Stellar orbits in adiabatically and impulsively evolving dark matter dominated potentials

Jan David Burger

Faculty of Physical Sciences
University of Iceland
2021

Stellar orbits in adiabatically and impulsively evolving dark matter dominated potentials

Jan David Burger

Dissertation submitted in partial fulfillment of a
Philosophiae Doctor degree in Astronomy

Advisor
Prof. Jesús Zavala Franco

PhD Committee
Prof. Jesús Zavala Franco
Prof. Gunnlaugur Björnsson
Dr. Miguel Ángel Sanchez Conde

Opponents
Prof. Justin Read
Dr. Nicola Amorisco

Faculty of Physical Sciences
School of Engineering and Natural Sciences
University of Iceland
Reykjavik, August 2021

Stellar orbits in adiabatically and impulsively evolving dark matter dominated potentials
Stellar orbits in evolving potentials
Dissertation submitted in partial fulfillment of a *Philosophiae Doctor* degree in Astronomy

Faculty of Physical Sciences
School of Engineering and Natural Sciences
University of Iceland
Dunhagi 5
107, Reykjavik
Iceland

Telephone: 525-4000

Bibliographic information:
Jan David Burger, 2021, *Stellar orbits in adiabatically and impulsively evolving dark matter dominated potentials*, PhD dissertation, Faculty of Physical Sciences, University of Iceland, 285 pp.

ISBN 978-9935-9614-2-6

Printing: Háskólaprent ehf.
Reykjavik, Iceland, August 2021

Abstract

In the ΛCDM concordance model cold and collision-less dark matter (CDM) accounts for about four fifths of the Universe's total matter content, while visible baryonic matter makes up the final fifth. The ΛCDM model successfully describes how large scale structure forms from tiny density fluctuations in the early Universe, explaining the growth of dark matter (DM) overdensities and their eventual collapse into gravitationally self-bound objects called haloes which host the visible galaxies. The most enduring challenges to the ΛCDM paradigm arise on the scale of dwarf galaxies whose internal kinematics are seemingly at odds with the radial density profiles of the corresponding DM haloes predicted by CDM *N*-body simulations. While the evidence for such mismatches remains controversial, several mechanisms have been proposed to reconcile the success of ΛCDM on large scales with these observations on smaller scales. Typically, such mechanisms rely on non-gravitational effects to redistribute mass within, and thus alter the gravitational potential of, DM haloes. Depending on the specific mechanism, this mass redistribution occurs on timescales which are either longer or shorter than the typical dynamical timescales of stars in dwarf galaxies. In this thesis, I show that stars can be used as kinematic tracers to differentiate between adiabatically, i.e., slowly evolving, and impulsively, i.e., fast evolving potentials. In adiabatically evolving potentials, the actions of kinematic tracers are conserved. For moderately fast evolving spherical potentials, I show that the evolution of radial actions is perturbative and oscillatory. In this regime, the evolution of radial action distributions of tracers is well described by a diffusion equation. The evolution of a tracer's radial action becomes non-perturbative as the rate at which the potential changes becomes large compared to the tracer's radial period. An impulsive change in the gravitational potential, under which actions are not conserved, changes the energy of kinematic tracers by an amount that depends on their orbital phase. Using controlled *N*-body simulations, I demonstrate that as a consequence of the latter, shell-like signatures related to phase mixing emerge in the phase space of kinematic tracers whose orbits where characterized by similar integrals of motion prior to the change in the potential. Such *orbital families* are reasonably approximated by sets of stars with similar ages and metallicites. Ongoing phase mixing in *orbital families* within dwarf galaxies with slow-rising rotation curves would indicate the presence of strong and impulsive supernova feedback which may have converted the central cusps of the dwarfs' host haloes into cores. Using a suite of hydrodynamic simulations of an isolated dwarf, I find that the kinematic properties of stars and gas differ in many ways between galaxies whose host haloes have an impulsive or an adiabatic core formation history. The complete absence of shell-like signatures of phase mixing in *orbital families* would be a strong indication that cores are likely formed adiabatically, for example through self-interactions between the DM particles.

Ágrip

Í hinu samræmda ΛCDM líkani er um fimmtungur alls efnis sýnilegar þungeindir og fjórir fimmtu hlutar hulduefni. Líkanið gefur góða lýsingu á því hvernig stórgerð alheimsins er til komin vegna lítilla þéttleikasveiflna í árdaga og hvernig yfirþéttleiki hulduefnis þróast yfir í þyngdarbundna hjúpa sem hýsa sýnilegar vetrarbrautir. Sterkasta áskorun líkansins kemur fram á skala dvergvetrarbrauta en gangfræði þeirra virðist stangast á við þéttleikaprófíl hulduefnishjúpanna sem spáð er í fjölagna hermireikningum. Þó vísbendingarnar um þetta misræmi séu umdeildar eru nokkrar mögulegar skýringar á því hvernig á þessu getur staðið. Oft byggjast þessar skýringar á þyngdarverkunarlausum hrifum sem breyta massadreifingu hulduefnishjúpsins og þar með einnig þyngdarmætti hans. Það fer svo eftir eðli hrifanna hvort massadreifingin breytist á tímaskala sem er styttri eða lengri en dæmigerður hreyfifræðilegur tími stjarna í dvergvetrarbrautum.

Í þessari ritgerð sýni ég að nota má stjörnur sem hreyfifræðilega spora (tracer) til að greina á milli annars vegar óverminna (hægfara) breytinga á mættinu og hins vegar snöggra breytinga á því. Í óvermnum breytingum er verkun hreyfifræðilegu sporanna varðveitt. Þá sýni ég að í kúlusamhverfu mætti sem þróast miðlungi hratt er radíalverkunin bæði truflin (perturbative) og sveiflukennd og að lýsa má þróun dreifingar radíalverkunarinnar með sveimjöfnu. Tímaþróun radíalverkunar sporanna verður ótruflin þegar tímaþróun mættisbreytinganna verður hraðari en sem nemur útþáttarlotu (radial period) sporanna. Við snögga breytingu í þyngdarmættinu, þar sem verkunin er ekki varðveitt, breytist orka sporanna um stærð sem er háð brautarfasa þeirra. Með stýrðri hermun fjölagnakerfa sýni ég að afleiðing hins síðarnefnda er að í fasarúmi sporanna koma fram merki um hvelamyndun í brautum þeirra. Þeim má lýsa allvel með fjölskyldum stjarna af svipuðum aldri og málminnihaldi. Fasablöndun fjölskyldna brauta stjarna í dvergvetrarbrautum með hægt vaxandi snúningshraðarit er vísbending um sterka og snögga sprengistjörnusvörun sem gæti hafa breytt þéttleikaprófíl hjúps dvergvetrarbrautarinnar. Með hermun einangraðrar dvergvetrarbrautar sýni ég að hreyfifræði gass og stjarna er ólík eftir því hvort kjarni hjúpsins hefur þróast óvermið eða snögglega. Ef engin merki er að finna um fasablöndnun brautarfjölskyldna er það sterk vísbending um að hjúpurinn hafi þróast á óverminn hátt, t.d. með eiginvíxlverkun á milli einda hulduefnisins.

Für meine Eltern und meine Großeltern:
Gisela & Robert,
Waltraud & Horst
und Elisabeth & Mathias

Table of Contents

Acknowledgments

This thesis marks the end of my time as a PhD student and, therefore, also my time as a student in general. I am thankful for all of the people who have studied with me, worked with me, or chosen to spend their free time with me over those last ten years. They are to many to name them all, but I would like to mention a few people without whom this thesis would not exist, or without whom my time in Iceland would not have been as exceptional as it was.

A big thank you goes to Jesús. For hiring me as his PhD student, for helping me get settled in Iceland, for guiding me in my work throughout the years, and for his incredible support, especially during the last year of my PhD. Working on this project has been a pleasure.

Next I would like to thank my PhD committee, Miguel-Angel and Gulli, with special thanks to Gulli for translating the abstract of this thesis.

I would also like to thank the other members of the Astrophysics group, for countless discussions, mostly over lunch and not always limited to physics. Mark deserves a special mention here. Maybe one day I will learn cricket terminology. Additionally, I would like to thank Sebastian, my colleague and office mate, for the very nice working atmosphere, the aimless but entertaining conversations about German politics, and for being my personal tech support on several occasions.

A great deal of people contributed to making the time around my work more active and more fun. This includes the members of Richa 116 and the Bandý / hiking group. Thanks for always keeping me running. Special thanks go to Emmeke and Watse, for keeping our group of friends together during Covid times.

Finally, I would like to thank Arngunnur, for all our time together, for encouraging me when my work did not go as planned, and for coming with me all the way to the top of Kristínartindar. I would also like to thank her family, for being as kind and welcoming as they are. And last but not least, I would like to thank the people who know me best, but have not seen nearly enough of me for the last four years, my family – and in particular my parents and my brothers: Gisela, Robert, Nils, and Tim – for their love and their support.

This work was supported by a Grant of Excellence from the Icelandic Research Fund (grant number 173929). The data from which the results presented in Chapters 5-9 are derived were generated on the large scale facility Garpur, provided by the Icelandic High Performance Computing Centre at the University of Iceland.

Chapter 1

Introduction

Humanity's fascination with the Universe dates back to the beginning of recorded history. For centuries, we have relied upon the sky for navigation, the sun's movement during the day and the position of stellar constellations at night. As I am writing this, humans have been to the moon, sent research probes to Mars, and relied upon the use of satellite technology for decades. The broad availability of telescopes has turned stargazing into a a pastime for a lot of people around the world.

Over the last century, our understanding of the Universe has evolved a lot, and scientific interest into the Universe has perhaps never been greater. It is, after all, the largest and most comprehensive laboratory available to mankind. In contrast to - and thanks to - scientists of past generations we now know the basic physics of the Universe fairly well, a circumstance that allows us to use ever more precise observations of radiation, gas, planets, stars, galaxies, and galaxy clusters to put theories of fundamental physics to the test. The biggest advancement throughout the last century has been the establishment of the standard concordance model of cosmology. After the cosmic microwave background (CMB) was first observed, the hypothesis that the Universe was created in a hot Big Bang quickly became the established consensus among cosmologists. Within the standard cosmological model, we can date the time at which the CMB radiation was emitted to ~ 400000 years after the Big Bang. The subsequent evolution of the Universe is governed by the gravitational growth of small density fluctuations. How fast these fluctuations grow depends on the composition of the Universe's energy density.

In ΛCDM, the standard model of structure formation (e.g. Peebles 1993), the energy density of the Universe today consists, overwhelmingly, of things that we cannot observe. While a cosmological constant (Λ), also referred to as dark energy and commonly interpreted as the Universe's vacuum energy, accounts for more than two thirds of this energy density and so-called cold dark matter (CDM) provides the bulk of the rest, visible matter (so-called baryonic matter) only accounts for $\sim 5\%$ of the total energy density in the Universe. This thesis discusses a few aspects of the late time evolution of self-gravitating objects. Within such objects, the relative density of dark energy is negligible. Between the two unknown components of the Universe, we will thus concentrate on dark matter (DM) for the remainder of this thesis.

There is a large body of evidence for the presence of DM in the Universe, although all of it remains purely gravitational to date. The presence of DM can, for example, be inferred from the observed anisotropies in the CMB (e.g. Planck Collaboration et al. 2020) or the lensed images of distant galaxies (e.g Tyson et al. 1998). Historically, however, most early evidence for DM originated from comparisons of the observed kinematics of stars and gas within galaxies and the dynamics of galaxies themselved to theoretical models that predict their expected dynamics from the observed mass of the galaxies or galaxy clusters to which they belong. The first such evidence dates back to Zwicky (1933), who discovered that the velocity dispersion of galaxies in the Coma cluster was much larger than was expected from the cluster's observed luminous mass. Since no isolated galaxies have been observed to move at speeds as large as the relative speeds between some galaxies in the Coma cluster, Zwicky deemed it likely that the galaxies in the Coma cluster were indeed gravitationally bound and concluded that the large velocity dispersion required the presence of a significant amount of missing matter (only later was it called dark matter). Further evidence for DM was discovered by investigating the rotation curves – the mean observed circular velocity of gas as a function of the radial distance from the center of the galaxy – of large spiral galaxies in the 1970's and 1980's. Freeman (1970) discovered that the rotation curves of M33 and NGC300 do not decline at large radii – contrary to what had been expected from their observed distributions of luminous matter. These observations suggested the presence of a substantial amount of non-luminous (dark) matter in those galaxies. Later, Ostriker et al. (1974) and Einasto et al. (1974) observed a linear increase of the enclosed mass of spiral galaxies, and Ostriker and Peebles (1973) gave a theoretical argument outlining why spheroidal haloes were required to stabilize galactic disks. Due to a multitude of flat (non-declining) rotation curves that were observed in large spiral galaxies (see e.g. Rubin et al. 1980, Bosma 1981a, Bosma 1981b, Rubin et al. 1985, and van Albada et al. 1985) the presence of smooth DM haloes around those galaxies soon became a widely accepted theory.

Initially, the DM was believed to consist of massive neutrinos (Szalay and Marx 1976, Cowsik and McClelland 1973). This was natural, as scientist knew of the existence of neutrinos, but their masses had yet to be measured[1]. When Lubimov et al. (1980) claimed to have measured a neutrino mass of $\sim 30\,\mathrm{eV}$ this idea gained further momentum, given that neutrinos of such mass implied a critical density large enough for the Universe to be closed (Cowsik and McClelland 1973). Lubimov's measurement was later shown to be wrong and the theory that the DM would consist mainly of such light particles was disproved by N-body simulations of structure formation not long after that (White et al. 1983). Subsequently, hypothetical particles which are significantly more massive than neutrinos and move significantly slower in the early Universe emerged as the preferred candidates, as the results of N-body simulations of structure formation within these so-called cold DM (CDM) models gave a much better fit to the large scale clustering of galaxies observed in the Universe today (Blumenthal et al. 1984). Several CDM particle candidates have been proposed, among them the lightest non standard model particle in supersymmetric theories (Ellis et al. 1984) and the QCD axion (Preskill et al. 1983), both of which are motivated by theoretical problems in particle physics. Despite extensive searches for the most promising particle candidates,

[1] This is still true today.

the particle nature of the CDM has not yet been uncovered (e.g. Bertone and Hooper 2018). Until that happens, CDM will remain a very well motivated – but unproven – model in which DM interacts only gravitationally. For that reason, alternative models in which DM is not dynamically cold and collisionless are currently an active and thriving field of research[2].

While CDM became the favoured paradigm in the early 80's, the first N-body simulations of structure formation worked within the so-called standard cold DM (SCDM) framework in which the Universe is flat and all of its energy density is provided by matter (see Frenk and White 2012 and references therein). The need for a non-zero cosmological constant was revealed later, through a combination of supernovae measurements (Riess et al. 1998,Perlmutter et al. 1999) which ruled out an Einstein-de-Sitter Universe at redshifts below $z = 1$, and the first measurements of the power spectrum of the cosmic microwave backgound's temperature anisotropies (de Bernardis et al. 2000,Hanany et al. 2000), which proved the Universe to be (approximately) spatially flat. Today, both the matter density parameter Ω_m and the energy density associated with the cosmological constant Ω_λ are known to unprecedented accuracy from precision measurements of the cosmic microwave background [3] (Planck Collaboration et al. 2020).

Within the ΛCDM model we can self-consistently describe the linear and non-linear evolution of initial perturbations in the density field consistent with the measured temperature anisotropies in the cosmic microwave background (CMB) into the large scale distribution of galaxies and galaxy clusters observed in galaxy redshift surveys. Any theory which aims to provide an alternative to ΛCDM must produce the same results on scales larger than ~ 1 Mpc. At smaller scales, where baryonic processes related to galaxy formation and evolution, become increasingly important, predictions from N-body simulations including only CDM become unreliable. It is at those scales that potential problems of ΛCDM arise from a comparison of the predictions of DM only (DMO) simulations to observations (Bullock and Boylan-Kolchin 2017). Many possible solutions have been proposed to these so-called small-scale challenges. Broadly, these can be divided into two different classes. In the first class of solutions, the physics of the DM differs significantly from CDM at scales smaller than ~ 1 Mpc. The construction of such exotic DM models is not trivial, as their predictions must be identical to those of CDM on large scales. Nonetheless, several interesting and feasible alternatives to CDM have been proposed in order to solve either one or all of CDM's small scale challenges. The most prominent of those include warm DM (WDM, Colombi et al. 1996, Bode et al. 2001), self-interacting DM (SIDM, Carlson et al. 1992, Spergel and Steinhardt 2000) and fuzzy DM (see Hui et al. 2017 for a review). The second class of solutions relies on non-gravitational, baryonic, processes – for instance feedback from supernovae – that change the distribution of gas in the interstellar medium. Such processes affect the DM indirectly through the gravitational coupling between DM and baryons. Therefore, they can alter the distribution of DM around galaxies, which can in turn explain the observed mismatch between DMO simulations and observations. However, as the

[2] It is worth mentioning that other theories, such as modified Newtonian dynamics (MOND, McGaugh 2015), seek to eliminate the need for DM altogether, but face severe theoretical and observational issues, both in the early and in the late Universe

[3] under the assumption of a ΛCDM cosmology

relevant physical processes – such as individual supernovae – occur on scales which are unresolved in large cosmological simulations, computational implementations of baryonic sub-grid processes typically rely on a number of effective parameters tuned to capture their net effect at the resolved scales.

As both exotic DM physics and baryonic feedback can potentially resolve the small-scale challenges of ΛCDM, it is difficult to differentiate between those two classes of models. Due to the importance of physical processes that occur below the resolution limit of (cosmological hydrodynamical) simulations, a definitive modeling of the exact impact of feedback processes in different cosmologies within large cosmological simulations unfortunately remains impossible for the time being. Thus, in light of the significant impact of the largely uncertain baryonic processes at scales below ~ 1 Mpc, identifying reliable observables that allow for a clear distinction between solutions to the small-scale challenges based on either exotic DM physics or baryonic feedback is a timely task for theoretical astrophysicists.

In this thesis, I develop a method to differentiate between mechanisms that modify the small-scale predictions of collisionless N-body simulations in a ΛCDM cosmology. The method relies on ensembles of so-called kinematic tracers, whose only relevant interaction is gravity and which react to an evolving gravitational potential. For many applications in astronomy, stars can be used as such tracers. I demonstrate that the evolution of the dynamical properties of such tracers depends on whether the potential evolves on timescales that are faster or slower than the typical dynamical times of those tracers, i.e. impulsively or adiabatically. Based on how much the conservation of radial actions of kinematic tracers is violated as the gravitational potential changes, I derive a clear-cut criterion that separates the adiabatic from the impulsive regime. I then show that impulsive changes of the potential induce characteristic features in the phase space distribution of ensembles of kinematic tracers which are not present for adiabatically changing potentials. Predicting such signatures is a theoretical achievement of its own, as they can be used to classify physical processes across dynamical scales and epochs. Here I show that since some promising alternatives to CDM with new DM physics change the gravitational potential on different timescales than baryonic feedback, we can tell them apart by the respective dynamical properties of tracers and in particular by whether or not signatures of an impulsively evolving potential are present.

The outline of this thesis is as follows. First, a review of linear and non-linear structure formation in ΛCDM is given, outlining how the initial matter perturbations are amplified and eventually collapse into virialized objects. Then I discuss how N-body simulations of CDM enable us to follow the non-linear evolution of structure in the Universe. Subsequently, the challenges ΛCDM faces on small scales, and how these challenges can possibly be resolved by taking into account baryonic feedback processes or exotic DM physics, are discussed. Afterwards, I outline the challenges that statistical physics faces when trying to model the evolution of dynamical systems in the presence of long range forces, followed by a short discussion of the key numerical methods used in this thesis and an introduction into integrals of motion, action angle coordinates, and phase mixing. The main body of this thesis consists of a series of research articles, both peer reviewed and in preparation. A brief discussion of the key results of each article is given and the articles themselves are included as individual chapters (5-9) at the end of the thesis. Finally, the key results of chapters 5-9 are summarized and possible future

projects are outlined. The appendix contains a brief description of the results of an article to which I contributed as a coauthor.

The research articles that are reproduced as chapters 5-9 are:

1. *The nature of core formation in dark matter haloes: adiabatic or impulsive?* by Jan D. Burger and Jesús Zavala, published in Monthly Notices of the Royal Astronomical Society, Volume 485, Issue 1,May 2019, Pages 1009-1028

2. *Conservation of radial actions in time-dependent spherical potentials* by Jan D. Burger, Jorge Peñarrubia, and Jesús Zavala, submitted to Monthly Notices of the Royal Astronomical Society

3. *SN-driven mechanism of cusp-core transformation: an appraisal* by Jan D. Burger and Jesús Zavala, in press at the Astrophysical Journal

4. *Degeneracies Between Self-interacting Dark Matter and Supernova Feedback as cusp-core transformation mechanisms* by Jan D. Burger, Jesús Zavala, Laura V. Sales, Mark Vogelsberger, Federico Marrinaci, and Paul Torrey, in preparation

5. *Kinematic signatures of impulsive supernova feedback in dwarf galaxies* by Jan D. Burger, Jesús Zavala, Laura V. Sales, Mark Vogelsberger, Federico Marrinaci, and Paul Torrey, in preparation

Chapter 2

Structure formation in ΛCDM

This chapter is dedicated to a brief review of structure formation in the ΛCDM model. In some instances, relevant deviations that arise in exotic DM models will be discussed for illustration.

2.1 Linear Growth of Perturbations

Almost immediately after the Big Bang, the Universe is believed to have undergone a period of accelerated expansion called inflation. Inflation is believed to be driven by an additional quantum field (Guth 1981, Linde 1982), producing nearly scale-invariant and adiabatic density perturbations within the early Universe (Guth and Pi 1982, Hawking 1982, Starobinsky 1982). The primordial power spectrum of scalar (matter) perturbations is commonly parametrized as

$$\mathscr{P}_0(k) = A_s \left(\frac{k}{k_0}\right)^{n_s}, \tag{2.1}$$

with a scalar tilt that has been determined from a combination of the CMB and baryon acoustic oscillations (BAOs) to be $n_s = 0.9665 \pm 0.0038$ (Planck Collaboration et al. 2020). At the time of last scattering, these initial density perturbations have grown to be of similar magnitude as the temperature anisotropies in the CMB, which are $\delta T/T \sim 10^{-5}$. However, the linear power spectrum at that time deviates significantly from being scale-invariant. The evolution of the power spectrum of CDM is affected by different physical processes than the evolution of the power spectrum of baryonic matter, and thus their power spectra at the time of recombination (or decoupling/ last scattering) are different. However, both the linear power spectra of CDM and baryons can be calculated using a linearized version of the Boltzmann equation that governs the evolution of the perturbations, i.e. with linear perturbation theory.

Before recombination, baryons are tightly coupled to radiation, and the radiation pressure can prevent the development of unstable modes, inducing acoustic oscillations. As the Universe cools down and photons begin to decouple, the average distance traveled between individual scatters of baryons and photons increases, giving rise to diffusion damping (also called Silk damping, Silk 1968) of small scale perturbations.

The efficiency of this mechanism, along with observations of galaxies at high redshifts, is another clear indication for the need for DM, i.e. a type of matter that does not couple to radiation and is thus not affected by collisional damping.

In the standard ΛCDM model, the gravitational potential after recombination is mainly sourced by the DM, and the baryons follow the DM by falling into the potential wells created by the growing DM perturbations. Before recombination, two physical mechanisms have a major impact on the growth of DM perturbations. The first one, the so-called Mészáros effect (Meszaros 1974), affects the evolution of sub-horizon scale perturbations during the epoch of radiation domination. While the expansion of the Universe is driven by radiation, the expansion rate is too fast to allow for the growth of perturbations in the matter density field. Superhorizon scales, however, are not affected by this mechanism and therefore the size of the comoving horizon at the time of matter-radiation equality marks a characteristic scale in the linear CDM power spectrum.

This causes a bend in the late time linear DM power spectrum, shown as a black line in figure 2.1. At larger scales, or smaller wavenumbers respectively, the dimensionless power spectrum, defined as $\Delta^2(k) = k^3 P_{\rm lin}(k)/2\pi$, is nearly scale invariant, $\Delta^2(k) \propto k^{3+n_s}$. At smaller scales, or larger wavenumbers, the slope of the power spectrum gradually decreases since the growth of perturbations on smaller scales has been halted for a longer time due to the Mészáros effect, given that they reenter the horizon earlier than perturbations on larger scales.

In the grey region in figure 2.1, DM is constrained to behave like CDM. These constraints come from the number counts of low mass galaxies at low redshift and observations of the Lyman alpha forest (Viel et al. (2013)) at high redshift. On smaller scales deviations from ΛCDM are harder to constrain (and probe), because baryonic effects can significantly alter the predictions of N-body simulations. Alternative theories that reduce the power on those scales can thus be feasible alternatives to CDM.

DM perturbations on small scales can be erased by two distinct mechanisms, collisional damping and collisionless damping, or free-streaming damping. The latter is the more known of the two, and is related to the time at which the DM particles become non-relativistic. Until matter becomes the dominant component of the Universe and density perturbations can grow, perturbations on scales smaller than $\lambda_{\rm fs}$ are erased due to the streaming of the DM particles. This so-called free streaming length,

$$\lambda_{\rm fs} = \frac{2ct_{\rm nr}}{a_{\rm nr}} \left[1 + \ln \frac{a_{\rm eq}}{a_{\rm nr}} \right], \tag{2.2}$$

depends on the time $t_{\rm nr}$ at which the DM particles become non-relativistic, the scale factor at that time, $a_{\rm nr}$, as well as the scale factor $a_{\rm eq}$ marking the epoch of matter-radiation equality. For collision-less DM, the free-streaming length scales roughly as $\lambda_{\rm fs} \propto m_\chi^{-1}$, and thus lighter DM particles are more strongly affected by free-streaming damping. DM particles with mass $m_\chi \sim 1\,\mathrm{keV}$ have $k_{\rm fs} \sim 3.8\,h/\mathrm{Mpc}$ and belong to a class of models called warm DM (WDM, see e.g. Zavala and Frenk 2019). As shown in red in figure 2.1, the damping of perturbations on small scales manifests itself in the linear power spectrum as a cutoff at large wavenumbers. WDM models with particle masses lighter than a few keV are ruled out by observations of the Lyman alpha forest (e.g. Viel et al. 2013). The free-streaming length directly corresponds to a lower limit on

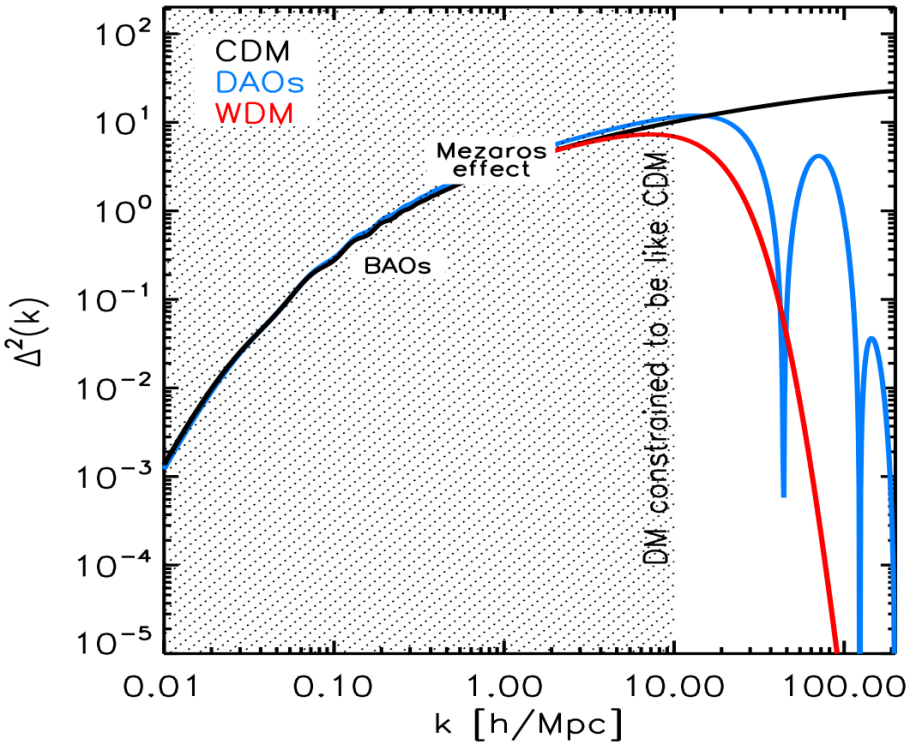

Figure 2.1. Late time dimensionless linear DM power spectrum for three different DM models, standard CDM, and two models with power suppression at small (sub-galactic) scales, warm DM (WDM) and DM matter which is coupled to a dark relativistic particle at early times, leading to dark acoustic oscillations (DAOs). Within the greyed in area, the linear DM power spectrum is constrained to look like that of CDM. Figure adopted from Zavala et al. (2019).

the size of gravitationally bound structures that can form from the collapse of growing DM density perturbations. For currently allowed WDM models, this scale corresponds roughly to the mass of dwarf galaxies, whereas for a commonly used DM candidate like the neutralino[1] the smallest gravitationally bound objects have roughly the mass of the Earth.

Collisional damping of DM occurs if the DM particles are coupled to a relativistic species at early times. This can either be a through a weak coupling between the DM and standard model photons or neutrinos (Boehm et al. 2002, Boehm et al. 2014), or through interactions between DM and so-called dark radiation, a relativistic species hidden in the dark sector (Buckley et al. 2014, Cyr-Racine et al. 2016). Supported by the relativistic species' radiation pressure, the DM density perturbations undergo dark

[1] In many supersymmetric models, the neutralino is the lightest stable particle and a so-called weakly interacting massing particle (WIMP), implying that it couples to the standard model via the weak force.

acoustic oscillations similar to those of the baryons at later times. The coupling between the DM particles and the relativistic species must be weaker than that between electrons and photons, since the DAOs can at most affect scales that correspond to $k > 10\,h^{-1}\mathrm{Mpc}$ (see blue line in figure 2.1). In such models, DM perturbations on very small scales are erased by collisional damping similar to Silk damping. In some regimes, the net effect of such DAOs on structure formation is similar to that of a simple WDM cutoff. However, in some cases DAOs can leave a distinct imprint in the matter distribution at late times (Bohr et al. 2020).

After photon-decoupling, the Universe enters the dark ages. During this epoch, the Universe is matter dominated and close to Einstein-de-Sitter. Radiation pressure and free-streaming are no longer relevant, and thus the perturbations grow, slowed down only by the Hubble drag. DM perturbations of all scales grow at the same rate during this epoch, $\delta(k,t) \propto D(t)/D(t_{\mathrm{dec}})\delta(k,t_{\mathrm{dec}})$, where t_{dec} is the photon decoupling time and the growth factor $D(t) \sim a(t)$ until the impact of the cosmological constant grows and causes the Universe's expansion to accelerate. The growth of the perturbations continues until their evolution becomes non-linear, i.e. when $\delta\rho/\rho \sim 1$. At that point, overdensities start to collapse into virialized, gravitationally bound structures with eventual overdensities of $\delta\rho/\rho \sim 178$, according to the spherical collapse model (Press and Schechter 1974). Clearly, the evolution of the DM distribution function is no longer reasonably predicted by linear perturbation theory at such overdensities. Thus, different methods are needed to describe the formation of gravitationally bound structures.

2.2 Cosmological N-body simulations

Several techniques have been developed to describe the non-linear evolution of DM overdensities and to follow their collapse into virialized structures. The mildly non-linear regime, i.e. the structures roughly corresponding to what is known as the cosmic web, can be explained fairly well by higher order perturbation theory, such as the *effective field theory of structure formation* (Baumann et al. 2012, Carrasco et al. 2012). For the collapse of individual structures, simplified theories based on assumptions of spherical collapse (Gunn and Gott 1972) or ellipsoidal collapse (Sheth et al. 2001) have been developed. They describe how overdensities grow, collapse, and eventually settle into virialized configurations called haloes. Under the assumption of ellipsoidal collapse, the so-called Press-Schechter formalism accurately (especially when working with the assumption of ellipsoidal collapse) predicts the abundance of virialized DM haloes as a function of their mass (the halo mass function), relying only on the additional assumption of a Gaussian initial distribution of perturbations (Sheth et al. 2001, Cooray and Sheth 2002). Since it is a fast and analytic model, Press-Schechter theory is still frequently used to predict halo mass functions of alternative DM models or the spatial clustering of galaxies.

N-body simulations, on the other hand, are used to self-consistently evolve the DM distribution function, starting from a time at which the perturbations were linear and up until today. The evolution of the standard CDM distribution function is governed by a system of coupled partial differential equations known as the Vlasov-Poisson equations. It consists of the collisionless Boltzmann equation, complemented by the

Poisson equation which states that the gravitational potential appearing in the Boltzmann equation is generated by the DM's self-gravity:

$$\frac{df}{dt} = \frac{\partial f}{\partial t} + \mathbf{v} \cdot \nabla f - \nabla \Phi \cdot \frac{\partial f}{\partial \mathbf{v}} = 0 \tag{2.3}$$

$$\nabla^2 \Phi(\mathbf{x}, t) = 4\pi G \int f(\mathbf{x}, \mathbf{v}, t)\, d^3\mathbf{v}. \tag{2.4}$$

Here f is the fine-grained distribution function of the DM. Instead of trying to solve the coupled equations 2.3 and 2.4 directly, cosmological N-body simulations follow the evolution of a set of N discrete simulation particles, which effectively represent the distribution function at randomly chosen phase space points. Equation 2.3 implies that the distribution function is conserved along the phase space flow, and thus a N-body representation of f at a subsequent time can simply be obtained by integrating the equations of motion of the N simulation particles. Alternatively, the discrete sampling of the fine-grained distribution function can be seen as an effective coarse-graining, and

$$\tilde{f}(\mathbf{x}, \mathbf{v}) = \sum_{i=1}^{N} m_i W(|\mathbf{x} - \mathbf{x}_i|; \varepsilon) \delta(\mathbf{v} - \mathbf{v}_i), \tag{2.5}$$

can be interpreted as the coarse-grained distribution function. In equation 2.5, m_i is the mass of an individual simulation particle, $\delta(\mathbf{v} - \mathbf{v}_i)$ is the Dirac delta function in velocity space, and W is a kernel function used to smooth interactions between simulation particles over a softening length ε. In cosmological N-body simulations, the evolution of this coarse-grained distribution function is calculated in a comoving reference frame, taking into account the expansion of the Universe through the scale factor which solves Friedmann's equations. Since for computational reasons the amount of simulation particles that can be used is orders of magnitude less than the amount of actual DM particles, the information one obtains from N-body simulations is always affected by resolution effects on small scales.

2.2.1 Initial conditions

For cosmological N-body simulations to predict a realistic distribution of structure in the Universe today, they need to start from realistic initial conditions, i.e. the initial set of N simulation particles must correspond to a real space representation of the linear DM power spectrum at the beginning of the simulation. Sampling of initial conditions for cosmological simulations is a two-step process. First, so-called pre-initial conditions are generated. Pre-initial conditions represent an entirely homogeneous Universe, with a mass density that corresponds to the Universe's critical density at the beginning of the simulation (see e.g. Sirko 2005). Thus, the N particles need to be uniformly distributed throughout the simulation box. In fact, it is important that these pre-initial conditions represent an entirely uniform distribution, simply drawing each particle's spatial position at random from a uniform distribution will introduce Poissonian shot noise and subsequent fast and artificial formation of structure (White 1994). There are two common ways of generating such entirely uniformly distributed particle configurations.

Suppose that the simulation volume in question is a periodic cubic box of side length L and the total number of simulation particles is N, and $N^{1/3}$ is a natural number. The easiest way to generate an entirely uniform particle distribution is to place the particles on a regular lattice, leading to a comoving pre-initial condition density

$$\bar{\rho}(\mathbf{x}) = m \sum_{i,j,k} \delta(x-(i+0.5)\Delta x)\delta(y-(j+0.5)\Delta x)\delta(z-(k+0.5)\Delta x), \tag{2.6}$$

where m is the mass of one simulation particle, (i,j,k) are integers running from 0 to $N^{1/3}$ and $\Delta x = L/N^{1/3}$ is the interparticle separation. One drawback of this approach is that it imprints Δx as a strong characteristic length into the simulations. Moreover, remnants of the grid pattern can skew the statistics of eventual low density regions, such as voids (White 1994).

An alternative to this regular grid is to generate what is called "glass" pre-inital conditions (White 1994, Baugh et al. 1995). To obtain those, the particle's positions within the periodic box are initially drawn at random from a uniform distribution. Subsequently, the system of particles is evolved with an N-body solver, using negative gravitational forces, until the particles reach an equilibrium configuration in which the net force on each particle is zero. While both of these methods produce sufficiently homogeneous pre-initial conditions, most modern simulations use glass pre-initial conditions as they are free of spurious grid patterns.

Once a homogeneous configuration of particles has been generated, the perturbations of the initial density field need to be taken into account. In order to do that, one samples the amplitude and the phase of the perturbations of individual modes in such a way that they are consistent with the linear power spectrum at the initial time of the simulation. The amplitude is drawn from a Rayleigh distribution with variance $\sigma^2 = L^3 P(k)/2$ and the phase is drawn at random from a uniform distribution (Sirko 2005). In this way, one obtains the field of overdensities in $\mathbf{k}$-space, $\hat{\delta}(\mathbf{k})$, which is then inverse Fourier transformed into the real space overdensity field $\delta(\mathbf{x})$. To obtain a particle representation of the perturbed density field, the Zel'dovich approximation (Zel'Dovich 1970, Efstathiou et al. 1985), or higher order Lagrangian perturbation theory (Jenkins 2010), are commonly used to calculate displacements and peculiar velocities of individual simulation particles. Using the Zel'dovich approximation, displacement and peculiar velocities are calculated from the density field as (Sirko 2005):

$$\mathbf{x} = \mathbf{q} + \Psi(\mathbf{q}) \tag{2.7}$$

$$\dot{\mathbf{x}} = \dot{\Psi}(\mathbf{q}) \tag{2.8}$$

where $\mathbf{q}$ is the particle's unperturbed Lagrangian coordinate and the gravitational field $-\nabla\cdot\Psi(\mathbf{q}) = \delta(\mathbf{q})$ and its time derivative is determined by the growth factor's evolution at the simulation's initial time[2] (or redshift, respectively), $\dot{\Psi} = \dot{D}(z)/D(0)\Psi$.

[2]The linear relation between the density field and the gravitational field therefore implies that the evolution of the gravitational field is also directly determined by the growth factor.

2.2.2 Force calculation and time integration

Once an accurate N-body representation of the power spectrum at a suitable initial time has been generated, the task of N-body simulations is to advance the state of the system in time until today, which means solving the system of coupled differential equations that emerges from the collisionless Boltzmann equation (2.3). In practice, this is done by calculating the phase space trajectory of each simulation particle over a series of appropriately chosen time steps. Since the equations of motion are linear, we can calculate each particle's acceleration at a fixed time by summing up the contributions from each of the other simulation particles. For large particle numbers, however, this direct summation approach adopted in early cosmological simulations (e.g. Frenk et al. 1983) is computationally exceedingly expensive, and more sophisticated methods to calculate the accelerations are needed. One way to reduce the number of computations is to calculate the smoothed density field on a pre-defined grid and then solve Poisson's equation in Fourier space (e.g. Melott 1983, Klypin and Shandarin 1983). Forces acting on individual particles are obtained by interpolation. This so-called particle mesh (PM) method is particularly advantageous for simulations in a periodic box, since this facilitates the evaluation of the Fourier transform. A potential drawback of the method is that the spatial resolution is inherently given by the mesh size. For very concentrated particle distributions, the algorithm is either increasingly inaccurate or computationally very expensive, as its complexity is $\mathcal{O}(n_{\rm grid}\log n_{\rm grid})$. This problem can be alleviated significantly by using an adaptive mesh (see Dehnen and Read 2011 for a review).

While the PM method described above solves the differential form of the Poisson equation, other methods have been developed to approximate the $\mathcal{O}(N^2)$ direct sum of Green's functions. Among the most commonly employed ones is the Barnes and Hut (1986) tree algorithm. Here, all the particles in the simulation are embedded in a hierarchical tree of cubic real space volumes. An example is a standard oct-tree, where a target number of particles per tree leave is given as an input parameter. Each cubic cell that contains more particles than this target number is then split into eight daughter cells, each of the same size and with side lengths equal to half that of the original cubic cell. The potential generated by each of the cubic cells is then calculated as a multipole expansion centered around the geometrical center of the cell – or its center of mass, in which case the dipole term vanishes. The expansion order can vary between codes, common choices are either a quadrupole expansion or a monopole expansion around the cell's center of mass. The power of the tree code comes from how the tree structure can be used to regulate the degree of accuracy with which gravitational interactions between particles are calculated, depending on their relative distances. The contribution of particles within a cubic cell C to the acceleration of a particle P is approximated as the multipole expansion of the potential in C if the distance between P and C's center of mass is a few times larger than the size of C (see Dehnen and Read 2011 for the full criterion). If this is not the case, the child cells of C are evaluated. The contributions from particles in the immediate vicinity of P are evaluated by direct summation. Note that the tree algorithm violates Newton's third law, since a multipole expansion is performed only on one end of the force calculation. A variation of the tree algorithm that circumvents this issue is the fast multipole method (FMM, e.g. Dehnen 2000, Dehnen 2002). Some modern codes use hybrid methods to calculate the particles'

accelerations. The TreePM method of GADGET-2 (Springel 2005), for instance, uses a particle mesh to evaluate the forces from distant particles and a tree code to add up the contributions from nearby particles.

One problem of cosmological N-body simulations of CDM is that the collisionless dynamics of a system of many actual DM particles is approximated by simulation particles which are fewer and heavier by many orders of magnitude. In a self-gravitating system of fixed total size, however, the relevant timescale for two body relaxation effects is $t_{\rm relax} \propto N/\ln N$ (Binney and Tremaine 2008). Thus, the pure dynamics of the N-body system may in fact not be collisionless and artificial two-body relaxation, such as large angle deflections and the formation of binaries, may occur. This can be prevented by introducing a gravitational softening scale to reduce the impact of close encounters and effectively render the system collisionless, which is done by modeling each simulation particle as an extended mass distribution. For instance, when using Plummer softening, the gravitational potential of each simulation particle becomes that of a Plummer sphere (Plummer 1911), yielding the total potential

$$\Phi(\mathbf{x}) = -G\sum_i \frac{m_i}{\sqrt{(\mathbf{x}-\mathbf{x}_i)^2+\varepsilon^2}}, \tag{2.9}$$

where ε is the softening length. In practice, Plummer softening is rarely used and one instead chooses a softening kernel for which the gravitational potential of each simulation particle becomes exactly Newtonian for $|\mathbf{x}-\mathbf{x}_i| > \varepsilon$ (Dehnen and Read 2011). The softening lengths of individual simulation particles are not always fixed to a single value throughout the simulation. Adaptive softening schemes are advantageous in simulations in which the magnitudes of the accelerations of simulation particles span several orders of magnitude. This includes simulations in which self-bound objects encompass substructure which can itself consist of gravitationally bound simulation particles. In such cases, choosing the wrong softening length may lead to undesired numerical effects, such as the artificial disruption of substructure (van den Bosch et al. 2018, van den Bosch and Ogiya 2018). A method to chose the numerical value of the softening length was put forward by Power et al. (2003). The authors argue that if one considers a self-gravitating, collisionless system of size R sampled with N particles, then the minimum gravitational softening length that should be used is $\varepsilon_{\rm min} \sim R/\sqrt{N}$, and the optimal value for the softening length is $\varepsilon = 4\varepsilon_{\rm min}$. For this choice of ε, moments of the distribution function that are measured on scales $\geq 3\varepsilon$ are typically converged.

After the accelerations of all particles have been calculated, the phase space positions of the particles are advanced in time, using their equations of motion. Since cosmological simulations aim to follow the evolution of the distribution function over about a Hubble time, accurate time integration is key to obtaining reliable results. Several time integration algorithms can be used, however, the one used in most simulations of collisionless dynamics is the so-called kick-drift-kick (KDK) leapfrog algorithm. For a fixed time-step, the KDK leapfrog is time-invariant and symplectic, meaning that total angular momentum and phase space volume are conserved (Saha and Tremaine 1992). In general, symplectic integrators are the method of choice when working with Hamiltonian systems, as they exactly solve an approximate Hamiltonian, implying that the energy associated with the integrated orbit does not exhibit a secular drift. Because of this property, the second order leapfrog outperforms even fourth order accu-

rate integrators when it comes to integrating the long term behaviour of Hamiltonian systems (see Springel 2016, in particular figure 2 therein). For a fixed time-step Δt, and a dissipationless differential equation of the form $\ddot{x} = f(x)$, the leapfrop integrator advances particles' phase space coordinates x and v between time-steps n and $n+1$ as follows (Springel 2016):

$$v_{n+\frac{1}{2}} = v_n + f(x_n)\frac{\Delta t}{2} \tag{2.10}$$

$$x_{n+1} = x_n + v_{n+\frac{1}{2}}\Delta t \tag{2.11}$$

$$v_{n+1} = v_{n+\frac{1}{2}} + f(x_{n+1})\frac{\Delta t}{2}. \tag{2.12}$$

However, the dynamic timescales of particles in cosmological N-body simulations can vary a lot between orbits, and thus a single fixed time-step hardly seems like the best choice for an efficient time integration scheme. Instead, cosmological N-body simulations use variable time-steps (Aarseth 2003), often based on the particles' accelerations. GADGET-2, for instance, uses a time-step (Springel 2005)

$$\Delta t = \sqrt{\frac{2\eta\varepsilon}{|\mathbf{a}|}}, \tag{2.13}$$

where ε is the gravitational softening length, $\mathbf{a}$ is the particle's acceleration, and η is a dimensionless accuracy parameter. Instead of assigning to each particle its own exact time-step, it is often preferable to restrict the simulation time-steps to values which are calculated from a base time-step Δt_0 as $\Delta t_n = \Delta t_0/2^n$ (Makino 1991) in order to assert that the simulation is fully synchronized after each base time-step. In such a scheme, each simulation particle is assigned to the time-step closest to the one obtained from equation 2.13. Unfortunately, the leapfrog integrator is no longer symplectic once variable time-steps are used. However, its accuracy hardly decreases if one ensures that the time integration scheme remains time-symmetric (Quinlan and Tremaine 1990, see also figure 2 of Dehnen and Read 2011).

2.3 Large scale structure and CDM haloes

2.3.1 Large scale structure

Perhaps the biggest success of the ΛCDM model for structure formation is how well the results of large cosmological N-body simulations agree with the large scale structure that is observed in galaxy redshift surveys (Springel et al. 2006). Starting from initial conditions consistent with the results of linear perturbation theory (see section 2.2.1) and then following the evolution of the distribution function until today with the methods outlined in section 2.2.2, N-body simulations provide a clear picture of how structure forms in a Universe with a given cosmogological model. Since the first N-body simulations of collisionless systems were performed (Davis et al. 1985, Frenk et al. 1988), the physics involved in the simulations, as well as the methods employed, have largely remained the same. However, the fast increase in computational power has

enabled computational astrophysicists to resolve the distribution function's evolution with ever higher precision. While early simulations followed the evolution of $\sim 10^4$ particles, the number of particles in the most expensive simulations to date is more than eight orders of magnitude larger than that (Potter et al. 2017, Ishiyama et al. 2020). The observed emergence of structure is similar in all CDM simulations and agrees remarkably well with the predictions of Zel'Dovich (1970). The collapse of growing DM perturbations happens along one spatial direction at a time. In consequence, the first structures to form are two-dimensional sheets called Zel'dovich pancakes. Collapse along a second spatial direction then leads to the formation of filaments. The final result of cosmological simulations is an interconnected web of filaments called the cosmic web. Along these filaments, collapse along the third spatial direction results in the formation of virialized objects which are decoupled from the Universe's expansion, so-called DM haloes. As a result of this gradual gravitational collapse, the clustering of matter today far exceeds the extrapolated linear power spectrum at redshift zero on nearly all scales. The distribution of visible matter within galaxy clusters and galaxies has been found to be biased relative to that of the DM. Several attempts have been made to model this so-called galaxy bias. One intuitive way is based on an extension of the Press and Schechter (1974) model of spherical collapse. In this so-called extended Press-Schechter theory (Bower 1991, Bond et al. 1991, Lacey and Cole 1993, Kauffmann and White 1993) one obtains the distribution of galaxies in two separate steps. First, the distribution of haloes as a function of their mass and phase space coordinates needs to be specified. This is usually measured from cosmological N-body simulations. One finds that the clustering of DM haloes, particularly heavy ones, is also biased high compared to the clustering of all DM particles in the simulation. This so-called halo bias can also be derived analytically from the original Press-Schechter theory (Cole and Kaiser 1989). All predictions of Press-Schechter and extended Press-Schechter theory improve significantly if ellipsoidal collapse (Sheth et al. 2001) is considered instead of spherical collapse.

Once the distribution of haloes is known, the galaxy bias follows if one has an adequate model for the distribution of galaxies as a function of halo mass. One way of obtaining such a model is to tune a simple statistical model such that it reproduces the total abundance and spatial clustering of galaxies (obtained from galaxy redshift surveys) from the abundance and spatial clustering of DM haloes obtained from an N-body simulation. Such a so-called Halo Occupation Distribution (HOD) model (Benson et al. 2000, Peacock and Smith 2000, Seljak 2000, Cooray and Sheth 2002, Smith et al. 2003) can provide a reasonable fit to observations. However, it can never match them perfectly as it fails to take into account that not only halo mass determines galaxy occupation – there is a residual dependence on intrinsic halo parameters such as spin, shape, or formation time, called assembly bias.

A natural way to tackle the challenges posed by this assembly bias is to go beyond a simple statistical matching between halo distributions and galaxy distributions and to instead directly model the baryonic processes driving the formation of galaxies at each step of the cosmological N-body simulation. Two different approaches to achieve such a modelling are currently being pursued. The most straight-forward approach is to include baryons into N-body simulations, along with models for all of the relevant

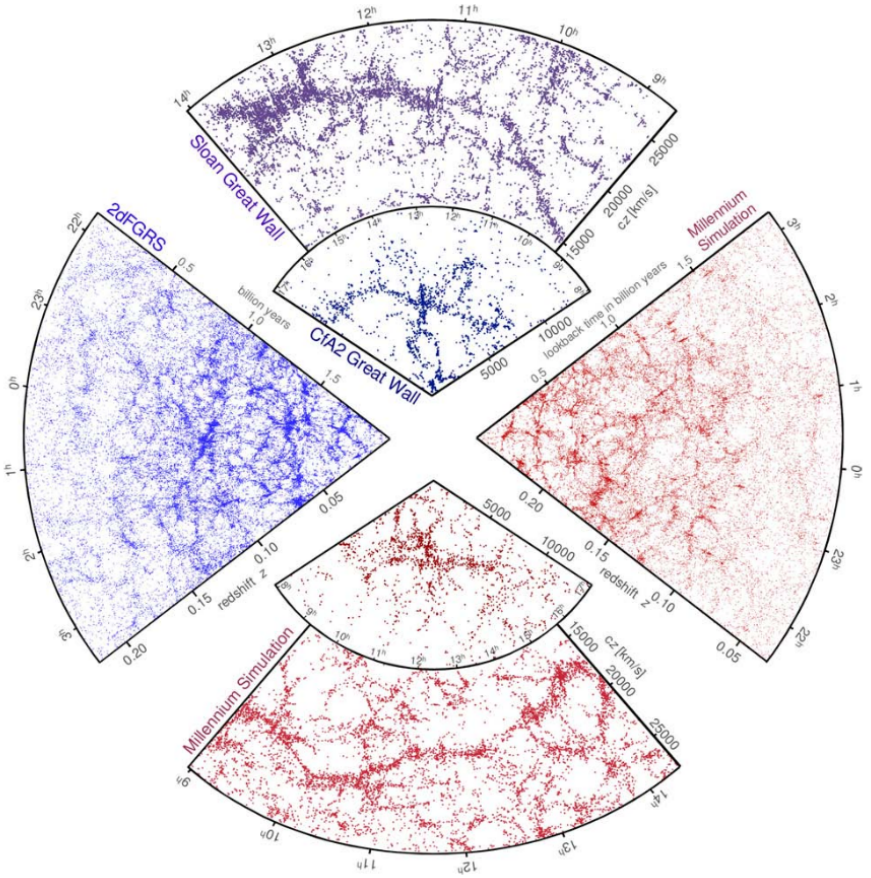

Figure 2.2. The large scale distribution of galaxies, either as measured in galaxy redshift surveys (left slice, two top slices) or as predicted by semi analytic models from the cosmological N-body Millenium simulation (Springel et al. 2005b)(two bottom slices, right slice). The scales in the simulation slices match those of the observations. The left slice shows one half of the results of the 2dFGRS (Colless et al. 2001), the small top slice is the result of the CfA2 survey (Geller and Huchra 1989), surrounding the Coma cluster at its center, and the large top slice shows the Sloan great wall as seen by SDSS (Gott et al. 2005). The mock surveys are constructed from the Millenium simulation and chosen such that their geometries and magnitude limits are similar to those of the real surveys. Adapted from Springel et al. 2006[3].

astrophysical processes that come along with them. Such hydrodynamic simulations are currently an active field of research in astrophysics, and they are becoming ever more sophisticated[4]. Yet, despite a fast increase in the resolution and the accuracy of hydrodynamic simulations, a full modelling of all the physical processes involved in galaxy formation remains computationally prohibitive, as the spatial resolution and the time resolution required to precisely follow all of these processes can currently not be achieved in cosmological simulations.

First introduced around 30 years ago (White and Frenk 1991), so-called semi analytic models are an alternative approach which has been extensively studied in the past. Semi analytic models calculate the formation of galaxies within DM haloes a postiori via a set of coupled differential equations modeling the impact of all the relevant baryonic processes during the haloes' assembly and merger history, obtained from a DMO cosmological simulation. These differential equations give a phenomenological description of gas cooling, star formation, feedback, and stellar evolution and are derived using results from either observations or simulations (Kauffmann et al. 1993, Cole et al. 1994). Applied to large cosmological CDM simulations, semi analytic models provide a remarkably accurate prediction of the statistical properties measured from galaxy redshift surveys (Kauffmann et al. 1997, Kauffmann et al. 1999, Benson et al. 2000, Springel et al. 2001). Figure 2.2 demonstrates just how well semi analytic models work in predicting the large scale abundance of galaxies. It shows a comparison of mock galaxy catalogs, obtained by applying a semi analytic model to the results of the Millenium simulation, to real galaxy catalogs obtained from galaxy redshift surveys. The match between the predicted and the observed large scale structure (LSS) is readily apparent. It should be remembered that this success, coming from the combination of ΛCDM and the phenomenological semi analytic models, is not trivially shared by other cosmological models, and that any alternative theory of structure formation is constrained to behave like ΛCDM on large scales.

2.3.2 DM haloes in ΛCDM

Semi analytic models solve for the formation of galaxies within virialized DM haloes. Indeed, since haloes are self-gravitating and decoupled from the expansion of the Universe, they provide the gravitational pull for gas to accumulate, cool, condense, and eventually form stars and thus galaxies. In this picture of galaxy formation, all visible galaxies are embedded within a DM host halo. The additional mass provided by the DM can explain the multitude of flat rotation curves (see chapter 1), as well as Zwicky's observation of the large velocity dispersion of galaxies in the Coma cluster. The statement that any competing theory of structure formation must match the predictions of ΛCDM on scales larger than $\sim 1\,\mathrm{Mpc}$ implies that the abundance and the clustering of large haloes is well constrained by observations. To put ΛCDM to the test is means focusing on its predictions on scales below 1 Mpc. Key small-scale predictions of ΛCDM that are obtained from cosmological simulations are the relative abundance of low mass haloes, i.e. the low-mass end of the halo mass function, and the inner structure

[3] Reprinted by permission from Springer Nature Customer Service Centre GmbH: Springer Nature; Nature; The large-scale structure of the Universe. Springel V., Frenk C. & White S, ©2006

[4] see Vogelsberger et al. (2020) and references therein for a current review

of DM haloes.

The halo mass function is easily obtained from N-body simulations, provided one uses a fixed definition of a DM halo. Several definitions are used throughout the literature (Davis et al. 1985, White 2001, Cuesta et al. 2008). One definition is based on the friends-of-friends algorithm used to identify haloes (Davis et al. 1985). In short, two particles are linked together if their spatial distance is shorter than a predefined linking length. This two-particle structure grows if another particle's distance to one of the two fulfills the above requirement and so on. Sets of particles which are interconnected in such a way are called haloes, and their collective mass is the friends-of-friends halo mass. Alternatively, the halo mass can be defined as the mass contained within a spherical region in which the mean density is a factor $\Delta(z)$ larger than the critical density of the Universe at redshift z (Cole and Lacey 1996). The sphere is centered around the most bound particle and the overdensity $\Delta(z)$ is related to the spherical collapse model. In a general cosmology the overdensity varies weakly with redshift (Eke et al. 1996, Bryan and Norman 1998), but in an Einstein-de-Sitter Universe $\Delta(z) = 178$, independent of redshift. Finally, using $\Delta = 200$ gives the virial mass M_{200}. Particles whose distances to the most bound particle are smaller than the associated radius, r_{200}, are typically in dynamical equilibrium, while particles further away from the halo's center are in a region which is still collapsing (Cole and Lacey 1996). Due to the simplicity of its definition, the virial mass is the halo mass which is most widely used in the literature.

With a fixed definition of a halo, the halo mass function is obtained from any large volume cosmological simulation simply by counting the abundance of haloes as a function of halo mass. The shape of the halo mass function, consistently obtained from several N-body simulations over the last decades (Boylan-Kolchin et al. 2009, Jenkins et al. 2001, Warren et al. 2006), Lukić et al. 2007, Tinker et al. 2008, Trujillo-Gomez et al. 2011, Hellwing et al. 2016, Ishiyama et al. 2020) is now well known for halo masses between $10^8\,\mathrm{M}_\odot$ and $10^{15}\,\mathrm{M}_\odot$. The ΛCDM halo mass function is a featureless power law at the low mass end,

$$\frac{dn}{dM} \propto M^{-1.9}, \tag{2.14}$$

with an exponential cutoff at large halo masses and a normalization that depends on the average density in the environment in which the halo mass function is measured (Crain et al. 2009). Below masses of $10^8\,\mathrm{M}_\odot$, the construction of the halo mass function is limited by the resolution of large cosmological simulations. However, theoretical arguments based on Press and Schechter (1974) and extended Press-Schechter (Bond et al. 1991, Bower 1991) theory suggest that the power law shape of the halo mass function will continue down to a cutoff mass which is related to the cutoff scale in the linear power spectrum (see figure 2.1) and thus depends on the particle nature of the DM. For a canonical WIMP CDM candidate, this cutoff mass is of the order of $10^{-6}\,\mathrm{M}_\odot$ (Bringmann 2009). The WDM cutoff in the power spectrum shown in figure 2.1, however, would correspond to a much larger cutoff mass (larger than $10^8\,\mathrm{M}_\odot$). The number of formed WDM haloes with masses which are up to $\sim$3 orders of magnitude larger than the cutoff mass is significantly reduced compared to the number of CDM haloes with similar masses (see e.g. Angulo et al. 2013, Lovell et al. 2014, Schneider et al. 2013, Bose et al. 2016). Thus, WDM models can be constrained, for instance, by

counting the number of satellites in the MW (e.g. Lovell et al. 2014). Other experiments designed to test WDM theories aim to indirectly detect low mass haloes devoid of gas and stars, e.g. by searching for gaps in tidal streams or flux anomalies in lensing signals.

A key result of cosmological N-body simulations of a ΛCDM Universe is the remarkable regularity of the inner structure of DM haloes over several orders of magnitude in halo mass. In particular, the spherically averaged density profile of CDM haloes is well described by a single two-parameter function[5] over more than 20 orders of magnitude in halo mass (Wang et al. 2020). This functional form of this so-called Navarro-Frenk-White, or NFW, profile is (Navarro et al. 1996b, Navarro et al. 1997),

$$\rho(r) = \frac{\rho_s r_s^3}{r\,(r+r_s)^2}. \tag{2.15}$$

According to equation 2.15, the halo's density asymptotically increases as r^{-1} towards the center of the halo, and bends to ever steeper slopes in the outskirts of the halo, with an asymptotic behaviour of $\rho \propto r^{-3}$ at large radii. The two intrinsic parameters of equation 2.15, ρ_s and r_s, can be expressed in terms of the halo mass and the so-called halo concentration, $c_{200} = r_{200}/r_s$. The scale radius then simply becomes $r_s = r_{200}/c_{200}$ and the scale density is fixed by the requirement that $M(< r_{200}) = M_{200}$, i.e. the definition of the virial mass. Integrating over the NFW density profile, and using the definition of the virial mass, then allows us to write

$$\rho_s = \delta_c \rho_{\rm crit}, \tag{2.16}$$

$$\delta_c = \frac{200\,c_{200}^3}{3\left(\ln(1+c_{200}) - \frac{c_{200}}{1+c_{200}}\right)}, \tag{2.17}$$

where $\rho_{\rm crit}$ is the critical density of the Universe. The halo's mass density profile is thus fully defined by the halo's virial mass[6] and its concentration. The concentration of a halo, in turn, depends on the average matter density in the Universe at the halo's assembly time. Since in ΛCDM there is no significant free streaming cutoff, the linear power spectrum of ΛCDM rises monotonically towards small scales (see figure 2.1). This leads to a hierarchical growth of structure, i.e. small haloes collapse first, at times when the matter density in the Universe is larger compared to when larger haloes collapse. Moreover, haloes form inside-out, meaning that the inner cusp is built at collapse, while the outer material is slowly accreted later on (Wang et al. 2011). Hence, the density in the halo's inner cusp is correlated with the halo's formation time, which is in turn correlated with the halo's virial mass today. In general, the density at a given radial distance r from the halo's center traces the average density of the Universe at the time at which material at this radius has been accreted on average (Ludlow et al. 2014). As a result of this assembly history of haloes in ΛCDM, less massive haloes are on average more concentrated than more massive haloes. The emerging correlation between halo mass and halo concentration is called the mass-concentration relation

[5] It is now known that a three-parameter Einasto profile (Navarro et al. 2010) provides a somewhat better fit. However, the accuracy with which the two-parameter NFW fit describes density profiles of haloes over several orders of magnitude in halo mass remains remarkable.

[6] Note the one-to-one correspondence between virial mass and virial radius.

(Navarro et al. 1997, Bullock et al. 2001, Eke et al. 2001, Wechsler et al. 2002, Neto et al. 2007, Zhao et al. 2009, Prada et al. 2012, Ludlow et al. 2014, Sánchez-Conde and Prada 2014, Diemer and Kravtsov 2015, Klypin et al. 2016, Pilipenko et al. 2017). The flattening of the linear power spectrum at small wavelengths causes a flattening of the concentration-mass relation at low halo masses. In addition, differences in the merger histories of haloes, along with an intrinsic scatter in the formation times of haloes with similar masses, induce a scatter in the concentration-mass relation.

The success of the NFW profile in describing the spherically averaged density profiles of DM haloes over many orders of magnitude in halo mass, combined with the concentration-mass relation, implies that the inner density profile of haloes is, to first order, characterized by a single parameter, the halo's virial mass M_{200}. This stunning universality of the inner density profile of CDM haloes, albeit empirically firmly established, is not well understood from fundamental principles. Pontzen and Governato (2013) have shown that the outer shape of DM haloes can be explained by methods of statistical physics. In particular, they show that when postulating that the sum of the radial actions of the DM particles in a halo is a conserved quantity, maximizing the halo's entropy results in a distribution function that closely matches that of an NFW halo. However, the density profile derived by Pontzen and Governato (2013) is underdense in the center of the halo and thus predicts CDM haloes to be cored instead of cuspy. The authors reconcile this mismatch by including a second population of particles (the cusp-particles), which is not in dynamical equilibrium with the rest of the halo. They justify this by noting that the dynamical timescales in the center of DM haloes are much shorter than in the outer region, and that for that reason dynamical mixing between those two regions has not yet occurred. In consequence, this implies that cusps must form at the time at which the halo collapses, and they must remain until redshift zero. This picture is essentially supported by simulations, although the haloes forming first in WIMP CDM appear to have a steeper cusp early on, $\rho \propto r^{-1.5}$, which is subsequently flattened until the ubiquitous r^{-1} slope is reached (Anderhalden and Diemand 2013, Ishiyama 2014, Delos et al. 2018).

DM particles within a relaxed halo are in dynamical equilibrium and obey the Vlasov-Poisson equations 2.3 and 2.4. If one assumes, to first order, that haloes are spherical and isotropic, both in position space and in velocity space, then the distribution function depends only on energy and is given by Eddington (1916)'s formula:

$$f(\mathscr{E}) = \frac{1}{\sqrt{8}\pi^2}\left\{\frac{1}{\sqrt{\mathscr{E}}}\left[\frac{d\rho}{d\Psi}\right]_{\Psi=0} + \int_0^{\mathscr{E}} \frac{d\Psi}{\sqrt{\mathscr{E}-\Psi}}\frac{d^2\rho}{d\Psi^2}\right\}, \tag{2.18}$$

which can be written as

$$f(\mathscr{E}) = \frac{1}{\sqrt{8}\pi^2}\int_0^{\sqrt{\mathscr{E}}} du \frac{d^2\rho}{d\Psi^2}(r(\Psi(u))) \tag{2.19}$$

if the first term in the brackets can be neglected. In equation 2.19, $u = \sqrt{\mathscr{E}-\Psi}$, where $\mathscr{E}$ is the negative specific energy and $\Psi(r)$ is the negative specific potential at the radius

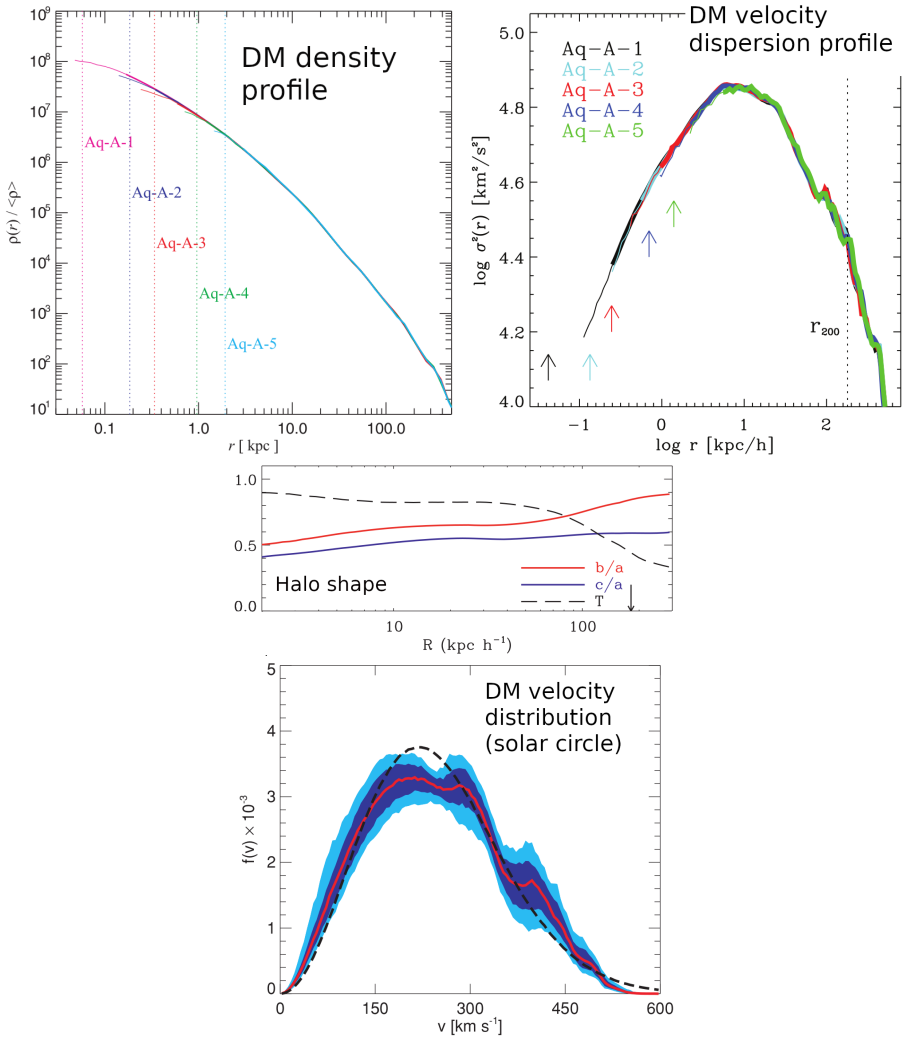

Figure 2.3. Structural properties of a cosmological Milky Way like DM halo from the Aquarius simulation[7]. The top left panel shows the spherically averaged density profile (adapted from Springel et al. 2008) and the top right panel the velocity dispersion profile (adapted from Navarro et al. 2010) of the Aquarius A halo (Springel et al. 2008) at different resolution levels. The middle panels shows shows the halo's axis ratios and the triaxiality of its moment of inertia tensor (Vera-Ciro et al. 2011), providing a measure of the halo's shape, and the bottom panel shows the distribution of velocities at the solar circle, i.e. in the radial range 2 kpc - 9 kpc (Vogelsberger et al. 2009).

r. For a halo with an NFW density profile[8] the resulting velocity distribution is not Maxwellian[9], but is instead a function of the radius. Equation 2.19 is particularly convenient if one wishes to numerically construct spherical systems in dynamical equilibrium, since the integrand can easily be expanded, i.e.

$$\frac{d^2\rho}{d\Psi^2} = -\frac{r^4}{(GM(<r))^2}\frac{d^2\rho}{dr^2} + \frac{2r^3M(<r) - 4\pi r^6\rho(r)}{(GM(<r))^2M(<r)}\frac{d\rho}{dr}. \tag{2.20}$$

If one can numerically invert the function $\Psi(r)$, calculating the distribution function reduces to numerically evaluating integral 2.19 on a grid of energies. At a given radius, velocities can then be drawn from the obtained distribution function using a rejection sampling algorithm.

In general, the density and velocity dispersion profiles of spherically symmetric, self-gravitating, collisionless haloes in dynamical equilibrium obey the Jeans equation (Binney and Tremaine 2008)

$$\frac{d(\rho\sigma_r^2)}{dr} + 2\frac{\beta}{r}\rho\sigma_r^2 = -\rho\frac{d\Phi}{dr} \tag{2.21}$$

$$\beta = 1 - \frac{\sigma_\theta^2 + \sigma_\phi^2}{\sigma_r^2}, \tag{2.22}$$

where β quantifies the anisotropy in the velocity dispersion tensor and the gravitational potential Φ is generated by the density distribution ρ. In the case of $\beta = 0$, equation 2.19 can be used instead. The velocity dispersion tensor of cosmological DM haloes, however, can be approximated as isotropic only in the halo's center. Towards the outskirts, they tend to be radially anisotropic (Navarro et al. 2010, Ludlow et al. 2011), and there is a tight correlation between the slope of the density profile and the anisotropy parameter β (Hansen and Moore 2006). The link between the phase space density and the velocity dispersion tensor is even more evident in the quantity ρ/σ^3, where $\sigma^2 = \sigma_r^2 + \sigma_\theta^2 + \sigma_\phi^2$. This so-called pseudo-phase-space density obeys an almost perfect power law in radius, $\rho/\sigma^3 \propto r^{-1.875}$ over many orders of magnitude in radius (Navarro et al. 2010, Taylor and Navarro 2001). Notably, the power law behaviour of ρ/σ^3 agrees exactly with the self-similar solution for infall onto a point mass in an EdS Universe (Bertschinger 1985). The remarkable regularity of the pseudo-phase-space density across all cosmological CDM haloes further highlights their universal structure, both in position space and in velocity space.

[7] The original figures are:
Fig. 4 of The Aquarius Project: the subhaloes of galactic haloes, Springel, V. et al., MNRAS 391
Fig. 10 of The diversity and similarity of simulated cold dark matter haloes, Navarro, J. et al., MNRAS 402
Fig. 3 of The shape of dark matter haloes in the Aquarius simulations, Vera-Ciro, C. et al., MNRAS 416
Fig. 2 of Phase space structure in the local dark matter distribution and its signature in direct detection experiments, Vogelsberger, M. et al., MNRAS 395

[8] Equation 2.19 is easily derived for all systems that are fully isotropic, but can only be solved if the system has a finite mass. For that reason, the distribution function for an analytic NFW halo can only be obtained once a reasonably smooth cutoff of the density profile is implemented for radii beyond the virial radius. For a certain number of potential-density pairs, the solution is analytic. Examples are, for instance, a Hernquist sphere or a Plummer sphere.

[9] The only density profile with an associated Maxwellian velocity distribution is that of the singular isothermal sphere

Figure 2.3 displays the different structural characteristics discussed above for a cosmological MW-like halo from the Aquarius project ($M \sim 1.8 \times 10^{12}\,\mathrm{M}_{\odot}$, $r_{200} \sim 250\,\mathrm{kpc}$). In the top left panel, the spherically averaged density profile is shown at several resolution levels. Down to the respective resolution limit, they are in excellent agreement with a NFW profile. The top right panel shows the corresponding velocity dispersion profile, whereas the bottom panel demonstrates that the halo's velocity distribution is not Maxwellian, using the velocity distribution of DM particles around the solar circle as an example.

While the anisotropy of the velocity dispersion tensor is highlighted in the bottom panel, the middle panel demonstrates that cosmological haloes are, in general, not spherically symmetric either. Shown are the axis ratios of the halo's moment of intertia tensor as a function of radius, with respect to the principal axis a. Furthermore, the triaxiality, defined as $T = (a^2 - b^2)/(a^2 - c^2)$ is shown, also as a function of radius. The trends observed here are common among cosmological CDM haloes. Haloes are triaxial rather than spherically symmetric (Frenk et al. 1988). Moreover, they are usually prolate close to the center and oblate in the outskirts (Frenk et al. 1988, Vera-Ciro et al. 2011, Jing and Suto 2002, Hayashi et al. 2007), where late accretion leads to a less anisotropic matter distribution (Vera-Ciro et al. 2011, Ganeshaiah Veena et al. 2018). Haloes which formed earlier, i.e. smaller haloes, tend to be closer to spherical symmetry (Bonamigo et al. 2015, Vega-Ferrero et al. 2017), while more recently formed larger haloes are more triaxial, retaining a memory of their recent collapse along a DM filament (Despali et al. 2014).

ΛCDM is by far the most studied cosmological model in the literature. However, it is important to remember that all the properties of DM haloes summarized in this subsection have been obtained from N-body simulations of CDM. Neither the number density of small haloes nor the inner dynamical structure of haloes are well constrained by observations. It is on those small scales that alternative DM theories are allowed to deviate from the behaviour of CDM. In some cases, alternatives to CDM have been specifically crafted to solve a supposed mismatch between the ΛCDM theory and observations. However, since galaxies are embedded within the centers of DM haloes, baryonic physics can also have a sizeable impact on the small-scale properties of DM haloes. Some of these small-scale problems, and how they can potentially be reconciled with the success of CDM on large scales, will be a subject of the next chapter.

Chapter 3

ΛCDM on small scales

As outlined in chapter 2, ΛCDM has become the standard cosmological model mainly due to its simplicity and its success in explaining the LSS of the Universe, namely the remarkable agreement of the observed matter power spectrum at wavenumbers $k \leq 1\,\mathrm{Mpc}^{-1}$ with the power spectrum inferred from CDM cosmological simulations. Yet, we have seen in section 2.3.2 that CDM cosmological simulations make specific predictions for the number count and the internal structure of virialized objects on smaller scales.

Structural properties of CDM haloes, such as its inner density profile, can be compared to observations only indirectly by modeling how DM is distributed within observed galaxies. Given the remarkable success of ΛCDM on large scales, current tests of the model focus mainly on potential mismatches between observations of dwarf galaxies and predictions from DMO cosmological simulations. Several so-called small-scale challenges to ΛCDM have emerged over the last decades, focusing either on mismatches between the abundance of low mass haloes and low mass galaxies or on kinematic observations that appear inconsistent with the ubiquitous NFW density profile (equation 2.15) of DM haloes established by cosmological simulations.

Since dwarf galaxies are rather faint, they are primarily observed in our immediate neighbourhood, most of them being so-called satellite galaxies of either Andromeda or the MW. In the ΛCDM picture, these galaxies are at the center of small DM haloes orbiting within the larger host halo, so-called subhaloes. The existence of subhaloes is well established from cosmological simulations. However, from studies of individual systems, we know that subhaloes are affected by gravitational effects which arise from interactions between the host halo and the subhaloes, such as dynamical friction and tidal stripping. It is thus crucial to be aware of these effects and of how they alter the distribution function of subhaloes when comparing observations to results of CDM cosmological simulations, lest we mistake them for a shortcoming of the ΛCDM paradigm.

In this chapter, I will briefly review subhaloes, both how many of them we can expect to find within a halo of a given mass, and how their structure is altered through gravitational interactions with the host halo and the host galaxy. Then I will discuss the small-scale challenges that ΛCDM faces, focusing in detail on the two problems

that are most relevant for this thesis, the too-big-to-fail problem and the cusp-core problem. Finally, I will discuss efforts that are being made to understand the small scale predictions of N-body simulations from fundamental principles. Such efforts aim to analytically model the long-term evolution of self-gravitating systems, using, for instance, tools of Hamiltonian perturbation theory or statistical physics.

3.1 Properties of subhaloes

In the hierarchical structure formation scenario of ΛCDM, haloes grow through mergers and accretion of both diffuse material and smaller haloes. As a consequence thereof, the smaller accreted haloes become subhaloes, moving on bound orbits around the center of the larger host. While orbiting in the host, environmental effects such as tides and dynamical friction strip material from the subhalo and change its mass profile, before eventually disrupting it entirely. This section briefly discusses the most relevant properties of subhaloes, i.e. their abundance and their inner structure, as well as the environmental effects responsible for changing their inner structure.

3.1.1 Definition of a subhalo

Since the outer structure of a subhalo is altered through gravitational interaction with the host halo, defining a subhalo is anything but trivial. Evidently the friends-of-friends mass is not a reasonable measure due to the halo being submerged in a larger host. Standard overdensity definitions also suffer from this limitation. Moreover, since some of the outer structure of the halo can be stripped away, the kinematics in the center of the halo are not necessarily linked to the inferred overdensity mass. Most definitions of subhaloes attempt to circumvent these issues by focusing either on properties of the halo before accretion or on properties that are less affected by environmental effects, while most halo finders look for local density maxima and then define the subhalo as the set of self-bound particles in the vicinity (e.g. SUBFIND, Springel et al. 2001, or the Amiga Halo Finder, Knollmann and Knebe 2009). To tag subhaloes, it is common to use either the peak mass of the halo before accretion, or the maximal circular velocity. The latter is related to the enclosed mass in the center of the halo and is thus less affected by tidal stripping. However, to some degree, tidal evolution affects even the maximal circular velocity v_{max} (D'Onghia et al. 2010). It may be preferable to use M_{peak} as a tag for subhaloes, especially when aiming to establish an unbiased relation between the total baryonic mass of a satellite galaxy and the halo mass of the surrounding dark subhalo. This is due to the fact that the baryonic mass is mostly concentrated to the center of the subhalo and is thus significantly less affected by environmental effects than the halo as a whole.

3.1.2 Subhalo abundance

In line with either of the two ways of tagging subhaloes outlined in section 3.1.1, the abundance of subhaloes within a host halo of a given mass can either be written as a subhalo mass function or a subhalo velocity function. In the low mass range, the subhalo mass function has an almost universal shape (i.e. independent of the host halo)

and a slope which is very similar to that of the global mass function of isolated haloes (Springel et al. 2008, Gao et al. 2004, Diemand et al. 2007, Gao et al. 2012, Garrison-Kimmel et al. 2014, Griffen et al. 2016). A natural cutoff of the subhalo mass function at large masses is given by the mass of the host halo. Moreover, the overall normalization of the subhalo mass function is also a function of the host halo mass, with larger host haloes hosting more subhaloes (Wang et al. 2011, Gao et al. 2012, Gao et al. 2011). As smaller haloes form earlier, this dependence of the normalization on halo mass arises because environmental effects have acted for a longer time on substructure within these smaller haloes, and thus have disrupted a larger fraction of the subhaloes compared to larger host haloes. The dependence of the subhalo abundance on the host halo mass can almost be fully captured by writing the subhalo mass function as a function of the mass fraction $\mu = M_{\rm sub}/M_{\rm host}$. For the total number of subhaloes, one obtains (Gao et al. 2011, Angulo et al. 2009, Boylan-Kolchin et al. 2010, Rodríguez-Puebla et al. 2016):

$$N(>\mu) = \left(\frac{\mu}{\tilde{\mu}_1}\right)^{1+\alpha} \exp\left[-\left(\frac{\mu}{\mu_{\rm cut}}\right)^{b}\right] \tag{3.1}$$

In equation 3.1, $\tilde{\mu}_1$ gives the typical mass fraction of the heaviest subhalo, which is about 0.01, b is a fitting parameter of order one (Boylan-Kolchin et al. 2010), and $\mu_{\rm cut}$ parametrizes the cutoff at large subhalo masses. $\tilde{\mu}_1$ has some small residual dependence on host halo mass, which is better captured by the subhalo velocity function. In terms of $x = v_{\rm sub}/v_{\rm h}$, the fraction of the maximal circular velocities (Klypin et al. 2011):

$$N(>x) \propto v_{\rm h}^{1/2} x^{-3}, x < 0.7 \tag{3.2}$$

In equation 3.2, the explicit dependence on the circular velocity of the host halo parametrizes this deviation from universality. Both the subhalo mass function and the subhalo velocity function arise as a result of the bottom-up structure formation in ΛCDM and the gravitational interaction between the host halo and its subhaloes. These environmental effects can change the inner structure of subhaloes and with time disrupt small subhaloes altogether.

3.1.3 Tidal stripping, tidal heating, and dynamical friction

Three main gravitational, environmental effects affect subhaloes orbiting within a larger host. Two of them are collisionless effects that occur because subhaloes are extended objects. The gravitational potential generated by the host halo and the host galaxy may have an appreciable gradient over the spatial extent of the subhalo. This is produces tidal forces, since one side of the subhalo experiences a stronger gravitational acceleration than the other. Depending on how strong this gradient is, and for how long the tidal forces act on the subhalo, the effect can either be approximated as continuous (tidal stripping) or as close to instantaneous (tidal heating). Dynamical friction, on the other hand, is not a mean field effect. Instead, it arises due to localized gravitational scattering. Here, the mass of the subhalo is the most important parameter. While orbiting through the host halo, the subhalo deflects diffuse DM particles, creating a wake of particles that follow its trajectory. This overdensity which lies opposite to the subhalo's direction of motion then creates a drag force, slowing down the subhalo on its orbit and eventually causing it to sink towards the center of the host halo.

Our current understanding of dynamical friction is in large part due to the work of Chandrasekhar (1943). Under a couple of simplifying assumptions, Chandrasekhar therein calculates the drag force acting on a compact object of mass M_S that moves with a velocity $\vec{v}_S$ through an ambient medium with constant homogeneous density ρ_m,

$$\vec{F}_{\rm df} = -4\pi \left(\frac{GM_S}{v_S}\right)^2 \ln\Lambda\, \rho(<v_S) \frac{\vec{v}_S}{v_S}. \tag{3.3}$$

In equation 3.3, $\ln\Lambda = \ln\left(1+(b_{\rm max}/b_{90})^2\right)$ is the Coulomb logarithm, with $b_{90} = G(M_S+m)/v_\infty^2$, where m is the mass of one particle in the ambient medium and v_∞ is the relative speed between this particle and the subject when the particle is at an infinite distance from the subject, i.e. the initial relative speed. $b_{\rm max}$, on the other hand, is the maximal impact parameter, and usually $b_{\rm max} \gg b_{90}$. $\rho(<v_S)$ denotes the mass density of particles in the ambient medium with a speed below v_S. The derivation of equation 3.3 assumes that the subject is a point mass (spatially non-extended) and the ambient medium is infinite, homogeneous, and consists of point particles with mass m. If the distribution of velocities in the ambient medium is Maxwellian, two limiting cases emerge from equation 3.3. If the subject moves at a speed which is smaller than the average speed of a particle in the ambient medium, $|\vec{F}_{\rm df}| \propto v_S$, i.e. the drag force increases proportional to subjects speed, similar to a standard friction. For large subject speeds, however, $\rho(<v_S) \approx \rho_m$ and thus $|\vec{F}_{\rm df}| \propto v_S^{-2}$, i.e. the dynamical friction force decreases at increasing subject speeds. Given the large number of assumptions entering into the derivation of equation 3.3, it is perhaps surprising that it provides a reasonable description of dynamical friction in more realistic scenarios as well. This is particularly true if one allows for modifications to the theory that take into account the extended density profile of subhaloes (see Zavala and Frenk 2019). The net effect dynamical friction has on subhaloes is that it causes them to sink towards the center of the host. The timescale for a subhalo to sink towards the host's center has been investigated by Boylan-Kolchin et al. (2008). They show that dynamical friction has an appreciable effect for mass ratios $0.025 \leq M_S/M_h \leq 0.3$. For lighter subhaloes the effect of dynamical friction is insignificant, whereas larger haloes simply fall into the host's center within a free fall time.

The effect of tidal forces on subhaloes can be broadly described in two different regimes. Subhaloes on circular orbits experience a constant tidal field, and material is adiabatically stripped from the subhalo in a process called tidal stripping. Subhaloes on more radial orbits experience a very strong tidal force close to the pericenter passage. This effect is more impulsive, i.e. it occurs on a timescale much shorter than the typical orbital timescale of the subhalo. The material in the subhalo then experiences what is called tidal shock heating.

In the adiabatic regime, a useful concept is the so-called tidal radius r_t. Particles that are closer to the subhalo's center than r_t remain bound to the subhalo, while particles that are further away from the subhalo's center will be removed by tidal stripping. Thus, the tidal radius is the radius at which the gravitational pull towards the subhalo's center equals the constant tidal force exerted by the host halo. For subhaloes on circular orbits

with orbital radius R, this tidal radius is (Tormen et al. 1998)

$$r_t = R \left[\frac{m_{\rm sub}(< r_t)}{(2 - d\ln M/d\ln r) M(< R)} \right]^{1/3}. \tag{3.4}$$

The tidal radius is hard to define in more general scenarios and predictions of the rate at which material is stripped from the halo usually do not match results from simulations very well, apart from the rather general statement that the timescale on which tidal stripping affects a subhalo is closely related to the subhalo's orbital time (Taylor and Babul 2001, Zentner and Bullock 2003, Zentner et al. 2005). Although the tidal stripping of subhaloes on circular orbits remains an open topic (van den Bosch et al. 2018, van den Bosch and Ogiya 2018), this may not be too serious since the orbits of most real subhaloes are far from circular. The tidal evolution of subhaloes on more realistic – and more radial – orbits is determined by tidal shock heating rather than tidal stripping (Zavala and Frenk 2019).

Tidal shock heating occurs due to strong tidal forces acting over a time that is very short compared to the subhalo's orbital timescale (Aguilar and White 1985), Aguilar and White 1986) and its impact is well explained by the so-called impulsive approximation (Spitzer 1958, Gnedin et al. 1999). Here, the energy change of a particle during a tidal shock is calculated as the square of the integral of the tidal acceleration experienced by the particle along the orbit of the subhalo. If this energy change is larger than the particle's initial binding energy, the particle is assumed to become unbound. This impulsive approximation is generally found to be in good agreement with N-body simulations (van den Bosch et al. 2018, Hayashi et al. 2003), in particular for subhaloes on very radial orbits.

To characterize the impact of tidal shock heating, Kazantzidis et al. (2004) carried out idealized simulations of subhaloes orbiting in a static external halo potential. They find that over multiple pericenter passages, their simulated subhaloes lose a substantial amount of their initial mass. While most of the mass is lost from the haloes' outskirts, the central densities are found to decrease as well. What is striking, however, is that the central slope of the subhaloes' mass density profiles is unaffected by tidal shock heating. In fact, Kazantzidis et al. (2004) find that the final density profiles of their subhaloes are very close to NFW (equation 2.15) at the center. In the outskirts of the shock heated subhaloes, however, the mass density profiles fall of significantly steeper than a NFW profile.

A detailed discussion of all the changes to the inner symmetry, the dynamical structure, and the radial distribution of subhaloes which are the result of dynamical friction, tidal stripping, and tidal heating can be found in Zavala and Frenk (2019) and the references therein. For our purposes, two things are important:

- Within ΛCDM, and in the absence of any baryonic effects, the total number of subhaloes whose masses exceed a certain fraction of the host halo mass is given by the subhalo mass function, equation 3.1.
- At the center of ΛCDM haloes, $\rho(r) \propto r^{-1}$. In subhaloes, this asymptotic behaviour persists and is unaffected by gravitational interactions with the host halo.

