation of the phase space structure of warm dark matter haloes (assuming we are able to make the stream-densities reliable). Further, we might aim to measure the halo-mass function and possibly the subhalo-mass function for warm dark matter cosmologies as a rescaleable function. This would for example be important to interpret the number of small scale perturbations in the arcs of strong gravitational lenses – and therefore for constraining the warmth of dark matter. Further, we could consider to explore the implications of the thermal smoothing in warm dark matter simulations – maybe initially as a pure post-processing step on existing simulations. Further, we can try to estimate whether the scheme that we have proposed for neutrino simulations would give significant advantages over existing methods. Finally, the concepts that we have developed for the classification of cosmological structures could be used to understand and interpret the cosmic web in new ways.

Appendix A

More images of structure classifications

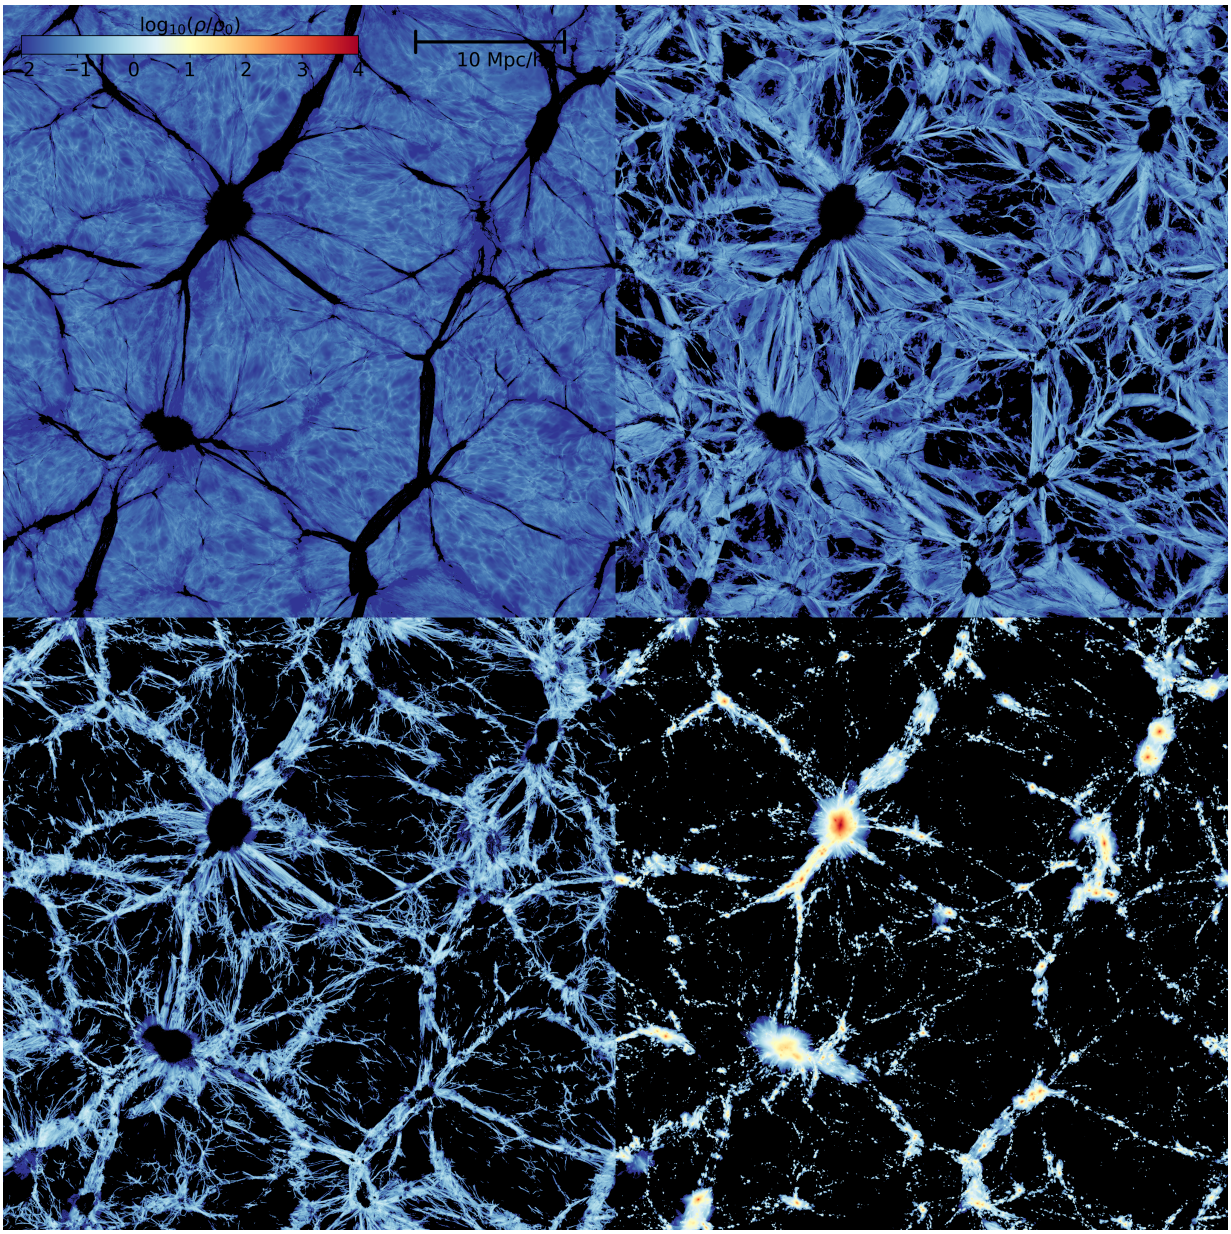

Figure A.1: Projections of the density field (depth = 0.8 Mpc/h) of a CDM simulation selected on different morphological types. Top left: only voids, top right: only pancakes, bottom left: only filaments, bottom right: only halos. Arguably the classification selects the structures close to the way one would intuitively classify them. We provide other versions of this Figure in Figure 2.26 and Figure A.2.