3.2 Small-scale challenges

The term "small-scale challenges to the ΛCDM paradigm" refers to mismatches between observations and predictions made by cosmological DMO simulations assuming a ΛCDM cosmological model. The task at hand for astrophysicists is then to find explanations for these mismatches without sacrificing the remarkable agreement between observations and theory on large scales. One way of addressing the small-scale challenges is to look for solutions within ΛCDM. These might either be related to baryonic physics, which is not included in DMO simulations, or they might be related to a limitation in the resolution of cosmological simulations due to which environmental effects are not properly described (see section 3.1.3). Alternatively, the small-scale challenges to ΛCDM may possibly be overcome by modifying the theory in such a way that only its predictions on small-scales change. For instance, DM may not be cold and collisionless. Instead, either significant free-streaming may have erased perturbations on small scales (see figure 2.1), or DM particles may collide with themselves at a sufficiently high rate to affect the DM distribution within the center of haloes. A more radical proposal aims to eliminate the need for DM altogether by posing that our understanding of gravity is fundamentally wrong on small scales. The work done in this thesis mainly focuses on baryonic feedback processes and alternative DM models, and for that reason, so-called Modified Newtonian dynamics (MOND) will not be reviewed here[1]. The following short review loosely follows Bullock and Boylan-Kolchin (2017). However, I will only briefly mention those small-scale challenges that are not relevant for the work of this thesis (missing satellites, plane of satellites, and the regularity of circular accelerations which stands opposed to the diversity of the rotation curves), while giving a more comprehensive and up-to-date discussion of the cusp-core problem and the too-big-to-fail problem, as well as a slightly modified version of too-big-to-fail that includes kinematic measurements of ultra-faint dwarfs. While possible solutions to those challenges relevant to the work of this thesis will be mentioned and explained, a more detailed descriptions of the methods used in my work will be given in chapter 4

3.2.1 Missing satellites

The problem of missing satellites is now understood to be not so much a challenge to ΛCDM as it is a constraint on the efficiency of galaxy formation at the faint end. The supposed contradiction between ΛCDM and observations is that according to equation 3.1, the number of DM subhaloes in the MW halo which should in theory have supported gas cooling at some point in their evolutionary history is much larger than the number of dwarf galaxies which have been observed to orbit the MW as satellites. Thousands of predicted subhaloes stand opposed to only ~ 50 observed MW satellites (Drlica-Wagner et al. 2015). Most likely, this is explained by galaxy formation becoming increasingly inefficient at smaller halo masses, for haloes with $M_h \leq 10^{12}\mathrm{M}_\odot$. This statement can be inferred from a statistical analysis called abundance matching, in which it is assumed that the brightest galaxies inhabit the most massive DM haloes (Frenk et al. 1988, Kazantzidis et al. 2004, Conroy et al. 2006, Moster et al. 2010, Behroozi et al. 2013). Matching observed galaxies to simulated haloes then results in a

[1] See McGaugh (2015) for a review on MOND.

statistical relation between halo mass and stellar mass. At the faint end, this abundance matching relation is rather unconstrained. However, the measured number of MW satellites is explained reasonably well by simply extrapolating the power law relation measured at the faint end of the stellar mass to halo mass relation down to smaller halo masses (see figures 6 and 8 of Bullock and Boylan-Kolchin (2017)). This suppression of galaxy formation at the faint end likely has several origins. For instance, while haloes with $M_{\rm peak} \geq 10^7 {\rm M}_\odot$ have likely supported molecular cooling at some point, atomic cooling has occurred only in haloes with $M_{\rm peak} \geq 10^8 {\rm M}_\odot$ (see e.g. Rees and Ostriker 1977), and gas accretion is suppressed by UV feedback from reionization in haloes with $M_{\rm peak} \geq 10^9 {\rm M}_\odot$ (Efstathiou 1992, Bullock et al. 2000, Benson et al. 2002, Bovill and Ricotti 2009, Sawala et al. 2016).

3.2.2 Satellite planes

Since the 1970's, it has been pointed out that the satellite galaxies of the MW appear to be distributed within a thin orbital disk which lies almost perpendicular to the MW (Kunkel and Demers 1976, Lynden-Bell 1976). More recent analyses (including more recently discovered satellite galaxies) essentially agree with this early statement (Pawlowski et al. 2012, Pawlowski and Kroupa 2013, Pawlowski et al. 2015), although one has to keep in mind that due to the brightness of the MW disk, surveys reach completeness only at high galactic latitudes. A similar disk of satellites has been observed around Andromeda, with 15 of 27 known Andromeda satellites belonging to a common plane (Conn et al. 2013, Ibata et al. 2013). Moreover, the observed kinematics seem to point towards a rotationally supported plane. However, it is unclear whether or not Andromeda's satellite disk truly is rotationally supported, given that the plane of satellites is seen almost directly edge-on. The question at hand is then whether or not such planes of satellites are consistent with ΛCDM structure formation scenarios. Considerable amount of work has gone into arguing that satellite planes are either consistent or inconsistent with ΛCDM structure formation (see discussion in Bullock and Boylan-Kolchin 2017). While a definitive answer has not yet been given, a few recent works indicate under which conditions the observed planes may be considered a serious challenge to ΛCDM in the future. On the one hand, data from the Sloan Digital Sky Survey shows that the observed satellite planes around M31 and the MW are not a common feature of MW-size galaxies (Phillips et al. 2015). On the other hand, it has been pointed out that while satellite planes are not particularly uncommon in ΛCDM simulations, they are very rarely rotationally supported (Bahl and Baumgardt 2014, Gillet et al. 2015, Buck et al. 2016). Thus, the satellite plane of M31 may be a serious issue for ΛCDM if future data confirms its rotational nature.

3.2.3 Regularity vs diversity

Observations of the internal dynamics of galaxies suggest a tight correlation between a galaxy's rotation curve and the enclosed baryonic mass. This correlation is particularly apparent when comparing the radial acceleration inferred from rotation curves, $g_{\rm obs} = V^2/r$ to the radial acceleration expected from the baryonic mass alone. At large accelerations, there is a one-to-one correspondence between the two. However, at

accelerations below $10^{-10}\,\mathrm{m\,s^{-2}}$, the observed accelerations begin to be systematically larger than the acceleration expected from baryons (McGaugh et al. 2016). In ΛCDM, this deviation at low accelerations is explained by DM. Yet, it is striking that the observed acceleration appears to be tightly correlated with the observed baryonic mass, even in the range where, under the assumption of Newtonian gravity, baryons alone cannot explain it. The observation of this radial acceleration relation is the key motivation behind MOND (McGaugh 2015), however, explanations can also be given within ΛCDM. Several authors have shown that the radial acceleration relation can naturally be obtained in hydrodynamical simulations of disk galaxies (e.g. Keller and Wadsley 2017). Navarro et al. (2017) have argued that based on results from abundance matching, a one-to-one correspondence between radial accelerations and observed baryonic mass is expected at accelerations above $10^{-10}\,\mathrm{m\,s^{-2}}$, and Di Cintio and Lelli (2016) used semi-analytic models to show that the full observed radial acceleration acceleration can be explained within ΛCDM, especially if the host haloes of dwarf galaxies are cored (see section 3.2.5).

While the properties of observed galaxy rotation curves are tightly correlated with the baryonic mass, there is a great observed diversity of rotation curves at a fixed DM halo mass. For instance, Oman et al. (2015) compared the rotational velocity measured at $r = 2\,\mathrm{kpc}$ from the center of galaxies to V_{flat} and demonstrated that at a fixed V_{flat}, there is a considerable amount of scatter in the rotational velocity at 2 kpc. This scatter is found to substantially exceed the amount that is expected when assuming DM distributed in NFW haloes as the dominating mass component at large radii. Oman et al. (2016a) and Santos-Santos et al. (2020) argue that the observed diversity of rotation curves may simply be the result of non-circular motion. However, if this diversity proves to be physical, it will pose a serious challenge to ΛCDM, in particular in combination with the radial acceleration relation. It remains to be seen whether the radial acceleration relation and the diversity of rotation curves can be self-consistently explained within ΛCDM once effects related to the formation and evolution of galaxies are taken into account.

3.2.4 Too-big-to-fail

The classical too-big-to-fail (TBTF) problem concerns the most massive subhaloes in high resolution cosmological simulations of MW-like haloes and the brightest satellite galaxies of the MW that have been detected. In short, the issue is that the central densities of the simulated subhaloes are too large to explain the observed internal kinematics of the satellites. Taken at face value, the CDM DM-only simulations would predict that no galaxies have formed within the heaviest subhaloes. On the other hand, since we do observe satellite galaxies, those must have formed within lighter subhaloes instead, which is a direct contradiction to the abundance matching assumptions that allowed us to “solve" the missing satellite problem. There is no theoretical reason for which the heaviest subhaloes should not host galaxies while the lighter ones do, they are too big to fail to form galaxies (Boylan-Kolchin et al. 2011). The classical TBTF problem is illustrated on the left panel of figure 3.1. Shown are the measurements of half light radius $r_{1/2}$ and circular velocity at the half light radius, $v_{1/2}$, of the 10 brightest MW satellites. Purple lines indicate the circular velocity curves of the heaviest subhaloes

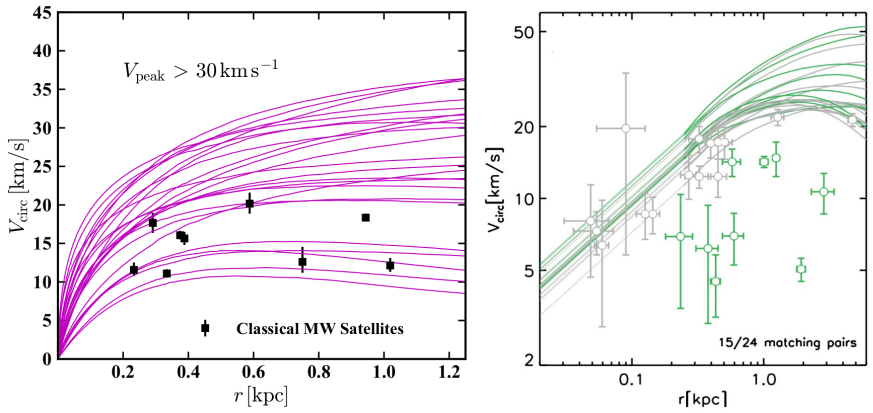

Figure 3.1. Classic TBTF (left panel[2]) and TBTF including mass measurements of ultra faint dwarfs (right panel, modified from Zavala et al. 2019). Both panels show the circular velocity profiles of the most massive subhaloes (after removing LMC and SMC analogs) of a zoom simulation of a MW-like CDM halo from the Aquarius simulation (Springel et al. 2008). The halo on the left panel is described in Boylan-Kolchin et al. (2012), the halo on the right panel in Springel et al. (2008). The data points on the left panel correspond to circular velocity measurements of the brightest dwarf spheroidal galaxies orbiting the MW. Only subhaloes with $v_{\rm peak} > 30\,{\rm km\,s^{-1}}$ *are shown. The most massive subhaloes are too dense in their center to explain the circular velocities of the brightest satellites. On the right panel, mass measurements of ultra-faint dwarfs (Errani et al. 2018) are included. Thus, the right panel compares the 24 heaviest subhaloes to the 24 dwarf galaxies less than* 300 kpc *away from the MW center. Only 15 data points are matched by the subhaloes and the heaviest subhaloes appear to host the ultra-faint galaxies.*

(with a cutoff at $v_{\rm peak} = 30\,{\rm km\,s^{-1}}$, where $v_{\rm peak}$ is the largest maximal circular velocity of the halo before it was accreted) found within a high-resolution zoom simulation of a MW-like halo (Boylan-Kolchin et al. 2012) from the Aquarius simulation (Springel et al. 2008). Haloes with such large peak circular velocities should theoretically have supported star formation. Evidently, though, the circular velocity curves of the heaviest subhaloes do not intersect any measurement of $(r_{1/2}, v_{1/2})$. In other words, they are too centrally dense to explain the kinematics observed in the brightest dwarfs.

TBTF appears to be an issue not only for MW satellites. Similar mismatches with the central density of simulated haloes have been identified for the satellites of Andromeda (Tollerud et al. 2014) and for field galaxies in the local group (Kirby et al. 2014, Garrison-Kimmel et al. 2014), as well as for more isolated galaxies (Ferrero et al. 2012, Papastergis et al. 2015, Papastergis and Shankar 2016). However, it has also been argued that rotational velocities may be systematically underestimated (Macciò et al. 2016), and that a lot of small galaxies may simply be too faint to be detected by current surveys (Brooks et al. 2017). Taken together, those two effects could explain TBTF in

the field. Errani et al. (2018) presented a mass estimator for ultra-faint dwarfs that the authors claim is unbiased and accurate to within $\sim 10\%$. Taking the measured half light radii and half light velocities of ultra-faints into account when comparing the measured circular velocities of dwarfs to the circular velocity profiles of simulated subhaloes changes the picture considerably. The resulting updated TBTF challenge is shown on the right panel of figure 3.1. Here, measurements of the half light radii and the circular velocities at the half light radii of all 24 known dwarf galaxies within the MW halo (less than 300 kpc away from the MW's center) are shown as data points, along with the circular velocity profiles of the 24 heaviest subhaloes (after removing LMC and SMC analogs) identified in a CDM zoom simulation (without baryons) of a MW analog from the Aquarius simulation suite. Strikingly, the largest CDM subhaloes now have galaxy counterparts. These counterparts, however, appear to be the ultra-faint dwarfs instead of, as one might expect, the brightest satellites. Moreover, the 24 heaviest subhaloes can account for at most 15 of the measured circular velocities. The 9 remaining satellites all have measured circular velocities which are smaller than would be expected if these galaxies were hosted by one of the 24 heaviest haloes. This updated version of TBTF is thus a statement about the diversity of the rotation curves of MW satellites. In similarity to the mismatch between the diversity of observed rotation curves of field galaxies and the much smaller scatter of circular velocity curves in NFW haloes pointed out by Oman et al. (2015) (see section 3.2.3), the scatter in the circular velocity curves of the 24 heaviest subhaloes within a MW-like halo cannot account for the diverse distribution of the circular velocities of the observed MW satellites.

It is currently unclear how much of a problem this is for ΛCDM. One may for instance ask whether the observed dwarfs need to indeed be hosted by the heaviest subhaloes. The right panel of figure 3.1 suggests that this standard assumption behind abundance matching does not apply to the MW satellites. One reason for that may be that galaxy formation becomes increasingly stochastic as one approaches the lower limit of halo masses for which galaxy formation is theoretically possible. Thus, it is possible that the heavy simulated subhaloes that appear to have stayed dark could be the hosts of faint dwarfs that have not yet been detected. However, considering only the currently detected MW satellites, the predicted heavy subhaloes without a galaxy counterpart do indeed present a theoretical challenge to ΛCDM. Since the accretion of gas onto subhaloes with $v_{\rm peak} > 30\,{\rm km\,s^{-1}}$ is not affected by UV feedback, such subhaloes should not have stayed dark.

Several mechanisms have been proposed to solve TBTF within ΛCDM. For instance, supernova (SN) feedback can lower the central DM density of dwarf-sized haloes by impulsively injecting energy into the halo's center, thereby causing the orbits of central DM particles to expand and creating constant density cores (see discussion and references in section 3.2.5 and chapter 4). However, since the resulting decrease in central DM mass is substantial only for galaxies with $M_\star > 10^6 {\rm M}_\odot$, SN feedback on its own is likely insufficient to explain TBTF (Bullock and Boylan-Kolchin 2017). As pointed out in section 3.1.3, gravitational interactions between the MW and the subhaloes, which are not accurately resolved in DM-only cosmological simulations,

[2]Republished with permission of Annual Reviews Inc., from Small-Scale Challenges to the ΛCDM Paradigm, Boylan-Kolchin, Michael; Bullock, James S., volume 55, edition 1, ©2017; permission conveyed through Copyright Clearance Center, Inc.

can substantially alter the inner structure of subhaloes. Such environmental effects are thus often invoked to explain the classical TBTF of MW satellites (e.g. Zolotov et al. 2012, Arraki et al. 2014, Brooks and Zolotov 2014, Brook and Di Cintio 2015, Wetzel et al. 2016, Tomozeiu et al. 2016, Sawala et al. 2016, Dutton et al. 2016). More recently, Fattahi et al. (2018) demonstrated that tidal stripping can significantly reduce the mass of the MW's subhaloes. The observed low circular velocites of relatively bright MW satellites can be explained if such dwarfs are hosted by subhaloes that have lost large fractions of their mass at infall to tidal stripping. Unresolved environmental effects could thus create an additional spread in the circular velocity curves of the most massive subhaloes, significantly alleviating the challenge that TBTF poses to ΛCDM. One way to assess whether such an explanation is indeed plausible is to look for systematic differences between field dwarfs and satellite dwarfs with $M_\star \lesssim 10^6 \mathrm{M}_\odot$, the mass range in which SN feedback is not traditionally[3] expected to have a significant impact on the central densities. If environmental effects have indeed reduced the central densities of DM subhaloes, they should have affected the dwarf galaxies residing in the central regions of those haloes as well. A study by Kirby et al. (2014) did not find any such systematic differences, but it should be noted that at the time, only very few systems had been detected in this mass range.

Another class of solutions to TBTF is based on modifications of the ΛCDM paradigm. If DM is not dynamically cold and collisionless, the abundance of small haloes can change, as can their inner structure. Similar to SN feedback for galaxies with $M_\star > 10^6 \mathrm{M}_\odot$, self-scattering between DM particles can reduce the central densities of haloes and subhaloes by redistributing mass and turning their central cusps into constant density cores (see discussion and references in section 3.2.5 and chapter 4). If TBTF were indeed solved by SN feedback or DM self-scattering, it would naturally be connected to the cusp-core problem (see section 3.2.5). Note, however, that simply having cored subhaloes may solve the classical TBTF challenge, but it in no way addresses the updated version (right panel of figure 3.1). The diversity of the observed rotation curves of satellite galaxies is still inconsistent with the assumption that they are hosted by the heaviest subhaloes. In fact, in Zavala et al. (2019) we show that the updated TBTF is a serious problem for self-interacting DM (SIDM) with a constant momentum transfer cross section $\sigma_T / m_\chi = 1\,\mathrm{cm}^2\mathrm{g}^{-1}$, since the central density of all SIDM subhaloes is too low to explain the circular velocities of some ultra-faints.

Modifying predictions from the linear regime of stucture formation theory can potentially explain both the classical and the updated TBTF challenge. As pointed out in section 2.1, perturbations in the early Universe can be erased through two distinct mechanisms, collisional damping or free streaming of thermal relics (see figure 2.1). The corresponding suppression in the linear power spectrum is mirrored in a reduced amount of small DM haloes at later times. For WDM, the free-streaming length translates into a free-streaming mass, corresponding to the smallest DM structures that can form in a WDM cosmology. For thermal relics with a mass of 3.5 keV, this free-streaming mass is roughly $M_{\mathrm{fs}} = 1.1 \times 10^5 \mathrm{M}_\odot$ (Bullock and Boylan-Kolchin 2017).

While no structure less massive than M_{fs} forms, the amount of haloes with masses that are larger by two to three orders of magnitude is also reduced significantly (Schnei-

[3] See Read et al. (2016) for arguments why the impact of SN feedback at those scales may be underestimated.

der et al. 2012). Generally, haloes which are lighter than the half mode mass are affected. For a thermal relic with mass 3.5 keV, this half mode mass is $\sim 3 \times 10^8 M_\odot$. The number of formed subhaloes is reduced significantly compared to CDM in the mass range between free-streaming mass and half mode mass. Moreover, since free-streaming damping affects the perturbations on those scales, haloes in this mass range form later on average. As a consequence, the central densities of these haloes are lower than in CDM, reflecting the decreased average density of the Universe at the time they collapse. Thus, WDM solves TBTF through a combination of two effects. Compared to CDM, the abundance of substructure within the MW halo is reduced, as are the central densities of corresponding subhaloes. WDM models can therefore be constrained by counting the total number of satellites detected in the MW[4] (e.g. Polisensky and Ricotti 2011, Lovell et al. 2014). The suppression of structure in models in which the early DM perturbations are affected by collisional damping on small scales works analogous to the WDM case[5]. However, such models can evade WDM constraints if the amplitude of the dark acoustic oscillations is sufficiently large (Bose et al. 2019b, Bohr et al. 2020).

3.2.5 Cusps or cores

The measured rotation curves of many observed dwarf galaxies, both field dwarfs and satellites, are inconsistent with DM only predictions of ΛCDM. Dwarf galaxies are DM dominated systems, and thus their internal kinematics trace the DM mass profile. The circular velocity at a given radius is directly related to the DM density profile,

$$V_{\mathrm{circ}}^2(r) = \frac{GM(r)}{r}. \tag{3.5}$$

For dwarf galaxies embedded in a NFW halo (equation 2.15), this implies $V_{\mathrm{rot}}(r) \propto r^{0.5}$ at small radii. However, some measured rotation curves in disk dwarf galaxies rise more slowly and are better described by $V_{\mathrm{rot}}(r) \propto r$ at small radii. Such rotation curves correspond to the circular velocity curves of DM haloes with a constant density core, i.e. $\rho(r) = const$ at small radii. This mismatch between observations and DM only ΛCDM predictions is referred to as the cusp core problem. It is illustrated in figure 3.2 where the measured rotation curves of two dwarf galaxies from the LITTLE THINGS survey (Oh et al. 2015) are compared to the circular velocity curves of a NFW halo and a cored Burkert (1995) halo with similar maximal circular velocity V_{max}. The data are better explained by the cored Burkert profile than the cuspy NFW profile.

Evidence indicating the presence of cored DM halo density profiles comes mostly from measurements of HI-rotation curves of field dwarfs (e.g. Moore 1994, de Blok et al. 2008, Kuzio de Naray et al. 2008, Read et al. 2019). Moreover, Walker and Peñarrubia 2011 also claim to have discovered that the DM haloes hosting Fornax and Sculptor, two of the classic MW dwarf spheroidals, have cored density profiles. The authors identify two distinct stellar subpopulations that differ systematically in metallicity within both of these satellites. Using stellar spectroscopic data, they then

[4] Other constraints on WDM come from the Ly-alpha forest

[5] The discussion here has been focused on thermal relic WDM. Different production mechanisms for WDM (for example resonant production) result in a slightly different free-streaming length at the same particle mass.

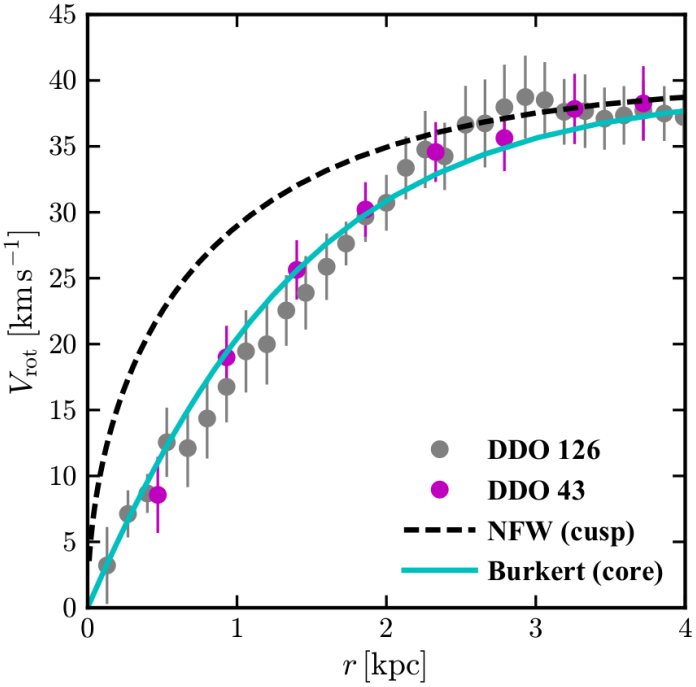

Figure 3.2. The cusp core problem in field dwarfs. The rotation curves of two dwarf galaxies from the LITTLE THINGS (Oh et al. 2015) survey are shown. The circular velocity curves of a cuspy NFW halo (dashed line) and a cored Burkert (1995) halo are overplotted, both tuned to fit the flat end of the rotation curves. Evidently, a cored DM density profile yields a better fit to the data. Figure adapted from Bullock and Boylan-Kolchin 2017[6].

measure the half-right radii and the velocity dispersion of both stellar population, and subsequently use those to calculate the slope of the DM mass profile. The inferred slopes are inconsistent with NFW haloes, but consistent with DM haloes with constant density cores. Possibly related to the evidence for cored DM profiles is the observation that simulated DM haloes have to much mass at their centers. The observed dynamics of galaxies that, following abundance matching, should be hosted by those haloes, are consistent with hosts with less central mass on average. This problem is related to TBTF and somewhat separate from the mismatch between observed rotation curves and simulated halo density profiles, as it concerns the overall normalization and not the inner slope of density profiles (Alam et al. 2002, Oman et al. 2015). However, the observed discrepancy would be resolved if DM haloes were cored instead of cuspy.

While several groups have independently made claims that their observations imply the presence of cored DM haloes, this evidence remains controversial. It has been argued, for instance, that a large diversity in observed rotation curves can be caused by non-circular gas motion (Oman et al. 2019, see also Santos-Santos et al. 2020). In

the presence of such non-circular motion, the rotation curves look different depending on the line of sight to the galaxy. Hence, observed slow rising rotation curves can potentially be mistakenly interpreted as evidence for cored DM haloes. The Walker and Peñarrubia (2011) claim that the DM host haloes of Fornax and Sculptor are cored has also been contested. According to Genina et al. (2018), their measurement of the slope of the DM mass profiles could be systematically flawed because it does not consider that spherical symmetry is violated in both of those dwarf spheroidals. Strigari et al. (2017) analyzed Sculptor's two stellar populations using a distribution function based framework. Contrary to Walker and Peñarrubia (2011), they found no evidence for a core.

Controversy notwithstanding, there are several ways in which the presence of cores in dwarf-sized DM haloes can be reconciled with the success of ΛCDM on large scales. Since any successful theory of structure formation must resemble ΛCDM on large scales, solving the cusp-core problem requires a mechanism that changes only the inner structure of DM haloes. Crucially, as pointed out in section 3.1.3, there are no environmental effects that cause this to happen, meaning that such a mechanism cannot be purely gravitational. This essentially leaves two options. Either cusps are transformed into cores through a mechanism related to baryonic physics, or the cusp-core transformation is the result of new DM physics, meaning that DM can not be approximated as cold and collisionless on small scales. It is indeed one specific baryonic effect and one particular modification to the properties of DM which are commonly invoked as the most viable methods to solve the cusp-core problem: supernova (SN) feedback on the one hand and self-interacting DM (SIDM) on the other hand.

SN feedback denotes the transfer of energy from supernovae to the DM particles through gravitational interaction between baryons and DM (e.g. Navarro et al. 1996a, Gnedin and Zhao 2002, Pontzen and Governato 2012, Di Cintio et al. 2014, Chan et al. 2015, Read et al. 2016, Tollet et al. 2016, Bose et al. 2019a, Benítez-Llambay et al. 2019, Lazar et al. 2020). The thermal energy of supernovae is converted into kinetic energy and deposited into the surrounding gas. This can cause rapid outflows of gas which substantially and impulsively change the local gravitational potential. Since mass is removed, the potential well becomes shallower, causing an expansion of the orbits of surrounding DM particles. Thus, if supernovae are very concentrated into the central region, they can, under some conditions, cause a net reduction of the surrounding DM halo's central density. That a rapidly decreasing gravitational potential in the center of a DM halo can transform a density cusp into a core has been known for more than two decades. Navarro et al. (1996a) showed that removing an external disk potential in a simulation of an idealized cuspy halo can cause this halo to form a core. A few years later, it was shown that such a single removal of mass is likely not sufficient to trigger a cusp-core transformation (Gnedin and Zhao 2002). The current picture of how exactly DM cusps are transformed into cores is largely due to Pontzen and Governato (2012). The authors showed that a periodic, SN-driven, removal of mass from the center of DM haloes drives the formation of a DM core and moreover demonstrated that this consistently explains how core formation occurs in hydrodynamic

[6]Republished with permission of Annual Reviews Inc., from Small-Scale Challenges to the ΛCDM Paradigm, Boylan-Kolchin, Michael; Bullock, James S., volume 55, edition 1, ©2017; permission conveyed through Copyright Clearance Center, Inc.

simulations. One condition of their model is that the removal of mass, and thus the change of the gravitational potential, happen instantaneously. In practice, this means that for SN feedback to be efficient at changing the DM mass distribution, so-called starburst events (the birth of a large number of massive stars - which is followed by supernovae some 10 Myrs later) must happen on timescales which are shorter than the typical dynamical timescales of DM particles in the inner halo. For massive dwarfs, observational evidence suggests that starbursts, and thus SN feedback cycles, happen indeed on timescales which are comparable to the dynamical timescales of the galaxies (Kirby et al. (2014)). However, for low mass MW satellite galaxies the situation is more uncertain, as observations currently lack the time resolution required to resolve the starburst cycle on the dynamical timescales of those dwarfs (Weisz et al. 2014). SN feedback cycles need not only be sufficiently bursty, the total energy that is deposited into the DM halo also needs to be enough to unbind the halo's cusp. As pointed out by Peñarrubia et al. (2012), the total energy deposited by supernovae scales linearly with a galaxy's stellar mass, while the energy required to unbind a DM halo's cusp is proportional to the squared halo mass. Whether or not SN feedback can transform a DM cusp into a core should thus be a function of the stellar-to-halo mass ratio. Most current hydrodynamic simulations confirm this picture and find that core formation through SN feedback is most efficient in the mass range of bright dwarfs (e.g. Fitts et al. 2017, Benítez-Llambay et al. 2019, Lazar et al. 2020, Dutton et al. 2020). However, Read et al. (2016)'s high-resolution simulations of isolated systems suggest even the faintest known dwarf galaxies should be efficient at forming constant density cores in their host haloes through SN feedback. Finally, while SN-feedback induced core formation is a feature in many hydrodynamic simulations of galaxy formation, this is a highly model-dependent result, and in some simulations no core formation occurs (Bose et al. 2019a). In chapter 4 I will take a closer look at SN feedback and how it is implemented, either in DMO or in hydrodynamic simulations.

SIDM (Spergel and Steinhardt 2000, Yoshida et al. 2000, Davé et al. 2001, Colín et al. 2002, Vogelsberger et al. 2012, Rocha et al. 2013) can solve the cusp-core problem through (strong) elastic self-scattering between DM particles. These self-scattering events cause an effective heat transfer towards the center of the halo and create an isothermal core. For SIDM to be efficient at forming a core, the cross section for self-interactions needs to be such that a DM particle which is located within a central sphere with radial extent of roughly the NFW scale radius will on average have scattered at least once with another DM particle between the halo's formation time and today. Since the velocity dispersion of NFW haloes is small in the haloes' center and rises out to roughly the scale radius (see figure 2.3), particles orbiting at the scale radius are initially much more energetic than particles at smaller radii. For that reason, elastic scattering events between particles in a sphere with radius $\sim r_s$ cause an outside-in energy redistribution in the center of haloes. The rate at which a core is formed depends on the self-interaction cross section, but the eventual core size is determined by the scale radius of the halo at its formation time (e.g. Vogelsberger et al. 2012, Rocha et al. 2013). If the cross section is small, the cores observed at the present time can be smaller, since the process of core formation is incomplete. Finally, it should be noted that the isothermal cores produced by SIDM are a transient stage, and SIDM haloes will eventually undergo gravothermal collapse (Balberg et al. 2002, Colín et al. 2002, Koda and Shapiro 2011,

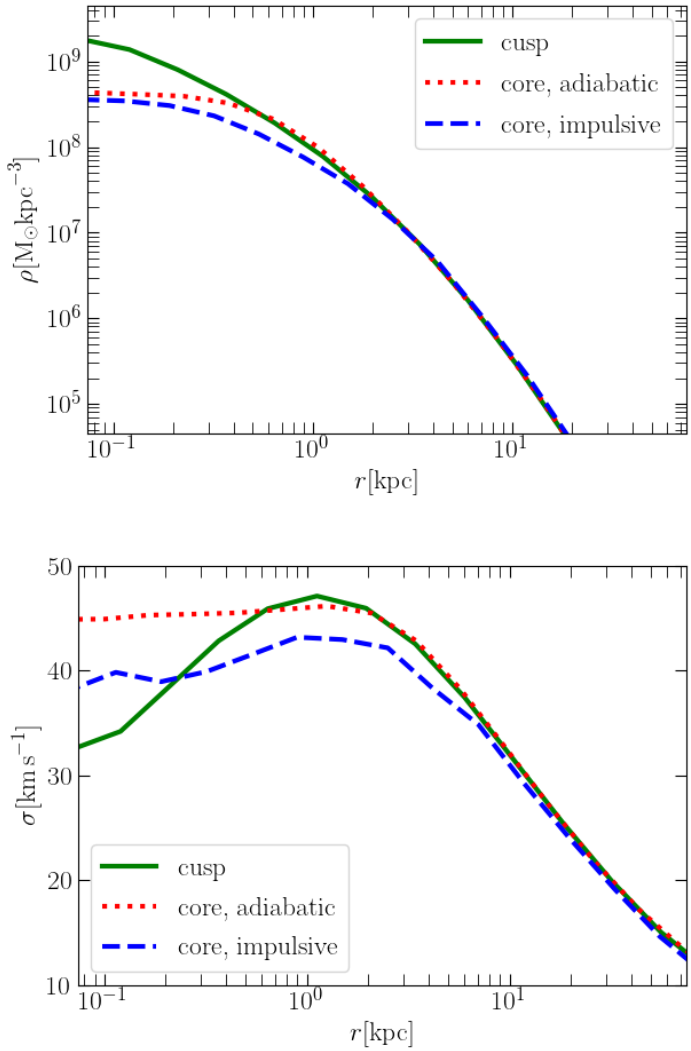

Figure 3.3. Core formation in a dwarf-size halo through two different mechanisms. Green solid lines show the density (upper panel) and velocity dispersion (lower panel) profile of an initially cuspy halo in dynamical equilibrium. The halo's central density is then reduced, either adiabatically through self-scattering between its DM particles (red dotted lines), or through a repeated impulsive removal of mass (blue dashed lines). The cusp-core problem can be solved in both ways, but the resulting distribution function of DM is different. Are these differences reflected in the dynamics of kinematic tracers ? Figure modified from Burger and Zavala (2019).

Pollack et al. 2015, Nishikawa et al. 2020), forming a very dense central region. In line with considerations above, this "gravothermal catastrophe", a phenomenon first described for globular clusters (Lynden-Bell and Wood 1968), happens faster for larger self-interaction cross sections. It is this strong dependence of SIDM predictions on the self-scattering cross section that allow astronomers to put constraints on viable SIDM models. Currently, the most stringent constraints on the SIDM momentum transfer cross section per unit mass, σ_T/m_χ, come from observations of clusters (Robertson et al. 2017, Robertson et al. 2019) and large elliptical galaxies (Peter et al. 2013). On the scales of clusters and elliptical galaxies, $\sigma_T/m_\chi \lesssim 1\mathrm{cm}^2\mathrm{g}^{-1}$ (see Tulin and Yu 2018) for a review on SIDM constraints). On the scale of dwarf galaxies, the SIDM cross section has only recently been constrained. Read et al. (2018) finds $\sigma_T/m_\chi \lesssim 0.6\mathrm{cm}^2\mathrm{g}^{-1}$ given that the host halo of the MW satellite Draco is cuspy. In section 3.2.4, we have shown that the mass measurements of ultra faint dwarfs cannot be explained if DM is self-interacting with a constant cross section of $\sigma_T/m_\chi \gtrsim 1\mathrm{cm}^2\mathrm{g}^{-1}$ and thus the ultra-faints put similar constraints on SIDM. This leaves a rather narrow parameter space for SIDM with a constant momentum transfer cross section, since the predictions of SIDM and CDM are virtually indistinguishable for $\sigma_T/m_\chi \lesssim 0.1\,\mathrm{cm}^2\mathrm{g}^{-1}$ (Zavala et al. 2013). However, particle physics models of SIDM are very versatile, and it is easily possible to construct models with velocity-dependent momentum transfer cross sections that evade current constraints and still solve the cusp core problem. Mildly velocity-dependent SIDM with $\sigma_T/m_\chi \sim 1\,\mathrm{cm}^2\mathrm{g}^{-1}$ on the scale of dwarfs and $\sigma_T/m_\chi \sim 0.1\,\mathrm{cm}^2\mathrm{g}^{-1}$ on the scale of galaxy clusters can potentially explain mass deficits observed across all scales (e.g. Kaplinghat et al. 2016). As shown in Zavala et al. (2019), a big advantage of models with velocity-dependent cross sections is that they can naturally explain the large diversity of circular velocities measured in the MW dwarf satellites thanks to the efficiency of the gravothermal catastrophe mechanism at small velocities/masses. A short introduction into how self-scattering is included in N-body simulations is given in chapter 4.

As stated above, SN-related mass removal must happen on timescales which are shorter than the DM particles' dynamical timescales for DM haloes to form cores. SN feedback induced core formation is therefore an impulsive process. Self-interactions between DM particles, on the other hand, change the gravitational potential generated by the DM halo itself on much longer timescales. Hence, core formation by SIDM is adiabatic. The resulting cores also look different. SIDM cores are fully isothermal, with a flat density profile and a flat velocity dispersion profile out to the halo's scale radius. SN feedback cores, on the other hand, are not necessarily isothermal, and the density profile is not necessarily flat across the radial range in which the density is reduced compared to an NFW profile. This is shown in figure 3.3, where we show simulated density (upper panel) and velocity dispersion (lower panel) profiles of a cuspy CDM halo (green lines), a cored SIDM halo (red dotted lines), and a halo that has formed a core through repeated impulsive mass removal (blue dashed lines), mimicking SN feedback. As we can see, core formation, and the related change of the gravitational potential, can occur adiabatically (as in the SIDM case) or impulsively (as in the SN feedback case). Both mechanisms can solve the cusp-core problem. However, figure 3.3 shows that the resulting distribution function of the DM is not the same in both cases, meaning that a similar mass distribution can correspond to a different distribution of

velocities. Unfortunately, we cannot observe the dynamics of DM particles. Stars within the central galaxy which orbit within the halo as it forms a core are, on the other hand, both observable and well approximated as collisionless. Thus, they act as kinematic tracers of the potential generated by the DM. One of the main goals of this thesis is to demonstrate that the orbits of such kinematic tracers are affected in different ways by adiabatically changing potentials and impulsively changing potentials. Dynamical signatures in the orbital data of kinematic tracers may thus be used to distinguish adiabatic from impulsive processes. In the context of the cusp-core problem this implies that we may be able to use stellar phase space data to investigate whether core formation in a halo is due to SN feedback or whether self-interacting DM is a more likely cause of observed cored density profiles.

3.3 Alternatives to N-body simulations

Aside from being hard to test observationally, the small-scale predictions of ΛCDM are also solely based on N-body simulations. On larger scales, there are well-motivated theoretical models that explain the results of N-body simulations. The clustering and the abundance of LSS, for instance, is reproduced remarkably well by Press-Schechter theory (see chapter 2). On smaller scales, the situation is significantly more complicated. While N-body simulations of CDM yield concordant results, these results are not realized one-to-one in nature, since the DM distribution on those scales is affected by baryonic processes which are not included in N-body simulations (see section 3.2) and whose implementation in full hydrodynamical simulation remains uncertain. Moreover, some small scale results of CDM N-body simulations have proven much more difficult to understand from fundamental principles than the large scale clustering of matter. In some cases, simplified models can provide good explanations of the evolution observed in simulations. The orbital decay of subhaloes, for instance, is explained remarkably well by Chandrasekhar's dynamical friction formula, and Read et al. (2006) showed that a simple analytical model can explain divergent behaviour that is seen in harmonic potentials. In general, however, an analytic derivation of the small scale results of CDM N-body simulations requires a theoretical model for the dynamical evolution of a large number of gravitationally interacting particles. At its core, this is an exercise in thermodynamics. Given a set of initial conditions, the task is to calculate how the system evolves towards a steady state. Unfortunately, there are a few long standing, unresolved problems one is faced with when attempting to model the evolution of a self-gravitating system using tools of statistical physics.

One key problem is that self-gravitating systems have negative specific heat (Antonov 1961, Lynden-Bell and Lynden-Bell 1977, Padmanabhan 1989; for a review see Lynden-Bell 1999). Hence, if two systems have different temperatures initially, their mutual gravitational interaction will only serve to increase this temperature difference, driving the system further away from thermodynamic equilibrium. This prohibits the use of the canonical or grand canonical ensemble when trying to obtain the system's most probable distribution function using statistical physics. The only remaining possibility is then to use the microcanonical ensemble.

However, one is faced with an additional theoretical problem when analyzing the

dynamical evolution of a collisionless system which is initially confined to a small volume in phase space, as is the case with tidal streams, for instance. Such systems will eventually populate all of the phase space which is available to them, however, this happens over some time in a process which is called phase mixing (see section 4.3.2). In some cases, phase-mixing may occur on time-scales which are much longer than the typical orbital time of a particle which is part of the self-gravitating system. Until the system is dynamically mixed, however, not all allowed microstates can be accessed and the system is restricted to a smaller region of phase space. Hence, self-gravitating systems may be non-ergodic, and knowledge of the full distribution of orbital phases is required to obtain a statistical model of their evolution.

A final issue concerns the process of statistical averaging in classical thermodynamics. Padmanabhan (1990) argued that in the presence of long range forces, it is not permissive to divide the system into non-interacting macro-cells, since the energy of such systems is a non-extensive parameter, i.e. the gravitational interaction between different macro-cells cannot be neglected. In consequence, Padmanabhan (1990) argues that the laws of equilibrium thermodynamics cannot be applied to gravitating systems.

All of these issues may explain why the results of early attempts at deriving an equilibrium distribution function of a gravitating system using methods of classical thermodynamics do not agree with the results of N-body simulations. A prominent example is Lynden-Bell (1967)'s calculation of the final equilibrium state of a collisionless system undergoing violent relaxation, which he found to be a weighted superposition of Fermi-Dirac distributions of different temperatures. In the non-degenerate limit, the well known Maxwell-Boltzmann distribution, i.e. the distribution function of an isothermal sphere, gives a reasonable approximation to this result. Nakamura (2000) obtained the Maxwell-Boltzmann distribution as an exact solution in their calculation of violent relaxation, where they used Jaynes (1957) information theory. Neither solution fits N-body experiments of violent relaxation (Arad and Johansson 2005), although Levin et al. (2008) and Levin et al. (2014) have shown that the final result of violent relaxation is likely a bivariate distribution in which particles in the innermost region of the potential are indeed distributed according to Lynden-Bell's theory, but the distribution is distinctly different in the outskirts. Perhaps even more troubling than the mismatch with simulations, however, is the fact that the derivation of the solutions of both Lynden-Bell (1967) and Nakamura (2000) require the system to have infinite mass. A more recent example of how statistical mechanics can be used to explain results of N-body simulations is the attempt by Pontzen and Governato (2013) to derive the ubiquitous inner density profile of DM haloes from first principles (see section 2.3.2). Most modern attempts focus on advancing the distribution function between equilibrium states, using a suitable approximation for the collisional operator in the Boltzmann equation (see Binney and Tremaine 2008). This approximation yields the so-called the Fokker-Planck equation, which is complemented by the accompanying Poisson equation. This system of coupled differential equations in six dimensions is then solved numerically by so-called Fokker-Planck codes. There are different methods to solve the Fokker-Planck equation, however, they all rely on a large number of assumptions and the results of Fokker-Planck codes are thus mainly used to gain an improved understanding of N-body simulations (see Dehnen and Read 2011 and references therein).

An alternative to both N-body simulations and statistical mechanics is Hamiltonian

perturbation theory (see e.g. Lynden-Bell and Kalnajs 1972, Tremaine and Weinberg 1984, Binney and Tremaine 2008). Hamiltonian perturbation theory is particularly successful if the Hamiltonian of the system is separable and can be expressed in terms of action-angle variables. The actions $\mathbf{J}$ are integrals of motion of the unperturbed Hamiltonian H_0. This means that they are conserved along the phase space path of a particle orbiting in H_0. The conjugate angles, however, vary along the phase space path of a particle and have a period that is related to the orbital period of the particle. Since the energy of a particle is conserved in the unperturbed system, the unperturbed Hamiltonian cannot be a function of the angles and must thus depend on the actions only. One may then expand the full Hamiltonian as a power series in some small parameter ε (with $\varepsilon \ll 1$),

$$H(\mathbf{J},\theta,t) = H_0(\mathbf{J}) + \varepsilon_1 H_1(\mathbf{J},\theta,t) + ... + \varepsilon^k H_k(\mathbf{J},\theta,t), \tag{3.6}$$

and solve the resulting equations of motion iteratively from low to high order. Hamiltonian perturbation theory is especially attractive due to the fact that in time-dependent potentials, actions are invariant if the evolution of the potential occurs sufficiently slowly and the symmetry of the system is not violated. Notably, Hamiltonian perturbation theory can be formulated into a fully self-consistent kinetic theory for the secular evolution of self-gravitating systems. The resulting system of equations resembles the Balescu (1960)-Lenard (1960) equations of Plasma physics, but in action-angle space instead of ordinary phase space (Heyvaerts 2010). This approach has been very successful, for instance in predicting the secular formation of a ridge in the distribution function of a rotating stellar disk (Fouvry et al. 2015) that can subsequently cause the formation of spiral arms (De Rijcke et al. 2019). Ultimately, Hamiltonian perturbation theory is limited by the the fact that it always relies on some degree of symmetry, since the gravitational potential must allow for the construction of independent action variables such that H_0 is well defined. Moreover, solving the coupled differential equations that arise in Hamiltonian perturbation theory is numerically challenging on its own.

The environmental effects which are most relevant for the evolution of DM subhaloes orbiting within a host are still most precisely known from the results of restricted N-body simulations. However, perturbative methods are becoming ever more sophisticated and in a few cases theoretical insight has been gained using tools of statistical mechanics. Currently, both perturbative methods and statistical approaches work only in the regime in which the dynamical evolution is adiabatic, i.e. actions of individual particles are (nearly) conserved (see section 4.3.1). In this thesis, I show that potentials with a global time-dependence lead to diffusion in radial action space. As long as the potential does not change on timescales which are much shorter than the typical dynamical times of the system, this diffusion can be calculated perturbatively. This implies that both, statistical mechanics and perturbation theory, may also be useful to describe the dynamical evolution of, for instance, haloes which accrete mass or haloes which undergo cusp-core transformation. However, if the potential changes too fast, the evolution of radial actions becomes discontinuous and non-perturbative.

Chapter 4

Methods and original articles

The previous chapters were dedicated to explaining the scientific context of the work presented in the articles that make up the body of this thesis. At the end of this chapter, I will briefly go over the main results obtained in each first-author paper. Before that, I will make a few general and very short remarks about the methods used to model SN feedback (in DMO simulations as well as in hydrodynamic simulations) and self-scattering between DM particles, as well as the methods used to analyze the dynamics of tracers in gravitational potentials.

In the interest of brevity, a few topics will have to be brushed over, since they are not fundamental towards understanding the work presented here. Most notably, hydrodynamic simulations will only be mentioned very briefly, and particle physics models of SIDM will not be discussed here. Hydrodynamic simulations do not need to be understood in detail in order to understand the physical impact of changing the parameters that regulate SN feedback in hydrodynamic simulations. Interested readers may consult Vogelsberger et al. (2020) and the references therein. Knowledge of any particular particle physics model for SIDM is not required to understand the physical implications of collisional DM. For the purpose of this work, it is sufficient to know that they exist and that there are models that can generate both constant and velocity-dependent self-scattering cross sections in the non-relativistic limit. For more information about SIDM, including some possible particle physics models, see the review by Tulin and Yu (2018).

4.1 Supernova feedback

4.1.1 Modeling SN feedback in DMO simulations

SN feedback affects the surrounding DM halo indirectly, through gravity. Supernova remnants sweep up gas and the combined effects of many supernova remnants can drive galactic scale winds. If active SN cycles happen near the halo center, they can rapidly decrease the gravitational potential in their surroundings and thereby increase the energy of the DM particles there, causing their orbits to expand. In DMO simulations, the effect of this rapid mass removal can be approximated by simply adding an external

acceleration term that mimicks the gravitational pull of baryonic matter to each particle in the halo, and then removing this acceleration term at once (thus mimicking the fast removal of mass). This method was originally used by Navarro et al. (1996a). A more modern version is the model by Garrison-Kimmel et al. (2013), which takes into account that SN feedback transforms cusps into cores by cyclic accumulation and subsequent removal of mass (Pontzen and Governato 2012). Usually, it is assumed that the acceleration is generated by some spherical mass distribution, centered at the halo's center of potential. Thus, one first calculates the halo's center of potential $\mathbf{R}$ from all N halo particles as

$$\mathbf{R} = \sum_{i=1}^{N} \frac{\mathbf{r}_i \Phi(\mathbf{r}_i)}{\Phi(\mathbf{r}_i)}, \tag{4.1}$$

where $\mathbf{r}_i$ is the position vector of particle i and Φ is the gravitational potential. This can be done in iterations to better capture the position of the center. A common method is to use shrinking spheres. In that method $\mathbf{R}_{\rm new}$ is calculated from all particles for which $|\mathbf{r}_i - \mathbf{R}| < r_{\rm th}$, where $r_{\rm th}$ is a threshold radius that is lowered in each iteration. Once the center of potential has been calculated to the desired precision, the external acceleration is added to the gravitational acceleration of each DM particle. For instance, to model the gravitational impact of a baryonic Hernquist sphere,

$$\mathbf{a}_{\rm i,ext} = -\frac{GM(t)}{(r_i + r_s)^2} \frac{\mathbf{r}_i}{r_i}. \tag{4.2}$$

Here, r_s is the scale radius of the Hernquist sphere and $M(t)$ is its time-dependent mass. In the scheme of Garrison-Kimmel et al. (2013), $M(t)$ is a periodic function that grows linearly up to some predefined value before immediately dropping to zero. The energy injection into the DM halo during each such explosion cycle is roughly equal to the binding energy of the external potential just before the drop in mass.

4.1.2 Modeling SN feedback in hydrodynamic simulations

Evolution of the ISM and star formation

Aside from the evolution of DM, hydrodynamic simulations also follow the evolution of baryons. In the early Universe, the baryonic content is a primordial composition of helium and hydrogen. The dynamical evolution of this primordial fluid is described by the Euler equations, coupled to the Poisson equation to take into account gravity. Solving these coupled differential equations is the main goal of cosmological hydrodynamical simulations and there are several numerical schemes that can be adopted to achieve this goal. Euler's equations can be expressed in several forms, and solved with different numerical schemes (Vogelsberger et al. 2020). AREPO (Springel 2010), the simulation code used in this thesis, numerically solves Euler's equations in the Arbitrary Eulerian-

Lagrangian formulation:

$$\frac{d}{dt}\int_V \rho\, dV = -\int_S \rho(\mathbf{v}-\mathbf{w})\cdot\mathbf{n}\, dS \tag{4.3}$$

$$\frac{d}{dt}\int_V \rho\mathbf{v}\, dV = -\int_S \rho\mathbf{v}(\mathbf{v}-\mathbf{w})\cdot\mathbf{n}\, dS - \int_S P\mathbf{n}\, dS \tag{4.4}$$

$$\frac{d}{dt}\int_V \rho e\, dV = -\int_S \rho e(\mathbf{v}-\mathbf{w})\cdot\mathbf{n}\, dS - \int_S P\mathbf{v}\cdot\mathbf{n}\, dS. \tag{4.5}$$

In equations 4.3, 4.4, and 4.5, ρ is the gas density, P is the pressure, $e = u + 1/2\mathbf{v}^2$ is the energy density, $\mathbf{n}$ is the unit vector which is orthogonal to the integration surface S, and u is the thermal energy per unit mass. The set of equations is closed by the equation of state $P = (\gamma-1)\rho u$, with the adiabatic index $\gamma = 5/3$. The first step towards solving Euler's equations in AREPO is the construction of Voronoi cells from a set of mesh-generating points which are allowed to move according to the flow of the gas. Equations 4.3, 4.4, and 4.5 are essentially conservation laws for mass, momentum, and energy, formulated in the finite volume V of one Voronoi cell with surface area S whose generating point is moving with a velocity $\mathbf{w}$. The right-hand sides are expressions for the mass, momentum, and energy flux out of the Voronoi cell. Of course, in order to evaluate the time evolution of mass, momentum, and energy, the fluxes into the cell (coming from neighbouring cells) also need to be taken into account. AREPO uses a second-order accurate MUSCL-Hancock scheme (e.g. Toro (1997)), in which a Riemann solver is used to estimate the time-averaged flux across cell surfaces for each time-step (Springel 2010).

Pressure opposes gravitational collapse. Since the pressure is proportional to the internal energy (and hence the temperature) of the gas, hydrodynamical simulations of galaxy formation need to include mechanisms for gas to cool down in order for it to collapse to large enough density for star formation to occur. Cooling processes like collisional excitation, ionization, inverse Compton, free free emission, and recombination are included in hydrodynamic simulations via cooling functions which couple to the energy equation 4.5. Primordial gas is usually assumed to be optically thin and in ionization equilibrium. For such a gas composition, the above mentioned processes dominate the cooling function. Later on, metal line cooling (cooling due to heavy elements) becomes important and dominates the cooling function for gas of temperatures between 10^5 K and $10^7\,K$. Most simulations do not explicitly model the cold phase ($T < 10^4\,K$) of the interstellar medium, since accurately following the evolution of such cold gas requires very small time-steps. Instead, they describe the dense gas phase through an effective polytropic equation of state, $T \propto \rho^{\gamma(\rho)}$, a form that captures the coexistence of a hot and voluminous gas phase with a cold and massive one (Springel and Hernquist 2003, Agertz et al. 2011, Dalla Vecchia and Schaye 2012). Recently, however, increasing effort has been made to resolve the multiphase structure of the interstellar medium (ISM) (Hopkins et al. 2012, Agertz et al. 2013, Rosdahl et al. 2015). ISM models that resolve the cold phase also need to include additional cooling processes that allow the gas to cool down to temperatures below $10^4\,K$, such as fine-structure cooling and molecular cooling. Moreover, self-shielding of molecular gas from radiation and several feedback channels need to be taken into account. Due to the large variety of timescales involved, physical processes that happen on scales below the resolution limit need to be

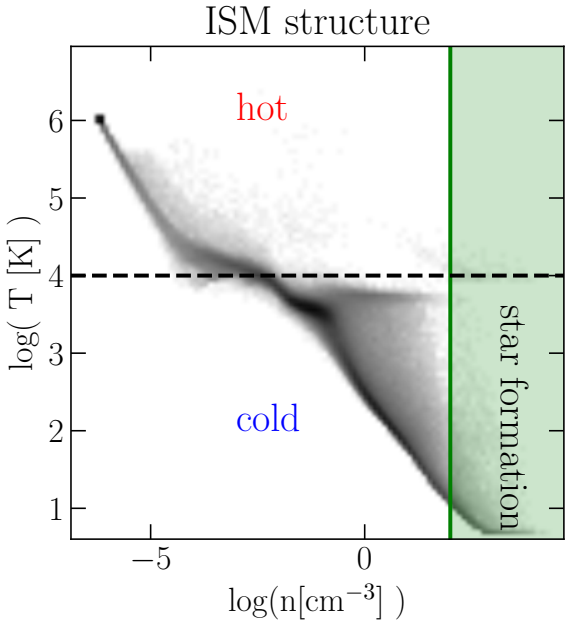

Figure 4.1. Structure of the ISM in a hydrodynamic simulation of an SMC-like system simulated using SMUGGLE (Marinacci et al. (2019)). Two gas phases coexist, a light and hot phase which occupies a large volume, and a heavy and cold phase which is very dense and concentrated. The vertical green line denotes the numerical density threshold for star formation in the simulation. Theoretically, gas of all temperatures whose density is larger than this threshold is eligible for star formation. However, only gravitationally collapsing gas can form stars, a criterion which is more likely to be fulfilled by cold gas.

averaged at the resolved scales and included in the simulation as *subgrid physics*. This averaging process varies between simulations and there is no clear bottom-up approach in which the descriptions of *subgrid physics* adopted in cosmological simulations are self-consistently derived from the results of simulations at smaller scales.

In the work of this thesis, I use the moving mesh code AREPO with the ISM model SMUGGLE (Marinacci et al. 2019), which aims at resolving the structure of the ISM by allowing the gas to cool down to $\sim 5\,K$. The resulting structure of the ISM in a system similar to the Small Magellanic Cloud (DM halo of $\sim 2 \times 10^{10} \mathrm{M}_{\odot}$) is shown in figure 4.1, where we show a $n-T$ diagram derived from the gas cells in the 3 Gyr snapshot of the CDM simulation with $n_{\mathrm{th}} = 100\,\mathrm{cm}^{-3}$ described in section 4. We can clearly see the co-existence of a hot ($T > 10^4\,K$) and a cold gas phase. As the gas cools down, it condenses to ever larger densities. In the simulation shown here, the gas can reach number densities $n > 100\,\mathrm{cm}^{-3}$, at which point it becomes eligible for star formation. As in many ISM codes, star formation in SMUGGLE occurs probabilistically, and

stellar mass is formed from gas with $n > n_{\rm th}$ at a rate

$$\dot{M}_\star = \varepsilon \frac{M_{\rm gas}}{t_{\rm dyn}}, \tag{4.6}$$

where $M_{\rm gas}$ is the mass of a given gas cell, $t_{\rm dyn}$ its dynamical time, and $\varepsilon = 0.01$ is an efficiency parameter. Aside from the requirement that the gas needs to be denser than a numerical threshold value, star formation is also limited to regions which are gravitationally collapsing. If both conditions are fulfilled, the fraction of the gas mass in the given cell that should be converted into stellar mass in a given time-step Δt is given by

$$p = 1 - \exp\left\{-\frac{\dot{M}_\star \Delta t}{M_{\rm gas}}\right\}. \tag{4.7}$$

In the code, whether or not a gas cell is converted into a star particle is decided by comparing this probability to a random number drawn from a uniform distribution in the interval $(0, 1)$. A star particle is formed if p is the larger of the two numbers. If a star particle is formed, it inherits velocity, position, and metallicity from the parent gas cell.

SN feedback and core formation

Whether or not SN feedback forms cores in hydrodynamic simulations of galaxy formation depends mostly on the assumed value of the star formation threshold. As displayed in figure 4.1, hot gas needs to cool down over time to reach large enough densities to be eligible for star formation. Shifting the star formation threshold to smaller values decreases this cooling time and makes star formation more continuous and less spatially concentrated. For larger star formation thresholds, gas tends to dominate the gravitational potential at the center of the galaxy between SN cycles. Rapid outflows of gas thus have a stronger impact on the evolution of the central gravitational potential. Recently, Benítez-Llambay et al. (2019) found that increasing the star formation threshold in the simulations of Bose et al. (2019a) will result in cores formed through SN feedback. Moreover, Dutton et al. (2020) have argued that above a certain numerical star formation threshold value, hydrodynamic simulations will converge towards predicting cored density profiles for dwarf-sized haloes.

After the first generation of stars has formed, star formation is regulated by SN feedback. The injection of energy, momentum, mass, and metallicity is modeled differently between codes. The energy, for instance, can be coupled kinetically or thermally (Vogelsberger et al. 2020). Aside from SN feedback, SMUGGLE also takes into account feedback from stellar winds and radiation. How exactly SN feedback is modeled in SMUGGLE is explained in detail in section 2.3 of Marinacci et al. (2019). In short, the number of type II supernovae is calculated directly from an assumed Chabrier (2001) initial mass function, whereas the number of type Ia supernovae is calculated from a delay-time distribution. If one or several supernovae occur, mass, momentum, energy, and metallicity is injected into neighbouring mass cells. The canonical thermal energy of one supernova is taken as $E_{\rm SN} = 10^{51}$ergs. In simulations of individual supernovae, the energy is found to be narrowly distributed around this value. In case the

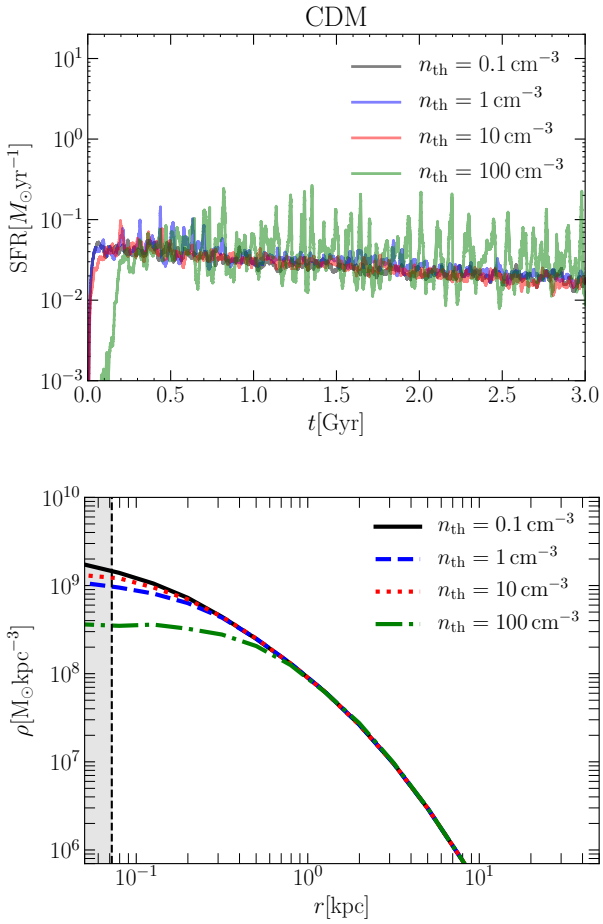

Figure 4.2. Core formation through SN feedback in hydrodynamic simulations. The top panel shows the star formation rate over 3 Gyr *of simulated time for three simulations with identical initial conditions (see text) and different star formation thresholds (adopted from chapter 8). The bottom panel shows the corresponding spherically averaged DM density profiles at* $t = 3\,{\rm Gyr}$. *Notably, the burstier star formation is, the more cored the DM halo is. Larger star formation rates generally correlate with burstier star formation. However, note that the star formation rate in the* $n_{\rm th} = 10\,{\rm cm}^{-3}$ *run settles to a rather steady and low value after a single initial spike. A likely explanation is that a significant fraction of the star-forming gas was expelled during this initial feedback cycle.*

cooling radius from individual supernovae cannot be resolved, the momentum injected into the neighbouring gas cells is boosted to account for the adiabatic Sedov-Taylor expansion. Ejected mass and metallicity are distributed among a pre-defined number of neighbours:

$$N_{\rm ngb} = \frac{4\pi}{3} h^3 \sum_i W(|\mathbf{r}_i - \mathbf{r}_s|, h). \tag{4.8}$$

Here, W is the cubic spline SPH kernel (Monaghan and Lattanzio 1985), $\mathbf{r}_s$ is the position of the exploding star particle, and $\mathbf{r}_i$ is the position of the i'th neighbouring gas cell. Equation 4.8 defines a search radius h that gives the distance between the star particle and the furthest away gas cell to which ejected mass and metallicity is distributed. In case h is larger than the radius of a superbubble, $R_{\rm SB} \sim 0.86$ kpc, feedback energy and momentum are only distributed among gas cells which lie within the superbubble radius. In summary, supernova energy and momentum are distributed among gas cells with $|\mathbf{r}_s - \mathbf{r}_i| < h_{\rm coupling}$, with $h_{\rm coupling} = \min(h, R_{\rm SB})$. Each one of those gas cells then receives a fraction $\tilde{w}_i = w_i / \sum w_i$ of the total feedback energy and momentum, where

$$w_i = \frac{1}{2} \left\{ 1 - \frac{1}{[1 + A_i/(\pi |\mathbf{r}_i - \mathbf{r}_s|^2)]^{1/2}} \right\} \tag{4.9}$$

is a weight based on the distance between star particle and gas cell and on the gas cell's area A_i (see also Hopkins et al. 2018, Smith et al. 2018). The injection of momentum, energy, mass, and metallicity is calculated in the rest frame of the star particle and the momenta are directed radially outward. After SN momentum and energy have been distributed among the neighbouring gas cells, the cells' final momenta and energies are transformed back into the simulation frame of reference.

Figure 4.2 illustrates SN feedback driven core formation in DM haloes. In the top (bottom) panel we show the star formation rate over time (DM density profile at $t = 3\,$Gyr) of the four CDM simulations from the simulation suite introduced in section 9.2. The simulations model the evolution of an isolated dwarf galaxy resembling the Small Magellanic Cloud within a live DM halo. Different lines in the respective panels correspond to runs with different star formation thresholds. In the upper panel, we see that star formation is significantly more bursty in the $n_{\rm th} = 100\,{\rm cm}^{-3}$ run than in any other run. As a consequence, we find in the bottom panel that the corresponding DM halo has formed a constant density core, while the DM haloes in the other runs have remained relatively cuspy. Notably, we find that star formation is more bursty for $n_{\rm th} = 1\,{\rm cm}^{-3}$ than it is for $n_{\rm th} = 10\,{\rm cm}^{-3}$. This is also reflected in the final DM profiles, with the latter being slightly cuspier than the former. The reason for that is that star formation becomes more stochastic at larger star formation thresholds. Starbursts can lead to gas depletion and in turn reduce the amount of further star formation significantly. Crucially, however, star formation tends to become burstier at larger star formation thresholds on average (see **Extended Data Figure** 9.1). Therefore, the efficiency of SN feedback as a core formation mechanism depends on the star formation threshold (see also Dutton et al. 2020).

4.2 Self-Interacting Dark Matter

Self-interacting DM (as proposed by Spergel and Steinhardt 2000), can interact with itself through some new dark sector force (see Tulin and Yu 2018), which does not have to imply that a coupling between the DM and baryonic matter exists. The strength of DM self-interactions can simply be characterized by means of an interaction cross section σ. This implies a mean free path $\lambda \sim 1/(n\sigma) = m_\chi/(\rho\sigma)$, where n is the (unknown) number density of DM particles and ρ is the DM mass density, which can be calculated from simulations. Since the mass of the DM particles is unknown it is common to parametrize the strength of DM self-interactions in terms of the cross section per unit mass, such that the mean free path is given by the inverse product of σ/m_χ and the known mass density ρ. Depending on the value of the so-called Knudsen number $\lambda(r)/r$, where r is the radial distance from the center of a DM halo, SIDM can either behave like a fluid (tightly coupled regime, $\lambda(r)/r \ll 1$) or like an optically thin gas (rarefied regime, $\lambda(r)/r \gg 1$). In the latter case, self-interactions are well described as elastic scattering events between DM particles (Koda and Shapiro 2011). This is the regime which is of interest for astrophysical SIDM. Modeling self-interactions in simulations becomes difficult if σ depends on the scattering angle. To avoid such complications, most simulations assume that the scattering is isotropic and then capture the dependence on the scattering angle by using the transfer cross section (e.g Kahlhoefer et al. 2014),

$$\sigma_T = 2\pi \int_0^\pi \frac{d\sigma}{d\Omega}(1-\cos\theta)\sin\theta\, d\theta, \tag{4.10}$$

which estimates the average momentum transfer during a collision (Tulin and Yu 2018). Here, θ is the scattering angle in the center of mass frame and $d\sigma/d\Omega$ is the standard differential cross section obtained from a particle physics model. Constraints in section 3.2.5 refer to σ_T/m_χ, though it is worth noting that in the s-wave limit $\sigma_T = \sigma$.

In my work, I use the algorithm of Vogelsberger et al. (2012) to model SIDM. In their algorithm, self-interactions are implemented as elastic scattering events between simulation particles. Whether and between which particles scatters occur in a given time-step is calculated using a Monte Carlo method. The probability that a particle i scatters with particle j (from a total of k neighbours) is

$$P_{ij} = m_i W(r_{ij}, h_i)\frac{\sigma_T}{m_\chi} v_{ij}\Delta t_i, \tag{4.11}$$

where m_i is the mass of one simulation particle, W is the standard cubic spline kernel, h_i is a smoothing length that encloses the k nearest neighbours, Δt_i is the time-step of particle i, r_{ij} is the distance between particles i and j, and $v_{i,j}$ is the relative velocity between those two particles. σ_T can either be constant or a function of the relative velocity. An extra accuracy criterion is used to ensure that the time-steps of individual particles are short enough to avoid multiple scattering. Whether particle i scatters with any one of its nearest neighbours during the time-step Δt_i is determined by comparing $P_i = 1/2\sum_j P_{ij}$, which is particle i's total probability for scattering during time-step Δt_i to x, which is a number drawn at random from a uniform distribution in $(0,1)$. If $P_i > x$, then particle i will scatter with one of its k neighbours, which are sorted by $r_{i,j}$,

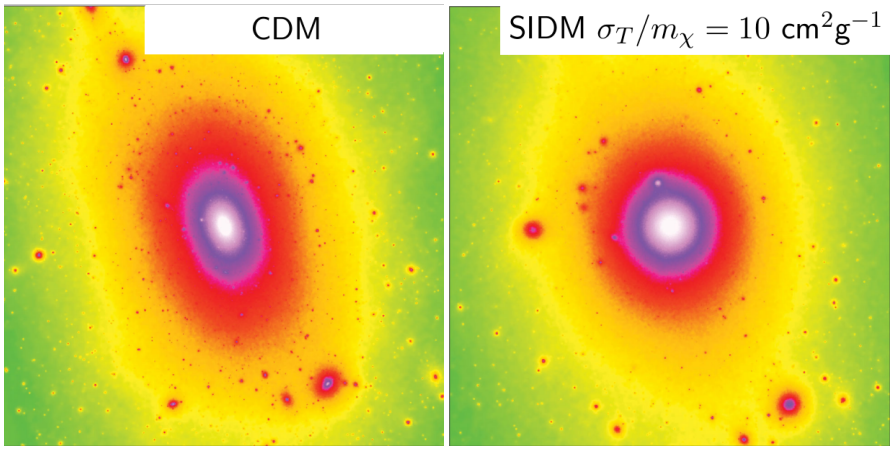

Figure 4.3. Comparison of the densities of a MW-size CDM halo (left panel) and the same halo in SIDM with $\sigma_T/m_\chi = 10\mathrm{cm}^2\mathrm{g}^{-1}$ (right panel), projected into 2d. The density of the CDM halo rises steeper towards the halo's center, whereas the SIDM halo has a constant density core. Moreover, the inner region of the SIDM halo is much closer to spherical symmetry than the inner region of the CDM halo. Figure adapted from Vogelsberger et al. (2012)[1].

starting with the one closest to particle i. Subsequently, the first neighbouring particle l for which $x \leq \sum_{j=1}^{l} P_{ij}/2$ is chosen as the scattering partner of particle i. Once two scattering partners have been chosen, they are assigned new velocities

$$\mathbf{v}_i = \mathbf{v}_{\mathrm{com}} + (v_{ij}/2)\mathbf{e} \tag{4.12}$$

$$\mathbf{v}_j = \mathbf{v}_{\mathrm{com}} - (v_{ij}/2)\mathbf{e}. \tag{4.13}$$

In equations 4.12 and 4.13, $\mathbf{v}_{\mathrm{com}}$ is the velocity of the center of mass of the two particles and $\mathbf{e}$ is a randomly selected unit vector. Total momentum and energy are conserved in this scheme, whereas conservation of angular momentum is violated.

Figure 4.3 shows how the inner structure of a MW-size halo is altered by SIDM with $\sigma_T/m_\chi = 10\mathrm{cm}^2\mathrm{g}^1$. Note that this value is excluded by the mass distribution in the Bullet cluster (Randall et al. 2008) and the measured ellipticities of galaxies (Peter et al. 2013) and is shown here for illustration only. The left panel of figure 4.3 shows the projected inner density of the Aquarius-A (CDM) halo (Springel et al. 2008), whereas the same quantity is shown on the right panel for a resimulated version of the same halo in SIDM. At the center, the SIDM halo is more isotropic and spherically symmetric than its CDM counterpart. Moreover, while the density of the CDM halo continuously increases towards the center, the SIDM halo has an extended constant density core.

[1] Original: Fig. 3 of Subhaloes in self-interacting galactic dark matter haloes, Vogelsberger, M., Zavala, J., and Loeb, A., MNRAS 423

4.3 Stellar dynamics

This section summarizes a few important concepts from stellar dynamics that form the basis of a substantial part of the work presented in this thesis. The discussion here is kept rather short and limited to the concepts needed to understand the analysis presented. For a more comprehensive overview and presentation of topics related to stellar dynamics see Binney and Tremaine (2008).

4.3.1 Action-angle variables

Suppose we follow the orbit of a single particle in a static gravitational potential. The orbit is defined by the six-vector $\mathbf{w}(t) = (\mathbf{x}(t), \mathbf{v}(t))$, which gives the particle's trajectory in phase space. Now suppose that $I(\mathbf{w}(t))$ is a function such that $I(\mathbf{w}(t_1)) = I(\mathbf{w}(t_2))$ for any t_1 and t_2. Such a function is called an integral of motion.

Integrals of motion can either be isolating or non-isolating. A careful definition of these terms can be found in Binney and Tremaine (2008). Effectively, each isolating integral reduces the dimensionality of the hypersurface in phase space to which an orbit is confined by one. Non-isolating integrals, on the other hand, have no such effect on the structure of orbits. In other words, if a three-dimensional potential admits a number n ($n < 6$) of isolating integrals, then the motion of particles orbiting in this potential is confined to a $6-n$-dimensional hypersurface in phase space. Examples for isolating integrals are for instance the particle's energy (in static potentials) and its angular momentum in case the potential is spherical. If a particle's orbit in a k-dimensional potential admits at least k independent isolating integrals of motion, it is called regular, otherwise irregular. For regular orbits, there is a set of canonical coordinates called action-angle coordinates. In general, regular orbits form orbital tori of dimension $6-k$ in phase space. Such tori have a set of $6-k$ "irreducible, irreconcilable closed paths" (Binney and Tremaine 1987, page 169). Actions are defined as integrals over one of these paths γ_i:

$$J_i = \frac{1}{2\pi} \oint_{\gamma_i} \mathbf{v} \cdot d\mathbf{x}. \tag{4.14}$$

Thus, a given set of actions represents a set of closed curves on the orbital torus. The location of the particle on each of these curves is given by the angles, the coordinates conjugate to the actions. Incrementing an angle by 2π one arrives at the same point on the curve. Written in action-angle coordinates, the Hamiltonian is a function of the actions only:

$$H(\mathbf{x}, \mathbf{v}) \rightarrow H(\mathbf{J}). \tag{4.15}$$

Then, following Hamilton's equations we get

$$\dot{J}_i = -\frac{\partial H}{\partial \theta_i} = 0 \tag{4.16}$$

$$\dot{\theta}_i = \frac{\partial H}{\partial J_i} = \Omega_i(\mathbf{J}) \tag{4.17}$$

for the evolution of the actions and the angles. Evidently, actions are integrals of motion. Angles grow linearly,

$$\theta_i = \theta_{i,0} + \Omega_i(\mathbf{J})t, \tag{4.18}$$

and since actions are constant, angles grow at a constant rate. Since angles are 2π-periodic, regular orbits can be expanded into a Fourier series:

$$\mathbf{x}(t) = \sum_{\mathbf{k}} \mathbf{X}_{\mathbf{k}}(\mathbf{J}) \exp\left[i\mathbf{k}\cdot\theta\right], \tag{4.19}$$

where $\mathbf{k}$ is a triple index. Such a Fourier expansion can be used to prove the important time averages theorem, which states that for incommensurable frequencies Ω_i, the "time that the phase point of a star on a regular orbit spends in any region D of its torus is proportional to the integral $V(D) = \int_D d^3\theta$ through D" (Binney and Tremaine 1987, page 171).

From this follows a statement which is of great importance in this thesis. Suppose a set of particles with regular orbits is on a common orbital torus, i.e. they have the same actions but different angles. Now suppose that the gravitational potential Φ_0 is continuously changed into a different potential Φ_1 whose orbits are also regular. In that case, particles which were initially on the same orbital torus of Φ_0 will in general end up on different orbital tori of Φ_1. However, if the change of the potential happens on a timescale which is much longer than $2\pi/\Omega_k$, where Ω_k is the lowest orbital frequency, then all particles which were on the same torus initially will also end up on the same torus. This is true because the time averages theorem shows that the fraction of time spent in a particular portion of the orbital torus is the same for all particles. Using the Poincaré invariant theorem, one can deduce from this statement that actions are invariant under sufficiently slow evolution of the gravitational potential. They are so-called adiabatic invariants.

While proving that actions exist for regular orbits is easy, actually calculating actions is rather difficult. In general, actions are most conveniently defined along paths γ_i that are related to the symmetry of the potential. For axisymmetric systems, for instance, one can easily define radial, vertical, and azimuthal actions. Our work deals with a changing potential in a DM halo. If we approximate such a halo as spherically symmetric, then we can work in spherical coordinates. Two of the three actions belonging to spherical coordinates are related to the angular momentum. We are interested in the third one, the radial action, which for spherical potentials $\Phi(r)$ is

$$J_r = \frac{1}{\pi}\int_{r_{\text{peri}}}^{r_{\text{apo}}} dr\,\sqrt{2(E-\Phi(r)) - L^2/r^2}. \tag{4.20}$$

Here, $r_{\text{peri/apo}}$ are the peri- and apocenter radius and functions of E and L, the energy and the magnitude of the angular momentum only. Motion in a spherically symmetric potential is essentially a one-dimensional problem, as the motion of a particle is confined into a plane perpendicular to the angular momentum vector and the relevant dynamics are captured by $r(t)$ and $\dot{r}(t)$, i.e. the radial coordinate and its conjugate momentum. Hence, particles with common integrals of motion E and L lie on a common closed curve in $r-\dot{r}$ space. According to equation 4.20, they also have the same radial action.

Now, if the underlying gravitational potential changes very slowly from $\Phi_0(r)$ to $\Phi_1(r)$ compared to all dynamical timescales, then we know that radial actions are conserved. Moreover, if spherical symmetry is not violated during the transition, angular momentum is conserved. From action conservation and equation 4.20 we then know that particles which initially had the same energy will also have the same energy after the transition. They will once again lie on a common closed curve in $r - \dot{r}$ space. This statement is not true if the potential changes very impulsively. In fact, Pontzen and Governato (2012) showed that in a one-dimensional potential which changes instantaneously the resulting energy change of kinematic tracers depends explicitly on the tracers orbital phase at the time of the change.

In this thesis I aim to characterize how the orbits of tracers in spherical potentials respond if the potential is changed adiabatically or impulsively, e.g. through SIDM in the former or SN feedback in the latter case. This behaviour is tested mostly in N-body simulations of an isolated halo. Since numerical inaccuracies can affect the calculation of a particle's energy and angular momentum, I work with the concept of an orbital family. Loosely defined, an orbital family is a set of particles with similar integrals of motion E and L. Phase mixed orbital families populate a slightly spread out, closed curve in $r - \dot{r}$ space.

4.3.2 Phase mixing

Consider an ensemble of collisionless tracers, i.e. particles that do not collide amongst each other and orbit in a fixed gravitational potential. The evolution of the distribution function of such an ensemble is fully determined by the orbits of all particles, and as such by the initial conditions. However, a process of collisionless relaxation can be observed in the distribution of particles in phase space. This process is called phase mixing and is best described by looking at the coarse-grained distribution function $\bar{f}$, which is a local mean of f, the fine-grained distribution function. Consider the ensemble of particles shown in the upper left panel of figure 4.4. They are initially set up with similar phase space coordinates q and p. Evidently, the initial distribution of particles covers only a small fraction of the phase space that can be accessed by the particles. It is initially only in this region that the coarse-grained distribution function is non-zero. Over time, however, the particles start to spread out in phase space, due to their slightly different energies, and hence orbital times. This leads to the emergence of hierarchical spiral features as seen in the top right and bottom left panels of figure 4.4 (see also Afshordi et al. 2009). Eventually, the particles are evenly distributed over the whole phase space that they can access. At that point, evolution of the coarse-grained distribution function stops and the final $\bar{f}$ is smooth over the whole accessible phase space volume. This is the final result of phase mixing.

Let us consider again the dynamical evolution of an orbital family in a spherical potential. If the potential changes slowly, particles on similar orbits should remain on similar orbits. However, if the potential changes impulsively the energy of individual tracers should change by a different amount depending on their orbital phases. This would then create a spread in the energy distribution and should lead to emerging signatures of phase mixing in the projection of the distribution function into radial phase space that resemble the spiral structures in figure 4.4. However, if these structures

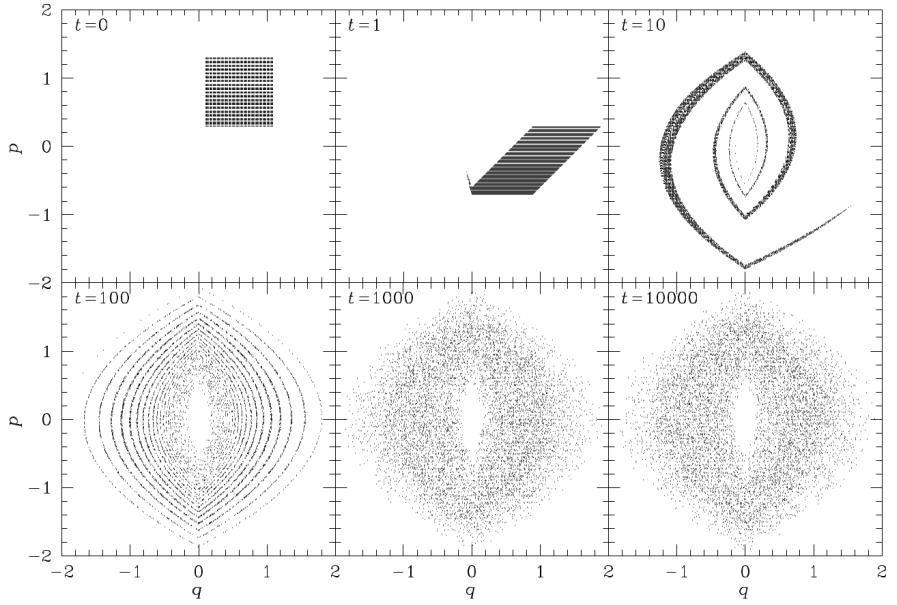

Figure 4.4. Phase mixing of 10^4 *particles in the 1d-Hamiltonian* $H = \frac{1}{2}p^2 + |q|$*. Phase mixing is best described as evolution of the coarse-grained distribution function. Initially, it is non-zero only within a tiny phase space volume. Eventually, it evolves towards a steady state and becomes smooth over the entire phase space area which is accessible to the* 10^4 *particles. Note the hierarchical formation of shells in the upper right and the lower left panel. The phase-mixed final state arises when the number of formed shells becomes comparable to the number of particles. Phase mixing cannot be described using the fine-grained distribution function since individual particles are never stationary. Figure adapted from Dehnen and Read (2011)*[2].

emerge, deviations from spherical symmetry in real haloes may result in them being rather short-lived (Pontzen et al. 2015).

4.4 Original articles

Here I briefly introduce the five articles that make up the main body of this thesis. This introduction is limited to stating the scope of the respective article and how the article relates to topics discussed in chapters 2, 3, and 4. The original articles constitute chapters 5-9 of the thesis.

[2]Reprinted by permission from Springer Customer Service Centre GmbH: Springer Nature, The European Physical Journal Plus, N-body simulations of gravitational dynamics, Dehnen, R., Read, J., ©2011

4.4.1 Chapter 5

This article is motivated by the cusp-core problem in dwarf galaxies (section 3.2.5). We explore two mechanisms of core formation in a dwarf size halo by means of high-resolution DMO simulations of an isolated, initially cuspy dwarf set up in dynamical equilibrium using Eddington's formula (2.19). We simulate core formation, either through self-interactions between the DM particles (section 4.2) or modeling SN feedback as a periodic impulsive removal of mass (section 4.1.1), and then investigate the differences in the response of three different distributions of kinematic tracers, namely a Gaussian distribution in radial action, an orbital family (section 4.3.1), and a Plummer sphere - initially in dynamical equilibrium - to model the stellar distribution of a dwarf galaxy. All three distributions evolve differently, depending on whether core formation is adiabatic or impulsive, in line with expectations from section 4.3.1. One important result is that we find clear signatures of the inital stages of phase mixing (section 4.3.2), namely shell-like features in the projection of the orbital family's distribution function into the $r - \dot{r}$ plane, as a result of the repeated mass removals in the impulsive core formation scenario, whereas no such features are found in the adiabatic case. Finally, we investigate the impact of deviations from spherical symmetry by repeating our experiments in a halo with a cosmological merger history. Consistent with Pontzen et al. (2015), we find that features of ongoing phase mixing in $r - \dot{r}$ space remain for only ~ 1 dynamical time if spherical symmetry is violated.

4.4.2 Chapter 6

We investigate how radial actions evolve in time-dependent spherical potentials, making use of the concept of dynamical invariants introduced in Peñarrubia (2013) which is based on a coordinate transformation to a time-independent frame originally introduced by Lynden-Bell (1982). We show that the action invariant is given by the action calculated in the static frame and relate this action invariant to the measured action in the original time-dependent frame by performing a Taylor expansion. The first-order correction to the action-invariant is an oscillatory term with zero mean. The period of the oscillation is simply the radial period of the particle and the first-order correction term oscillates in-phase with the particle's radial motion. Since actions are known to be invariant under adiabatically changing potentials, we can define the adiabatic regime as the regime in which the first-order correction to the action invariant ΔJ_r is much smaller than the action invariant, $J_{r'}$, itself, $\Delta J_r \ll J_{r'}$. We find that this definition recovers the condition under which the time averages theorem holds (section 4.3.1), namely that the potential must change on timescales which are much longer than the dynamical timescale.

Akin to Peñarrubia (2015), we develop a diffusion theory for radial action distributions in the regime in which $\Delta J_r \lesssim J_{r'}$. Radial action distributions of tracer particles are treated as a microcanonical ensemble whose evolution is fully determined by the average evolution of particles with a given radial action. This is a first step towards a statistical theory for the evolution of particle ensembles in the presence of long range forces (see section 3.3)[3] We find that the results of numerical experiments simulating the evolution

[3] In order to obtain a complete theory, we need to be able to self-consistently predict the evolution of the

of the radial action distributions of ensembles of kinematic tracers in a time-dependent Kepler potential are well described by our theory. The accuracy of the diffusion theory decreases as the first order correction becomes comparable in magnitude to the action invariant and eventually breaks as we enter the impulsive regime in which the evolution of radial actions cannot be calculated perturbatively.

Using our novel diffusion theory, we estimate that diffusive effects due to an increasing potential well should be significant within a sizeable but monotonously decreasing fraction of the phase space covered by the MW-halo's distribution function. Diffusion of radial actions could, for instance, be an additional source of uncertainty in the analysis of tidal streams in the MW. Moreover, we hypothesize that the rapid accretion at the time at which haloes first collapse could create a "cusp" particle population with low angular momentum, and that this may be the reason that the ubiquitous cusps (see equation 2.15) are not recovered when trying to predict the structure of DM haloes using a maximum entropy approach (Pontzen and Governato 2013). Finally, we find that radial action distributions of kinematic tracers are invariant under SIDM induced core formation, provided the self-interaction cross section lies within current constraints (see section 3.2.5).

4.4.3 Chapter 7

In this article we develop an effective model for SN feedback in DMO simulations (see section 4.1.1). Changing the model parameters allows us to investigate how the total feedback energy, the timescale of SN cycles, and the spatial distribution of baryons affect how effective SN feedback is at forming a core in a dwarf-sized DM halo. We test our effective model in a suite of high-resolution simulations of an isolated dwarf-sized halo. In the halo center, we add the potential of either an external Plummer sphere or an external disk. We then simulate how SN feedback from those two galaxies affects the DM density in the central halo for different values of the nominal coupling efficiency of feedback energy to the ISM, the timescale of SN cycles, and the half-light radii of the galaxies. We find that a more energetic, more impulsive, and more concentrated SN feedback energy injection is more effective at forming cores. Crucially, whether or not cores form is primarily determined by the total feedback energy injected into the DM. Finally, we also look at the evolution of an orbital family of kinematic tracers. We find signatures of phase mixing in the $r - \dot{r}$ projection of the tracer's distribution function if SN feedback is sufficiently impulsive[4]. In general, we find that SN feedback strong enough to form a core is always accompanied by signatures of phase mixing, implying that the potential needs to change impulsively for SN feedback to form cores. However, signatures of impulsive SN feedback do not guarantee the presence of a core.

gravitational potential. It may be possible to do this in a MCMC-like manner, i.e. by calculating the response of the particle distribution to a changed potential and then updating the potential in short alternating steps (see also the appendix of Peñarrubia 2015).

[4]Similar to the results of article 1, we find that symmetry matters. Phase mixing progresses much slower if the external potential is spherically symmetric. The disk potential causes much faster mixing and thus erases the shell-like features of phase mixing in the $r - \dot{r}$ plane.

4.4.4 Chapter 8

This article is focused on SIDM and SN feedback as possible solutions to the cusp-core problem. Using the ISM model SMUGGLE (Marinacci et al. 2019, see section 4.1.2) and Vogelsberger et al. (2012)'s SIDM code (see section 4.2), we run a suite of 16 high-resolution hydrodynamic simulations of an isolated SMC-like dwarf galaxy embedded within an initially cuspy, dwarf-sized DM halo. DM and baryons are set up in dynamical (and hydrostatic) equilibrium and the evolution of the system is simulated over several gigayears. Each simulation has a unique combination of two independent model parameters, the (velocity-independent) SIDM transfer cross section per unit mass σ_T/m_χ and the gas density threshold for star formation $n_{\rm th}$. As discussed in section 4.2 and section 4.1.2, core formation is expected to be more effective for larger σ_T/m_χ and larger $n_{\rm th}$, since SN feedback becomes more impulsive for larger star formation thresholds. We choose $(0, 0.1, 1, 10)\,{\rm cm}^2/{\rm g}^{-1}$ for σ_T/m_χ and $(0.1, 1, 10, 100)\,{\rm cm}^{-3}$ for $n_{\rm th}$. We find that the DM halo forms constant density cores in several of these simulations and that different combinations of σ_T/m_χ and $n_{\rm th}$ can lead to similar-size cores.

The remainder of the article discusses how observations of the structural properties of dwarf galaxies, as well as the kinematics of stars and gas in such galaxies, may be used to break this degeneracy between adiabatically formed (SIDM) cores and impulsively formed (SN feedback) cores. We find that the galaxies in haloes with adiabatically formed cores are spatially more extended than those in haloes with impulsively formed cores (baryons follow an adiabatically evolving DM potential). Moreover, impulsive SN feedback tends to increase the amount of random motion in the gas, leading to (on average) smaller values of $v_{\rm los}/\sigma_{\rm los}$ in the center of the galaxy. Finally, we show that stellar age and metallicity gradients can be used to break the degeneracy between adiabatic and impulsive core formation. In particular, stars in simulations with larger star formation thresholds have steeper age gradients and shallower metallicity gradients than stars in simulations with smaller star formation thresholds.

4.4.5 Chapter 9

In this article we present a more in-depth analysis of the kinematics of the stars formed in the simulations presented in article 4. We demonstrate that stars with similar metallicities and similar formation times constitute constitute orbital families. We then show that in simulations with larger star formation thresholds, we can, in the aftermath of strong supernova activity, find shell-like signatures of phase mixing (similar to the features found in article 1) in the $R - \dot{R}$ slice of phase space, provided we look at the phase space distribution of stars of similar age and metallicity. While we model the DM halo as spherically symmetric, the SMC analog in our simulations is modeled as a rather massive disk galaxy. The shell-like signatures of early stage phase mixing are present despite this significant violation of spherical symmetry. Throughout the entire simulation time, we find no such signatures of phase mixing in the phase space of stars formed in simulations with less impulsive SN feedback, i.e. smaller star formation thresholds. On the other hand, they are ubiquitous if the simulated star formation is bursty, and SN feedback impulsive as a consequence. We conclude that not observing

such phase mixing signatures in dwarf galaxies with cored DM haloes and ongoing star formation would be a clear sign that SN feedback is not sufficiently impulsive to transform DM cusps into cores. At last, we show that such signatures of phase mixing should be clearly identifiable as clusters of high metallicity stars in the rotation curves of galaxies that are observed edge-on.

Chapter 5

Impulsive and adiabatic core formation

This chapter is based on the following article:

The nature of core formation in dark matter haloes: adiabatic or impulsive?

Published in Monthly Notices of the Royal Astronomical Society, Volume 485, Issue 1, May 2019, Pages 1008–1028

Authors:
Jan D. Burger[1] and Jesús Zavala[1]
[1]centre for Astrophysics and Cosmology, Science Institute, University of Iceland, Dunhagi 5, 107 Reykjavik, Iceland

It is well established that the central deficit of dark matter (DM) observed in many dwarf galaxies disagrees with the cuspy DM haloes predicted in the collision-less and cold DM (CDM) model. Plausible solutions to this problem are based on an effective energy deposition into the central halo with an origin that is either based on baryonic physics (e.g. supernova-driven gas blowouts; SNF) or on new DM physics (e.g. self-interacting DM, SIDM). We argue that the fundamental difference between the two is whether the process is impulsive or adiabatic, and explore novel ways to distinguishing them by looking at the response of stellar orbits. We perform idealised simulations of tracers embedded in a $1.48 \times 10^{10} \mathrm{M}_{\odot}$ spherical halo, and model the creation of a $\sim$ 1kpc DM matter core in SIDM with $\sigma_T / m_\chi = 2 \mathrm{cm}^2 \mathrm{g}^{-1}$ and injecting energy into the halo with a sudden mass removal equivalent to $\mathscr{O}(10\%)$ of the halo's potential energy. Choosing idealised initial orbital configurations for the tracers, we find that radial actions are conserved (changed) in the SIDM (impulsive) case. The adiabaticity of the

SIDM case prevents tracers from changing their orbital family during core formation whereas SNF separates tracers of initially the same family to a variety of orbits. We show that these key features remain in a cosmological halo, albeit for a few dynamical timescales. The number density and velocity dispersion profile of a Plummer sphere with $r_{1/2} = 500$ pc change only marginally under adiabatic core formation, whereas SNF causes a substantial expansion of the sphere driving it out of Jeans equilibrium. Our results point towards promising ways of differentiating adiabatic from impulsive core formation.

5.1 Introduction

The vast amount of evidence pointing towards the existence of dark matter (DM) remains purely gravitational and thus the nature of DM as a particle remains a mystery. Among the different avenues followed down to search for clues about the DM nature, a promising one is that of looking for the dynamical signature of non-gravitational DM physics in the physical properties of galaxies. In particular, dwarf galaxies are among the best targets to look for new DM physics since they are the most DM-dominated self-gravitating systems in the Universe. These galaxies have also been a constant point of contention over the last decades in the otherwise successful Cold Dark Matter (CDM) model of structure formation, where gravity is the only relevant DM interaction in the physics of galaxies (for a review see Bullock and Boylan-Kolchin 2017). It remains an open question whether the CDM model coupled with a complete description of gas and stellar physics, commonly referred to as *baryonic physics*, can fully explain the properties of dwarf galaxies. To this date, the possibility remains that non-gravitational DM physics plays a major role in galaxy formation and evolution.

A prominent feature in dwarf galaxies that requires explanation is the DM deficit within their innermost regions relative to the predictions of the CDM model (without baryonic physics). Traditionally, two distinct issues are commonly associated with this deficit, both of which might share a common solution. The better known of these is that many dwarf galaxies most likely reside in cored DM haloes as opposed to the centrally cusped CDM haloes, which is the so-called core-cusp problem (e.g. Moore, 1994; de Blok et al., 2008; Kuzio de Naray et al., 2008; Walker and Peñarrubia, 2011). The second, most recent issue is known as the too-big-to-fail-problem (Boylan-Kolchin et al. (2011); Papastergis et al. (2015)). It refers to an over-abundance of mass found within $\sim$1 kpc of a significant number of dwarf-size haloes in CDM, making their enclosed mass inconsistent with the internal kinematics of the dwarf galaxies they are naturally assumed to host.

Among the most viable mechanisms to explain the observed inner DM deficit in dwarf galaxies are those which efficiently deposit energy into the centre of DM haloes. In particular, we focus on two distinct mechanisms that have been studied extensively in the past and that represent solutions based on either baryonic physics or new DM physics. In both, dwarf-size haloes are initially cuspy and dynamically cold at their centre and evolve to develop a hotter constant density core within the innermost regions. The first is a (gravitational) mechanism of energy transfer to DM particles in the centre of haloes following the rapid removal of gas during outflows caused by supernovae

(supernova feedback; SNF) (e.g. Navarro et al., 1996a; Gnedin and Zhao, 2002; Pontzen and Governato, 2012; Di Cintio et al., 2014; Chan et al., 2015; Read et al., 2016; Tollet et al., 2016). The second is a mechanism of energy redistribution in the centre of haloes from the outside-in caused by (strong) elastic self-scattering between DM particles (SIDM; Spergel and Steinhardt, 2000; Yoshida et al., 2000; Davé et al., 2001; Colín et al., 2002; Vogelsberger et al., 2012; Rocha et al., 2013).

Both of these mechanisms of cusp-core transformation are currently plausible and efforts are being made to validate their key ingredients. In the case of SNF, although the evidence of gas outflows in galaxies is substantial (e.g. Martin et al., 2012), the efficiency of this mechanism in dwarf galaxies remains unclear. In particular, the efficiency depends on the bulk of the energy deposition occurring at much shorter time scales than the characteristic dynamical time of the galaxy Pontzen and Governato (2012). There is evidence that bursts of star formation (and subsequent supernovae) happen in time scales at least comparable to the dynamical time scale of galaxies with 10^8 $M_\odot$ in stellar mass Kauffmann (2014), but the supernova mechanism of cusp-core transformation requires the latter time scale to be considerably larger than the former. The situation is notably more uncertain for low-mass dwarf galaxies where the CDM problems mentioned above are more severe, and where the validation of the supernova mechanism requires a time resolution for bursts of star formation that cannot be achieved currently Weisz et al. (2014).

On the other hand, the viability of DM self-scattering as a cusp-core transformation mechanism depends on whether or not the cosmological SIDM model is observationally allowed overall. Substantial effort has been put in constraining SIDM using different astrophysical probes, with the stringent (and more reliable) constraints being at the scale of elliptical galaxies Peter et al. (2013) and galaxy clusters (e.g. Robertson et al., 2017, 2019) with the transfer cross section per unit mass limited to $\sigma_T/m_\chi \lesssim 2$ cm^2 g^{-1}; for a recent review on SIDM constraints see Tulin and Yu (2018). To put this value into the context of the CDM challenges in dwarf galaxies, it has been shown that for $\sigma_T/m_\chi \lesssim 0.1$ cm^2 g^{-1}, SIDM would be essentially indistinguishable from the CDM predictions Zavala et al. (2013). Until very recently, the SIDM model remained essentially unconstrained at the scale of dwarf galaxies, but now a couple of works suggest that $\sigma_T/m_\chi \lesssim 1$ cm^2 g^{-1} given the presence of a cusp in the Milky-Way satellite Draco Read et al. (2018), and the diversity of inner DM densities in the Milky-Way satellites (Zavala et al. in prep.). Both of these works, however, are based on SIDM simulations without taking into account baryonic physics. Although the effect the latter has on the DM mass distribution within the most massive Milky-Way subhaloes is expected to be small, it remains to be seen if these stringent SIDM constraints remain after this has been taken into account. Even if these constraints are validated, reaching the lower limit for the cross section for SIDM to be an alternative to CDM, 0.1 cm^2 g^{-1}, might prove to be quite challenging since full galaxy formation and evolution models within CDM and SIDM are expected to have global predictions for the galaxy population that have degeneracies (e.g. Harvey et al., 2018).

Given the current situation, it is crucial to identify ways in which the dark and baryonic mechanisms of energy transfer into the centre of dwarf-scale DM haloes can be differentiated in a definitive way. This objective has remained largely unexplored

and we here take the first steps towards making progress in this direction by identifying distinct dynamical signatures in stellar orbits as they respond to these mechanisms. We focus on the fundamental difference that arises between a transfer of energy that occurs adiabatically (as in SIDM) and one that occurs impulsively (as in SNF). Using controlled simulations of an isolated halo with an embedded stellar population, we show that it is possible to form a DM core with similar global properties (size and central density) through both mechanisms. Yet, the responses in the kinematics of stars exhibit dramatic differences.

This paper is organized as follows. In Section 5.2 we provide the physical basis of our work with a simple one-dimensional toy model presented in appendix 5.8, which illustrates the key differences between the adiabatic and impulsive cases. In Section 5.3, we describe the initial conditions of the DM halo and the modeling of the two cusp-core transformation scenarios in our simulations. In Section 5.4 we construct three different initial configurations for tracers (stars) that are added to our simulations in order to explore the insights obtained in Section 5.2. In Section 5.5, we discuss or main results. In Section 5.6 we discuss how abandoning the premise of an initially spherically symmetric isolated halo affects these results by analysing a cosmological halo. We then draw our conclusions in Section 5.7.

5.2 Adiabatic vs impulsive cusp-core transformation in DM haloes

In order for supernovae to efficiently transfer energy to the surrounding DM particles, the removal of gas following the explosion needs to occur impulsively, i.e., the gravitational potential needs to change rapidly relative to the orbital period of the DM particles. This phenomenon was extensively investigated in Pontzen and Governato (2012) and we follow their results in detail in this paper. We start by considering their toy model of a one-dimensional harmonic oscillator which transitions between two different frequencies (modelling a change in gravitational potential) during a fixed transition time interval. If this time interval is much longer (shorter) than the oscillation period, then the change in the potential is adiabatic (impulsive). This toy model is a nice illustration of the key physical differences between the adiabatic and impulsive cases and has the advantage of having solutions that are easily derived. We describe the model and its solution in both the impulsive and adiabatic cases in Appendix 5.8, while in the following we discuss the main insights we obtain from this exercise.

In the case in which the potential changes adiabatically, the oscillation transitions smoothly from the initial to the final frequency. The ratio of the final to initial amplitudes and the shift in the phase of the oscillation are both independent of the phase of the oscillation prior to the change in potential. They are completely determined by the frequency transition (see Eqs. 5.39 and 5.40), i.e., by the energy change in the potential. By contrast, if the potential changes impulsively, the transition is sudden and the change in amplitude and phase depends not only on the frequency (energy) change, but also explicitly on the initial phase (see Eqs. 5.24 and 5.25). This implies that for an ensemble of oscillators, the adiabiatic case will preserve the original distribution of amplitudes and phases, with just a constant re-scaling of the amplitude and a constant phase shift,

whereas in the impulsive case, the original distributions are modified (see Eq. 5.26).

If we extend the implications of this toy model to particle orbits in a spherical potential, we expect different orbital responses to the reduction of the central potential in either the adiabatic or the impulsive case. In both cases we expect a net expansion of the initial orbits of the tracer particles. In the adiabatic case, however, we expect the expansions to be fully determined by the integrals of motion associated with the trajectories of the tracers. In the impulsive case, on the other hand, we expect the phase of the radial oscillation to be of key importance, leading to a different response to the change in the potential depending on whether particles are, for instance, moving toward the halo centre or away from it. We expect this behaviour as a consequence of the time averages theorem for stellar orbits (see for instance Binney and Tremaine, 2008), which states that averaging physical quantities associated with the orbit of a star over a time longer than the orbital period is equivalent to calculating an ensemble average over the phase space volume which is accessible to the star. This theorem can only be applied if the accessible phase space volume does not change in a time scale of the order of an orbital period. Under this condition, the time averages theorem implies that the actions of the stars are conserved quantities, so-called adiabatic invariants. However, if the change in the potential is impulsive, the time averages theorem does not hold, and stars originally in a similar orbital family will split during the explosive event and end up on widely different orbits.

This fundamentally different nature of the adiabatic and the impulsive mechanisms of core formation offers a promising possibility to distinguish the SIDM and supernova-driven scenarios. The distribution of stellar orbits in a cored halo might contain signatures of the type of process that created the core. Our objective in this work is to model these core formation scenarios using simulations of isolated haloes, and study in detail the consequence of the adiabatic (SIDM) and impulsive (supernova-driven) cases for the orbits of collision-less tracers (stars).

5.3 Initial conditions and core formation modelling

We perform idealised (non-cosmological) simulations of isolated DM haloes using the AREPO code Springel (2010) with an added algorithm that implements elastic DM self-scattering (as described in detail in Vogelsberger et al., 2012). For all our main results, the simulated halo has a total mass of 1.48×10^{10} $M_\odot$ (see Section 5.3.4 for more details). This value is chosen as it is a characteristic mass for the haloes of the progenitors of the most massive satellites of the Milky Way. The halo has a virial radius[1] $r_{200} = 52\,\mathrm{kpc}$ and a concentration $c_{200} = r_{200}/r_s = 15$ where r_s is the scale radius (see Eq. 5.2 below). The concentration value we choose is within the 1σ range of the concentration-mass relation for haloes of this mass in a Planck cosmology (e.g. Pilipenko et al., 2017). The Plummer equivalent softening length for the DM particles is chosen following Power et al. (2003) (see also Appendix 5.9):

$$\varepsilon = 4\frac{r_{200}}{\sqrt{N_{200}}} \tag{5.1}$$

[1]Defined as the radius where the mean enclosed DM density is 200 times the critical density.

where N_{200} is the number of enclosed DM particles within r_{200}. After testing for stability and convergence (see appendix 5.9), and given the fact that we want to observe effects related to the formation of a $\approx$ 1kpc-sized core, we have chosen to run all of the simulations with $N = 10^7$ equal-mass DM particles. The gravitational softening according to Eq. 5.1 is then 60 pc.

5.3.1 Initial conditions: equilibrium configuration

All our haloes are spherically symmetric isotropic Hernquist haloes Hernquist (1990) with an exponential cutoff beyond the virial radius that are initially set up in dynamical equilibrium following the algorithm described in Kazantzidis et al. (2004). Specifically, we generate a distribution of simulation particles with the following density profile:

$$\rho(r) = \begin{cases} \frac{\rho_s r_s^4}{r(r+r_s)^3} & r \leq r_{200} \\ \frac{\rho_s}{c_{200}(1+c_{200})^3} \left(\frac{r}{r_{200}}\right)^{\varepsilon_d} \exp\left(-\frac{r-r_{200}}{r_{\rm decay}}\right) & r > r_{200}. \end{cases} \tag{5.2}$$

Here, ρ_s is the scale density and $r_{\rm decay}$ is a parameter that determines how sharp the exponential cutoff is. The choice of the latter parameter is somewhat arbitrary, although a sensible choice is for it to be smaller than r_{200} in order to guarantee a reasonable convergence of the equilibrium configuration, and to avoid having large additional mass outside r_{200}. Choosing the exponent ε_d to be:

$$\varepsilon_d = \frac{-1-4c_{200}}{1+c_{200}} - \frac{r_{200}}{r_{\rm decay}} \tag{5.3}$$

renders a continuous density profile with an equally continuous logarithmic slope.

To achieve (isotropic) dynamical equilibrium we first calculate the phase space distribution using Eddington's formula Eddington (1916):

$$f(\mathscr{E}) = \frac{1}{\sqrt{8}\pi^2} \int_0^{\sqrt{\mathscr{E}}} du \frac{d^2\rho}{d\Psi^2}(r(\Psi(u))), \tag{5.4}$$

where $u = \sqrt{\mathscr{E} - \Psi}$ and $\mathscr{E}$ and $\Psi(r)$ are the (negative) specific energy and gravitational potential, respectively, both shifted by the zero point of the gravitational potential (set at $r \to \infty$). We then assign velocity vectors to all the DM particles in the halo by sampling their magnitude from the phase space distribution calculated above using a rejection sampling algorithm and then randomly drawing directions from the unit sphere. See Appendix 5.9 for our tests on the stability of the equilibrium configuration.

5.3.2 Modelling supernova-induced core formation

In order to transform an initially cuspy density profile into a cored one, a certain amount of energy needs to be deposited into the centre of the halo. An estimate of the required energy is obtained by comparing the potential energy of the halo before and after the cusp-core transformation (see Peñarrubia et al. 2012):

$$\Delta E = \frac{W_{\rm core} - W_{\rm cusp}}{2}. \tag{5.5}$$

The potential energy of the halo is

$$W = -4\pi G \int_0^\infty \rho(r) M(<r)\, r\, dr, \tag{5.6}$$

where $M(<r)$ is the enclosed mass. Neglecting the exponential cutoff of the profile in Eq. 5.2, we can take the potential energy of the Hernquist halo as an approximation for $W_{\rm cusp}$:

$$W_{\rm cusp} \sim W_{\rm HQ} = -\frac{GM_\infty^2}{6r_s}. \tag{5.7}$$

where M_∞ is the total mass of the halo $M_\infty = (16/15)^2 \times M_{200}$. In order to have an estimate of $W_{\rm core}$ we can model the final cored halo with the following density profile:

$$\rho_{\rm cored} = \frac{\rho_s r_s^4}{(r+r_c)(r+r_s)^3}, \tag{5.8}$$

where r_c is the core radius. With these considerations we estimate that the required energy to form a $r_c = 1$ kpc core in a $1.48 \times 10^{10} \mathrm{M}_\odot$ halo with a concentration $c_{200} = 15$ is roughly 11% of the halo's initial potential energy. We take this energy estimate into consideration in the modelling of supernova gas blowouts below to form a core of the relevant size.

We follow the procedure described in Garrison-Kimmel et al. (2013) to implement a simple model of an episode of gas accretion followed by a sudden gas blowout. This is done by adding an external time-dependent Hernquist potential centered at the halo centre, which represents an effective baryonic mass distribution $M_{\rm gal}(t)$ with scale radius $r_{\rm s,gal}$, and induces an additional acceleration of the particles in the simulation:

$$\mathbf{a}_{\rm ext} = -\frac{GM_{\rm gal}(t)}{(r+r_{\rm s,gal})^2}\frac{\mathbf{r}}{r}, \tag{5.9}$$

where $\mathbf{r}$ is the position vector relative to the centre of mass of the halo. Contrary to Garrison-Kimmel et al. (2013), we do not introduce an extra gravitational softening since the acceleration does not diverge at $r \to 0$, but rather reaches a constant value. We explore two simple models for the behaviour of $M_{\rm gal}(t)$:

- *single explosion:* starting from a halo in equilibrium (set up as described in Section 5.3.1), we add a constant external potential $M_{\rm gal}(t) = M_{\rm single}$ at $t = 0$ and let the halo contract adiabatically into the centre until it reaches a new equilibrium. Once the configuration is stable, we set the mass to zero instantaneously, mimicking a single large mass blowout.

- *multiple explosions:* starting from a halo in equilibrium (set up as described in Section 5.3.1) at $t = 0$, we grow the mass distribution linearly with time until a maximum value $M_{\rm mult}$ is reached at $t = \tau$ ($M_{\rm gal}(t) = M_{\rm mult} \times t/\tau$), afterwards we set the mass to zero instantaneously. This procedure is repeated periodically, which mimics several consecutive mass blowouts.

Each time the mass is set to zero, a total amount of energy:

$$E = \frac{GM_{\rm gal}^2}{6r_{\rm s,gal}} \tag{5.10}$$

is deposited into the halo, which implies that if we keep $r_{\rm s,gal}$ fixed, the transferred energy scales as the square of the baryonic mass distribution. This means that in order to achieve the same effect with N smaller explosions as with a single explosion we require

$$M_{\rm mult} = \frac{M_{\rm single}}{\sqrt{N}}. \tag{5.11}$$

5.3.3 Core formation in SIDM

DM haloes naturally develop cores with a size that is comparable to the scale radius of the halo if the cross section for self-scattering is sufficiently large to thermalize the inner halo. In particular, collisions redistribute energy among the DM particles from the dynamically hotter intermediate regions near the scale radius (where the velocity dispersion profile approximately peaks) to the central colder regions (see e.g. Fig. 2 of Colín et al., 2002). The time scale for full thermalisation of the halo is approximately given by the condition of having on average one collision per particle in the lifetime of the halo centre ($t_{\rm age}$):

$$\rho_{\rm p}(\sigma_T/m_\chi)v_{\rm p}t_{\rm age} \approx 1 \tag{5.12}$$

where ρ_p and $v_{\rm p}$ are the density and average relative velocities between DM particles, respectively, at the peak of the velocity dispersion profile. If the average number of collisions is less than one, then the cusp core transformation is incomplete. In fact, for dwarf-sized SIDM haloes, if $\sigma_T/m_\chi \sim 0.1$ cm^{2}g^{-1}, then the resulting density profile is only slightly less cuspy than the corresponding CDM halo Zavala et al. (2013). We can consider this as a lower limit of the cross section for SIDM to be a distinct alternative to CDM. On the other hand, if the average number of collisions is much larger than one, then the unavoidable gravothermal collapse is triggered and the SIDM halo eventually becomes even cuspier than its CDM counterpart (e.g. Koda and Shapiro, 2011; Pollack et al., 2015). In the case of dwarf-sized haloes, the threshold for collapse is reached when σ_T/m_χ is of $\mathcal{O}(10)$ cm^{2}g^{-1} (Elbert et al. 2015 and Zavala et al. in prep.). Thus, for dwarf-sized haloes, a cross section of $\mathcal{O}(1)$ cm^{2}g^{-1} provides the right amplitude for fully thermalized SIDM haloes. In our SIDM simulations we pick a value of 2 cm^{2}g^{-1}.

As mentioned at the beginning of Section 5.3, we model the effect of DM self-interactions by using the probabilistic algorithm for elastic self-scattering in AREPO described in detail in Vogelsberger et al. (2012).

5.3.4 Cored haloes in SIDM and CDM

Fig. 5.1 shows a comparison of the density profile (top panel) and velocity dispersion profile (bottom panel) of a halo that develops a core impulsively (SNF) and adiabatically (SIDM; dotted red line) starting from the same initial condition. For the former case we show the single explosion case (dot-dashed black line) and a case with 10 explosions

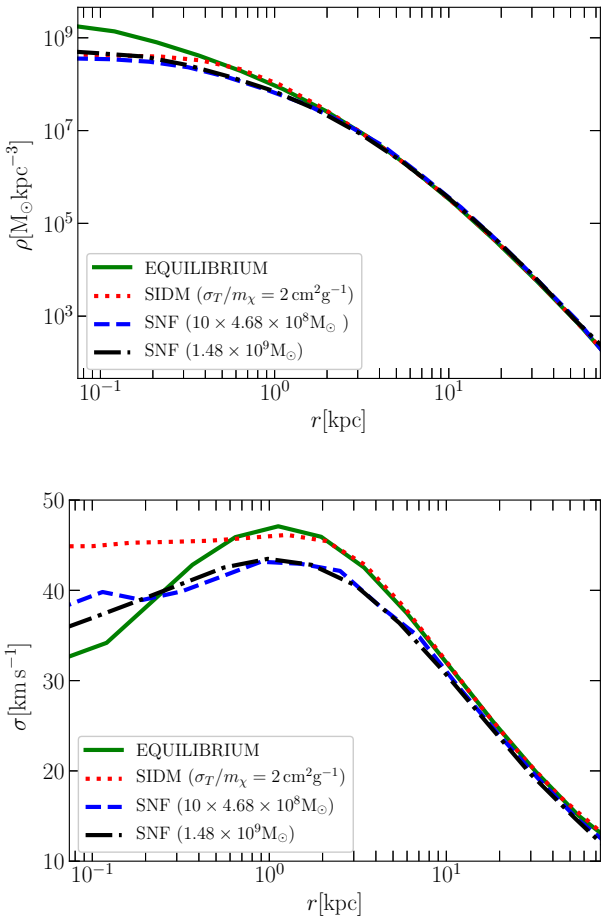

Figure 5.1. DM core formation in the impulsive (SNF) and adiabatic (SIDM) scenarios: density (top) and velocity dispersion profiles (bottom) for a $1.48 \times 10^{10}\ M_\odot$ *halo. The initial condition is that of a Hernquist halo with a cutoff at large radii (see Eq. 5.2) set in dynamical equilibrium. Different lines belong to the different simulations: the green solid line denotes the equilibrium case in which we evolve the halo profile set up as in section 5.3.1 including only self-gravity between the DM particles; the dashed black line refers to the SNF model with a single large explosion removing a mass of* $1.48 \times 10^9\,\mathrm{M}_\odot$*, the dashed blue line refers to the SNF model with ten consecutive explosions depositing a total amount of energy equivalent to the single explosion case; the red dotted line refers to the case where DM is self-interacting with* $\sigma_T/m_\chi = 2\ \mathrm{cm}^2\mathrm{g}^{-1}$*. Except for the equilibrium case, the profiles are shown in their stable configuration after the cusp-core transformation has finalized.*

with an injected energy equivalent to the single explosion case (dashed blue line); see Section 5.3.2. We also show the case where there is no energy deposition or redistribution (solid green line). This simulation serves as a reference for the degree to which the equilibrium method described in Section 5.3.1 is stable across the simulation time. It is also a baseline for comparison with the cuspy Hernquist profile, which is the initial condition for all the simulations. Except for this reference case, all the other profiles are shown after the halo has reached a new equilibrium configuration after the cusp-core transformation. We remark that tn the SIDM case, the equilibrium is only quasi-stable. Thermalization is fully complete after 1.2 Gyrs have passed from the start of the simulation. Afterwards, the central density increases due to the beginning of the gravothermal collapse phase, which is a stage that we want to avoid since our goal is to compare DM cores formed due to self-interactions with cores formed due to supernovae.

For the cases of impulsive core formation (SNF), the scale radius of the external baryonic distribution is fixed to 300 pc (a typical half-light radius among bright Milky-Way satellites), while its mass is chosen to be larger than the mass of typical dwarf galaxies in order to create a core with an approximate size of 1 kpc (comparable to the SIDM case explored) according to the energy requirement outlined in Section 5.3.2. In the single explosion explosion case, this corresponds to $M_{\rm gal} = 1.48 \times 10^9 {\rm M}_\odot$. As mentioned in Section 5.3.2, this baryonic mass is added as an external Hernquist potential, and thus, the halo responds to it by adiabatically contracting substantially given the large baryonic mass. Once the halo is fully relaxed, the explosion is triggered and we follow the evolution of the halo until it reaches its final equilibrium state, which is the one shown in Fig. 5.1. In the multiple explosion case, the baryonic mass distribution is added periodically as a linear function of time with a maximum at $M_{\rm mult} = \sqrt{10} \times 1.48 \times 10^8 {\rm M}_\odot$ (see Eq. 5.11) and a period of 0.3 Gyrs followed by an explosion. A total of 10 explosions are simulated in this way to achieve the same energy deposition as in the single explosion case. After the last explosion, we let the halo evolve into its final equilibrium configuration. Although this episodic model of accretion and mass blowout also results in adiabatic contraction of the halo during the accretion phase, it is naturally milder than the single explosion case.

Fig. 5.1 shows that both of the impulsive (SNF) cases result in a cored halo that looks very much alike, with both of them having very similar density and velocity dispersion profiles. This shows that the energy requirement is fulfilled regardless of whether the explosion occurs only once, or is episodic as long as the energy transferred during the explosion(s) is equivalent. This result is a validation of the scaling of the deposited energy on the square of the baryonic mass mentioned in Section 5.3.2 (based on Peñarrubia et al., 2012), which supports our modelling of explosions as an accurate implementation. We note that this scaling with mass is apparently in disagreement with the results in Fig. 4 of Garrison-Kimmel et al. (2013). There, the authors compare the single and multiple explosion cases; and although they mention that the single explosion case is more efficient than the multiple explosion case (for the same amount of blowout mass), it is clear from their Fig. 4 that the difference is not very strong. We have found, however, that if the total blowout mass is set to the same value as in the single explosion case ($10 \times 1.48 \times 10^8$ M$_\odot$), then the size of the formed core reduces significantly. Given the energetic arguments in Section 5.3.2, we are confident that our implementation is

accurate.

Looking at the adiabatic (SIDM) case in Fig. 5.1, we can see that the adiabatic core is fully isothermal within the radius where the velocity dispersion peaks (~ 1 kpc). As we mentioned in Section 5.3.3 this is a signature of full thermalization of the core. We note that the isothermal region extends further outwards from the halo centre than the constant density plateau, which agrees with results found previously (e.g. Koda and Shapiro, 2011; Vogelsberger et al., 2012). Comparing the SIDM and SNF cases, it is apparent that although we have achieved our goal of creating a DM core with the same asymptotic central density and similar core size, there are clear differences in the density and velocity dispersion profiles. The density core is flatter in SIDM and the impulsive case is not fully isothermal. Instead, it retains a similar slope - albeit shallower - as in the initial cuspy halo. Within the central region ($\lesssim 15$ kpc), both the average density and velocity dispersion of DM particles are larger in the SIDM case than in the SNF case. This is a signature of the difference between the central energy redistribution taking place in SIDM, which bounds the expansion of orbits within the central regions, versus the energy transfer in the SNF case, which causes the orbits to expand to much larger radii. We study this difference in detail below where we explore the impact of these two mechanisms of core formation on the orbits of tracer particles.

5.4 Simulations with tracer particles

In order to study the effect of a time-dependent potential on the orbits of stars in detail, we need to introduce (collision-less) tracer particles. The reason for this is twofold: (i) since we need to track the orbits accurately in time, it will be computationally prohibitive (and unnecessary) to use all the $N = 10^7$ DM particles as tracers; (ii) for the SIDM case, the DM particles do not act as tracers since they are collisional. We require that the population of tracer particles fulfill the following requirements:

1. The combined mass of all the tracer particles should be smaller than the mass of one DM particle. In this way, they have a negligible impact on the motion of DM particles. We choose $1.48 \times 10^{-1} \mathrm{M}_{\odot}$ as the mass of a single tracer.

2. The orbits of the particles are computed exactly from the sum of the gravitational forces exerted by all DM particles[2].

3. The softening length of the tracer particles is small (we pick $\varepsilon_{\mathrm{tracer}} = 15$ pc), and we neglect self-gravity to ensure detailed time integration.

An exact calculation of the forces exerted on the tracer particles is computationally expensive and thus we cannot afford to include a very large number of tracers, still, the sample needs to be large enough to obtain statistically significant results. We have found that a sample of 2000 particles is sufficient for our purposes.

Following the arguments in Section 5.2 we expect that the actions of tracer particles will be conserved when the change in the potential is adiabatic (SIDM case), whereas

[2]This is a direct summation approach that requires N partial forces per tracer particle, while for DM particles, gravitational forces are computed using a tree algorithm that requires $\mathcal{O}(\log N)$ calculations per DM particle.

we do not expect that in the impulsive (SNF) case. In particular, the radial action of bound orbits will be an adiabatic invariant:

$$J_r = \frac{1}{\pi}\int_{r_{\rm peri}}^{r_{\rm apo}} dr\sqrt{2E - 2\Phi(r) - \frac{L^2}{r^2}} \tag{5.13}$$

where E is the energy of the tracer particle, L the magnitude of its angular momentum, Φ is the gravitational potential, and $r_{\rm apo}$ and $r_{\rm peri}$ are the apocentre and pericentre of the orbit of the tracer.

Our strategy to investigate the differences between an impulsive and an adiabatic change in potential is to put the tracer particles in configurations that probe quantities related to the radial action. In the rest of this Section we describe the setup for three different configurations we have chosen to accomplish this objective.

5.4.1 A single orbital family

Following Meiron and Kocsis (2018), we expect different orbital families to respond differently to a change in the gravitational potential, as well as to effects of diffusion and mixing. For this reason, we decided to investigate how one particular orbital family responds to the impulsive and adiabatic core formation processes. An orbital family is defined by having similar values of E and L. Given the spherical symmetry of our problem, these two requirements are sufficient and it is not necessary that the orientation of the angular momentum vector be the same for orbits to be of the same family. Since we want to compare the adiabatic and impulsive cases of core formation, we need to pick an orbital family with (E,L) values which corresponds to particles in orbits around the region where the core will form. In the case of spherically symmetric problems there is a natural way to make a sensible choice by connecting (E,L) to the apocentre and pericentre of the orbit. In the stable case we can write the energy for a central potential as:

$$E = \frac{1}{2}\left(\frac{dr}{dt}\right)^2 + \frac{1}{2}\frac{L^2}{r^2} + \Phi(r) \tag{5.14}$$

In a bound orbit, $dr/dt = 0$ at the apocentre and pericentre of the orbit. Therefore,

$$0 = 2\left[E - \Phi(r_{\rm apo(peri)})\right] - \frac{L^2}{r_{\rm apo(peri)}^2} \tag{5.15}$$

which can be solved to yield E and L in terms of the apocentre and pericentre radii:

$$L = r_{\rm apo} r_{\rm peri}\sqrt{\frac{2\left[\Phi(r_{\rm apo}) - \Phi(r_{\rm peri})\right]}{r_{\rm apo}^2 - r_{\rm peri}^2}} \tag{5.16}$$

$$E = \Phi(r_{\rm apo}) + \frac{1}{2}\frac{L^2}{r_{\rm apo}^2} \tag{5.17}$$

In this way, we can chose the orbital family we wish to inspect by simply choosing an interesting radial range, spanning from radii that are heavily affected by the change in

potential to radii that are barely affected. Specifically, we draw pericentre and apocentre radii for all the 2000 tracers from Gaussian distributions with a spread of 0.075 kpc centered at 0.75 kpc and 3 kpc, respectively. We then calculate the associated angular momentum and energy pairs and subsequently populate the radial range with orbits at different radial positions (phases), sampling from a flat distribution between pericentre and apocentre. The direction of rotation is randomly chosen.

5.4.2 Gaussian distribution in radial action

To directly investigate the degree to which radial actions (Eq. 5.13) are (not) conserved in the (impulsive) adiabatic case, we construct a Gaussian distribution of radial actions $p(J_r)$ (with a mean of 45 $\mathrm{km\,s^{-1}kpc}$ and a standard deviation of 3.75 $\mathrm{km\,s^{-1}kpc}$) and assign J_r values to each of the 2000 tracers using a rejection sampling algorithm von Neumann (1951) as follows:

1. draw a random radius r^* between two reference radii (0.15 kpc and 3 kpc).
2. draw a random energy E^* value from a flat prior between the potential r^* and $E^* = 0$, the energy at which the tracer is no longer gravitationally bound.
3. draw a random L^* value from a flat prior bounded by the maximum possible value given by the radius and energy drawn in (i) and (ii).
4. calculate the value of the radial action J_r^* for the set (r^*, E^*, L^*) using Eq. 5.13.
5. accept J_r^* if $p(J_r^*)$ is larger than a random variable drawn in the interval $[0, \max(p(J_r))]$, reject otherwise

In case J_r^* is accepted, we assign a velocity vector to the tracers which is in agreement with the pair (E^*, L^*).

5.4.3 A Plummer sphere

Finally, in order to model a stellar population with a more realistic distribution, we also sample the positions of the 2000 tracers from a radial Plummer profile with a half-light radius of 500pc, which is a typical value for the bright Milky Way satellites. This is the only free parameter of the Plummer distribution since by construction the stellar population is a set of tracers that does not contributes to the gravitational potential. We set the total mass of the Plummer profile to be equal to $N_\star m_\star$, where $N_\star = 2000$ and $m_\star = 1.48 \times 10^{-11}\mathrm{M}_\odot$. The total mass needs to be consistent in order to be able to sample the tracers' positions from a Plummer density profile with the same method we used to sample the DM particles' positions from a Hernquist profile. Since in this setting the gravitational potential is fully determined by the DM distribution, we use the spherical Jeans equation (assuming isotropic orbits for the tracers) to compute the equilibrium velocity dispersion profile of the Plummer sphere. The velocities of the tracers are then sampled from a Maxwellian distribution with a radially dependent velocity dispersion.

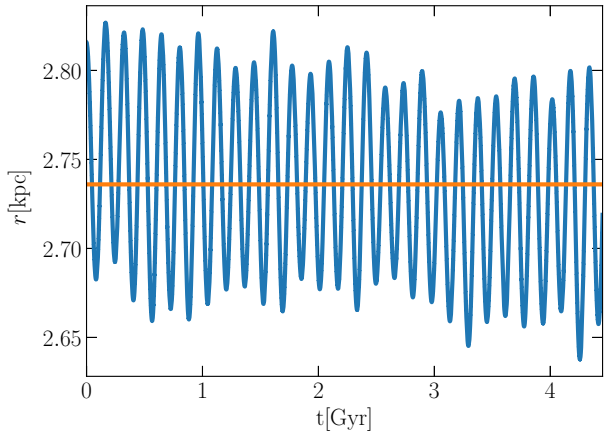

Figure 5.2. Numerical errors from circular orbits. The radial distance of one particular tracer particle in a Hernquist halo relative to the centre of mass of the halo as a function of time (the red line is the time average). Despite the tracer initially being on a circular orbit, the radial distance oscillates in time, which also implies a non-zero radial velocity. We use these oscillations to quantify the baseline errors (in the simulation where the halo is in equilibrium) in the time integration scheme we use.

5.5 Results

5.5.1 Numerical errors in circular orbits

Before looking at the evolution of the tracer particles as they respond to the two mechanisms of cusp-core transformation, we will first study the amplitude of the time integration errors, force errors, and discreteness effects which are inherent to the code and setup we are using to track the orbits of tracers. Looking at simulated circular orbits in the case where the potential is fixed (equilibrium) is a simple way to estimate the magnitude of these numerical errors. In particular, we look at deviations from a true circular orbit where the radius is constant and the radial velocity is zero. This will serve as a benchmark when studying the orbital response in a time-varying potential. To construct this benchmark, we set up the 2000 tracers within the halo in equilibrium on circular orbits whose radii are sampled from a flat distribution between 0.75 and 3 kpc, and evolve the simulation over 4.5 Gyrs.

Fig. 5.2 shows one typical circular orbit, albeit at a rather large radius within the range we probe. We can see that the orbit is not completely circular, with the orbital radius oscillating in time by $\sim \pm 2\%$ around an average value. We look at the statistical behaviour of the ensemble of circular orbits to estimate the baseline integration errors. To give a measure to the errors in the radius of a single orbit, we calculate the mean and

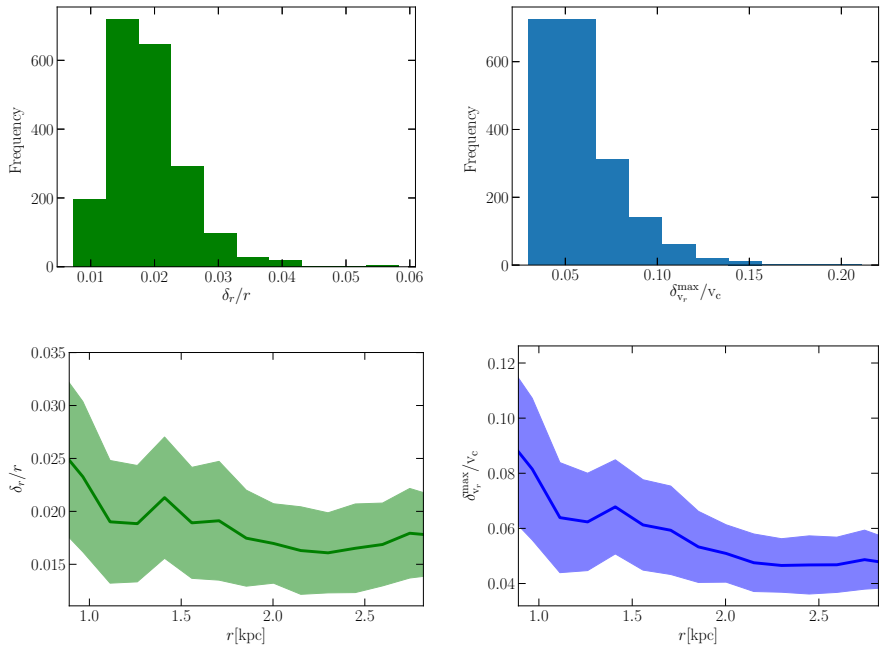

Figure 5.3. Baseline integration errors for circular orbits. On the left panels, histogram and radial profile of the relative errors in the orbital radius, and on the right panels, histogram and radial profile of the maximum errors in the radial velocity divided by the circular velocity at the radius where the maximum error occurs. The relative error in the orbital radius is centered around 2 per cent, whereas the absolute error in radial velocity is centered around 4 $\mathrm{km\,s^{-1}}$*, which translates to a relative error of around 7 per cent. The bottom panels show a clear dependence of the orbital errors on the orbital radius, which can be mainly attributed to the number of enclosed particles being smaller at smaller radii and hence the sampling of the gravitational potential being poorer towards the halo centre. The Poissonian sample variance is included in the lower panels and it exhibits no dependence on radius.*

variance as:

$$\langle r \rangle = \frac{1}{N_{\mathrm{steps}}} \sum_{i=0}^{N_{\mathrm{steps}}} r_i \tag{5.18}$$

$$\delta_r^2 = \frac{1}{N_{\mathrm{steps}} - 1} \sum_{i=0}^{N_{\mathrm{steps}}} (r_i - \langle r \rangle)^2 , \tag{5.19}$$

where r_i is the orbital radius at each timestep in the simulation and N_{steps} is the total number of timesteps. For the velocity, we provide a worst-case estimate of the errors by recording the maximal radial velocity encountered throughout the whole simulation time for each tracer.

Fig. 5.3 shows the orbital errors of all the 2000 tracer particles. The upper panels show the histograms of the relative error in radius on the left ($\delta_r/\langle r\rangle$) and the worst-case error in the radial velocity divided by the circular velocity at the radius where the maximum error occurs ($\delta_{\mathrm{v_r}}/\mathrm{v_c}$); the latter has been normalised to the circular velocity, which, representing the characteristic magnitude of the tracers' velocities at a given radius, is the relevant velocity to compare the errors to. We see that on average, the relative error in the radius is 2 per cent, with a distribution tail extending to 4 per cent. On the other hand, the mean maximum value of $\delta_{\mathrm{v_r}}/\mathrm{v_c}$ is around 7 per cent with a distribution tail extending to 15 per cent. We emphasize that in the case of velocities, these are worst-case errors over the entire simulation time.

The lower panels of Fig. 5.3 show the dependence of the orbital errors on the (time averaged) radius of the orbit. The sample of tracers is split in radial bins according to the individual (time averaged) radius with the solid line in the left (right) panel showing the mean value of $\delta_r/\langle r\rangle$ ($\delta_{\mathrm{v_r}}^{\mathrm{max}}/\mathrm{v_c}$) over the ensemble of tracers, while the shaded area shows the Poisson sample variance. We find that the orbital errors are larger at smaller radii, where the impact of particle discreteness is larger. However, over the radial range shown in the plots, the mean of the orbital errors changes by less than a factor of 2. Thus, we conclude that the central value of the upper left histogram in fig. 5.3 suffices as the leading order estimate for the relative baseline error in the orbital radius. In the case of errors in orbital velocities, we can see that relative to circular velocities, the median of the errors is around 5-7 per cent. We note that in absolute terms the distribution of errors is centered at around 4 $\mathrm{km\,s^{-1}}$. The approximate magnitudes of these errors will help us to quantify the significance of any detected discrepancy between the behaviour of tracer particles when undergoing adiabatic or impulsive core formation.

5.5.2 A single orbital family

To study the evolution of the orbital family defined in Section 5.4.1, we will analyse the radial 2-dimensional phase space density of orbits in the $r - v_r$ plane. Although this is only a projection of the six-dimensional phase space density, it contains the relevant information to describe the orbits of the tracers since the equation of motion in a spherically symmetric potential is two-dimensional. Fig. 5.4 shows the initial (top) and final (bottom) radial 2-dimensional phase space density of the 2000 tracers embedded in the halo in equilibrium after 4.2 Gyrs. As we can see in the top panel, the phase space of the orbital family is not sampled completely even initially, particularly in the vicinity of the apocentre and the pericentre of the orbits where the sampling is sparse. This is due to a non-perfect random number generator for the uniform distribution, which undersamples the extrema. Nevertheless, the distribution of points essentially fills a closed area in phase space, which was our intention in Section 5.4.1 when introducing the orbital family configuration. After 4.2 Gyrs of evolution within a halo that essentially remains in equilibrium (see Appendix 5.9), we find that the orbital family remains within essentially the same phase space area. There is some minor evolution, however, most noticeably the phase space area is more evenly distributed and has now spread slightly compared to the initial configuration. This is a consequence of the baseline errors in the orbits described in Section 5.5.1. Another difference is that there are noticeably more stars near the apocentre than near the pericentre, which is due to the tracers spending

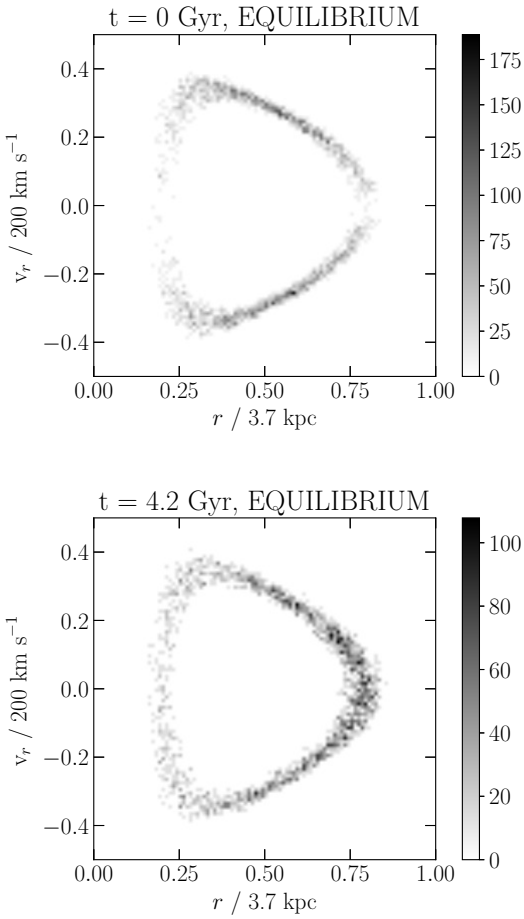

Figure 5.4. Initial and final configuration (after 4.2 Gyrs) of the radial 2D phase space density (sampled with 2000 tracers) of the orbital family described in 5.4.1 in a purely gravitational simulation of a dwarf-size halo set in equilibrium. The radial coordinate is in units of 3.7 kpc, whereas the radial velocity is in units of 200 km s^{-1}*. The colour bar refers to the 2D phase space density in units of* $[\mathrm{s\,km^{-1}kpc^{-1}}]$*. In the long term, the tracer particles are preferentially located near the apocentre of their orbits, which is in general agreement with expectations from collision-less kinematics. Furthermore, the spread of the distribution near the apocentre increases, which can be attributed to the numerical integration errors discussed in section 5.5.1 (see fig. 5.3). Despite this, the overall evolution in phase space is minimal, which allows us to use the final distribution as a reasonable baseline of comparison with the simulations in the adiabatic and impulsive core formation scenarios in Figs. 5.5 and 5.6.*

a longer period of time in the vicinity of the apocentre where the radial velocities are lower (see also Peñarrubia, 2013). We take this baseline numerical diffusion in phase space as the benchmark for the extent of the irreducible artificial evolution that occurs in our simulations, and move on to discussing the evolution of the orbital family under the core formation scenarios.

Fig. 5.5 shows the 2-dimensional phase space distributions of the adiabatic (SIDM, top panel) and impulsive (SNF with 10 explosions, bottom panel) cases after 4.2 Gyrs. The final configuration in the SIDM case is remarkably similar to the equilibrium case shown in the bottom panel of Fig. 5.4, both in the area occupied by the orbital family and in the phase space density values across this area (notice that the grey scale is the same in both). The main difference is found in the mean apocentre of the orbits, which has contracted sligthly by $\sim 10\%$. This is a result of the conservation of radial actions as the central potential is reduced adiabatically. This result is a nice illustration of the time averages theorem (see Section 5.2) being valid for the SIDM case, as we can directly see in the top panel of Fig. 5.5 that tracers (stars) that are part of the same orbital family end up being part of the same orbital family after the adiabatic cusp-core transformation.

Looking at the impulsive case of core formation, shown in the bottom panel of Fig. 5.5, we find an entirely different picture. The tracers are no longer in a common orbital family, but instead they have separated into different families occupying a wider range of radial velocities and orbital radii. Some of the tracers are in bound orbital families with apocentres near $5-6$ kpc, twice as large as in the original orbital family. Other particles have moved to even larger radii out to almost 45 kpc, near the virial radius of the halo. This clear separation of the orbits is a result of the time averages theorem not being valid in the impulsive case, and can be explained from the results of the 1-dimensional harmonic oscillator toy model (see Section 5.2 and Appendix 5.8). From this model, we expected that the expansion of an orbit (amplitude of an oscillator) after an explosion is determined by the magnitude of the change in potential (energy deposited) as well as the particular location of the tracer along its orbit (phase of the oscillator). It is due to this dependence on the phase of the orbit that the tracers group into different orbital families as the explosions occur. We notice that a substantial fraction of the tracers is found in new orbits with a smaller apocentre than the initial orbital family (shown in the top panel of Fig. 5.4). This is likely due to the periodic cycles of growth of the baryonic mass distribution in the multiple explosion case (see Section 5.3.2), which cause an adiabatic contraction of the collision-less tracers and DM particles.

To look at the impact of the impulsive case on the orbital family in a more clear way, we show in Fig. 5.6 the single explosion model of core formation. The 2-dimensional phase space density is shown 1.05 Gyrs after the explosion. Here the segregation of orbits into distinct orbital families (shells in phase space) is more clear than in the bottom panel of Fig. 5.5, where the combined effect of multiple explosions has diluted the effect. We note, however, that as time progresses, more shells appear in this single explosion case as well, reducing the contrast between the different orbital families and thus diluting the overall effect. We have chosen the time shown in Fig. 5.6 to highlight the appearance of the shell-like structure of orbits. Finally, we also identify a clear infall signature towards the halo centre, starting beyond 10 kpc, which is also visible, albeit barely, in the multiple explosion case. The tracers in this orbit are those

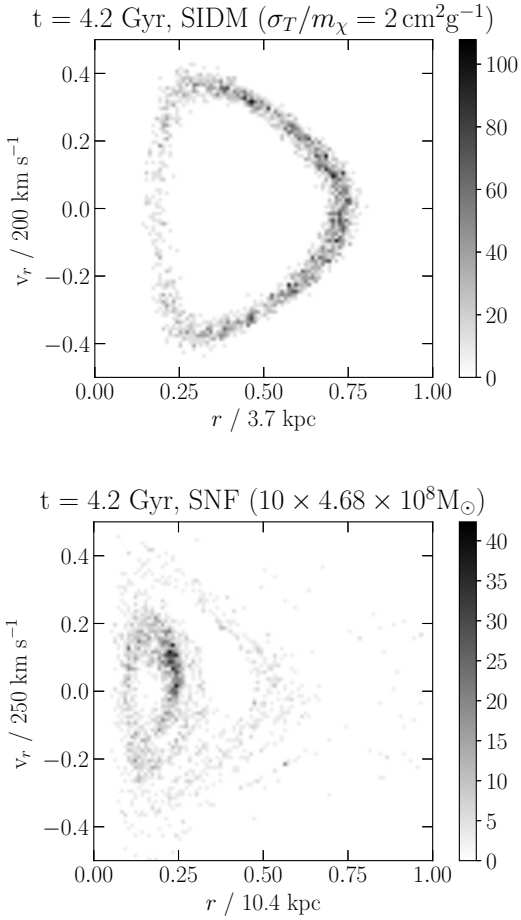

Figure 5.5. The same as in fig. 5.4 but showing the final configuration (at 4.2 Gyrs) of the adiabatic (SIDM) and impulsive (SNF with 10 explosions) simulations on the top and bottom, respectively. Notice that in the SNF case (bottom) the axes have been scaled differently to the SIDM (top) and equilibrium (Fig. 5.4) cases. This is because in this case the tracer orbits expand to considerably larger radii in response to the impulsive mass removal, with some even expanding to radii beyond the range shown in the figure, up to 4 times as large. We choose to plot this radial range in order to focus on the most densely populated region of the phase space. Although both mechanisms of core formation differ from the equilibrium case depicted in Fig. 5.4, the evolution in the SIDM (adiabatic) case is not very far from it, while the SNF (impulsive) case shows striking differences.

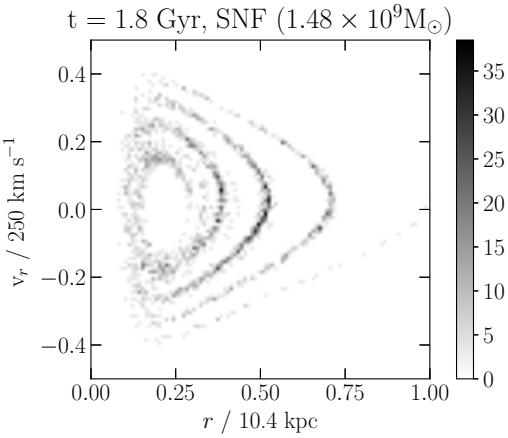

Figure 5.6. The same as the bottom panel of Fig.5.5 but for the simulation with one single explosion. The time of the simulation is 1 Gyr after the explosion occurred (1.8 Gyrs from the simulation start). We can see that the initial orbital family (top panel of Fig. 5.4) has been split into at least 3 distinct shells (orbital families), and a more diffuse area closer to the initially occupied phase space area. Furthermore, we can identify a clear infall orbit starting at the largest radii occupied by tracers.

that were in an initial orbital phase that maximised the orbital expansion when the potential suddenly changed. They were thus expelled to very large radii and are now falling back into the halo. This resonant amplification of the energy transferred to the orbit is a distinct feature of the impulsive case. Along with the segregation of the original orbital family, it causes the final configuration of the SNF (single explosion) case to differ significantly from the adiabatic case. These effects are, however, not as striking in the more realistic episodic explosion core formation scenario. Nevertheless, the preservation of the original orbital family in the adiabatic (SIDM) is very clear in general, and exhibits a dramatic contrast to the destructive nature of the impulsive (SNF) case.

5.5.3 Gaussian distribution in radial action

To analyse the degree to which the conservation of radial action is satisfied in the adiabatic (SIDM) case and violated in the impulsive (SNF) case, we examine the evolution of the tracers originally sampled from a Gaussian distribution in radial action (see Section 5.4.2). Fig. 5.7 shows the final distributions of radial actions for the tracers in the benchmark equilibrium case (green), SIDM (red) and SNF (10 explosions, blue) cases.

To quantify the differences between the distributions, we calculate the mean and standard deviations of each distribution. We then interpret the final distribution as a Gaussian with parameters μ and σ given by the calculated mean and standard deviation.

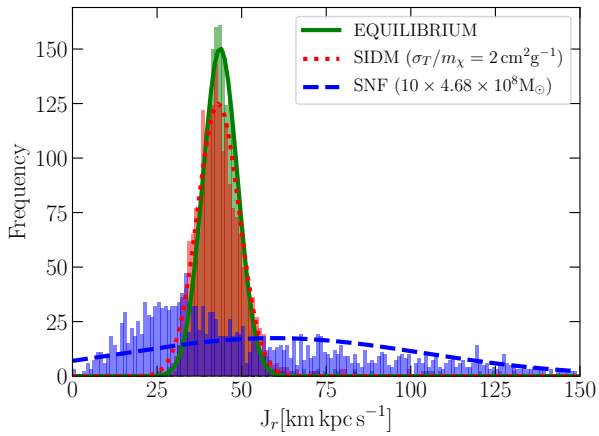

Figure 5.7. Distributions of radial actions after 4.2 Gyrs for the 2000 tracers originally set up following a Gaussian distribution in J_r (see Section 5.4.2). In green, red, and blue we show the equilibrium, adiabatic (SIDM), and impulsive (SNF with 10 explosions) cases respectively. The bin size for J_r is 1 km kpc s^{-1}. The dashed lines are Gaussian functions with the mean and standard deviation of each distribution (see Table 5.1) as their parameters).

We also quantify how much the actual final distribution deviates from a Gaussian with the calculated mean and spread by calculating χ^2/dof (see Table 5.1). We also investigate the impact of the baseline numerical errors presented in Section 5.5.1 (see Fig. 5.3) on the numerical diffusion of the initially seeded Gaussian distribution. We do this by calculating the mean of the relative error in radius $\langle \delta_r/\langle r \rangle \rangle$ and the mean of the absolute error in radial velocity $\langle \delta_{v_r}^{\mathrm{max}} \rangle$. We then assume that the errors are Gaussian and use the mean errors to re-sample the tracers' phase space coordinates in the following way:

1. start with the initial Gaussian distribution $p(J_r)$ for the tracers.
2. replace the orbital radius r of the tracers by a randomly sampled value from a Gaussian distribution with a mean equal to r, and a standard deviation equal to the $\langle \delta_r/\langle r \rangle \rangle \times r$.
3. replace the radial velocities v_r of the tracers by a randomly sampled value from a Gaussian distribution with a mean equal to v_r, and a standard deviation equal to the $\langle \delta_{v_r}^{\mathrm{max}} \rangle$.
4. re-calculate the action with the new coordinates and velocity vector components.
5. repeat for every tracer and construct the new radial action distribution.

Since we are including the (average) maximum radial velocity baseline errors[3], we expect that this estimate will overestimate the numerical diffusion, if there are indeed no other sources of diffusion.

In Table 5.1 we give an overview of the main characteristics of the radial action distributions in the different cases. Since it is difficult to achieve a full convergence of the iterative rejection sampling algorithm in the tails of the Gaussian distribution, we note that we have removed $\lesssim 5\%$ of the tracers that could not be assigned a sensible value of J_r. In particular, we limit our analysis to those tracers that are initially at least within 5σ of the mean of the initial Gaussian distribution. We then tag these particles and follow their evolution throughout the various simulations.

By looking at the results in Fig. 5.13 and Table 5.1 we draw several interesting conclusions. Firstly, by comparing the result of the simulation that evolves in equilibrium (green in Fig. 5.13) with the re-sampled distribution explained above (bottom row of Table 5.1), we can be certain that the numerical diffusion of the radial action distribution is fully contained within the expectations of the baseline orbital errors described in Section 5.5.1 (notice that the spread of the re-sampled distribution is larger than that of the final configuration of the equilibrium simulation). Thus, the growth of the spread in the equilibrium case by the end of the simulation represents the degree to which the conservation of radial action is violated purely from numerical errors. This is quantified by looking at the difference in the standard deviations of the equilibrium and initial distributions shown in Table 5.1, which amounts to $\sim 50\%$. Looking at the SIDM case, we find that the result is consistent with our expections based on the adiabaticity of the change in potential in this case. The radial action distribution is still well-fitted by a Gaussian that is quite close to the equilibrium case, albeit with slight distortions and a larger spread. The degree to which the radial actions are conserved in the SIDM case is remarkable.

The key result in this Section is the comparison between the final distributions of the impulsive (blue line) and adiabatic (red line) cases of core formation. It is clear that the initial Gaussian distribution is entirely dissolved in the impulsive case, with a final distribution considerably broader and skewed towards low J_r values, relative to the mean of the original distribution, and with a long tail of very large J_r values. The spread is in fact larger than in the SIDM case by a factor of ~ 7.5 and the interpretation of the final distribution as a Gaussian is significantly worse than in the SIDM case (see the values of χ^2/dof in table 5.1). The striking difference between the SIDM and SNF case in the degree to which radial actions are conserved in the former, and violated in the latter offers a promising guideline to look for observational quantities in cored dwarf galaxies that are related to radial actions, which could probe the adiabatic or impulsive nature of the core formation mechanism.

5.5.4 A Plummer sphere

For the final configuration of tracers, we investigate how a Plummer sphere with a half-light radius of 500 pc (typical of the bright Mikly Way satellites) evolves in the three scenarios. To obtain an adequate comparison to the behavior of the DM particles,

[3]We recall that these are the maximum values over the tracers' orbits in the baseline benchmark, averaged over all tracers; see Section 5.5.1.

Simulation	σ_T/m_χ [$\mathrm{cm^2 g^{-1}}$]	Exploding mass [$\mathrm{M_\odot}$]	mean J_r [$\mathrm{km\,kpc\,s^{-1}}$]	spread in J_r [$\mathrm{km\,kpc\,s^{-1}}$]	Number of tracers	χ^2/dof
Equilibrium	–	–	43.7	5.3	1928	1.2
SIDM	2.0	–	43.1	5.9	1928	1.7
SNF	–	$10 \times 4.68 \times 10^8$	59.7	**44.2**	1928	**5.7**
Initial	–	–	44.9	3.5	1928	1.8
Re-sampled	–	–	50.0	7.2	1901	1.8

Table 5.1. The columns list the main properties of the final radial action distributions for the tracers in each of the simulations shown in Fig. 5.13, as well as the initial and re-sampled Gaussian distributions. Column(1): simulation case; Column (2): self-interaction transfer cross-section per unit mass for the SIDM case; Column(3): the total mass blowout in the SNF case (in 10 explosions); Columns (4-5) the mean and standard deviation of the radial action distribution; and Column (6): the amount of tracers taken into account in these distributions. The reason for not taking all of the 2000 particles into account each time lies in inaccuracies of the initial rejection sampling algorithm (see text for details). Finally, Column (7) shows the χ^2/dof with respect to a Gaussian with the calculated mean and spread in each case. Notice the large deviation of the SNF over the Gaussian distribution (highlighted in bold font).

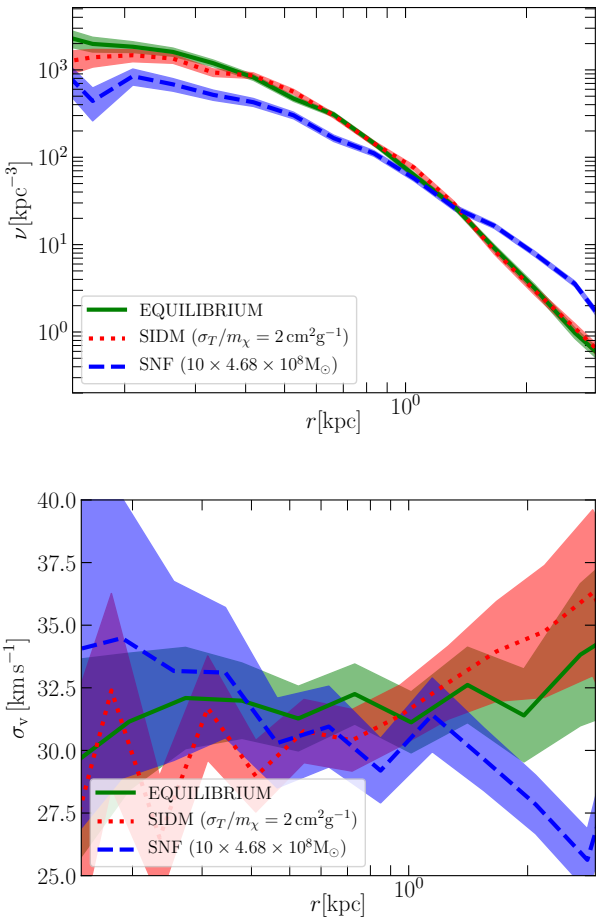

Figure 5.8. Impact on the stellar distribution (Plummer sphere) due to cusp-core transformation in the impulsive (SNF) and adiabatic (SIDM) scenarios. We show the number density and velocity dispersion profiles of the stars in the top and bottom panels, respectively. The stars are modelled as tracers initially following a Plummer density profile with a half light radius of 500 pc. The initial setup as an equilibrium configuration works as in the case of the Hernquist halo. The SIDM case is shown after the isothermal core is fully formed, while the SNF and equilibrium cases are shown after 4.2 Gyrs, which is the time where the SNF DM core has the same central density as in the SIDM case (see Fig. 5.1). The different cases are colour-coded as in previous figures. The error bands correspond to Poisson counting errors, which are larger than the (systematic) numerical integration errors (see Fig. 5.3). The number density in the impulsive (SNF) case is modified significantly relative to the equilibrium and adiabatic (SIDM) cases. The stellar kinematics (velocity dispersion profile) imply that the final SNF configuration is likely not in equilibrium.

we show the profiles at the same times as in Fig. 5.1 (after 4.2 Gyrs for SNF and equilibrium, after 1.2 Gyrs for SIDM). Fig. 5.8 shows the number density (top panel) and the velocity dispersion (bottom panel) profiles. The solid lines show the calculated mean of the radial bins while the shaded area in the top panel shows the (Poisson) counting variance. The shaded area on the bottom panel shows the sampling errors on the velocity dispersion calculated from the biased estimator of the variance when assuming an underlying Gaussian distribution (see, for instance Lehmann and Casella 1998),

$$\delta_{\sigma_v} = \frac{\sigma_v}{\sqrt{2(N_{\rm shell}-1)}}, \tag{5.20}$$

where $N_{\rm shell}$ is the number of stars within the radial shell in which we calculate the velocity dispersion.

As in the other two configurations of tracers, we find that there is a significant difference in the final tracer distribution depending on whether core formation proceeds adiabatically or impulsively. In the adiabatic (SIDM) case, the radial distribution of the Plummer sphere remains close to the equilibrium case, showing only a slight reduction in central density (comparing the red and green lines in the top panel of Fig. 5.8). The velocity dispersion profile also responds adiabatically, moving from an almost flat profile (in the radial range shown in Fig. 5.8) at the beginning (green line) to a slightly steeper profile with positive slope. The adiabatic formation of the DM core drives the slow evolution of the tracers towards a new state of Jeans equilibrium. These results roughly agree with previous ones from SIDM cosmological simulations (see Figs. 6 and 8 of Vogelsberger et al. 2014). Comparing both the number density profiles and the velocity dispersion profiles of the tracers in the SIDM and the equilibrium cases, it is evident that the Plummer sphere evolution in the SIDM case is quite mild relative to the evolution of the DM profile (see Fig. 5.1). In fact, just as we have found in the previous Sections 5.5.2 and 5.5.3, there is no clear observational signature of the formation of the core by just looking at the spatial and kinematic distribution of tracers (stars).

The impulsive case of core formation (SNF with 10 explosions), on the other hand, triggers a strong evolution of the Plummer sphere. Particularly, the number densities are considerably lower than in the equilibrium case within the half light radius of the sphere with the tracers being pushed out towards radii beyond ~ 1 kpc, which would correspond to a more extended galaxy with a lower central surface brightness. The velocity dispersion profile of the SNF case has similar central values as in the other cases. However, it has a negative slope with radius, which is particularly apparent beyond 1 kpc. Although the sampling errors are relatively large, the evidence for this trend seems sufficient. The behaviour of the velocity dispersion profile seems to indicate that the tracers are not in isotropic equilibrium. This is most clearly seen within ~ 400 pc, where the DM density profiles are very similar in all simulations (see top panel of Fig. 5.1), while the profile of the tracers has essentially the same shape, but with considerable deficit in the SNF case relative to the other simulations (top panel of Fig. 5.8). Under the assumption of isotropic Jeans equilibrium, these conditions would imply that the central velocity dispersion in the SNF case should be lower than in the other simulations. Although the sampling errors are large in the centre, this is not what is seen in the bottom panel of Fig. 5.8, implying that the tracers are not in

isotropic equilibrium. This circumstance could be a transient state resulting from the strong fluctuations caused by the multiple accretion and explosion events in this SNF configuration.

Overall, we can conclude that the response of a Plummer sphere to the formation of a DM core *of the same central density* is considerably different in the adiabatic and impulsive cases. For the former, the properties of the sphere remain remarkably unchanged, while the destructive nature of the latter one causes the centre of the sphere to expand considerably. This expansion proceeds beyond the extent of the central DM core, thus putting the orbits of the tracers out of equilibrium (at least within the time scales we simulated).

5.6 Behaviour of a cosmological halo

The results so far are based on the assumption of an isolated, spherically symmetric DM halo. This was done in order to test the theoretical predictions of a direct link between conserved quantities derived from the initial symmetry of the halo (angular momentum and radial action) and the nature of the process responsible for triggering the formation of the DM core. However, it is evident that realistic DM haloes are better described as being triaxial rather than spherically symmetric, and that they are clearly not isolated but rather assembled by a hierarchical merging process throughout their cosmological formation history. When looking at real haloes we thus do not expect all of the predictions made in section 5.5 to remain the same. In particular, one can argue that the results shown in Fig. 5.7 will be quite different in realistic haloes because radial actions will no longer be conserved in haloes with a significant degree of triaxiality, and thus the differences between the adiabatic and impulsive core formation scenarios that are obvious in Fig. 5.7 will be diluted. This lack of conservation of radial actions for orbits in non-spherical potentials was for instance shown explicitly in Pontzen et al. (2015) for a 3D anisotropic harmonic oscillatior, and for a cosmological dwarf-size halo (see in particular their figures A1 and A2).

Regarding the predictions of Fig. 5.8, we argue that they will remain qualitatively the same, as long as orbits expand significantly in the SNF case and stay relatively localized in the case of SIDM. The most striking result we presented in the spherically symmetric case is that of the appearance of distinct features (shells) in phase space (discussed in Section 5.5.2 and shown in Figs. 5.5 and 5.6) due to the impulsive episodes of energy injection that lead to core formation in the SNF case. To test whether or not these features survive in a realistic halo, we took the zoom-in dark-matter-only simulation of a dwarf-size halo presented in Vogelsberger et al. (2014). This halo (labelled dA-CDM-B-hi) is similar in mass and size to the one we analysed here and it also has a similar numerical resolution; its main properties are given in tables 2 and 3 of Vogelsberger et al. (2014). The relevant ones for our purposes are: $M_{200} = 1.07 \times 10^{10} \mathrm{M}_{\odot}$, $r_{200} = 45.8\,\mathrm{kpc}$, $\varepsilon = 34.2\,\mathrm{pc}$, and a particle mass $m_{\mathrm{DM}} = 9.7 \times 10^{2} \mathrm{M}_{\odot}$ (equivalent to $\sim 1.1 \times 10^{7}$ particles within r_{200}).

To perform our analysis, we only consider the particles in the original simulation at $z = 0$ within a $(500\mathrm{kpc})^3$-sized box around the centre of potential of the halo which will be the origin of the reference frame used in the subsequent analysis. We calculate the

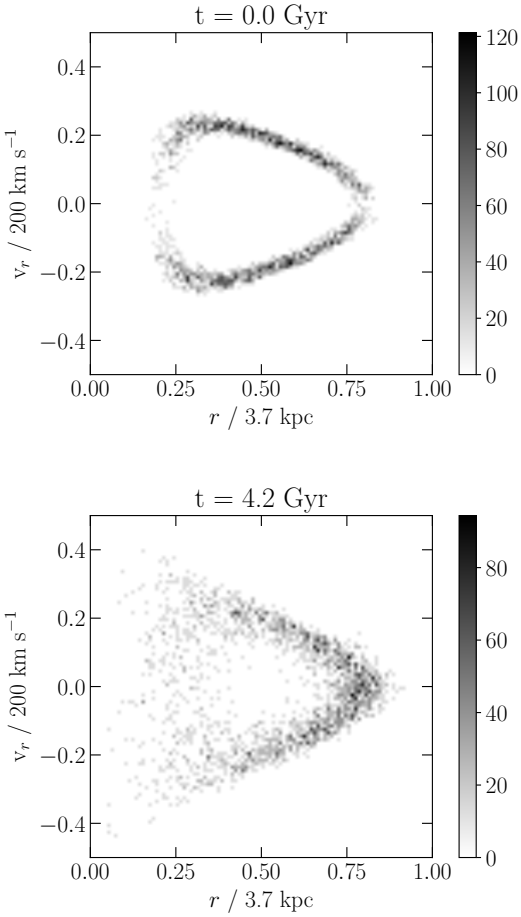

Figure 5.9. Evolution of the radial 2D phase space density of the orbital family described in Section 5.5.2, but set up within the $z = 0$ cosmological dwarf-size halo taken from Vogelsberger et al. (2014) and without any core formation mechanism. The top panel shows the initial configuration whereas the bottom panel shows the result after 4.2 Gyrs. As in figure 5.4, the scale on the x-axis is unchanged and all of the tracers remain within our initially chosen patch of phase space. However, there is a striking evolution in the phase space occupation which we attribute mainly to the triaxiality of the halo and the resulting non-conserved angular momentum. In particular, the final result exhibits a heavily broadened and scarce density in the pericentre vicinity relative to the original distribution of the orbital family.

centre of the potential iteratively considering all particles within concentric shells with radii of 75, 7.5 and 0.75 kpc using a weighted mean with the potential at the position of each particle as a weight. This is done during the initial setup of the tracer particles as well as during each timestep of the orbital evolution. By using the centre of the potential instead of the center of mass as the origin of the reference frame we are taking into account the asymmetric distribution of mass in the halo. We note that the halo chosen has had most of its major merger activity in the past and is relatively quiescent by $z = 0$.

5.6.1 Orbital family in a cosmological halo

To analyse the impact of having a cosmological halo in our results, we take the halo described above at $z = 0$ and calculate its spherically averaged enclosed mass profile and radially dependent potential and set up an orbital family of tracer particles as described in section 5.4.1. We then evolve the system for 4.2 Gyrs in time in three configurations: no core formation mechanism (baseline), SIDM and SNF. The first case will serve as our new baseline in which we can measure the impact on the evolution of the orbital family caused purely by the lack of spherical symmetry and having a subhalo distribution rather than just a smooth halo. We show the radial 2D phase space density evolution in this baseline case in Fig. 5.9: the top panel where we set the orbital family at $t = 0$, while on the bottom we show the configuration after $t = 4.2$ Gyr. Comparing this to Fig. 5.4 (the spherically symmetric smooth halo), we find that the final distribution is quite different in the two cases. In particular, we find a much stronger evolution - and dilution - of the initial phase space density in the cosmological halo than was the case in the idealized spherical halo. This is especially evident in the vicinity of the pericentres of the initial orbital family, where the initially narrow distribution is heavily broadened. Qualitatively however, we still observe that the area around the orbital apocentre of the orbits is more populated than other segments along the trajectories. Furthermore, even after 4.2 Gyrs, the occupied region in phase space still exhibits key features of the initial distribution. The most notable feature is that there is still a gap in phase space between the largest pericentre and the smallest apocentre radius. Moreover, we see no considerable net expansion of the radial range occupied by the orbital family.

Most likely, the main reason for the observed discrepancy between the results in the idealized case and in the case of the cosmological halo is the triaxiality of the latter. Since spherical symmetry is broken, angular momenta are no longer conserved, and thus, a configuration based on the premise of such conservation will not remain stable while evolving in a triaxial potential. We notice however, that for tracers orbiting in a potential which does not strongly violate spherical symmetry, a quantity closely related to the angular momentum is an integral of motion Binney and Tremaine (2008). This is particularly relevant in our case since we are interested in dwarf-size haloes in our analysis, which have been shown to be closer to spherical symmetry than larger haloes (see e.g. Fig. 13 of Bonamigo et al. 2015). This is likely the reason why the radial phase space density of the orbital family still remains confined to a relatively narrow region, which retains a similar shape to the original despite evolving significantly within a couple of dynamical times. This picture is supported by the fact that once we start the simulation, most of the evolution takes place within the first 300 Myrs, indicating that the orbital family quickly relaxes to a new steady state set by the actual shape of the

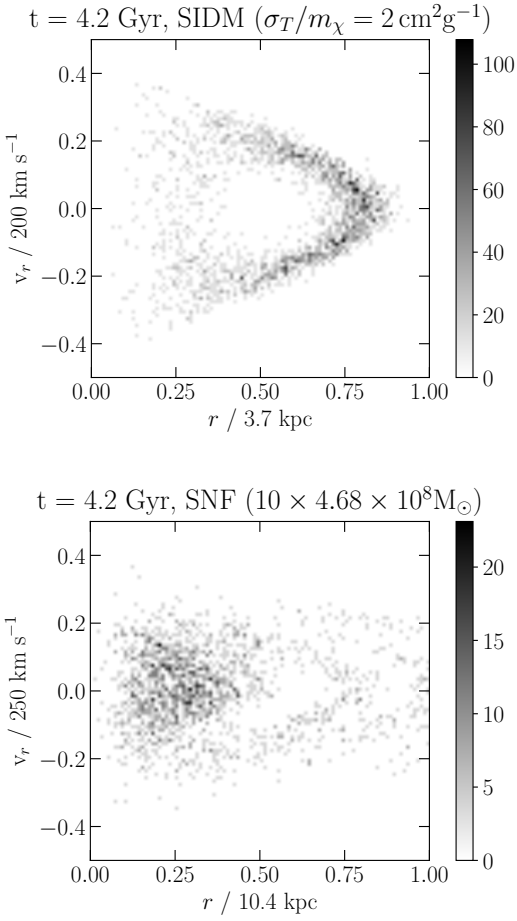

Figure 5.10. As Fig. 5.9 but in the case of core formation driven by SIDM (adiabatic, top panel) and by SNF (impulsive multiple explosions, bottom panel) in a cosmological halo. The final profile in the case of adiabatic core formation can hardly be distinguished from the final phase space density in the baseline case shown in the bottom panel of Fig. 5.9. The SNF case, however, clearly shows a strong average radial expansion compared to the SIDM case. Notice the change of scales on the x-axis in the bottom panel. The shell-like structures, clearly visible in Fig. 5.5 are no longer as distinct when the host halo is no longer idealised (spherical and smooth).

cosmological halo and then remains relatively unchanged until the simulation ends.

Fig. 5.10 shows the same comparison between the adiabatic and impulsive (multiple explosions) core formation scenarios in the case of a cosmological host halo as Fig. 5.5 did in the case of an idealized spherically symmetric halo. Just as in the idealized case, we identify strong differences between adiabatic and impulsive core formation in terms of the response of tracer particles: (i) the expansion of the orbits to considerably larger radii in response to impulsive mass removal remains a clear SNF signature; (ii) the final orbital distribution in the SIDM case can hardly be distinguished from that of the baseline case.

On the other hand, tn the case of SNF we find that, contrary to the idealised case, the pronounced shell structures that were clearly visible in Fig. 5.5 have disappeared almost entirely, with the final radial phase space density in the bottom panel of Fig. 5.10 being a mixture of several families overlapping and occupying most of the available phase space. Keeping in mind that in the multiple explosion case, the last of the 10 explosions happens at 3 Gyrs, after which the system relaxes into a new stable configuration, the question arises on which time-scales the shell-like structures dilute and how this compares to the dynamical time-scale of the tracers. As a benchmark, we estimate the dynamical time as the time one radial oscillation takes under the assumption of spherical symmetry, which is

$$t_{\rm dyn} = 2\int_{r_{\rm peri}}^{r_{\rm apo}} {\rm d}r \frac{1}{\sqrt{2\left[E-\Phi(r)\right]-\frac{L^2}{r^2}}}, \tag{5.21}$$

where all quantities are calculated from the initial peri- and apocentre radii and the spherically averaged potential. For our cosmological halo, we calculate a dynamical time of $t_{\rm dyn} \approx 170$Myrs.

To check how fast the shell-like features dilute, we looked at several outputs from different simulation times and assessed the respective phase space densities in the SNF case. We show the 2D radial phase space density of the orbital family at two representative times in Fig. 5.11. In the top panel, we see that at an early stage of the simulation (after two explosions and right around the time of the third) the shell-like structures are indeed there and -moreover- are very pronounced, akin to what we have seen in the case of the idealized halo at the end of the simulation in Fig. 5.5. Furthermore, we clearly identify an infall branch in the phase space density. It is thus the process of letting the tracer stars settle onto new orbits after a last sudden energy injection that dilutes these features in a non-spherically symmetric potential. To investigate how long it takes for these features to disappear, we looked at the evolution of the orbital family at several times after the explosion occurred. As an example we show the configuration at $t = 3.3$ Gyr (0.3 Gyr after the last explosion) in the bottom panel of Fig. 5.11. At this time, there are at least two distinct shells that can be identified, one with apocentre at around 10 kpc and a very faint one containing only the most energetic tracers and extending out of the radial range shown in the plot. At the same time, however, it is apparent that orbits have already mixed and features which emerged as a result of the impulsive events have been diluted, particularly within the phase space region close to the halo centre. When letting the system relax further, we find that this trend quickly continues outwards, with most of the features being unidentifiable after 3.6 Gyrs. We

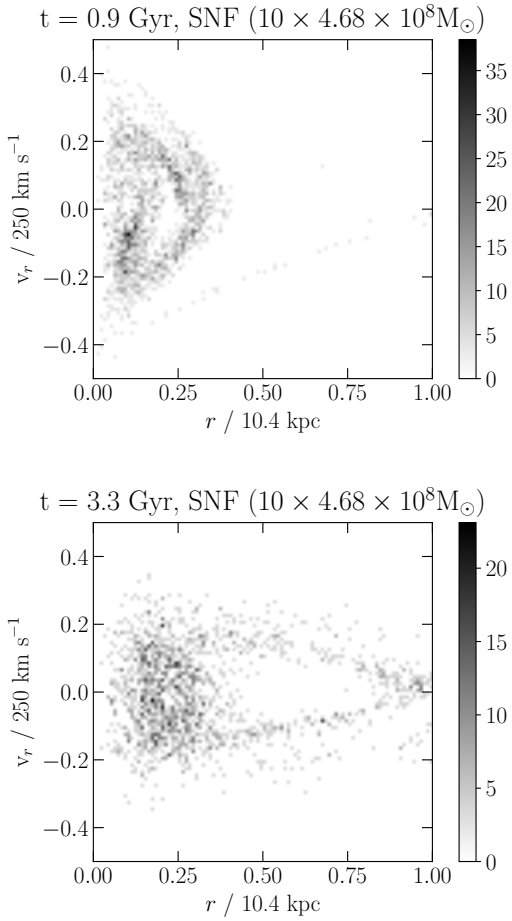

Figure 5.11. 2D radial phase space density of the orbital family at two different times in the evolution of the SNF (10 explosions) core formation case in a cosmological halo (prior stages to that shown in the bottom panel of Fig. 5.10). In the top panel, the time is $t = 0.9$ Gyr (around the time of the third explosion), whereas in the bottom panel $t = 3.3$ Gyr which is 300 Myr after the last explosion. In the top panel we can see shell-like features and a clear infall signature, which are the characteristic signatures of an impulsive event, as we pointed out in Fig. 5.5 for the idealised halo. The bottom panel shows that after letting the system relax for 300 Myrs after the last explosion, the phase mixing that appears in Fig. 5.10 (1.2 Gyr after the last explosion) is less advanced and impulsive features still remain.

thus estimate that the time during which effects of (impulsive) supernova-driven core formation could potentially be observed is $\lesssim 0.5$ Gyrs, which equals roughly 3 orbital revolutions of the tracer particles, i.e. three dynamical times.

We have shown that the key differences between the impulsive (SNF) and adiabatic (SIDM) core formation scenarios remain even when we abandon the assumption of a spherically symmetric isolated halo, and take instead a halo assembled in a cosmological setting, despite the fact that radial actions are no longer conserved in a non-spherical potential, an assumption that is at the core of most of our analysis. The fundamental complication in the more realistic setting of a cosmological halo is that phase mixing occurs within a few orbital time scales after an (impulsive) mass blowout episode in the SNF case, which tends to erase the remarkable shell-like features in phase space that are the clearest signatures of the impulsive scenario. Despite this complication, our results suggest that the differences in the stellar orbits between the adiabatic and impulsive core formation scenarios are relevant enough to warrant further investigation in a more realistic modelling (cosmological simulations with a full treatment of baryonic physics) and with actual stellar kinematic data in dwarf galaxies. We plan to move forward along both of these avenues in future work.

5.7 Conclusion

The surprising deficit of dark matter (DM) in the innermost regions of dwarf galaxies relative to the naive expectations of the Cold Dark Matter (CDM) model is a notable observational feature that remains poorly understood. Perhaps the most likely possibility to explain this DM mass deficit is to have a mechanism that efficiently deposits energy into the inner halo, effectively transforming the CDM cusp into a central core. Among the various possibilities, we focus on two which have been invoked extensively: supernovae and DM self-interactions (SIDM). In those two cases, the energy required to transform a cusp into a core originates from energy sources of a completely different nature, yet both are viable causes of core formation as they are capable of lowering the inner densities of CDM haloes to the observed levels (for references see Section 5.1).

In the case of supernovae, the viability of this process to explain the ubiquity of DM cores in dwarf galaxies has been discussed in terms of whether there is sufficient energy to achieve this goal (e.g. Peñarrubia et al., 2012) and whether the energy injection is efficient enough (e.g. Tollet et al. 2016 versus Bose et al. 2019a). On the other hand, in the case of SIDM, the discussion has focused on whether the SIDM transfer cross section per unit mass σ_T/m_χ needed to generate DM cores in dwarf galaxies, which is $\mathcal{O}(1)\,\mathrm{cm}^2\mathrm{g}^{-1}$, is consistent with a wide range of other astrophysical observations (see Table 1 of Tulin and Yu 2018 for a recent compilation of SIDM constraints).

In this paper we take a different perspective and focus on the problem of how to distinguish these two mechanisms of cusp-core transformation. We approach this question in a fundamental way by considering the time scales of energy deposition relative to the dynamical time scales of DM orbits in the inner halo. If the former is much larger than the latter, as in SIDM, then the process is adiabatic, while in the opposite regime the process is impulsive and irreversible. This is the regime where supernova-driven mass outflows are effective in forming DM cores Pontzen and Governato (2012).

We argue that by looking for signatures of the adiabatic or impulsive nature of observed DM cores in dwarf galaxies, we can provide evidence for or against the necessity of new DM physics. This paper is a first exercise towards accomplishing this goal in which we analyse in detail the different evolution of stellar orbits in an idealized (spherical) potential that changes in an adiabatic or impulsive way. Furthermore, we also investigate the implications of having a realistic cosmological dwarf-sized halo providing the background potential instead of an idealised spherically symmetric halo.

Our idealised halo model consists of a truncated Hernquist sphere with a virial mass of $M = 1.48 \times 10^{10} M_{\odot}$, a virial radius of 52 kpc and a scale radius of 3.5 kpc. These parameters are typical of the haloes corresponding to the dwarf galaxies where DM cores of $\mathcal{O}(1)$ kpc are common. The halo is discretised with $N = 10^7$ simulation particles and set up initially in equilibrium in a self-consistent way (see Section 5.3.1 and Appendix 5.9). To analyse what happens in a more realistic setting we also take the DM-only zoom-in simulation of a dwarf-size halo described in Vogelsberger et al. (2014), which has a virial mass of $1.1 \times 10^{10} M_{\odot}$, and a mass and spatial resolution similar to ours (1.1×10^7 particles within the virial radius, and $\varepsilon = 34$ pc.) We then add the two different mechanisms of core formation into either one of the haloes and follow the evolution using the AREPO code Springel (2010):

- *supernova-driven mass blowouts* are modeled by adding an external centrally concentrated Hernquist potential, which is then suddenly (impulsively) removed. This is either done in a single event that removes 10% of the idealised halo mass, or in 10 consecutive explosions, each removing $\sqrt{10}$% of the halo mass. These masses are chosen to satisfy the energy requirements to form a core of ~ 1 kpc in either the single or multiple explosion cases (see Section 5.3.2).

- *DM self-interactions* are added by letting the simulation particles collide elastically in a probabilistic way following the algorithm described in Vogelsberger et al. (2012). The transfer cross section is fixed to $\sigma_T / m_\chi = 2\,\mathrm{cm}^2\mathrm{g}^{-1}$. This cross section is chosen to form a DM core of the same central density and similar size as in the supernova case (see Section 5.3.4 and Fig. 5.1).

To study the impact of the cusp-core transformation on stellar orbits, we add 2000 tracer particles into the simulations in three different idealised configurations. The first two are constructed to study key fundamental differences between the adiabatic and impulsive transformations and follow from simple theoretical expectations of the behaviour of the stellar orbits (see Section 5.2 and Appendix 5.8). We test for the conservation of radial actions under an adiabatic change of potential in the idealized halo which is predicted by the time averages theorem. Furthermore, we look at the segregation of stars that were initially in a similar orbit (but with distinct orbital phase) into distinct orbits after the impulsive transformation (i.e. violation of the time averages theorem; see e.g. Binney and Tremaine 2008). We first test this in the idealised halo and then check by how much the results change in the cosmological halo. We find strong evidence in support of these theoretical expectations. The setup for the three configurations and our main results are as follows.

1. *Orbital family.*- The tracers are originally set up with a similar energy and (magnitude of) angular momentum, which is equivalent to sampling similar orbits

(similar apo- and pericentre but different phases). This leads to a small region in $r - v_r$ space being populated at a rather large density. This region remains the same, within the numerical errors in our simulations, if there is no change in the potential (see Fig. 5.4). The response of the orbits (and the evolution of the populated area in phase space) is dramatically different in the SIDM and supernova-driven core formation scenarios. In the SIDM case, we find that the region in phase space which is occupied by the tracer particles remains remarkably compact, proving that the SIDM case is fully in the adiabatic regime (see top panel of Fig. 5.5). In the case of supernova-driven mass blowouts, the situation is drastically different: the orbital family is split into distinct orbital families, depending on the location of the tracer on its orbit (phase) when the explosion(s) happen (a peculiar consequence of the violation of the time averages theorem). In the idealised spherically symmetric halo, the orbital families form a distinctive structure of concentric shells in the 2-dimensional phase space (see bottom panel of Fig. 5.5 and Fig. 5.6). Although these shell-like structures are more evident in the single explosion case, traces of them clearly remain even in the more chaotic multiple explosion case. When looking at the evolution in the cosmological halo, we find that there is a significant amount of phase mixing after roughly three dynamical times. This phase mixing leads to a broadening of the originally occupied region in phase space (see Fig. 5.9) and a relatively quick dilution of the shell-like structures in the impulsive case (see Figs. 5.10 and 5.11). We attribute this to the triaxiality of the host halo, which implies that angular momentum is not conserved. Nonetheless, the final occupied region in phase space remains remarkably compact in the case of SIDM and the aforementioned shell-like signatures of an impulsive event such as a supernova are present, albeit only for about 0.5 Gyrs.