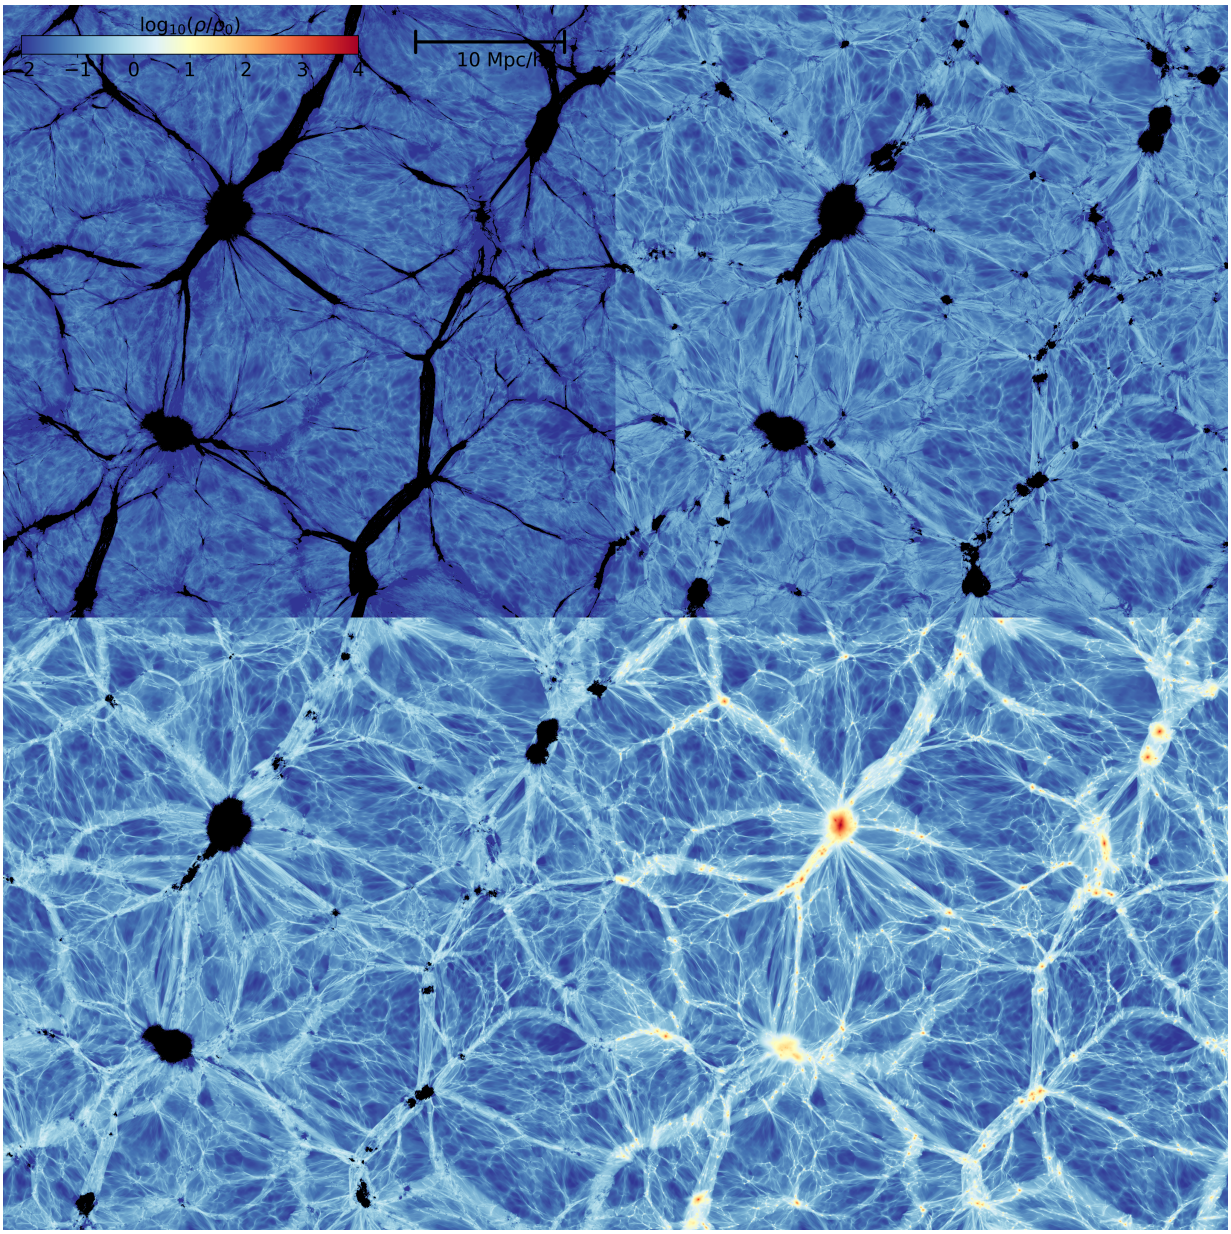

Figure A.2: Projections of the density field (depth = 0.8 Mpc/h) of a CDM simulation selected on different morphological types. Top left: only voids, top right: voids + pancakes, bottom left: voids + pancakes + filaments, bottom right: all mass. Arguably the classification selects the structures close to the way one would intuitively classify them. We provide other versions of this Figure in Figure 2.26 and Figure A.1.

Bibliography

Abel T., Hahn O., Kaehler R., 2012, MNRAS, 427, 61

Adkins G. S., McDonnell J., Fell R. N., 2007, Phys. Rev. D, 75, 064011

Angulo R. E., White S. D. M., 2010, Monthly Notices of the Royal Astronomical Society, 401, 1796

Angulo R. E., Hahn O., Abel T., 2013, MNRAS, 434, 3337

Arnold V. I., Shandarin S. F., Zeldovich I. B., 1982, Geophysical and Astrophysical Fluid Dynamics, 20, 111

Barreira A., Schmidt F., 2017, Journal of Cosmology and Astro-Particle Physics, 2017, 053

Blas D., Lesgourgues J., Tram T., 2011, Journal of Cosmology and Astro-Particle Physics, 2011, 034