2. *Gaussian distribution in radial action.*- Since the radial action of tracer orbits is only conserved in a spherically symmetric host potential, we perform this test only in the idealised halo to confirm our underlying theoretical expectations. The orbits of the tracers are initially sampled to have a narrow Gaussian distribution in radial action. If the potential does not change, then this distribution should remain the same, and indeed it does in our simulation with a small level of numerical diffusion. The evolution of the radial action distribution is drastically different in the two mechanisms of core formation. In the SIDM case, radial actions are remarkably close to being conserved, with the final distribution remaining very close to the initial one (see Fig. 5.7). There is, however, a small distortion in this adiabatic case towards a broader distribution, which is worth studying in the future (e.g. along the lines presented in Peñarrubia 2013). In the supernova-driven core formation scenario, however, there is a substantial evolution of the distribution of radial actions, which emphasizes the degree to which actions are not conserved in the impulsive regime The final distribution is considerably broader with a long tail of high radial action values, and no longer consistent with a Gaussian distribution (see Fig. 5.7). This confirms the original expectation that the fundamental difference between adiabatic and impulsive core formation in a spherically symmetric host potential is whether or not the radial actions of tracer

(star) particles are conserved.

3. *Plummer sphere.-* To test how a more realistic stellar distribution would react to the different mechanisms of core formation, we initially arrange the tracers in an isotropic Plummer sphere with a half-light radius of 500 pc set up in Jeans equilibrium within the idealised halo (the sphere remains a collection of tracers, i.e., it exerts no gravity). We find relevant differences between the two core formation scenarios. In the SIDM case, the sphere responds by (slightly) expanding adiabatically towards a new state of Jeans equilibrium, with only mild deviations from the original density and velocity dispersion profiles (see Fig. 5.8). This agrees with previous results based on SIDM cosmological simulations Vogelsberger et al. (2014). In the supernova-driven case, on the other hand, the Plummer sphere expands considerably in response to the formation of the DM core, which we emphasize has the same central density as in the SIDM case. This eventually results in a more extended sphere with a lower central density (see Fig. 5.8), corresponding to a galaxy with lower central surface brightness. Furthermore, we find that the episodic supernova events tend to drive the stellar configuration out of Jeans equilibrium for a rather long time (our simulation stops 1.2 Gyr after the last explosion and the sphere remains out of equilibrium). This implies that, contrary to SIDM, repeated supernovae (in the impulsive regime) have a lasting impact on the spatial distributions, as well as the Jeans equilibrium, of stars in dwarf galaxies.

In summary, we have quantified in detail up to which degree the process of re-distributing energy can be considered adiabatic in the SIDM case of core formation. Moreover, we have quantified the amount of energy (mass removal) which is required to form a core with similar characteristics in the impulsive regime. We have found that with the two processes constituting extremes of the adiabatic-impulsive spectrum (SIDM being very close to adiabatic and SNF being entirely impulsive), the consequences for stellar orbits are severe. In the idealised cases we have studied, we find that the impulsive case preserves information about the initial locations and velocities (phases) of the tracers before the cusp-core transformation. This information proves to have a sizeable impact on the final stellar orbit. On the other hand, radial actions are not conserved in realistic cosmological haloes, which are not spherical and have a DM distribution that is not smooth. Moreover, the distinctive phase space signatures in the impulsive (SNF) case are diluted due to phase mixing within a few dynamical times in a cosmological halo. We nevertheless have shown that the impact of different core formation mechanisms on the orbits of tracer orbits remains substantial. We argue that these signatures in the stellar kinematics constitute a promising avenue towards distinguishing the different types of core formation mechanisms.

The interplay between different physical processes and the core formation mechanisms is expected to be quite complex, particularly for the SIDM case (for dwarf galaxies e.g. Vogelsberger et al., 2014; Fry et al., 2015; Robles et al., 2017). This interplay remains essentially unexplored when it comes to its effect on stellar orbits, and the type of signatures we have found here. We aim to comprehensibly explore this interplay in the near future using cosmological hydrodynamical simulations with full baryonic physics.

Acknowledgements

We thank Volker Springel for giving us access to the AREPO code, and thank Chaichalit Srisawat for useful discussions. We thank Andrew Pontzen for insightful comments that encouraged us to significantly improve our work. This work was supported by a Grant of Excellence from the Icelandic Center for Research (Rannís; grant number 173929−051). We acknowledge that the results of this research have in part been achieved using the PRACE Research Infrastructure resource CURIE based in France at CEA. The simulations in this paper were carried out on the Garpur supercomputer, a joint project between the University of Iceland and University of Reykjavík with funding from Rannís.

5.8 Appendix A - Impulsive vs adiabatic change in the potential: 1D harmonic oscillator toy model

To develop an intuition about the conceptual differences between impulsive and adiabatic core formation, we use the 1-dimensional harmonic oscillator toy model introduced in Pontzen and Governato (2012). For the impulsive regime, we simply reproduce the key results and equations derived by the authors, while we develop a detailed description of the adiabatic regime, which was only mentioned vaguely in Pontzen and Governato (2012).

The 1-dimensional model consists of a massless test particle responding to an external potential:

$$V(x;t) = V_0(t)x^2. \tag{5.22}$$

Here only the normalization of the potential V_0 is explicitly time-dependent while the functional form remains the same. Depending on whether the change in potential is impulsive or adiabatic, the solution to the equation of motion has to be obtained differently.

5.8.1 The impulsive regime

To model an instantaneous change in the potential, we impose a sudden frequency change from ω_0 to ω_1 at time $t = 0$. Before and after the frequency change, the test particle's equation of motion is that of an ordinary harmonic oscillator, which has the general solution

$$x(t) = A\cos(\omega t + \psi), \tag{5.23}$$

where the phase ψ and amplitude A have different values before and after the frequency change. The final amplitude A_1 and phase ψ_1 of the oscillator can then be computed from the initial ones by requiring that $x(t)$ and $\dot{x}(t)$ be continuous at $t = 0$, which results

in Pontzen and Governato (2012):

$$A_1^2 = A_0^2 \left[1 + \frac{\omega_0^2 - \omega_1^2}{\omega_1^2} \sin^2 \psi_0 \right] \tag{5.24}$$

$$\tan \psi_1 = \frac{\omega_0}{\omega_1} \tan \psi_0 \tag{5.25}$$

where A_0 and ψ_0 are the initial amplitude and phase, respectively. The key feature of the solution is that there is an explicit dependence of the final amplitude and phase on the initial conditions. In particular, if we imagine two oscillators moving with the same amplitude A_0 initially but with different initial phases, they will end up oscillating with different amplitudes and phases after an impulsive change in the potential. This is in essence a consequence of the time averages theorem not being valid for the impulsive case.

For 3-dimensional stellar orbits, this result suggests that stars that are initially on a common orbit (but with different phases) will in general end up on very different orbits after an impulsive gas outflow produced by supernovae. If we assume a flat (random) distribution of radial phases in the beginning and a subsequent substantial and explosive change in the potential ($\omega_1 \ll \omega_0$ in Eq. 5.25), then after the explosion, as the orbits expand, most of the particles will be significantly closer to the pericentre of their new orbit than to the pericentre of their previous orbit. This implies that the initially flat distribution of radial phases changes to a non-flat distribution with a bias towards phases close to π. In the case of the harmonic oscillator, starting from a flat prior $2\pi p(\psi_0) = 1$ results in a final distribution of phases given by Pontzen and Governato (2012):

$$2\pi p(\psi_1) = \left(\frac{\omega_0}{\omega_1} \cos^2 \psi_1 + \frac{\omega_1}{\omega_0} \sin^2 \psi_1 \right)^{-1}. \tag{5.26}$$

5.8.2 The adiabatic regime

Although a numerical solution for the adiabatic case is shown in Fig. 3 of Pontzen and Governato (2012), no details are given on how this solution was obtained. Thus, in the following we provide a full description of this regime.

To model an adiabatic change in the harmonic oscillator potential we use a smooth function for the time dependence of the frequency $\omega(t)$ which, contrary to the impulsive case, changes gradually from an initial value ω_0 to ω_1 over several oscillation periods. To be precise, the solution is subject to the conditions:

$$\lim_{t \to -\infty} \omega(t) = \omega_0 \tag{5.27}$$

$$\lim_{t \to \infty} \omega(t) = \omega_1 \tag{5.28}$$

and must obey the equation of motion:

$$\ddot{x} + \omega(t)^2 x = 0 \tag{5.29}$$

We propose as an ansatz that the solution has the general form of the harmonic oscillator at all times:

$$x(t) = A(t) \cos(\omega(t) t + \psi(t)), \tag{5.30}$$

where $\omega^2(t) = 2V_0(t)$, $A(t)$, and $\psi(t)$ are the frequency, amplitude and phase of the oscillation. In addition, we define a new variable:

$$\zeta = \omega + \dot{\omega}t + \dot{\psi}, \tag{5.31}$$

which is the total time derivative of the argument of the cosine function. Demanding that Eq. 5.29 be fulfilled leads to the following set of two equations:

$$\frac{\ddot{A}}{A} = \left(\zeta^2 - \omega^2\right), \tag{5.32}$$

$$\frac{\dot{A}}{A} = -\frac{1}{2}\frac{\dot{\zeta}}{\zeta}. \tag{5.33}$$

Now, since

$$\frac{d}{dt}\left(\frac{\dot{A}}{A}\right) = \frac{\ddot{A}}{A} - \left(\frac{\dot{A}}{A}\right)^2, \tag{5.34}$$

we can combine Eqs. 5.32 and 5.33 to find the following second order non-linear differential equation for the variable ζ:

$$\frac{\ddot{\zeta}}{\zeta} - \frac{3}{2}\left(\frac{\dot{\zeta}}{\zeta}\right)^2 = -2\left(\zeta^2 - \omega^2\right), \tag{5.35}$$

which can be solved numerically for a given functional form of $\omega(t)$. We can then solve for the amplitude and phase of the oscilator:

$$A(t) = A_0 \exp\left[-\frac{1}{2}\int_{-\infty}^{t}\frac{\dot{\zeta}}{\zeta}dt'\right], \tag{5.36}$$

$$\psi(t) = \psi_0 + \int_{-\infty}^{t}\left(\zeta - \omega - \dot{\omega}t'\right)dt'. \tag{5.37}$$

Assuming now that the frequency transition is fully adiabatic, we can theoretically halt it at any point and continue it later on. This implies that if the ansatz in equation 5.30 is correct, the motion of our test particle is described by an ordinary harmonic oscillator at all times during the transition with an (instantaneous) period given approximately by $T \sim 2\pi/\omega(t)$. As the variable ζ defined in equation 5.31 is the time derivative of the cosine's argument, adiabaticity thus demands that at any point in time

$$\zeta(t) \sim \omega(t). \tag{5.38}$$

This immediately implies that

$$\begin{aligned} A(t) &\sim A_0 \exp\left[-\frac{1}{2}\int_{-\infty}^{t}\frac{\dot{\omega}}{\omega}dt'\right] = A_0 \exp\left[-\frac{1}{2}\int_{\omega_0}^{\omega(t)}\frac{d\omega}{\omega}\right] \\ &= A_0\sqrt{\frac{\omega_0}{\omega(t)}} \end{aligned} \tag{5.39}$$

which has the asymptotic behaviour $A_{t\to\infty} = A_0\sqrt{(\omega_0/\omega_1)}$, and agrees with the results shown in Fig. 3 of Pontzen and Governato (2012) (specifically, for the numerical values used there, $A_{t\to\infty} \approx 1.78$). Using Eq.5.38 we can also approximate Eq. 5.37 for the phase $\psi(t)$ as:

$$\psi(t) \sim \psi_0 - \int_{-\infty}^{t} \dot{\omega} t' \, dt' \tag{5.40}$$

and see that the change in phase depends explicitly on the time coordinate, as different transitioning times demand a different phase change to ensure a smooth behaviour of the oscillator. Note also that, contrary to the impulsive case, the phase shift in no way depends on the initial phase.

Eqs. 5.39 and 5.40 are exact if the change in potential occurs fully adiabatically, i.e., in the limit of the frequency transition taking an infinite amount of time. In practice, this is never exactly the case. However, the adiabatic approximation in general works very well if the change in energy (or equivalently frequency) occurs slowly compared to the dynamical time of the system, which in this case is given by the instantaneous frequency of oscillation. For adiabaticity, we thus demand that

$$\left| \frac{\Delta E}{E_{\min}(\Delta t)} \right| \ll \frac{\omega_{\min}}{2\pi}, \tag{5.41}$$

where we choose the smallest characteristic time-scale and energy to obtain the most conservative bound. On the other hand, we have that the initial energy of the harmonic oscilator is

$$E_{\rm ini} = \frac{1}{2} A_0^2 \omega_0^2 \tag{5.42}$$

and using Eq. 4 of Pontzen and Governato (2012), the final energy in the adiabatic limit is:

$$E_{\rm fin} = \frac{1}{2} A_0^2 \omega_1 \omega_0. \tag{5.43}$$

If we thus lower the frequency of the harmonic oscillator by $\Delta\omega$ starting from ω_0 and ending at ω_1 during a time Δt, the condition for adiabaticity 5.41 can then be written as:

$$\left| \frac{\Delta\omega}{\omega_1 \Delta t} \right| \ll \frac{\omega_1}{2\pi}. \tag{5.44}$$

We will see that this condition will have to be fulfilled by any parametrization $\omega(t)$ in order for our ansatz 5.30 to provide a reasonable solution to the differential equation for the time-dependent harmonic oscilator (Eq. 5.29).

To calculate explicit solutions to Eq. 5.35, we use a few specific functions to model the change in frequency. In particular, we consider three different parametrizations for $\omega(t)$. The first one is a linear transition:

$$\omega(t) = \begin{cases} \omega_0 & t \le t_0 \\ \omega_0 + \frac{\omega_1 - \omega_0}{t_1 - t_0}(t - t_0) & t_0 < t \le t_1 \\ \omega_1 & t_1 < t \end{cases} \tag{5.45}$$

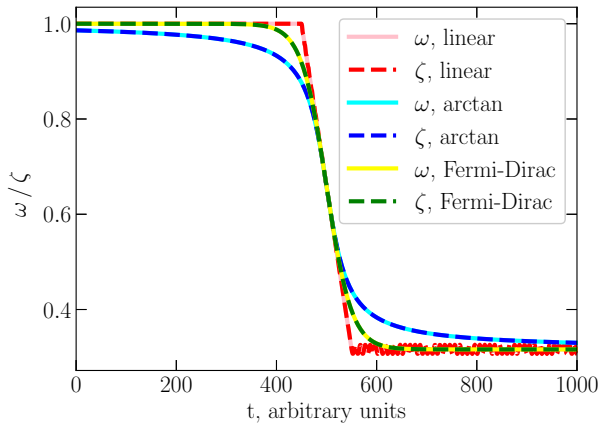

Figure 5.12. The solutions $\zeta(t)$ to equation 5.35 are shown alongside $\omega(t)$ for three parametrizations of $\omega(t)$ (as given in the legend; see Eqs. 5.45-5.47) with $t_0 = 450$, $t_1 = 550$, $\omega_0 = 1$ and $\omega_1 = \sqrt{0.1}$. Notice that the transition period between frequencies is strictly confined to the time interval $t_1 - t_0$ only in the 'linear' case, but in this case, the solution oscillates around a constant value after the frequency transition is complete. In the 'arctan' case on the other hand, the solution is stable but the transition takes a long time to be complete. The 'Fermi-Dirac' case constitutes the best parametrization we found in terms of representing a sharp transition with a stable long term solution.

which has the advantage that the frequency change is strictly confined to the transition period between t_1 and t_0. However, its obvious disadvantage is that at these two times the derivative of $\omega(t)$ is not continuous, which leads to numerical issues when solving the differential equation 5.35 as we show below.

To circumvent this problem, we explore two different parametrizations which do not have sharp transitions at the boundaries but where the frequency change is still confined mostly to the time interval from t_0 to t_1. The other requirement we impose for these parametrizations is that at the time $t = (t_0 + t_1)/2$, the slope of the frequency function be $(\omega_1 - \omega_0)/(t_1 - t_0)$, just as in the linear case. Following these conditions, we explore these two functions:

$$\omega(t) = \frac{\omega_0 + \omega_1}{2} + \frac{\omega_1 - \omega_0}{\pi} \arctan\left(\frac{\pi}{t_1 - t_0}\left[t - \frac{t_0 + t_1}{2}\right]\right) \tag{5.46}$$

$$\omega(t) = \omega_0 + \frac{\omega_1 - \omega_0}{\exp\left(-\frac{4}{t_1 - t_0}\left[t - \frac{t_1 + t_0}{2}\right]\right) + 1} \tag{5.47}$$

which we refer to as the 'arctan' and 'Fermi-Dirac' cases.

The solutions to equation 5.35 are shown in fig. 5.12 (solid lines) along with the corresponding functions $\omega(t)$ (dashed lines) for each parameterization and for the parameters $t_0 = 450$, $t_1 = 550$, $\omega_0 = 1$ and $\omega_1 = \sqrt{0.1}$. We find that $\omega = \zeta$ just as

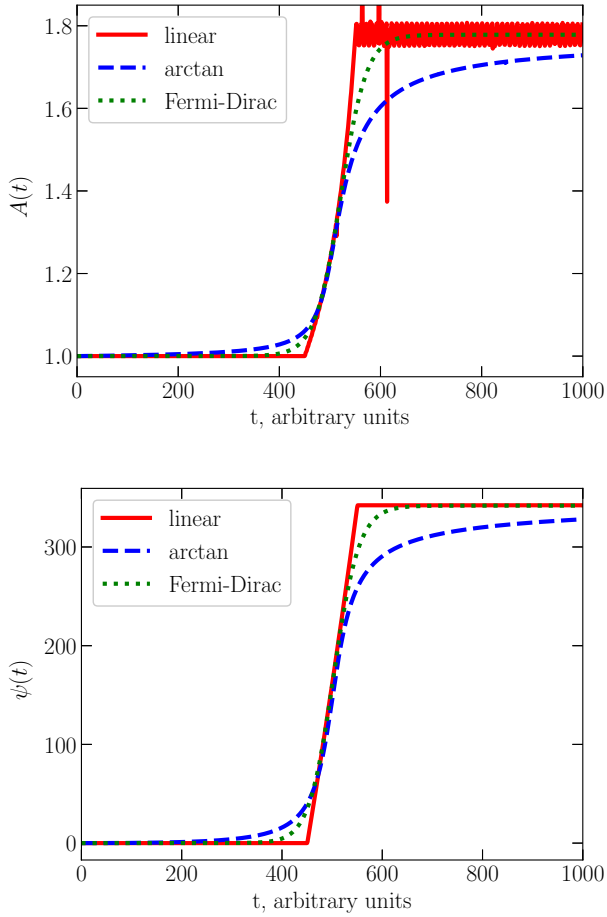

Figure 5.13. The solutions for the amplitude A (top panel) and phase ψ (bottom panel) of the harmonic oscillator under a change of frequency corresponding to the three different parametrizations shown in Fig. 5.12 (see Eqs. 5.45-5.47) for $A_0 = 1$ and $\psi_0 = 0$. For the phase, the transition is quickly reached in the 'linear' case, closely followed by the 'Fermi-Dirac' case, while the 'arctan' case takes much longer to complete the transition. When looking at the amplitude, it is clear that the 'linear' case is not appropriate since the amplitude shows numerical oscillatory features. Overall, the 'Fermi-Dirac' case is the most appropriate as the amplitude and phase reach their final values quickly and without numerical issues.

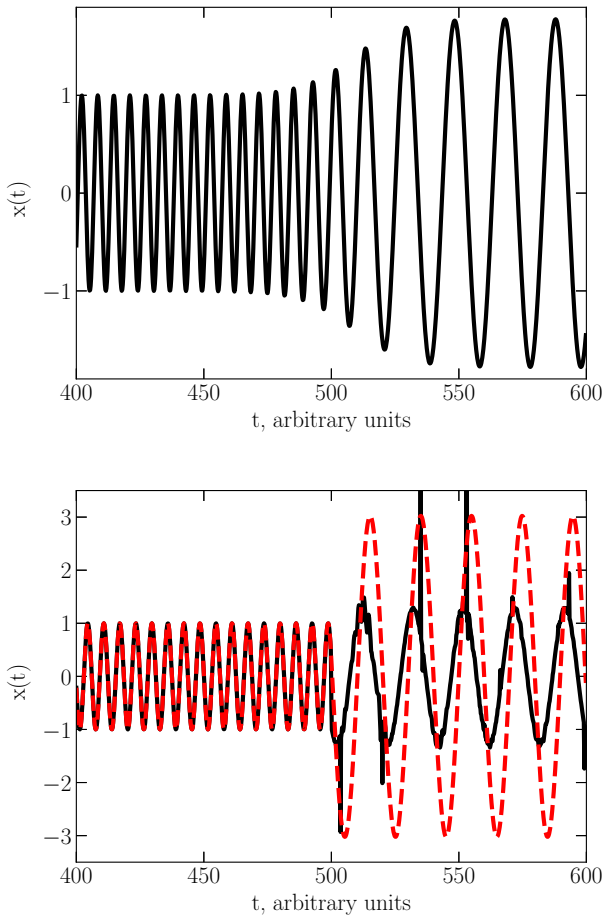

Figure 5.14. Solution to the harmonic oscillator equation under a transition in frequency with a 'Fermi-Dirac' parametrization (see Eq. 5.47). The top panel shows the case in which the frequency is altered over exactly two long orbital periods, whereas the bottom panel shows the case in which the transition period $t_1 - t_0$ is too small for Eq. 5.44 to be fulfilled. In the top panel the solution is smooth and behaves just as expected from an adiabatic case (see Fig. 3 of Pontzen and Governato 2012). In the bottom panel, on the other hand, we observe random oscillations of the calculated solution and a very non-harmonic behavior, which implies that the original ansatz that led to equation 5.35 is no longer applicable. Instead, the impulsive approximation is appropriate (see Eqs. 5.24 and 5.25), which is shown as a red dotted line. We caution the reader that in Eq. 5.24, the final amplitude of oscillation depends explicitly on the exact time of the transition.

predicted, and looking at the 'arctan' and 'Fermi-Dirac' cases, we find that precisely the condition 5.44 needs to be fulfilled in order for the left-hand-side of equation 5.35 to be small and thus to have $\omega \sim \zeta$ at all times. In the 'linear' case, however, we find the predicted numerical effects occuring due to $\ddot{\omega}/\omega$ not being small at the times t_0 and t_1. Since in the 'Fermi-Dirac' case the solution approaches the required asymptotic values significantly faster than in the 'arctan' case, we will focus on the 'Fermi-Dirac' case when conducting further analysis such as evaluating the solution to the original equation of motion. Fig. 5.13 shows the results for the amplitude (top) and phase (bottom) of the oscillations in all of the three cases. As anticipated, the discontinuity in the first derivative of the 'linear' parametrization of the frequency causes undesired numerical issues in the amplitude $A(t)$ of the oscillator. The phase shift, however, is not affected. Moreover, it is evident from fig. 5.13 that the solution to the 'arctan' case displays a convergence behavior which is much slower than that of the 'Fermi-Dirac' case. This implies that out of all the three cases, the 'Fermi-Dirac' one turns out to be the best one to describe a smooth transition between two frequencies within a fixed interval of time. This is further confirmation that all reconstructions of $x(t)$ should be based on solutions of the 'Fermi-Dirac' case.

The final result for the solution $x(t)$ is a smooth function which transitions continuously from a harmonic oscillator with a frequency ω_0 to a final harmonic oscillator with a frequency ω_1. The amplitudes and phases at the end of the transition are consistent with the approximations in Eqs. 5.39 and 5.40. In order to reproduce the results in Fig. 3 of Pontzen and Governato (2012), we need to choose the transition time interval to be equal to twice the longest orbital period possible:

$$\Delta t = \frac{4\pi}{\omega_1}. \tag{5.48}$$

We note that for our choice of parameters this transition time interval no longer strongly satisfies condition 5.44, since the left-hand-side and the right-hand-side of the equation are of the same order. Nevertheless, the adiabatic approximation still holds as can be seen in the top panel of fig. 5.14, where we show the solution to the harmonic oscillator equation in this case.

It is interesting to check at what point the adiabatic approximation breaks down in this simple toy model. Following Eq. 5.44, this should certainly happen when the period of time over which we change the potential is much shorter than the dynamical time-scale of the system:

$$\frac{\Delta t}{2\pi} \ll \frac{\Delta\omega}{\omega_1^2}. \tag{5.49}$$

To test this, we integrate Eq. 5.35 for $\Delta t = t_1 - t_0 = 1$, and leave the rest of the parameters unchanged with respect to the fully adiabatic case shown in the top panel of Fig. 5.14. The bottom panel of fig. 5.14 shows the numerical solution in this case (solid black line). It is apparent that when the change in potential is impulsive, the solution to equation 5.35 that we obtain by numerical integration is no longer sensible, which confirms that we can no longer use the ansatz 5.30. Instead we need to resort to the arguments summarized in Section 5.8.1 and derived in Pontzen and Governato (2012). A solution based on the latter is shown in the bottom panel of Fig. 5.14 as a red dashed line. We

caution the reader that a specific choice of phase underlies this solution, in particular one that maximises the growth of the amplitude.

5.9 Appendix B - Stability of haloes simulated in equilibrium

In order to asses the intrinsic differences between the adiabatic and impulsive core formation scenarios, we need to ensure that numerical artefacts are well under control in the central regions of the halo where core formation takes place. In particular we want to ensure that two effects are negligible in the timescales of our simulations within the region of interest for our work: (i) the growth of numerical errors coming from the initial setup of a DM halo in equilibrium, and (ii) the development of collision-less relaxation due to particle discreteness. In order to assess the magnitude of these two effects, we look at the stability of the halo set up in equilibrium in a self-consistent way in Section 5.3.1 (referred to as the equilibrium case throughout this work). Regarding (i) above, the method we have used (developed in Kazantzidis et al., 2004) has been used and tested extensively and we can be confident that numerical errors are under control as long as the amount of simulation particles sampled in the central region is large enough, which is connected to the second issue of particle discreteness in (ii).

Regarding collision-less relaxation, the relaxation time scale is given by Binney and Tremaine (2008):

$$t_{\rm relax} \approx \frac{N}{8 \ln N} t_{\rm cross}. \tag{5.50}$$

where N is the number of collision-less particles and $t_{\rm cross}$ is the crossing time. To ensure that an $N-$body representation of a DM halo does not suffer from gravitational two-body relaxation, the relaxation time needs to be longer than the simulation time. In order to achieve that, we need to avoid very large accelerations arising due to close encounters between simulation particles. Choosing an appropriate softening length ensures this Power et al. (2003). In particular, by demanding that the maximal gravitational force due to two-body encounters (equal to the maximum stochastic *softened* force) is smaller than the minimum mean field force in the simulation (which occurs at the virial radius), we get

$$a_{\varepsilon} \leq a_{\rm min} \equiv \frac{Gm}{\varepsilon^2} \leq \frac{GM_{200}}{r_{200}^2}. \tag{5.51}$$

From this it follows that a lower limit to the softening length required to avoid strong discreteness effects is:

$$\varepsilon \geq \frac{r_{200}}{\sqrt{N_{200}}} \tag{5.52}$$

where $M_{200} = N_{200} m$.

In CDM simulations, a safe choice is to set the softening length to 4 times the minimal one in Eq. 5.52, which is the choice we use in our work (see Eq. 5.1). With this choice, it has been shown by many works in the past that the density profile of

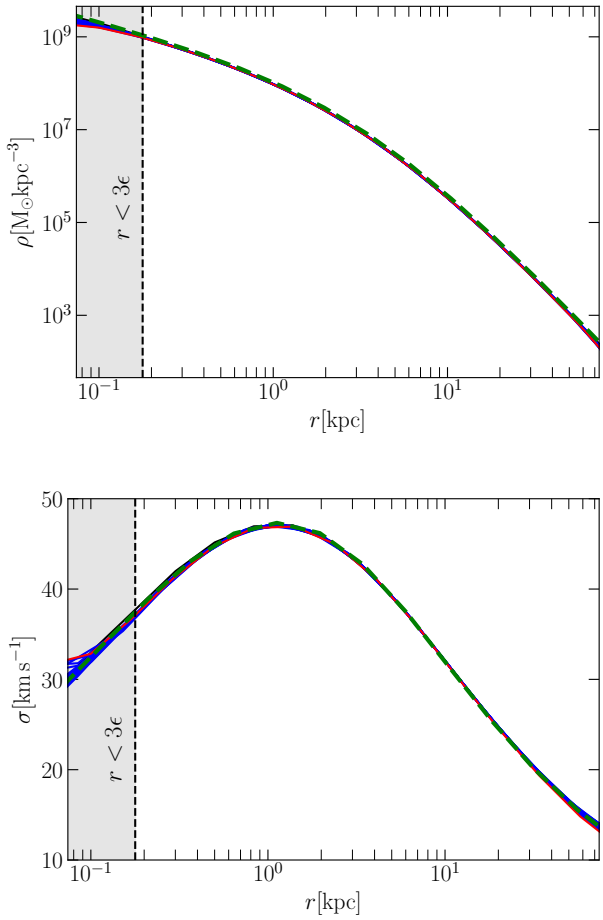

Figure 5.15. Stability of the the equilibrium configuration for the (dwarf-size) Hernquist halo used in this work. The evolution of the density (top) and velocity dispersion profiles (bottom) are shown over a simulation time of 4.2 Gyrs with outputs at every 0.3 Gyrs shown as blue solid lines. The initial profile is highlighted in black, whereas the final profile is shown in red. The vertical dashed black line marks the radius beyond which we expect collision-less relaxation to be negligible (the softening length ϵ is given by Eq. 5.1). The region where collision-less relaxation effects may play a role is shaded in grey. A halo in perfect equilibrium would match the analytic formulae for the Hernquist profile shown as dashed green lines. The equilibrium configuration setup we have used is stable for $r > 3\epsilon$.

CDM haloes converges at $r \sim 3\varepsilon$, across a broad range of resolution levels (particle number). This implies that two-body relaxation becomes negligible for $r > 3\varepsilon$. In SIDM simulations, the physical collisional relaxation produced by self-scattering and the corresponding development of a central density core results in SIDM haloes with density profiles that converge at even smaller radii than in CDM (e.g. Vogelsberger et al., 2012). We thus concentrate on the CDM case in its equilibrium configuration to explicitly check the stability of this configuration in the timescale of our simulations. Specifically, we investigate the stability of both the radial density and the velocity dispersion profile of the dwarf-size halo we have used as an initial condition for all our simulations when the core formation mechanisms are switched off (i.e. the equilibrium case). The properties of the halo are: $\mathrm{M}_{200} = 1.48 \times 10^{10}\,\mathrm{M}_{\odot}$, $r_{200}/r_{\mathrm{s}} = 15$, $N = 10^7$ and $3\varepsilon = 178$ pc. In this test, we run the simulation for a total of 4.2 Gyrs, and take snapshots every 0.3 Gyrs. The density and velocity dispersion profiles for each snapshot are computed in the centre of mass frame of the halo at the corresponding time of the snapshot. Fig. 5.15 shows the evolution of the radial profiles, density on the top panel and velocity dispersion on the bottom panel. Each simulation output is shown with blue solid lines, while the initial (final) times are shown with black (red) solid lines. We also show the analytical expectations for a Hernquist profile Hernquist (1990) as green dashed lines. The radial range in which the profiles deviate from each other is the range where the halo is no longer in equilibrium. In line with our expectations, only the innermost regions (well within 3ε, shown with a vertical dashed line and shaded in grey) are the ones that become visibly unstable within the simulation time. For $r > 3\varepsilon = 178$ pc, we find essentially perfect stability in both the density and the velocity dispersion profiles. This is 5 times smaller than the ~ 1 kpc size cores we study in this paper. Thus, we conclude that the halo equilibrium configuration we have set up is not affected by particle discreteness or numerical errors that could affect its stability at the relevant radii over the simulated time.

Chapter 6

Perturbative evolution of radial actions

This chapter is based on the following article:

Conservation of radial actions in time-dependent spherical potentials

Submitted to Monthly Notices of the Royal Astronomical Society, revised version

Authors:
Jan D. Burger[1 1], Jorge Peñarrubia[2] and Jesús Zavala[1]
[1]Centre for Astrophysics and Cosmology, Science Institute, University of Iceland, Dunhagi 5, 107 Reykjavík, Iceland
[2]Institute for Astronomy, University of Edinburgh, Royal Observatory, Blackford Hill, Edinburgh EH9 3HJ, UK

In slowly evolving spherical potentials, $\Phi(r,t)$, radial actions are typically assumed to remain constant. Here, we construct dynamical invariants that allow us to derive the evolution of radial actions in spherical central potentials with an arbitrary time dependence. We show that to linear order, radial actions oscillate around a constant value with an amplitude $\Delta J_r \propto \dot{\Phi}/\Phi\, P(E,L)$. Using this result, we develop a diffusion theory that describes the evolution of the radial action distribution of ensembles of tracer particles orbiting in generic time-dependent spherical potentials. Tests against restricted N-body simulations in a varying Kepler potential indicate that our linear theory is accurate in regions of phase-space in which the diffusion coefficient $\tilde{D}(J_r) < 0.01\, J_r^2$. For illustration, we apply our theory to two astrophysical processes. We show that the median mass accretion rate of a Milky Way (MW) dark matter (DM) halo leads to slow global time-variation of the gravitational potential, in which the evolution of

radial actions is linear (i.e. either adiabatic or diffusive) for ~ 84 per cent of the DM halo at redshift $z = 0$. This fraction grows considerably with lookback time, suggesting that diffusion may be relevant to the modelling of several Gyr-old tidal streams in action-angle space. As a second application, we show that dynamical tracers in a self-interacting DM (SIDM) dwarf halo (with $\sigma/m_\chi = 1\,\mathrm{cm}^2\mathrm{g}^{-1}$) have invariant radial actions during the formation of a cored density profile.

6.1 Introduction

Cosmological N-body simulations of the gravitational growth and collapse of primordial cold dark matter (CDM) density perturbations have been very successful in reproducing the observed large scale structure of the Universe (e.g Springel et al. 2005b). Such CDM N-body simulations also make concordant predictions on smaller scales, including for the abundance of small haloes and the inner structure of haloes in general (see Zavala and Frenk 2019). Moreover, N-body simulations are frequently used to model the dynamical evolution of virialized self-gravitating systems, such as globular clusters, individual DM haloes, or galaxies of different shapes and sizes. However, an interpretation of the results of such N-body simulations is not always straightforward. The behaviour of collisionless CDM haloes on small scales cannot be directly compared against observations owing to poorly understood baryonic effects. Therefore, it is difficult to assess whether mismatches between simulations and observations pose serious challenges to the CDM paradigm or not (Bullock and Boylan-Kolchin 2017). Another issue is that on small scales, the results of N-body simulations, while concordant, often stand unaccompanied by a theoretical explanation derived from fundamental principles. To understand the small scale predictions of CDM simulations and to disentangle baryonic effects from the long term gravitational evolution of virialized systems, it is thus desirable to derive from fundamental principles a suitable theoretical description of the latter.

One approach to modeling the evolution of self-gravitating systems is Hamiltonian perturbation theory (see e.g. Lynden-Bell and Kalnajs 1972, Tremaine and Weinberg 1984, Binney and Tremaine 2008). Hamiltonian perturbation theory is most useful if the evolution is due to a small perturbation over a smooth and stationary distribution of particles whose symmetry allows the Hamiltonian to be written as a function of a set of invariant actions. Such a Hamiltonian may then be expressed in terms of a power series of ε:

$$H(\mathbf{J},\theta,t) = H_0(\mathbf{J}) + \varepsilon H_1(\mathbf{J},\theta,t) + ... + \varepsilon^k H_k(\mathbf{J},\theta,t), \tag{6.1}$$

where $\varepsilon \ll 1$, t is the time coordinate, $\mathbf{J}$ is the three-vector of action variables, and θ is the three-vector of conjugate angles. The equations of motion derived from this expanded Hamiltonian are solved iteratively starting at the lowest order. Formulating Hamiltonian perturbation theory in action-angle space is particularly attractive, given that actions are adiabatic invariants (Binney and Tremaine (2008)). Recently, Hamiltonian perturbation theory has been formulated into a fully self-consistent kinetic theory in which the evolution of a self-gravitating system is governed by a set of equations akin to the Balescu (1960)-Lenard (1960) equations of Plasma physics (Heyvaerts 2010). Hamiltonian perturbation theory in general, and this formalism in particular, have been

quite successful in describing the secular evolution of self-gravitating systems. For instance, (Fouvry et al. 2015) and (De Rijcke et al. 2019) have modeled the formation of spiral arms in isolated rotating disk galaxies. While it is a very powerful approach from a conceptual point of view, it can prove rather difficult to solve the differential equations arising in Hamiltonian perturbation theory. In some cases more physical insight may be gained using a different approach, particularly if the evolution of the dynamical system is governed by a globally evolving gravitational potential instead of a localized perturbation.

A possible alternative to Hamiltonian perturbation theory is to treat self-gravitating objects as thermodynamical ensembles of particles. However, attempts at predicting the evolution of self-gravitating systems with the tools of statistical mechanics face a number of well-known difficulties. For one, particles interacting with each other gravitationally have negative specific heat (Antonov 1961,Lynden-Bell and Lynden-Bell 1977,Padmanabhan 1989; for a review see Lynden-Bell 1999). As a consequence, the evolution of systems of particles which are subject to long range forces cannot be described using canonical or grand canonical ensembles. Another property which is specific to large systems of particles which interact via the gravitational force and are not in dynamical equilibrium is non-ergodicity (Lynden-Bell, 1999), which implies that the time-average and the ensemble average of dynamical quantities are not equivalent. Yet a further problem was highlighted by Padmanabhan (1990) who argued that the long range nature of the gravitational force forbids the division of the system into non-interacting macrocells, since the energy of the system is now non-extensive. As a result of this, gravitating systems cannot be described by standard thermodynamics, a conclusion which has been supported by Levin et al. (2008) and Levin et al. (2014).

Despite these challenges, several attempts have been made to derive a valid statistical description of collisionless systems under gravity. Lynden-Bell (1967) sought to construct the equilibrium distribution function of a particle ensemble subject to a strongly time-dependent gravitational force of a newly formed galaxy. The author finds that in general, the most probable coarse-grained distribution function of said particle ensemble is that of a Fermi-Dirac gas, save a normalization factor. In the non-degenerate limit, which is applicable for galaxies, this distribution function can be approximated as a Maxwell-Boltzmann distribution, i.e. the distribution function of an isothermal sphere. The Maxwell-Boltzmann distribution is also obtained by Nakamura (2000) in a different approach using Jaynes (1957) information theory. While this result is accurate in the center of the potential, it fails to reproduce numerical experiments of violent relaxation (Arad and Johansson, 2005) in the outskirts. Moreover, it also implies that the system of particles has infinite total mass.

More recently, Pontzen and Governato (2013) attempted to derive the equilibrium distribution function of a virialized DM halo. Following arguments by Jaynes (1957), the authors state that statistical mechanics can still be applied to gravitating systems, provided that additional physical constraints other than just the conservation of energy are taken into account. In their formalism, they maximize the entropy of the system to derive its equilibrium configuration, using the additional constraint that the ensemble average of the DM particle's radial actions is approximately conserved. The resulting distribution function matches the properties of simulated haloes over several orders of magnitude. However, the authors need to include a second population of particles in

order to accurately predict the abundance of particles with very low angular momentum in DM haloes (see their Fig. 4). Without this second population their formalism fails to explain the inner density cusps that are ubiquitous in simulated DM haloes (Wang et al. 2020).

N-body simulations consistently show that a cuspy density profile is formed immediately after gravitational collapse of DM haloes. Subsequently, haloes evolve towards a universal mass distribution, which is well described by the single two parameter Navarro-Frenk-White (NFW, Navarro et al. 1996b, 1997) profile for virtually all simulated haloes[2], through minor mergers and diffuse accretion. It thus appears that Pontzen and Governato (2013)'s formalism captures the late evolution of the halo, but cannot explain how particles with low radial actions and low angular momenta (the "cusp" particles) are created during – or immediately after – the initial, impulsive gravitational collapse. The mechanism(s) that drive the accumulation of low angular momentum material in the central density cusp are not yet understood.

While central density cusps are ubiquitous in N-body CDM simulations of structure formation, the observed kinematics of several dwarf galaxies (Moore 1994, de Blok et al. 2008, Kuzio de Naray et al. 2008, Walker and Peñarrubia 2011) favour DM haloes with constant-density cores. Several scenarios have been proposed to reconcile the success of the CDM paradigm at explaining the large scale structure of the Universe with the apparent failure on smaller scales. Two frequently discussed mechanisms are supernova (SN) feedback (e.g. Navarro et al., 1996a; Pontzen and Governato, 2012) and self-interacting DM (SIDM, Spergel and Steinhardt 2000, Yoshida et al. 2000, Davé et al. 2001, Colín et al. 2002, Vogelsberger et al. 2012, Rocha et al. 2013). For SN feedback to be a feasible mechanism of cusp-core transformation, star formation needs to be bursty and cyclical (Pontzen and Governato 2012). Moreover, supernovae need to be energetic enough to unbind the cusp (Peñarrubia et al. 2012) and feedback is more efficient if baryons are more concentrated towards the center of the galaxy (Burger and Zavala 2021). SIDM is a feasible core formation mechanism if the momentum transfer cross section per unit mass σ/m is of the right magnitude $\sim 1\mathrm{cm}^2\mathrm{g}^{-1}$ (Zavala et al. 2013, Kaplinghat et al. 2016). Burger and Zavala (2019) demonstrated that the key difference between those two mechanisms is that cores are formed impulsively through SN feedback and adiabatically through SIDM. Hence, luminous tracers conserve radial actions in SIDM haloes, but not in CDM haloes with impulsive SN feedback.

In this article, we aim to develop a statistical theory for the evolution of the radial action distribution of an ensemble of tracer particles orbiting in a generic time-dependent spherical potential. In contrast to Hamiltonian perturbation theory, our statistical theory does not focus directly on the evolution of the distribution function but rather on individual particles; and then derives the evolution of the distribution function by treating the particles as a microcanonical ensemble. With this approach, we aim to clearly characterize the difference between the adiabatic and the impulsive regime at the level of individual particles. We further seek to determine how distributions of radial actions behave in potentials whose rate of evolution lies between those two regimes. By characterizing the behaviour of tracers in those three regimes, we aim to qualitatively understand how cusps form in collapsing haloes, how radial action

[2] A slightly improved fit can be obtained with the three parameter Einasto profile (Navarro et al. 2010), but the NFW profile still works remarkably well.

distributions evolve in a typical MW halo at the current time, and whether cusp-core transformation due to SIDM is truly an adiabatic process. At the current stage, our formalism does not take into account deviations from isotropy and instead focuses on globally evolving potentials. To develop our theory, we closely follow the work presented in Peñarrubia (2013) and Peñarrubia (2015). Peñarrubia (2013) generalized an argument of Lynden-Bell (1982), who found a coordinate transformation relating the equations of motion in Dirac's cosmology with a time-dependent gravitational constant G to the standard equations of motion in a frame in which G is constant. Peñarrubia (2013) showed that using a similar coordinate transformation, the equations of motion of particles orbiting in any time-dependent central potential can be solved in a frame in which the potential is static, provided that one is able to solve an auxiliary differential equation for a scale factor R that relates the spatial coordinates in the static frame to the original time-dependent frame. The evolution of a particle's energy in the time-dependent frame is then fully determined by $R(t)$ and the particle's phase space coordinates, while the energy in the static frame is a constant of motion (a.k.a. dynamical invariant). In a subsequent paper, Peñarrubia (2015) showed that if the evolution of the potential is slow enough, the evolution of the energy distribution of a set of tracers is diffusive and can be calculated statistically by treating the tracers as a microcanonical ensemble. In fact, the evolution of the energy distribution of tracers is fully determined by the drift and diffusion coefficients $\tilde{C}(E,t)$ and $\tilde{D}(E,t)$, which are related to microcanonical averages of the difference between a particle's energy and its dynamical energy invariant. To perform this average, it is not necessary to assume a phase-mixed particle distribution and hence this approach to statistical physics does not rely on the assumption of ergodicity.

Our paper is structured as follows: In Section 6.2, we derive the first order Taylor expansion of the time-dependent radial action J_r in the parameter $\dot{R}/R$. We show that to first order, the radial action oscillates with an amplitude ΔJ_r around a dynamical action invariant J_r' – the radial action in the static frame. The oscillation amplitude depends linearly on the radial period and can be calculated in general time-dependent spherical potentials, provided $R(t)$ is known. We test our model on a time-dependent Kepler potential, where analytic expressions for both $R(t)$ and the radial period are known. In Section 6.3, we derive the diffusion equation in radial action space and define drift and diffusion coefficients $\tilde{C}(J_r,t)$ and $\tilde{D}(J_r,t)$. Furthermore, we show that we can classify the evolution of radial action distributions as linear or non-linear, depending on whether $\sqrt{\tilde{D}} \ll J_r$ or $\sqrt{\tilde{D}} \geq J_r$. In Section 6.4, we test our diffusion theory, using restricted N-body simulations of five different tracer particle ensembles in a time-dependent Kepler potential. In particular, we test whether the diffusion formalism yields accurate results, both in cases where $\sqrt{\tilde{D}} \ll J_r$ (linear) and in cases where $\sqrt{\tilde{D}} \sim J_r$ (non-linear) on average. Based on the results obtained in Section 6.4, in Section 6.5 we apply our theory to two different astrophysical processes. We discuss if radial actions can be considered conserved quantities in Milky-Way (MW) size CDM haloes. To this end, we apply our formalism using the median mass accretion history of MW size haloes reported in Boylan-Kolchin et al. (2010) to estimate the fraction of DM particles within the MW halo whose radial actions are expected to show a linear rather than a non-linear evolution. We discuss implications of our results for the analysis of tidal streams in the MW today and the formation of central density cusps in DM haloes shortly after

gravitational collapse.

Furthermore, we also simulate core formation in a dwarf size SIDM halo. For $\sigma/m_\chi = 1\,\mathrm{cm^2 g^{-1}}$, we quantify how adiabatic core formation proceeds by determining the fraction of DM particles whose radial actions evolve linearly and applying the diffusion formalism developed in Section 6.3 to model the evolution of an initially Gaussian radial action distribution of tracer particles in the SIDM halo. We draw our conclusions in Section 6.6. Appendix 6.7 outlines how the scale factor is calculated numerically, Appendix 6.8 discusses a modification to the diffusion formalism developed in Sections 6.2 and 6.3, and in Appendix 6.9 we discuss an approximation to the scale factor in potentials that are not scale-free.

6.2 Radial actions in a time-dependent potential

6.2.1 Time-dependent radial action distributions

The radial action is an integral of motion in spherical static potentials (Binney and Tremaine, 2008). It is defined as

$$J_r = \frac{1}{\pi}\int_{r_\mathrm{peri}}^{r_\mathrm{apo}} dr \sqrt{2\left[E-\Phi(r)\right]-\frac{L^2}{r^2}}, \tag{6.2}$$

where E denotes the particle's energy, L its angular momentum and $r_\mathrm{peri,apo}$ the peri- and apocentre of the particle's orbit.

If the potential varies with time while retaining its spherical symmetry, radial actions are approximately conserved insofar as the change is 'slow' (Binney and Tremaine 2008). The question of exactly how slow the change in potential has to be for actions to be adiabatic invariants, however, is non-trivial. If the evolution of the gravitational potential is too fast, radial actions can no longer be considered adiabatic invariants and this can cause an asymmetric drift of radial action distributions. In Fig. 6.1 we illustrate this process on a population of tracer particles orbiting in a time-dependent Kepler potential,

$$\Phi(r,t) = -\frac{GM(t)}{r}. \tag{6.3}$$

Using a Kepler potential facilitates the calculation of radial actions as Eq. (6.2) has an analytic solution (e.g. Goodman and Binney 1984):

$$J_r = \frac{GM}{\sqrt{-2E}} - L. \tag{6.4}$$

We run a simple test simulation of 10^5 tracer particles orbiting in a time-dependent Kepler potential with a linear time dependence of the mass

$$\Phi(r,t) = -\frac{GM_0\,(1+\varepsilon t)}{r}. \tag{6.5}$$

The tracers' initial phase space coordinates are chosen at random. We require the initial orbital radius to be smaller than 300 kpc and that the angular momentum of each

particle is smaller than $150\,\mathrm{kpc\,km\,s^{-1}}$ to avoid very extended and energetic orbits. For definiteness, we choose $M_0 = 10^8 M_\odot$ and $\varepsilon = 1/30\,\mathrm{Gyr}^{-1}$ in Eq. (6.5). The mass of the Kepler potential grows by 10 per cent over 3 Gyr. In the following, this is the adopted benchmark model whenever we use simulations in a time-dependent Kepler potential to test our theory.

Using the N-body code AREPO (Springel 2010) we follow the tracer orbits and compute the distribution of the tracers' radial actions at different times to determine its time evolution.

In Fig. 6.1, we show the evolved radial action distribution at two different times: $t = 300\,\mathrm{Myr}$ (top panel), and $t = 3\,\mathrm{Gyr}$ (bottom panel). In the upper part of each panel, we show the final distribution $N(J_r,t)$ as a red solid line and the distribution at $t = 0$ as a black dotted line. We also show the final radial action distribution in a simulation in which the host potential is kept constant as a green dashed line. In the bottom part of each panel, we present the fractional change in the distribution $\delta N(J_r,t)$, defined as

$$\delta N(J_r,t) = \frac{N(J_r,t) - N(J_r,0)}{N(J_r,0)}. \tag{6.6}$$

We find almost no evolution when using a static Kepler potential, as should be the case, since actions are integrals of motion. In the simulation with a time-dependent host potential, however, we see a striking evolution of the radial action distribution. As time goes by, we observe a progressive flattening of the distribution's tail at $J_r > 100\,\mathrm{kpc\,km\,s^{-1}}$. At larger values, we find that $\delta N(J_r,t)$ slowly tends towards the limiting value of -1 (at $J_r \sim 300\,\mathrm{kpc\,km\,s^{-1}}$), indicating that almost no particles with radial actions larger than that remain. As a result of this, we find that $\delta N(J_r,t)$ is continuously larger than 0 at radial actions smaller than $100\,\mathrm{kpc\,km\,s^{-1}}$. Overall, this indicates a significant net drift of the initial distribution towards smaller radial actions.

6.2.2 First-order expansion of time-dependent radial actions

To attempt to understand the evolution observed in Fig. 6.1, we start by deriving the time evolution of the radial actions of individual tracers in time-dependent spherical potentials. Our calculation closely follows the one for energies presented in Peñarrubia (2013), which is based on a coordinate transformation found by Lynden-Bell (1982).

In a time-dependent potential, particles are subject to a time-dependent force

$$\ddot{\mathbf{r}} = \mathbf{F}(\mathbf{r},t). \tag{6.7}$$

If the force is conservative, Peñarrubia (2013) and Lynden-Bell (1982) show that for each phase space trajectory there exists a canonical transformation $\mathbf{r} \to R(t)\mathbf{r}'$ and a complementary transformation of the time coordinate $dt \to d\tau R^2(t)$ such that the time-dependence vanishes from the equation of motion, which can then be written as

$$\frac{d^2\mathbf{r}'}{d\tau^2} = \mathbf{F}'(\mathbf{r}') \tag{6.8}$$

The scale factor $R(t)$ is then a solution to the differential equation

$$\ddot{R}R^3\mathbf{r}' - R^3\mathbf{F}(R\mathbf{r}',t) = -\mathbf{F}'(\mathbf{r}'). \tag{6.9}$$

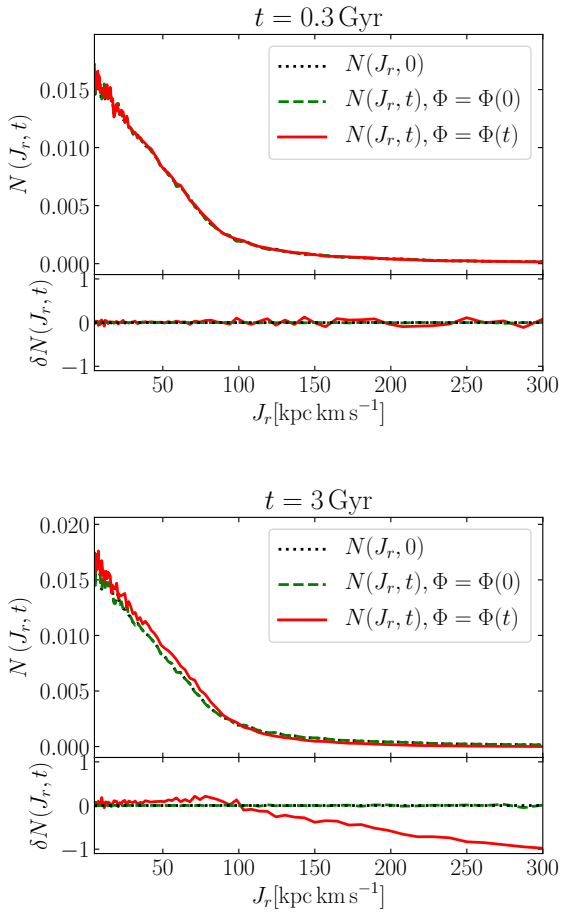

Figure 6.1. Evolution of the radial action distribution calculated from the orbits of 10^5 *tracer particles in a time-dependent Kepler potential (see Eq. 6.5). The black dotted lines show the initial distribution, the dashed green lines show the distribution evolved in a frame with a static potential, and the solid red lines show the distribution in the time-dependent frame. The simulation times at which the distributions are measured are indicated in the title of each panel. The top of either panel shows the distribution, whereas the bottom shows the fractional change with respect to the initial distribution. When evolving the distribution from 0.3 to 3* Gyr*, the radial actions of tracers in the simulation with a time-dependent potential tend to systematically drift towards smaller values.*

Notice that here, the time-evolution of the scale factor is coupled to the phase space trajectory. Peñarrubia (2013) shows that in the case of a slowly changing potential, an approximate analytic first-order solution to Eq. (6.9) exists for scale-free potentials. In Appendix 6.9, we demonstrate that a modification of this solution may be used for general spherical potentials.

Up to first order in $\dot{R}/R$, the energy of a particle which is subject to the force in Eq. (6.7) is

$$E \approx \frac{I}{R^2} + \frac{\dot{R}}{R}\left(\mathbf{r}\cdot\mathbf{v}\right), \tag{6.10}$$

where I is a dynamical invariant equal to the energy in the frame in which Eq. (6.8) is the equation of motion, i.e., the "time-independent", or static, frame.

Eq. (6.10) expresses the energy in the time-dependent frame as a function of the invariant energy and a first order correction that depends on the orbit of each particle. In the following, we aim to derive a similar expression for the radial action. A possible way to do so, presented in Peñarrubia (2013), is to simply insert Eq. (6.10) into Eq. (6.4) and identify the first order correction from there. However, since the radial action is not usually analytic in generic spherically symmetric potentials, it is desirable to derive a more general expression.

We start by integrating Eq. (6.9) on both sides to define

$$\tilde{\Phi}(\mathbf{r}') = -\int d\mathbf{r}'\cdot\mathbf{F}'(\mathbf{r}') = \frac{1}{2}\ddot{R}R^3 r'^2 + R^2\Phi(R\mathbf{r}',t). \tag{6.11}$$

The radial action in the time-independent frame is

$$J_{r'} = \frac{1}{\pi}\int_{r'_{\rm peri}}^{r'_{\rm apo}} dr'\sqrt{2I - 2\tilde{\Phi}(\mathbf{r}') - \frac{L^2}{r'^2}}, \tag{6.12}$$

where $J_{r'}$ is a dynamical invariant akin to I. We now seek to relate Eq. (6.12) to the radial action that one would de-facto measure in a time-dependent potential. To measure J_r in a time-dependent frame, one usually fixes the gravitational potential at the time one measures J_r. For this reason, we shall also refer to J_r as the "instantaneous" action. Fixing Φ at the time we measure J_r, we can define effective apo- and pericentre radii in the time-dependent frame. Formally, this means that we consider quantities in the time-dependent frame at a fixed time t_1. At this time, the invariant energy is

$$I = R^2(t_1)E(t_1) - R^2(t_1)\frac{\dot{R}(t_1)}{R(t_1)}\left(\mathbf{r}(t_1)\cdot\mathbf{v}(t_1)\right). \tag{6.13}$$

We then obtain

$$J_{r'} = \frac{1}{\pi}\int_{r'_{\rm peri}}^{r'_{\rm apo}} dr'\,R(t_1)\sqrt{f(\mathbf{r}',t_1)}, \tag{6.14}$$

where

$$f(\mathbf{r}',t_1) = 2E(t_1) - 2\frac{\dot{R}(t_1)}{R(t_1)}\left(\mathbf{r}(t_1)\cdot\mathbf{v}(t_1)\right) - \ddot{R}(t_1)R(t_1)r'^2 - 2\Phi(R(t_1)\mathbf{r}',t_1) - \frac{L^2}{R^2(t_1)r'^2}. \tag{6.15}$$

Now, we can perform the coordinate transformation $\mathbf{x} = R(t_1)\mathbf{r}'$, which is simply a shift of the radial coordinate. We write the effective apo- and pericentre radii in the time-dependent frame as

$$R(t_1)r'_{\mathrm{apo,peri}} = r_{\mathrm{apo,peri}}(t_1) \tag{6.16}$$

which is physically exact in the fully adiabatic limit. With these definitions,

$$J_{r'} = \frac{1}{\pi}\int_{r_{\mathrm{peri}}(t_1)}^{r_{\mathrm{apo}}(t_1)} dx\sqrt{f(\mathbf{x},t_1)} \tag{6.17}$$

with

$$f(\mathbf{x},t_1) = 2E(t_1) - 2\frac{\dot{R}(t_1)}{R(t_1)}\left(\mathbf{r}(t_1)\cdot\mathbf{v}(t_1)\right) - \frac{\ddot{R}(t_1)}{R(t_1)}x^2 - 2\Phi(\mathbf{x},t_1) - \frac{L^2}{x^2}.$$

In the limit $|\ddot{R}/R| \ll |\dot{R}/R|$ we can perform a Taylor expansion to first order in $\dot{R}/R$ to find

$$\begin{aligned} J_{r'} = {} & \frac{1}{\pi}\int_{r_{\mathrm{peri}}(t_1)}^{r_{\mathrm{apo}}(t_1)} dx\sqrt{2E(t_1) - 2\Phi(\mathbf{x},t_1) - \frac{L^2}{x^2}} \\ & - \frac{\dot{R}(t_1}{R(t_1)}\left(\mathbf{r}(t_1)\cdot\mathbf{v}(t_1)\right) \\ & \times \frac{1}{\pi}\int_{r_{\mathrm{peri}}(t_1)}^{r_{\mathrm{apo}}(t_1)} dx\,\frac{1}{\sqrt{2E(t_1) - 2\Phi(\mathbf{x},t_1) - \frac{L^2}{x^2}}} \\ & + \mathscr{O}\left(\left[\frac{\dot{R}(t_1)}{R(t_1}\right]^2, \frac{\ddot{R}(t_1)}{R(t_1)}\right). \end{aligned} \tag{6.18}$$

Dropping all the higher order terms, this leads to

$$J_{r'} \approx J_{r,t_1} - \frac{\dot{R}(t_1)}{R(t_1)}\left(\mathbf{r}(t_1)\cdot\mathbf{v}(t_1)\right)\frac{P(E,L,t_1)}{2\pi} \tag{6.19}$$

In this equation, J_{r,t_1} and $P(E,L,t_1)$ refer to instantaneous actions and radial periods measured at the fixed time t_1. Now, if the system is adiabatic, we can chose t_1 to be any time t and simply write

$$J_r \approx J_{r'} + \frac{\dot{R}}{R}(\mathbf{r}\cdot\mathbf{v})\frac{P(E,L)}{2\pi} \tag{6.20}$$

up to linear order in perturbation theory. As long as the change in the potential is slow enough, the quantity $J_{r'}$ in Eq. (6.20) is a dynamical invariant. In case of a faster change in the potential, higher order terms of the perturbative expansion have to be taken into account until at some point the evolution becomes non-perturbative.

6.2.3 Numerical tests of the first-order expansion

We test the performance of Eq. (6.20) by following the orbits of three tracers in our benchmark Kepler potential. An example of the performance of Eq. (6.20) in a more general potential is shown in Appendix 6.9. The three particles are initially at a distance of 5 kpc from the host's centre, but have different initial energies and angular momenta.

Since negative energies correspond to gravitationally bound particles, we define $\mathscr{E} = -E$ as our energy variable of reference. In a Kepler potential, the radial period and the azimuthal period coincide and can, in terms of the energy, be written as

$$P(E,L) \equiv P(\mathscr{E}) = 2\pi \frac{GM}{\sqrt{2\mathscr{E}^3}} \tag{6.21}$$

Eq. (6.20) implies that the amplitude with which J_r oscillates around the dynamical invariant $J_{r'}$ is directly related to the orbital period. Given that the first order correction to $J_{r'}$ has to be small compared to the value of $J_{r'}$ itself for the linear approximation to be valid, we expect from Eq. (6.21 that the first order Taylor expansion will become progressively less accurate as $\mathscr{E} \to 0$. This region is known as the "fringe" of a self-gravitating system.

The initial energies and angular momenta of the three test particles are

$$(\mathscr{E}[\mathrm{km^2 s^{-2}}], \mathrm{L}[\mathrm{kpc\,km\,s^{-1}}]) = \{(50,33),(20,55),(2,15)\}. \tag{6.22}$$

The first two tracers initially have similar radial actions but different energies. The third tracer has a much larger initial radial action and represents a particle in the "fringe".

The upper subpanels of Fig. 6.2 show the time evolution of the radial actions of the three particles. Black lines in each figure denote an exact calculation of the radial action following Eq. (6.4). Red-dotted and blue-dashed lines both refer to the linear approximation for the radial action given in Eq. (6.20), using expression (6.21) to calculate the radial period. The difference between the latter two cases lies in the way $\dot{R}/R$ is estimated. The scale factor used to obtain the red line is calculated directly from Eq. (6.9), using a KDK leapfrog algorithm (see Appendix 6.7) to solve the differential equation. The scale factor's evolution is thus "coupled" to the particle's phase space trajectory. The dashed blue lines refer to an approximate analytic solution to the scale factor that can be calculated for scale-free potentials (see Peñarrubia (2013)). In this approximation the equation of motion is "de-coupled" from Eq. (6.9).

In the three lower subpanels of Fig. 6.2 we show $\chi = \Delta/J_r$, where $\Delta \equiv (\dot{R}/R)\,(\mathbf{r}\cdot\mathbf{v})\,P/(2\pi)$. Notice that χ quantifies the fractional size of the first order correction relative to the radial action, and can be positive or negative depending on the orbital phase.

Comparison of the top and middle panels of Fig. 6.2 shows that the orbital period is shorter in the top panel, which is in line with our expectation from Eq. (6.21). Furthermore, the results confirm our expectation that the amplitude of the first order

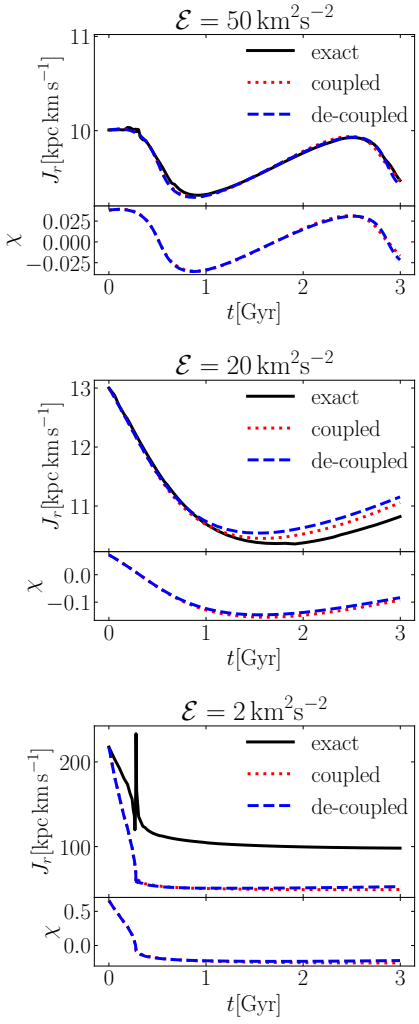

Figure 6.2. The upper subpanels show the time evolution of the radial action of three different particles orbiting in the benchmark Kepler potential. Initial energies are shown at the top of each plot. Black solid lines denote a direct measurement of the radial action from energies and angular momenta according to Eq. (6.4). Red dotted lines and blue dashed lines both correspond to the linear approximation given by Eq. (6.20), using two different estimates for $\dot{R}/R$. "Coupled" and "de-coupled" refer to an exact (dotted red) or approximate (dashed blue) way of solving Eq. (6.9). Angular momenta are chosen such that the initial radial actions are $J_r \sim 10\,\mathrm{kpc\,km\,s^{-1}}$ in the top and middle panel, and $J_r \sim 200\,\mathrm{kpc\,km\,s^{-1}}$ in the bottom panel. The lower subpanels show the ratio χ between the linear approximation and the exact value of the radial action as a function of time.

correction increases with the energy at a fixed invariant action. We find that for the most bound particle (left panel), the first order Taylor expansion is an excellent fit to the measured radial action. We furthermore find that there is almost no difference between the red and blue lines, which indicates that the "de-coupled" approximation to the scale factor provides a good approximation to its true value. The success of the Taylor expansion here is reflected in χ as well, which peaks at ~ 2.5 per cent when the magnitude of the radial velocity is maximal. As the period increases, the agreement between the measured radial action and the prediction from Eq. (6.20) starts to deteriorate after roughly 1 Gyr of integration (notice that in the middle panel $|\chi|$ has risen to ~ 10 per cent). Nonetheless, the evolution of the radial action is still captured well by this approximation. We furthermore find that the numerical solution to Eq.(6.9) now provides a slightly more accurate result than the approximate analytic solution for $\dot{R}/R$.

The bottom panel of Fig. 6.2 shows the evolution of the radial action of the third tracer. That this particle is in the "fringe" of the potential is evident by χ reaching values as large as 70 per cent. Due to that, we find that the dynamical invariant J_r' is not properly defined by the linear approximation – J_r changes impulsively. Note that the orbital period of this particle is much longer than the total simulation time, meaning we only resolve a fraction of an orbital revolution. Within this time interval, the radial action undergoes a strongly non-linear fluctuation when the particle passes its orbital pericentre and then quickly settles to a value which is roughly half of its initial value. The linear approximation completely fails to capture this behaviour.

6.3 Diffusion formalism for radial action distributions

Based on the first order Taylor expansion for time-dependent radial actions derived in the previous section, we now develop a diffusion formalism similar to the one in Peñarrubia (2015) to analytically describe the evolution of radial action distributions in time-dependent spherical potentials. We furthermore show how this diffusion formalism can be used to differentiate regimes in integral of motion space in which the evolution of radial action distributions is linear from regimes in which it is not.

6.3.1 The diffusion equation for radial action distributions

To derive a diffusion equation for radial actions, we closely follow the formalism presented in Peñarrubia (2015) based on Einstein (1905).

Let us define $\phi(\Delta|J_{r'})d\Delta$ as the conditional probability that a particle with the dynamically invariant action $J_{r'}$ will change its action by an amount Δ in the range $d\Delta$ during the time interval $(t, t+\tau')$; ϕ is normalized such that

$$\int \phi(\Delta|J_{r'})d\Delta = 1. \tag{6.23}$$

Subsequently, we define $p(J_r, t_0+\tau'|J_{r'}, t_0)dJ_r$ as the probability that a particle with action $J_{r'}$ at t_0 will have an action in the interval $[J_r, J_r + dJ_r]$ at a later time $t = t_0 + \tau'$.

As with the equivalent in energy space presented in Peñarrubia (2015), these two functions obey Einstein's master equation

$$p(J_r, t_0+\tau'|J_{r'}, t_0)dJ_r = \\ dJ_r \int p(J_r-\Delta, t_0|J_{r'}, t_0)\phi(\Delta|J_{r'})d\Delta. \tag{6.24}$$

Now if we are in the adiabatic limit where Eq. (6.20) applies, or equivalently $\chi = \Delta/J_r \ll 1$, we can expand Eq. (6.24) in Δ to find

$$p(J_r, t_0+\tau'|J_{r'}, t_0) \approx p(J_r, t_0|J_{r'}, t_0) \\ - \left.\frac{\partial p}{\partial J_r}\right|_{\Delta=0} \int \phi(\Delta|J_{r'})\Delta d\Delta \\ + \frac{1}{2}\left.\frac{\partial^2 p}{\partial J_r^2}\right|_{\Delta=0} \int \phi(\Delta|J_{r'})\Delta^2 d\Delta. \tag{6.25}$$

Since we can also expand the lhs of Eq. (6.25) around $t = t_0$ to first order

$$p(J_r, t_0+\tau'|J_{r'}, t_0) \approx p(J_r, t_0|J_{r'}, t_0) + \frac{\partial p}{\partial t}\tau',$$

we then have

$$\frac{\partial p}{\partial t} = C \left.\frac{\partial p}{\partial J_r}\right|_{\Delta=0} + D \left.\frac{\partial^2 p}{\partial J_r^2}\right|_{\Delta=0}, \tag{6.26}$$

where we define the drift coefficient

$$C(J_{r'}, t) = -\frac{1}{\tau'}\int \phi(\Delta|J_{r'})\Delta d\Delta \tag{6.27}$$

and the diffusion coefficient

$$D(J_{r'}, t) = \frac{1}{2\tau'}\int \phi(\Delta|J_{r'})\Delta^2 d\Delta. \tag{6.28}$$

Analogous to the discussion in Peñarrubia (2015), the initial condition to solve Eq. (6.26) is $p(J_r, t_0|J_{r'}, t_0) = \delta(J_r - J_{r'})$ and the general solution to this equation is a Green function

$$p(J_r, t|J_{r'}, t_0) = \frac{1}{\sqrt{4\pi\tilde{D}(J_{r'}, t)}} \\ \times \exp\left\{-\frac{\left[J_r - J_{r'} + \tilde{C}(J_{r'}, t)\right]^2}{4\tilde{D}(J_{r'}, t)}\right\}, \tag{6.29}$$

the properties of which are discussed in detail in Peñarrubia (2015). In Eq. (6.29) we have implicitly defined scaled drift and diffusion coefficients $\tilde{C} = C\tau'$, $\tilde{D} = D\tau'$. Note that the fact that $\tilde{C}$ and $\tilde{D}$ are independent of time implies that there is no divergence for short transition times in Eq. (6.29).

In the perturbative regime, we can use the transition probability given by Eq. (6.29) to calculate the radial action distribution, $N(J_r,t)$, from the invariant action distribution. To calculate the invariant distribution $N(J_{r'})$ from the initial distribution $N(J_r,t_0)$, we also need $p(J_{r'}|J_r,t_0)$. The derivation of this probability works analogous to the calculation above, and one obtains:

$$p(J_{r'}|J_r,t_0) = \frac{1}{\sqrt{4\pi\tilde{D}(J_r,t_0)}} \times \exp\left\{-\frac{\left[J_{r'}-J_r-\tilde{C}(J_r,t_0)\right]^2}{4\tilde{D}(J_r,t_0)}\right\}. \tag{6.30}$$

We can then calculate $N(J_{r'})$ from $N(J_r,t_0)$ as

$$N(J_{r'}) = \int dJ_r\, p(J_{r'}|J_r,t_0) N(J_r,t_0). \tag{6.31}$$

Contrary to the time average of a particle's energy in a time-dependent potential, the time average of the radial action is constant according to Eq. (6.20), and thus we do not need an extra convolution analogous to Eq. 28 of Peñarrubia (2015). The radial action distribution at a time t is thus obtained from

$$N(J_r,t) = \int dJ_{r'}\, p(J_r,t|J_{r'},t_0) N(J_{r'}), \tag{6.32}$$

where the transition probability is the one defined in Eq. (6.29).

6.3.2 Drift and diffusion coefficients

Similar to Section 2.3 of Peñarrubia (2015) we now briefly discuss how to calculate the drift and diffusion coefficients defined in Eqs. (6.27) and (6.28). To this end, we define the microcanonical distribution function $w(J_r,E)$ depending on both energy and radial action,

$$w(J_r,E) = \int \delta\left(J_r - X\right)\delta(E-H)\mathbf{d^3r}\,\mathbf{d^3v}, \tag{6.33}$$

where

$$X = \int_{r_{\rm peri}}^{r_{\rm apo}} dr \sqrt{2(E-\Phi(r)) - \frac{L^2}{r^2}}. \tag{6.34}$$

We then define the drift and diffusion coefficients as microcanonical averages over the particle distribution as follows

$$\tilde{C}(J_r,E,t) = -\frac{1}{w} \times \int \frac{\dot{R}}{R}(\mathbf{r}\cdot\mathbf{v})\frac{P(H,X,t)}{2\pi}\delta(J_r-X)\delta(E-H)\mathbf{d^3r}\,\mathbf{d^3v} \tag{6.35}$$

$$\tilde{D}(J_r,E,t) = \frac{1}{2w} \times \int \left(\frac{\dot{R}}{R}(\mathbf{r}\cdot\mathbf{v})\frac{P(H,X,t)}{2\pi}\right)^2 \delta(J_r-X)\delta(E-H)\mathbf{d^3r}\,\mathbf{d^3v}. \tag{6.36}$$

Notice that here we have written the period as a function of both the energy and the radial action, which is possible in general spherical potentials where the angular momentum is a function of energy and radial action. To obtain drift and diffusion coefficients that depend only on the radial action, however, we have to integrate over the energies

$$\tilde{C}(J_r,t) = \frac{1}{w(J_r)} \int dE\, w(J_r,E)\, \tilde{C}(J_r,E,t) \tag{6.37}$$

$$\tilde{D}(J_r,t) = \frac{1}{w(J_r)} \int dE\, w(J_r,E)\, \tilde{D}(J_r,E,t), \tag{6.38}$$

where

$$w(J_r) = \int dE\, w(J_r,E) \tag{6.39}$$

is the microcanonical distribution function depending only on the radial action.

6.3.3 Adiabatic, diffusive, and impulsive evolution

We can use the diffusion coefficient defined in Eq. (6.36) to characterize the rate at which a gravitational potential evolves. As we argued in Section 6.2.2, the regime in which the evolution of radial actions becomes non-perturbative corresponds to the regime in which the evolution of the gravitational potential is impulsive. Mathematically, this means that in Eq. (6.20), $\Delta = |J_r - J_{r'}| \sim J_r$. For a given time-dependent potential, there is always some part of integral of motion space in which this condition is fulfilled on average. Whether the evolution of a dynamical system is adiabatic, diffusive, or impulsive then depends on how populated this part of integral of motion space is. Since $\tilde{D} \sim 1/2\Delta^2$, a possible way to estimate whether distributions of particles occupying a particular integral of motion space volume evolve either adiabatically to diffusively (linearly) or impulsively (non-linearly) is by the ratio $\sqrt{\tilde{D}}/J_r$.

This is particularly useful for phase-mixed particle ensembles. In such cases $\tilde{C} = 0$ and thus the drift cannot be used to estimate whether the evolution of such ensembles is diffusive or impulsive. The diffusion coefficient, on the other hand, can be calculated analytically provided one can make the simplifying assumption that in Eqs. (6.35) and (6.36) the radial period is approximately constant within a small region in $J_r - \mathscr{E}$ space. In this case,

$$\tilde{C}(J_r,E,t) \approx -\frac{\dot{R}}{R} \frac{P(E,L)}{2\pi} \langle (\mathbf{r}\cdot\mathbf{v}) \rangle \tag{6.40}$$

$$\tilde{D}(J_r,E,t) \approx \frac{1}{2} \left(\frac{\dot{R}}{R}\right)^2 \left(\frac{P(E,L)}{2\pi}\right)^2 \left\langle (\mathbf{r}\cdot\mathbf{v})^2 \right\rangle \tag{6.41}$$

where $\langle . \rangle$ denotes an ensemble average. For time-dependent spherical potentials, all factors appearing in Eq. (6.41) can be calculated if $\Phi(r)$ is known. Appendix 6.9 discusses how $\dot{R}/R$ can be calculated and $\langle \mathbf{r}\cdot\mathbf{v} \rangle$ can be obtained following Peñarrubia (2019):

$$\left\langle (\mathbf{r}\cdot\mathbf{v})^2 \right\rangle = \frac{1}{2P(E,L)} \int_{r_{\rm peri}}^{r_{\rm apo}} dr\, r^2 \sqrt{2(E-\Phi(r)) - \frac{L^2}{r^2}}. \tag{6.42}$$

In general potentials, equation 6.41 has to be evaluated numerically. From the value of $\sqrt{\tilde{D}}/J_r$ we can then estimate whether the dynamical evolution of tracers with a given set of integrals of motion (E, J_r) is likely to be diffusive or impulsive for a given evolving potential.

A special case is again the Kepler potential in which the right hand side of Eq. (6.41) can be evaluated analytically. We therefore illustrate the above point on our benchmark potential. Peñarrubia (2013) shows that in a scale-free potential with a time-dependent force

$$F(r,t) = -\mu(t) r^n, \tag{6.43}$$

a good analytic approximation to the scale factor is

$$R(t) \approx \left(\frac{\mu(t)}{\mu(0)} \right)^{-1/(n+3)}. \tag{6.44}$$

and for a phase-mixed distribution of particles orbiting in a time-dependent Kepler potential ($n = -2$), Eq. (6.41) yields

$$\tilde{D}(J_r, E, t) = \frac{1}{4} \left(\frac{\dot{R}}{R} \right)^2 (GM)^2 \frac{GMJ_r - \sqrt{2\mathscr{E}} J_r^2}{(2\mathscr{E})^{\frac{7}{2}}}. \tag{6.45}$$

We identify the linear (adiabatic) region in $\mathscr{E} - J_r$ space with the region corresponding to all combinations of energy and radial action for which $\sqrt{\tilde{D}} \ll J_r$. Based on the ratio between the (theoretical) diffusion coefficient calculated according to Eq. (6.45) and the radial action we define three different areas in the $J_r - \mathscr{E}$ space, characterized by $\sqrt{\tilde{D}} < 0.1 J_r$, $\sqrt{\tilde{D}} < J_r$ and $\sqrt{\tilde{D}} > J_r$.

Fig. 6.3 shows a plot of $\mathscr{E} - J_r$ space where the different correspond to the above-defined regimes in our benchmark potential at $t = 0$. For reference, we show the radius of a circular orbit with energy $\mathscr{E}$ on the right y-axis.

The grey area is dictated by Eq. (6.4) and its border corresponds to the line at which $L = 0$. Any allowed orbits are located to the left in one of the coloured areas.

Particles inhabiting the green area have an average diffusion coefficient which is smaller than one per cent of the squared radial action. In this (linear, i.e. adiabatic to diffusive) area, the diffusion formalism is applicable.

For particles in the cyan area, the average diffusion coefficient is larger than one per cent of the squared radial action, yet smaller than the squared radial action itself. This represents a “transition" regime between the diffusive, perturbative regime and the impulsive “fringe".

The red area is the “fringe" of the potential. Here, the evolution of radial actions is highly non-linear and the diffusion formalism is not applicable.

The exact locations of the areas vary with time, as the mass of the Kepler potential directly impacts all the three boundaries shown in this picture. Furthermore, as not all relevant distributions of tracer particles are phase-mixed and in virial equilibrium, using Eq. (6.45) along with $\tilde{C} = 0$ is not always a valid approximation. Nonetheless, the areas defined in Fig. 6.3 give a good indication as to which radial action - energy combinations imply a linear, diffusive evolution and which do not.

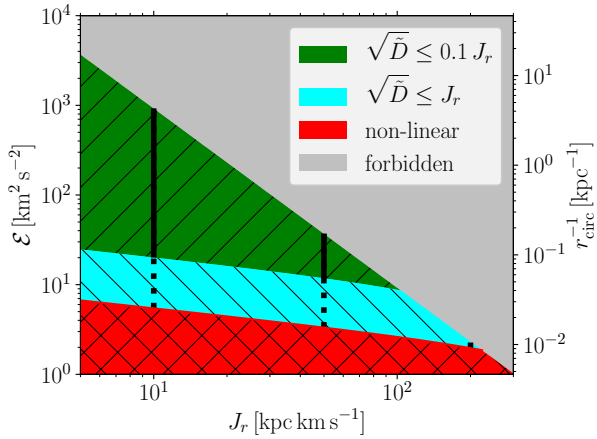

Figure 6.3. Areas in $\mathscr{E} - J_r$ space in which the expected dynamical evolution of tracer distributions is distinctly different, calculated for our benchmark time-dependent Kepler potential at $t = 0$. The grey area is not populated, as no particles can exist there according to Eq. (6.4). The green area denotes the perturbative "linear" regime in which $\sqrt{\tilde{D}} \ll 0.1 J_r$. The cyan area denotes a "transition" regime between the "linear" regime and the "fringe". Here, $0.1 < \sqrt{\tilde{D}}/J_r < 1$. In red, we show the non-linear ("fringe") part of the phase space. Here, the first order variation can be bigger than the action itself, meaning that the invariant action defined in Eq. (6.20) is not a dynamical invariant and the diffusion formalism is not applicable. On the right y-axis we show the radius of a circular orbit of a particle with energy $\mathscr{E}$ as an indication of the scales involved.

In Section 6.4 we validate the predictions of Fig. 6.3. We then apply the above analysis to investigate whether the flattening of the tail in Fig. 6.1 is a linear or a non-linear phenomenon. Given that the tail is comprised mainly of particles inhabiting the "fringe" area in Fig. 6.3, our expectation is that the effect is a non-linear one. To investigate this further, we construct five restricted simulations and then compare their evolution to the evolution predicted by the diffusion formalism.

6.4 Tests on a time-dependent Kepler potential

In this section, we test the diffusion formalism developed in Sections 6.2 and 6.3 in a series of numerical simulations performed using AREPO. The simulations follow the evolution of five different initial tracer particle populations in our benchmark time-dependent Kepler potential. We start by describing the initial conditions of each run. We also test the analysis outlined in Section 6.3.3, using it to forecast whether the diffusion formalism will adequately describe the evolution of the radial action distribution of the tracers for each simulation. Then we compare the evolved radial action distributions

to the predictions of the diffusion formalism and evaluate if our prior assessment of whether the evolution of each distribution will be diffusive or impulsive was accurate. Finally, we use our results to discuss whether the net drift observed in Fig. 6.1 is a linear or a non-linear effect.

6.4.1 Initial Conditions

We set up five different Gaussian distributions in radial action to follow their evolution in our benchmark Kepler potential. The black lines in Fig. 6.3 indicate the mean of the different distributions, highlighting the respective values of 10,50 and 200 $\mathrm{kpc\,km\,s^{-1}}$. The solid black lines indicate two initial configurations in which we confine the particles to the integral of motion space area in which the evolution of the radial action is expected to be linear. The dotted lines indicate three simulations that are confined to the “linear" and the “transition" regimes. The simulation whose central value is 200 $\mathrm{kpc\,km\,s^{-1}}$ is a special case, as here only a tiny area of non-impulsive integral of motion space is available which is entirely part of the “transition" area. As a consequence, the average diffusion coefficient is rather large. As the occupied integral of motion space is in close vicinity to the “fringe", we anticipate that the diffusion formalism may fail for this simulation. In table 6.1 we show the parameters defining the initial conditions.

Fig. 6.4 shows a comparison of the initial coarse-grained distribution functions in $\mathscr{E} - J_r$ space between the “10-linear" case on the top panel and the “10-transition" case on the bottom panel. The Gaussian shape in the direction of J_r is apparent in both cases and reflects the input parameters given in Table 6.1. The distribution in energy is largely determined by the additional cuts we impose (see Table 6.1). In particular, the lower limit on the pericentre radius effectively introduces a lower limit on the angular momentum. Hence, some theoretically possible combinations of $\mathscr{E}$ and J_r which correspond to low angular momenta are forbidden. The energy cuts stated in table 6.1 are clearly reflected in the high energy end of the distribution functions.

The distribution functions presented in Fig. 6.4 are coarse-grained versions of Eq. (6.33), which we calculate as

$$\begin{aligned}\tilde{N}_{i,j}(J_{r,i},E_j) = \sum_{k=0}^{N_{tot}} &[\theta(\log(E_k)-\log(E_{j,min}))\\ &\times\theta(\log(E_{j,max})-\log(E_k))\theta(J_{r,k}-J_{r,i,min})\\ &\times\theta(J_{r,i,max}-J_{r,k})\big]\,.\end{aligned} \tag{6.46}$$

The coarse-grained versions of the drift and diffusion coefficients defined in Eqs. (6.40) and (6.41) are given by

$$\tilde{C}_{i,j}(J_{r,i},E_j) = \frac{1}{N_{i,j}} \sum_{k=0}^{N_{tot}}\Bigg|_{\substack{J_{r,i,min}<J_{r,k}<J_{r,i,max}\\ E_{j,min}<E_k<E_{j,max}}} \left[-(\mathbf{r}_k\cdot\mathbf{v}_k)\frac{\dot{R}_k}{R_k}\frac{P(E_k)}{2\pi}\right] \tag{6.47}$$

Distribution	central J_r [kpc km s^{-1}]	σ_{J_r} [kpc km s^{-1}]	$\mathscr{E}_{\rm min}$ [km s^{-1}]	$r_{\rm peri,min}$ [kpc]	$r_{\rm ini}$ [kpc]	$L_{\rm ini}$ [kpc km s^{-1}]
10-transition	10	2	5.44	0.1	0.1-15	5-150
50-transition	50	3	3.32	-	0-15	5-150
200-transition	200	12	1.96	-	0-15	5-150
10-linear	10	2	20.06	0.1	0.1-15	2-50
50-linear	50	3	11.78	0.1	0.1-15	5-150

Table 6.1. The relevant parameters of the five initial particle distributions All distributions consist of 20000 particles which are sampled from a Gaussian distribution in radial action. "Linear" refers to distributions which are confined to the green area of Fig. 6.3 (i.e. the area in which the linear approximation of J_r is accurate). "Transition" refers to distributions which are confined to the green and cyan area, implying that the linear correction can be of the same order as the action itself for some particles. In column 2 we show the mean of the Gaussian distribution in each of the cases, whereas the spread of the distribution is given in column 3. The energy limit which is applied to confine the particles to the integral of motion space areas specified above is shown in column 4. The last three columns show various additional cuts that have been imposed for numerical reasons. Column 5 shows minimum values for the pericenter radii, column 6 the range of initial radii at which we set up the tracers. The last column shows the initial range of angular momenta. Notice that in some cases, these collide with other cuts such as the pericenter radius cut. In fact, the reason for the "10-linear" cut to be different from the others is that with the imposed energy cut, no angular momenta larger than 50 kpc km s^{-1} *are allowed in the sampled range of radial actions.*

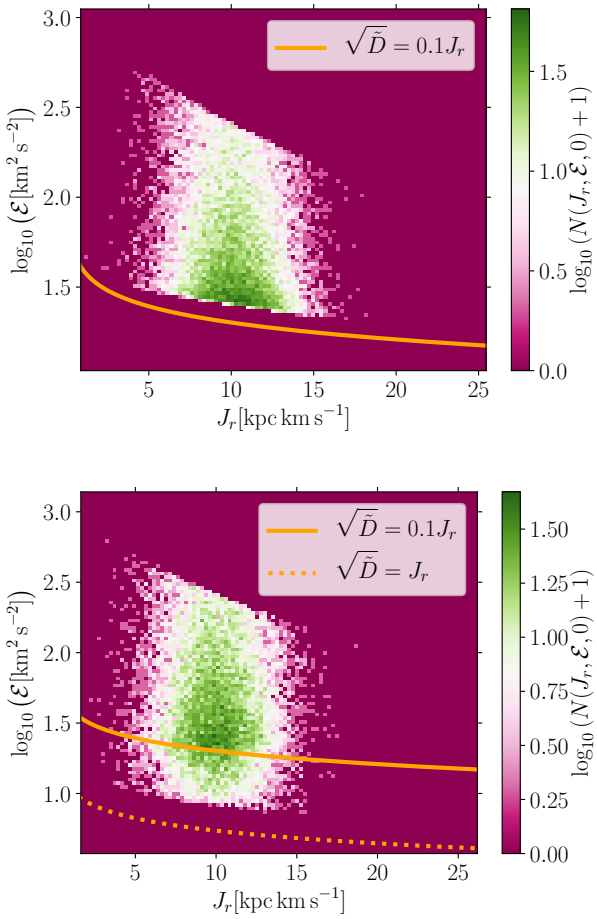

Figure 6.4. Comparison of the initial integral of motion space distributions in the 10-linear (top) and the 10-transition (bottom) cases. The lines defining the limits between different areas in Fig. 6.3 are shown as orange lines here. The solid orange line is the border between the "linear" and the "transition" regime, whereas the dashed orange line is the limit to the "fringe" regime. The logarithmic colour scale indicates the number of particles.

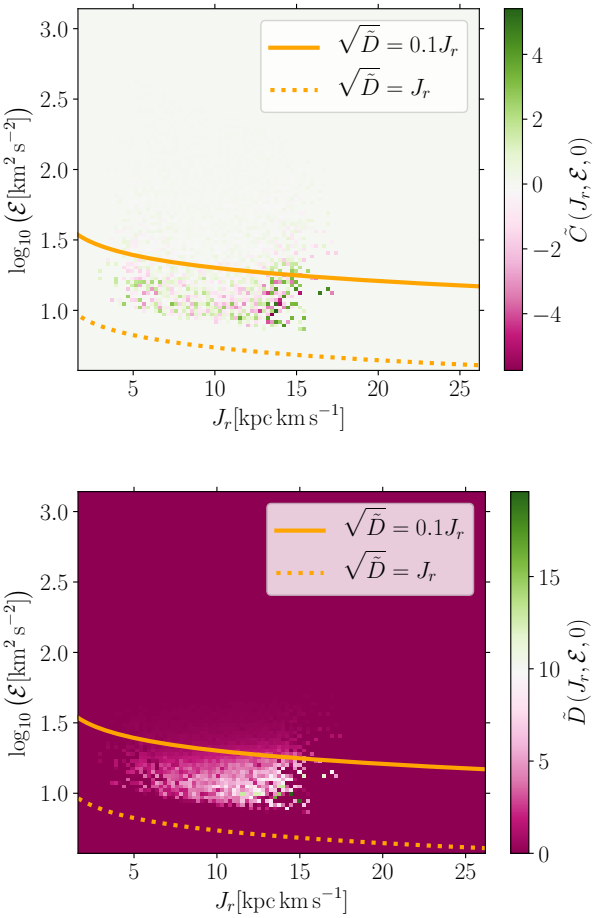

Figure 6.5. Initial drift (top) and diffusion (bottom) coefficients in the "10-transition" case as functions of energy and radial action. The colour scales indicate the drift (diffusion) coefficient in units of $\mathrm{kpc\,km\,s^{-1}}$ *(*$\left(\mathrm{kpc\,km\,s^{-1}}\right)^2$*) in the left (right) panel. The orange lines are the same as in Fig. 6.4. Both the drift and diffusion coefficients increase in magnitude for larger energies and radial actions. In the "transition" area, the drift coefficients can reach values which are of the same order of magnitude as the radial action at which they are measured.*

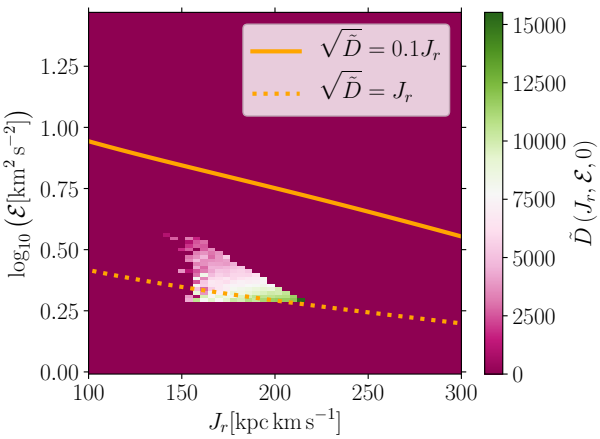

Figure 6.6. Diffusion coefficients for the "200-transition" distribution (see Table 6.1) at $t = 0$ (units as in Fig. 6.5). Only a narrow area in integral of motion space is populated, which is very close to the border of the "fringe" (dotted orange curve here). The upper limit of the populated integral of motion space area in the $\mathcal{E}$-direction is given by the $L = 0$ curve. The measured diffusion coefficients are very large over the entire populated integral of motion space area.

and

$$\tilde{D}_{i,j}(J_{r,i}, E_j) = \frac{1}{2N_{i,j}} \sum_{k=0}^{N_{tot}} \Bigg|_{E_{j,min}<E_k<E_{j,max}}^{J_{r,i,min}<J_{r,k}<J_{r,i,max}} \left[(\mathbf{r}_k \cdot \mathbf{v}_k) \frac{\dot{R}_k}{R_k} \frac{P(E_k)}{2\pi} \right]^2 \tag{6.48}$$

In Fig. 6.5 we show the coarse-grained drift and diffusion coefficients of the "10-transition" distribution in $\mathcal{E} - J_r$ space at $t = 0$. The drift coefficients (in units of $\mathrm{kpc\,km\,s}^{-1}$) are shown in the top panel, while the diffusion coefficients (in units of $\left(\mathrm{kpc\,km\,s}^{-1}\right)^2$) are shown on the bottom panel. Both the drift and the diffusion coefficients show the same trend with energy and radial action. The largest drift and diffusion coefficients are well within the "transition" area in phase space. Moreover, the largest drift coefficients are quite substantial in magnitude, up to the same order of magnitude as J_r itself. Furthermore, the average drift and diffusion coefficients increase for larger values of $\mathcal{E}$ and J_r. Since drift and diffusion coefficients are larger in the "transition" area than in the "linear" area, the analysis of Section 6.3.3 suggests that the diffusion formalism will perform better at predicting the evolution of the "10-linear" distribution than that of the "10-transition" distribution.

Fig. 6.6 shows the initial diffusion coefficients of the "200-transition" distribution. Since the populated integral of motion space area lies in close vicinity to the "fringe", the measured diffusion coefficients are very large. Given the magnitude of the averaged

diffusion coefficients, we expect the evolution of the “200-transition" radial action distribution to be impulsive and non-perturbative. We thus do not expect our diffusion formalism to give an accurate prediction. Across all five simulations, we expect the predictions of the diffusion formalism to deteriorate when a larger fraction of particles occupies the increasingly non-linear “transition" area in integral of motion space.

6.4.2 Evolution of the “10-linear" distribution

In this Section we compare the evolved “10-linear" distribution to predictions from the diffusion formalism. We calculate the drift and diffusion coefficients in two different ways. In the first (numerical) method, we calculate them directly from the particle's phase space coordinates using Eqs. (6.47) and (6.48). To obtain coefficients that depend solely on radial action, we marginalize over the energy as

$$\tilde{C}_i(J_{r,i}) = \frac{1}{N_i}\sum_{j=0}^{N_{bin}} \tilde{C}_{i,j}(J_{r,i},E_j) \tag{6.49}$$

$$\tilde{D}_i(J_{r,i}) = \frac{1}{N_i}\sum_{j=0}^{N_{bin}} \tilde{D}_{i,j}(J_{r,i},E_j), \tag{6.50}$$

where we have used

$$N_i = \sum_{j=0}^{N_{bin}} N_{i,j}. \tag{6.51}$$

In the second (analytic) method, we assume the particle distribution to be phase-mixed, set the drift coefficient to zero, and calculate the diffusion coefficient using Eq. (6.45).

The upper panel of Fig. 6.7 shows the invariant distribution $N(J_{r'})$ obtained at the start of the “10-linear" simulation, while the lower panel shows the drift coefficient $\tilde{C}(J_r,t)$ that was used to calculate $N(J'_r)$ using Eq. (6.31). We combine sparsely populated adjacent bins until the combined bin contains more than 50 particles. The center of the new bin is calculated as a weighted mean of the centres of the combined bins, using the particle number as a weight. The grey area marks the resulting radial action range in which bins are empty, suggesting that the drift obtained in this area is not a robust measurement due to insufficient particle sampling. The prediction of the diffusion formalism using drift and diffusion coefficients calculated according to the first (second) method is shown as a green dashed (red dotted) line. The “measured" distribution (black line) is obtained by calculating J'_r for each particle individually and then calculating a histogram in J'_r. Comparing the “measured" distribution with the results of the diffusion formalism, we find that they all coincide remarkably well. The agreement between the two results of the diffusion formalism reveals that our initial sampling algorithm created a fully phase-mixed distribution in radial action with no net drift. This is confirmed on the lower panel, where we show $\tilde{C}(J_r,0)$ and find it to be consistent with zero over the full radial action range.

In Fig. 6.8 we $N(J_r,t)$ at $t = 300\,\mathrm{Myr}$ on the top panel and at $t = 3\,\mathrm{Gyr}$ on the bottom panel. The colours used are the same as in Fig. 6.7. In the lower part of each panel we show the drift coefficients as functions of the invariant action, $\tilde{C}(J_{r'},t)$. Black

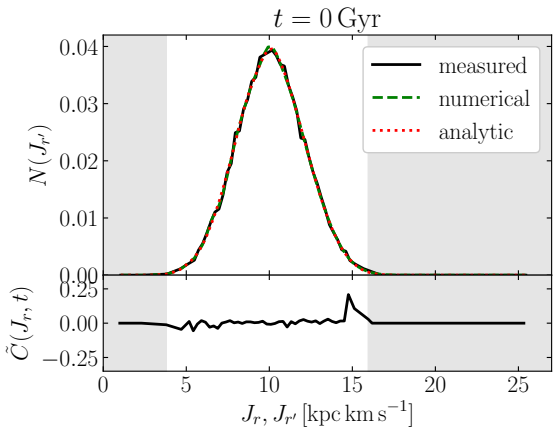

Figure 6.7. Invariant distribution $N(J_{r'})$ in the upper panel, and the drift coefficient $\tilde{C}(J_r,t)$ in the lower panel for the "10-linear" simulation at $t = 0$ (see Table 6.1). The black solid line is the measured distribution, obtained by calculating $J_{r'}$ for each particle and binning the results. The green dashed line is the result of the convolution (Eq. 6.31) using drift and diffusion coefficients obtained numerically from the data. The drift coefficient shown in the lower panel is the one depending on the instantaneous action J_r, i.e., the one that was used to do the convolution resulting in the green dashed line of the upper panel. The red dotted line in the upper panel shows the result of Eq. (6.31) using the analytic diffusion coefficient from Eq. (6.45). In order to avoid sampling noise in the data we combine bins with less than 50 particles in the tails of the distribution. In the grayed-in areas $\tilde{C}(J_r)$ is too poorly sampled.

solid lines denote the drift at the initial time, whereas blue dashed lines indicate the drift at the time displayed in the top panels. We combine bins as in Fig. 6.7.

Although at $t = 300\,\mathrm{Myr}$, $\tilde{C}(J'_r,t)$ deviates strongly from zero at large radial actions, by $t = 3\,\mathrm{Gyr}$, it has decreased considerable over the entire range of invariant actions. The impact of this evolution in the drift is evident when comparing between the two results of the diffusion formalism. If numerically calculated drift and diffusion coefficients are used, the result of the diffusion formalism is in perfect agreement with the 'measured" distribution at both times. Using the "analytic" method to calculate the diffusion coefficients results in a worse match at $t = 300\,\mathrm{Myr}$, as this method assumes $N(J_{r'})$ to be phase-mixed.

Overall, we conclude that the diffusion formalism outlined in Section 6.3 provides a remarkably good description of the adiabatic time evolution of $N(J_r,t)$ in the "10-linear" simulation. In the next section, we will investigate if and how the accuracy of this formalism deteriorates when analyzing the more impulsive evolution of the other distributions from Table 6.1.

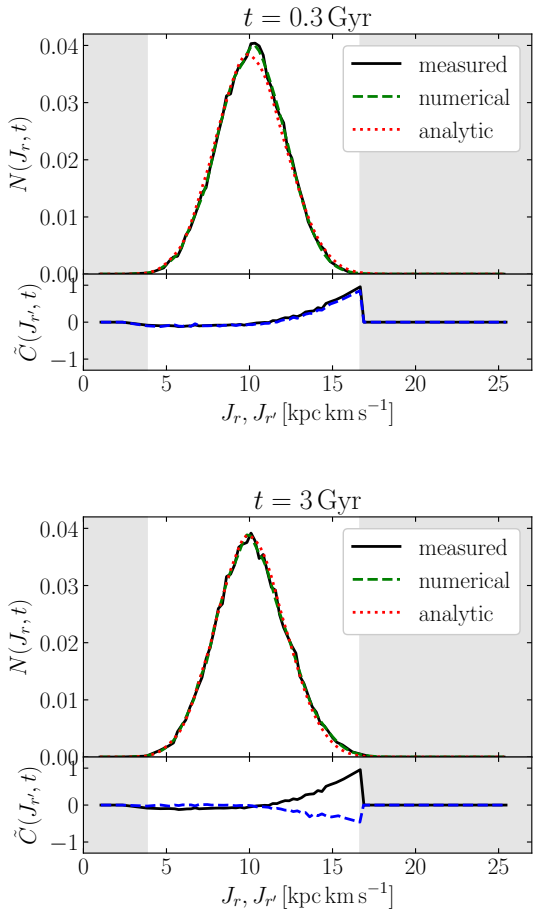

Figure 6.8. $N(J_r,t)$ *at* $t = 0.3\,\mathrm{Gyr}$ *on the top panel and at* $t = 3\,\mathrm{Gyr}$ *on the bottom panel for the "10-linear" simulation (see Table 6.1). As in Fig. 6.7, the black line in the upper panels is the measured distribution, the green dashed line is obtained using the diffusion formalism with coefficients calculated from the simulation directly and the red dotted line is calculated using the analytic diffusion approach assuming phase mixing (no drift). In the lower subpanels, we show the drift coefficient used in the convolution (Eq. 6.32) to obtain the green dashed line in the upper subpanels. The black line in the lower panels is the drift coefficient* $\tilde{C}(J_{r'},t)$ *at* $t = 0$*, whereas the blue dashed line is the one calculated at time t (shown in the upper legend). We combine bins as in Fig. 6.7. Initially, the drift coefficient is non-zero, especially towards the tail of the distribution to the right, but it decreases considerably with time. In line with this, we find that the analytic approximation is not a perfect match at initial times, but improves significantly as phase mixing in* $N(J'_r)$ *progresses.*

6.4.3 The diffusion formalism in different regimes

In this section we present results obtained when applying the diffusion formalism to the “10-transition", “50-linear" and “50-transition" distributions from Table 6.1. For each of those cases, we show a comparison between the directly measured invariant distribution and the result of Eq. (6.31) at $t = 0$. We furthermore show a measurement of $N(J_r,t)$ after 3 Gyr of simulation time, as well as the result of the diffusion formalism using Eq. (6.32) and the initial radial action distribution $N(J_r,0)$. The drift and diffusion coefficients used in Eqs. (6.31) and (6.32) are calculated numerically from the phase space coordinates of the tracers.

Fig. 6.9 shows the invariant distributions at time $t = 0$ in the “10-transition" (top), “50-linear" (middle) and “50-transition" (bottom) cases in the left column. In solid black lines, we show the “measured" distribution obtained by calculating the invariant action using Eq. (6.20) for each particle. The green dashed lines are the “convolved" distributions obtained from Eq. (6.31). The match between the measured and the calculated distributions is accurate in the “10-transition" simulation (top left) and only marginally worse in the “50-linear" case (mid left). However, in the “50-transition" case (bottom left), the match between the result of Eq. (6.31) and the direct measurement of $N(J_r')$ is substantially worse. In particular, the result of the convolution is does not resolve the peak in the measured invariant distribution.

The right column of Fig. 6.9 displays $N(J_r,t)$ after 3 Gyr in the “10-transition" (top), “50-linear" (middle) and “50-transition" (top) cases. Black solid lines are direct measurements, green dashed lines are the results of the convolutions (Eqs. 6.31 and 6.32) and red dotted lines show $N(J_r,t = 0)$ for comparison.

In the "10-transition" simulation, shown in the top right panel of Fig. 6.9, we find a very good agreement between the measured distribution and the result of the diffusion formalism. Furthermore, we find that there is only relatively little evolution in the shape of $N(J_r,t)$, with the most obvious effect being the formation of a tail towards smaller values of J_r.

In the middle right panel, we show the “50-linear" distribution at the end of the simulation. We find good agreement between the result of the diffusion formalism and the measured distribution, despite the rather substantial evolution with respect to $N(J_r,0)$. A slight mismatch can be observed in the tail of the distribution towards large values of J_r, where the convolution overpredicts the true measured distribution. Note that the diffusion formalism nonetheless provides a substantial improvement over the assumption that radial actions are invariant.

To explain why the evolution of the “10-transition" distribution appears to be better captured by the diffusion formalism than the “50-linear" distribution, we consult Fig. 6.3. While some particles in the “10-transition" simulation inhabit the cyan “transition" area, most particles are far within the “linear" regime (see bottom panel of Fig. 6.4). Furthermore, the linear energy range is much larger for $J_r = 10\,\mathrm{kpc\,km\,s^{-1}}$ than it is for $J_r = 50\,\mathrm{kpc\,km\,s^{-1}}$. In the former case particles are on average more bound and have shorter periods (see Eq. 6.21), and thus the evolution of $N(J_r,t)$ is closer to adiabatic and the diffusion formalism is more accurate.

In the bottom right panel of Fig. 6.9 we show the “50-transition" case. Fig. 6.3 suggests that the fraction of “transition" integral of motion space available to the

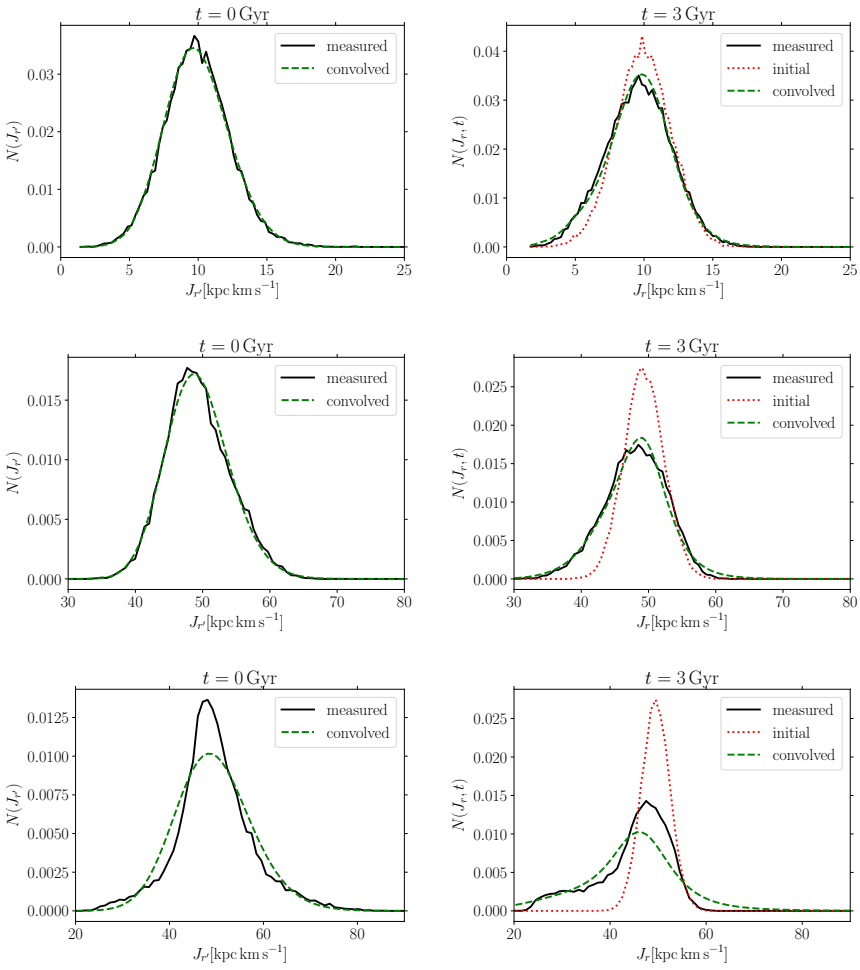

Figure 6.9. Invariant distributions (on the left) and radial action distributions after a simulation time of 3 Gyr *(on the right) of the "10-transition" simulation (top row), the "50-linear" simulation (middle row) and the "50-transition" simulation (bottom row); see Table 6.1. The green dashed lines marked "convolved" are the result of applying the diffusion formalism (Eqs. 6.31 and 6.32) using coarse-grained drift- and diffusion coefficients (see Eqs. 6.49 and 6.50). The black lines marked "measured" are direct measurements of the distributions. On the right panels, we also show $N(J_r, t = 0)$ as red dotted lines.*

particles is much larger here than in the “10-transition" case. The consequences of this can immediately be seen in the mismatch between the measured and convolved invariant distributions in the bottom left panel of Fig. 6.9. In the bottom right panel, we find that $N(J_r,t)$ has evolved substantially at the end of the simulation, with an extended tail towards lower actions being present in the final distribution. While the peak of the distribution remains around its initial value of $50\,\mathrm{kpc\,km\,s^{-1}}$, its mean shifts towards smaller radial actions and the resulting distribution is non-Gaussian. The diffusion formalism does not fully capture the evolution of the radial action distribution. Most notably, it resolves neither the peak of the distribution nor the tail at large J_r values. This is to some extent expected, given that the calculated invariant distribution already deviates significantly from the measured invariant distribution (bottom left panel of Fig. 6.9). Interestingly, the tail of $N(J_r,t)$ towards smaller radial actions is captured fairly well by the diffusion approximation, indicating that its formation is a linear effect. We note that the diffusion formalism still yields a considerable improvement over the assumption that $N(J_r,t)$ is invariant, yet the growing mismatch between the predictions of the formalism and the actual measurement is a strong hint that non-linear effects are becoming increasingly significant.

Finally, we note that the “50-transition" distribution drifts towards smaller radial actions. Since this drift is roughly captured by the diffusion formalism, it appears to be a linear effect. Given that the number of particles around $J_r = 50\,\mathrm{kpc\,km\,s^{-1}}$ grows with time in Fig. 6.1, a more substantial drift towards smaller actions must occur at initially larger actions, where the available integral of motion space becomes increasingly non-linear according to Fig. 6.3.

6.4.4 Evolution of the “200-transition" distribution

According to Fig. 6.3, there is no “linear" integral of motion space available for particles with $J_r = 200\,\mathrm{kpc\,km\,s^{-1}}$. We thus expect the evolution to be impulsive. Indeed, we find that the diffusion formalism fails and that the “invariant" distribution measured using Eq. (6.20) evolves strongly with time. To understand this better, we look at the time evolution of both $N(J_r,t)$ and $N(J_r',(t))$, where J_r' is defined by Eq. (6.20) – and cannot be considered a dynamical invariant. Fig. 6.10 shows $N(J_r,t)$ (lower panel) and $N(J_r')$ as calculated using the linear approximation (Eq. 6.20) (upper panel) at different times. Both evolve strongly with time. The fact that $N(J_r')$ is not time-invariant indicates that the evolution of the “200-transition" distribution is highly non-adiabatic (and clearly non-linear). However, the evolution of the “invariant" distribution explains some of the time evolution of $N(J_r,t)$ in the bottom panel of Fig. 6.10. While $N(J_r,t)$ is initially a Gaussian distribution set up as in Table 6.1, the initial “invariant" distribution is bimodal, with one peak around $50\,\mathrm{kpc\,km\,s^{-1}}$ and a second peak around $350\,\mathrm{kpc\,km\,s^{-1}}$.

This initial bimodal shape of the distribution can be explained by Fig. 6.6. A large fraction of the particles in the “200-transition" distribution populates an integral of motion space region with diffusion coefficients of the order of J_r^2. This implies that for most individual particles, the linear variation is of the order of the central radial action itself. Since its sign depends on the direction of the particle's radial velocity, the linear correction can be either positive or negative. Therefore, particles with $J_r = 200\,\mathrm{kpc\,km\,s^{-1}}$ can have a linear “invariant" action which is either larger or

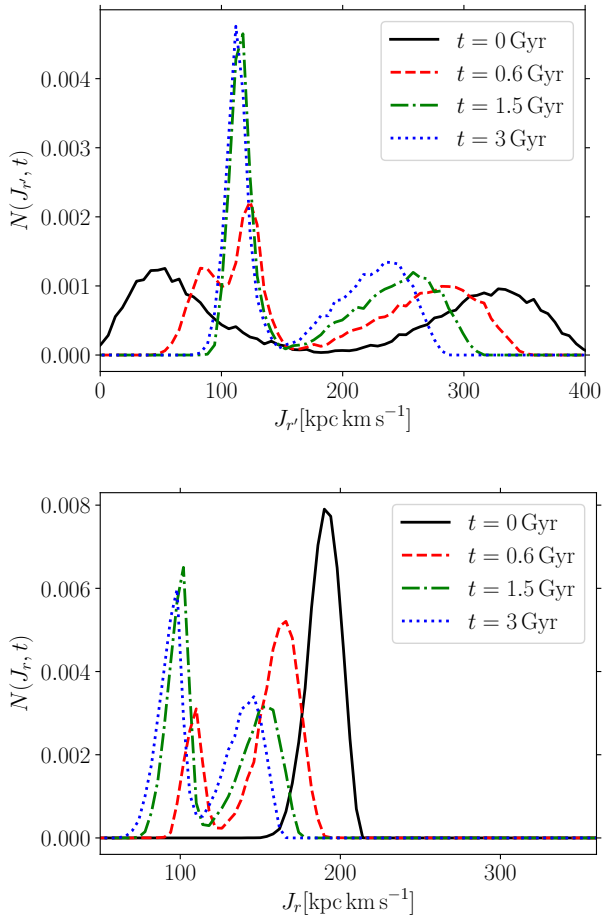

Figure 6.10. Time evolution of $N(J'_r,(t))$ as defined in Eq. (6.20) (upper panel) and of $N(J_r,t)$ (lower panel) for the "200-transition" case (see Table 6.1). Different lines correspond to distributions at different times as described in the legends.

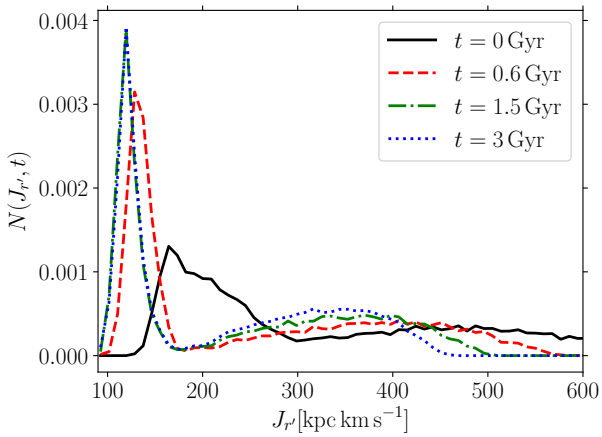

Figure 6.11. Time evolution of the "invariant" distribution for the "200-transition" case (see Table 6.1) calculated using a second order expansion of the radial action (Eq. 6.94). Different lines correspond to distributions at different times as indicated in the legend.

smaller by an amount roughly equal to J_r itself. As a result the diffusion formalism breaks and we here attempt to qualitatively explain the evolution of the radial action distribution.

Right after the start of the simulation, $N(J_r,t)$ develops a bimodal structure. This happens because the initial 200-transition distribution consists mainly of tracers at their orbital apo- and pericentres, respectively. Their subsequent evolution splits the radial action distribution into the populations corresponding to the two peaks of the "invariant" distribution, which can be understood from Eq. (6.20). To explain the subsequent evolution, we calculate $N(J'_r)$ at various times keeping the second order terms of Eq. (6.18) (see Appendix 6.8 for the detailed calculation). The resulting distributions are shown in Fig. 6.11 where J'_r is now defined by Eq. (6.94). Initially, the second order version of $N(J'_r,t)$ is more populated around $J'_r = 200\,\mathrm{kpc\,km\,s^{-1}}$ than the first order (linear) version shown on the upper panel of Fig. 6.10. The bimodal shape of the linear version of $N(J'_r)$ vanishes when including second order corrections. Instead, the second order $N(J'_r)$ obtains an extended tail at large values of J'_r, which demonstrates that higher order corrections are rather significant for the "200-transition" distribution.

Both $N(J_r,t)$ and the second order version of $N(J'_r)$ exhibit a net drift towards smaller radial action values over the course of the simulation. As the second order "invariant" distribution is still not a true invariant, we conclude that the evolution of the "200-transition" distribution is truly non-linear and impulsive. Therefore, the drift observed in Fig. 6.1 is a non-linear effect which cannot be captured by the diffusion formalism developed in Section 6.3. Moreover, we conclude that the analysis presented in Section 6.3.3 is a fairly good indicator of whether the radial action distribution of a set of tracer particles will evolve linearly or non-linearly.

6.5 A couple of consequences for dark matter haloes

So far, we have calculated the evolution of radial actions of tracer particles in generic time-dependent spherical potentials and developed a diffusion formalism for radial action distributions. We then tested our formalism in a time-dependent Kepler potential and found that it accurately describes the time-evolution of radial action distributions in a large part of integral of motion space but breaks down near the “fringe". Hereafter, we discuss a couple of examples where the diffusion formalism developed in Section 6.3, and its limitations (i.e. in the impulsive regime), provide new physical insights. These examples are i) the mass accretion history of a Milky-Way (MW) like halo and ii) core formation in a self-interacting dark matter (SIDM) halo. For both of these cases, we discuss whether or not the rate of change of the self-gravitating potential of the halo implies an adiabatic evolution of the radial action distributions of tracers. Notice that our discussion is solely focused on the evolution that arises due to a global time-dependence of the gravitational potential. Potential resonant diffusion that arises due to local perturbations is not included in our current formalism, but could easily be included, either using Hamiltonian perturbation theory or an extension of the diffusion theory presented here into two dimensions, similar to e.g. Peñarrubia (2015, 2019). The key point of the calculations here is to assess whether conservation of radial actions is a plausible assumption on average, given the rate at which i) the MW accretes mass or ii) an SIDM halo forms a core.

6.5.1 Mass accretion in dark matter haloes

Pontzen and Governato (2013) developed a theoretical formalism to derive the distribution function of a DM halo by maximizing the entropy of an ensemble of collisionless self-gravitating particles, while imposing the extra constraint that the ensemble average of the particles' radial actions is conserved, $\langle J_r \rangle \approx const$. This was motivated by the observation that the evolution of $\langle J_r \rangle$ in simulated haloes is much slower than the variation of the radial action of individual DM particles, J_r. The authors find that the distribution function of simulated haloes is accurately reproduced over several orders of magnitude in J_r. However, the number of particles with low angular momentum is underpredicted in the central regions of the halo (see their Fig. 4), and in turn Pontzen and Governato (2013)'s formalism does not reproduce the ubiquitous CDM cusps. The authors show that this mismatch can be alleviated by including a second, dynamically decoupled population of particles into the formalism, the so-called "cusp" particles.

Collision-less N-body simulations show that centrally-divergent cusps arise during the early build up of DM haloes, when the gravitational potential of these systems undergo impulsive changes. After that epoch, central DM cusps are retained throughout the hierarchical accretion history of the parent halo. In galactic haloes that are initialized with a cored DM profile, a DM cusp can re-grow through dry mergers with massive cuspy dark matter subhalos that fall into the central regions of the host via dynamical friction (Laporte and Peñarrubia, 2015).

Here, we use the analysis presented in Section 6.3.3 to assess if a MW-like halo today – and at a characteristic early time – contains a significant sub-population of DM particles whose radial actions are not conserved on average – which would be a plausible

explanation for the mismatch between Pontzen and Governato (2013)'s analysis and the results of simulations. In our formalism, this sub-population consists of particles that inhabit the "fringe".

To this aim, we apply Section 6.3.3's analysis to the mass accretion of MW-like haloes as follows. We self-consistently construct an idealised MW-size halo following a Hernquist density profile at two different redshifts. We use the mean accretion history reported in Boylan-Kolchin et al. (2010) to determine the mass of the halo at a given redshift and use the rejection sampling scheme described in Burger and Zavala (2019) based on Eddington's formalism (Eddington 1916, Binney and Tremaine 2008) to sample a particle representation of the DM halo in dynamical equilibrium. We then calculate the (theoretical) diffusion coefficient of each individual DM (simulation/sampled) particle from Eq. (6.41). The MW-like halo we consider here has a virial mass of $M_{200} = 1.43 \times 10^{12} M_{\odot}$ and a concentration[3] of $c_{200} = r_{200}/r_{-2} = 12$ at redshift $z = 0$. These are typical values for MW-like haloes (e.g. Boylan-Kolchin et al. 2010).
For the scale factor $R(t)$ (formally given by Eq. 6.9), we use an approximate formula derived in analogy to the scale-free case (Eq. 6.44)[4]:

$$R(t) = \frac{1}{2+\alpha(r,t)} \left(\frac{\Psi(r,t)}{\Psi(r,0)} \right)^{-\frac{1}{2+\alpha(r,t)}}, \tag{6.52}$$

where

$$\alpha(r,t) = \frac{d\log\left(\Psi(r,t)\right)}{d\log(r)} \tag{6.53}$$

is the logarithmic slope of the potential at a given radius and

$$\Psi(r,t) = \Phi(r,t) - \Phi(0,t) \tag{6.54}$$

is the potential shifted such that $\Psi(0,t) = 0$. To calculate the diffusion coefficient, we need the time derivative of the scale factor, namely

$$\begin{aligned} \frac{\mathrm{d}R}{\mathrm{d}t} = &- \frac{1}{\left(2+\alpha(r,t)\right)^2} \times \left(\frac{\Psi(r,t)}{\Psi_0(r)} \right)^{-\frac{1}{2+\alpha(r,t)}} \\ &\times \left\{ \left[1 - \frac{1}{2+\alpha(r,t)} \ln\left(\frac{\Psi(r,t)}{\Psi_0(r)} \right) \right] \times \frac{\mathrm{d}\alpha}{\mathrm{d}t} + \frac{\dot{\Psi}}{\Psi} \right\}. \end{aligned} \tag{6.55}$$

In our model scenario of mass accretion into a MW-size halo, we assume that mass accretes primarily into the outer parts of the halo, leaving the inner mass content unchanged. This is in line with the simple picture of cosmological halo mass assembly in layers/shells in which the concentration parameter $c_{200} = r_{200}/r_{-2}$ evolves with redshift only due to the evolution of $r_{200} \propto (1+z)$ in an expanding Universe. This simple picture is approximately validated by full cosmological N-body simulations (e.g.

[3] With r_{-2} being the radius at which the logarithmic slope of the halo's density profile equals -2 and r_{200} the radius at which the enclosed density equals 200 times the critical density of the Universe.

[4] For a discussion of the validity of this approximation see Appendix 6.9)

Sánchez-Conde and Prada 2014, Ludlow et al. 2014). Under these assumptions, the scale radius in a Hernquist halo can be written as a function of redshift as

$$r_s(z) = 2r_{-2}(z) = \frac{2(1+z)}{c_{200,0}} \left(\frac{GM(z)}{100H^2(z)} \right)^{1/3}, \tag{6.56}$$

where $c_{200,0}$ is the concentration parameter at the current time, $M(z)$ is the redshift-dependent halo mass and $H(z)$ is the Hubble rate. In a Hernquist halo,

$$\alpha(r,z) = 1 - \frac{r}{r + r_s(z)} \tag{6.57}$$

$$\Psi(r,z) = \frac{GM(z)r}{r_s(z)(r + r_s(z))} \tag{6.58}$$

and we can thus write the relevant time derivatives in equation 6.55 as

$$\frac{\dot{\Psi}}{\Psi} = \frac{\dot{M}}{M} - \frac{r+2r_s}{r_s(r+r_s)}\dot{r}_s \tag{6.59}$$

$$\frac{\mathrm{d}\alpha}{\mathrm{d}t} = \frac{r}{(r+r_s)^2}\dot{r}_s, \tag{6.60}$$

where we suppress the redshift-dependence and derivatives are taken with respect to cosmic time. In our model of mass accretion, the density profile is that of a Hernquist halo at all redshifts. Therefore, for most times and radii, we expect that the evolution of the potential's logarithmic slope is negligible. For the cases we study in the following, we have explicitly verified that $\dot{\Psi}/\Psi \gg \dot{\alpha}$. Independent of our definition of the static frame in which we define the action invariant, and hence independent of the definition of $\Psi_0(r)$, we therefore find that

$$\frac{\dot{R}}{R} = -\frac{1}{2+\alpha(r)} \left(\frac{\dot{M}}{M} - \frac{r+2r_s}{r_s(r+r_s)}\dot{r}_s \right). \tag{6.61}$$

In order to evaluate Eq. (6.61) at different times, we adopt the median mass accretion history for MW-like haloes reported in Boylan-Kolchin et al. (2010)

$$M(z) = M_0\,(1+z)^{\eta} \exp\left(-\kappa' \left(\sqrt{1+z} - 1 \right) \right), \tag{6.62}$$

with $\kappa' = 4.9$ and $\eta = 2.23$. Thus

$$\frac{\dot{M}}{M} = \left(\frac{\eta}{1+z} - \frac{\kappa'}{2\sqrt{1+z}} \right) \times \frac{\mathrm{d}z}{\mathrm{d}t} \tag{6.63}$$

$$= \left(\frac{\eta}{1+z} - \frac{\kappa'}{2\sqrt{1+z}} \right) \times (-(1+z)H(z)) \tag{6.64}$$

where $H(z)$ is the Hubble rate at redshift z. At the current time ($z = 0$) we find that for $M(z=0) = 1.43 \times 10^{12} M_\odot$ and $c_{200,0} = 12$,

$$\frac{\dot{M}}{M}(z=0) \approx 1.57 \times 10^{-2}\,\mathrm{Gyr}^{-1}. \tag{6.65}$$

In the past, however, the amplitude of the specific accretion rate may have been very different. As an example, let us look at the redshift $z_{\rm max}$ maximizing $\mathrm{d}M/\mathrm{d}z/M$, which is easily shown to be

$$1+z_{\rm max} = \frac{16\eta^2}{\kappa'^2}. \tag{6.66}$$

At this redshift, we find that $M(z=z_{\rm max}) \approx 3.7\times10^{11} M_\odot$ and

$$\frac{\dot{M}}{M}(z=z_{\rm max}) \approx 0.35\,\mathrm{Gyr}^{-1}. \tag{6.67}$$

The second term in equation 6.61 can be written as

$$\dot{r}_s = -(1+z)H(z)\frac{\mathrm{d}r_s}{\mathrm{d}z}. \tag{6.68}$$

Throughout our calculations, we use $\Omega_{\Lambda,0}=0.7$ and
$\Omega_{m,0}=0.3$. Using Eq. (6.61), we can now calculate the theoretical diffusion coefficients at different redshifts using Eq. (6.41). To that end, we calculate energy, angular momentum, radius and radial velocity of each (simulation) particle in our halo. We then numerically calculate the radial action, the radial period, and $\langle \mathbf{r}\cdot\mathbf{v}\rangle$ (according to Eq. 6.42) of each particle. From the ratio $\sqrt{\tilde{D}}/J_r$ we then construct the "linear", "transition" and "fringe" areas introduced in Fig. 6.3.

Fig. 6.12 shows the integral of motion space distribution of one million particles at different times. In the top (bottom) panel, we show the distribution at the redshift $z_{\rm max}$ (today). In each panel, we show the fraction of particles inhabiting the different integral of motion space areas, which correspond to the ones introduced in Fig. 6.3 for the time-dependent Kepler potential. To determine the boundaries between the different regions, we sort the particles within different radial action bins by their respective values of $\sqrt{\tilde{D}}/J_r$.

In the upper panel of Fig. 6.12, we see that at $z=z_{\rm max}$ only $\sim 41\%$ of the particles in the DM halo inhabit the "linear" integral of motion space area in which the diffusion formalism applies, while $\sim 18\%$ are located in the "fringe", where we expect impulsive evolution. Just as in the Kepler case, at large values of J_r, we find that the integral of motion space area of the "linear" regime shrinks and eventually disappears with all particles at very large radial actions being in the "fringe", where the evolution of the radial action is non-linear. We note that in the Kepler case, particles with these large actions tend to drift towards smaller values of J_r and get "locked in" on short to intermediate time-scales (see Fig. 6.1). Since such drift does not seem to occur in the opposite direction in the Kepler case, radial actions which are too large to satisfy a "linear" evolution are effectively erased from the distribution. Assuming that this quantitative behaviour will be the same in a DM halo (with a potential closer to that of a NFW distribution) an integral of motion space distribution of DM particles such as the one seen at $z=z_{\rm max}$ could have significant consequences. In particular, if the "fringe" is as populated as in the upper panel of Fig. 6.12, the radial action distribution is expected to drift towards smaller radial action at later times, i.e., the amount of particles with smaller radial actions increases and the tail of the distribution is erased.

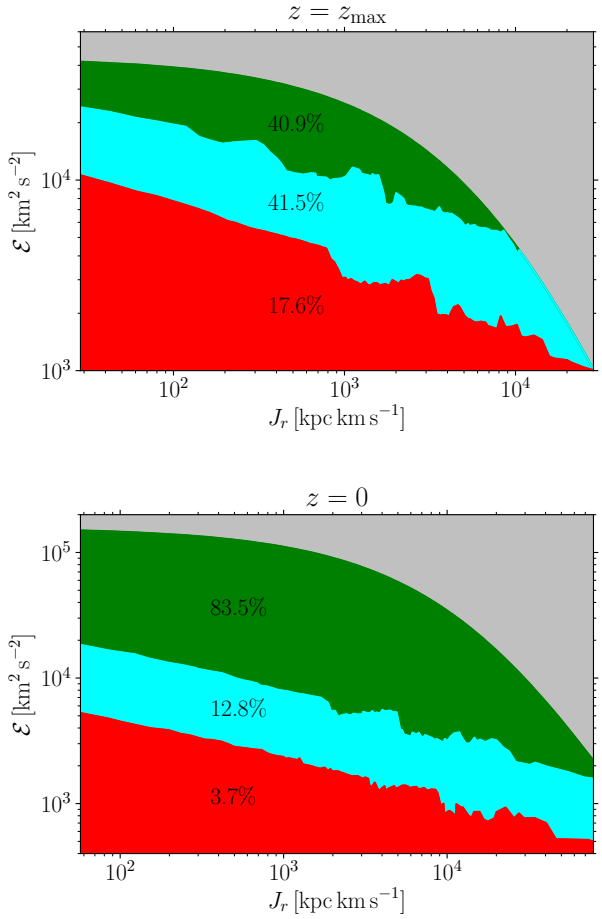

Figure 6.12. Occupancy of the $\mathscr{E} - J_r$ space in an idealised MW-size halo accreting mass according to the mass assembly history obtained in cosmological simulations (see Eq. 6.64). The silver area is not populated, the green area is the "linear" regime in which $\sqrt{\tilde{D}}/J_r < 0.1$. The area coloured in cyan is the "transition" area in which $\sqrt{\tilde{D}}/J_r < 1$ and the red area is the "fringe". The fraction of particles in the different regions is shown as a percentage of the total amount of particles. The top panel shows the integral of motion space occupancy at $z = z_{\rm max}$ (see text), the bottom panel at $z = 0$.

The lower panel of Fig. 6.12 shows the integral of motion space occupancy today. In contrast to the upper panel, the evolution of radial actions in a MW-size halo today is on average significantly more adiabatic, with 83 per cent of particles inhabiting the “linear" regime, as opposed to only 3.7 per cent that inhabit the “fringe". This implies that *at the current time*, it is fair to assume that the ensemble average of radial actions is conserved in a MW-size halo, as shown in the bottom panel of Fig. 6.12. The results of our simple analysis confirm that the assumptions made by Pontzen and Governato (2013) are well founded. Yet, the top panel suggests that the (mean) evolution of the gravitational potential becomes gradually more impulsive the further we go back in time. Since we know that cusps form immediately after the gravitational collapse of the parent halo, this likely points to a link between the formation of the cusp and the impulsive evolution of the gravitational potential. Fig. 6.1 shows that an impulsive evolution leads to a net evolution of radial action distributions towards smaller actions. A possible explanation for that may come from Eqs. (6.10) and (6.20). In fast evolving potentials, infalling particles have lower energies and thus shorter “instantaneous radial periods". Such infalling tracers may thus be “locked" into the regime in which $\Delta/J_r \lesssim 1$ and Eq. (6.20) applies. Fig. 6.12 demonstrates that at a fixed radial action, the linear oscillation amplitude (see also Eq. 6.20) is (on average) smaller for particles with smaller energies. Comparison between the two panels also reveals that faster evolution of the gravitational potential shifts the “linear” integral of motion space towards lower radial actions and energies. At a fixed radial action, particles with smaller energies also have smaller angular momenta. In very fast evolving potentials, this lock-in mechanism is therefore most efficient for particles with low angular momenta. This way, impulsively evolving potentials themselves would create populations of particles with small radial actions and low angular momenta. These may be the “cusp particles" of Pontzen and Governato (2013). However, since this is an impulsive/non-linear effect, a complete understanding cannot be obtained using the diffusion formalism presented here. A first step towards an overall better understanding will be to expand our formalism beyond the analysis of tracers and to self-consistently evolve the radial action distribution and the gravitational potential of an ensemble of gravitating particles at subsequent time-steps.

Finally, we note that diffusion formalism may be useful to improve the analysis of tidal streams in the MW. Buist and Helmi (2015) have analyzed the evolution of radial action distributions of tidal streams in a time-dependent Aquarius potential. In tidal streams whose orbits are not adiabatic, they find an increased spread in radial action between the stream particles when compared to the final distribution in a static potential. In the bottom panel of Fig. (6.12), we find that $\sim 17\%$ of DM particles today inhabit a region in integral of motion space in which diffusion of radial actions due to the time-dependence of the MW is a relevant phenomenon. Stars in tidal streams are essentially tracers of the gravitational potential and thus our results suggest that the observed increased spread in radial action is the result of a diffusion process in radial action space caused by the time-dependence of the potential. This diffusion in radial action space constitutes an additional baseline error source when using the clustering of tidal streams in action space to constrain the potential of the Galaxy, as was done, e.g, by Sanderson et al. (2015) and Sanderson et al. (2017). In particular, our results suggest that the accuracy of the assumption that the true potential implies the most tightly clustered radial action distribution depends on both the accretion history of the

MW and the orbit of the tidal stream's progenitor.

6.5.2 Cusp-core transformation in SIDM

Burger and Zavala (2019) looked at the evolution of a set of tracer particles with a Gaussian distribution in radial action orbiting in the potential of a dwarf-sized DM halo developing a constant density core in one of two different ways[5], through elastic self-scattering between the DM particles (SIDM) or through impulsive energy injections into the system akin to supernova feedback. Fig. 7 in Burger and Zavala (2019) shows a comparison of the final radial action distributions of the tracer particles in these two reference scenarios, as well as a third baseline scenario in which the host halo retains its cusp. From this figure, it is clear that $N(J_r,t)$ at the end of the SIDM simulation, which has a cusp-core transformation, is very close to the final distribution in the benchmark simulation (without a cusp-core transformation). Presumably, the difference between those two distributions can be explained by some small amount of diffusion in radial action space. Hence, we expect that a large part of the $\mathscr{E} - J_r$ space area occupied by particles orbiting in a SIDM halo belongs to the "linear" regime of radial action evolution.

To test this hypothesis, we perform the analysis introduced in Section 6.3.3 in a SIDM halo similar to the one used in Burger and Zavala (2019) following Eq. (6.41). The only difference is that we re-run the SIDM simulation with a smaller self-interaction cross section of $\sigma/m_\chi = 1\mathrm{cm}^2\,\mathrm{g}^{-1}$. We save snapshots (simulation outputs) every 0.14 Gyr and calculate the potential $\Psi(r)$ by averaging the potential of tracer particles in logarithmically spaced spherical shells. To obtain $\Psi(r,t)$ we interpolate $\Psi(r)$ between snapshots. The scale factor itself is calculated according to Eq. (6.52). We calculate $\alpha(r,t)$ numerically by interpolating the function $\log(\Psi)(\log(r))$ and calculating its derivative at each (simulation) particle's radius. Since the SIDM halo changes shape we cannot neglect the time-dependence of $\alpha(r,t)$ and the derivative of the scale factor at a given radius given by Eq. (6.55). $\dot{\alpha}$ and $\dot{\Psi}$ are calculated numerically at each particle's position. We then calculate $\tilde{D}(J_r,\mathscr{E})$ from Eq. (6.41) as in Section 6.5.1, but using $\Phi(r,t)$ of the SIDM halo. Core formation in a SIDM halo is a gradual process. Initially, self-interactions are most efficient in the halo's centre and cause a relatively fast decrease of the central density. Subsequently, the forming core slowly thermalizes. In this latter stage, due to the prior mass redistribution, the density is completely flat (isothermal core) in the centre followed by a small region outside the inner core (but still within the scale radius) where the density rises slightly above that of the original profile (see e.g Vogelsberger et al. 2012, 2014). At the end of our simulation, the halo's core is fully thermalized and in a transient quasi steady state. Notably, including baryons into the simulation can change the phenomenology of SIDM (see Santos-Santos et al. 2020 or Robles et al. 2017). We briefly note that there is an additional phase called gravothermal collapse (e.g. Zavala et al. 2019 and Turner et al. 2020) in which the core collapses to very large densities. However, this phase is only relevant for very large cross sections.

Fig. 6.13 shows the evolution of the central part of the SIDM-halo's density profile

[5]The halo used was a dwarf-sized halo which initially has a Hernquist density profile with mass $M_{200} = 1.43 \times 10^{10} M_\odot$; see Table 1 of Burger and Zavala (2019) for the relevant simulation parameters.

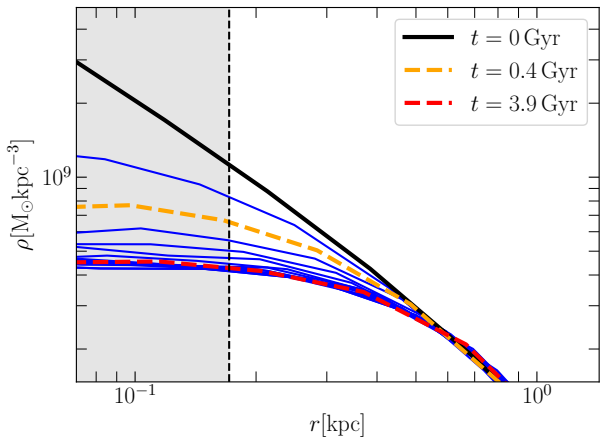

Figure 6.13. Evolution of the density profile in the inner region of a dwarf-size SIDM halo with $\sigma/m_\chi = 1\mathrm{cm}^2\mathrm{g}^{-1}$. *Lines denote the halo density profile at different times. The black solid line marks the density profile at the beginning of the simulation , whereas the orange and red dashed lines are the density profiles after* 0.4 Gyr *and* 3.9 Gyr *of simulation time respectively. Blue lines indicate the evolution in between other snapshots. The gray shaded area denotes the resolution limit of the equivalent CDM-only simulation (see Burger and Zavala 2019).*

in our simulation. The black solid line is the initial density profile. In dashed orange we show the density profile after 0.4 Gyr and the red dashed line is the density profile after 3.9 Gyr. Blue lines denote snapshots between or after the ones highlighted in the legend. The grey shaded area shows the resolution limit for CDM simulations (Power et al. 2003). We note that SIDM profiles are usually converged to well within this so-called Power radius (Vogelsberger et al., 2012; Rocha et al., 2013; Vogelsberger et al., 2014). The orange density profile corresponds to a time during the initial stage of core formation when the mass density in the innermost part of the halo decreases rapidly. The red line, on the other hand, corresponds to a time at which the core is fully formed. We have chosen those two times because they represent different stages of the core formation process in a SIDM halo.

Fig. 6.14 shows the integral of motion space occupancy for the SIDM halo at the times corresponding to the orange dashed line (0.4 Gyr) and the red dashed line (3.9 Gyr) in Fig. 6.13. The areas in the plot correspond to the equivalent areas in Fig. 6.12. The change in gravitational potential due to core formation (induced by DM self-interactions) causes an almost adiabatic evolution of radial action distributions over the entire simulation time. In both panels of Fig. 6.14, more than 97 per cent of DM particles are in the “linear" integral of motion space area and only around 1 per cent inhabit the “fringe". As a consequence, we expect the radial action distribution of tracer particles to be altered very little, if at all, by the change in gravitational potential triggered by SIDM, provided that the numerical values of the self-interaction cross

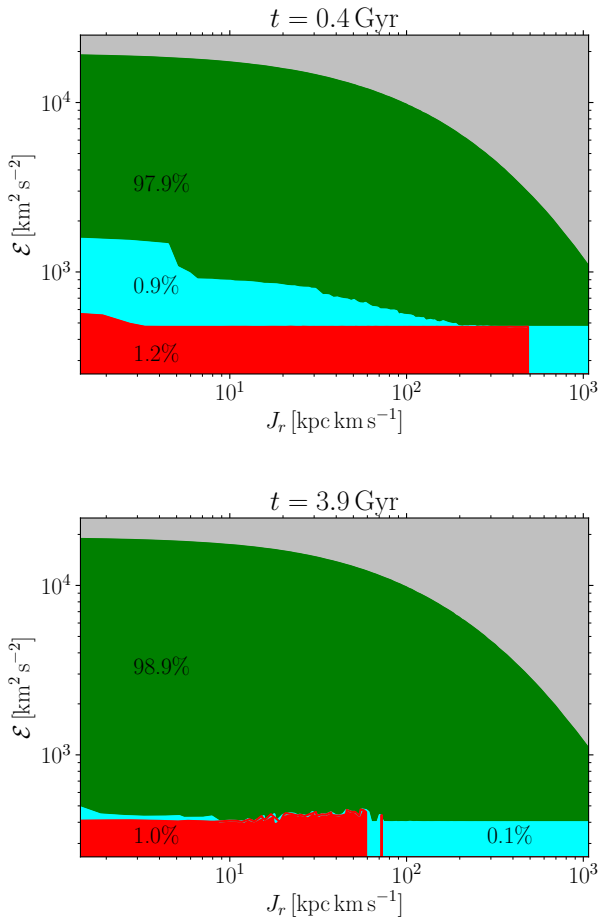

Figure 6.14. The $\mathcal{E} - J_r$ space occupancy in a dwarf-size halo forming a core due to self-scattering between its DM particles. The different curves and areas are equivalent to those in Fig. 6.12. The fraction of DM particles populating each area of integral of motion space is shown as a percentage in the respective areas. The upper panel corresponds to the orange dashed line in Fig. 6.13 (early-stage of core formation), while the lower panel corresponds to the red dashed line in Fig. 6.13 (late-stage of core formation).

section lies within the range that is allowed by astrophysical constraints.

We remark that the statement that the evolution of radial action distributions in an SIDM halo is expected to be linear and can be described by our diffusion formalism applies only to kinematic tracers. For the DM particles themselves, SIDM-induced elastic collisions are a different matter as the radial action of both collision partners changes in a random fashion, since their post-collision orbits are entirely different from the pre-collision orbits.

We apply the diffusion formalism presented in Sections 6.3 and 6.4 to describe the evolution of a Gaussian radial action distribution comprised of 20000[6] tracer particles orbiting in the core-forming SIDM halo. We use the approximations given by Eqs. (6.52) and (6.55) for $R(t)$ to calculate the drift and diffusion coefficients. In Appendix 6.9, we demonstrate that the evolution of radial actions of kinematic tracers in a potential with a time-dependent shape is accurately described by this approximate scale factor. Since the rate at which the shape of the potential in Appendix 6.9 changes is larger than the rate at which the potential of the SIDM halo changes, we conclude that we can use Eqs. (6.52) and (6.55) to obtain accurate predictions of the evolution of $N(J_r,t)$ in a core-forming SIDM halo.

Fig. 6.15 shows the “measured" radial action distribution (black solid line) after 2.1 Gyr of simulation time in the top panel and after 3.9 Gyr in the bottom panel, along with the result of the diffusion formalism (green dashed line) and the initial distribution (red dotted line). Since the calculation of the drift and diffusion coefficients involves numerical derivatives of time-dependent measured quantities, we take $N(J_r,t)$ at $t = 0.4$ Gyr as the initial distribution in order to assure that the time-derivatives are numerically stable. The result of the diffusion formalism is overall in good agreement with the measured distribution in Fig. 6.15 (particularly in the early stages of core formation; see upper panel). However, in the late stage of core formation (between 2.1 Gyr and 3.9 Gyr) the measured distribution gradually forms a slightly more extended tail towards larger radial actions. This evolution, albeit not very sizeable, is not predicted by our formalism. In fact, our diffusion formalism predicts hardly any evolution at all, which might imply that this is a significant deviation. In Appendix 6.9, however, we show that the amplitude of the oscillation of radial actions in a potential developing a core is in general well approximated by our diffusion formalism. We thus surmise that the small but gradually increasing mismatch between the measured radial action and the result of the diffusion formalism is caused by numerical effects arising due to the discrete sampling of the SIDM halo. In fact, since the DM simulation particles are more massive than the tracers, individual close encounters between tracers and DM simulation particles can occasionally lead to a small sudden increase in the tracer's energy, leading to a small increase in radial action as well.

Nonetheless, as we find in Appendix 6.9, a shape-changing host potential can cause sizeable oscillations of the radial action values if the change in shape of the potential occurs fast enough. The fact that our diffusion theory does not predict any sizeable evolution of $N(J_r,t)$ implies that the change of the shape of $\Psi(r,t)$ is rather slow between $t = 0.4$ Gyr and $t = 3.9$ Gyr. In Fig. 6.16, we look at the time evolution

[6]To enable us to derive smoother drift and diffusion coefficients, we have added 18000 tracers to the original 2000 used in Burger and Zavala (2019). They are set up in exactly the same way as the initial 2000 tracers and thus obey the same Gaussian distribution.

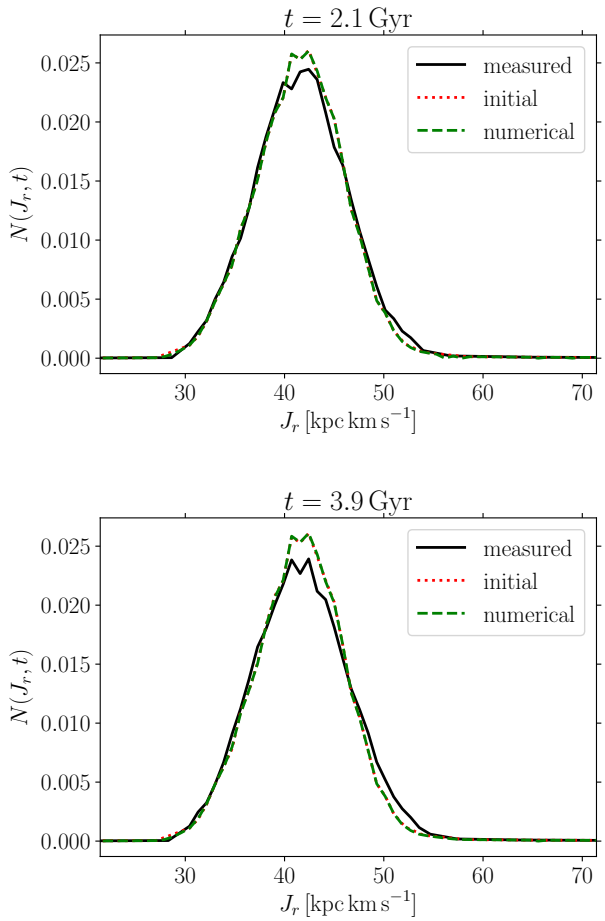

Figure 6.15. Radial action distribution of 20000 tracer particles after 2.1 Gyr *of simulation time (top panel) and after* 3.9 Gyr *of simulation time (bottom panel) in the core-forming SIDM halo shown in Fig. 6.13. The black lines show the measured distribution. The red dotted line is the initial distribution and the green dashed line denotes the result of the diffusion formalism.*

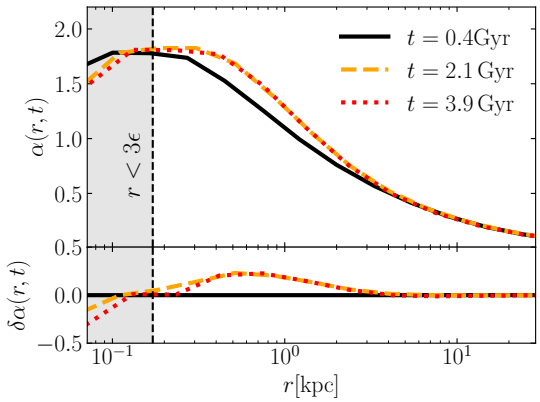

Figure 6.16. Logarithmic slope $\alpha(r,t)$ of the potential of the SIDM halo in Fig 6.13 as a function of radius at different times (upper panel) and change of the logarithmic slope with respect to $t = 0.4$ Gyr, also as a function of radius (lower panel). The times shown match the ones in Fig. 6.15. The grey shaded area highlights the resolution limit as in Fig. 6.13.

of the logarithmic slope $\alpha(r,t)$ of the SIDM halo's potential as a function of radius. We show the logarithmic slope of the potential at three different times, $t = 0.4$ Gyr (solid black line), $t = 2.1$ Gyr (dashed orange line) and $t = 3.9$ Gyr (dotted red line). In the lower panel, we furthermore show the change of the logarithmic slope with respect to $t_0 = 0.4$ Gyr, defined as

$$\delta\alpha(r,t) = \alpha(r,t) - \alpha(r,t_0). \tag{6.69}$$

We find that the evolution of $\alpha(r,t)$, which directly affects the ratio $\dot{R}/R$ through the time-derivative $\dot{\alpha}(r,t)$, is most significant at earlier times and at smaller radii. In fact, in Fig. 6.16 we show that most of the evolution in $\alpha(r,t)$ occurs between 0.4 Gyr and 2.1 Gyr. Afterwards, the shape of the potential is essentially constant, implying that at that point the core is fully formed. Considering that Fig. 6.15 hardly shows any evolution of both predicted and measured $N(J_r,t)$ between $t = 0.4$ Gyr and $t = 2.1$ Gyr, we conclude from Fig. 6.16 that there should not be any physical evolution afterwards either. We take this as confirmation that the small observed drift in this time interval is due to discreteness effects, i.e., it occurs because the sampled potential is not perfectly smooth.

In conclusion, we find that SIDM with a cross section of $\sigma/m_\chi = 1\,\mathrm{cm^2 g^{-1}}$ causes virtually no evolution in $N(J_r,t)$ and can thus be considered a fully adiabatic process. In a cosmological SIDM halo the impact of mass accretion on the radial action distribution of tracers will thus far outweigh the impact of SIDM.

6.6 Conclusion

Cosmological simulations of CDM yield concordant results on all resolved scales. An example is the inner density profile of collapsed DM haloes, which is found to be universal over 20 orders of magnitude in halo mass (Wang et al. 2020). Simulations of the dynamics of virialized self-gravitating systems have also provided valuable insight into the long-term evolution of such systems, modeling, for instance, the shapes of observed galaxies or the formation of tidal streams in the galactic halo.

However, many of the results of N-body simulations have not yet been understood from fundamental principles. An approach that is frequently used to model the secular evolution of bound systems is Hamiltonian perturbation theory (e.g. Lynden-Bell and Kalnajs, 1972; Tremaine and Weinberg, 1984). In this method, the (time-dependent) Hamiltonian is written as a power series expansion in some small parameter ε and the zero'th order term is a time-independent Hamiltonian that can be written in terms of canonical action variables (see Eq. 6.1). Hamiltonian perturbation theory can be formulated into a kinetic theory for the evolution of self-gravitating systems (Heyvaerts 2010), an approach that has proven quite successful in recent years (e.g. Fouvry et al. 2015). However, a crucial assumption behind Hamiltonian perturbation theory is that actions are integrals of motion of the zero'th order term H_0 and that they evolve very slowly. The theory performs well for potentials in which these conditions are fulfilled, but breaks once actions evolve significantly with time.

An alternative to Hamiltonian perturbation theory is to use statistical mechanics. Historically, however, deriving a thermodynamic model for the evolution of self-gravitating systems has proven to be difficult for a variety of reasons. Negative specific heat and non-ergodicity of gravitating systems forbid the use of canonical or grand canonical ensembles. On top of that, Padmanabhan (1990) shows that in gravitating systems the energy is not an extensive parameter, thus the system cannot be divided into non-interacting macrocells and the laws of standard thermodynamics do not apply. Despite these challenges, several attempts have been made to describe the evolution of self-gravitating systems within the framework of statistical physics. Pontzen and Governato (2013) attempted to derive the distribution function of a virialized CDM halo, using a maximum entropy approach and the additional constraint that the ensemble average of the DM particle's actions be conserved. Their formalism accurately recovers the distribution function of simulated haloes over several orders of magnitude, but does not predict the centrally-divergent density cusps without adding a second, dynamically decoupled particle population with low angular momenta.

Peñarrubia (2013) developed a coordinate transformation relating the equations of motion in a frame with a time-dependent potential to the dynamics in a different frame in which the potential is static. The integrals of motion in this new reference frame expressed in the coordinates of the original frame are so-called dynamical invariants. Subsequently, Peñarrubia (2015) developed a statistical theory for the diffusion in energy space of a particle ensemble in the time-dependent potential based on those dynamical invariants.

In this paper, we have developed a similar theory for the diffusion of particle ensembles in radial action space. For now, our theory is restricted to spherically symmetric potentials but we note that a modified theory that includes diffusion along

resonant spatial directions can be obtained as discussed in Peñarrubia (2015). In our theory, the diffusion is derived by treating the tracers as a microcanonical ensemble and relating the drift and diffusion coefficients to the evolution of the radial actions of individual tracers. This makes our approach fundamentally different from Hamiltonian perturbation theory, in which perturbed systems do not depart strongly from a state of dynamical equilibrium. To first order, we show that the time-evolution of the radial action of a particle moving in a time-dependent spherical potential is given by an oscillation around the invariant action, with an amplitude that is fully determined by the phase space coordinates of the particle and the transformation between the static and the time-dependent frames. Since actions are traditionally considered adiabatic invariants, we can use the ratio between the amplitude of the linear oscillation and the radial action to determine whether the evolution takes place in an adiabatic or impulsive regime. We illustrate this theory with two examples i) an idealised Milky-Way (MW) size galaxy with a time-dependent potential that grows according to the average mass accretion history of MW-like ΛCDM haloes, and ii) the time-dependent potential that of a dwarf-size self-interacting DM (SIDM) halo forming a constant density core. A summary of our main methods and key results is as follows:

1. *Calculating the linear variation of J_r.*- The starting point is to find a numerical (or approximate) solution to the differential equation for the scale factor $R(t)$ which defines the coordinate transformation introduced in Peñarrubia (2013). We then write the action invariant J_r' defined in the new time-independent reference frame in terms of the phase space coordinates of the particle in the original frame. Under the assumption that the potential evolves slowly, we make a first order Taylor expansion of J_r' in $\dot{R}/R$ (see Section 6.2.2). We find that to zero'th order, the dynamical invariant J_r' and the radial action J_r in the original reference frame are the same. The first order correction to this equality is an oscillation with an amplitude proportional to $\dot{R}/R$ – and thus to $\dot{\Phi}/\Phi$ – around $J_{r'}$. This oscillation happens in phase with the radial motion and is proportional to the radial period of the particle's orbit. We have here for the first time presented an analytical calculation of the linear oscillation amplitude of radial actions in time-dependent potentials (Eq. 6.20). As a numerical test of this linear model, we follow the orbits of three tracer particles in a time-dependent analytic Kepler potential using the N-body code AREPO (Springel, 2010) (see Section 6.2.3). In case the oscillation amplitude is small compared to J_r itself our linear model provides an excellent fit to the time-evolution of J_r (see top panel of Fig. 6.2). For orbits in the "fringe" of the potential ($E \approx 0$), such as the one shown on the bottom panel of Fig. 6.2, the linear oscillation amplitude is of the order of J_r itself and the radial action's evolution becomes non-perturbative.

2. *Formulating a diffusion formalism in radial action space.* - Having calculated the linear variation of J_r for individual particles in a time-dependent potential, we then aim to statistically predict the evolution of the radial action distribution $N(J_r,t)$ of a particle ensemble. In Section 6.3.1, we show that if the time evolution of the gravitational potential is sufficiently slow, $N(J_r,t)$ is related to the invariant distribution of dynamical action invariants $N(J_r')$ through a diffusion equation. The drift and diffusion coefficients $\tilde{C}(J_r,t)$ and $\tilde{D}(J_r,t)$ which appear in the

diffusion equation are proportional to the ensemble average and the variance of the first order Taylor correction, respectively. Having computed $\tilde{C}(J_r',t)$ and $\tilde{D}(J_r',t)$, we can calculate $N(J_r,t)$ from the invariant distribution $N(J_r')$ through a simple Gaussian convolution. We define a numerical criterion based on $D(J_r,t)$ to determine whether the expected evolution of radial action distributions can be described as a diffusion process. Dependent on the ratio between $\sqrt{\tilde{D}}$ and J_r we define three different regions in integral of motion space:

- In the **linear** regime, $\sqrt{\tilde{D}} < 0.1\,J_r$. We expect that the diffusion formalism will provide a good prediction of the evolution of $N(J_r,t)$ for particle distributions inhabiting this area in phase space.
- In the **transition** regime, $0.1 < \sqrt{\tilde{D}}/J_r < 1$ we expect (average) deviations from linear evolution, and thus our expectation is that the diffusion formalism will yield only rough results when applied to particle distributions in this part of phase space.
- In the **fringe**, $\sqrt{\tilde{D}} > J_r$, and thus we expect that the diffusion formalism does not apply.

To confirm these predictions, we perform a set of 5 restricted simulations of 20000 tracer particles orbiting in the same time-dependent Kepler potential as the three individual tracers before and covering different integral of motion space regions (see Table 6.1 and Fig. 6.3 for the simulation parameters and the restrictions in phase space). A theoretical calculation of the diffusion coefficients in energy-action space agrees well with the coefficients measured from the initial tracer distributions (see Figs. 6.3−6.6). For the “10-linear" simulation we show that if the drift and diffusion coefficients are measured directly from the particle distribution, the predictions of the diffusion formalism are in perfect agreement with direct measurements of both the invariant distribution $N(J_r')$ and $N(J_r,t)$ at the end of the simulation (see Figs. 6.7, 6.8). Moreover, we find that a theoretical calculation of the diffusion coefficient yields equally sound results provided the drift coefficient is negligible, i.e., for fully phase-mixed particle ensembles (see Section 6.4.2 for a discussion). In general, we find that the agreement between the measured distribution and the prediction of the diffusion formalism deteriorates the larger the fraction of particles populates the “transition" region in integral of motion space as opposed to the “linear" region (Section 6.4.3, Fig. 6.9). As a consequence, the evolution of $N(J_r,t)$ in the ‘200-transition" simulation is highly non-linear and cannot be modelled as a diffusion process (see Section 6.4.4). We attempt to explain the observed evolution using higher-order perturbation theory (see Fig. 6.11 and Appendix 6.8) but we have to concede that while we manage to obtain an improved understanding of the shortcomings of the first order approximation, a full understanding of the non-linear evolution of $N(J_r,t)$ in the “fringe" remains beyond our reach. Phenomenologically, however, we observe that on average, there is a clear trend for particles whose evolution is non-linear to drift towards considerably smaller radial actions on the simulation time-scale (see Figs. 6.1, 6.2, 6.10). In fact, particles whose actions are initially so large that there is no “linear" integral of motion space available (see Fig. 6.3)

tend to quickly drift towards values of J_r where the evolution becomes linear. We conclude that fast changes in the gravitational potential occurring over a short period of time can permanently alter the radial action distribution of particle ensembles in a non-trivial way. This effect is particularly important if the fraction of particles in the “fringe" area is significant while the change in gravitational potential occurs.

3. *Implications for halo mass accretion and cusp-core transformation.* - Following our arguments from Section 6.3.3 we can decide from the ratio $\sqrt{\tilde{D}}/J_r$ whether the evolution of a particle ensemble's radial action distribution is linear, and hence whether actions are approximately conserved on average. For illustration, we analyze whether radial action distributions evolve linearly during the accretion history of a MW-size halo and during the process of cusp-core transformation in a dwarf-size SIDM halo.

Mass accretion in a MW-size halo

To assess whether actions evolve linearly during the mass accretion history of a typical MW-size halo we calculate the relative accretion rate $\dot{M}/M$ as a function of redshift from the mass-redshift relation reported in Boylan-Kolchin et al. (2010) and determine the redshift $z_{\rm max}$ that maximizes $dM/dz/M$. We then self-consistently sample the phase space coordinates of 10^6 particles from the distribution function of a Hernquist halo with mass $M(z)$ and a scale radius $r_s(z)$ - both today and at $z_{\rm max}$ (see Section 6.5.1). Calculating $\tilde{D}(J_r,t)$ for each particle, we determine the fraction of particles inhabiting the “linear" part, the “transition" part and the “fringe" of the available integral of motion space (see Fig. 6.12). We find that at the current time, the radial action distribution of particles in a MW-size halo evolves rather linearly, with 83 per cent of particles being in the “linear" regime of phase space. However, at $z = z_{\rm max}$, we find that radial actions could not be considered conserved quantities on average. In fact, with 18 per cent of particles inhabiting the “fringe" and another 41 per cent being in the “transition" area in phase space, we expect a highly non-linear evolution on average. Furthermore, we find that at $z = z_{\rm max}$ no linear integral of motion space is available at the largest radial actions.

From our discussion of the Kepler case, we expect that the radial action distributions of DM particles in haloes drift strongly towards smaller actions at early times (high redshifts). This impulsive drift is likely related to the formation of primordial cusps. At later times, our results support Pontzen and Governato (2013)'s statement that the ensemble average of radial actions is approximately conserved. We note, however, that there is a sparsely populated volume in phase space where the evolution of radial actions is not completely linear, even at late times. Our results suggest that the action distribution of tidal streams in Milky-Way like galaxies is not invariant, and that the hierarchical growth of the host DM halo necessarily causes diffusion and drift in radial actions. This may explain the results of Buist and Helmi (2015), who show that streams in time-dependent potentials have a systematically larger spread of radial actions than streams in static potentials. Treating the evolution of radial actions as a diffusive process may

thus be a way to improve the analysis of tidal streams in the MW and constrain the time evolution of the Galaxy potential.

It is important to bear in mind that our diffusion formalism provides approximate results. The most significant caveat of our theory is that we consider diffusion in only one dimension. While this is sufficient to describe the impact of a global time-dependence of the potential, we cannot model resonant diffusion that may arise due to local perturbations caused by mergers, for instance. Taking these effects into account requires modeling the diffusion in multiple dimensions (see Peñarrubia 2019 for how this may be done). For now, our main achievement is that by modeling the diffusion of radial action distributions in time-dependent potentials we can successfully describe ensembles of tracer particles in potentials whose rate of evolution lies between adiabatic and impulsive. Since we do not resolve the effect of localized perturbations, the diffusion of stars in tidal streams predicted by our formalism would constitute a lower limit to the observed diffusion.

Cusp-core transformation in a dwarf-size SIDM halo

Finally, we use the analysis presented in Section 6.3.3 to investigate whether cusp-core transformation in a dwarf-size SIDM halo leads to linear or non-linear evolution of radial action distributions of tracer particles (see Section 6.5.2). We re-run the SIDM simulation of table 1 in Burger and Zavala (2019), increasing the number of tracers by a factor of ten and changing the self-scattering cross section to $\sigma/m_\chi = 1\,\mathrm{cm^2 g^{-1}}$, a value which is around the current constraint on SIDM at dwarf scales (Read et al., 2018). We take snapshots every 143 Myr and closely follow the orbits of the tracers at each time step. We use the same functional form for the scale factor as for the MW-size halo above. However, we now take the changing halo shape into account when calculating its time derivative. At each snapshot, we calculate the fraction of DM particles in the “linear" and the “transition" regime as well as the “fringe" by calculating $\tilde{D}$ and J_r for each particle. The fraction of particles in the linear regime is never below 97 per cent (see Fig. 6.14) and we conclude that cusp-core transformation due to SIDM leads to a very slow change in the halo’s potential, implying that the evolution of radial action distributions of particles orbiting in the SIDM halo is adiabatic.

Our diffusion formalism predicts that $N(J_r,t)$ remains essentially constant during core formation. However, the measured radial action distribution gradually develops a tail towards larger radial actions that extends beyond the prediction made by the diffusion formalism (see Fig. 6.15). For that reason, we test in Appendix 6.9 whether the use of the approximate scale factor $R(t)$ is valid in potentials with a time-dependent shape. Modeling core formation in a smooth external potential rather than a live halo, and using the approximate scale factor, we find the linear Taylor expansion of the time evolution of J_r to be in very good agreement with the directly measured evolution. Furthermore, we show that during the time in which the extended tail forms in the measured distribution $N(J_r,t)$, the shape of the potential is largely constant (see Fig. 6.16), implying

that radial actions should be conserved during that time. This leads us to conclude that the observed evolution of $N(J_r,t)$ is caused by individual close encounters between tracers and DM (simulation) particles, and is thus a numerical effect. Therefore, we take the prediction of our diffusion formalism at face value and deduce that core formation in dwarf-sized haloes due to SIDM with $\sigma/m_\chi \sim 1\,\mathrm{cm}^2\mathrm{g}^{-1}$ has no detectable impact on the radial actions of kinematic tracers. In other words, it is fully adiabatic.

In summary, we have here for the first time analytically calculated the linear time evolution of radial actions in time-dependent spherical potentials. Based on that, we have developed a diffusion theory for radial action distributions of tracers in time-dependent gravitational potentials. This diffusion theory relates the evolution of the distribution $N(J_r,t)$ to an invariant distribution $N(J_r')$ in a frame which is related to the time-dependent frame via a coordinate transformation defined by a scale factor. We have demonstrated the validity of our formalism, and furthermore provided a discussion of its limitations, by performing restricted N-body simulations of tracer particles with different initial distributions in radial action and applying the diffusion formalism to model how these distributions evolve with time. We have shown that the size of the diffusion coefficient $\tilde{D}(J_r,t)$ relative to the square of the radial action itself is a good indicator for whether radial action distributions will evolve linearly or not, and thus for whether the diffusion formalism we developed is applicable. We have furthermore shown that highly non-linear evolution causes a drift of radial action distributions. We have then applied our mechanism to two distinct cases of interest in astrophysics.

First we have argued that while the radial actions of DM particles in a MW-size halo are likely conserved on average at late times, this was not the case shortly after the gravitational collapse of these objects. At the current time, we expect the evolution of radial action distributions of gravitational tracers to be mildly diffusive. Diffusion becomes increasingly important at earlier times, and we predict a significant asymmetric drift towards smaller radial actions for $z \gtrsim 2$. The impulsive evolution of haloes at very early times may be linked to the formation of primordial cusps, and taking into account diffusion due to the time-dependence of the MW halo's potential can potentially improve the analysis of tidal streams. As a second application, we have demonstrated that radial actions are conserved in core-forming dwarf-size SIDM haloes.

In future work, we aim to extend our diffusion theory into two dimensions to take into account spatial resonances and deviations from spherical symmetry. Moreover, applying our diffusion theory to massive particles instead of tracers is an important future step towards developing a fully self-consistent statistical theory for the dynamical evolution of self-gravitating systems

Data Availability

The data underlying this article were accessed from the Garpur supercomputer. The derived data generated in this research will be shared on reasonable request to the corresponding author.

Acknowledgements

We thank Volker Springel for giving us access to the AREPO code. We thank Andrew Pontzen for helpful comments and we thank the anonymous referee for suggestions that have helped us clarify some important aspects of our theory. JB and JZ acknowledge support by a Grant of Excellence from the Icelandic Center for Research (Rannís; grant number 173929). The simulations in this paper were carried out on the Garpur supercomputer, a joint project between the University of Iceland and University of Reykjavík with funding from Rannís.

6.7 Appendix A - Numerical calculation of the scale factor

In this section we discuss how we calculate the scale factor $R(t)$ numerically using Eq. (6.9). We start by recalling that Eq. (6.9) has an analytic solution if the evolution is adiabatic. In this case, we can neglect the term proportional to the second derivative of the scale factor to have

$$R^3\mathbf{F}(R\mathbf{r}',t) \approx \mathbf{F}'(\mathbf{r}'). \tag{6.70}$$

Substituting Eqs. (6.3) and (6.5) into Eq. (6.70), we then obtain

$$-R^3\frac{GM_0(1+\varepsilon t)}{(Rr')^2} = -\frac{GM_0}{(r')^2} \tag{6.71}$$

which leads to

$$R(t) = \frac{1}{(1+\varepsilon t)}, \tag{6.72}$$

or in terms of the ratio $\dot{R}/R$, which is the relevant quantity in the linear expansion of the radial action (Eq. 6.20)

$$\frac{\dot{R}}{R} = -\frac{\varepsilon}{\varepsilon+t}. \tag{6.73}$$

This implies that up to first order, the scale factor only depends on time, and there is no residual dependence on the orbital trajectory of individual tracer particles. In this paper, however, we are interested in both the fully adiabatic regime and the regime in which the evolution of radial actions is non-adiabatic. Thus, adopting Eq. (6.72) as the scale factor for all tracer particles might be a bad approximation, in particular when $R/\ddot{R}$ is of the order of the radial period squared. To avoid this issue, we solve Eq. (6.9) numerically along the phase space trajectory of each particle, using a KDK leapfrog algorithm. First, we rewrite Eq. (6.9) for the Kepler case as

$$\ddot{R}R^3r' + R^3\frac{GM(t)}{R^2r'^2} = \frac{GM(0)}{r'^2}, \tag{6.74}$$

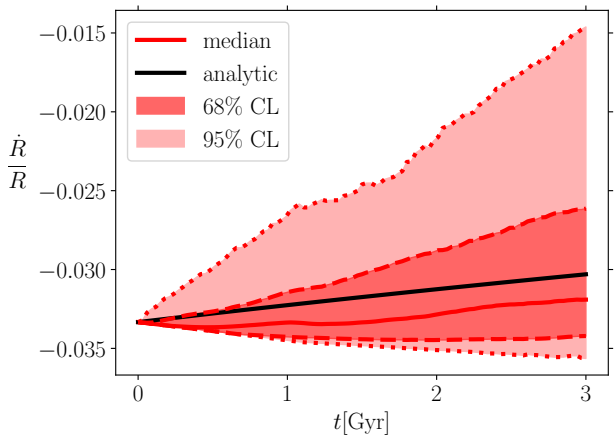

Figure 6.17. Confidence level regions of the distribution of $\dot{R}/R$ versus time, taken from the "10-linear" simulation (see Table 6.1). The solid red line shows the median value, the dark (light) red shaded region covers 68% (95%) CL of the distribution. The analytic solution from Eq. (6.73) is shown as a solid black line.

which using $r = Rr'$ can be written as

$$\ddot{R} = -R\frac{GM(t)}{r^3} + \frac{GM(0)}{r^3} = G(R(t), r(t), t). \tag{6.75}$$

During the simulation, we record phase space coordinates at each time step, as well as the simulation time. We evolve Eq. 6.75 between times t and $t + \Delta t$, where Δt is the time step in the simulation, as follows:

$$\dot{R}\left(t + \frac{\Delta t}{2}\right) = \dot{R}(t) + \frac{\Delta t}{2} G(R(t), r(t), t) \tag{6.76}$$

$$R(t + \Delta t) \quad = R(t) + \Delta t\, \dot{R}\left(t + \frac{\Delta t}{2}\right) \tag{6.77}$$

$$\dot{R}(t + \Delta t) \quad = \dot{R}\left(t + \frac{\Delta t}{2}\right) + \frac{\Delta t}{2} G(R(t + \Delta t), r(t + \Delta t), t + \Delta t)). \tag{6.78}$$

Since Eqs. (6.76−6.78) explicitly depend on the radius of the tracer at a given time, there is now an orbit-dependent part that enters into the solution for $R(t)$ (contrary to Eqs. 6.72 and 6.73). We thus expect that the true values of $R(t)$ will be distributed around the analytic solution given by Eq. (6.72).

Fig. 6.17 shows the median of $\dot{R}/R$ versus time as a red solid line, calculated from the "10-linear" case defined in table 6.1. The 68% (and 95%) confidence level regions are shown as dark (and light) red shaded areas. The analytic approximate solution given

by Eq. (6.73) is shown as a black solid line. Both the median and the analytic solution lie well within the 68% confidence level. Overall, we find that the analytic solution is a reasonable approximation for most particles. A small fraction of particles tends towards significantly (50 per cent) larger values of $\dot{R}/R$ as time progresses, yet the median stays close to the analytic solution at all times. This suggests that in the "10-linear" case, which is the closest to being adiabatic, the analytic solution to Eq. (6.74), i.e. Eq. (6.73), provides a good approximation of the true scale factor.

6.8 Appendix B - Higher order perturbation theory

The formalism presented in Sections 6.3 and 6.4 is based on linear perturbation theory. Here we calculate the radial action evolution to higher orders in perturbation theory, allowing for a slightly more accurate description of radial actions inhabiting the "transition" area between the "linear" and "fringe" regimes.

Eq. (6.18) is a general expansion of the invariant radial action J_r', which equates the difference between J_r' and J_r to an infinite Taylor expansion. In Eq. (6.20) we truncate this series to first order by assuming that the perturbations to the radial action are sufficiently small on average. However, as we have seen throughout Section 6.4, the linear prediction gradually deteriorates in the "transition" regime and is not valid in the "fringe" of the potential. Hereafter, we calculate the second order expansion and investigate how it changes the results of the diffusion formalism.

To calculate J_r to second order in perturbation theory, we first need to use the exact (higher order) version for the energy of a particle (Eq. 6.10), which is (see Peñarrubia 2013)

$$\frac{I}{R^2} = E - \frac{\dot{R}}{R}(\mathbf{r}\cdot\mathbf{v}) + \frac{1}{2}\left(\frac{\ddot{R}}{R} + \left(\frac{\dot{R}}{R}\right)^2\right). \tag{6.79}$$

The invariant action in the time-independent reference frame is

$$J_r' = \int_{r'_{\rm peri}}^{r'_{\rm apo}} dr' \sqrt{2I - 2\tilde{\Phi}(\mathbf{r}') - \frac{L^2}{r'^2}} \tag{6.80}$$

Since I is a true dynamical invariant, changing between coordinate systems and writing I in terms of time-dependent quantities cannot result in I depending on phase space variables. The integral in Eq. (6.80) is thus always defined solely in the static frame, but we are allowed to write the constants and functions appearing in the integral as functions of time-dependent quantities, as long as the integrand depends on the integration variable in the same way. The defining equation relating the potentials in both frames is

$$\tilde{\Phi}(\mathbf{r}') = \frac{1}{2}R\ddot{R}r^2 + R^2\Phi(\mathbf{r},t) \tag{6.81}$$

which we can rewrite in terms of r' as

$$\tilde{\Phi}(\mathbf{r}') = \frac{1}{2}R^3\ddot{R}r'^2 + R^2\Phi(R\mathbf{r}',t) \tag{6.82}$$

If the potential is scale-free, this enables us to write the integrand of Eq. (6.80) in terms of time-dependent quantities while still keeping all dependence on the integration variable explicit and analytic. Note that to do so we have to use $\ddot{R} = 0$ in Eq. (6.82) for consistency, since we originally used this assumption in Eq. (6.9) to obtain scale factors which do not depend on the phase-space trajectory of each particle.

In potentials that are not scale-free, the scale factor cannot, in general, be treated as independent of the phase space trajectory. This implies that writing the potential in the time-independent frame in terms of time-dependent quantities introduces an implicit and thus non-analytic dependence on the integration variable (the radius) through the scale factor. This does not alter the linear correction, provided we take the following (adiabatic) limit.

The key assumption is that even if the potential is not scale-free, we can approximate the scale factor as independent of the particle's trajectory on time-scales which are of the order of the dynamical time. This means that the scale factor, and hence the potential, evolves slowly compared to the particle's orbital period, which is exactly how the adiabatic regime is usually defined. Mathematically, this means that since we evaluate Eq. (6.79) at a fixed time and we treat R in Eq. (6.82) as dependent on time only, we can treat R as a simple scalar when evaluating the integral 6.80.

To expand Eq. (6.80), we first introduce a few definitions. We define

$$f'(r') = 2I - 2\tilde{\Phi}(\mathbf{r}') - \frac{L^2}{r'^2} \tag{6.83}$$

as the square of the integrand in the static frame. Then we define the functions $g'(I,L)$ and $h'(I,L)$ as the peri- and apocentre radii of particles with energy I and angular momentum L, respectively. Next, we define

$$\begin{aligned} f(\mathbf{r}',y,z) = 2E - 2y\,(\mathbf{r}\cdot\mathbf{v}) + r^2\left(z+y^2\right) \\ - R^2 z r'^2 - 2\Phi(R\mathbf{r}',t) - \frac{L^2}{R^2 r'^2}, \end{aligned} \tag{6.84}$$

which is $f = R^2 f'$ with the energy and the potential written in terms of time-dependent quantities. Here we have further introduced $y = \dot{R}/R$ and $z = \ddot{R}/R$ as the variables upon which the Taylor expansion will be made. Finally, we define $g(E,L,y,z)$ and $h(E,L,y,z)$ as $g = Rg'$, $h = Rh'$, which reduce, respectively, to the peri- and apocenter radii in the time-dependent frame when taking the limit $y \to 0$, $z \to 0$. Under the assumption that R is independent of the orbital trajectory, we can write $R\,dr' = dx$ and thus Eq. (6.80) is

$$J'_r = \int_{g(E,L,y,z)}^{h(E,L,y,z)} dx\, k(\mathbf{x},y,z), \tag{6.85}$$

where $k(\mathbf{x},y,z) = \sqrt{f(\mathbf{x},y,z)}$ and

$$\begin{aligned} f(\mathbf{x},y,z) = 2E - 2y\,(\mathbf{r}\cdot\mathbf{v}) + r^2\left(z+y^2\right) \\ - z x^2 - 2\Phi(\mathbf{x},t) - \frac{L^2}{x^2}, \end{aligned} \tag{6.86}$$

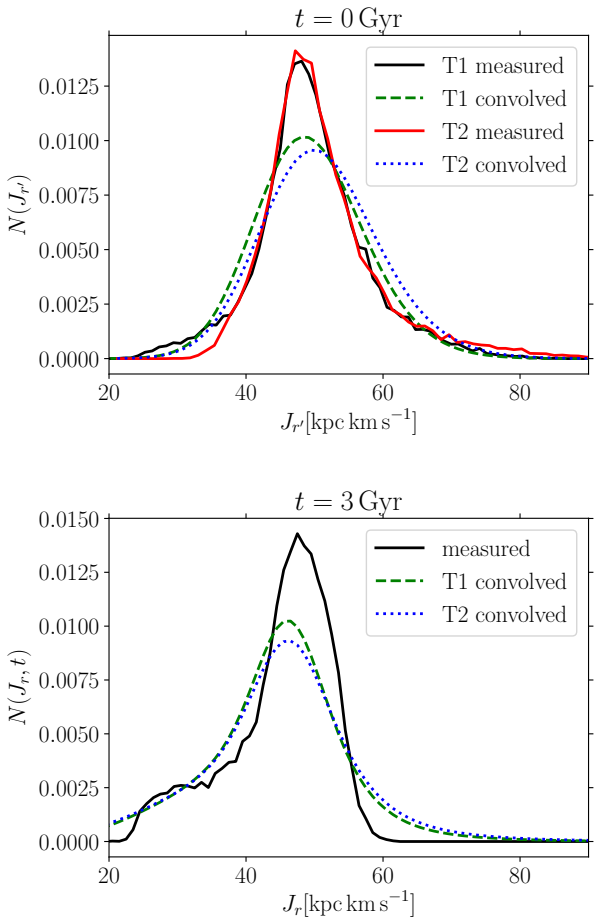

Figure 6.18. Comparison between the first (T1) and second (T2) order Taylor expansions (T1) in the "50-transition" case (see Table 6.1 and Section 6.4). The top panel shows the invariant distribution calculated at $t = 0$. The black and red lines are distributions measured directly by calculating J'_r for each particle and binning the results. The red line is calculated applying the second order Taylor expansion to calculate J'_r, whereas the black line is calculated using the first order expansion. The green dashed line is the result of a Gaussian convolution, with drift and diffusion coefficients calculated from the first order Taylor expansion. The blue dotted line is the result of a convolution using the second order Taylor expansion. On the bottom panel, we see the measured radial action distribution in the time-dependent frame in black, calculated at $t = 3\,\mathrm{Gyr}$. The result of the first (second) order diffusion formalism is shown as a green dashed (blue dotted) line.