Bode P., Ostriker J. P., Turok N., 2001, ApJ, 556, 93

Bond J. R., Myers S. T., 1996, ApJS, 103, 1

Bond J. R., Cole S., Efstathiou G., Kaiser N., 1991, ApJ, 379, 440

Bond J. R., Kofman L., Pogosyan D., 1996, Nature, 380, 603

Bouchet F. R., Colombi S., Hivon E., Juszkiewicz R., 1995, A&A, 296, 575

Boylan-Kolchin M., Springel V., White S. D. M., Jenkins A., Lemson G., 2009, Monthly Notices of the Royal Astronomical Society, 398, 1150

Buehlmann M., Hahn O., 2018, arXiv e-prints: 1812.07489,

Chappell J. M., Chappell M. J., Iqbal A., Abbott D., 2013, Physics International, 3, 50

Chiueh T., Lee J., 2001, ApJ, 555, 83

Eisenstein D. J., Hu W., 1999, ApJ, 511, 5

Falck B., Neyrinck M. C., 2015, MNRAS, 450, 3239

Falck B. L., Neyrinck M. C., Szalay A. S., 2012, ApJ, 754, 126

Frenk C. S., White S. D. M., 2012, Annalen der Physik, 524, 507

Green A. M., Hofmann S., Schwarz D. J., 2005, J. Cosmology Astropart. Phys., 8, 003

Hahn O., Abel T., 2011, MNRAS, 415, 2101

Hahn O., Angulo R. E., 2016, MNRAS, 455, 1115

Hahn O., Abel T., Kaehler R., 2013, MNRAS, 434, 1171

Hernquist L., 1990, ApJ, 356, 359

Hidding J., Shandarin S. F., van de Weygaert R., 2014, MNRAS, 437, 3442

Hockney R. W., Eastwood J. W., 1981, Computer Simulation Using Particles

Hubble E., 1929, Proceedings of the National Academy of Science, 15, 168

Hummer G., 1996, Journal of Electrostatics, 36, 285

Iannuzzi F., Dolag K., 2011, MNRAS, 417, 2846

Kates-Harbeck J., Totorica S., Zrake J., Abel T., 2016, Journal of Computational Physics, 304, 231

Lewis A., Bridle S., 2002, Phys. Rev. D, 66, 103511

Lin H., et al., 2009, ApJ, 699, 1242

Lovell M. R., Frenk C. S., Eke V. R., Jenkins A., Gao L., Theuns T., 2014, MNRAS, 439, 300

Macmillan W. D., 1958, Dover Publications, New York

Marsh D. J. E., 2016, Phys. Rep., 643, 1

Navarro J. F., Frenk C. S., White S. D. M., 1996, ApJ, 462, 563

Neyrinck M. C., 2013, MNRAS, 428, 141

Pandey B., White S. D. M., Springel V., Angulo R. E., 2013, Monthly Notices of the Royal Astronomical Society, 435, 2968

Peebles P. J. E., 1980, The large-scale structure of the universe

Planck Collaboration et al., 2016, A&A, 594, A13

Planck Collaboration et al., 2018, arXiv e-prints, p. arXiv:1807.06205

Powell D., Abel T., 2015, Journal of Computational Physics, 297, 340

Press W. H., Schechter P., 1974, ApJ, 187, 425

Price D. J., Monaghan J. J., 2007, MNRAS, 374, 1347

Ramachandra N. S., Shandarin S. F., 2015, MNRAS, 452, 1643

Ramachandra N. S., Shandarin S. F., 2017, MNRAS, 470, 3359

Robertson A., Massey R., Eke V., 2017, MNRAS, 465, 569

Rubin V. C., Ford Jr. W. K., 1970, ApJ, 159, 379

Sahni V., Shandarin S., 1996, MNRAS, 282, 641

Sandvik H. B., MÃűller O., Lee J., White S. D. M., 2007, Monthly Notices of the Royal Astronomical Society, 377, 234

Schmidt A. S., White S. D. M., Schmidt F., Stücker J., 2018, MNRAS, 479, 162

Schneider A., Smith R. E., Macciò A. V., Moore B., 2012, MNRAS, 424, 684

Shandarin S. F., Medvedev M. V., 2014, preprint, (`arXiv:1409.7634`)