It is thus evident that $J_r'|_{y,z\to 0} = J_r$. For the expansion to first order in y, we have to calculate

$$\frac{\partial J_r'}{\partial y} = \int_{g(E,L,y,z)}^{h(E,L,y,z)} dx \frac{\partial k(\mathbf{x},y,z)}{\partial y} + k(h(E,L,y,z),y,z)\frac{\partial h}{\partial y} - k(g(E,L,y,z),y,z)\frac{\partial g}{\partial y}, \tag{6.87}$$

which satisfies the following limit

$$\left.\frac{\partial J_r'}{\partial y}\right|_{y,z\to 0} = -(\mathbf{r}\cdot\mathbf{v})\frac{P(E,L)}{2\pi} \tag{6.88}$$

To continue the expansion to higher orders, we have to calculate

$$\frac{\partial J_r'}{\partial z} = \int_{g(E,L,y,z)}^{h(E,L,y,z)} dx \frac{\partial k(\mathbf{x},y,z)}{\partial z} + k(h(E,L,y,z),y,z)\frac{\partial h}{\partial z} - k(g(E,L,y,z),y,z)\frac{\partial g}{\partial z} \tag{6.89}$$

and

$$\begin{aligned}\frac{\partial^2 J_r'}{\partial y^2} = &\int_{g(E,L,y,z)}^{h(E,L,y,z)} dx \frac{\partial^2 k(\mathbf{x},y,z)}{\partial y^2} \\ &+ 2\frac{\partial k(h(E,L,y,z),y,z)}{\partial y}\frac{\partial h}{\partial y} \\ &- 2\frac{\partial k(g(E,L,y,z),y,z)}{\partial y}\frac{\partial g}{\partial y} \\ &+ \frac{\partial k(h(E,L,y,z),y,z)}{\partial h(E,L,y,z)}\left(\frac{\partial h}{\partial y}\right)^2 \\ &- \frac{\partial k(g(E,L,y,z),y,z)}{\partial g(E,L,y,z)}\left(\frac{\partial g}{\partial y}\right)^2 \\ &+ k(h(E,L,y,z),y,z)\frac{\partial^2 h}{\partial y^2} \\ &- k(g(E,L,y,z),y,z)\frac{\partial^2 g}{\partial y^2}.\end{aligned} \tag{6.90}$$

We can then write to second order

$$J_r' = J_r + A\frac{\dot R}{R} + B\frac{\ddot R}{R} + \frac{1}{2}C\left(\frac{\dot R}{R}\right)^2, \tag{6.91}$$

with

$$A = \left.\frac{\partial J_r'}{\partial y}\right|_{y,z\to 0},\ B = \left.\frac{\partial J_r'}{\partial z}\right|_{y,z\to 0},\ C = \left.\frac{\partial^2 J_r'}{\partial y^2}\right|_{y,z\to 0}. \tag{6.92}$$

We note that in Eq. (6.91), we implicitly assume that $\ddot{R}/R$ is of the same order as $(\dot{R}/R)^2$. In the Kepler example, we can verify this and find from Eq. (6.72) that $\ddot{R}/R = 2\dot{R}/R$ in the adiabatic limit.

In general, the integrals defining the coefficients A, B, C in Eq. (6.91), have to be evaluated numerically, which is a complicated task, especially considering that all we hope to achieve is a slightly better understanding of the evolution of radial action distributions in the "transition" regime of Fig. 6.3. Fortunately, in the case of the Kepler potential, we can calculate the invariant action explicitly as

$$J_r' = \frac{GM_0}{\sqrt{-2I}} - L \tag{6.93}$$

If we Taylor expand the invariant action up to the second order, we find

$$\begin{aligned} J_r = J_r' &+ \frac{GM}{\sqrt{2\mathscr{E}}^3} (\mathbf{r}\cdot\mathbf{v}) \frac{\dot{R}}{R} \\ &- \frac{1}{2} \frac{GM}{\sqrt{2\mathscr{E}}^3} r^2 \left(\left(\frac{\dot{R}}{R}\right)^2 + \frac{\ddot{R}}{R} \right) \\ &- \frac{3}{2} \frac{GM}{\sqrt{2\mathscr{E}}^5} (\mathbf{r}\cdot\mathbf{v})^2 \left(\frac{\dot{R}}{R}\right)^2 . \end{aligned} \tag{6.94}$$

From Eq. (6.94) we see that second order corrections become important at larger energies, i.e., closer to the "fringe". Here we investigate the impact of including second order corrections on the "50-transition" case of Section 6.4 (see Table 6.1).

In Fig. 6.18 we compare the results of the diffusion formalism for the "50-transition" distribution when using first order coefficients (T1) and second order coefficients (T2). The top panel compares the invariant distributions at the beginning of the simulation. The "measured" curves refer to a direct measurement of the distributions, where in order to measure the T1 (T2) distribution we have applied the first order (second order) correction to the action of each individual particle. The "convolved" curves are the result of applying the Gaussian convolution defined in Eq. (6.31), using drift and diffusion coefficients obtained from the first or second order Taylor expansions. Overall, we find that "T1 measured" and "T2 measured" look very similar when compared to each other. The same is true for "T1 convolved" and "T2 convolved", implying that including second order corrections does not improve the accuracy of our formalism.

The effect of including the second order coefficients is mainly to reduce (increase) the tail at the lower (higher) end of the "measured" invariant distribution (top panel of Fig. 6.18). This is to some extent expected, as the second order corrections all increase the value of J_r' (see Eq. 6.94). Comparing the "T1" curves to the "T2" curves, we find no significant improvement of the match between the "measured" and the "convolved" curves. Since the tails of the invariant distribution seem to be quite sensitive to the order of the perturbative expansion, we expect higher orders to have a sizeable impact as well.

On the bottom panel of Fig. 6.18 we show the final result for the distribution of radial actions in the time-dependent frame obtained from Eq. (6.32), using the distributions on the left panel and drift and diffusion coefficients calculated at $t = 3\,\mathrm{Gyr}$. Once more, T1 (T2) refers to the first (second) order Taylor expansion. Evidently, the agreement

between the measured distribution and the result of the diffusion formalism does not improve significantly when using second order coefficients. The most obvious effect is that the second order result has a slightly more extended tail towards larger radial actions, which is, however, hardly significant. The tail on the lower action side is still resolved quite well in both cases. Overall, the diffusion formalism's prediction of the evolution of the '50-transition" distribution does not improve when including second order terms.

We conclude that the evolution of the "50-transition" distribution is is at least partly non-perturbative and cannot be fully captured by the diffusion formalism derived in Section 6.3. Nonetheless, the second order correction still represents a more accurate approximation of the invariant action of individual particles inhabiting the transition regime. We can thus use it to gain additional insight into the non-linear evolution of $N(J_r,t)$ in the "200-transition" simulation.

6.9 Appendix C - An approximate scale factor for generic spherical potentials

For the case of scale-free power law potentials, Peñarrubia (2013) showed that in the adiabatic limit, an approximate solution to Eq. (6.9) is given by Eq. (6.44). In Appendix 6.7, we compare this approximate scale factor to a distribution of numerically calculated scale factors of a set of tracer particles orbiting in a time-dependent Kepler potential; overall, we find a good agreement. For several applications of physical interest, it is desirable to obtain an approximate scale factor for generic spherical potentials which are not scale-free. Under certain assumptions, we can find an approximate solution to Eq. (6.9) as we show in the following. If we assume that the force can be written as

$$\mathbf{F}(\mathbf{r},t) = F(t)\, r^{\beta(r,t)}\, \frac{\mathbf{r}}{r}, \tag{6.95}$$

where $\beta(r)$ is a local power law index, we can find an approximate solution to Eq. (6.9) in the adiabatic limit, which is

$$R_1(t) = \left(\frac{F(r,t)}{F(r,0)}\right)^{-\frac{1}{(3+\beta(r,t))}}, \tag{6.96}$$

where $F(r,t)$ denotes the magnitude of the local time-dependent force. Peñarrubia (2013) shows that Eq. (6.9) can also be integrated, leading to an equation that relates the time-dependent potential to the constant potential in the time-independent frame. An approximate solution to this integrated equation is

$$R_2(t) = \frac{1}{2+\alpha(r,t)} \left(\frac{\Phi(r,t)}{\Phi(r,0)}\right)^{-\frac{1}{2+\alpha(r,t)}}, \tag{6.97}$$

where

$$\alpha(r,t) = \frac{d\log(-\Phi(r,t))}{d\log r}. \tag{6.98}$$

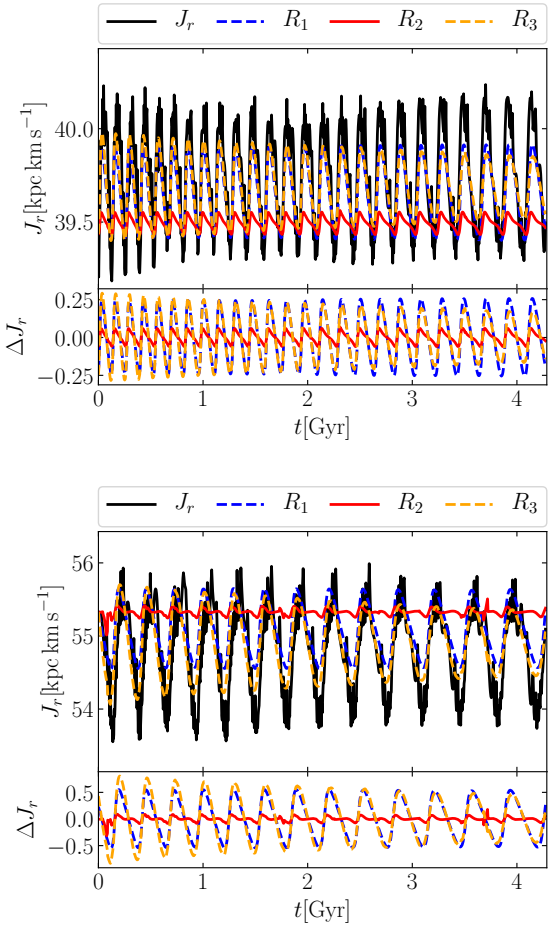

Figure 6.19. The top subpanels show the evolution of the radial action of two tracers in a time-dependent Dehnen potential; see text for details. The black line denotes a direct measurement of J_r. The dashed blue line, the solid red line, and the dashed orange line indicate the first order Taylor expansion of J_r (Eq. (6.20), using different approximations for the scale factor $R(t)$ (Eqs. 6.96, 6.97 and 6.100). In the bottom subpanels, we show the first order corrections ΔJ_r corresponding to each of the three scale factors.

It is worth noting, however, that the potential $\Phi(r,t)$ is only defined up to an integration constant. In most N-body codes, this constant is used to fix the limit of the potential at large radii to $\Phi(r \to \infty, t) = 0$. For the potentials of most self-gravitating systems, e.g. Plummer, Hernquist or NFW potential, this implies $\Phi(r \to 0, t) = \text{const}$, and thus a local power law index of $\alpha(0,t) = 0$, irrespective of the behaviour of the force. However, to make Eq. (6.97) consistent with Eq. (6.96), we demand that $\alpha(0,t) = 1 + \beta(0,t)$, in order to match the case of scale-free potentials. We can recover this behaviour by setting the integration constant to zero when calculating the potential, or equivalently, by defining

$$\Psi(r,t) = \Phi(r,t) - \Phi(0,t). \tag{6.99}$$

Then, we can calculate the scale factor as

$$R_3(t) = \frac{1}{2+\alpha(r)} \left(\frac{\Psi(r,t)}{\Psi(r,0)} \right)^{-\frac{1}{2+\alpha(r)}}, \tag{6.100}$$

where

$$\alpha(r) = \frac{d\log(\Psi(r,t))}{d\log r}. \tag{6.101}$$

When calculating the first order correction term to the radial action in Eq. (6.20), the relevant quantity is not the scale factor itself, but its normalized time derivative $\dot{R}/R$. For cusp-core transformation in a SIDM halo, which we discuss in Section 6.5.2, both the time evolution of the enclosed mass at a given radius and the time-evolution of the potential's logarithmic slope can be significant. This can cause two potential problems. Firstly, the approximate solutions (Eqs. 6.96, 6.97 and 6.100) are derived under the assumption of constant local power law slopes, and it is thus not clear whether they will hold in a potential with a time-dependent shape. Moreover, the time-dependence of the power law slopes differs between Eqs. (6.96), (6.97) and (6.100), both due to the impact of the integration constant on the power law index of the potential, and because in generic spherical potentials, there is not a single one-to-one correspondence between the power law indices of the force and the potential, i.e., $\alpha = 1 + \beta$ is never true for all radii in potentials that are not scale free.

Here we perform a short numerical test to determine whether we can use the approximate scale factors introduced above to describe the evolution of potentials with a time-dependent shape. We follow the orbits of tracers in the potential generated by a time-dependent Dehnen density profile (Dehnen, 1993):

$$\rho(r,t) \propto r^{-\gamma(t)} \left(r + r_s\right)^{\gamma(t)-4} \tag{6.102}$$

We fix the normalization of the density profile by enforcing mass conservation. At each point in time we normalize the density profile such that

$$M(< r_{200}) = 1.43 \times 10^{10} M_\odot, \tag{6.103}$$

where $r_{200} = 50.1\,\text{kpc}$. The scale radius r_s is then fixed by setting $c_{200} = 15$. We model the transition from a cuspy to a cored profile with the following time dependence for

the power law slope:

$$\gamma(t) = 1 - \frac{t}{t_{\max}}, \tag{6.104}$$

where $t_{\max} = 4.3\,\mathrm{Gyr}$ is the total simulation time. We then closely follow the orbits of kinematic tracers and numerically calculate the radial action at each time-step (a similar exercise as in Fig. 6.2). Additionally, we also calculate the first order Taylor expansion of J_r defined in Eq. (6.20) using the scale factors R_1, R_2 and R_3 defined above (Eqs. 6.96, 6.97 and 6.100). To that end, we estimate their time derivatives as

$$\begin{aligned}
\frac{\mathrm{d}R_1}{\mathrm{d}t} &= -\frac{1}{(3+\beta(r,t))} \times \left(\frac{F(r,t)}{F(r,0)}\right)^{-\frac{1}{3+\beta(r,t)}} \\
&\quad \times \left\{ \frac{\dot{F}}{F} - \frac{1}{(3+\beta(r,t)} \ln\left(\frac{F(r,t)}{F(r,0)}\right) \times \frac{\mathrm{d}\beta}{\mathrm{d}t} \right\} \qquad (6.105) \\
\frac{\mathrm{d}R_2}{\mathrm{d}t} &= -\frac{1}{(2+\alpha(r,t))^2} \times \left(\frac{\Phi(r,t)}{\Phi(r,0)}\right)^{-\frac{1}{2+\alpha(r,t)}} \\
&\quad \times \left\{ \left[1 - \frac{1}{2+\alpha(r,t)} \ln\left(\frac{\Phi(r,t)}{\Phi(r,0)}\right)\right] \times \frac{\mathrm{d}\alpha}{\mathrm{d}t} + \frac{\dot{\Phi}}{\Phi} \right\} \qquad (6.106) \\
\frac{\mathrm{d}R_3}{\mathrm{d}t} &= -\frac{1}{(2+\alpha(r,t))^2} \times \left(\frac{\Psi(r,t)}{\Psi(r,0)}\right)^{-\frac{1}{2+\alpha(r,t)}} \\
&\quad \times \left\{ \left[1 - \frac{1}{2+\alpha(r,t)} \ln\left(\frac{\Psi(r,t)}{\Psi(r,0)}\right)\right] \times \frac{\mathrm{d}\alpha}{\mathrm{d}t} + \frac{\dot{\Psi}}{\Psi} \right\}, \qquad (6.107)
\end{aligned}$$

where we have taken into account that the local power law slopes are time-dependent.

Fig. 6.19 shows the evolution of the radial action of two tracers, along with the prediction of the first order Taylor expansion using each of the three scale factors. Apart from some numerical issues that occur when calculating the radial action close to the apo- and pericentre of the particle's orbits, the evolution of J_r for both tracers is accurately described by the Taylor expansion up to linear order, provided we either use R_1 (Eq. 6.96) or R_3 (Eq. 6.100) as the scale factor. In both of these cases, the amplitude and oscillation period of the radial actions are well described by the Taylor expansion. Comparing the linear correction terms directly (lower subpanels of Fig. 6.19) reveals that the corrections obtained when using R_1 and R_3 are almost identical. The result of the Taylor approximation using R_2, however, provides a rather bad description of the evolution of J_r for both tracers. This confirms that the most important effect is the shape evolution of the potential. Given that in the inner part of the potential, where the impact of the modelled core formation is most substantial, there is a one-to-one correspondence between the logarithmic slopes of R_1 and R_3, it is no surprise that they yield very similar values for ΔJ_r. R_2 fails because the logarithmic slope is not related to the radial dependence of the force in the innermost region of the potential.

Since there is no physical preference of one scale factor over the other, we conclude that for the purpose of modeling core formation in a SIDM halo, as done in Section 6.5.2, both R_1 and R_3 can be used to obtain a valid approximation of the evolution of radial actions of particles orbiting in the time-dependent host potential.

Chapter 7

Cusp-core transformation through supernova feedback

This chapter is based on the following article:

SN-driven mechanism of cusp-core transformation: an appraisal

to appear in the Astrophysical Journal

Authors:
Jan D. Burger[1] and Jesús Zavala[1]
[1]centre for Astrophysics and Cosmology, Science Institute, University of Iceland, Dunhagi 5, 107 Reykjavik, Iceland

We present and test an effective model for N-body simulations that aims at mimicking the impact of supernova (SN) feedback on the dark matter (DM) distribution of isolated halos hosting dwarf galaxies. Although the model is physically decoupled from the cosmological history of both the DM halo and the dwarf galaxy, it allows us to study the impact of different macroscopic parameters such as galaxy concentration, feedback energy and energy injection time in the process of SN-driven core formation in a physically clear way. Using our effective model in a suite of N-body simulations of an isolated halo with different SN feedback parameters, we find that whether or not a DM core forms is determined by the total amount of SN feedback energy that is transferred to the DM particles. At a fixed injected energy, the amount of transferred energy is bigger – and the size of the DM core is larger – the faster the energy injection occurs and the more compact the dwarf galaxy is. Analyzing the orbital evolution of kinematic tracers, we demonstrate that a core forms through SN feedback only if the energy injection is impulsive relative to the dynamical timescale of particles in the inner

halo. However, there is no fundamental link between the total amount of injected energy and the injection rate. Consequently, the presence of signatures of impulsive changes of the gravitational potential is not a sufficient condition for dwarf-size halos to have cored density profiles.

7.1 Introduction

The overall successful cold dark matter (CDM) paradigm of structure formation faces potentially severe challenges on small, sub-galactic scales (see Bullock and Boylan-Kolchin 2017 for a review), particularly in its comparison with the properties of the population of satellites in the Milky Way (MW). For instance, inferred circular velocities of the MW satellites are inconsistent with the properties of the subhalos found in MW-like halos identified in dark matter only (DMO) CDM cosmological simulations; a mismatch that is known as the too-big-to fail (TBTF) problem (Boylan-Kolchin et al. 2011, 2012). An updated (extended) version of the TBTF problem highlights that MW satellites reside within DM subhalos with a distribution of circular velocity profiles that is too diverse to be easily reconciled with the narrow distribution seen in (DMO) CDM simulations (Zavala et al., 2019).

A second, possibly related, problem of the DMO CDM model is that it predicts halos to have central density cusps, which is potentially in disagreement with observations of some dwarf galaxies. This so-called cusp-core problem has been put forward e.g. through observations of HI-rotation curves of field dwarfs (e.g. Moore 1994, de Blok et al. 2008, Kuzio de Naray et al. 2008, Read et al. 2019) and through a kinematic analysis of two stellar populations of different age in the dwarf spheroidal galaxies Fornax and Sculptor (Walker and Peñarrubia, 2011). However, there is some debate about whether these observations truly imply cored and isothermal DM density profiles. For instance, Oman et al. (2019) argue that non-circular motion can lead to a large diversity in rotation curves (and hence to a misleading inference of the presence of a cusp or a core), depending on the observer's line of sight (see also Santos-Santos et al. 2020), while Genina et al. (2018) argue that the measurement by Walker and Peñarrubia (2011) could potentially be flawed due to a violation of spherical symmetry in Fornax and Sculptor. While the evidence of the presence of the ubiquity of cored density profiles remains controversial, several mechanisms of cusp-core transformation have been investigated. New DM physics can, for instance, potentially solve both the TBTF and the cusp-core problem simultaneously (see Zavala and Frenk 2019 and references therein). As the kinematics of dwarf galaxies are largely determined by their DM content, DM candidates which deviate from the cold and collision-less CDM hypothesis – and offer solutions to CDM's small scale problems – should be considered a feasible alternative to CDM, provided they manage to evade constraints derived from observations on larger cosmological scales. One such candidate is self-interacting dark matter (SIDM, Spergel and Steinhardt 2000, Yoshida et al. 2000, Vogelsberger et al. 2012, Davé et al. 2001, Colín et al. 2002, Rocha et al. 2013). In SIDM halos, elastic scattering between individual DM particles causes a redistribution of energy from the outer parts of the halo into its center, creating an isothermal core in a fully adiabatic way. SIDM is a feasible DM candidate if the self-interaction cross section evades

current astrophysical constraints (e.g. from elliptical galaxies Peter et al. 2013, clusters Robertson et al. 2017, 2019, and dwarf galaxies Read et al. 2018).

Given that the potential CDM challenges at small scales are only strongly supported from comparisons between observations and results of DMO simulations, the idea that an adequate modelling of baryonic physics may provide a solution is appealing. In particular, supernova feedback has been invoked to solve both the TBTF and cusp-core problems. Navarro et al. (1996a) showed that the sudden removal of an external disk potential can cause initially cuspy halos to form cores in DMO simulations. Gnedin and Zhao (2002) repeated the experiment concluding that a single mass removal equivalent to one entire galaxy is not sufficient to trigger the formation of isothermal cores despite significantly reducing the inner halo density. A decade later, Pontzen and Governato (2012) developed a model to explain core formation through SN feedback as observed in modern hydrodynamical simulations. They show that repeated SN-driven outflows of gas can cause halos to form a core, provided that the energy injection occurs much faster than the typical dynamical timescale in the inner part of the halo. From observations, there is some evidence that star formation histories in dwarf galaxies at the high mass are indeed "bursty" (and the subsequent supernovae feedback cycles impulsive), i.e., they happen on timescales that are comparable to the dynamical time of the galaxy (Kauffmann, 2014), but observations have yet to reach the time resolution needed to resolve the starbust cycle on timescales smaller than the dynamical timescales of the (low-mass) MW dwarf satellite population (Weisz et al., 2014). Based on Pontzen and Governato (2012)'s idea of core formation through periodic SN-driven mass removal, Garrison-Kimmel et al. (2013) slightly altered the model used by Navarro et al. (1996a) and showed that SN feedback alone cannot solve the classical TBTF problem (not taking into account ultra faint dwarfs). A general discussion of the coupling between the cusp-core problem and the TBTF problem was presented in Peñarrubia et al. (2012), where, most notably, the authors calculate the energy required to form a core as a function of halo mass.

From hydrodynamical simulations, there is growing consensus that cores can form as a result of episodic impulsive SN feedback, and that whether or not they do form depends on the ratio of stellar mass to halo mass in a given halo (Di Cintio et al. 2014, Tollet et al. 2016, Chan et al. 2015, Fitts et al. 2017, Lazar et al. 2020). In particular, the universal finding is that there is a limited range of stellar to halo mass ratios for which SN feedback is efficient at forming cores, corresponding for the most part to the mass range of bright/massive dwarfs (see, e.g. Figure 2 of Lazar et al. 2020). Performing high resolution hydrodynamic simulations of isolated halos, Read et al. (2016) find that SN feedback can cause core formation even in ultra faint dwarfs if star formation is sustained for long enough (see also Amorisco et al. 2014 for the relevant impact of star formation histories on the plausibility of forming cores in MW satellites). The origin of this disagreement is not entirely clear, but a possible reason may be that the baryon fraction and/or star formation histories assumed in the controlled/isolated simulations of Read et al. (2016) cannot be realized in a full (CDM) cosmological setting, or vice versa, the latter assumptions are present in nature but cannot yet be obtained/modelled in full (CDM) cosmological simulations. To date, the question of whether or not SN feedback can cause core formation in ultra faint dwarfs remains a subject of debate (see e.g. Genina et al. 2020, Wheeler et al. 2019, Orkney et al. 2021).

Two commonly identified conditions for core formation in the hydrodynamical simulations mentioned above are i) that the energy condition of Peñarrubia et al. (2012) for cusp-core transformation must be fulfilled and ii) that the dwarf galaxies have a sufficiently bursty star formation history, leading to gaseous outflows, and thus fluctuations of the gravitational potential, on rather short timescales. However, Bose et al. (2019a) bring attention to the fact that no cores are formed in dwarfs within the AURIGA and APOSTLE simulations, despite their simulated dwarf galaxies' bursty star formation history. The solution, as presented in Benítez-Llambay et al. (2019), is that whether or not cores form through SN feedback also depends on the numerical value of the star formation threshold adopted in the simulations, a fact that had already been discussed in Pontzen and Governato (2012). Benítez-Llambay et al. (2019) state that larger star formation thresholds lead to higher gas densities in the inner halo, eventually causing the gas to dominate the gravitational potential in the halo's center. As SN feedback couples to DM solely through the gravitational interaction between DM and baryons, the impact of SN feedback is maximized if the contribution of gas to the gravitational potential is significant, as SN feedback would then cause stronger fluctuations in the potential. While Benítez-Llambay et al. (2019) find that there is a range of star formation threshold values which corresponds to a "sweet spot" for core formation, Dutton et al. (2020) argue that the final DM profiles of dwarf-sized halos converge towards having an isothermal core for ever larger numerical star formation thresholds, provided that the right softening length is chosen in the simulation.

Therefore, it appears that whether or not cores are formed in hydrodynamical simulations of dwarf galaxies depends on three different macroscopic conditions:

- Is the injected feedback energy sufficient to transform a cuspy profile into a cored one?
- Are the baryons sufficiently concentrated in the halo center prior to the first starburst?
- Is the galaxy's star formation history bursty enough to cause energy injection on timescales smaller than the dynamical time in the halo center?

The answer to these questions depends on a complex regulation of different properties in hydrodynamical simulations, which are ultimately limited by subgrid (unresolved) physics, most notably, the (effective) star formation threshold and the implementation of SN feedback itself (its local energy/momentum deposition in the surrounding gas elements/particles). In this article, we follow the philosophy of Garrison-Kimmel et al. (2013) and test the impact of changing these three properties directly in a controlled way by introducing an effective model for SN feedback, consisting of an external potential mimicking the gravitational pull of a galaxy, as well as a scheme to periodically inject energy into the halo. In our scheme, the spatial distribution of "supernovae" is determined by the mass distribution corresponding to the external potential and the time over which energy is injected into the halo is free parameter. This simple model allows us to separate the effects of the three key properties connected to core formation, enabling a more transparent physical interpretation than in full hydrodynamical simulations.

In a suite of DMO simulations of an isolated DM halo with identical initial conditions, we look at the effects of including three different (external) Plummer galaxies

of equal mass but with different half-light radii, as well as three different exponential disk galaxies, also of equal mass but with different half-light radii. We furthermore vary the total injected energy, and the time over which the energy is injected. To analyze whether our SN feedback scheme is adiabatic or impulsive compared to the dynamical timescale in the center of the halo, we monitor the evolution of the phase space density of an orbital family; a method that was first introduced in Burger and Zavala (2019).

This article is structured as follows. We present our effective model for supernova feedback in Section 7.2. The simulations we ran to test our model are introduced and discussed in Section 7.3. The main results are presented in Section 7.4 while their implications are discussed in Section 7.5, along with a discussion of some of the approximations made in our model. Finally, we summarize in Section 7.6.

7.2 Method

Our effective model of supernova feedback is built upon the one introduced in Burger and Zavala (2019). Therein, supernova feedback was modelled as a periodic addition and subsequent sudden removal of an external potential located at the halo's center. We modelled the external potential as a Hernquist (1990) sphere and thus, in a coordinate system with origin at the halo's center, including the extra potential amounts to an additional external acceleration for each particle in the simulation:

$$\mathbf{a}_{\text{ext}} = -\frac{GM(t)}{(r+r_s)^2}\frac{\mathbf{r}}{r}, \tag{7.1}$$

where $M(t)$ is a function defining the time-dependent mass corresponding to the external potential and r_s is the scale length of the Hernquist sphere. While this model serves the purpose of approximating a very sudden central – star-burst like – injection of energy into the halo, it is clear that it serves as a rather coarse approximation to the true effects of SN feedback. Nonetheless, it is a time-efficient method to investigate kinematic signatures of tracer particles in a halo which undergoes impulsive core formation - without the need to perform full hydrodynamical simulations. In this article, we aim to expand upon the Burger and Zavala (2019) method with two distinct goals. Firstly, we would like to model the dependence of SN feedback on a set of three relevant parameters that can be connected to both observations and hydrodynamical simulations: galaxy size, total feedback energy, and the time over which the SN energy deposition occurs (hereafter to be called injection time). Secondly, we attempt to couple the SN feedback energy release to the energy budget that is actually available in a galaxy of the modelled size.

In this Section we describe our improved method by first outlining how the external potential is placed within the center of the halo. Then we present our implementation of both a spherically symmetric external Plummer potential and the external potential generated by an axisymmetric flat exponential disk. Finally, we discuss our implementation of SN feedback-like energy input and how we relate it to the total stellar mass and the size of observed dwarf galaxies.

7.2.1 Determination of the halo's center of potential

To place our external, galaxy-mimicking, potential in such a way that its position is consistent with the halo's self-gravitating potential, we need to calculate the halo's center of potential at each time-step. We do this in an iterative manner using a shrinking spheres method. As we are concerned with a single isolated halo in this work, our first step takes into account all DM particles in the simulation[1]. In the first step, the location of the center of potential is estimated as

$$\mathbf{R} = \frac{\sum_{i=1}^{N} \mathbf{r}_i \Phi(\mathbf{r}_i)}{\sum_{i=1}^{N} \Phi(\mathbf{r}_i)}, \tag{7.2}$$

where N is the number of DM particles in the halo, $\mathbf{r}_i$ is each particle's position vector, $\Phi(\mathbf{r}_i)$ is the potential at the particle's position and it is understood that the potential is defined such that

$$\Phi(\mathbf{r}) = -\sum_{i=1}^{N} \frac{Gm_i}{|\mathbf{r}_i - \mathbf{r}|}, \tag{7.3}$$

where m_i is the mass of the i'th DM particle. After calculating an initial estimate of $\mathbf{R}$, we repeat the calculation in Equation (7.2), limiting the sum to particles for which $|\mathbf{r}_i - \mathbf{R}_\mathrm{e}| < r_\mathrm{target}$, where $\mathbf{R}_\mathrm{e}$ is the center of potential estimated in the previous step, and r_target is a target radius which we decrease each iteration. In this paper we use three values $r_\mathrm{target} = 50\,h^{-1}\mathrm{kpc}, 5\,h^{-1}\mathrm{kpc}$ and $0.5\,h^{-1}\mathrm{kpc}$. The result of the last iteration is then taken to be the halo's center of potential. During the last iteration, we also calculate the velocity of the halo's center of potential in analogy to Equation 7.2, restricting the sum as outlined above and replacing the particles' position vectors with their velocities. We note that the implementation of our effective model of SN feedback does not require knowledge of the velocity of the halo's center of potential. However, we do use the velocity to place the halo's center "at rest" when determining the phase space distribution of tracer particles (see Section 7.4).

7.2.2 External potentials

In a more realistic setting, the total amount of energy injected and coupled to the DM due to SN feedback depends on the amount and distribution of stellar mass within the DM halo. This stellar mass, however, can also cause an adiabatic contraction of the DM halo. When investigating core formation due to SN feedback, it is important to take this effect into account, as it can counteract, at least in part, the cusp-core transformation triggered by an impulsive energy injection. We model the net effect of a baryonic component by including an external potential into the simulation, centered at the halo's center of potential. We examine two cases which are of importance in dwarf-sized halos, a spherically symmetric Plummer potential (to mimic a bulge/spheroid) and the potential of an axisymmetric, flat exponential disk.

[1] We note that the method can easily be extended to halos within a larger simulation that have been identified by, for instance, a friends-of-friends algorithm.

An external Plummer sphere

The gravitational effect of a spherically symmetric Plummer profile can be approximated by adding an acceleration to each particle:

$$\mathbf{a}_{i,\mathrm{ext}} = -\frac{GM_{\mathrm{Pl}}}{\left(a^2 + (\mathbf{r}_i - \mathbf{R})^2\right)^{3/2}} (\mathbf{r}_i - \mathbf{R}), \tag{7.4}$$

where M_{Pl} is the total mass of the galaxy modelled by the external potential and a is the scale length of the Plummer sphere. To connect our model to observations of dwarf galaxies, we note that the half-mass radius of the Plummer profile, which is $r_{1/2} \sim 1.3a$, can be compared to observed half-light radii.

An external axisymmetric flat disk

Including an analytic potential to model a disk galaxy is a somewhat more complicated task. In fact, for a vertically extended disk, the calculation of the external acceleration generated by the disk cannot be solved analytically. For that reason, we here model the disk to be infinitely flat, in which case the Poisson equation can be solved up to an integral and the vertical and radial force components can be calculated through numerical differentiation of the potential generated by the disk. We further assume the disk to be homogeneous and thus, the force generated by the disk has no azimuthal component. The flat disk is then fully characterized by its surface density profile

$$\Sigma(R) = \frac{M_{\mathrm{d}}}{2\pi H^2} \exp\left(-\frac{R}{H}\right), \tag{7.5}$$

with M_{d} being the total mass of the disk, H its scale length and R is the polar (cylindrical) radius. The total mass (volume) density is $\rho(r) = \Sigma(R)\delta_D(z)$, and thus we can solve Poisson's equation to obtain the potential:

$$\Phi(R,z) = -\int_0^\infty dk \, \frac{GM_{\mathrm{d}} J_0(kR) \exp(-k|z|)}{\sqrt{1 + (kH)^2}^3}. \tag{7.6}$$

In Equation 7.6, $J_0(kR)$ is a Bessel function of the first kind. We evaluate Equation 7.6 at the start of our simulations on a grid of (R,z) values after testing whether the integral is converged by systematically varying the upper integration limit. During the simulation, the potential at a given point in space can then be calculated through interpolation over the values calculated at the grid points. The force at a certain point in space is then easily obtained as the (directional) numerical derivative of this two-dimensional interpolated potential.

7.2.3 Spatial distribution of supernova locations

In the model of Burger and Zavala (2019), the SN-driven outflow is located in the centre of the halo and modelled as the sudden removal of a spherically symmetric mass distribution. Here we develop this model further by implementing a probabilistic way of assigning (fixed) positions to the superbubbles created by the outflow events

Galaxy (external potential)	Parameter	Units	Description
Plummer/Disk	$\varepsilon_{\rm DM}$	-	Nominal SN energy coupling efficiency
Plummer/Disk	$f_\star$	-	Stellar mass fraction
Plummer	$M_{\rm Pl}$	$h^{-1}M_\odot$	Mass of external Plummer potential
Plummer	$a_{\rm Pl}$	h^{-1}kpc	Scale of external Plummer potential
Disk	$M_{\rm d}$	$h^{-1}M_\odot$	Mass of external disk potential
Disk	H	h^{-1}kpc	Radial scale of external disk potential
Disk	z_0	-	Vertical scale used for SNF distribution in disk
Plummer/Disk	P	h^{-1}Gyr	Period of one SNF cycle
Plummer/Disk	$f_{\rm g}$	-	Fraction of SNF period over which $m(t)$ increases
Plummer/Disk	$a_{\rm SNF}$	h^{-1}kpc	Scale of individual SNF Plummer spheres
Plummer/Disk	$N_{\rm SNF}$	-	Number of SNF spheres in each period

Table 7.1. Parameters of the effective model for supernovae feedback (SNF; see Section 7.2).

across the modelled external potential (galaxy), mimicking the fact that supernovae can occur wherever there are stars. In the following, we will occasionally refer to individual explosions within our model as "supernovae", implemented as outlined below. To be precise, these individual explosions are to be interpreted as approximations of superbubbles – regions devoid of gas that are created through a spatially concentrated series of (actual) supernovae that follows a (local) episode of bursty star formation. When determining the positions of individual "supernovae" we make the approximation that they are more likely to occur in regions of larger stellar mass. In reality, the local density of type II supernovae is strongly correlated with the local star formation rate density. Our model cannot – by construction – capture local bursts in star formation. However, we note that the star formation rate is on average larger in regions where the gas is denser and since in our model "gas" density and "stellar" density have the same functional form, this means that our way of assigning "supernova" locations is consistent with the approximations made within our model. Assuming that the SN density follows the stellar density implies that if we wish to distribute SN feedback probabilistically across the mock galaxy - modeled by either an external Plummer sphere or an external disk - then we can use the (normalized) differential mass profile associated with either density profile as a probability density from which to sample the positions of individual "supernovae". For the Plummer sphere, this means that radii of explosion centers can be obtained from a random number $X \in (0,1)$ via

$$r = \frac{a}{\sqrt{X^{-2/3}-1}}. \tag{7.7}$$

Once the radius is calculated, we determine the exact position vector by randomly selecting the angular position. In the case of a flat external disk, the enclosed mass profile cannot be analytically inverted and our task is thus to find a cylindrical radius for which

$$\frac{X - M(R)}{dM/dR} < \tau R, \tag{7.8}$$

where τ is some numerical threshold. While we are working with the potential of a flat disk, distributing all of the "supernovae" exactly within the x-y plane may lead to an overestimate of the impact of SN feedback in disk galaxies. To prevent that, we assign a vertical offset to each explosion center in a probabilistic manner. To that end, we assume that in the vertical direction the mass is distributed according to a probability density $\propto \cosh(z/z_0)^{-2}$, where z_0 is a scale length. From a random number $X \in (0,1)$, we can then calculate the z-coordinate as

$$z = \frac{z_0}{2} \ln\left(\frac{X}{1-X}\right). \tag{7.9}$$

The azimuthal angle is chosen at random from a uniform distribution. Using Equation 7.7 for a Plummer sphere, and Equations 7.8 and 7.9 for a disk, we can sample any desired number of "supernova" locations. In the limit of a very large number of "supernovae", their cumulative spatial distribution will be closely related to the external potentials introduced in Section 7.2.2 (for a Plummer sphere) and Section 7.2.2 (for a

flat disk)[2].

7.2.4 Implementation of SN feedback

In our model, "supernovae" occur simultaneously and with the same impact at all of the positions determined as outlined in Section 7.2.3. At each SN location, we "cut a hole" into the external potential, by essentially subtracting an external acceleration generated by a Plummer sphere at each SN location. The masses of these Plummer profiles are time-dependent and identical at each location. Thus, the net effect of our effective SN feedback model at a given time t can be written as

$$\mathbf{a}_{\rm SNF}(\mathbf{r}) = \sum_{i=1}^{N_{\rm SNF}} \frac{Gm(t)}{\left(a_{\rm SNF}^2 + (\mathbf{r}-\mathbf{r}_{{\rm SNF},i})^2\right)^{3/2}} (\mathbf{r}-\mathbf{r}_{{\rm SNF},i}). \tag{7.10}$$

where $\mathbf{r}_{{\rm SNF},i}$ denotes the position vector of the i'th SN, $a_{\rm SNF}$ is a Plummer scale that prevents exceedingly short time-steps for particles that get too close to a particular SN and $m(t)$ is the time-dependent mass of the Plummer profile used to model SN feedback. In our model, we increase $m(t)$ linearly over a fixed amount of time and then decrease it linearly over a longer time. This process happens periodically throughout the simulation and for each new period new SN positions are sampled as described in Section 7.2.3. The energy release period P is a free input parameter, as well as the fraction of each period during which $m(t)$ increases until it reaches its maximum value (f_g). The number of "supernovae" during each period, $N_{\rm SNF}$, is a further input parameter of our model. The positions of individual "supernovae" are fixed during the energy release period P. This constitutes a simplifying approximation since the positions of actual superbubbles are time-dependent – due to the streaming motion of the surrounding gas. In principle, assuming static "supernova" locations may introduce an additional degree of asymmetry into the system, an effect that we aim to minimize by choosing a sufficiently large number of "supernovae", $N_{\rm snf}$, during each explosion cycle.

To connect the energy that is injected through our SN feedback model to the total stellar mass modeled by the external Plummer/disk potential, we attempt to connect the maximum value of $m(t)$ to the model parameters in a self-consistent way. To estimate the amount of energy that is injected into the interstellar medium through SN explosions by a galaxy of stellar mass $M_\star$ we follow Equation 6 of Peñarrubia et al. (2012),

$$\Delta E = \frac{M_\star}{\langle m_\star \rangle} \xi(m_\star > 8{\rm M}_\odot) E_{\rm SN} \varepsilon. \tag{7.11}$$

Equation 7.11 renders the available SN feedback energy as a function of the mean stellar mass $\langle m_\star \rangle$, the fraction of stars with a mass larger than $8{\rm M}_\odot$, $\xi(m_\star > 8{\rm M}_\odot)$, the typical energy of one supernova, $E_{\rm SN}$, and the effective coupling efficiency of SN feedback to the interstellar medium, ε. In our model, the coupling efficiency is a free parameter. For the other parameters in Equation 7.11, we follow Peñarrubia et al. (2012) and set $\langle m_\star \rangle = 0.4{\rm M}_\odot$, $\xi(m_\star > 8{\rm M}_\odot) = 0.0037$, and $E_{\rm SN} = 10^{51}$erg, where the first two values originate from using a Kroupa (2002) initial mass function, while $E_{\rm SN}$ is the canonical

[2] In the case of the disk, we are also allowing for the possibility to introduce an additional spread in the vertical direction.

kinetic energy released in SN type II explosions (e.g. Utrobin, V. P. and Chugai, N. N., 2011).

Equation 7.11 estimates the total energy that is returned to the interstellar medium (ISM) by SN feedback. The total increase in the virial energy of the DM halo determines how large the eventual formed core can be (see Peñarrubia et al. 2012). However, how much of the energy that is injected into the interstellar medium couples to the DM is still uncertain. In principle, the coupling efficiency between the energy injected by supernovae and the DM, $\varepsilon_{\rm DM}$, depends on many factors. Since this coupling is purely gravitational, it can depend on the positions of individual superbubbles, their lifetime, and the local DM density. Since we can measure the energy of individual DM particles throughout the simulated time, we can approximately determine how $\varepsilon_{\rm DM}$ depends on the various settings of our model. However, to choose sensible values for the parameters of our effective SN feedback model, we need to obtain a reasonable a priori guess for the total energy that is injected into the interstellar medium (Equation 7.11) by SN feedback. To that end, we assume that the stellar mass $M_\star$ within the external Plummer (disk) baryonic potential is given by $M_\star = f_\star M_{\rm Pl}$ ($M_\star = f_\star M_{\rm d}$). If there are $N_{\rm P}$ SN feedback periods (cycles) during our simulation, the injected energy during each period is $\Delta E_{\rm P} = \Delta E / N_{\rm P}$. The energy associated with each individual "superbubble" is then equal to $E_{\rm SB} = \Delta E_{\rm P} / N_{\rm SNF}$. We then identify this energy with the gravitational binding energy of one of the Plummer spheres that we use to model superbubbles, and use this correspondence to fix the maximum removed "mass":

$$E_{\rm SB} = \frac{3\pi}{32} \frac{G m_{\rm max}^2}{a_{\rm SNF}}. \tag{7.12}$$

We note that Equation 7.12 is an approximation that can only provide a rough order of magnitude estimate for the energy injected by each one of the "supernovae". The central assumption is that the energy of the "supernovae" is equal to the gravitational binding energy of the removed baryonic material of mass $m_{\rm max}$, distributed following a Plummer profile as implemented in our model. For this approximation to be (at least approximately) applicable, the density associated with the removed material needs to be significantly larger than the baryonic density in the surroundings. If this is not the case, the gravitational interaction between the removed material and the surrounding mass contributes significantly to the total energy budget, and the total injected energy will be significantly larger than the nominal energy quoted in Equation 7.11. A way to achieve that Equation 7.12 approximates the injected energy is if the local gas density within the (fully formed) "superbubbles" is negative. While negative gas densities are clearly unphysical, they are not a problem within our model as long as the surrounding DM particles remain gravitationally bound at the end of each supernova cycle – otherwise the DM halo would be artificially disrupted. It is thus advisable to choose $N_{\rm SNF}$ in such a way that the local gas density within fully formed "superbubbles" is negative without unbinding the neighbouring DM particles. Once all of the model parameters are set, we use Equation 7.12 to calculate the maximal mass $m_{\rm max}$ of the individual Plummer spheres that mimic localized starbursts/outflows across the modeled galaxy. The mass

$m(t)$ defined in Equation 7.10 is then given by

$$m(t) = \begin{cases} m_{\max} \frac{t}{f_g P} & t \leq f_g P \\ m_{\max} \frac{P-t}{P(1-f_g)} & t > f_g P \end{cases} , \tag{7.13}$$

where t is the simulation time modulo the period P. Having fixed the external potential of the Plummer (disk) galaxy, as well as the SN feedback associated with either of those potentials, our effective model for SN feedback is now almost fully defined. An explanation of how we determine the number of SN feedback periods from a simulation's parameter file will follow when we discuss the setups of our simulations. Moreover, we discuss the energy that is actually injected into the DM halo, and how it compares to the energy that is nominally injected into the ISM (defined as outlined above), in Section 7.4.4.

A summary of the parameters defining the effective model for SN feedback, along with a brief description of the role of each parameter, can be found in Table 7.1.

7.3 Simulations

Through a series of controlled simulations of an isolated dwarf-sized DM halo, we aim to both test our effective model for SN feedback and investigate the impact of varying three key parameters: ε, f_g, as well as a (in the Plummer case) or H (in the disk case). In this Section, we outline the steps taken to conduct this series of simulations. First, we discuss how we obtain initial conditions of a halo in approximate dynamical equilibrium. Then, we outline how we set up orbital families of tracer particles (akin to how it was done in Burger and Zavala 2019) to track whether changes in the halo's gravitational potential are adiabatic or impulsive. Finally, we discuss the adopted model parameters (see table 7.1) for each of the simulations in our simulation suite.

7.3.1 Initial conditions

For each of our initial conditions, we start by self-consistently sampling a live halo of collision-less DM. We use Eddington (1916)'s formalism to construct a distribution function for a NFW halo (Navarro et al., 1996b) with a DM mass of $M_{200} = 10^{10} h^{-1} \mathrm{M}_\odot$ and an initial concentration of $c_{200} = 13$, where $c_{200} = r_{200}/r_s$, i.e., the ratio between the halo's virial radius and its scale radius; for definiteness we refer to virial quantities for the halo properties corresponding to a virial radius r_{200} enclosing an average density equal to 200 times the critical density of the Universe today. Beyond r_{200}, we exponentially cut off the density profile in order to avoid an infinitely massive halo. Having constructed the distribution function, we draw radii from the analytic differential mass profile of such a halo and then use a rejection sampling method to self-consistently sample the DM particles' velocities from the distribution function (assuming an isotropic velocity dispersion tensor).

After self-consistently sampling position and velocity vectors of 10^7 DM particles, we obtain an N-body representation of an isolated halo in approximate dynamical equilibrium[3]. However, adding an external – Plummer or disk – potential will put the

[3]The equilibrium can never be perfect due to the softened force law used in collisionless N-body simula-

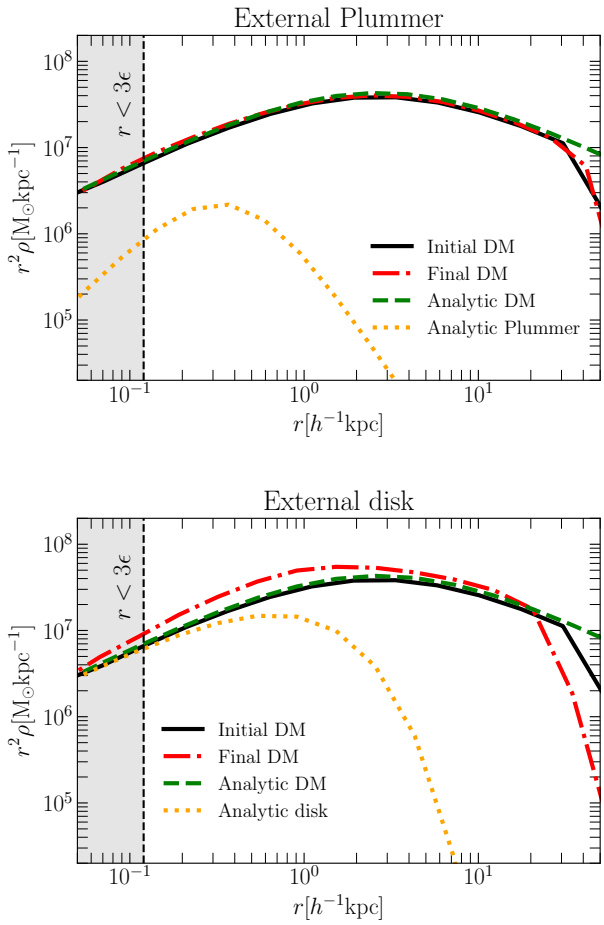

Figure 7.1. Construction of the (equilibrium) initial conditions for our simulations, exemplified on the halo including a benchmark external Plummer (disk) potential on the top (bottom) panel. The grey shaded area indicates the region affected by numerical resolution. The black solid lines show the initial density profile of the simulated DM particles, calculated in spherical shells around the halo's center of potential. This profile overlaps with an analytic NFW profile (green dashed lines) for radii smaller than r_{200}. The dotted orange lines show the (spherically averaged) density profile equivalent to adding the external Plummer (disk) potential to the simulation. The red dash-dotted lines show the contracted DM profiles after relaxing for $1h^{-1}\mathrm{Gyr}$ of simulation time. These final profiles represent equilibrium configurations that are used as initial conditions for subsequent simulations. To highlight the effect of the adiabatic contraction, all profiles are multiplied by r^2.

system out of dynamical equilibrium. In order to start our simulations from controlled initial conditions, we thus let the halo (plus external potential) evolve until the system settles into a new state of dynamical equilibrium, using the code AREPO (Springel, 2010) to calculate the self-gravity between the DM particles.

Since we want to investigate the impact of varying the concentration of the galaxy, we perform several preparation runs – one for each variation of the external potential. For illustration, we chose external potentials that are inspired by the properties of observed MW satellites. In particular, we chose a benchmark setup for both the external Plummer potential and the external exponential disk. The parameters of the benchmark Plummer potential are inspired by observed properties of Fornax. Following McConnachie (2012), Fornax has a stellar mass of $2 \times 10^7 \mathrm{M}_\odot$ and a half-light radius of $r_{1/2} = 0.7\,\mathrm{kpc}$. Moreover, Yuan et al. (2016) show that while gas in Fornax is fully depleted today, the gas mass in Fornax has been on the order of or slightly larger than its stellar mass for a large part of its evolutionary history. With our benchmark Plummer model, we approximate a Fornax-like dwarf by choosing the parameters $M_{\mathrm{Pl}} = 2 \times 10^7 h^{-1}\mathrm{M}_\odot$, $f_\star = 0.5$, and $a = 0.4\,h^{-1}\mathrm{kpc}$ (see Table 7.1 for an explanation). To investigate the effect of a more (less) compact galaxy, we halve (double) the scale radius a. With our benchmark disk model, we aim to roughly capture the gravitational effect of the Small Magellanic Cloud (SMC). Hopkins et al. (2012) models an SMC-like galaxy as a disk with a total baryonic mass of $M_{\mathrm{d}} = 8.9 \times 10^8 \mathrm{M}_\odot$. The stellar disk has a scale length of $H = 0.7\,\mathrm{kpc}$. However, the gaseous component extends much further outwards and is the dominant component in terms of mass (see Table 1 of Hopkins et al. 2012) while our model does not allow for two different disks with different scale lengths. Moreover, Hopkins et al. (2012) assumes a host halo mass which is significantly larger than the mass of our fiducial DM halo. For these reasons, we chose to use slightly different parameters in our benchmark SMC-like disk: $M_{\mathrm{d}} = 4 \times 10^8 h^{-1}\mathrm{M}_\odot$, $f_\star = 0.15$, and $H = 0.7\,h^{-1}\mathrm{kpc}$. Aside from the benchmark model, we also investigate the effect of having a more (less) compact disk by halving (doubling) the disk's scale length. All external potentials are static, meaning that they are added instantaneously at the beginning of the preparation runs.

The initial conditions for our eventual simulations are the final snapshots of the preparation runs conducted for a total time of $1h^{-1}\mathrm{Gyr}$[4], in which the DM particles have had the time to respond to the addition of the external potential and settle into a new dynamical equilibrium.

In Figure 7.1, we show the process of generating the initial conditions of our simulations, exemplified on the two benchmark cases (Plummer on the top panel, disk on the bottom panel). The grey shaded area indicates the radial range in which the measured density profiles are not fully reliable, i.e. radii which are smaller than three times the gravitational softening length of the simulation (Power et al., 2003). As coloured lines we show the density profiles corresponding to the external potentials, as well as the DM density profiles at different times, scaled by r^2. The green dashed lines and the black solid lines are the same across both panels. In green-dashed, we show an analytic NFW profile with the target virial mass and concentration. In black, we

tions.

[4]For a NFW profile with $c_{200} = 13$, this corresponds to the orbital period of a particle on a circular orbit with radius $r_{\mathrm{circ}} \approx 7\,h^{-1}\mathrm{kpc}$.

show the halo's initial DM density profile – before adding the external potential. Over a large range of radii those two lines coincide, but at radii larger than r_{200} we see the exponential cutoff included to have a numerical solution to Eddington's equation. The spherically averaged density profile corresponding to the added external Plummer (disk) potential is shown as an orange dotted line on the top (right) bottom. Finally, the red dash-dotted lines show the DM density profiles measured after letting the system (DM halo plus external galactic potentials) relax for $1h^{-1}$Gyr. When adding the external Plummer potential, we find that the final and the initial DM density profiles are almost exactly identical: the relatively small additional mass causes no significant contraction in the DM. In the case of an external disk potential, however, the mass of the galaxy causes a significant contraction in the DM density profile. We thus anticipate that the gravitational pull of the external disk can be of significant importance when it comes to whether (and how quickly) the DM cusp can be restored after impulsive removal of baryonic mass from the halo's center. To ensure that the final, adiabatically contracted DM density profile is fully determined by the structural parameters of the external potential – and does not depend on the rate at which the external potential is added – we have conducted one additional preparation run. In this run, instead of instantaneously adding the benchmark external disk potential, the "mass" of the external disk is linearly increased over $600\,h^{-1}$Myr, before we let the system relax for another $400\,h^{-1}$Myr. We found no obvious differences between the final DM density profile in this additional preparation run and the final DM density profile shown on the right panel of Figure 7.1.

7.3.2 Orbital Families and Explosion Times

In order to track whether the sizes of cores formed through SN feedback correlate with how impulsive our implementation of SN feedback is, we set up orbital families as in Burger and Zavala (2019) to investigate how they respond to the changes in potential. Depending on whether the orbital family remains united or splits into several families of orbits, we can tell whether the implemented SN feedback was adiabatic or impulsive.

To the initial conditions described in Section 7.3.1 for the Plummer spheres, we add a family of orbits as described in Section 4.1 of Burger and Zavala (2019). We sample 2000 tracers with pericenter radii of $r_{\rm peri} = 0.5 \pm 0.05\,h^{-1}$kpc and apocenter radii of $r_{\rm apo} = 2 \pm 0.05\,h^{-1}$kpc. Note that in a spherically symmetric potential, this is equivalent to sampling orbits with similar energies and angular momenta. Since in a fixed spherically symmetric potential the radial action is only a function of energy and angular momentum, asking whether or not an orbital family of tracers remains united (and thus whether the SN feedback is adiabatic or impulsive) is equivalent to asking whether or not radial actions are conserved.

Setting up the orbital family works slightly differently for the initial conditions corresponding to the three external disk potentials. Since these potentials are axisymmetric, we have to restrict the orbits of the tracers to be within the plane of the disk. Hence, instead of $\Phi(r)$ we now consider the in-plane potential $\Phi(R,0)$, and we initialize all tracers in-plane and with no vertical velocity component ($z = 0, v_z = 0$). We note, however, that a slight deviation from cylindrical symmetry throughout the simulation can cause perturbations to the orbits that may cause the orbital families to diffuse, irrespective of whether these perturbations occur on adiabatic or impulsive timescales.

Galaxy type (1)	Parameter impact (2)	$\varepsilon_{\rm DM}$ (3)	$f_{\rm g}$ (4)	a h^{-1}kpc (5)	H h^{-1}kpc (6)
Plummer	adiabatic	0.01	0.33	0.8	-
	benchmark	0.05	0.017	0.4	-
	impulsive	0.4	0.0017	0.2	-
Disk	adiabatic	0.01	0.46	-	1.4
	benchmark	0.05	0.023	-	0.7
	impulsive	0.4	0.0023	-	0.35

Table 7.2. Parameters that vary between simulations: amount of injected energy ($\varepsilon_{\rm DM}$, column 3), fraction of explosion cycle over which the energy is injected ($f_{\rm g}$, column 4) and concentration of the external potential ($a_{\rm Pl}$ for the Plummer sphere, column 5, and H for the exponential disk, column 6). For each parameter, we investigate a benchmark value, as well as a value that should cause a more adiabatic (impulsive) SN feedback.

With that in mind, we will see if signatures of impulsive SN feedback will still be apparent in the case of an external disk potential, but we do not require that the orbital family stay fully united in order to classify a particular setup of the effective model for SN feedback as adiabatic.

In our simulation suite, we also aim to investigate the impact of changing the time over which the energy from SN feedback is injected into the DM distribution. To enable a comparison between the different external potentials, we express the injection times as fixed fractions of the radial period of a particle which is part of the orbital family. The radial periods $\mathscr{P}_r$ for the Plummer potentials and $\mathscr{P}_R$ for the disk potentials are

$$\mathscr{P}_r = \int_{r_{\rm peri}}^{r_{\rm apo}} \frac{dr}{v_r} \qquad \mathscr{P}_R = \int_{R_{\rm peri}}^{R_{\rm apo}} \frac{dR}{v_R}. \tag{7.14}$$

Setting up our simulations, we then chose growth fractions $f_{\rm g}$ such that the mass $m(t)$ grows as outlined in Section 7.2.4 for a time which equals 1 per cent, 10 per cent, or 200 per cent of the radial period of a particle with pericenter radius $r_{\rm peri} = 0.5\,h^{-1}$kpc and apocenter radius $r_{\rm apo} = 2\,h^{-1}$kpc.

7.3.3 Simulation Settings

Having generated all the required initial conditions, we now look to run a simulation suite in order to test the effective model for SN feedback and to investigate how changing the total energy input, the injection time, and the concentration of the external potentials affects the final DM density profiles, as well as the orbital family. Table 7.2 shows the different numerical values that we adopt for the parameters regulating energy input, injection time, and concentration. For each parameter, we define one benchmark value, as well as one value which should make SN feedback more adiabatic and one value which should make SN feedback more impulsive. We investigate each possible combination of these parameters in order to determine which one of them

causes larger cores in the DM profiles and whether core size relates directly to how adiabatic / impulsive the change in gravitational potential induced by our effective model for SN feedback is. This means that in total we run $3^3 = 27$ simulations for each of the two galaxy-like potentials.

All of the other model parameters (see Table 7.1) are fixed to benchmark values in all simulations. Let us here briefly introduce and discuss their numerical values:

- $M_{\rm Pl} = 2 \times 10^7\, h^{-1}{\rm M}_\odot$ and $M_{\rm d} = 4 \times 10^8\, h^{-1}{\rm M}_\odot$ as discussed in Section 7.3.1.
- $f_\star$ is 0.5 and 0.15 for the Plummer sphere and disk, respectively. This choice is to mimic the effect of a Fornax-like dwarf in the former case and an SMC-like galaxy in the latter case (see Section 7.3.1).
- $a_{\rm SNF}$ is set to $10\,h^{-1}$pc. This sets the scale of individual "superbubbles" to be smaller than the gravitational softening in our simulation, assuring that the energy injection is effectively point-like and thus that the local density contrast introduced by the "superbubbles" is large (see relevant discussion in Section 7.2.4).
- z_0 is always determined as $0.2H$, in agreement with Hopkins et al. (2012).
- For definiteness, P is set to $0.6h^{-1}$Gyr in all simulations. From the period, we can fix the number of explosion periods N_P in all simulations. Before the first wave of explosions, we wait for one period in order to monitor whether the orbital family remains united in the absence of SN feedback. Additionally, we want the system to relax at the end of the simulations, in order for the final DM density profiles to not be affected directly by the gravitational impact of the explosion centers. To that end, no explosions are implemented during the last $1\,h^{-1}$Gyr of the simulations. We run all simulations for a total time of $4\,h^{-1}$Gyr, which implies that the number of explosion cycles is $N_P = 3/0.6 - 1 = 4$. SN feedback is thus distributed over four explosion cycles of equal duration, with $N_{\rm SNF} = 20$ explosion centers during each cycle.

As stated in Section 7.3.1, we use the AREPO code to determine the self-gravity and the time evolution of the isolated systems. For collisionless particles, AREPO uses adaptive time-steps. The adopted time-step criterion is based on the softening lengths of individual particles (see Equation 34 of Springel 2005). Gravitational forces between different simulation particles are calculated using a hierarchical multipole expansion. A relative cell opening criterion is used (see Equation 18 in Springel 2005). In all our simulations and preparation runs, we choose softening lengths of $40\,h^{-1}$pc ($1h^{-1}$pc) for the DM (tracer) particles and an accuracy parameter for the cell opening criterion $\alpha = 0.0005$.

The calculation of each particle's time-step is based on its total acceleration, taking into account both self-gravity and the external forces generated by our effective model of supernova feedback (see Section 7.2). In our model, the external accelerations can change rapidly over a short time and as a consequence, the time-steps of particles that are near a "supernova location" may occasionally be too long just before the "explosion", and thus, their acceleration may not be updated fast enough. To verify that this does not significantly affect our results, we have repeated one of our simulations (the Plummer

run with $\varepsilon = 0.05$, $a = 0.2\,h^{-1}$kpc, and $f_g = 0.017$), but this time fixing the time-step of all DM particles to the minimum value reached in the run with adaptive time-steps. We have found that both the size of the final core and the time evolution of the DM density profile are in good agreement between the two simulations.

7.4 Results

First, we present examples of runs in which a core has formed and compare them to cases in which the density profile remains cuspy. Subsequently, we discuss how the symmetry of the system affects the kinematics of the orbital family of tracers. Then, we compare the final profiles of all simulations to discuss the impact of changing the nominal energy coupling, the injection time and the concentration of the external potentials. Thereafter, we compare the effective change in energy of the DM particles to the nominally injected energy (see Equations 7.11 and 7.12). Based on this comparison, we discuss the accuracy of our a priori guess for the amount of energy that is injected into the ISM (based on Equation 7.12), as well as how the (gravitational) coupling of the injected energy to the DM depends on the energy injection time and the size and shape of the external galaxy. Finally, we briefly discuss whether impulsive energy injection is a necessary condition for core formation through SN feedback.

7.4.1 Cored vs. cuspy profiles

Here we compare the evolution of the DM density profiles, as well as the final phase space distribution of the orbital family, between a simulation in which the DM halo forms a core and a simulation in which it retains its cusp. We focus separately on the cases of an external Plummer sphere and exponential disk.

Feedback from a Plummer sphere

In Figure 7.2, we compare the results of two different simulations including an external Plummer potential, with differently regulated SN feedback. The upper two panels show results of the simulation in which all of the parameters introduced in Table 7.2 lean towards an adiabatic change in the potential. The left panel shows the evolution of the DM density profile after each $1\,h^{-1}$Gyr of simulation time. The grey shaded area denotes the range in which the initial DM profile cannot be considered stable according to the Power et al. (2003) stability criterion. For radii that lie outside this range, however, we hardly detect any evolution in the DM density profile. The halo retains its cusp with our SN feedback model having no significant impact on the DM distribution. In agreement with that, the final phase space distribution of the orbital family of tracers (upper right panel) remains united by the end of the simulation. Relative to the initial phase space distribution – which is distributed very closely around the black dashed line – we detect a non-negligible diffusion of the orbital family, which is due to the fact that the gravitational potential is in fact changing (slowly) with time. However, we see no signatures of an impulsive change in the gravitational potential, i.e., neither does the orbital family split up, nor do the orbits expand to larger radii on average (see Burger and Zavala 2019 for a more detailed analysis).

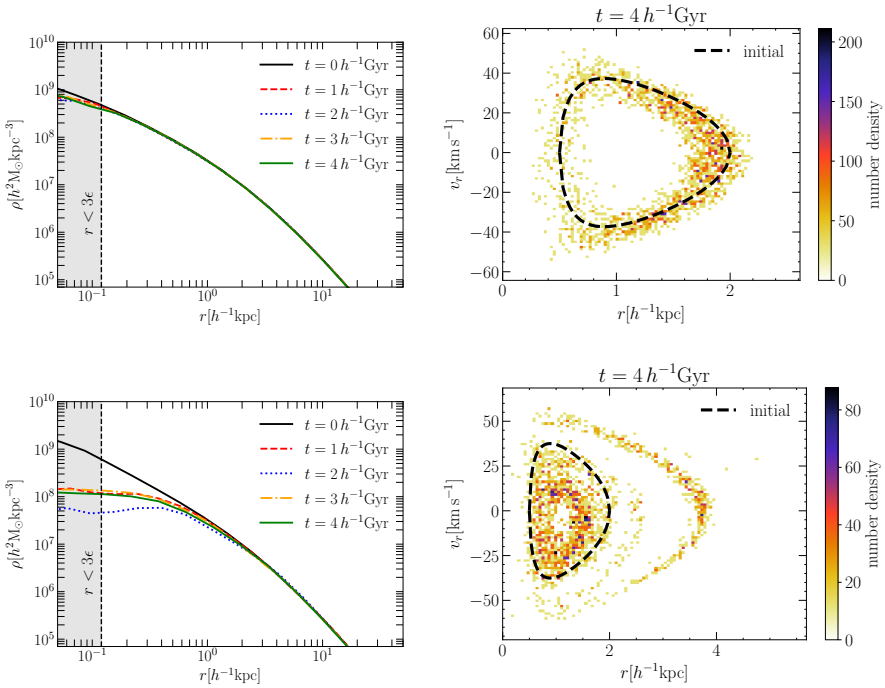

Figure 7.2. The results of two DMO simulations including an external Plummer sphere potential are shown. In the left column, we show the measured DM density profiles at different simulation times in intervals of $1\,h^{-1}$Gyr. *The grey shaded area indicates the region affected by numerical resolution. In the right column, we show the final phase space density distribution of tracer particles that were initially set up as an orbital family. The black dashed lines show the initial "central" phase space trajectory of the orbital family. The upper panels correspond to the simulation in which* $\epsilon = 0.01$ *(adiabatic),* $f_\mathrm{g} = 0.33$ *(adiabatic), and* $a = 0.8\,h^{-1}$kpc *(adiabatic). The lower panels correspond to the simulation with* $\epsilon = 0.05$ *(benchmark),* $f_\mathrm{g} = 0.017$ *(benchmark), and* $a = 0.2\,h^{-1}$kpc *(impulsive).*

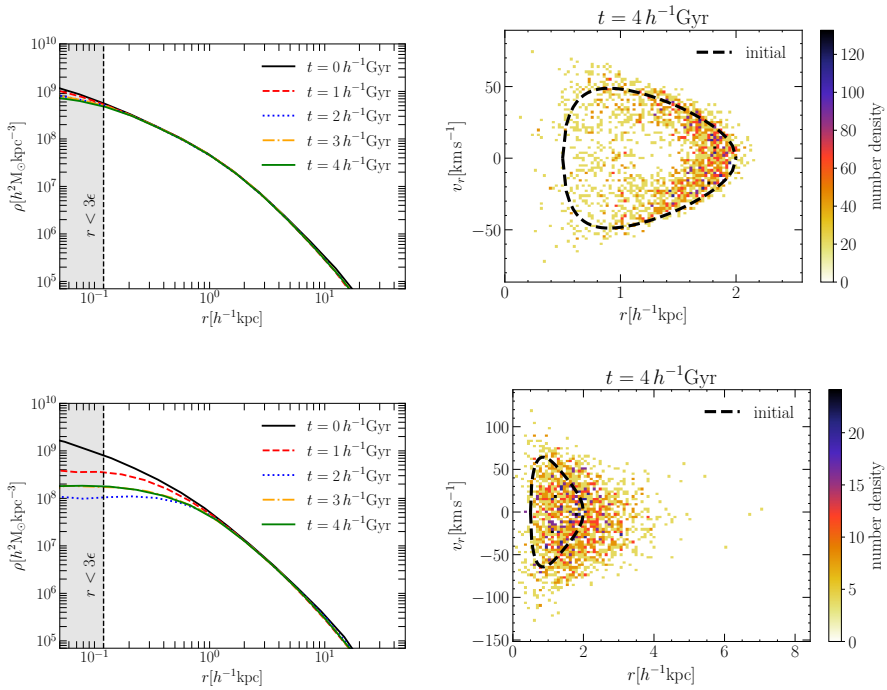

Figure 7.3. As Figure 7.2 but for an external disk potential. The upper panels correspond to the simulation in which $\epsilon = 0.01$ (adiabatic), $f_\mathrm{g} = 0.46$ (adiabatic), and $H = 1.4\,h^{-1}$kpc (adiabatic). The lower panels correspond to the simulation with $\epsilon = 0.05$ (benchmark), $f_\mathrm{g} = 0.023$ (benchmark), and $H = 0.35\,h^{-1}$kpc (impulsive).

The lower panels of Figure 7.2 correspond to a simulation in which the Plummer sphere is more compact (see Table 7.2), whereas the injection time and the amount of injected energy assume their benchmark values. As we can see, the results of this simulation are vastly different from those in the upper panel. In the lower left panel we see how the halo forms a core of size $\sim 1\,h^{-1}$kpc already during the first explosion cycle, and how it retains this core until the end of the simulation. The final phase space distribution of the orbital family of tracers (lower right panel) shows clear signs of an impulsive change in the gravitational potential. The orbital family has split up into several shells and the radial range occupied by the tracers has expanded to significantly larger radii compared to the initial distribution, which is roughly given by the black dashed line. This implies that radial actions are not approximately conserved throughout the simulation (as they are in the upper right panel).

Figure 7.2 suggests that there is a link between whether or not periodic SN feedback-like energy injection is a feasible core formation mechanism and whether the induced change in the central gravitational potential is adiabatic or impulsive.

Feedback from an exponential disk

Figure 7.3 shows the same kind of comparison as Figure 7.2, but for an external disk potential. In the upper panels we show the results of the simulation in which all of the relevant parameters in Table 7.2 correspond to an adiabatic configuration. Just as in the Plummer case, we hardly see any evolution in the DM density profile beyond the resolved radii. The final phase space density of the orbital family (upper right panel) does, however, exhibit a degree of orbital diffusion that clearly exceeds the diffusion observed in the upper right panel of Figure 7.2. Still, the range of radii covered by the tracers' orbits does not expand much further than the original configuration, which means that there is no net migration outward. We attribute the additional diffusion in the orbital family to the additional challenges of preserving cylindrical symmetry. In fact, the setup of the orbital family relies on a potential which has been measured in the $x-y$ plane. Thus, it should in theory be applicable only to particles whose orbits are confined within the disk plane at all times. However, the slightest perturbation into the vertical direction (notably, our implementation of SN feedback can introduce these) can cause the tracers to obtain a non-zero vertical velocity. The tracers' plane of motion then changes, and the in-plane potential (measured as a function of polar radius) no longer determines the tracers' orbits. This deviation from cylindrical symmetry is thus cause of an additional diffusion of the orbital family.

The lower panels of Figure 7.3, shows the case in which the disk scale length is a factor of two smaller than the benchmark value whereas the injection time and the amount of injected energy take their benchmark values. As in the lower panels of Figure 7.2, we can now observe the formation of a $\sim 1\,h^{-1}$kpc core in the DM halo by tracking the evolution of its density profile. However, in this case it takes at least two explosion cycles for the core to fully form, indicating that core formation is slightly less efficient for this disk configuration than for the Plummer sphere case. Moreover, we observe a clear difference between the "impulsive" simulation that includes the external disk potential and the one including the external Plummer potential when looking at the final phase space distribution of the tracers that were initially part of one orbital family. While the range of radii covered by the tracers expands roughly by the same amount, we do not observe any emergent shell-like patterns in phase space, i.e., a split into several orbital families, in the simulation including an external disk potential. Instead, we find that the final phase space distribution is essentially phase-mixed, indicating that significant diffusion has occurred. As in the adiabatic case, we attribute this to the growing impact of deviations from cylindrical symmetry that accumulate throughout the simulation. We take a closer look at the role of symmetry in Section 7.4.2. We thus conclude that while radial migration outwards is a clear signature of impulsive changes to the underlying gravitational potential, shell-like structures, as seen in Figure 7.2, are only relatively long-lived (and thus apparent at the end of our simulations) if the potential's underlying spatial symmetry is closely preserved as the potential changes.

7.4.2 The role of symmetry

The results of Sections 7.4.1 and 7.4.1 suggest that the symmetry of the external "baryonic" potential – and in turn the distribution of individual "supernovae" – affects the final phase space structure of orbital families. In particular, we find that impulsive

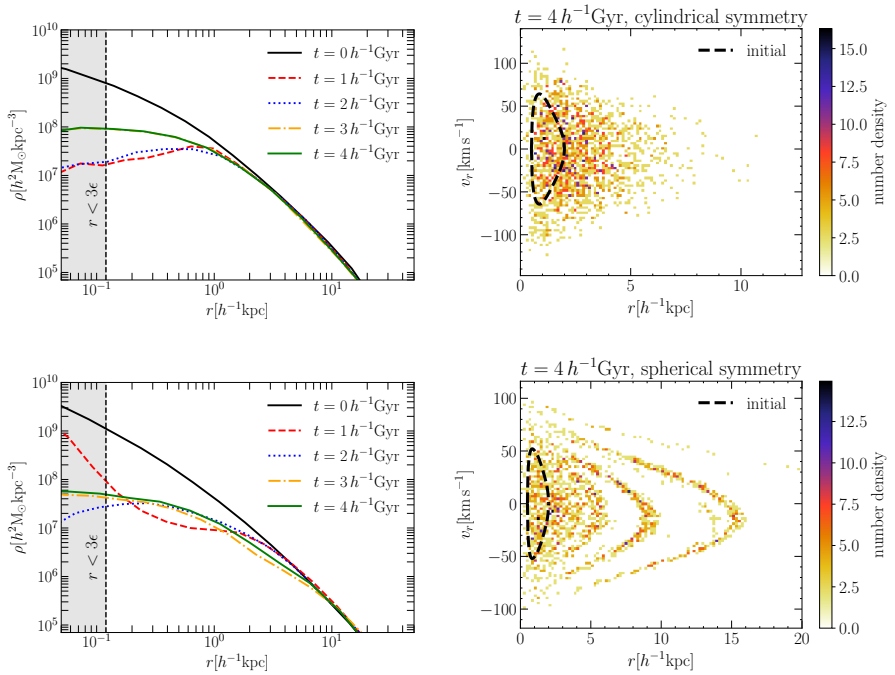

Figure 7.4. Comparison of the evolution of the DM density profile (left column) and the final phase space distribution of an orbital family of kinematic tracers (right column) between two simulations with external potentials that correspond to identical enclosed mass profiles. The top panels show results from the run including an external disk potential with $\epsilon = 0.4$ (impulsive), $f_g = 0.46$ (adiabatic), and $H = 0.35\,h^{-1}kpc$ (impulsive). The bottom panels show results from a simulation in which the "baryonic" enclosed mass profile is idential, but generated by a spherically symmetric mass distribution. The distribution of individual "supernovae" is sampled from the spherically symmetric "baryonic" mass distribution. Notice that while an extended core forms in both cases, the final phase space distribution of the kinematic tracers is distinctly different.

SN feedback in simulations with a spherically symmetric external potential gives rise to shell-like features in the phase space of kinematic tracers that initially belonged to the same orbital family; these features are long-lived and remain evident at the end of our simulations (see Figure 7.2). However, such long-lasting shell-like features are not present at the end of our impulsive SN feedback simulations with a disk-like external potential (see Figure 7.3).

In Section 7.4.1 we stated that this divergent behaviour of the kinematic tracers can likely be explained by the difference in spatial symmetry. However, there is also a significant mass difference between the Plummer spheres and the disks. To verify that the decisive factor is symmetry – and not the mass of the external potential – we performed an additional set of simulations. For this set, the setup of the external potentials is such that the spherically averaged mass profiles associated with them are identical to the mass profiles in the corresponding runs with external disk potentials, i.e.,

$$M(r) = M_{\rm d}\left\{1 - \left(1 + \frac{r}{H}\right)\exp\left(-\frac{r}{H}\right)\right\}, \tag{7.15}$$

where $M_{\rm d}$ is the total mass and H is the scale length of the equivalent disk potential. The radii of individual "supernovae" are randomly sampled from the normalized mass profile. Since the external potential is now spherically symmetric, the disk height parameter is superfluous. All other parameters are kept as in Table 7.2.

Across simulations, we find that a spherically symmetric external potential induces a qualitatively different contraction of the halos' density profiles. While an axisymmetric disk potential leads to shallower central slopes (see Figure 7.1), an equivalent spherically symmetric potential gives rise to steeper density profiles of all simulated DM halos, resulting in deeper potential wells, and thus, requiring more energy to unbind the cusps. As a consequence, we find that in such (spherical) cases, no cores form in runs in which $\varepsilon = 0.01$ or $\varepsilon = 0.05$. Only for the largest choice of the energy coupling parameter, $\varepsilon = 0.4$, do cores form. This is the case we choose to make the comparison between the axisymmetric and spherical potentials. The top row of Figure 7.4 shows results from the run with an external disk potential, using the parameters $\varepsilon = 0.4$, $f_g = 0.46$, and $H = 0.35\,h^{-1}$kpc. The bottom row shows results of the corresponding simulation including a spherically symmetric external potential and a spherical distribution of "supernovae". In the left column, we show the evolution of the spherically averaged DM density profiles, while the final phase space distribution of the orbital family of kinematic tracers is shown in the right column. Although extended constant density cores form in both cases, it is evident that core formation is slower in the spherically symmetric case – a strong cusp-restoring contraction effect due to the external potential can be observed after $1\,h^{-1}$Gyr. Nevertheless, the final core is somewhat larger in the spherically symmetric case. The final phase space distribution of the kinematic tracers is remarkably different between the two simulations. In the axysimmetric case, we observe a considerable radial expansion on average, in line with the radial expansion of the DM particles. Moreover, the final distribution in radial phase space is largely featureless, i.e. completely phase-mixed. In the spherical case, the radial expansion is accompanied by the emergence of prominent shell-like structures. We thus conclude that the divergent behaviour of the kinematic tracers between Figures 7.2 and 7.3 is due to the difference in spatial symmetry – and not due to the difference in the baryonic

mass.

The reason for this difference is as follows. Initially, all particles are on orbits defined by nearly identical actions (J_r in the case of a spherically symmetric external potential and J_R in the case of a disk), but with different orbital phases. As outlined in Pontzen and Governato (2012) and Burger and Zavala (2019), a sudden change in the gravitational potential causes a change in the energies of kinematic tracers that depends on their respective orbital phases. Thus, impulsive mass removal, as in our effective model of SN feedback, turns the original orbital family into a particle distribution which is not phase-mixed. The shell-like features seen in the bottom right panel of Figure 7.4 are signatures of early-stage phase mixing (see e.g. Binney and Tremaine 2008). As long as the underlying symmetry is preserved, i.e, J_r or J_R are integrals of motion for individual tracers, phase mixing progresses relatively slowly. Deviations from spherical (or cylindrical) symmetry, however, can cause orbital diffusion along resonant directions. As a result, actions associated with the broken symmetry are no longer integrals of motion (see Pontzen et al. 2015) and phase mixing progresses much faster[5]. As we mentioned in Section 7.4.1, the cylindrical symmetry in our runs with an external disk potential is only exact in the $x-y$ plane – and is easily broken by the distribution of "supernovae" during a given explosion cycle. As a result, tracers acquire non-zero vertical velocities and their plane of motion becomes tilted with respect to the disk plane. At that point J_R is no longer an integral of motion and resonant diffusion can occur. Therefore, the divergent behaviour of the tracers in Figures 7.2, 7.3, and 7.4 is due to the difference in both the spatial symmetry of the external potential and the spatial distribution of individual "supernovae".

7.4.3 Net SN feedback impact on the inner DM distribution

Figure 7.5 shows a comparison of the results of all the simulations described in Table 7.2. Following Santos-Santos et al. (2020, see also Oman et al. 2016b), we classify halo density profiles with a single number by their galaxy's circular velocity curves through the ratio $v_{\rm fid}/v_{\rm max}$ where $v_{\rm fid}$ is the circular velocity at a radius $r_{\rm fid} = 2(v_{\rm max}/70\,{\rm km\,s^{-1}})$kpc and $v_{\rm max}$ is the maximal circular velocity. As demonstrated in Santos-Santos et al. (2020), smaller ratios between the two velocities correspond to more cored DM density profiles. The value for a NFW profile is typically $v_{\rm fid}/v_{\rm max} \sim 0.7$. We note however that there is no absolute correspondence between the ratio $v_{\rm fid}/v_{\rm max}$ and how cored the DM halo's density profile is since this ratio depends on halo concentration (for a fixed $v_{\rm max}$), and the precise shape of the profile (e.g. a Hernquist profile, albeit cuspy, has slightly different values of $v_{\rm fid}/v_{\rm max}$). More importantly, the baryonic galaxy can have a major impact on this value in several ways: i) massive and concentrated galaxies contribute significantly to the total circular velocity curve in the inner region and can contract the DM halo, making it cuspier in the center (and thus increasing its contribution to the circular velocity curve); ii) SN-driven outflows expel gas from the center and redistribute the DM from the inside out, reducing the value of $v_{\rm fid}/v_{\rm max}$, which is the focus of this work. In our simulations, the mass of the external potential is the same across all simulations of a given type (Plummer or disk) and thus, since the scale length of

[5] In such cases, these signatures can still be observed in the immediate aftermath of an impulsive energy injection. However, observing them requires a more in-depth analysis (see Burger et. al., in prep.)

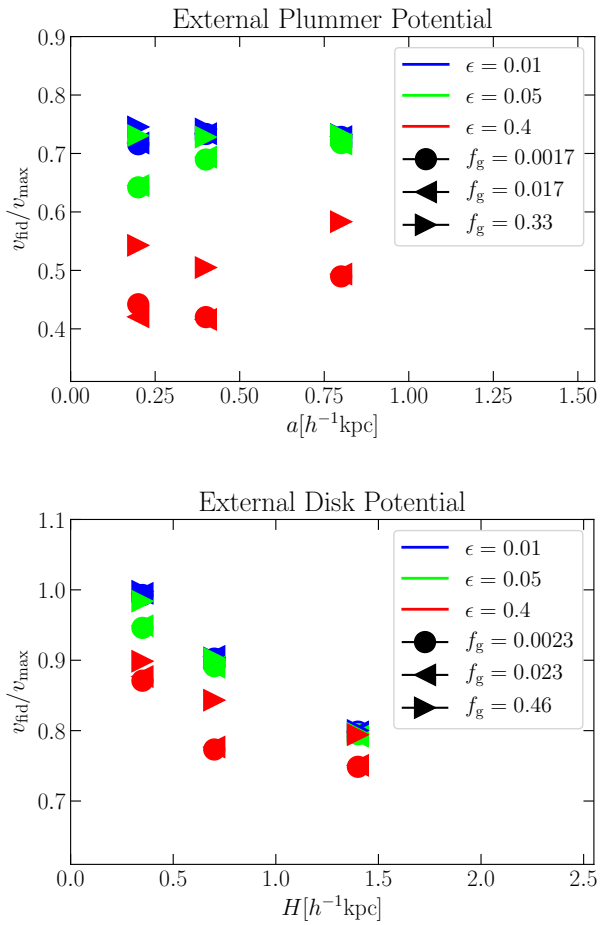

Figure 7.5. A comparison showing the impact of SN feedback on the inner DM distribution in all 27 simulations including an external Plummer and disk potential on the top and bottom panels, respectively. The y−axis quantifies how cuspy/cored the total matter density profile is, i.e., how fast the galaxy circular velocity curve rises, depending on the SN feedback configuration. Lower values of $v_{\rm fid}/v_{\rm max}$ correspond to more cored profiles. However, this ratio can also be altered by contraction of the halo due to the external potential (prior to the first SN feedback cycle; see 7.1), an effect that is of particular importance for the SMC-like galaxy (bottom panel). For that reason, we plot the scale length of the external potentials on the x−axis. At fixed values of the scale length, we can then make a relative comparison between the different SN feedback configurations. Besides galaxy size, we also show how $v_{\rm fid}/v_{\rm max}$ depends on total injected energy and energy injection time, using color-coded symbols as explained in the legends. Across both panels, we find that injected energy and injection time together largely determine whether the DM halo's density profile changes or not. Still, a residual impact of the galaxy's concentration can also be observed.

the galactic potential appears on the $x-$axis of Figure 7.5, all measured ratios $v_{\rm fid}/v_{\rm max}$ with the same $x-$coordinate correspond to simulations with the same initial conditions. Hence, while the value of $v_{\rm fid}/v_{\rm max}$ cannot tell us in absolute terms which DM halo is more cored *across all simulations*, we can use the difference between measured values of $v_{\rm fid}/v_{\rm max}$ with the same $x-$coordinate to establish which combination of parameters in our SN feedback model (see Table 7.2) is more efficient at forming a core.

On the top panel of Figure 7.5, we show the results of the 27 simulations including an external Plummer sphere. As can be seen in Figure 7.1, the mass of the Plummer sphere chosen to mimic a Fornax-like dwarf galaxy is not large enough to cause a significant contraction of the inner DM halo, allowing for an easy comparison of the circular velocity curves between simulations. We can see that the measured final values of $v_{\rm fid}/v_{\rm max}$ are very similar for all choices of a, as long as the injected energy is small and the injection time is long, i.e., small ε and large $f_{\rm g}$ (see Table 7.2). Thus, for this combination of ε and $f_{\rm g}$ (both assuming their "adiabatic" values) the halo essentially retains its cusp for all choices of the galaxy's concentration. The most important factor in determining whether or not a core is formed is the amount of injected energy. In fact, if the a priori energy coupling parameter is small ($\varepsilon = 0.01$), a (very small) core is formed only if the galaxy is very concentrated and the injection time is quite short. However, if the energy coupling is very large ($\varepsilon = 0.4$), cores are formed for virtually every combination of the other two parameters. Injection time and concentration of the galaxy play somewhat smaller roles, with injection time appearing to be slightly more important: the core becomes more significant the shorter the injection time is and/or the more concentrated the Plummer sphere is. We note that varying the injection time between 1 per cent and 10 per cent of the orbital family's radial period ($f_g = 0.0017$ and $f_g = 0.017$, respectively) produces virtually no difference. This indicates that for an intermediate value of the energy coupling parameter, the relevant timescale for energy injection is roughly set by the orbital period of particles in the halo's center (and not much smaller than that), suggesting that although SN feedback needs to be impulsive for a core to form, the requirement for the degree of impulsiveness (and thus for how bursty star formation should be) is not that severe.

The bottom panel of Figure 7.5 shows the case of the external disk potential introduced to mimic a SMC-like galaxy. From Figure 7.1 we know that this far more massive external potential causes a significant contraction of the DM halo, which is reflected in the values of $v_{\rm fid}/v_{\rm max}$ (compared to top panel of Figure 7.5). In fact, all of the measured values lie above the value 0.7, the value for a typical unperturbed NFW halo. However, it is obvious from Figure 7.3 that cores are formed in our simulations with a disk as long as SN feedback is sufficiently energetic and impulsive. As we discussed at the beginning of this subsection, the absolute values of $v_{\rm fid}/v_{\rm max}$ do not imply on their own whether a halo is cored or not. We have explicitly verified that no cores form in the runs in which $\varepsilon = 0.01$ and $f_{\rm g} = 0.46$. This implies that for a fixed value of H, we can take the blue right-pointing triangle as a 'cuspy' baseline and assess how cored the DM profile is for the other simulations at the same H (but with different parameters regulating injected energy and injection time). With that in mind, we observe very similar trends as the ones for the Plummer sphere on the top panel of Figure 7.5. In general, however, the final profiles are somewhat less cored on average for the SMC-like case than for the Fornax-like case, relative to the baselines (blue right-pointing triangle). If cores form

for a given combination of ε and H, the core size (significance) is again regulated by how short the injection time is. The same is true for the disk's scale length H, with more significant cusp-core transformation (relative to the baseline) for more concentrated disks.

7.4.4 Actual and nominal energy change

As shown in Section 7.4.3, the sizes of the final DM cores vary for a fixed value of the nominal energy coupling parameter ε. Assuming that the final DM halos are in dynamical equilibrium and fulfill the virial theorem, this implies that the total energy that has been transferred to the DM particles in the halo is different between those simulations (Peñarrubia et al. 2012). The effective coupling between SN feedback and DM, $\varepsilon_{\rm DM}$, is thus distinctly different from the nominal energy coupling ε defined in Section 7.2.4.

In Figure 7.6 we show the energy that is actually injected into the DM halo for the nine simulations with an external Plummer profile and $a = 0.4\,h^{-1}$kpc on the top panel and the nine simulations with an external disk potential and $H = 0.35\,h^{-1}$kpc on the bottom panel. The energies displayed correspond to the total energy of all DM particles in the halo at a given time measured from snapshots taken every $200\,h^{-1}$Myr. For orientation, we also show as horizontal lines final energies corresponding to several values of the effective coupling parameter $\varepsilon_{\rm DM}$, which we define as

$$\varepsilon_{\rm DM} = \frac{\langle m_\star \rangle \Delta E_{\rm DM}}{M_\star \xi (m_\star > 8 M_\odot) E_{\rm SN}}, \tag{7.16}$$

where $\Delta E_{\rm DM}$ is the increase in the DM halo's total energy and all other quantities are as in Equation 7.11. The values shown for $\varepsilon_{\rm DM}$ are 1.0 (solid black line), 0.4 (dashed black line), 0.05, and 0.01 (thin dotted black lines). For individual runs, we can compare the nominal energy coupling (to the ISM) ε with the actual, effective energy coupling $\varepsilon_{\rm DM}$ at the end of the simulation.

In the Plummer case (top panel of Figure 7.6), we see that our effective model underestimates the energy that is truly injected into the halo in all nine simulations. The discrepancy between ε and $\varepsilon_{\rm DM}$ is particularly large for small nominal energy couplings and decreases for larger values of ε. This is in line with our remarks in Section 7.2.4. The approximation that the injected energy is given by the binding energy associated with one of the Plummer spheres performs better if the local density contrast generated by individual "supernovae" is large.

This picture is validated in the disk case (bottom panel of Figure 7.6), where the agreement between nominal and actual injected energy is again better in the three simulations in which the nominal energy coupling $\varepsilon = 0.4$. We can explain this behaviour by recapitulating how we implement the energy injection in our effective model of SN feedback. To fix the mass of an individual "superbubble", we assume that the energy that is required to unbind a Plummer sphere of said mass is given by Equation 7.12, i.e., the gravitational binding energy of such a Plummer sphere in vacuum. However, within the gravitational potential of the halo – and the external disk or Plummer potential – removing a Plummer sphere of mass $m_{\rm max}$ is equivalent to an energy injection that is larger than just the binding energy given in Equation 7.12. This is because the

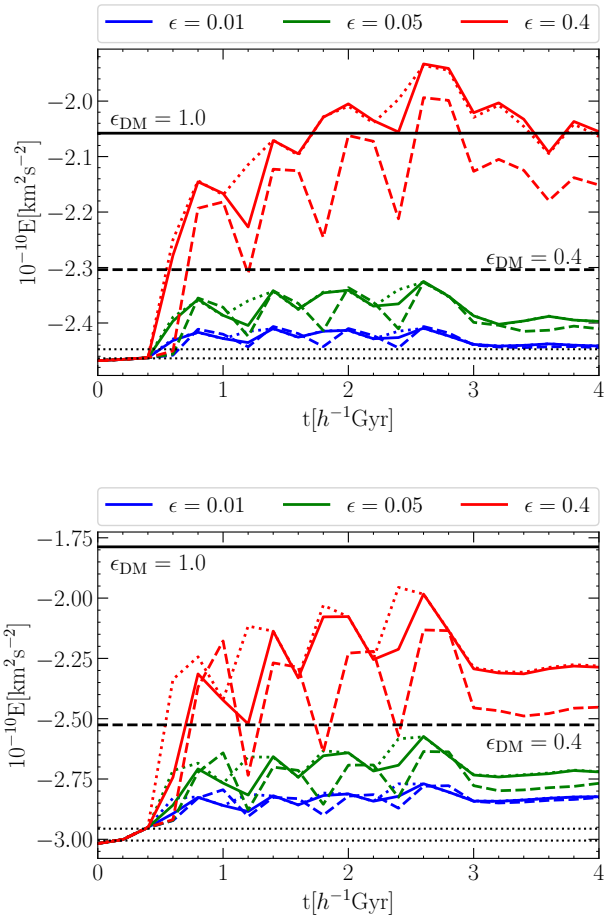

Figure 7.6. The total energy (per unit mass) of DM partices in the halo as a function of time in the nine simulations including an external Plummer sphere with $a = 0.4\,h^{-1}kpc$ (top panel) and an external exponential disk with $H = 0.35\,h^{-1}kpc$ (bottom panel). The colours of the lines refer to simulations with nominal energy couplings ε as indicated in the legends above the panels. Solid coloured lines refer to the runs using the benchmark value of f_g, dashed (dotted) coloured lines to the runs using the adiabatic (impulsive) values of these parameters (see Table 7.1). Black horizontal lines show the final energies corresponding to "effective" energy couplings ($\varepsilon_{\rm DM}$, see Equation 7.16) of 1.0, 0.4, 0.05, and 0.01, respectively. The "effective" energy coupling is larger than the nominal coupling in all simulations, suggesting that our order of magnitude estimate for the injected energy systematically underestimates the actually injected energy. Note that the summed energy of the DM particles is approximately stable during the final $1\,h^{-1}Gyr$ for almost all shown simulations, with the exception of the three runs with the largest nominal energy coupling and an external Plummer potential (red lines on the top panel).

gravitational pull of the surrounding matter needs to be overcome as well. The relative contribution of this extra energy injection, which is associated with the interaction with the surrounding matter and unaccounted for in our model, is larger if the local density contrast generated by the "superbubble" is small. As a consequence, the ratio between $\varepsilon_{\mathrm{DM}}$ and ε is closer to unity in simulations with a larger nominal energy coupling, as can be seen on both panels of Figure 7.6. In line with these considerations, similar trends emerge when comparing between simulations with external Plummer (disk) potentials of different sizes. The effective energy coupling is smaller in simulations with more extended potentials since the ambient density in the surroundings of the "superbubbles" is smaller.

A few further general statements can be made from Figure 7.6:

- The order of magnitude estimate of the total injected energy is roughly consistent with the nominally injected energy with the maximum discrepancy being a factor of a few in the Plummer case.
- In the disk case, the comparison between ε and $\varepsilon_{\mathrm{DM}}$ is complicated by some residual evolution in the total energy at the beginning of the simulation – particularly for low nominal energy couplings.
- The qualitative scaling of the injected energy with the nominal energy injection parameter is as intended, i.e., there are sizeable differences between the three cases.
- In almost all simulations, the energy is relatively stable after $3\,h^{-1}\mathrm{Gyr}$, i.e., when the last SN cycle ends. An exception are the Plummer runs with $\varepsilon = 0.4$. The residual fluctuation here is likely a numerical effect. Violent behaviour of particles in the halo's center can impede the accurate determination of the position and the velocity of the halo's center of potential.
- Gravitational coupling of energy to the DM particles is not perfect. This is immediately evident from the fact that for $\varepsilon = 0.05$ and $\varepsilon = 0.4$, the increase in energy of the DM particles is smaller if the energy is injected over a longer time.

Related to the last point, we notice that DM halos whose final density profiles are more cored (see Figure 7.5) also have larger final energies (for the same external potential). Given that the energy is – in most cases – stable after the last explosion cycle, we can assume that the DM halos are once more in virial equilibrium at the end of the simulation. The correlation between injected energy and final core size is therefore expected (see Peñarrubia et al. 2012).

In summary, our formalism provides a reasonable (within a factor of a few) estimate for the mass that needs to be removed by individual "supernovae" to match the nominally injected energy (Equation 7.11). In most cases, the actual injected energy is somewhat larger. The model can safely be used to test the impact of SN feedback of different strengths. However, if knowledge of the exact amount of injected energy is required, then the sum over the energies of the DM particles at the end of the simulation will have to be manually compared to the equivalent sum at the beginning of the simulation.

7.4.5 Core formation and impulsive energy injection

We have seen so far that core formation always coincides with at least one signature of an impulsive change of the gravitational potential in the phase space plots of the orbital family (radial expansion or formation of shells). We should note that while observing these signatures is a necessary condition for core formation, it is not sufficient. For instance, we observe shell-like structures in the Plummer sphere simulations with small energy coupling if the energy injection time is short and $a < 0.8\,h^{-1}$kpc, but this system does not develop a significant core. This implies that the most important criterion for an effective cusp-core transformation is whether or not the total injected energy is sufficient, which is a requirement calculated in Peñarrubia et al. (2012). If there is enough SN feedback energy deposition, then the significance of the core formed depends on how impulsively the energy injection proceeds (see Figure 7.5). In light of that, perhaps the most striking result of Figure 7.5 is that if the injected energy is large enough, cores can form even if the injection time is of the order of a typical radial period in the halo's central region (red right-pointing triangles in Figure 7.5). If impulsive energy injection is, as stated above, a prerequisite for core formation, then this implies that the rate at which an injection of energy changes the gravitational potential must be fast enough to be perceived as impulsive by a sizeable amount of particles in the halo's center. Below, we aim to provide a more detailed discussion of this in light of our results in Figure 7.5.

Burger et al. (2020) presented a theory for the diffusion of radial actions in time-dependent spherical potentials. One key result is that radial actions J_r in slowly evolving potentials can be written to first order as

$$J_r = J_{r'} + (\mathbf{r}\cdot\mathbf{v})\frac{\dot{R}}{R}\frac{P_r(E,L)}{2\pi} \equiv J_{r'} + \Delta J_r, \tag{7.17}$$

where $J_{r'}$ is a dynamical invariant (the radial action in a frame in which the potential is time-independent), $P_r(E,L)$ is a particle's orbital period and $\dot{R}/R$ describes the rate at which the gravitational potential changes. Another important result of Burger et al. (2020) is that whether radial action distributions display adiabatic or impulsive evolution can be roughly determined by the ratio $\sqrt{\tilde{D}}/J$, where $\tilde{D}$ is the diffusion coefficient of the distribution of radial actions, which for phase-mixed particle ensembles can be calculated as

$$\tilde{D}(J_r,E,t) \approx \frac{1}{2}\left(\frac{\dot{R}}{R}\right)^2\left(\frac{P_r(E,L)}{2\pi}\right)^2\left\langle(\mathbf{r}\cdot\mathbf{v})^2\right\rangle. \tag{7.18}$$

where the brackets represent an ensemble average. The radial action distributions of ensembles of tracer particles (such as the orbital families here) evolve adiabatically if $\sqrt{\tilde{D}}/J_r \leq 0.1$, but their evolution becomes increasingly non-linear as this ratio rises.

In Figure 7.7, we aim to link this measure of whether or not radial actions evolve impulsively on average to the observed final DM density profiles in the simulations with an external Plummer potential and $f_g = 0.33$ (corresponing to the nominally adiabatic energy injection; see Table 7.2). We show a bilinear interpolation of $v_{\rm fid}/v_{\rm max}$ as a function of a and the logarithm of the energy coupling parameter, ε, using the nine simulations corresponding to all the right-pointing triangles in the top panel of Figure 7.5. As we have shown earlier, the core becomes more significant with larger energy coupling. To test how this relates to our above-introduced measure of how impulsive the

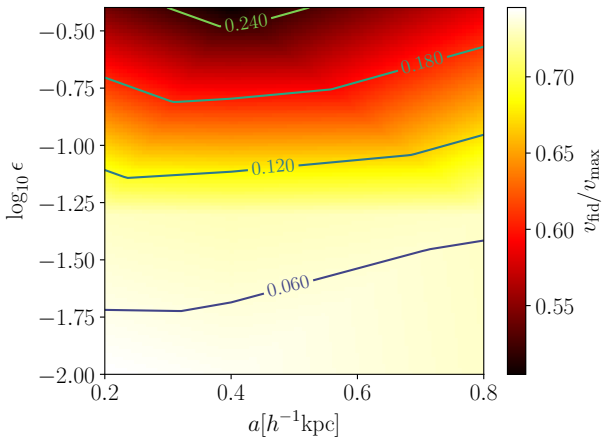

Figure 7.7. Bilinear interpolation of $v_{\mathrm{fid}}/v_{\mathrm{max}}$ (a measure of how cuspy/cored the total density profile is; see Figure 7.5) as a function of the scale of the external Plummer potential, a, and the logarithm of the energy coupling, $\log_{10}\varepsilon_{\mathrm{DM}}$. The interpolation is done using the nine simulations including an external Plummer profile and with $f_{\mathrm{g}} = 0.33$ (see Table 7.2). The lines correspond to a theoretical measure of how impulsive SN feedback is as seen by the tracers comprising the orbital family. It measures the amplitude at which the radial action of a typical tracer oscillates, normalized by the action itself. Larger values correspond to particles whose actions evolve more impulsively on average.

energy injection is, we calculate an estimate of $\sqrt{\tilde{D}}/J_r$ for each of those nine simulations. To do that, we evaluate J_r, $P_r(E,L)$ and $\left\langle(\mathbf{r}\cdot\mathbf{v})^2\right\rangle$ from Equation 7.18 for particles which are part of the initial orbital family of tracers. To obtain a rough estimate of $\dot{R}/R$, we make use of Equations C6 and C13 in Burger et al. (2020), but assume here that the change in the amplitude of the potential dominates over the change in its logarithmic slope[6]. We estimate $\dot{R}/R$ at the pericenter radius of the orbital family, since the change in gravitational potential is more disruptive in the halo's center:

$$\frac{\dot{R}}{R} \approx \frac{1}{2+\alpha(r_0,t_0)} \frac{\Psi(r_0,t_0+f_{\mathrm{g}}P)-\Psi(r_0,t_0)}{f_{\mathrm{g}}P\Psi(r_0,t_0)}, \tag{7.19}$$

where $r_0 = 0.5\,h^{-1}\mathrm{kpc}$, t_0 is the time at which the first SN cycle starts, P is the explosion cycle period (see Table 7.1), and α denotes the logarithmic slope of the shifted potential Ψ, which is defined as $\Psi(r,t) = \Phi(r,t) - \Phi(0,0)$.

Using Equations 7.18 and 7.19, we can estimate the typical diffusion coefficient $\tilde{D}$ for the orbital family of tracers in each of the simulations with $f_{\mathrm{g}} = 0.33$ (slow energy

[6]We note that this might not be a good approximation since the central slope of the gravitational potential can change substantially during the energy injection. However, taking this effect into consideration would require an in-depth analysis of individual DM particle's orbits, which is beyond the scope of this paper. Hence, the calculation of $\dot{R}/R$ presented here can be taken as a rough estimate.

injection). From a bilinear interpolation of the values $\sqrt{\tilde{D}}/J_r$ obtained from each of those runs we calculate the contour lines shown in Figure 7.7. We find that larger cores correspond to larger values of $\sqrt{\tilde{D}}/J_r$. More importantly, our estimate is that $\sqrt{\tilde{D}}/J_r \sim 0.2$ in the simulations with $\varepsilon = 0.04$. According to Burger et al. (2020), this corresponds to the regime in which the evolution of radial action distributions transitions from adiabatic to impulsive behaviour. In particular, it can in fact be impulsive for a significant subset of particles, which would explain the signatures of impulsive energy injection seen in the orbital family for these configurations in our simulations (see also Figure 7.4).

7.5 Discussion

In this article, we have presented an effective model for SN feedback that can be adopted in DMO simulations of isolated halos in order to quickly investigate the impact of changing the total energy budget of SN feedback, the timescale over which the energy is injected, and the spatial concentration of the baryonic matter within the DM halo. In this Section we discuss some key elements and assumptions of our model. In particular, we first focus on some strengths and weaknesses of our chosen galaxy models and of how we determine the total SN feedback energy. We then discuss the differences in the results we obtain depending on the concentration of the baryonic galaxy and on how impulsive the energy deposition is. Eventually, we discuss the implications of our results for how we can use kinematic tracers to detect the imprint of impulsive SN feedback in dwarf galaxies.

7.5.1 Distribution of baryonic matter

All of our simulations are of a DM halo of mass $M_{200} = 10^{10} h^{-1} \mathrm{M}_{\odot}$ with an initial concentration of $c_{200} = 13$. The halo is initially set in approximate dynamical equilibrium, having a NFW density profile for radii smaller than r_{200}, and an exponential cutoff at larger radii. In a first step, we add an external galactic potential to the halo and wait for the halo to contract before applying our SN feedback model. The parameters defining all of the six external potentials used in our work are listed in Table 7.2 and described in Section 7.3.1. The benchmark cases have been chosen in order to approximate Fornax (Plummer sphere) and the SMC (flat exponential disk).

The values of a and M_{Pl} in our benchmark Plummer sphere are inspired by measured properties of Fornax reported in McConnachie (2012), as well as Fornax's historic gas content reported in Yuan et al. (2016). Moreover, the size of our halo corresponds to roughly the measured size of Fornax's host halo. It should be noted that our Plummer model is spherically symmetric, and since Fornax is flattened (e.g. Genina et al. 2018), our model does not give a true representation of Fornax's gravitational potential. However, the contraction of the initial DM halo due to the baryonic mass estimated for Fornax is negligible in all three configurations we explore. Critically for our purposes, a spherically symmetric potential allows for a direct investigation of whether our implementation of SN feedback is impulsive or adiabatic by following the evolution of the phase space distribution of an orbital family as presented in Burger and Zavala (2019).

Our benchmark disk model has been chosen in order to approximate the SMC following Hopkins et al. (2012), but scaling down the mass of the disk by a factor of 2, since our host halo is smaller in mass than the assumed SMC's host halo by exactly that factor. We note that due to the same reason the scale length of our benchmark exponential disk potential should in principle be reduced as well. However, it is unclear how to make this correction since the presence of a prominent gaseous component (far more extended than the stellar disk) in a SMC-like galaxy complicates matters. In order to not over-estimate the gravitational effect of the external disk, we chose to fix its scale length to the value reported in Hopkins et al. (2012). Departures from this choice are nevertheless considered when we discuss the impact of the scale length on how effective SN feedback is at forming a core. We note that since our SMC-like system is simply a scaled down version of the SMC and its host halo reported in Hopkins et al. (2012), its stellar-to-halo mass ratio is somewhat inconsistent with observations (e.g. Moster et al., 2010). We thus emphasize that our disk-galaxy models should be taken as a case study of how the impact of SN feedback changes if the modelled dwarf galaxy is heavier (and axisymmetric instead of spherically symmetric).

7.5.2 SN feedback energy deposition

The nominal SN feedback energy injected into the DM halo from each mock galaxy is calculated from Equation 7.11, taken from Peñarrubia et al. (2012). This total energy budget depends on the choice of the initial mass function and on the effective energy coupling of SN feedback to DM ($\varepsilon_{\rm DM}$). The latter remains a subject of debate with a broad range of values between 0 and 1 effectively used across diverse SN feedback implementations. In this work, we do not model $\varepsilon_{\rm DM}$ directly, but rather measure it at the end of our simulations. To determine the mass of individual "supernovae", we explore the range of values between 0.01 and 0.4 for the nominal coupling of the SN feedback energy to the ISM, i.e., the values that Peñarrubia et al. (2012) considered plausible (see Section 7.4.4 for a discussion of our energy injection scheme and associated caveats). We remark that the available energy for SN feedback depends linearly on the stellar mass in the galaxy, whereas the energy required for cusp-core transformation depends on the square of the DM halo mass. Thus, it is not surprising that Di Cintio et al. (2014), Tollet et al. (2016), Chan et al. (2015), Fitts et al. (2017) and Lazar et al. (2020) find that the inner slope of DM halos in hydrodynamic simulations is a function of the stellar-to-halo mass ratio. In addition to Section 7.4.4, we here make a few further remarks about how we implement the energy injection through SN feedback in our effective model:

- Contrary to how SN feedback occurs in hydrodynamical simulations, the "superbubbles" in our model are stationary over $600\,h^{-1}$Myr. This is larger than the typical orbital times at the relevant radii and may introduce artificial asymmetries into the halo. This effect can be reduced by increasing $N_{\rm SNF}$, and thus creating, on average, a more isotropic distribution of "supernovae".

- Negative local "gas" densities can occur in the "superbubbles". This is not an issue within our model unless the associated acceleration becomes large enough to disrupt the halo. In fact, a large density contrast is desirable, as it leads to a

better agreement between the nominal and actual injected energy (see Section 7.4.4).

- The two points above represent opposing requirements for the parameter $N_{\rm SNF}$. Having a better handle on the injected energy requires $N_{\rm SNF}$ to be small, while an increased symmetry is obtained for larger $N_{\rm SNF}$. We have checked for a few benchmark simulations that varying $N_{\rm SNF}$ by a factor of 2 up or down does not significantly affect our results.

- Since our idealized simulations are initialized from halo properties today, we are not modeling the cosmological history of halo assembly. Keeping this caveat in mind, our analysis is more appropriate for dwarf galaxies with a fairly recent dominant star formation activity, some of which have been associated to have cored DM density profiles (Read et al., 2019).

- The very large cores that form for $\varepsilon = 0.4$ in the simulations with an external Plummer potential and short injection time (red circles and left pointing triangles in Figure 7.5) are fairly unphysical with SN feedback disrupting the halo's density profile out to radii of almost 10 kpc. Simulations with this combination of parameters have nevertheless provided valuable insight into the nature of the SN-driven mechanism of cusp-core transformation.

- Finally, we briefly note that the parameters $f_\star$ and ε are degenerate in our model. We include $f_\star$ as a parameter to facilitate a comparison of our results with observations and simulation results in the literature and to relate the feedback energy to the total stellar mass.

7.5.3 Impact of the concentration of baryons on SN feedback

We investigate the impact of how concentrated baryonic matter is within the halo in the SN-driven mechanism of cusp-core transformation by varying the scale length a (H) of the external Plummer (exponential disk) potential. The mass distribution corresponding to those potentials directly determines the spatial distribution of explosion centers in our SN feedback model (see Section 7.2.3).

In the case of the disk potential, our modelling has a couple of caveats, which could lead to underestimating the SN feedback impact. The first one is that in order to analytically calculate the potential, we assume the galaxy to be an infinitely flat exponential disk (see Section 7.2.2). However, we distribute the locations of individual "SN centers" in a more realistic way; they are not placed exactly in the disk plane but follow a $\cosh^{-2}$ distribution in the vertical direction. This makes our implementation of SN feedback slightly inconsistent with the calculation of the external potential, but we accept this inaccuracy in the interest of distributing the "SN locations" more realistically. The second reason we may underestimate the impact of SN feedback in the disk case is that we do not differentiate between gas and stars in the baryonic disk. Modelling a two-component disk potential with a more extended gaseous disk would thus reduce the contraction of the DM halo (see Figure 7.1), while keeping the amount of injected SN energy, as well as its spatial distribution, the same. These effects should lead to an increased impact of SN feedback. However, as we mentioned above, the mass of our

disk potential lies above the Moster et al. (2010) stellar-to-halo mass relation and thus the calculated energy available for SN feedback is rather large. This over-estimate of the SN feedback energy can potentially cancel some of the suppression effects outlined above.

Irrespective of how well our results can be compared to observations of real dwarf galaxies, a clear trend emerges when changing the scales of the external disk or Plummer potential. In case of a nominal energy coupling $\varepsilon = 0.05$, this trend becomes particularly obvious in Figure 7.5. The more concentrated the baryonic distribution is, the larger the reduction in central DM density, provided the injection time is shorter than a typical dynamical time. This result is in general agreement with the findings of Benítez-Llambay et al. (2019) and Lazar et al. (2020). Namely, these works find that increasing the star formation threshold in hydrodynamical simulations leads to denser and more concentrated gas in the halo center, and subsequently to a more concentrated energy injection into the halo, making SN feedback more efficient at forming DM cores. This does not mean that an inversion of this trend, as reported in Benítez-Llambay et al. (2019), can be disproved by our model. In our model, "gas" is removed from the central halo instantaneously. In hydrodynamic simulations, gas can only be removed from the center of the galaxy if the energy injected through SN feedback is sufficiently large to overcome the gravitational pull of the combined potential of the DM and the baryonic matter in the center. In the simulations of Benítez-Llambay et al. (2019), this condition may not be fulfilled for very large star formation thresholds, due to the concentrated accumulation of baryonic mass in the center of the galaxy. This would explain the "sweet spot" range of star formation thresholds for core formation reported by the authors. While we cannot conclude from our results whether cores can form in very dense systems or not, we can make a different statement. If cores are formed in these systems, then DM cusps are unlikely to be restored due to the gravitational pull of gas that re-accumulates in the center of the galaxy. In other words, the DM cores that are formed in our simulations are stable – even in the presence of a very centrally concentrated "baryonic" potential.

7.5.4 Impulsiveness of SN feedback

There is an overall consensus in the literature that star formation needs to be "bursty" for SN feedback to be efficient at forming a DM core. In this article, we find that the most important criterion is the amount of energy that is deposited into the system by SN feedback. Nonetheless, we find that the energy injection timescale plays an important role as well, either in determining the core significance (for high energies, $\varepsilon \sim 0.4$), or whether or not a core is formed at all (for low energies, $\varepsilon \sim 0.05$, see Figure 7.5). We find virtually no difference in our results (core significance and the presence of signatures of impulsive SN feedback in the final phase space distribution of an orbital family of tracers) between simulations with injection times that are $\sim 1\%$ of the dynamical time and injection times that are $\sim 10\%$ of the dynamical time. The situation changes significantly once the injection time becomes comparable to the dynamical time. For small and medium energy couplings, DM cores do not form for such long injection times and no signatures of impulsive SN feedback can be detected in the phase space distribution of tracers. The case of a large nominal energy coupling ($\varepsilon = 0.4$) is

different, both DM cores and their signatures in the phase space of orbits are present, albeit reduced in significance. As explained in Section 7.4.5, the reason for this is that if $\varepsilon = 0.4$, the change in the gravitational potential is substantial enough to be perceived as impulsive by particles in the halo's center, even if the SN feedback energy is injected on a timescale comparable to the dynamical time.

Overall, we find that signatures of impulsive SN feedback can always be detected in the phase space distribution of an orbital family of tracers if the DM halo has formed a core. The reverse statement, however, is not true. If the energy coupling is rather low ($\varepsilon = 0.01$) cores do not form (see Figure 7.5), yet we can see some signatures of an impulsively changing gravitational potential imprinted in the phase space of the orbital family, provided the energy is injected on timescales shorter than the dynamical time. In general, we find that while these signatures are present, the net radial expansion and diffusion of orbits is less significant for SN feedback that is not energetic enough to form a core. Still, this raises the question of whether and how such kinematic signatures can be used to differentiate between adiabatic and impulsive core formation, as suggested in Burger and Zavala (2019). One possibility is to look at differences in age and metallicity gradients of stars between a system that undergoes adiabatic core formation and one that undergoes impulsive core formation. We are currently investigating this issue (Burger et.al., in prep.) by means of a suite of hydrodynamical simulations of a single isolated halo, including different star formation thresholds using the stellar feedback model SMUGGLE (Marinacci et al., 2019) incorporated into AREPO (Springel, 2010).

7.6 Summary

We have presented a new effective model of SN-driven cusp-core transformation that can be included in N-body simulations of isolated halos. Our model consists of two main components, an external potential that approximates the distribution of baryons in a dwarf galaxy, and a scheme to inject energy into the DM particle distribution in a manner that is approximately consistent with the stellar distribution modelled by the external potential. In a series of simulations, we have tested how the effect of SN feedback depends on the baryonic concentration, the amount of injected SN feedback energy, and the timescale on which this energy is injected into the halo. We have used simulations of a dwarf-size halo to examine the cases of a Plummer potential (to mimic a Fornax-like system) and a disk potential (to mimic a SMC-like system). We find that the most important factor determining whether SN feedback can form cores in dwarf galaxies is whether or not enough energy is available to transform the halo's density profile. If the available energy is close to maximal, cores form even if the SN injection time is longer than the dynamical timescale in the halo center and/or baryons are concentrated or not. If less energy is available, whether or not cores form depends on how fast the energy is injected and on how concentrated the baryonic matter is within the halo. For minimal values of the energy, cores cannot form. For a fixed amount of feedback energy, larger cores form for faster injection times and more concentrated galaxies. Cores formed in very concentrated galaxies are stable – adiabatic contraction due to the centrally concentrated baryonic potential does not restore the cusp, even if no further supernovae occur. Analyzing the phase space distribution of tracer particles, we find clear signatures

of impulsive SN feedback in all simulations in which the DM halo develops a core. However, we also find these signatures in a few simulations in which the halo's density profile remains cuspy as the amount of SN feedback energy is insufficient to trigger core formation. The longevity and the appearance of those signatures are closely linked to the spatial symmetry of the halo and the external potential.

Acknowledgments

JB and JZ acknowledge support by a Grant of Excellence from the Icelandic Research Fund (grant number 173929). The simulations in this paper were carried out on the Garpur supercomputer, a joint project between the University of Iceland and University of Reykjavík with funding from the Icelandic Research Fund.

Chapter 8

Signatures of SIDM and SNF

This chapter is based on the following article:

Degeneracies Between Self-interacting Dark Matter and Supernova Feedback as cusp-core transformation mechanisms

in preparation

Authors:
Jan D. Burger[1], Jesús Zavala[1], Laura Sales[2], Mark Vogelsberger[3], Federico Marinacci[4], and Paul Torrey[5]

[1]Centre for Astrophysics and Cosmology, Science Institute, University of Iceland, Dunhagi 5, 107 Reykjavik, Iceland

[2]University of California Riverside, 900 University Ave., Riverside CA 92521, USA

[3]Department of Physics, Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

[4]Department of Physics & Astronomy, University of Bologna, via Gobetti 93/2, I-40129 Bologna, Italy

[5]Department of Astronomy, University of Florida, 211 Bryant Space Sciences Center, Gainesville, FL 32611, USA

We present a suite of 16 high-resolution hydrodynamic simulations of an isolated dwarf galaxy (gaseous and stellar disk plus a stellar bulge) within an initially cuspy dark matter (DM) halo, including self-interactions between the DM particles (SIDM), stochastic star formation and subsequent supernova feedback (SNF). SIDM momentum transfer cross section and star formation threshold are varied between simulations. The DM halo forms a constant density core of similar size and shape for several combinations of those two parameters. Haloes with cores which are formed due to SIDM (adiabatic cusp-core transformation) have velocity dispersion profiles which are closer to isothermal than those of haloes with cores which are formed due to SNF in

simulations with bursty star formation (impulsive cusp-core transformation). Impulsive SNF can generate steep stellar age gradients and increase random motion in the gas at the centre of the galaxy. Simulated galaxies in haloes with cores which were formed adiabatically are spatially more extended, with stellar metallicity gradients that are shallower (at late times) than those of galaxies in other simulations. Such observable properties of the gas and the stars, which indicate either an adiabatic or an impulsive evolution of the gravitational potential, may be used to determine whether observed cores in DM haloes are formed through self-interactions between the DM particles or in response to impulsive SNF.

8.1 Introduction

Precision measurements of the cosmic microwave background (Planck Collaboration et al. 2020) reveal that the matter distribution in the early Universe was almost completely homogeneous, perturbed only by small density fluctuations. The ΛCDM concordance model, in which ~ 80 per cent of the matter content in the Universe consists of collisionless, cold dark matter (CDM), successfully explains the growth of these small fluctuations into the large scale structure we observe today (Springel et al. 2006). Collisionless N-body simulations predict the hierarchical collapse of overdensities into sheets, filaments, and eventually self-gravitating virialized dark matter (DM) haloes. Galaxies, consisting of ordinary baryonic matter, are hosted by such DM haloes. Under the assumption that the brightest observed galaxies are hosted by the most massive DM haloes, the clustering and the abundance of observed galaxies are well explained by the spatial distribution of DM haloes in large cosmological simulations (Frenk et al. 1988, Kazantzidis et al. 2004, Conroy et al. 2006, Moster et al. 2010, Behroozi et al. 2013). Moreover, the observed rotation curves of large spiral galaxies are well explained by the combined mass of visible matter and DM (see e.g. van Albada et al. 1985).

On the scale of dwarf galaxies, however, the situation is far more uncertain. The dynamical properties of some observed dwarf galaxies appear to be inconsistent with predictions from collisionless *N*-body simulations in regards to the abundance and the inner structure of low-mass CDM haloes. These mismatches between simulations and theory are longstanding issues that have become known as the small-scale challenges to ΛCDM (see Bullock and Boylan-Kolchin, 2017, for a review). To date, it remains unclear whether these challenges are a manifestation of known but uncertain non-gravitational baryonic physics, which is not present in DM only N-body simulations, or whether a modification of the ΛCDM concordance cosmogony is needed to tackle them.

One of these challenges is the so-called cusp-core problem. Cosmological CDM N-body simulations predict that the spherically-averaged density profiles of DM haloes can be uniquely described by a single two-parameter fitting function, the so-called Navarro-Frenk-White (NFW) profile (Navarro et al., 1996b, 1997). This universality has been demonstrated over 20 orders of magnitude in halo mass (Wang et al., 2020). Importantly, the spherically-averaged density of NFW haloes rises inversely proportional with radius close to the halo's centre; NFW haloes are cuspy. However, the observed rotation curves of some dwarf Irregulars and Low Surface Brightness galaxies in the field (e.g., Moore, 1994; de Blok et al., 2008; Kuzio de Naray et al., 2008; Read

et al., 2019), and at least two Milky Way dwarf spheroidals (Fornax and Sculptor) (Walker and Peñarrubia, 2011), are seemingly inconsistent with the assumption that these galaxies are hosted by cuspy DM haloes. Instead, the slow-rising nature of their rotation curves suggests that these galaxies may be hosted by DM haloes with extended central cores of constant density. A potentially related issue is that some observations suggest that the mass enclosed within the central kiloparsec of dwarf galaxies may be overpredicted by collisionless CDM N-body simulations (Alam et al. 2002, Oman et al. 2015). Currently, there is ongoing debate about whether measurements of the HI-rotation curves of field dwarfs are interpreted correctly (e.g. Oman et al. 2019, Santos-Santos et al. 2020), and whether strong deviations of spherical symmetry in the dwarf spheroidals Fornax and Sculptor may have wrongfully led to the conclusion that the inner density profiles of their host haloes are cored (Genina et al. 2018). However, it is clear that if observations of slow-rising rotation curves in dwarf galaxies stand the test of time, a non-gravitational physical mechanism that transforms central density cusps into cores is needed to reconcile them with the success of ΛCDM on larger scales. Several such mechanisms of cusp-core transformation have been proposed and while some of them invoke baryonic physics to flatten out the central density profile of dwarf-size haloes, others require abandoning ΛCDM for a different cosmogony that resembles ΛCDM on large scales.

Among the mechanisms of cusp-core transformation that work within ΛCDM, the most viable one is core formation induced by supernova (SN) feedback (Navarro et al., 1996a; Gnedin and Zhao, 2002; Pontzen and Governato, 2012; Garrison-Kimmel et al., 2013; Di Cintio et al., 2014; Tollet et al., 2016; Chan et al., 2015; Fitts et al., 2017; Lazar et al., 2020; Burger and Zavala, 2021). Repeated energy injection from supernovae in the dwarf galaxy can give rise to galactic-scale gas outflows, causing rapid fluctuations of the enclosed baryonic mass, and hence of the total gravitational potential within the inner DM halo. As shown in detail by Pontzen and Governato (2012), repeated impulsive changes in the gravitational potential cause a net radial expansion of the orbits of particles that move within. In the case of core formation induced by SNF, this means that the strongly fluctuating gravitational potential causes a radial expansion of the orbits of individual DM particles in the halo centre, thus flattening the central density profile.

To be a feasible mechanism of cusp-core transformation, SNF needs to fulfill a number of conditions. First and foremost, the total energy that is released by supernovae has to be sufficient to unbind the DM halo's central cusp (Peñarrubia et al., 2012). A secondary condition is that SNF needs to be impulsive, i.e., SN-driven gas outflows need to give rise to sizeable changes of the gravitational potential on timescales which are shorter than the typical dynamical times of DM particles in the inner halo (Pontzen and Governato, 2012; Burger and Zavala, 2021). From the observational side, there is evidence that starbursts in bright dwarfs, and thus, their associated supernova cycles, happen on timescales that are comparable to the typical dynamical times of those galaxies (Kauffmann, 2014). However, observations still lack the time resolution required to resolve starburst cycles on the smaller dynamical timescales of the low-mass MW dwarf spheroidals (Weisz et al., 2014). In general, the more energy is injected during a SNF cycle, the shorter the time is over which that energy is injected, and the more concentrated the baryonic mass is to the centre of the DM halo (Burger and Zavala

2021), the more efficient the CM cusp-core transformation will be. In hydrodynamic simulations of galaxy formation, the implementations of SNF are calibrated to the resulting structural properties of larger galaxies. Thus, the efficiency of SNF at flattening the cusps of dwarf-size DM haloes is mainly determined by one model parameter, the gas density threshold for star formation (Benítez-Llambay et al. 2019, Dutton et al. 2020). Larger star formation thresholds lead to more bursty star formation, more concentrated and impulsive feedback, and a stronger contribution of baryons to the central potential, and hence to enhanced core formation (Bose et al., 2019a; Benítez-Llambay et al., 2019).

Among the most viable mechanisms of cusp-core transformation that require changes to the assumed cosmogony is one that was proposed specifically as a possible solution to the cusp-core problem. It proposes that the DM is in fact not collisionless but self-interacting (SIDM, Spergel and Steinhardt 2000, Yoshida et al. 2000, Davé et al. 2001, Colín et al. 2002, Vogelsberger et al. 2012, Rocha et al. 2013, see Tulin and Yu 2018 for a review). In SIDM, particles can exchange energy and momentum through elastic scattering, causing an outside-in energy redistribution within the centre of DM haloes, resulting in the formation of an isothermal core. The timescale on which an initially cuspy SIDM halo forms a flat and isothermal core is roughly given by the time it takes for each DM particle in the inner halo to scatter at least once (Vogelsberger et al., 2012; Rocha et al., 2013). The strength of the self-interaction in SIDM models is parametrized in terms the momentum transfer cross section per unit mass, σ_T/m_χ. Depending on the specific SIDM model, σ_T/m_χ can either be constant or dependent on the relative velocity between the two scattering DM particles. SIDM is an efficient mechanism of cusp-core transformation in dwarf-size haloes for $\sigma_T/m_\chi \gtrsim 1\,\mathrm{cm^2 g^{-1}}$, whereas SIDM haloes are virtually indistinguishable from CDM haloes if $\sigma_T/m_\chi \lesssim 0.1\,\mathrm{cm^2 g^{-1}}$ (Zavala et al. 2013). The most stringent and precise constraints on the self-interaction cross section have been put on the scales of galaxy clusters (e.g. Robertson et al. 2017, Robertson et al. 2019) and large elliptical galaxies (Peter et al. 2013), where observations require that $\sigma_T/m_chi \lesssim 1\,\mathrm{cm^2 g^{-1}}$. On smaller scales, Read et al. (2018) concluded that $\sigma_T/m_\chi \lesssim 0.6\,\mathrm{cm^2 g^{-1}}$, based on their findings that the central density profile of the MW dwarf spheroidal galaxy Draco is cuspy. Moreover, based on a DM only analysis of the updated too-big-to-fail problem, Zavala et al. (2019) concluded that SIDM models with a constant cross section of $\sigma_T/m_\chi \sim 1\,\mathrm{cm^2 g^{-1}}$ fail to explain the apparently large central densities of the host haloes of the ultra-faint satellites of the MW (Errani et al., 2018). It should be pointed out that the constraints on σ_T/m_χ on the scale of dwarf galaxies are affected by significantly larger systematic uncertainties than on the scales of galaxy clusters or elliptical galaxies. Moreover, Zavala et al. (2019) demonstrate that SIDM with a strongly velocity-dependent self-interaction cross section may provide a natural explanation for the observed diversity in the rotation curves of the MW dwarf spheroidals (see also Correa 2021). The strong dependence of the self-interaction cross section on the typical DM velocities would create a bimodal distribution of rotation curves in the MW satellites in which the heavier haloes have constant density cores while the lighter haloes have undergone gravothermal collapse and have very steep central cusps as a consequence . The same mechanism of gravothermal collapse might be accelerated by tidal interactions in the environment of the MW leading to an agreement between constant cross section SIDM models with $\sigma_T/m_\chi \sim 3\,\mathrm{cm^2 g^{-1}}$ and the internal

kinematics of MW satellites (e.g. Kahlhoefer et al. 2019; Sameie et al. 2020).

Burger and Zavala (2019) have shown that while both SNF and SIDM can transform cusps into cores in dwarf-size haloes, the two mechanisms leave distinct signatures in the dynamical properties of kinematic tracers. This difference is related to the different timescales on which SNF and SIDM affect the gravitational potential. While SNF is a viable mechanism for cusp-core transformation only if it causes strong and impulsive fluctuations in the central potential, SIDM thermalizes the central region of DM haloes on timescales which are comparable to or larger than the typical dynamical timescales at distances of $\sim$ 1kpc from the centre of dwarf galaxies. In other words, SIDM haloes form cores adiabatically, while SNF forms them impulsively. Stars, which approximately act as tracers of the gravitational potential, respond differently to impulsively changing potentials than they do to adiabatically changing potentials. In particular, while the actions of tracers on regular orbits are conserved in adiabatically evolving potentials (e.g. Binney and Tremaine 2008), this is not the case in impulsively evolving potentials. Moreover, the orbits of tracers in adiabatically changing potentials quickly adapt to the evolution of the potential, while an ensemble of tracer particles can be put out of dynamical equilibrium in impulsively changing potentials. Hence, the dynamical properties of the stars may differ considerably between *i*) dwarf galaxies with cuspy haloes, *ii*) cored haloes with an adiabatic core formation history, and *iii*) cored haloes with an impulsive core formation history.

In this article, we aim to identify such differences using 16 high-resolution hydrodynamical simulations of an isolated dwarf galaxy with global parameters resembling the Small Magellanic Cloud (SMC) embedded within a live halo (similar to Hopkins et al. 2012). Starting from idealized initial conditions, we simulate the evolution of the system over roughly half a Hubble time using the moving-mesh code AREPO (Springel 2010) with the interstellar medium (ISM) and stellar evolution model "Stars and MUltiphase Gas in GaLaxiEs" (*SMUGGLE*) introduced in (Marinacci et al. 2019) and the Monte-Carlo code for self-interactions between DM particles described in Vogelsberger et al. (2012). Core formation within colissionless dark matter haloes in *SMUGGLE* is investigated in detail in Jahn et al., (*in-prep*).

All simulations start from identical initial conditions and are carried out using different combinations of the momentum transfer cross section per unit mass σ_T/m_χ and the gas density threshold for star formation $n_{\rm th}$. We investigate which combinations of those two parameters cause the DM halo hosting the SMC-analogue to form a constant-density core, and for which parameter combinations the halo retains its initial central cusp. Subsequently, we look for observable quantities that are characteristically different between simulations in which the DM halo forms a core of similar size. In other words, we look for ways in which we can break the degeneracy between SNF and SIDM as cusp-core transformation mechanisms. To that end, we compare three observable quantities, which are derived from the dynamical properties of either the stars or the gas: i) the spatial extent of the visible galaxy, ii) the amount of random motion in the line-of-sight gas velocity, and iii) the age and metallicity gradients of the stars formed throughout the simulation.

This article is structured as follows. We describe the simulations and initial conditions in Section 8.2, present our results in Section 8.3, and summarize our findings in Section 8.4. In Appendix 8.5 we discuss several caveats that arise because of the

stochastic nature of star formation.

8.2 Simulations

Our goal is to investigate core formation in dwarf-size DM haloes and how we can differentiate between cores that have been formed adiabatically and cores that have been formed impulsively. To that end, we perform a suite of 16 different hydrodynamical simulations starting from the same idealized system. In the suite of 16 simulations, we investigate the impact of different SNF models on the final DM density profile. Specifically, we use the ISM and stellar feedback model *SMUGGLE*(Marinacci et al., 2019) with four different values of the gas density threshold for star formation. We also investigate how elastic self-scattering between DM particles affects the final DM density profile, testing four different values of the momentum transfer self-scattering cross section. In this Section, we briefly outline how the initial conditions of our simulations are generated and how an orbital family of kinematic tracers is included into the initial conditions. Then, we will discuss the ISM model used in our simulations, as well as the algorithm employed to model DM self-scattering.

8.2.1 Initial conditions

We set up an isolated DM halo in dynamical equilibrium containing a baryonic galaxy consisting of a stellar disk, a gaseous disk and a stellar bulge. The structural parameters of our initial conditions are similar to the SMC-like galaxy presented in table 1 of Hopkins et al. (2012).

The DM halo is modeled as a Hernquist sphere whose structural parameters are defined by its circular velocity v_{200} at the virial radius r_{200} and its concentration parameter c_{200}[1]. Here, we use $v_{200} = 36.3\,\mathrm{km\,s^{-1}}$ and $c_{200} = 18$. The virial mass, radius and scale length (M_{200}, r_{200}, and r_s, respectively) can be derived from these two parameters.

The baryonic components are defined by their mass fractions, relative to the DM halo mass, and by their density profiles. Both the stellar and gaseous disks combined have a mass of $0.0445 \times M_{200}$ (84% of the disk mass is gas while the rest consists of collisionless star particles). The stellar bulge consists fully of collisionless star particles and makes up for a fraction of 0.005 of the total mass M_{200}. The gas disk has an exponential surface density profile,

$$\Sigma(R) \propto \exp\left(-\frac{R}{H_{\mathrm{gas}}}\right), \tag{8.1}$$

with a scale length $H_{\mathrm{gas}} = 2.1\,\mathrm{kpc}$, where R is the cylindrical radius. As we are interested in a late time dwarf galaxy, we consider a fully ionized gas composition. The gas is initially isothermal with a temperature of $10^4\,\mathrm{K}$, and has solar metallicity. The vertical structure of the gaseous disk is initialized such that the gas is in hydrostatic equilibrium (see Hernquist 1993 and Springel et al. 2005a).

[1] r_{200} is defined through the equation $M_{200} = 200 \times 4\pi/3\rho_{\mathrm{crit}} r_{200}^3$, where M_{200} is the halo's virial mass, ρ_{crit} is the critical density of the Universe, and $c_{200} = r_{200}/r_{-2}$ is the halo concentration, with r_{-2} being the radius at which the logarithmic slope of the halo's density profile equals -2.

The stellar disk also has an exponential surface density profile, but with a smaller scale length of $H_\star = 0.7\,\mathrm{kpc}$, and a vertical distribution given by:

$$\rho(R,z) \propto \Sigma(R) \left[\cosh\left(\frac{z}{z_0}\right)\right]^{-2}, \tag{8.2}$$

with a scale height $z_0 = 0.14\,\mathrm{kpc}$. The bulge is modelled as a Hernquist sphere with a scale length $A = 0.233\,\mathrm{kpc}$.

Due to the spherical symmetry of halo and bulge, we can make use of Eddington's equation (Eddington, 1916) to calculate the full distribution functions of both the halo and bulge particles. The velocities of different particles are subsequently sampled directly from the distribution function. However, due to the presence of the baryonic disk, the total gravitational potential deviates from spherical symmetry. For definiteness, we calculate the distribution function by performing Eddington's integral along the direction perpendicular to the plane of the disk and note that this procedure introduces a small degree of inaccuracy.

For the stellar disk, we calculate the velocity dispersion tensor on a logarithmic grid of R, z values using the Jeans equation in cylindrical coordinates and the streaming velocity from the enclosed mass profile using the epicyclic approximation (see Springel et al., 2005a, for details). The velocities of individual disk particles are then comprised of the streaming velocity and an added random component which is calculated using a local Maxwellian velocity distribution based on the calculated velocity dispersion tensor. The velocities of individual gas cells in the gaseous disk are set to the gas' streaming velocities (calculated taking into account both gravity and the gas pressure gradient) at the position of the respective cell.

We initially set up 1.2×10^7 DM particles, 4×10^5 gas cells, 8×10^4 collisionless disk particles and 8×10^3 bulge particles. The mass of each particle is then approximately $1.3 \times 10^3 M_\odot$. The gravitational softening length is $\varepsilon_g = 24\,\mathrm{pc}$ for all particle species.

Since the algorithm we use to set up our initial conditions relies on assuming spherical symmetry to calculate the distribution functions of the bulge and the halo, the resulting distribution of particles is not fully in dynamical equilibrium due to the presence of the axisymmetric disk component. To remedy this, we evolve the system for a time of 1 Gyr, solving for the dynamical evolution of the gas and the collisionless particles but disabling cooling processes and deactivating star formation and stellar feedback. After letting the system relax for 1 Gyr, we take the final snapshot as our new initial conditions.

8.2.2 The stellar evolution model

We use the *SMUGGLE* stellar feedback and ISM model (Marinacci et al., 2019) for the moving mesh code AREPO (Springel, 2010). We refer the reader to the original paper for details about the model and its implementation. Here, we briefly review two key components that are of key importance in our work, namely the stochastic implementation of star formation and the implementation of SNF.

The formation of star particles proceeds stochastically and is based on the star

formation rate within a given gas cell, which is given by (Marinacci et al., 2019):

$$\dot{M}_\star = \begin{cases} 0 & \rho < \rho_{\rm th} \\ \varepsilon \dfrac{M_{\rm gas}}{t_{\rm dyn}} & \rho \geq \rho_{\rm th} \end{cases}, \tag{8.3}$$

where $M_{\rm gas}$ is the gas mass in a given gas cell, $t_{\rm dyn}$ is the dynamical time of the gas cell, and ε is the star formation efficiency parameter, set to a value of 0.01 in all of our runs. It is evident from Eq. (8.3) that star formation can only proceed if the gas density in a given gas cell is larger than the threshold density $\rho_{\rm th}$. This in itself is of key importance, as it implies that changing this parameter can significantly impact the distribution of gas densities throughout the simulation. In particular, increasing the threshold will lead to more concentrated gas and therefore to more concentrated star formation. Apart from the density criterion, gas cells are also required to be gravitationally bound, meaning that they cannot overcome their self-gravity through gas motion and thermal energy. If both of these criteria are fulfilled, a gas cell is stochastically converted into star particles with a probability of $p = 1 - \exp(-\dot{M}_\star \Delta t / M_i)$, where M_i is the mass in the gas cell i and Δt here denotes a simulation time step. This probability is then compared to a random number x in the interval $(0,1)$ drawn from a uniform distribution. The gas cell is converted into a star particle if $p \geq x$. The formed star particles represent stellar populations with a Chabrier (2001) initial mass function.

The implementation of SNF is explained in great detail in Section 2.3 of Marinacci et al. (2019). The algorithm differentiates between type II supernovae and type Ia supernovae. The total momentum injected into the ISM is boosted if the cooling radius, the radius at which the SN remnant transitions from an adiabatic Sedov-Taylor phase to a momentum conserving phase, cannot be resolved in the simulation. This is the case for most simulations, given that the cooling radius is of the order of a few pc, well below the scales that are resolved in galaxy formation simulations. The expected values of ejected mass, energy, and the total number of supernovae are self-consistently calculated at each time-step and for each star particle. Time steps are chosen such that the expected number of supernovae is below one at essentially all times. A discrete number of supernovae is then sampled from a Poisson distribution with the expected number of supernovae as the distribution's mean. Once the number of supernovae, the ejected energy, momentum, mass, and metallicity have been determined, these quantities are distributed over a fixed number of nearest neighbour gas cells. Fixing the number of nearest neighbours implies defining a search radius h via

$$N_{\rm ngb} = \frac{4\pi}{3} h^3 \sum_i W(|\mathbf{r}_i - \mathbf{r}_s|, h), \tag{8.4}$$

where $\mathbf{r}_i$ is the position vector of the i'th neighbouring gas cell, $\mathbf{r}_s$ is the star particle's position vector, and W is the cubic spline kernel. If the radius h determined in this way is larger than $R_{\rm SB}$, the typical radius of a super bubble ($\sim 1\,$kpc), then the feedback energy and momentum are distributed amongst cells within $R_{\rm SB}$, while mass and metallicity are distributed amongst the $N_{\rm ngb}$ nearest neighbours within the search radius h. If there are $N_{\rm ngb}$ nearest neighbouring gas cells within the super bubble radius, no distinction is made. The SN ejecta are divided amongst cells using weights that are proportional

to the solid angle covered by the cells as seen from the stellar particle's position (see Eq. (35) in Marinacci et al. 2019).

A key parameter of the model is $n_{\rm th}$, the number density threshold for star formation. Together with the average mass per gas cell this determines the density threshold $\rho_{\rm th}$ in Eq. (8.3). Its value is therefore directly related to how clustered the stellar populations that form are, and hence, how clustered SNF is. A larger star formation threshold leads to more bursty star formation, leading in turn to more energetic and impulsive SNF. We thus expect $n_{\rm th}$ to play a key role in determining whether SNF is effective at forming cores in our simulations.

8.2.3 The SIDM algorithm

To model SIDM, we use the algorithm introduced in Vogelsberger et al. (2012) and described in Section 2.2 therein. In this algorithm, the probability for scatter between two DM particles i and j is given by

$$P_{ij} = m_i \frac{\sigma_T}{m_\chi} v_{ij} W(r_{ij}, h_i) \Delta t_i, \tag{8.5}$$

where m_i is the i'th DM simulation particle's mass, σ_T/m_χ is the SIDM momentum transfer cross section per unit mass, v_{ij} is the relative velocity between particles i and j, and Δt_i is the time step of particle i. The scattering probability is smoothed by the cubic spline kernel W, whose arguments are the distance between particles i and j and the smoothing length h_i, denoting the radius of a sphere around simulation particle i which contains a predetermined number of nearest neighbours. The total probability for a scatter is given by a sum of the probabilities calculated according to Eq. (8.5) over all the nearest neighbours and multiplied by $1/2$. Whether and with which of the nearest neighbours a scatter occurs in a given timestep is determined stochastically as outlined in Vogelsberger et al. (2012). To model an elastic scatter, the two colliding particles are assigned new velocities in a way that conserves both total momentum and total energy. The total number of scattering events is regulated by σ_T/m_χ; if each DM simulation particle in the halo centre takes part in ~ 1 scattering event, the inner halo forms an isothermal constant density core. The size of this core and the timescale at which it forms depend on the strength of the interaction, and hence on the SIDM transfer cross section.

8.2.4 Simulation suite parameter space

The focus of our work is whether or not cores are formed in dwarf-sized DM haloes through either the adiabatic (SIDM) or the impulsive (SNF) mechanism, and if so, how quickly they form. Moreover, we aim to find observables that can differentiate between the two cusp-core transformation mechanisms. By changing the parameters of both the star formation and stellar feedback model, and the SIDM algorithm, we are able to regulate the relative importance of the impulsive and adiabatic processes, respectively, in our simulations. As outlined above, the key parameter determining the impact of DM self-interactions is the momentum transfer cross section σ_T/m_χ, while the density threshold for star formation $n_{\rm th}$ determines how impulsive SNF is. Our simulation suite consists of 16 simulations, for each of which we adopt a different combination of these

two model parameters. We thus cover a four by four grid in parameter space, running one simulation for each combination of $\sigma_T/m_\chi = \{0, 0.1, 1, 10\}$ (in units of $\mathrm{cm^2 g^{-1}}$) and $n_{\mathrm{th}} = \{0.1, 1, 10, 100\}$ (in units of $\mathrm{cm^{-3}}$).

The numerical values of the other parameters of the star formation and stellar feedback model are given in table 3 of Marinacci et al. (2019). For the SIDM algorithm, we adopt $N_{\mathrm{ngb}} = 32 \pm 5$ for the nearest neighbour search.

8.3 Results

In this Section, we present the results of our simulations. We start by showing how the numerical value of the star formation threshold affects the star formation history, both in a CDM and in an SIDM halo. Then, we compare the evolution of the density and velocity dispersion profiles of the DM halo for all 16 combinations of n_{th} and σ_T/m_χ. Thereafter, we focus on several dynamical quantities that help to break the degeneracy between simulations in which the final density profiles (and in turn the galaxy rotation curves) look nearly identical. Our simulations are run for a total of 4 Gyr. Most results presented in this Section are derived from snapshots taken after 3 Gyr, except for a few relevant cases in which we present results for 4 Gyr. We have verified that the differences between simulations with different transfer cross sections and star formation thresholds persist at later times.

8.3.1 Star formation histories

Across simulations, we find that the burstiness of star formation is regulated by the star formation threshold n_{th}. In Fig. 8.1, we show the star formation rates measured in eight different simulations over a simulated time of 3 Gyr. The top panel shows the star formation rates of all CDM simulations, while the bottom panel shows the star formation rates of all SIDM simulations with $\sigma_T/m_\chi = 1\,\mathrm{cm^2 g^{-1}}$.

On the top panel, we find quasi-periodic bursty star formation cycles only in the CDM simulation with $n_{\mathrm{th}} = 100\,\mathrm{cm^{-3}}$. In all other simulations, star formation decreases monotonously after ~ 300 Myr. We can identify a single star burst in the simulation with $n_{\mathrm{th}} = 10\,\mathrm{cm^{-3}}$, after ~ 250 Myr. The star formation histories in the two simulations with low star formation thresholds are smooth over the entire simulated time.

On the bottom panel, we find bursty star formation in both simulations with larger star formation thresholds, i.e. for $n_{\mathrm{th}} = 10\,\mathrm{cm^{-3}}$ and for $n_{\mathrm{th}} = 100\,\mathrm{cm^{-3}}$. For $n_{\mathrm{th}} = 10\,\mathrm{cm^{-3}}$, bursty episodes of star formation start appearing after ~ 1 Gyr. For $n_{\mathrm{th}} = 100\,\mathrm{cm^{-3}}$, on the other hand, we identify massive bursts of star formation only during the first gigayear of simulated time. After a particularly strong burst, the star formation rate drops significantly and does not recover. This drop is directly related to the strong star burst before. The large number of supernovae that occur shortly after this star burst drive a large amount of gas out of the galaxy, effectively shutting off star formation. In the two simulations with lower star formation thresholds we once again observe a smooth star formation history throughout the simulations, with a steadily decreasing star formation rate as more of the gas is converted into stars.

Overall, we find that bursty star formation can only occur in simulations with

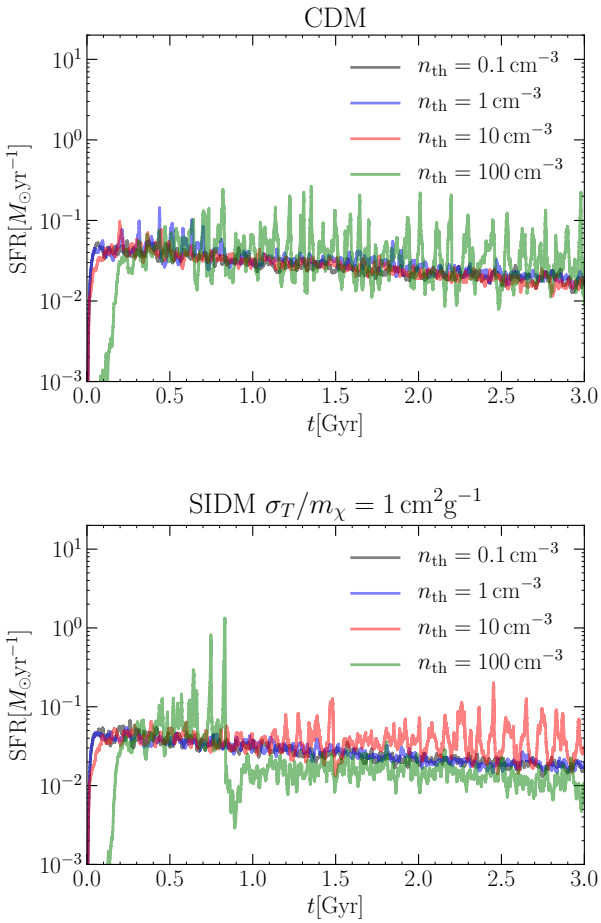

Figure 8.1. Star formation history in two different haloes over 3 Gyr of simulated time. On the top (bottom) panel we show the star formation rates in the simulations of the CDM (SIDM, $\sigma_T/m_\chi = 1\,{\rm cm}^2{\rm g}^{-1}$) halo with four different star formation thresholds as indicated. On average, larger star formation thresholds lead to burstier star formation, independent of the self-interaction cross section. However, star formation can also shut down in simulations with large star formation thresholds if gas is ejected from the galaxy early on (see for instance the green line on the right panel).

large star formation thresholds. For $n_{\rm th} \lesssim 1\,{\rm cm}^{-3}$, the star formation rate is smooth and monotonously decreases with time. The larger the star formation threshold, the burstier star formation can be. However, star formation also becomes more stochastic in simulations with larger star formation thresholds. In particular, SNF following a massive star burst can result in star formation being completely shut off, due to a large amount of gas being removed from the galaxy in large-scale galactic winds.

8.3.2 Density profiles and final galaxy rotation curves

We find striking differences in the evolution of the density and velocity dispersion profiles of the DM halo between simulations with different momentum transfer cross sections and star formation thresholds. Fig. 8.2 compares their evolution (density to the left, velocity dispersion to the right) for three different simulations. The top panels correspond to the CDM run with $\rho_{\rm th} = 0.1\,{\rm cm}^{-3}$, in the middle panels $\sigma_T/m_\chi = 1\,{\rm cm}^2{\rm g}^{-1}$ and $\rho_{\rm th} = 0.1\,{\rm cm}^{-3}$, and in the bottom panels we show results of the CDM run with $\rho_{\rm th} = 100\,{\rm cm}^{-3}$. Several profiles are shown in each panel, calculated from snapshots that are spaced apart by 1 Gyr of simulation time as labelled in the legend.

The DM density and velocity dispersion profiles show almost no evolution in the case in which both the star formation threshold and momentum transfer cross section are small (upper panels). In fact, the DM halo remains cuspy down to the smallest resolved radius. In the other two cases however, a constant density core forms in the inner halo. For $\sigma_T/m_\chi = 1\,{\rm cm}^2{\rm g}^{-1}$ and $n_{\rm th} = 0.1\,{\rm cm}^{-3}$ (middle panels), a $\sim 1\,$kpc core forms quickly and is fully formed after $\sim 2\,$Gyr. The corresponding velocity dispersion profile is flat out to approximately the scale radius of the initial halo. Density and velocity dispersion profiles of the CDM simulation with $n_{\rm th} = 100\,{\rm cm}^{-3}$ are displayed in the bottom panels. After $\sim 3\,$Gyr, the density profile closely resembles the SIDM density profile shown in the middle panel. However, the cusp-core transformation occurs more gradually and we see that the corresponding velocity dispersion profile is not yet fully isothermal at the end of the simulation. Thus, while the timescales for impulsive (SNF driven) and adiabatic (SIDM related) core formation are slightly different, the resulting cored density profiles look remarkably similar. As a consequence, we cannot differentiate between those two core formation scenarios by means of their final DM density profiles.

Crucially, DM density profiles are not observed directly but instead reconstructed from the measured rotation curves of observed galaxies. Santos-Santos et al. (2020) introduced a method to categorize rotation curves by comparing the maximal circular velocity $v_{\rm max}$ with the circular velocity v_C at a fiducial radius $r_{\rm fid} = 2(v_{\rm max}/70\,{\rm km/s})\,{\rm kpc}$ (see also Oman et al. 2015). The authors state that a value of $v_C(r_{\rm fid})/v_{\rm max} \sim 0.7$ is typical for cuspy NFW haloes and that larger values correspond, on average, to adiabatically contracted haloes, whereas smaller values correspond to cored haloes. However, Santos-Santos et al. (2020) also mention that while this ratio is a useful statistical measure to characterize rotation curves, it cannot be used to decide whether individual DM density profiles are cored or cuspy. Burger and Zavala (2021) investigated a system with similar distributions of dark and baryonic mass components as the one simulated here. The authors find that the baryonic component dominates the rotation curve at the fiducial radius, making it difficult to tell whether the DM density profile

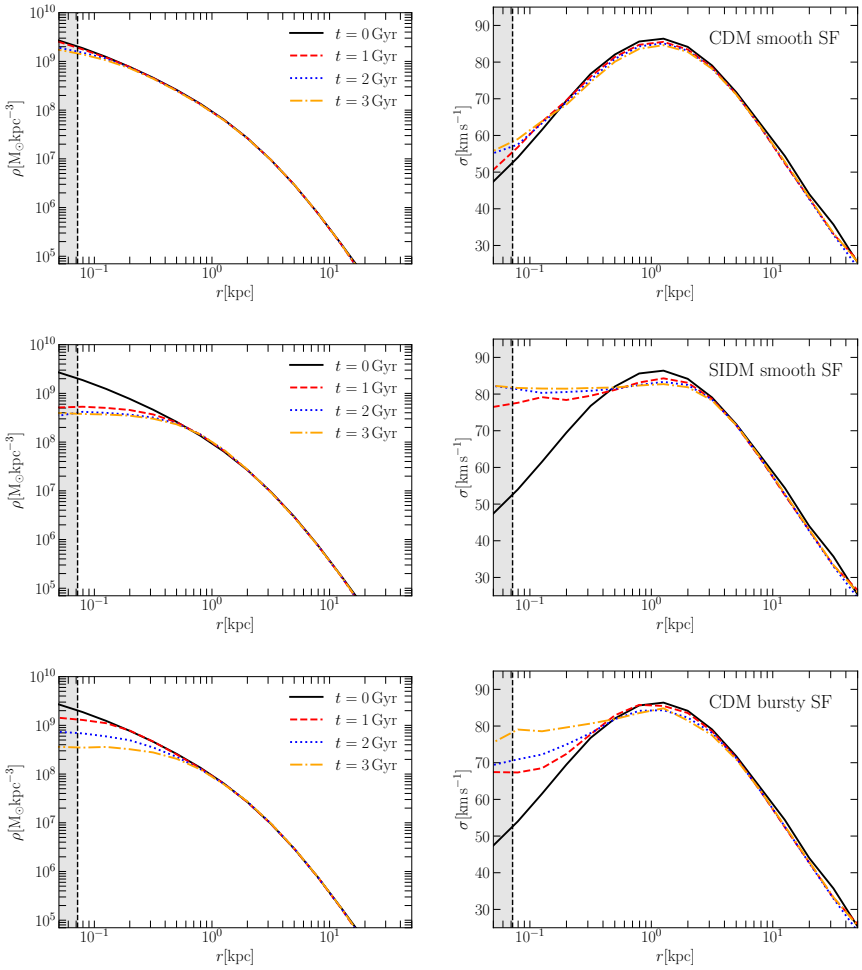

Figure 8.2. Evolution of density profile (left panels) and velocity dispersion profile (right panels) of the DM halo in three different simulations. In the top row we show results of a CDM run with $\rho_{\mathrm{th}} = 0.1\,\mathrm{cm}^{-3}$*, in the middle row* $\sigma_T/m_\chi = 1\,\mathrm{cm}^2\mathrm{g}^{-1}$ *and* $\rho_{\mathrm{th}} = 0.1\,\mathrm{cm}^{-3}$ *and in the bottom row we show results of a CDM run with* $\rho_{\mathrm{th}} = 100\,\mathrm{cm}^{-3}$*. We show spherically averaged profiles measured at the times indicated in the legends. The grey area corresponds to the radial range in which the profile is not converged according to the Power et al. (2003) criterion.*

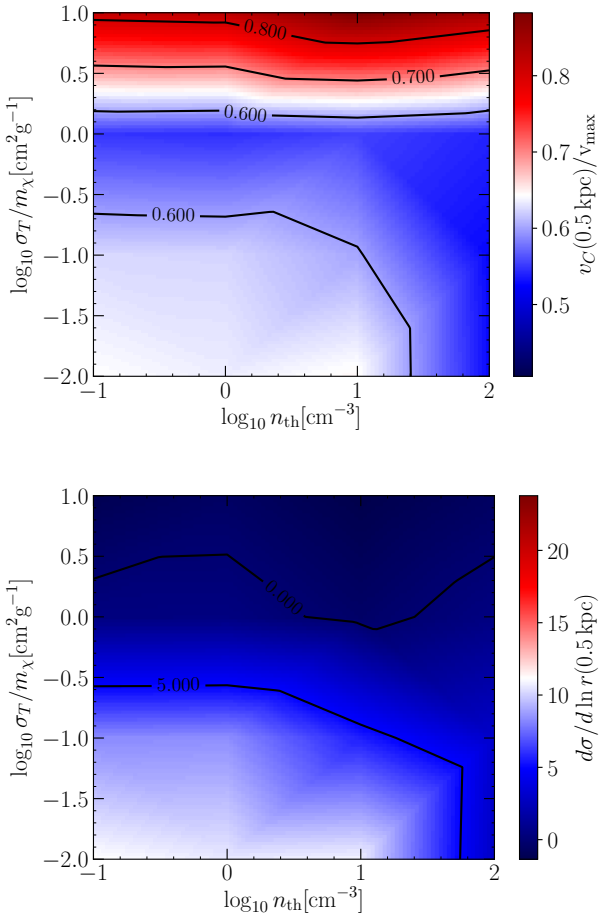

Figure 8.3. We show two measures that characterize cored and cuspy DM haloes at the characteristic radius of $0.5\,\mathrm{kpc}$ *in our simulations: the circular velocity profile relative to its maximum value (upper panel), and the logarithmic slope of the velocity dispersion profile. These are shown as a function of (logarithmic) star formation threshold and (logarithmic) transfer cross section per unit mass. We adopt a value of* $\sigma_T/m_\chi = 0.01\,\mathrm{cm}^2\mathrm{g}^{-1}$ *to represent CDM in this plot. The colour maps show bilinear interpolations over all 16 simulations. The contour lines show degenerate curves in parameter space along which* $v_C(0.5\,\mathrm{kpc})/v_{\mathrm{max}}$ *(upper panel) or* $d\sigma/d\ln r(0.5\,\mathrm{kpc})$ *(lower panel) assume the indicated values. The upper panel quantifies the (total) mass deficit in the inner part of the halo, whereas the lower panel focuses directly on whether the DM distribution in the inner halo is isothermal or not. Both panels correspond to results after* $3\,\mathrm{Gyr}$*. The colour map is chosen such that white colour corresponds to the CDM simulation with* $n_{\mathrm{th}} = 0.1\,\mathrm{cm}^{-3}$ *in which the DM halo remains cuspy (lower left corner). Deviations from this benchmark value in either direction are then coloured in red or blue as indicated by the colour bar.*

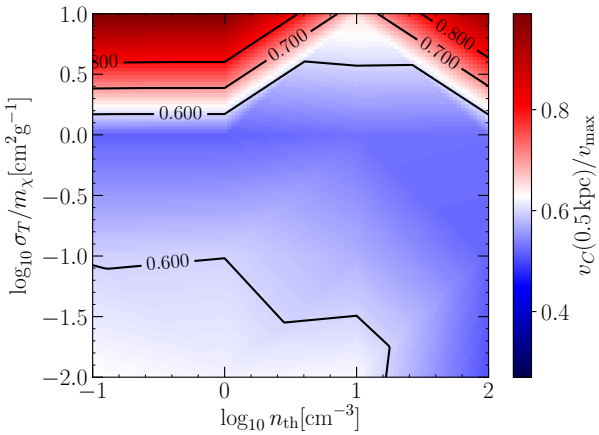

Figure 8.4. Same as upper panel of figure 8.3, but after 4 Gyr *of simulation time.*

is cuspy or cored from the value of $v_{\mathrm{fid}}/v_{\mathrm{max}}$. Thus, instead of v_{fid}, we here calculate the circular velocity at a smaller radius $r = 0.5\,\mathrm{kpc}$ since Fig. 8.2 suggests that cored profiles start to differ strongly from cuspy ones at this radius for our simulations. We therefore adopt $v_C(0.5\,\mathrm{kpc})/v_{\mathrm{max}}$ as a measure for how cored or cuspy our simulated DM halo is at a given time.

Fig. 8.3 shows $v_C(0.5\,\mathrm{kpc})/v_{\mathrm{max}}$ (upper panel) and $d\sigma/d\ln r(0.5\,\mathrm{kpc})$ (lower panel) as a function of (logarithmic) star formation threshold and (logarithmic) transfer cross section per unit mass after 3 Gyr of simulation time. The colour map is a bilinear interpolation in the (logarithmic) parameter space between all 16 simulations (see Section 8.2.4) and we have assigned a "self-interaction cross section" of $0.01\,\mathrm{cm}^2\mathrm{g}^{-1}$ to the CDM runs to be able to include them in the Fig.[2]. Fig. 8.3 demonstrates that there are curves in the σ_T/m_χ - n_{th} parameter space along which the measured values of $v_C(0.5\,\mathrm{kpc})/v_{\mathrm{max}}$ (or $d\sigma/d\ln r(0.5\,\mathrm{kpc})$) are degenerate. Some of these curves are highlighted by the contour lines. We have constructed all colour maps here such that they refer to deviations from the CDM simulation with $n_{\mathrm{th}} = 0.1\,\mathrm{cm}^{-3}$, in which the halo remains cuspy (see Fig. 8.2). Quantities measured for this benchmark simulation are assigned white colour, while deviations into either direction are coloured blue or red.

A couple of interesting trends emerge in the upper panel of Fig. 8.3. For low star formation thresholds, the final mass distribution is solely determined by the self-interaction cross section. Cross sections up to $\sim 0.1\,\mathrm{cm}^2\mathrm{g}^{-1}$ are rather ineffective at forming a core within a simulation time of $\sim 3\,\mathrm{Gyr}$. The most prominent cores are formed at values $\gtrsim 1\,\mathrm{cm}^2\mathrm{g}^{-1}$. However, for much larger cross sections an inversion of

[2] As mentioned above, CDM and SIDM are virtually indistinguishable for $\sigma_T/m_\chi \ll 0.1\,\mathrm{cm}^2\mathrm{g}^{-1}$. We thus expect no important differences between CDM runs and potential runs with $\sigma_T/m_\chi = 0.01\,\mathrm{cm}^2\mathrm{g}^{-1}$ and thus assume that setting them equal is justified for the sake of presentation.In the remainder of the article, CDM runs will be interpreted as SIDM runs with $\sigma_T/m_\chi = 0.01\,\mathrm{cm}^2\mathrm{g}^{-1}$ whenever interpolations over parameter space are presented.

this effect occurs. In fact, for $\sigma_T/m_\chi \sim 10\,\mathrm{cm^2 g^{-1}}$ we find that the final enclosed mass within 0.5 kpc is larger than in the baseline cuspy CDM case. This is due to the onset of the gravothermal collapse phase (Balberg et al., 2002; Colín et al., 2002; Koda and Shapiro, 2011; Pollack et al., 2015; Nishikawa et al., 2020).

When increasing the star formation threshold, we find that at some value between $10\,\mathrm{cm^{-3}}$ and $100\,\mathrm{cm^{-3}}$, the SNF mechanism becomes concentrated and impulsive enough to form a core in the DM profile that is of roughly the same size as the largest cores formed by SIDM. Combining a large star formation threshold with SIDM cross sections that would by themselves lead to the formation of cores does not change the value of $v_C(0.5\,\mathrm{kpc})/v_{\mathrm{max}}$ by much. Hence, the measured rotation curves are truly degenerate at this characteristic radius indicating that DM cores of similar size form all across for cross sections around $1\,\mathrm{cm^2 g^{-1}}$ (regardless of the star formation threshold) and for smaller cross sections as long as the star formation thresholds is large enough. On the other hand, if the SIDM transfer cross section per unit mass is large enough, the effect of gravothermal collapse always outweighs the effect of SNF, meaning that even at large star formation thresholds the circular velocity measured at the end of the simulation is always larger than in the benchmark simulation.

The lower panel of Fig. 8.3 aims to provide a measure of the dynamical differences between the 16 DM haloes after 3 Gyr. We show the derivative $d\sigma/d\ln r$ at $r = 0.5\,\mathrm{kpc}$, interpolated across the parameter space shown in the figure. For fully isothermal cores we expect values of $d\sigma/d\ln r \sim 0$. If we focus on the contour line in parameter space along which we found cored halo profiles in the upper panel of Fig. 8.3, we see that the behaviour of the velocity dispersion profiles to some degree breaks this degeneracy between SIDM cores and SNF cores. SIDM cores are in general more isothermal than their SNF counterparts. This is a generalization of the statement that SIDM and SNF can lead to similar core sizes, but their DM components have a different dynamical structure (see Burger and Zavala 2019), at least over the simulated time interval. While the bottom panels of Fig. 8.2 suggests that the core in CDM run with impulsive star formation becomes increasingly isothermal, we do not know whether a steady state similar to the quasi-equilibrium state of cored SIDM haloes will eventually be reached[3]. Throughout the simulations presented here, the dynamical structure of the DM haloes is different for different core formation scenarios. Finally, we note that larger SIDM cross sections can lead to slightly negative gradients in $\sigma(r)$ at 0.5 kpc, indicating that the velocity dispersion profiles are no longer fully isothermal after the onset of gravothermal collapse.

In some of our simulations, the DM haloes continue to evolve after 3 Gyr, and thus the picture presented in Fig. 8.3 changes slightly. For illustration, we show $v_C(0.5\,\mathrm{kpc})/v_{\mathrm{max}}$ measured after 4 Gyr as a function of σ_T/m_χ and n_{th} in Fig. 8.4. Two trends are apparent when comparing Fig. 8.4 to the upper panel of Fig. 8.3. Firstly, the SIDM haloes with $\sigma_T/m_\chi = 0.1\,\mathrm{cm^2 g^{-1}}$ are more cored after an extra 1 Gyr of evolution since despite being a weak value for the self-interaction, it can still cause DM haloes to form cores, on longer timescales. Secondly, the gravothermal collapse has progressed, heavily altering the dynamical structure of SIDM haloes with $\sigma_T/m_\chi = 10\,\mathrm{cm^2 g^{-1}}$, making them "cuspier" on average. The CDM runs do not change appreciably in the additional 1 Gyr, implying that residual evolution due to SNF occurs on longer

[3] We remark that no such steady state is found after 4 Gyr of simulated time either.

timescales. Importantly, the degeneracy contours in parameter space along which haloes with a flat constant density core are located very similarly in the upper panel of Fig. 8.3 and Fig. 8.4.

8.3.3 Galaxy sizes

The orbits of stars change in response to an evolving gravitational potential. However, their response may differ depending on whether the change in the potential is adiabatic or impulsive. Hence, the stellar mass distribution may differ between galaxies whose host haloes have formed a core either through SNF or through self-interactions between the DM particles.

Fig. 8.5 shows different measures of the size of the simulated galaxy (after 3 Gyr) as a function of star formation threshold and transfer cross section per unit mass. In the upper panel we show the half mass radius, whereas in the lower panel we show the quarter mass radius of the simulated galaxies. These radii are determined by calculating the enclosed stellar mass in spherical shells around the halo's centre of potential and then determining (using nested intervals and gsl "akima" interpolation) the radius of the spherical shell which contains half (a quarter of) the total stellar mass. When calculating the stellar mass profile, we take into account all collisionless disc and bulge particles (see Section 8.2.1, as well as newly formed "star" particles. The half mass (upper panel) radius is usually taken as a characteristic scale of galaxies. Given the simulation setup we have, in particular the values of the scale lengths for the gaseous and stellar discs chosen for the initial conditions in our simulations (see Section 8.2.1), the half mass radius ends up being larger than the typical DM core radius by a factor of ~ 1.5 and is therefore not ideal to analyse how core formation affects the stellar distribution. Hence, we also look at the quarter mass radius (lower panel of Fig. 8.5), which probes exactly the radial range of interest.

A few trends are similar across both panels of Fig. 8.5. At small SIDM cross sections ($\sigma_T/m_\chi \leq 0.1\,\mathrm{cm^2 g^{-1}}$), the star formation threshold hardly has any impact on the final galaxy size. Another key trend that can be observed in both panels is that the galaxy size contracts significantly in those simulations in which the self-interaction cross section is large enough to trigger the gravothermal catastrophe. The most significant contraction is observed for the case in which $\sigma_T/m_\chi = 10\,\mathrm{cm^2 g^{-1}}$ and $n_{\mathrm{th}} = 10\,\mathrm{cm^{-3}}$, where the gravothermal collapse proceeds somewhat faster than in simulations with the same self-interaction cross section and different star formation thresholds.

The most interesting feature of Fig. 8.5 appears only in the bottom panel. For $n_{\mathrm{th}} \leq 1\mathrm{cm^2 g^{-1}}$ the galaxy becomes more extended if $\sigma_T/m_\chi \sim 1\,\mathrm{cm^2 g^{-1}}$, precisely the case in which *i*) the DM density profile forms a core in an adiabatic way due to SIDM and *ii*) SNF does not cause impulsive changes in the gravitational potential. In simulations in which the DM halo adiabatically forms a core, the stellar tracers follow the adiabatic evolution of the gravitational potential, resulting in a less bright and more extended galaxy. Vogelsberger et al. (2014) performed SIDM simulations with $\sigma_T/m_\chi \sim 1\,\mathrm{cm^2 g^{-1}}$ with a baryonic physic implementation having effectively a low star formation threshold. The authors find that the stellar distribution of their simulation traces the evolution of the DM, forming a core that is related to the DM core (see Fig. 8 in Vogelsberger et al. (2014)), which is in very good agreement with the results we find

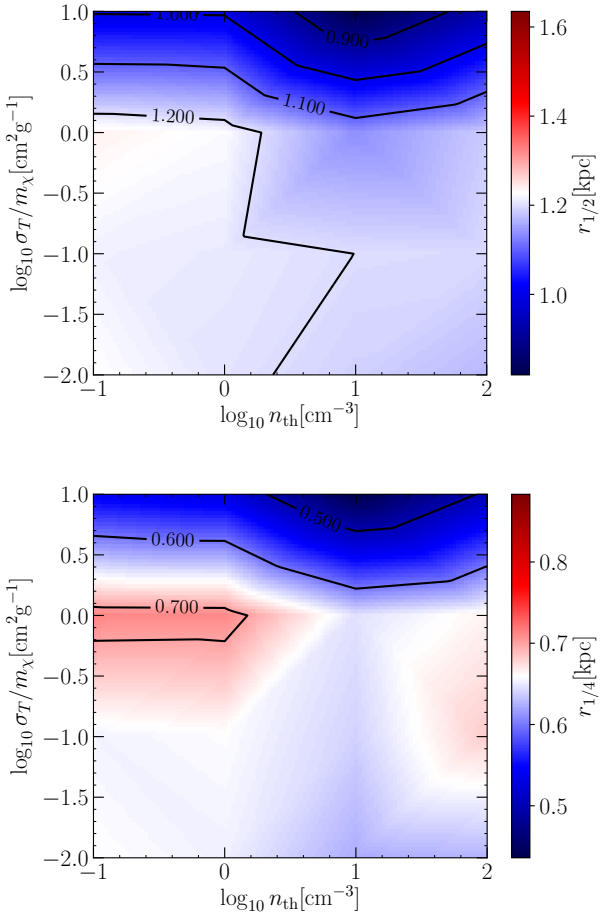

Figure 8.5. The galaxy half (stellar) mass radius $r_{1/2}$ (top panel) and quarter (stellar) mass radius $r_{1/4}$ (bottom panel) as a function of σ_T/m_χ and n_{th}, measured after 3 Gyr. *As in Fig. 8.3, the parameter space shown is filled by interpolation across the 16 simulations in our suite. Contour lines indicate degeneracies (i.e. equal values of either $r_{1/2}$ or $r_{1/4}$) in parameter space. Notice that the half mass radius in all cases is larger than the typical core radius of* ~ 0.7 kpc *found in Fig. 8.2. Thus, core formation has a much stronger effect on the quarter mass radius of the stars than on the half mass radius.*

here.

The fact that with respect to the bechmark simulation, no global changes in the stellar distribution are observed in simulations with impulsive SNF may indicate that impulsive fluctuations in the gravitational potential prevent the stars from following the adiabatic evolution of the DM component. Thus, even if the evolution of the DM is governed by the effect of SIDM when the cross section is $\gtrsim 1\,\mathrm{cm^2 g^{-1}}$, SNF remains an important perturber to the dynamics of the stars. The galaxy sizes measured after 4 Gyr are very similar to the ones presented in Fig. 8.5. The only difference is a strong additional spatial contraction of the galaxies in the simulations with $\sigma_T/m_\chi = 10\,\mathrm{cm^2 g^{-1}}$.

An important conclusion from Fig. 8.5 is that the spatial distribution of stars in dwarf galaxies can break the degeneracy in the rotation curves of galaxies within haloes with impulsive or adiabatic core formation histories (seen in Fig. 8.3). In order to use observations of the structure and size of galaxies to determine whether their inferred DM cores formed adiabatically or impulsively, it would be necessary to first establish a statistical relation between the dynamical mass of the system, the core formation mechanism and the average structure of the stellar distribution through an expanded suite of simulations like the one presented here but in a full cosmological setting within a large cosmic volume.

Inspired by the results of Burger and Zavala (2019), we now focus on the dynamical properties of the stars and the gas as a more promising avenue to search for observable and truly distinct signatures of either the adiabatic or impulsive cusp-core transformation scenarios.

8.3.4 Line-of-sight gas dynamics

Impulsive injection of energy, momentum, and ejecta mass from supernovae into the surrounding interstellar medium (ISM) causes random motion in the gas, perturbing its circular streaming motion. Fig. 8.6 shows the ratio between the line-of-sight speed and the line-of-sight rms (root mean squared) velocity of the gas in the CDM simulation in which $n_{\mathrm{th}} = 100\,\mathrm{cm^{-3}}$ (bursty star formation), calculated at different times as indicated in the legend. To calculate this ratio, we first determine the centre of potential from all simulation particles with a shrinking spheres method and subsequently calculate the total angular momentum vector from all particles which are part of the rotating disc (gas cells, “disc" particles, and “star" particles). We then use the normalized total angular momentum vector to rotate the galaxy into a coordinate system whose origin is the centre of potential and whose vertical axis is aligned with the normalized total angular momentum vector.

Afterwards, we calculate $|\bar{v}_{\mathrm{los}}|/\sigma_{\mathrm{los}}$ for an edge-on galaxy configuration. Since our initial conditions were set up with an axisymmetric baryonic disc, we here assume that our simulated discs remain axisymmetric by the end of the simulations. Without loss of generality, we can then assume that the line of sight is aligned with the y-axis in our new coordinate system.[4] Averaged line-of-sight speed and velocity dispersion are then calculated in logarithmic bins of cylindrical radius, with $\bar{v}_{\mathrm{los}} = \bar{v}_y$ and $\sigma_{\mathrm{los}} = \sigma_{v_y}$. Initially, the gas streams without any random motion, and thus the observed dispersion in Fig. 8.6 is simply a consequence of the projection of the circular motion performed

[4] We have verified that the results do not change significantly when choosing the x-axis instead.

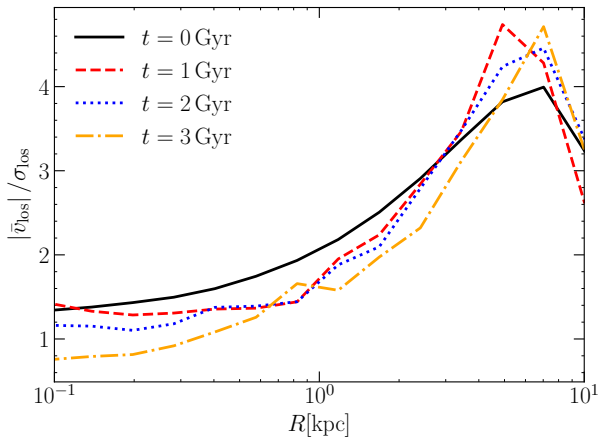

Figure 8.6. Evolution of the ratio of the average line-of-sight speed to the line-of sight velocity dispersion as a function of cylindrical radius for the CDM simulation with $n_{\mathrm{th}} = 100\,\mathrm{cm}^{-3}$ (bursty star formation). Different lines correspond to different times as indicated in the legend.

by the gas into the line-of-sight direction. As the simulation progresses, successive injection of momentum, energy and mass into the ISM causes additional random motion in the gas, in particular near the centre of the galaxy. After 3 Gyr, we see that the inner value of $|\bar{v}_{\mathrm{los}}|/\sigma_{\mathrm{los}}$ has dropped below one. The effect of the reduction is particularly significant at very small radii $\lesssim 200\,\mathrm{pc}$.

If this increase in random gas motion is in fact a direct consequence of SNF, then the different values of n_{th} used in our simulation suite should lead to a systematic difference in $|\bar{v}_{\mathrm{los}}|/\sigma_{\mathrm{los}}$ across simulations. In the upper (lower) panel of Fig. 8.7, we show $|\bar{v}_{\mathrm{los}}|/\sigma_{\mathrm{los}}$ measured at a cylindrical radius of 0.2 kpc from the centre of the galaxy after 3 Gyr (4 Gyr) as a function of star formation threshold and SIDM transfer cross section. Over large parts of the parameter space, the degree of random motion in the gas close to the centre of the galaxy is nearly constant throughout the simulations. However, in some cases with large star formation thresholds (impulsive SNF), the central value of $|\bar{v}_{\mathrm{los}}|/\sigma_{\mathrm{los}}$ is reduced after 3 Gyr (relative to the baseline cuspy CDM case with smooth star formation), indicating an increase in random motion within the gas.

This increase in random motion is particularly strong for $n_{\mathrm{th}} = 100\,\mathrm{cm}^{-3}$ and $\sigma_T/m_\chi \leq 0.1\,\mathrm{cm}^2\mathrm{g}^{-1}$. The strength of this distinction between smooth and bursty star formation (in CDM and SIDM with $\sigma_T/m_\chi \leq 1\,\mathrm{cm}^2\mathrm{g}^{-1}$) is however quite dependent on the simulation time. For instance, in the lower panel of Fig. 8.7, which corresponds to $t = 4\,\mathrm{Gyr}$, the difference across different star formation thresholds is much smaller than in the upper panel. This implies that impulsive SNF can lead to a significant increase in the random motion of the gas that is rather short lived. Thus, while hypothetical observations of very chaotic gas motion in dwarf galaxies with cored host haloes would hint at a recent impulsive starburst event, SNF cannot be ruled out as

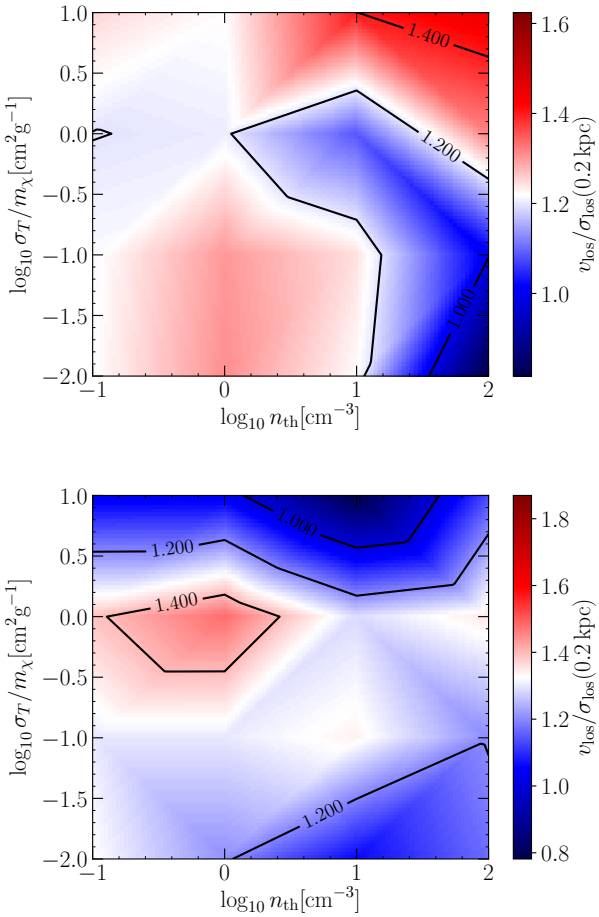

Figure 8.7. A measure of the relative impact of random motions over rotation: the ratio $|\bar{v}_{\mathrm{los}}|/\sigma_{\mathrm{los}}$*, after* 3 Gyr *(top panel) and* 4 Gyr *(bottom panel). We show this ratio at a cylindrical distance of* 200 pc *from the centre of the galaxy as a function of the parameters* n_{th} *and* σ_T/m_χ*. A tendency towards smaller ratios (i.e. larger random motion) arises in simulations with cored haloes and bursty star formation (*$n_{\mathrm{th}} \geq 10\,\mathrm{cm}^{-3}$*), i.e., roughly within the contour line of value 1.2, which is roughly the baseline cuspy CDM value with smooth star formation. The strength of this trend is however transitory, being significantly stronger at* 3 Gyr *(top panel) than at* 4 Gyr *(bottom panel). In the latter case, we find that random gas motion increases significantly in the central region of galaxies whose host haloes undergo gravothermal collapse.*

the cause of a DM core if no such increased random motion is observed.