Shandarin S. F., Zeldovich Y. B., 1989, Reviews of Modern Physics, 61, 185

Shandarin S., Feldman H. A., Heitmann K., Habib S., 2006, MNRAS, 367, 1629

Shandarin S., Habib S., Heitmann K., 2012, Phys. Rev. D, 85, 083005

Sheth R. K., Mo H. J., Tormen G., 2001, MNRAS, 323, 1

Sousbie T., Colombi S., 2016, Journal of Computational Physics, 321, 644

Springel V., 2005a, MNRAS, 364, 1105

Springel V., 2005b, MNRAS, 364, 1105

Springel V., et al., 2005a, Nature, 435, 629

Springel V., et al., 2005b, Nature, 435, 629

Stücker J., 2015, Master's thesis, `https://www.usm.uni-muenchen.de/CAST/cms/media/theses/master/MA_Stuecker.pdf`

Stücker J., Busch P., White S. D. M., 2018, MNRAS, 477, 3230

Vegetti S., Czoske O., Koopmans L. V. E., 2010, MNRAS, 407, 225

Vogelsberger M., White S. D. M., 2011, MNRAS, 413, 1419

Vogelsberger M., White S. D. M., Helmi A., Springel V., 2008, MNRAS, 385, 236

Wang J., White S. D. M., 2007, MNRAS, 380, 93

White S. D. M., Vogelsberger M., 2009, MNRAS, 392, 281

White S. D. M., Frenk C. S., Davis M., 1983, ApJ, 274, L1

Yang L. F., Neyrinck M. C., Aragón-Calvo M. A., Falck B., Silk J., 2015, MNRAS, 451, 3606

Yoshikawa K., Yoshida N., Umemura M., 2013, ApJ, 762, 116

Zel'dovich Y. B., 1970, A&A, 5, 84

Zeldovich I. B., Einasto J., Shandarin S. F., 1982, Nature, 300, 407

Zentner A. R., 2007, International Journal of Modern Physics D, 16, 763

Zwicky F., 1933, Helvetica Physica Acta, 6, 110

van de Weygaert R., Platen E., 2011, in International Journal of Modern Physics Conference Series. pp 41–66 (`arXiv:0912.2997`), doi:10.1142/S2010194511000092

Acknowledgements

My PhD was a long journey and I am very grateful that I had the opportunity to share this period of my life with so many wonderful people.

I want to thank everyone who supported and inspired me while I was working on my PhD. First of all I am very thankful to my supervisor Simon White who has made this PhD project possible. I have learned a lot from Simon and his ideas and questions inspired me a lot. I loved it to work in an environment where new ideas were greatly appreciated and where I was constantly encouraged to think independently and critically. Beyond this Simon has always supported me in visiting various conferences and workshops and very importantly in building up a collaboration with Oliver Hahn and Raul Angulo. I would like to thank Oliver Hahn and Raul Angulo for the welcoming, exciting and scientifically stimulating atmosphere in our collaboration. I cannot stress enough, how much I have learned from you two, and how much the exciting discussions have motivated me during my PhD! Further, I would like to thank Philipp Busch with whom I have worked together on the project about the median density of the universe. This was arguably one of the most fun projects that we have ever worked on. Beyond this, daily discussions with Philipp were always helpful and motivated me on a daily basis.

Next I would like to thank my fellow PhD (and non-PhD) students (in anti-alphabetical order) Zsolt, Vlas, Tamas, Timo, Philipp, Ricard, Ninoy, Max, Matthias, Matteo, Martha, Manuel, Miranda, Linda, Julian, Jiamin, Francesco, Aoife, Andreas, Alessandro and many others. Thanks for the great time, guys! Without you I would have never made it! I hope we will always stay in touch, even when thousands of kilometers separate us.

Moreover, I want to thank my girlfriend Ece, who has filled my heart with meaning when I felt somewhat empty otherwise. Further, I want to thank my family, who never stopped supporting me and always gave me the most warm welcome when I came back home.