We notice as well that due to the increasing impact of the gravothermal collapse phase on the central properties of the galactic system, there is a strong increase in random gas motion in the simulations with $\sigma_T/m_\chi = 10\,\mathrm{cm^2 g^{-1}}$ by $t = 4\,\mathrm{Gyr}$. Therefore, we speculate that hypothetical observations of compact dwarf galaxies with very fast-rising rotation curves and a large amount of random gas motion would be good targets to look for the presence of haloes that have gravothermally collapsed due to very large SIDM cross sections.

By comparing the lower panel of Fig. 8.5 with the results of Fig. 8.7 we can make an interesting remark: while an average expansion of the central stellar distribution appears to be a distinct signature of galaxies embedded in DM haloes with adiabatically formed cores which are not significantly affected by SNF ($r_{1/4} \sim 0.7\,\mathrm{kpc}$ in Fig. 8.5), a significant increase in random motion of the gas in the centre of the galaxy can only be observed in galaxies with cored host haloes if SNF is impulsive (right corner in both panels of Fig. 8.7 relative to the CDM case with smooth star formation baseline in the lower left corner).

8.3.5 Age and metallicity gradients

The stellar evolution module of *SMUGGLE* keeps track of several properties of individual star particles, among them, their formation time and metallicity. Fig. 8.8 shows projections of the phase space distribution of the stellar age and metallicity of newly formed stars after 3 Gyr in the benchmark CDM simulation with $n_{\mathrm{th}} = 0.1\,\mathrm{cm^{-3}}$ (smooth star formation). In the upper (lower) panels, we show the average stellar age (metallicity) as a function of the phase space coordinates R and v_R (left panels) and z and v_z (right panels). For presentation purposes, the data is averaged in 100×100 equally spaced bins in phase space. The average stellar age appears to be approximately independent of the cylindrical radius R (upper left panel of Fig. 8.8). However, older stars seem to be on orbits with a relatively large vertical extent (upper right panel of Fig. 8.8). This hints at a slow migration of stars out of the disc plane over the course of the simulation. In contrast, we observe a clear radial and vertical gradient in metallicity, which has a straightforward physical explanation. As outlined in Section 8.2.2, metals that are ejected by supernovae are distributed among the neighbouring gas cells. Since the gas is initially densest in the centre of the galaxy, this is where most stars form and hence where most supernovae happen. This larger supernova rate, leads to a more metal-rich ISM in the inner galaxy and in turn to second generation stars with larger metallicities than in the the outskirts of the galaxy, explaining the observed metallicity gradient.

The stellar age and metallicity distributions shown in Fig. 8.8 are calculated after 3 Gyr in the CDM simulation with smooth star formation in which the DM halo does not form a core (see Fig. 8.2), i.e., the gravitational potential remains approximately constant. On the other hand, in a halo with an evolving gravitational potential, (adiabatic or impulsive) cusp-core transformation can alter the phase space distribution of the stars in dwarf galaxies (Burger and Zavala, 2019). Thus, we surmise that the age and metallicity distributions will look distinctly different in our simulations in which the DM haloes' final density profiles are cored.

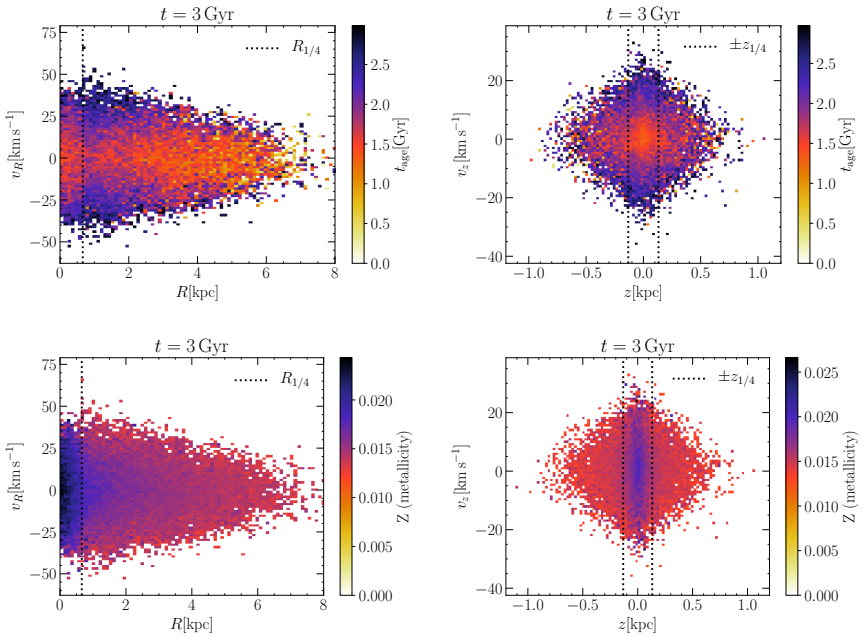

Figure 8.8. Average age (upper panels) and metallicity (lower panels) of stars formed after $3\,\mathrm{Gyr}$ *in the CDM simulation with* $n_{\mathrm{th}} = 0.1\,\mathrm{cm}^{-3}$ *(smooth star formation) as a function of phase space coordinates: cylindrical radius R and velocity* v_R *(left panels), vertical coordinate z and the vertical velocity* v_z *(right panels). The data is averaged in phase space bins. Black dotted vertical lines denote either the quarter mass radius (*$R_{1/4}$*, left column) or* $z_{1/4}$ *(right column, see text).*

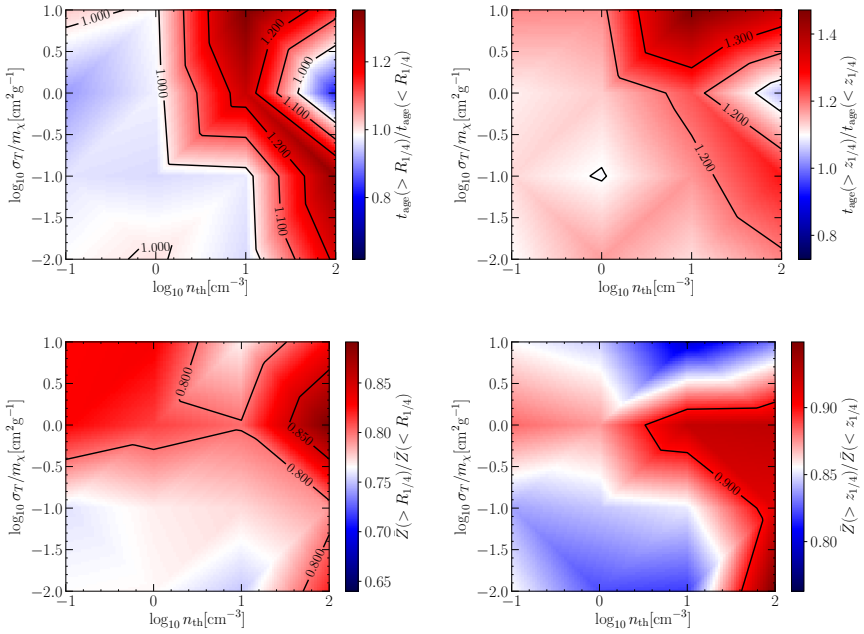

Figure 8.9. Different ratios characterizing age gradients and metallicity gradients in our simulations as a function of star formation threshold and SIDM transfer cross section, calculated after 3 Gyr. *In the upper left panel, we show the ratio between average stellar ages in the outer and inner regions of the simulated galaxy across the disc plane, where the quarter mass (cylindrical) radius (of the baryonic distribution) is used as the boundary. The upper right panel shows a similar plot, but vertically, perpendicular to the plane of the disc, with a boundary for inner and outer regions given by* $z_{1/4} = 0.2R_{1/4}$ *(see also Fig. 8.8). The lower panels are as the upper panels but for stellar metallicities instead of ages. Since the quarter mass radius lies well within the radial range which is strongly affected by core formation (see Fig. 8.2), these ratios are a good probe of how adiabatic or impulsive core formation mechanisms affect the stellar phase space distributions.*

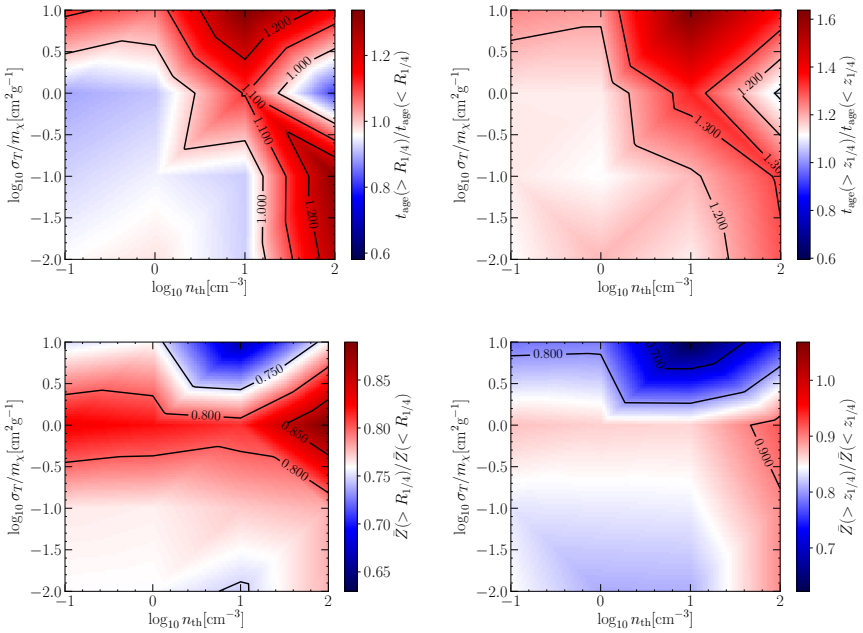

Figure 8.10. Same as Fig. 8.9, but after 4 Gyr *of simulation time.*

Figs. 8.9 and 8.10 compare the age and metallicity gradients measured after 3 Gyr and 4 Gyr, respectively across the parameter space of the simulations. In order to facilitate a comparison between different values of σ_T/m_χ and $n_{\rm th}$, we quantify the steepness of the gradients by taking the age/metallicity ratios between stars with radial/vertical distances from the centre of the galaxy which are larger than a certain characteristic scale to those that are smaller than this scale. From Figs. 8.3 and 8.5, we know that the enclosed mass within the baryonic quarter mass radius ($r_{1/4}$) is a good proxy for whether the DM halo is cored or cuspy, and thus, we choose this radius as the characteristic scale that characterizes the age/metallicity gradients in Figs. 8.9 and 8.10. Since we are interested in the gradients along the disc plane and perpendicular to it, we define the characteristic cylindrical radius $R_{1/4} = r_{1/4}$ at $z = 0$, and the characteristic vertical scale $z_{1/4} = 0.2R_{1/4}$ in line with the initial ratio between vertical and radial scale height of the disc (see also Fig. 8.8).

The upper left panel of Fig. 8.9 shows the ratio between the average age of stars with cylindrical radii larger than $R_{1/4}$ and stars with radii smaller than $R_{1/4}$ at $t = 3$ Gyr. For $n_{\rm th} \leq 1\,{\rm cm}^{-3}$ this function is featureless, while at larger star formation thresholds, the age gradient is positive and significantly larger than in the benchmark case for almost all simulations, particularly for those with $n_{\rm th} \geq 10\,{\rm cm}^{-3}$ (bursty star formation) in which the DM halo forms a core. However, for $n_{\rm th} = 100\,{\rm cm}^{-3}$ and $\sigma_T/m_\chi = 1\,{\rm cm}^2{\rm g}^{-1}$, the age gradient is reversed, becoming negative at the scale of the quarter mass radius: $t_{\rm age}(> R_{1/4})/t_{\rm age}(< R_{1/4})$. The general trend of simulations with larger star formation thresholds having (on average) older stars in the outer parts of the galaxy than in the

inner parts can be explained by the mechanism of impulsive SNF.

The rapid change in the gravitational potential caused by the supernova cycle triggered in early starbursts that causes the formation of the DM core also results in an outward migration of some of the stars that were present in the inner regions at the time. The exception seen for the simulation with $n_{\rm th} = 100\,{\rm cm}^{-3}$ and $\sigma_T/m_\chi = 1\,{\rm cm}^2{\rm g}^{-1}$ is likely caused by a massive starburst occurring just before $t = 1\,{\rm Gyr}$ (see Figs. 8.1 and 8.11), which leads to a SN-driven gas outflow that effectively shuts off star formation for a long time. At later times, the supernova feedback mechanism is not energetic enough to cause older stars to migrate into the outskirts of the galaxy. The net result is that at the end of this simulation, the stellar population within $R = R_{1/4}$ is older (on average) than outside of it. The upper left panel of Fig. 8.10, showing the age ratio after 4 Gyr of simulation time, essentially confirms the same picture. However, given the growing impact of the gravothermal collapse phase for large cross sections, the age gradient is noticeably different (larger than one) than for smaller cross sections, even for low star formation thresholds (non-impulsive supernova feedback).

The upper right panels of Figs. 8.9 and 8.10 show a similar ratio as in the upper left panel, but in this case perpendicular to the plane of the disc instead of along it. Displayed is the average age of stars with $|z| > z_{1/4}$ divided by the average age of stars with $|z| < z_{1/4}$. At both times shown in Figs. 8.9 and 8.10, the age gradient in the vertical direction (upper right panels) shows a very similar pattern as in the (cylindrical) radial direction on the upper left panels.

The lower left panels of Figs. 8.9 and 8.10 show the ratio of average stellar metallicities outside of $R_{1/4}$ to average stellar metallicities within $R_{1/4}$ after 3 Gyr and 4 Gyr, respectively. After 3 Gyr, this ratio is a mostly featureless constant across the parameter space of the simulations, except for a slightly reduced metallicity gradient in the simulation with $\sigma_T/m_\chi = 1\,{\rm cm}^2{\rm g}^{-1}$ and $n_{\rm th} = 100\,{\rm cm}^{-3}$, where star formation is strongly reduced for a while and almost no stars with high metallicities are formed. The spatial metallicity distribution is fairly even in this case, compared to the other simulations. In general, metallicity gradients are negative, since the ISM is more metal-rich in the centre of galaxies.

This is reflected in metallicity ratios (measured with respect to the quarter mass radius) which are smaller than 1. In simulations with $\sigma_T/m_\chi = 1\,{\rm cm}^2{\rm g}^{-1}$, where the DM halo forms an isothermal core due to DM self-interactions, the observed metallicity gradients (along the cylindrical radial direction) are shallower than in other simulations. This is particularly apparent at $t = 4\,{\rm Gyr}$ where a clear difference emerges between simulations in which the DM haloes have formed cores adiabatically, and simulations in which they have not. A potential explanation for this arises from Fig. 8.5. Galaxies which are hosted by haloes with SIDM-induced cores have stellar distributions in which the central density of stars is smaller (i.e. the galaxies are less compact) than in the other simulated galaxies. Thus, the ejected metals from SNF are more evenly distributed within a larger volume around the centre of the galaxy.

Finally, the lower right panel of Fig. 8.9 (Fig. 8.10) shows the ratio between the average metallicity of stars with $|z| > z_{1/4}$ and stars with $|z| < z_{1/4}$ after 3 Gyr (4 Gyr). As in the case of the age gradients in the upper right panels of Figs. 8.9 and 8.10, we find that the metallicity gradients along the vertical direction track more or less closely the trends and patterns of the (cylindrical) radial metallicity gradients (compare lower

right to lower left panels of Figs. 8.9 and 8.10). Overall, however, the differences across the parameter space in the vertical direction are not as pronounced as in the (cylindrical) radial direction.

In summary, we can draw two conclusions. Very large stellar age gradients (to a lesser extent tracked by metallicity gradients) are found in the simulations in which supernova feedback is impulsive. Observing them in galaxies whose DM haloes have cored density profiles does not rule out SIDM as the mechanism responsible for the formation of the DM cores, but it suggests that supernova feedback is the dominant mechanism of cusp-core transformation. If very small metallicity gradients were ubiquitously observed in galaxies with cored host haloes, these haloes would likely have formed their cores adiabatically, and thus SIDM would be preferred over supernova feedback as a mechanism of cusp-core transformation.

8.4 Summary

We have explored under which conditions either SNF or SIDM are feasible mechanisms of cusp-core transformation in the haloes of isolated dwarf galaxies. Particularly, we focused on how the dynamical properties of gas and stars in dwarf galaxies might be affected by either impulsive (SNF driven) or adiabatic (SIDM driven) core formation.

To that end, we performed a suite of 16 high-resolution hydrodynamical simulations, evolved for 4 Gyr, of an idealized SMC-size galaxy embedded within a live DM halo with an initially cuspy Hernquist density profile. Our simulations included both a stellar evolution and feedback prescriptions using the *SMUGGLE* model (Marinacci et al., 2019) and self-interactions between the DM particles (Vogelsberger et al., 2012), all within the framework of the AREPO code (Springel, 2010). We present a detailed description of core formation in *SMUGGLE* for CDM haloes in a companion paper (Jahn et al., *in-prep*) and focus here on the comparison between dark matter cores formed due to SNF versus due to dark matter self-interactions, all evolved with the same *SMUGGLE* baryonic treatment.

Starting from identical initial conditions, each simulation was performed with a different combination of SIDM momentum transfer cross section (σ_T/m_χ) and star formation threshold ($n_{\rm th}$). These two parameters effectively control the efficiency of the SIDM-driven and SNF-driven mechanisms of cusp-core transformation, respectively. The values of these parameters were chosen in order to probe star formation regimes from smooth (low $n_{\rm th}$) to bursty (high $n_{\rm th}$, see Fig. 8.1), as well as to probe the regimes from collisionless DM (CDM; $\sigma_T/m_\chi = 0$) to strong self interactions as large as $10\,{\rm cm}^2{\rm g}^{-1}$. Our results indicate for which combinations of self-interaction cross section and star formation threshold, the initially cuspy halo develops a $\mathcal{O}(1)$ kpc size constant density core (see Fig. 8.3). In particular, we identified a degenerate line in the $\sigma_T/m_\chi - n_{\rm th}$ parameter space plane along which the final simulated DM haloes are cored.

Moreover, we found that adiabatically formed cores (SIDM cores) tend to be fully isothermal, while those formed through impulsive SNF are not, at least within the timescales of our simulations. To be more precise, our results indicate that SIDM cores fully thermalize significantly faster than those formed through SNF (see bottom panel

of Fig. 8.3). SIDM haloes with $\sigma_T/m_\chi \sim 10\,\mathrm{cm^2 g^{-1}}$ undergo gravothermal collapse after ~ 2.5 Gyr in our simulations. Their density profiles are cored for a short while before they collapse and form very steep central density cusps (see Figs. 8.3 and 8.4).

To differentiate between SIDM and SNF as core formation mechanisms, we compare several observable quantities between simulations. A few clear trends emerge. Galaxies within cored host haloes form extended stellar distributions that follow the gravitational potential of the host halo if *i*) the core was formed adiabatically through SIDM and *ii*) star formation is smooth instead of bursty, i.e., SNF is not impulsive (Fig. 8.5). Impulsive SNF can cause steep stellar age gradients (Figs. 8.9 and 8.10 and increased random motion in the gas (Fig. 8.7). Ubiquitous observations of turbulent gas or steep stellar age gradients within cored DM haloes would therefore suggest that impulsive SNF has caused the cusp-core transformation.

On the other hand, the metallicity gradients of stars in dwarf-size SIDM haloes with $\sigma_T/m_\chi \sim 1\,\mathrm{cm^2 g^{-1}}$ are systematically shallower than the metallicity gradients of stars in CDM haloes (Figs. 8.9 and 8.10). The dynamical properties of galaxies embedded in haloes which have undergone gravothermal collapse are systematically different from the dynamics of all other simulated galaxies.

The parameter space plots in our results are to be indicative of general trends, not precise predictions due to the fact that our simulations are idealised to some extent. To obtain an accurate quantitative understanding of the degeneracies/interplay between SNF and SIDM cross section in the inner structure of dwarf-size haloes, as well as the role of mergers, a comprehensive exploration of the $(n_{\rm th}, \sigma_T/m_\chi)$ parameter space in a full cosmological setting is required. Although such an undertaking is computationally expensive, requiring large suites of cosmological simulations with sufficiently high resolution, our idealized runs strongly suggest it would be fruitful, leading to detailed predictions regarding the properties of the visible components (gas and stars) that are truly distinct between these mechanisms of cusp-core transformation based on either baryonic physics or new DM physics.

On the observational front, searching for the trends we have found in this work could prove to be quite significant to understand how dwarf-size DM haloes develop cores. For instance, if large age gradients were observed in most dwarf galaxies with cored host haloes, SNF would likely be impulsive. This would strongly suggest that SNF is the main mechanism that drives the cusp-core transformation in these galaxies. Finally, studies of the dynamical properties of kinematic tracers (e.g. Burger and Zavala 2019) may reveal whether SNF is impulsive enough to be a feasible mechanism of cusp-core transformation, provided we have a way of identifying orbital families of stars in observational data.

8.5 Appendix A - Caveats

In this article, we have presented a suite of 16 high-resolution hydrodynamical simulations of an isolated SMC-size dwarf galaxy in a live DM halo. Across simulations, we have changed the star formation threshold $n_{\rm th}$, regulating the "burstiness" of star formation and hence the ability of SNF to drive large-scale gas outflows that rapidly change the gravitational potential. Furthermore, we have tested the impact of self-interactions

between the DM particles by changing the momentum transfer cross section per unit mass across simulations. By exploring this two-dimensional parameter space, we can study the two most viable mechanisms of core formation in DM haloes: adiabatic due to the impact of SIDM, and impulsive due to supernova feedback. We have then searched for differences in the distribution function of gas and stars across the simulations. In the following, we briefly discuss some caveats of our analysis. In particular, we outline how a different choice of initial conditions, and of the values of the gravitational softening may impact our results. Moreover, we discuss how to interpret the results presented in Section 8.3 in light of the fact that a single snapshot (time output) of each simulation in our simulations suite corresponds to only a single possible realization of the evolved distribution function.

8.5.1 Initial conditions

Since our simulations are not cosmological, the simulated halo has no cosmological assembly history. In cosmological simulations, haloes and galaxies form from initial conditions that are not arbitrary, but (statistically) fixed by the assumed cosmological model and constrained on large scales by the observed perturbations in the cosmic microwave background, which set the cosmic fractions of dark matter and baryons. The relative amount of DM or baryons in a galaxy is thus a prediction of cosmological simulations which is obtained from the full structure formation and evolution process, coupled to the baryonic physics model, and not an initial condition. With our choice of initial conditions, we aim to mimic an isolated virialized dynamical system that is similar (in scale) to the Small Magellanic Cloud, with initial structural parameters as in Hopkins et al. (2012).

With our choice of initial conditions, the structure of the halo and the galaxy, as well as the relative amounts of DM, gas, and stars are fixed. Since the efficiency of DM self-interactions as a mechanism of cusp-core transformation depends on the density of DM in the central halo, whereas the efficiency of SNF depends on the relative amount of baryonic matter in the central galaxy, our choice of initial conditions can thus have a large impact on the efficiency of SIDM or SNF as mechanisms of cusp-core transformation For that reason, our isolated simulations can make no definitive quantitative statements about the exact properties (including formation timescales) of DM cores formed through these mechanisms in realistic dwarf-size haloes formed in a cosmological setting.

For SIDM, the observed core formation for $\sigma_T/m_\chi \sim 1\,\mathrm{cm^2 g^{-1}}$ (see middle panels of Fig. 8.2) is approximately in agreement with cosmological simulations of SIDM haloes (see e.g. Vogelsberger et al. 2012 and Vogelsberger et al. 2014). We therefore conclude that the predictions of our SIDM simulations, including the timescales for core formation and the adiabatic nature od the cusp-core transformation, are reasonable.

For SNF, the situation is more complex. Most hydrodynamical cosmological simulations find that SNF can form cores only in the mass range of bright dwarfs (see e.g. Lazar et al. 2020). However, Read et al. (2016) find, based on high-resolution simulations of isolated haloes, that SNF can form cores event at the scale of ultra faint galaxies. Crucially, it is unclear whether the results of cosmological simulations are more correct in this regime or not. The baryon fractions assumed by Read et al. (2016)

are rather large, and it is possible that they cannot be realized in haloes with a cosmological formation history. However, due to the high resolution of Read et al. (2016)'s simulations, the effects of SNF can be modeled much more accurately than in large cosmological simulations. using modern ISM and stellar feedback models. As far as our initial conditions are concerned, it is worth mentioning that the assumed stellar mass of our SMC-size system is at the upper end of what is allowed by abundance-matching results for the stellar-to-halo mass relation (see e.g. Moster et al. 2010). In principle, this may imply increased efficiency of SNF. However, our galaxy is also very extended, and hence star formation is not very concentrated towards the centre of the halo, which has an adverse effect on the core formation efficiency of SNF (Burger and Zavala 2021).

8.5.2 Gravitational softening and concentration of baryons

The choice of the gravitational softening length can significantly affect how DM haloes respond to SNF with a large star formation threshold. Dutton et al. (2020) suggest that for large star formation thresholds, small softening lengths need to be adopted in order for SNF to efficiently transform cusps into cores. For collisionless simulations of gravitationally self-bound haloes, Power et al. (2003) have conducted a convergence study and derived an optimal force softening length

$$\varepsilon_{\rm opt} = 4\frac{r_{200}}{\sqrt{N_{200}}}, \tag{8.6}$$

where r_{200} is the halo's virial radius and N_{200} is the number of DM simulation particles contained within r_{200}. For our simulated halo, this corresponds to $\varepsilon_{\rm opt} \sim 50\,{\rm pc}$.

In hydrodynamical simulations, however, the choice of the force softening length is less clear. Before running our final simulation suite, we tested different simulation settings. In particular, we ran the CDM simulation with $n_{\rm th} = 100\,{\rm cm}^{-3}$ with different choices for the gravitational softening length and found that the cusp-core transformation does not occur for $\varepsilon = \varepsilon_{\rm opt}$. Our final choice of $\varepsilon = 24\,{\rm pc} \sim 0.5\,\varepsilon_{\rm opt}$ is on the higher end of the force softening lengths for which core formation does occur in the CDM run with $n_{\rm th} = 100\,{\rm cm}^{-3}$. The reason for this dependence on the force softening is simple. In runs with larger star formation thresholds, gas needs to be denser for stars to form. Hence, larger gravitational forces need to be resolved on small scales, for which smaller softening lengths are required.

Apart from the gravitational softening length, how baryonic matter is initially distributed within the inner DM halo can also significantly change the impact of SNF on the inner DM distribution. We performed a CDM test run with $n_{\rm th} = 100\,{\rm cm}^{-3}$ in which we omitted the stellar bulge when setting up the initial conditions (see Section 8.2.1). In this simulation, significantly less stars formed compared to the simulation in which the stellar bulge is included in the initial conditions. Moreover, star formation was less concentrated towards the centre of the halo. As a consequence, the DM density profile remained cuspy. In part, this is explained simply by the reduction in star formation which inevitably means reduced SN activity. Additionally, SNF which is less concentrated is less effective at forming cores (Burger and Zavala 2021).

The question then arises why the inclusion of the stellar bulge changes the picture this much, despite it accounting for only half a per cent of the total baryonic mass in the simulation. The answer must be that without it, the gradient of the gravitational

potential in the inner halo is too shallow, since both the stellar and the gaseous disk are very extended and have no appreciable density gradient towards the centre. Including the very concentrated bulge generates a steeper gradient in the central potential, and causes cooling gas to fall into the centre and reach the large densities required for star formation.

We thus stress that our results should not be understood as absolute predictions. Instead, our goal is to study in controlled/idealised simulations the (key) parameter space of the two cusp-core transformation mechanisms: SNF and DM self-interaction, and explore the similarities and differences between DM cores formed in these scenarios.

8.5.3 Interpreting our results

Most of our results are presented in the $n_{\rm th} - \sigma_T/m_\chi$ parameter space, bilinearly interpolated from the 16 simulations in our suite with each interpolation point corresponding to an estimate of the outcome of an actual simulation in that point in parameter space. There is one significant caveat to this way of presentation which affects how our results should be interpreted. We use hydrodynamical simulations to evolve the distribution function of a self-gravitating system in time, starting from well-defined initial conditions. Each snapshot which is taken at a later time corresponds to a single realization of the distribution function at that time, when in reality an ensemble average of different realizations would be required to determine the most likely evolved state of the system.

In DM-only simulations, this is not a very big issue since gravity is fully deterministic in the sense that different DM-only simulations will produce essentially the same results, above the (coarse-grained) scale resolved in the simulations. However, the stellar evolution model used in the suite of hydrodynamical simulations presented in this article is stochastic. In particular, star formation and supernovae are implemented as probabilistic random processes (see Sections 8.2.2). Therefore, two simulations defined by identical parameters and starting from the same initial condition are not guaranteed to produce the same results. Instead, the evolved state of a system at a fixed time may differ between such simulations. It is thus possible that some of the signatures presented in Section 8.3 may be statistical outliers whose occurrence is related to the stochastic implementation of star formation and supernova feedback.

Within our suite of 16 simulations, two of them yield results that are "peculiar" when compared to the other simulations. The first one is the simulation in which $\sigma_T/m_\chi = 10\,{\rm cm^2 g^{-1}}$ and $n_{\rm th} = 10\,{\rm cm^{-3}}$. Here, the onset of the gravothermal collapse phase is triggered earlier than in the other simulations with the same self-interaction cross section. As a result, essentially all observables that we consider in this article are very different between this simulation and all other simulations, in particular after 4 Gyr of simulation time (see Fig. 8.4 for the rotation curve, the lower panel of Fig. 8.7 for the gas random motion, and Fig. 8.10 for age and metallicity gradients). Particularly striking are the steep age and metallicity gradients of the stars in this simulation.

The second simulation whose results are strikingly different from that of the other runs is the one with $\sigma_T/m_\chi = 1\,{\rm cm^2 g^{-1}}$ and $n_{\rm th} = 100\,{\rm cm^{-3}}$. In particular, it is the only simulation in which SNF is impulsive and the age gradient of stars in the simulated galaxy is negative: on average, stars at larger radii are younger than at smaller radii. This is a rather unexpected feature, since stars should form at all radii all the time and

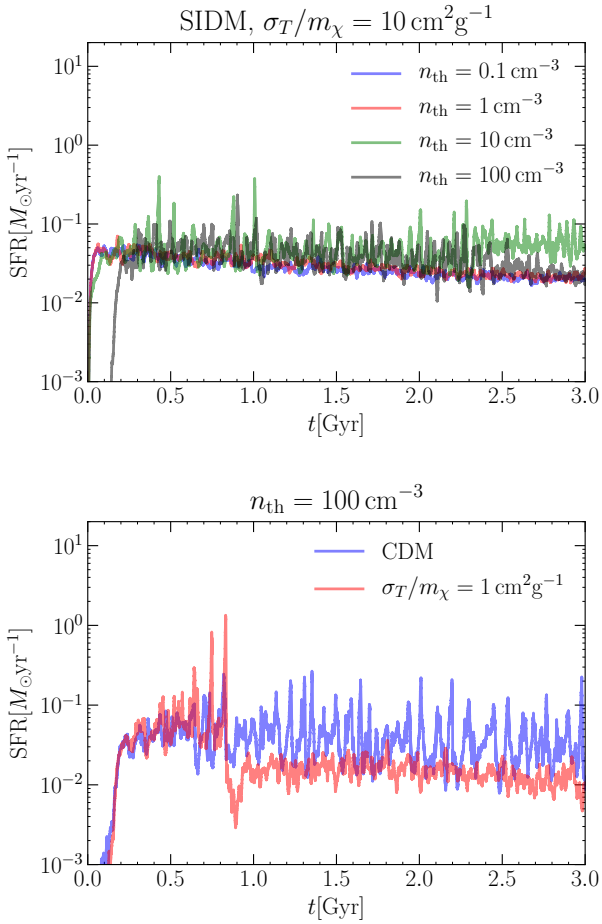

Figure 8.11. Star formation rate as function of time (star formation history) for different simulations. On the top panel we show the star formation histories for all SIDM simulations with $\sigma_T/m_\chi = 10\,\mathrm{cm}^2\mathrm{g}^{-1}$. The star formation thresholds are as indicated in the legend. Notice the bursty star formation history of the simulation with $n_{\rm th} = 10\,\mathrm{cm}^{-3}$, with four star burst events during the first gigayear. On the bottom panel, we compare the star formation histories of two simulations with $n_{\rm th} = 100\,\mathrm{cm}^{-3}$. The blue line corresponds to the CDM run, whereas the red line is the run with $\sigma_T/m_\chi = 1\,\mathrm{cm}^2\mathrm{g}^{-1}$. Star formation in the latter is strongly suppressed at $\sim 0.9\,\mathrm{Gyr}$ after a strong series of starbursts.

we expect that impulsive SNF causes a migration of older stars from the centre into the outskirts of the galaxy, an expectation that is generally confirmed by the results of all other simulations in which SNF is impulsive (see Figs. 8.9 and 8.10).

In Fig. 8.11, we attempt to explain the behaviour observed in those two simulations by looking at their star formation histories. On the top panel, we compare the star formation histories of all the simulations with very large SIDM momentum transfer cross section ($\sigma_T/m_\chi = 10\,\mathrm{cm^2 g^{-1}}$. Interestingly, we find that star formation is very bursty early on in the simulation with $n_\mathrm{th} = 10\,\mathrm{cm^{-3}}$. In particular, we can identify four strong star burst events within the first gigayear of the simulation. Overall, star formation in this run is even burstier than in the simulation with $n_\mathrm{th} = 100\,\mathrm{cm^{-3}}$. While we do not know the exact reason for this behaviour, we argue that it implies that baryons are very concentrated towards the centre of the galaxy in the beginning of the simulation. The early gravothermal collapse observed in this simulation, along with all the "odd" signatures outlined above, is thus likely the result of a complex interplay between baryonic physics and DM self-interactions.

The bottom panel of Fig. 8.11 compares the star formation history of the CDM simulation with $n_\mathrm{th} = 100\,\mathrm{cm^{-3}}$ (bursty star formation) to the SIDM simulation with $\sigma_T/m_\chi = 1\,\mathrm{cm^2 g^{-1}}$ and $n_\mathrm{th} = 100\,\mathrm{cm^{-3}}$. We see that initially stars in the two simulations form at a similar rate. After $\sim 700\,\mathrm{Myr}$ however, the star formation histories start to deviate significantly and at $t \sim 900\,\mathrm{Myr}$, there is a very large spike in the star formation rate of the SIDM simulation, followed by a sharp drop down to a steady, smooth, and rather low rate. In the CDM case, on the other hand, star formation continues in bursty cycles, with spikes that occur every $\sim 100\,\mathrm{Myr}$ on average. The impulsive SNF episodes following those spikes in star formation cause the migration of older stars into the outskirts of the galaxy. Since they are absent in the SIDM simulation, no strong stellar age gradient forms. Most likely, the large spike in SNF activity that followed the large peak in star formation in the SIDM simulation has driven most of the gas out of the centre of the galaxy. Subsequently, gas no longer accumulates in the centre of the galaxy and star formation proceeds more or less smoothly at random locations in the galaxy. Hence, the final stellar age gradients are significantly different from simulations with similarly strong SNF mainly due to the very strong initial star formation activity, a consequence of the stochastic implementation of the stellar evolution model.

Chapter 9

The impact of impulsive SNF on stellar orbits

This chapter is based on the following article:

Kinematic signatures of impulsive supernova feedback in dwarf galaxies

in preparation

Authors:
Jan D. Burger[1], Jesús Zavala[1], Laura Sales[2], Mark Vogelsberger[3], Federico Marinacci[4], and Paul Torrey[5]

[1]Centre for Astrophysics and Cosmology, Science Institute, University of Iceland, Dunhagi 5, 107 Reykjavik, Iceland

[2]University of California Riverside, 900 University Ave., Riverside CA 92521, USA

[3]Department of Physics, Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

[4]Department of Physics & Astronomy, University of Bologna, via Gobetti 93/2, I-40129 Bologna, Italy

[5]Department of Astronomy, University of Florida, 211 Bryant Space Sciences Center, Gainesville, FL 32611, USA

One of the most persevering small-scale challenges (Bullock and Boylan-Kolchin, 2017) to the collisionless cold dark matter (CDM) paradigm concerns the inner density profiles of DM haloes that host dwarf galaxies. Observational hints for constant density cores (Moore, 1994; Kuzio de Naray et al., 2008; Read et al., 2019; Walker and Peñarrubia, 2011; Oh et al., 2011, 2015) appear to be at odds with the ubiquitous cusps predicted by CDM N-body simulations (Navarro et al., 1996b; Wang et al., 2020). A potential way to flatten the central density profile of haloes is through strong and

impulsive fluctuations in the gravitational potential caused by supernova-driven episodes of gas removal (Navarro et al., 1996a; Pontzen and Governato, 2012; Di Cintio et al., 2014; Benítez-Llambay et al., 2019; Lazar et al., 2020; Dutton et al., 2020; Burger and Zavala, 2021). Although there is evidence of bursty star formation in dwarf galaxies at the high mass end (Kauffmann, 2014), it is still uncertain if the duration of bursts is sufficiently short in the intermediate and low mass end (Weisz et al., 2014; Emami et al., 2019) to explain the formation of large cores within the CDM paradigm. Here we show that the impulsive supernova cycles that follow episodes of bursty star formation leave distinct features in the distribution function of stars: groups of stars with similar ages and metallicities – which are associated with orbital families (Burger and Zavala, 2019; Price-Whelan et al., 2021) – develop shells in phase space that are characteristic of the initial stages of phase mixing (Binney and Tremaine, 2008; Dehnen and Read, 2011). If cores are formed through supernova feedback, such features should be present in star-forming dwarf galaxies with cored host haloes. On the other hand, the absence of such features would hint at an adiabatic core formation mechanism, giving strong support for deviations from CDM, for instance self-interacting DM (Spergel and Steinhardt, 2000; Vogelsberger et al., 2012; Rocha et al., 2013; Tulin and Yu, 2018).

9.1 Main

To reconcile the success of the cold dark matter (CDM) paradigm at predicting the properties of the large-scale structure of the Universe with observations on the scale of low-mass galaxies suggesting the presence of cored DM density profiles (Moore, 1994; Kuzio de Naray et al., 2008; Read et al., 2019; Walker and Peñarrubia, 2011; Oh et al., 2011, 2015), a physical mechanism is required to remove the central DM cusps predicted by CDM N-body simulations (Navarro et al., 1996b; Wang et al., 2020). Supernova feedback (SNF) is a cusp-core transformation mechanism whose effectiveness depends on the stellar-to-halo mass ratio and on how bursty and spatially concentrated star formation (SF) is within dwarf galaxies (Burger and Zavala, 2021). For SNF to be effective, supernovae must occur in quasi-periodic cycles in which the bulk of the energy injection happens on timescales that are shorter than the typical dynamical time in the galaxy, i.e. impulsively (Pontzen and Governato, 2012). In cosmological simulations, SNF is most efficient at forming cores on the scale of bright dwarfs (Di Cintio et al., 2014; Benítez-Llambay et al., 2019; Lazar et al., 2020). The relevant energy coupling occurs at the scales of the supernova remnant's interaction with the surrounding interstellar medium and such scales remain beyond the resolution of most cosmological galactic-scale simulations. Because of this, how impulsive the SNF mechanism is, and by extension how effective it is at forming cores, depends on a single effective parameter set at the resolved scales in the simulations, the so-called SF threshold $n_{\rm th}$, which gives the minimum density that gas needs to reach before it is eligible to form stars (Dutton et al., 2020). Adopting larger values of $n_{\rm th}$ results in burstier SF (and more impulsive SNF) until eventually a threshold for core formation is reached.

Another possible mechanism, in which cores are formed adiabatically, requires abandoning the assumption that DM is completely collisionless. In self-interacting

DM (SIDM), elastic scattering between DM particles redistributes energy outside-in leading to the formation of a ~ 1 kpc size DM core in dwarf-size haloes, provided the self-interaction cross section is $\sigma_T/m_\chi \sim 1\,\mathrm{cm^2 g^{-1}}$ on the scales of dwarf galaxies (Vogelsberger et al., 2012; Peter et al., 2013; Kaplinghat et al., 2016). Such value of the cross section also evades current constraints (Robertson et al., 2017, 2019; Read et al., 2018), while for cross sections smaller by about an order of magnitude, SIDM is indistinguishable from CDM (Zavala et al., 2013).

Below we present results of a suite of 16 high-resolution hydrodynamic simulations of an isolated dwarf galaxy, with structural properties similar to those of the Small Magellanic Cloud, embedded within a live DM halo. In each simulation, a different combination of $n_{\rm th}$ and σ_T/m_χ is adopted. Hereafter, we focus on comparing the results of two of these simulations with collisionless DM physics, namely the CDM runs with $n_{\rm th} = 0.1\,\mathrm{cm^{-3}}$ (representative of smooth SF and thus adiabatic SNF) and $n_{\rm th} = 100\,\mathrm{cm^{-3}}$ (representative of bursty SF and thus impulsive SNF). We discuss results of other runs whenever relevant.

Projections of the gas distribution and the emitted stellar light after 3 Gyr of simulation time are shown in **Figure** 9.1 for both benchmark simulations. We observe a few striking differences between the two runs, in particular towards the centre of the galaxies. In the simulation with bursty SF, the central gas density is lower than in the immediate surroundings due to a supernova-driven gas outflow extending out of the galactic disc, which can be clearly appreciated as a nearly spherical bubble in the edge-on projection. In contrast, in the simulation with smooth SF, the edge-on projection of the gas appears rather regular and the face-on projection has no distinct features. Interesting differences arise also in the stellar light projections shown in **Figure** 9.1. The face-on projections shows that the central surface brightness is larger in the run with a larger SF threshold (a central starburst is clearly apparent). The edge-on projection suggests also a difference in galaxy morphology with the galaxy with smooth SF being essentially bulge-less while the central starbursts have led to the development of a small bulge in the case with bursty SF. These morphological differences are related to the SF threshold: a large value of $n_{\rm th}$ generates a bursty SF history and induces a significant scatter between simulations due to the stochastic nature of the SF implementation, while a small value of $n_{\rm th}$ results in a smooth SF history with little variation across different simulations (see **Extended Data Figure** 9.1).

Figure 9.2 shows the final DM density and velocity dispersion profiles of our two benchmark runs plus a SIDM run with $n_{\rm th} = 0.1\,\mathrm{cm^{-3}}$ (smooth SF) and $\sigma_T/m_\chi = 1\,\mathrm{cm^2 g^{-1}}$. In the CDM run with smooth SF, the DM halo remains cuspy and in dynamical equilibrium. The inner density profile of the halo rises steeply towards the centre and the velocity dispersion profile is bell-shaped, with a maximum at a radius of ~ 2 kpc. The DM haloes simulated in the other two runs have constant density cores, which are remarkably similar, albeit the inner velocity dispersion profile of the SIDM halo (with smooth SF) is flat/isothermal, while that of the CDM halo (with bursty SF) is not fully isothermal, i.e., there is a residual non-zero slope in the central velocity dispersion profile. This demonstrates that while both impulsive SNF and SIDM can create similar kpc-size cores, their distinct nature ultimately affects the dynamics of the DM particles differently (Burger and Zavala, 2019). **Figure** 9.1 suggests that impulsive SNF also has a strong impact on the dynamical properties of gas and stars, creating

unique features that are potentially observable as we show below.

A striking dynamical signature of impulsive SNF is presented in **Figure** 9.3, which shows (for the benchmark CDM runs) the metallicity distribution and the mass-weighted distribution function of mono-age stars which are $0.8-0.9\,\mathrm{Gyr}$ old projected onto the $R-v_R$ plane at the end of the simulation. For the simulation with smooth SF, we observe a steady decrease of the average metallicity of stars with increasing cylindrical radius, a natural consequence of the centrally concentrated star-forming gas. Statistically, more stars form in environments with higher gas densities, i.e., towards the centre of galaxies. Thus, the subsequent SNF cycles cause a metal enrichment of the ISM that is larger in the central regions. Therefore, stars of subsequent generations (like the ones shown in **Figure** 9.3 that form after $\sim 2.2\,\mathrm{Gyr}$ of simulation time) acquire a negative metallicity gradient. Moreover, the radial velocities of stars are rather small in magnitude, $v_R \sim 25\,\mathrm{km\,s^{-1}}$ at most. A different picture emerges in the centre of the galaxy for the bursty SF case. Instead of a monotonic stellar metallicity gradient, a pattern of several shells in $R-v_R$ space emerges in the metallicity distribution – and the mass-weighted distribution function – of stars with similar ages. The shells are comprised of star particles with high metallicities, some of which move at radial velocities of more than $50\,\mathrm{km\,s^{-1}}$. These high metallicity shells intersect phase space regions inhabited by more metal-poor star particles whose radial speeds are smaller on average. Such features are transient for a given group of stars, but occur at various times in the evolution, and are not unique to the CDM run with $n_{\mathrm{th}} = 100\,\mathrm{cm^{-3}}$; we find them also in other simulations (both in CDM and SIDM) as long as $n_{\mathrm{th}} \geq 10\,\mathrm{cm^{-3}}$ (see e.g. **Extended Data Figure** 9.2). The exception are simulations in which no bursty SF occurs (see **Extended Data Figure** 9.1).

To quantify the difference between the final stellar distributions in the bursty SF case and in the smooth SF case, we estimate the likelihood of randomly finding, in the smooth SF case, an overdensity similar to that associated with the clear shell in the bursty SF simulation. We re-sample the smooth distribution (left column of **Figure** 9.3) 10^7 times, each time calculating the overdensity within the "shell region" of the bursty SF simulation and the largest overdensity within a pre-defined signal region, corresponding to the area in which we found shells in other simulations with bursty SF. From the obtained overdensity distributions we estimate that the overdensity associated with the shell shown in Figure 9.3 has a local (global) significance of more than $5\,\sigma$ ($3.3\,\sigma$) compared to the smooth SF case (see **Extended Data Figure** 9.5). These are conservative estimates for the significance of the shell-like feature since they are based on the stellar density distribution only and do not take into account information on the metallicity of stars. Finding an overdensity of such amplitude, combined with the observation that the overdensity consists mainly of high-metallicity stars, would be a smoking gun signature of an impulsive SNF cycle following an episode of bursty SF.

Figure 9.4 shows the metallicity distributions and the mass-weighted distribution functions of the same stars as in **Figure** 9.3, but projected into $|x|-v_y$ space to identify how the shell-like features would appear in the line-of-sight phase space of galaxies which are observed edge-on. In the featureless smooth SF case, the emerging distribution of stars is U-shaped, with larger line-of-sight velocities at larger transverse distances (i.e. tracking the rotation curve of the galaxy) and a monotonic decrease of metallicity with distance. In the bursty SF case, we observe two isolated overdense clusters consisting

mainly of high metallicity stars. The distance of those clusters to the centre of the galaxy strongly suggests that they consist of the same stars as the shell in **Figure** 9.3 (see also **Extended Data Figure** 9.4). Those overdense high-metallicity regions are thus a distinct (and potentially observable) feature of simulations with bursty star formation (and impulsive SNF), while smooth star formation leads to smooth rotation curves.

The shell-like features shown in **Figure** 9.3 (and **Extended Data Figure** 9.2) arise in the immediate aftermath of starburst events (followed by impulsive episodes of SNF) and are signatures of the first stages of phase mixing (Binney and Tremaine, 2008; Dehnen and Read, 2011). Phase mixing is usually described as a process that gradually smooths out the coarse-grained distribution function of particles which initially occupy a narrow volume in phase space but have different integrals of motion (Dehnen and Read, 2011). However, phase mixing can also occur in *orbital families*, sets of orbits defined by similar integrals of motion, if at least one of those integrals of motion changes discontinuously due to an impulsive change in the gravitational potential (Burger and Zavala, 2019). Notably, an impulsive change in the potential changes the energy of a particle by an amount that depends on its orbital phase (Pontzen and Governato, 2012; Burger and Zavala, 2019) and can thus split an initially phase mixed *orbital family*. In dynamical systems in which orbits are regular and stars act as dynamical tracers of the gravitational potential, the average metallicity of stars can only depend on their actions (Price-Whelan et al., 2021). Thus, we can approximate *orbital families* by groups of stars with similar ages and metallicities. We find that in simulations with impulsive SNF (following bursty SF), the energy distribution of such *orbital families* is wider than in simulations with smooth SF. This widening is a result of episodic impulsive changes in the potential, as has been shown using a toy model of an ideal *orbital family* of kinematic tracers (Burger and Zavala, 2019) (see also **Extended Data Figure** 9.3). The early-stage phase mixing observed here is thus a direct consequence of the SNF induced widening of the energy distribution of *orbital families* that occurs in simulations with bursty SF.

If shell-like features similar to the ones presented here are found in cored dwarf galaxies, this implies that those haloes have experienced episodes of impulsive SNF, establishing SNF as a feasible core formation mechanism. However, the absence of such features in star-forming galaxies with cored host haloes would favour adiabatic mechanisms of cusp-core transformation such as SIDM. In cosmological halos, we expect additional diffusion caused by environmental effects to erase the shells within ~ 1 dynamical time (Burger and Zavala, 2019) and thus the above conclusion cannot be applied directly to quenched galaxies. Nevertheless, it is instructive to evaluate the potential of detecting such signatures of bursty SF in the Milky Way satellites. Among those, the best target for our analysis is Fornax (Walker et al., 2009; de Boer et al., 2012; Rusakov et al., 2021), since it has been claimed to have a core (Walker and Peñarrubia, 2011) (albeit this remains controversial (Genina et al., 2018)) and information on line-of-sight kinematics, metallicity (Walker et al., 2009), and age (de Boer et al., 2012) is available for a sub-sample of its member stars. However, the ages of individual stars carry uncertainties of ~ 1 Gyr – too large to to conclusively define *orbital families*. Moreover, the number of stars with available metallicity information is currently too low in Fornax to find signatures of impulsive SNF that are sufficiently significant. Sampling

the final stellar distribution of our CDM simulation with smooth SF (right column of **Figure** 9.3) with a decreasing number of stars, we estimate that we require precise age (with a resolution of around a dynamical timescale) and metallicity information of $\mathcal{O}(10^4)$ stars to determine whether signatures of impulsive SNF are present at a $2\,\sigma$ level. Ideal future targets to look for impulsive SNF signatures are star-forming field dwarfs within the Local Group. Within the next decade, the Roman Space Telescope will provide precise photometric data of individual stars in dwarf galaxies within the Local Volume (Khan et al., 2018). Combined with spectroscopic data from the ground, this will enable the precision needed to determine the ages, metallicities, and kinematics of a sufficient number of stars to conclusively establish whether the characteristic shell-like signatures of impulsive SNF presented here are ubiquitously present or absent. The systematic absence of such signatures in isolated, star-forming field dwarfs would deliver a blow to the CDM model, giving strong support to alternative DM models.

Acknowledgements

JB and JZ acknowledge support by a Grant of Excellence from the Icelandic Center for Research (Rannís; grant number 173929). The simulations presented here were carried out on the Garpur supercomputer, a joint project between the University of Iceland and University of Reykjavík with funding from Rannís. PT acknowledges support from NSF grant AST-1909933, AST-200849, and NASA ATP Grant 80NSSC20K0502. FM acknowledges support through the program "Rita Levi Montalcini" of the Italian MUR.

Author Information

Author contributions

JB performed the simulations, created the figures (except figure 9.1), and wrote an initial version of the manuscript. JZ assisted with writing the initial manuscript. PT created the stellar light density projections in figure 9.1. FM created the gas density projections in figure 9.1. LS, MV, and JZ contributed to the preparation and execution of the simulations. All authors reviewed the manuscript.

Materials and correspondence

Correspondence and material requests should be addressed to Jan Burger (jdb5@hi.is).

Competing interests

The authors declare no competing interests.

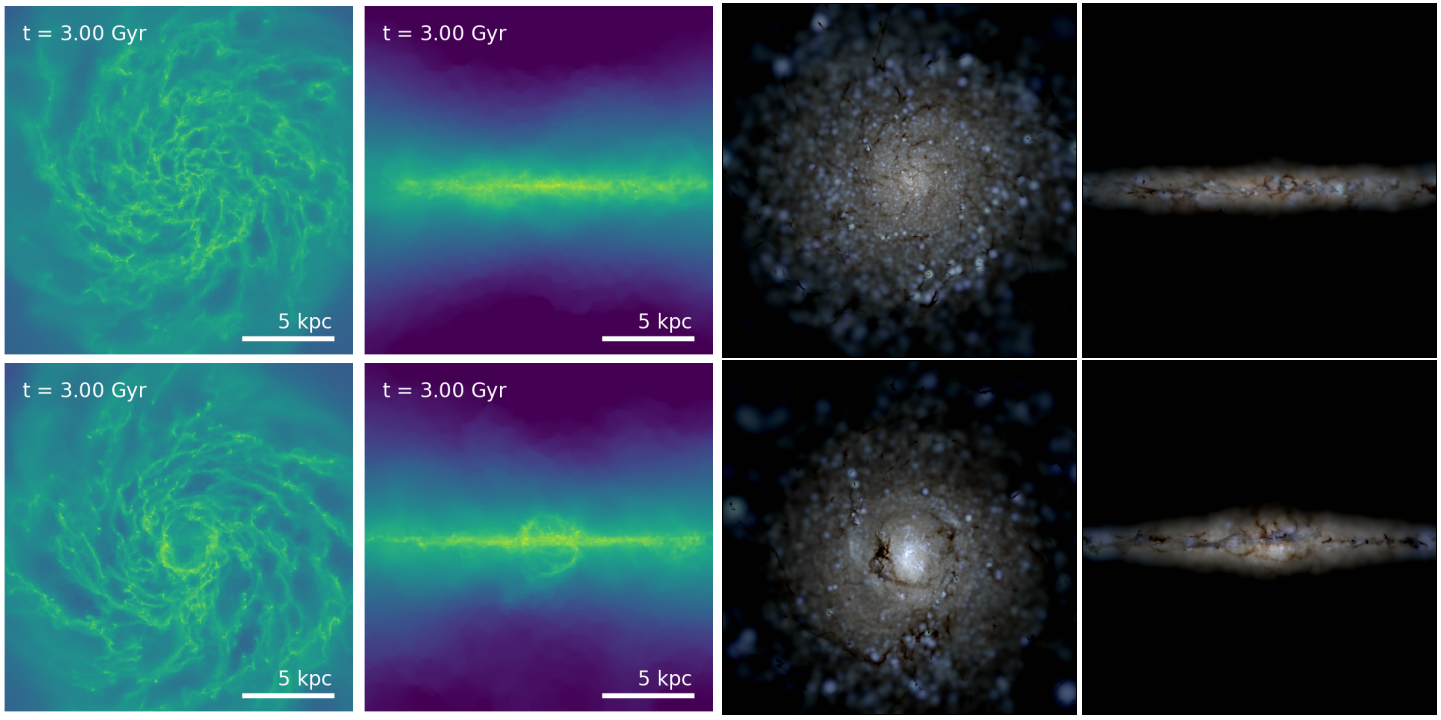

Figure 9.1. **Resulting galaxies after 3 Gyr of simulation time for two simulations with different star formation thresholds.** *In the top panels, we show face-on and edge-on projections of the gas density (first and second column, respectively) and stellar light (third and fourth column, respectively) for the CDM simulation with smooth star formation ($n_{\rm th} = 0.1\,{\rm cm}^{-3}$). The side-length of the field of view is 20 kpc in each panel. In the bottom panels we show the CDM simulation case with bursty star formation ($n_{\rm th} = 100\,{\rm cm}^{-3}$). Notice how central gas is vertically expelled out of the disc plane in the simulation with a bursty star formation. Such galactic outflows are characteristic of violent and impulsive events of supernova-driven energy release.*

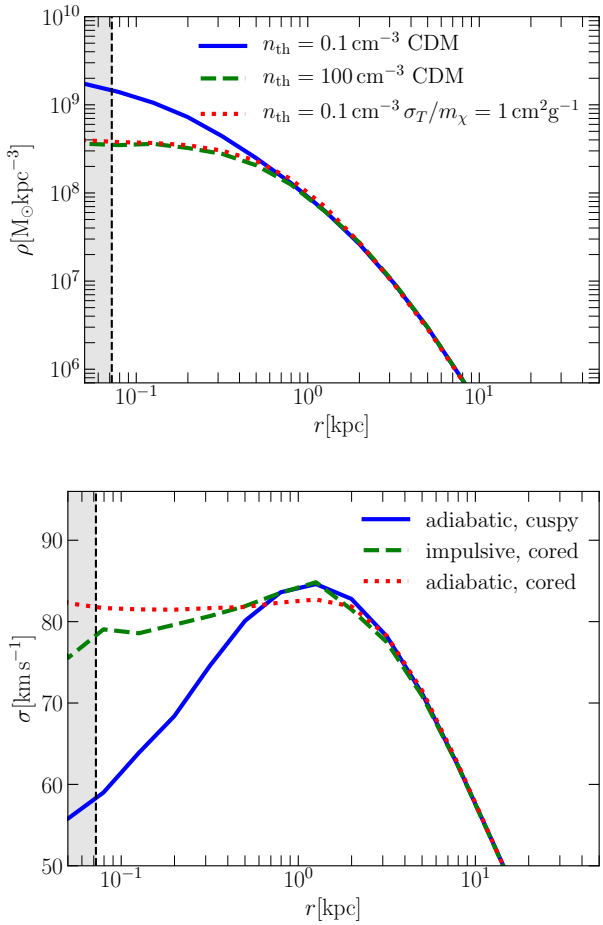

Figure 9.2. **Adiabatic and impulsive cusp-core transformation: Spherically averaged radial density (top panel) and velocity dispersion (bottom panel) profiles of three different simulated dwarf-size haloes after 3 Gyr of simulation time**. *The solid blue (green dashed) lines correspond to the CDM simulation with smooth (bursty) star formation The red dotted lines show the SIDM simulation with smooth star formation and a self-interaction cross section* $\sigma_T/m_\chi = 1\,\mathrm{cm}^2\mathrm{g}^{-1}$ *as red dotted lines. In CDM the cycles of energy injection by supernovae into the DM halo need to happen impulsively (i.e., caused by bursty star formation) in order to generate a core, while for relevant SIDM models the development of an adiabatic core is a natural stage in DM haloes regardless of the influence of supernovae. Notice that the inner DM cusp of a dwarf-size halo can be transformed into a nearly identical* $\sim \mathcal{O}(1)$ *kpc size core through both, an impulsive supernova feedback scenario, or an adiabatic SIDM scenario.*

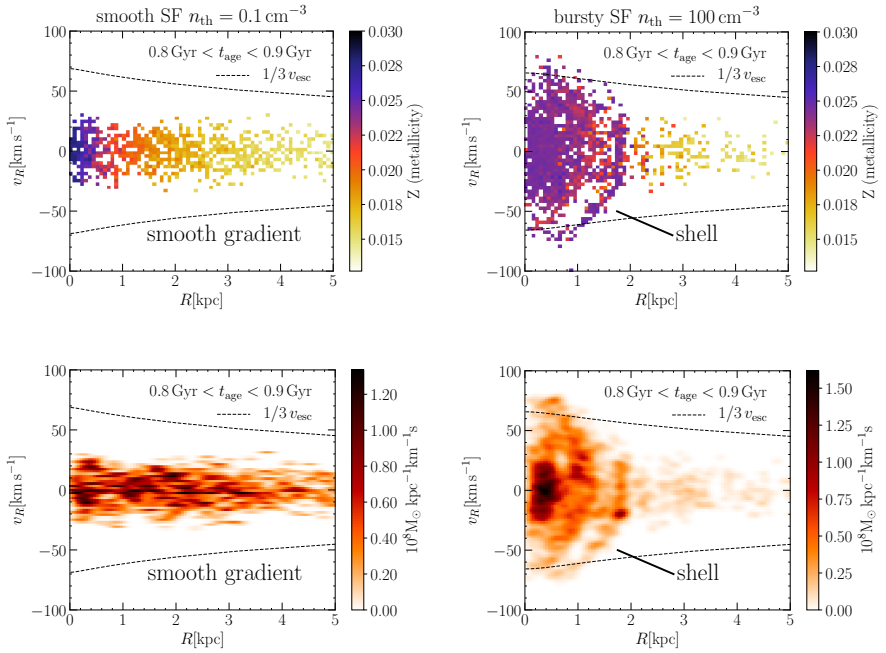

Figure 9.3. **Average metallicity distribution (top panels) and mass-weighted distribution function (bottom panels) of star particles calculated after 3 Gyr of simulation time for two different simulations, projected into the radial phase space ($R - v_R$).** *Star particles are subject to a cut in stellar age, only stars which are* $0.8 - 0.9\,Gyr$ *old are shown in these plots. The metallicities of the star particles are averaged in* 70×70 *bins in* $R - v_R$ *space. The averaged stellar metallicity (mass-weighted distribution function) is colour-coded according to the coulour scale to the right of each panel. The mass-weighted distribution function is obtained using a Gaussian kernel density estimate. For scale, we show a third of the escape velocity as a function of radius as black dashed lines. The left (right) column corresponds to the CDM simulation with smooth (bursty) star formation. Smooth star formation results in a smooth metallicity gradient with the enriched/younger starts in the centre. Bursty star formation results in an overall weaker gradient, along with the presence of a shell of stars with high metallicity that intersects a low metallicity population at* $R \sim 2kpc$*. This shell appears as a distinct overdense region in the mass-weighted distribution function.*

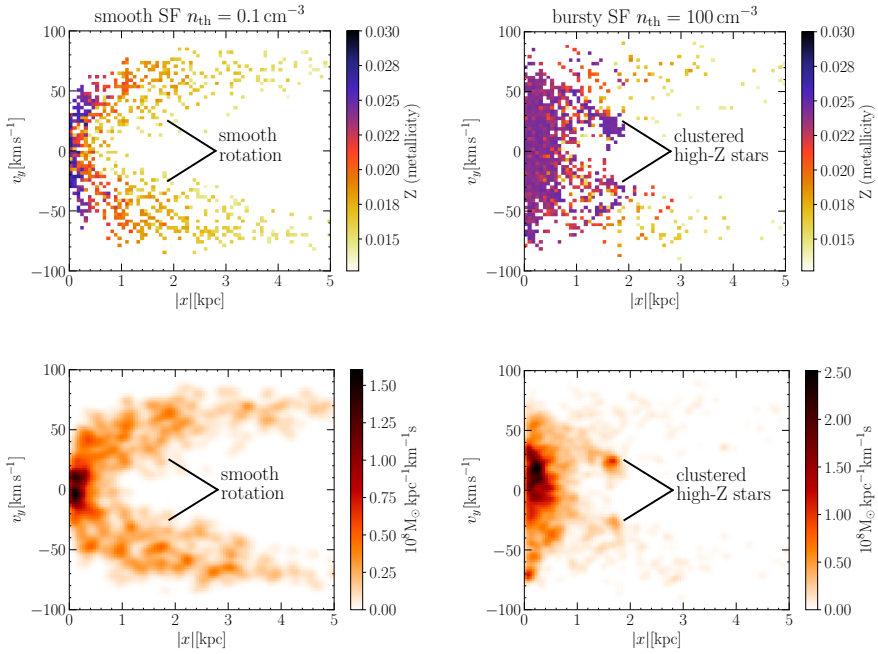

Figure 9.4. **Average metallicity distribution (top panels) and mass-weighted distribution function (bottom panels) of star particles calculated after 3 Gyr of simulation time for two different simulations, projected into $|x| - v_y$ space.** *The same as figure 9.3 but along the disk plane: in the x direction spatially, and in the y direction in velocity space. This gives an idea of the expected stellar kinematics projected into the line-of-sight phase space. Notice how in the left panels (smooth star formation), a smooth, but spread/scattered, rotation curve can be observed, while in the right panels, features associated to the impulsive case (bursty star formation), are apparent, albeit their structure is harder to discern (relative to the radial phase space shown in figure 9.3, where these features are clear). In particular, two distinct overdense clusters of high metallicity stars appear at* $|x| \sim 2\,{\rm kpc}$ *in the impulsive (bursty star formation) case.*

9.2 Methods

This Article is based on a suite of 16 hydrodynamic simulations of an isolated dwarf galaxy in a live DM halo. Here we briefly outline how we set up the initial conditions, run the simulations, and analyze the simulation results.

9.2.1 Initial conditions

We generate initial conditions of a self-gravitating system of a similar scale and global properties as the Small Magellanic Cloud (Hopkins et al., 2012) at redshift zero with the important distinction that our simulated system is isolated. The initial conditions consist of four separate mass components: a spherically-symmetric DM halo with a Hernquist (Hernquist, 1990) density profile, an exponential gas disc, an exponential "stellar" disc consisting of collisionless particles, and a "stellar" bulge following a Hernquist profile, also consisting of collisionless particles. We fix the structural parameters of the halo through its circular velocity at the virial radius r_{200} and its concentration parameter c_{200}[1]. In all our simulations, $v_{200} = 36.31\,\mathrm{km\,s^{-1}}$ and $c_{200} = 18$. The corresponding virial mass of the system is $M_{200} \sim 1.6 \times 10^{10}\mathrm{M}_{\odot}$ and we note that total halo mass, virial radius, and scale radius are fully determined by specifying v_{200} and c_{200}. The gas disc is defined by its surface density profile,

$$\Sigma(R) = \frac{M_{\infty}}{2\pi H_{\mathrm{gas}}^2} \exp\left(-\frac{R}{H_{\mathrm{gas}}}\right), \tag{9.1}$$

where M_{∞} is the total mass of gas in the system, H_{gas} is a scale length, and R is the polar radius in a coordinate system whose origin is at the centre of the galaxy and whose vertical axis is normal to the galactic plane. The scale length of the gas disc is $H_{\mathrm{gas}} = 2.1\,\mathrm{kpc}$. M_{∞} is fixed as follows: we assume that the gas and stellar discs combined contain 4.45% of the total mass of the dynamical system and that gas makes up for 84% of the disc mass. In turn, this also fixes the mass of the stellar disc. The vertical structure of the gas disc is fixed by demanding that the gas is initially in hydrostatic equilibrium. We assume that the pressure of the gas is given by the equation of state of an isothermal ideal gas with a temperature of $T = 10^4$ K and that initially, the gas mixture has solar metallicity, $Z_{\odot} = 0.0127$ (Asplund et al., 2009), and is fully ionized. Under these conditions, and at a fixed polar radius R, the vertical structure of a gas disc in hydrostatic equilibrium is determined by the following set of equations:

$$\frac{d\rho}{dz} = -\frac{d\Phi}{dz}\frac{\rho^2}{P\Gamma} \tag{9.2}$$

$$\frac{d\Sigma}{dz} = \rho. \tag{9.3}$$

where ρ is the gas density, Φ is the total gravitational potential, P is the gas pressure, Γ is the ratio of specific heats, and $\Sigma = \Sigma(z)$ denotes the cumulative gas surface density.

[1] r_{200} is defined through the virial mass, $M_{200} = 200 \times 4\pi/3\rho_{\mathrm{crit}} r_{200}^3$ with ρ_{crit} the critical density of the Universe today, and $c_{200} = r_{200}/r_{-2}$, with r_{-2} the radius at which the logarithmic slope of the halo's density profile equals -2.

Using an initial guess for $\rho(R,0)$, we integrate equations 9.2 and 9.3 upwards from the midplane and calculate $\rho(R,z)$ and $\Sigma(R,z)$ on a grid of (R,z) values. The initially guessed value for $\rho(R,0)$ is multiplied by a correction factor and the vertical integration is repeated. The applied correction factor is the ratio between the known surface density (equation 9.1) and the cumulative surface density at the largest z value on the grid, $\Sigma(R,z_{\rm max})$, as calculated from equation 9.3. This vertical integration procedure is iterated five times. The surface density profile of the stellar disc has the same functional form as that of the gas disc, i.e., equation 9.1, but with a scale length of $H_\star = 0.7$ kpc. The star particles in the disc follow a sech^2-distribution in the vertical direction,

$$\rho(R,z) \propto \Sigma(R) \left[\cosh\left(\frac{z}{z_0}\right)\right]^{-2}, \tag{9.4}$$

with a vertical scale height $z_0 = 0.14\,$kpc. The star particles that make up the bulge are distributed following a Hernquist profile with scale length $A = 0.233\,$kpc.

Since both the DM halo and the bulge are spherically symmetric, we can use Eddington sampling to calculate their full distribution functions (Eddington, 1916). Note, however, that due to the presence of the baryonic discs, the total gravitational potential of the system, which appears in Eddington's integral, is not fully spherically symmetric. Therefore, we have to chose a preferred axis along which we evaluate the integral. Here, we chose to evaluate Eddington's integral along the vertical direction and neglect the inaccuracy that is introduced by assuming isotropic distribution functions[2]. Having calculated the distribution functions, we assign velocities to the halo and bulge particles using a rejection sampling scheme.

The velocities of the disc particles are calculated in a different way. First we calculate the streaming velocity and the velocity dispersion tensor on a logarithmic grid of radial and vertical coordinates, using the Jeans equation in cylindrical coordinates and the epicyclic approximation (Springel et al., 2005a; Hernquist, 1993). Individual disc particles are then assigned velocities which are a sum of the streaming velocity at the position of the particle and a random component sampled from a local Maxwellian velocity distribution. The motion of gas cells is given by the local streaming velocity of the gas. The streaming velocity field of the gas is calculated taking into account both gravity and pressure gradient (Springel et al., 2005a).

Initially, our system consists of 1.2×10^7 DM particles, 4×10^5 gas cells, 8×10^4 disc particles, and 8×10^3 bulge particles. Each particle has an approximate mass of $1.3 \times 10^3 {\rm M}_\odot$. In a subsequent step, the system is enclosed within a cubic volume with a side length of 100 kpc, into which a grid of background gas cells with low density and high temperature is introduced (Springel, 2010). In the same step, previously constructed gas cells can be refined or de-refined if they contain a mass which is either larger than twice the average mass of a gas cell or smaller than half the average mass of a gas cell before introducing the background grid. Finally, since the resulting dynamical equilibrium is incomplete, we evolve the system for 1 Gyr in a simulation in which gas cooling and star formation are turned off. We adopt the final output of this preparation run as the relaxed initial conditions for our simulation suite.

[2] We later let the system relax into a steady state. Thus, it is not crucial that the dynamical equilibrium is perfect at this stage.

9.2.2 Simulations

We run all our 16 simulations using the code Arepo (Springel, 2010) which uses a TreePM algorithm to calculate the gravitational forces, while the hydrodynamics is solved on a moving mesh, using a quasi-Lagrangian scheme based on a Voronoi tessellation. We use the SMUGGLE (Marinacci et al., 2019) stellar feedback model to model gas cooling, star formation, stellar evolution, supernova feedback, and radiative feedback. Self-interactions between DM particles are implemented using the Monte Carlo approach described in Vogelsberger et.at. (2012) (Vogelsberger et al., 2012). Both SMUGGLE and the SIDM modules incorporated in Arepo can be calibrated through a number of parameters. Apart from the star formation threshold, all SMUGGLE model parameters are fixed to the values reported in table 3 of the original code paper (Marinacci et al., 2019). For the DM self-interactions, we adopt a fixed velocity-independent transfer cross section that varies between simulations. The amount of particles included in the nearest neighbour search for a scattering partner is set to $N_{\rm ngb} = 36 \pm 5$. The model parameters that determine whether our modeled DM halo forms a constant density core or not are the SF threshold $n_{\rm th}$ and the self-interaction cross section σ_T/m_χ.

In SMUGGLE, gas cells are converted into star particles following a standard probabilistic approach (Springel and Hernquist, 2003). The local star formation rate in a given gas cell is

$$\dot{M}_\star = \begin{cases} 0 & \rho < \rho_{\rm th} \\ \varepsilon \frac{M_{\rm gas}}{t_{\rm dyn}} & \rho \geq \rho_{\rm th} \end{cases} . \tag{9.5}$$

where ε is an efficiency parameter which is set to 0.01 in line with observations of slow SF in dense gas (Krumholz and Tan, 2007), $M_{\rm gas}$ is the gas mass in the cell and

$$t_{\rm dyn} = \sqrt{\frac{3\pi}{32 G \rho_{\rm gas}}} \tag{9.6}$$

is the free-fall time of the gas. Stars can only form in gas cells in which the gas density is larger than $\rho_{\rm th}$, which is calculated from the numerical number density threshold $n_{\rm th}$ and the mean molecular weight of the gas. Moreover, large densities are reached faster in regions in which the potential well is deeper, i.e., towards the centre of galaxies. For larger SF thresholds, the gas must accumulate until it reaches the density threshold to become eligible for star formation, leading to more concentrated and more bursty SF, both of which increases the efficiency of SNF at transforming DM cusps into cores (Burger and Zavala, 2021).

In the stochastic SIDM model adopted here, the probability that a simulation particle i scatters with one of its nearest neighbours j in a given timestep is proportional to σ_T/m_χ. Thus, the scattering rate increases if larger transfer cross sections are adopted. An increased scattering rate increases the rate at which the inner DM halo thermalizes, and thus reduces the time it takes to form the DM core. There is a relatively narrow window for the value of σ_T/m_χ to lead to fully isothermal cores in haloes today. If the cross section is too small, $\sigma_T/m_\chi < 0.1 {\rm cm}^2 {\rm g}^{-1}$, only the innermost regions of the halo thermalize leading to a halo structure that is nearly indistinguishable from CDM

(Zavala et al., 2013; Rocha et al., 2013). On the other hand, adopting very large cross sections, $\sigma_T/m_\chi \gtrsim 10\mathrm{cm}^2\mathrm{g}^{-1}$, will trigger the gravothermal collapse phase of SIDM (Balberg et al., 2002; Colín et al., 2002; Koda and Shapiro, 2011; Pollack et al., 2015; Nishikawa et al., 2020), which eventually results in SIDM haloes that are even cuspier than CDM haloes.

In our simulation suite, we adopt 4×4 combinations of n_{th} and σ_T/m_χ. The SF threshold takes values of $n_{\mathrm{th}} = (0.1, 1, 10, 100)\,\mathrm{cm}^{-3}$, whereas for the SIDM transfer cross sectionm we adopt values of $\sigma_T/m_\chi = (0, 0.1, 1, 10)\,\mathrm{cm}^2\mathrm{g}^{-1}$.

In all our production runs we adopt a softening length $\varepsilon = 24\,\mathrm{pc}$ for all particle species. In order to properly resolve the effects of supernova feedback in runs with large SF thresholds (Dutton et al., 2020), this softening length is slightly smaller than the optimal softening length for a system of this size (Power et al., 2003).

9.2.3 Post processing

The gas density projections and the stellar light projections shown in Figure 1 are created as in Marinacci et.al. (2019) (Marinacci et al., 2019). The density projection panels are obtained by integrating for each pixel of the image the total gas density along the line of sight for a depth equal to the image side-length. The stellar light images are generated using stellar population synthesis models coupled to a line-of-sight dust extinction calculation assuming a constant dust-to-metals ratio.

Here we briefly discuss how we process simulation snapshots to create the remaining figures. We start by calculating the centre of potential $\mathbf{R}$ and velocity $\mathbf{V}$ of the DM system, using a shrinking spheres method. A first guess is obtained by summing over all N DM particles,

$$\mathbf{R}^0 = \sum_{i=1}^{N_{\mathrm{DM}}} \frac{\Phi(\mathbf{r}_i)\,\mathbf{r}_i}{\Phi(\mathbf{r}_i)} \tag{9.7}$$

$$\mathbf{V}^0 = \sum_{i=1}^{N_{\mathrm{DM}}} \frac{\Phi(\mathbf{r}_i)\,\mathbf{v}_i}{\Phi(\mathbf{r}_i)}, \tag{9.8}$$

where $\mathbf{r}_i$ and $\mathbf{v}_i$ are the position and the velocity vector of particle i and Φ is the gravitational potential. In subsequent steps, the sums are restricted to DM particles which are closer to the currently estimated centre of potential than a given threshold radius r_{th}. For instance, in step $n+1$, particles need to satisfy $|\mathbf{r} - \mathbf{R}^n| < r_{\mathrm{th},n+1}$ to be included in the sum. Here we use three iterations, with threshold radii of 50 kpc, 5 kpc, and 0.5 kpc. Once the final values of $\mathbf{R}$ and $\mathbf{V}$ have been calculated, we shift the phase space coordinates of all particles and gas cells, defining a new coordinate system in which the centre of potential of the system is at $(0,0,0)$ and the system has no bulk motion.

Subsequently, since we simulate the evolution of a rotationally supported galaxy, it is convenient to perform a coordinate transformation such that the net rotation occurs in the $x-y$ plane of the new coordinate system. To that end, we calculate the total angular momentum from all particles which are part of the rotating disc, i.e., the collisionless disc particles and the formed star particles, as well as the gas cells. Since individual simulation particles represent stellar populations and gas cells represent a

extended volume of moving gas, we have to give a different weight to the angular momenta of particles/gas-cells with different masses. Hence, we calculate the total angular momentum as a mass-weighted sum of the angular momenta of all rotating particles and gas cells. We then rotate the coordinates and velocities of all particles into a coordinate system in which the vertical axis is aligned with the direction of the total angular momentum vector. All figures presented here are constructed from dynamical quantities measured in this coordinate system.

The density and velocity dispersion profiles shown in **Figure** 8.2 are calculated in 20 bins which are equally spaced in in logarithmic radius. **Figures** 9.3 and 9.4 and **Extended Data Figure** 9.2 are subject to the age cuts stated in the respective captions. The phase space binning adopted in **Figure** 9.3 and **Extended Data Figure** 9.2 is as follows. We focus on the radial range $(0\,\mathrm{kpc} < R < 5\,\mathrm{kpc})$ and the radial velocity range $(-100\,\mathrm{km\,s^{-1}} < v_R < 100\,\mathrm{km\,s^{-1}})$. To produce the two-dimensional colour plots shown in the top row of **Figure** 9.3 and in **Extended Data Figure** 9.2, we divide the $R - v_R$ plane into 70×70 bins which equidistantly cover the adopted ranges of radii and radial velocities. The metallicity values shown in **Figure** 9.3 and **Extended Data Figure** 9.2 are mass-weighted averages over all stellar particles within a given bin. The two-dimensional colour plots in the top row of **Figure** 9.4 are constructed in the same way, using the phase space coordinates $|x|$ and v_y instead of R and v_R. To construct the two-dimensional colour plots shown in the bottom rows of **Figures** 9.3 and 9.4 we have used the scipy routine *stats.gaussian_kde* with a fixed scalar bandwidth of 0.07. The escape velocities shown in **Figure** 9.3 and **Extended Data Figure** 9.2 are calculated from the gravitational potential as

$$v_{\mathrm{esc}} = \sqrt{-2\Phi(r)}, \tag{9.9}$$

where the gravitational potential is calculated in the same 20 bins as the density profile and the velocity dispersion profile. The energy distributions displayed in the middle and right panel of **Extended Data Figure** 9.3 are obtained from the 150 most metal-rich stellar particles of a given age in three different simulations as stated in the figure caption. The energy of each stellar particle is comprised of its kinetic energy, which we calculate using the shifted velocities, and its potential energy, which is calculated by Arepo on the fly. In the left panel of **Extended Data Figure** 9.3 we show the energy distribution of *orbital families* of kinematic tracers in a DM halo that has formed a core either adiabatically or impulsively simulated using a spherically symmetric toy model (Burger and Zavala, 2019). Similar to the energies of the stellar particles formed in the hydrodynamic simulations presented here, the energies of the tracers are calculated in a shifted coordinate system in which the centre of potential is at $(0,0,0)$ and the system has no bulk motion. However, since the simulated toy model systems are spherically symmetric and isotropic, there is no net angular momentum and hence we do not need to rotate into a different coordinate frame. Finally, we note that the SF histories shown in **Extended Data Figure** 9.1 are calculated by SMUGGLE on the fly and that **Extended Data Figure** 9.4 is created by simply plotting the phase space coordinates of all stars of the indicated age and that all stars that fulfill the cuts stated in the left panel are coloured in red.

9.2.4 Determining the significance of the shell overdensity

We aim to obtain a conservative estimate for the significance of the overdense shell in the CDM run with bursty SF (see right column of **Figure** 9.3). To that end, we define the shell overdensity as the ratio between the mass contained within the rectangle in phase space defined by $[1.7 < R/\mathrm{kpc} < 1.9125, -30 < v_R/(\mathrm{km\,s^{-1}}) < 30]$ and the mass contained within the rectangle $[1.4875 < R/\mathrm{kpc} < 1.7, -30 < v_R/(\mathrm{km\,s^{-1}}) < 30]$ and find that the overdensity $\Delta_{\mathrm{shell}} = 2.89$ (see **Extended Data Figure** 9.5). We now look to determine the likelihood for such an overdensity to arise as a random fluctuation of the smooth SF distribution. To that end, we take the normalized cumulative radial mass distribution, $\bar{M}(R)$, of the CDM run with smooth SF (left column of **Figure** 9.3) as a target distribution for random sampling. We then re-sample this distribution a total of 10^7 times, using the same total number of stars as in the original (simulated) distribution. For each sampled distribution, we calculate the density ratio between the same two rectangles as in the bursty SF case, defining the "local" overdensity. We also look for "global" overdensities, i.e. the largest overdensities within a signal region, A_{signal}, which is defined by $[0.5 < R/\mathrm{kpc} < 3, -30 < v_R/(\mathrm{km\,s^{-1}}) < 30]$ (see middle panel of **Extended Data Figure** 9.5) and corresponds to the phase space region in which we found shell-like overdensities across all our runs with bursty SF. We look for overdensities in this region by calculating the ratios between adjacent phase space rectangles of the same size as above, and covering the same range of radial velocities. We scan the radial range of A_{signal} in steps which are equal to the softening length adopted in the simulations, i.e. the resolution limit. In the right panel of **Extended Data Figure** 9.5, we present the resulting distribution of local (global) overdensities as a blue (green) line. We also show the shell overdensity in red. We calculate the local (global) significance from the number of sampled distributions in which the local (global) overdensity $\Delta < \Delta_{\mathrm{shell}}$ by assuming that the local (global) overdensities follow a Gaussian distribution. To determine the minimum number that is required for the shell to have a global significance which is $> 2\sigma$ we repeat the re-sampling procedure several times, each time with a smaller number of (sampled) stars.

Data availability

Raw data were generated at the Garpur large-scale facility. Derived data supporting the findings of this study are available from the corresponding author upon request.

Code availability

The script used to analyze the simulation snapshots is available from the corresponding author upon reasonable request.

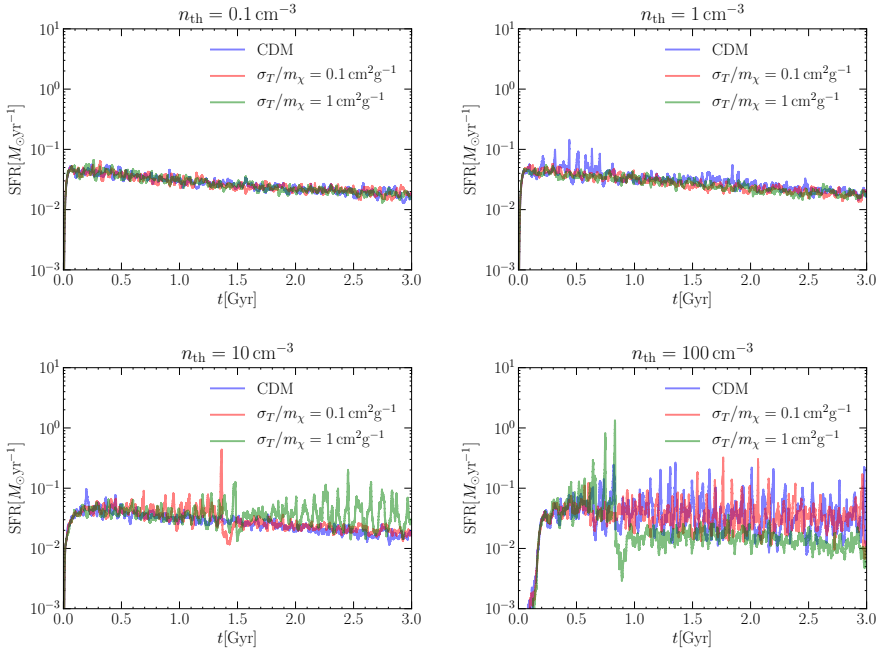

Extended Data Figure 9.1. **SF histories for 12 different simulations within CDM and SIDM.** *Different panels correspond to simulations with different SF thresholds, $n_{\rm th}$, according to the legend on the top of each panel. Within each panel, differently coloured lines correspond to simulations with different DM self-interaction cross sections as indicated in the figure legends. As the SF threshold is increased, SF bursts become more prominent and episodic. After particularly intense bursts, the SF history can diverge across simulations due to the stochastic nature of the SF implementation. The benchmark runs of this Article are the CDM run with $n_{\rm th} = 0.1\,{\rm cm}^{-3}$ (blue line upper left panel) and the CDM run with $n_{\rm th} = 100\,{\rm cm}^{-3}$ (blue line lower right panel). The SIDM run used in* **Figure** *8.2 has $n_{\rm th} = 0.1\,{\rm cm}^{-1}$ and $\sigma_T/m_\chi = 1\,{\rm cm}^2{\rm g}^{-1}$ (green line upper left panel).*

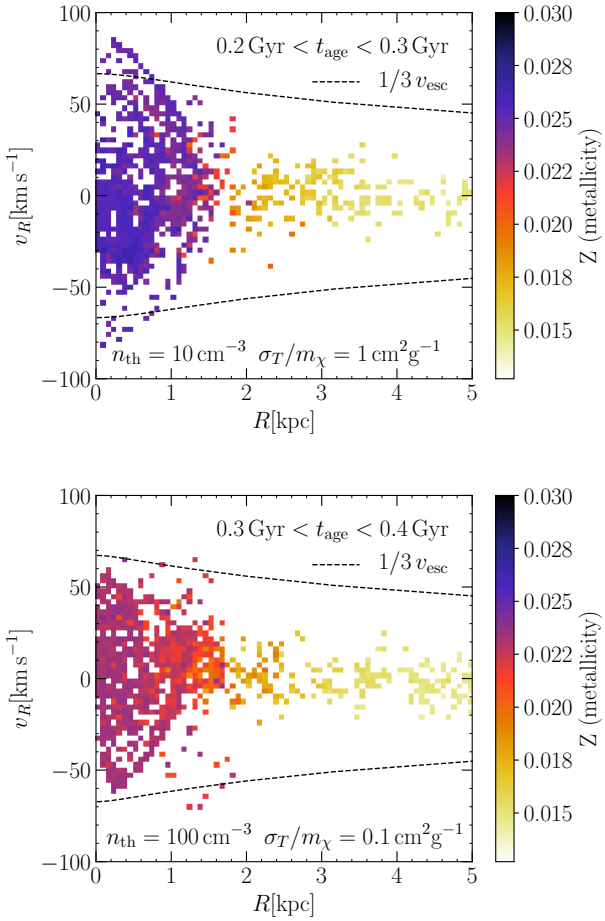

Extended Data Figure 9.2. **Radial phase space projections of the metallicity distribution of star particles within a given age range formed in two different simulations, calculated after different simulation times.** *The top panel corresponds to the simulation with "mildly" bursty SF ($n_{th} = 10\,cm^{-3}$) and $\sigma_T/m_\chi = 1\,cm^2g^{-1}$ and only shows stars which are $0.2-0.3$ Gyr old, calculated from a 2.4 Gyr snapshot. The bottom panel corresponds to the simulation with bursty SF ($n_{th} = 100\,cm^{-3}$) and $\sigma_T/m_\chi = 0.1\,cm^2g^{-1}$ and only shows stars which are $0.3-0.4$ Gyr old, calculated from a 2.4 Gyr snapshot. Stellar metallicities are averaged in 70×70 bins in $R-v_R$ space and the displayed average metallicity per bin is colour-coded as indicated to the right of each panel. Shell-like features such as the ones presented here and in* **Figure** *9.3 are found at different times across all simulations with bursty SF.*

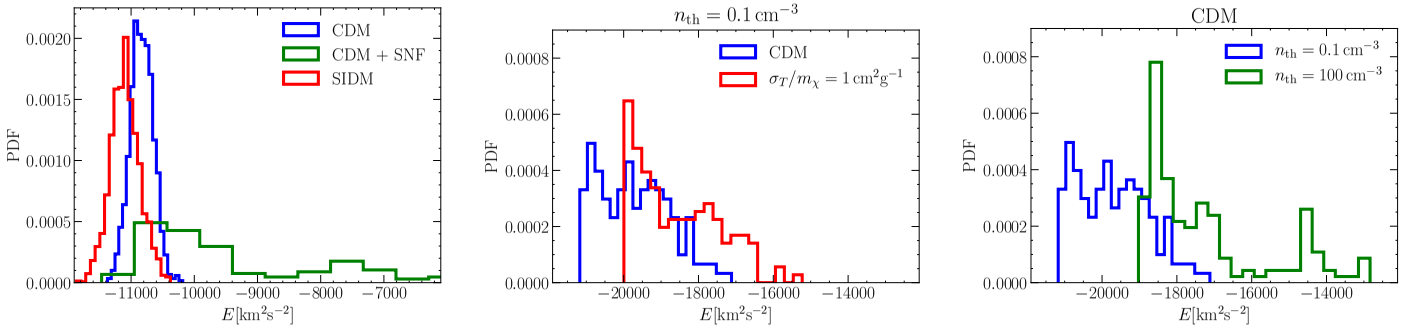

Extended Data Figure 9.3. **Comparison of energy distributions of tracer (stellar) particles from different simulations after 3 Gyr of simulation time.** *The left panel shows the energy distribution of the orbital families of tracer particles whose $R - v_R$ space distributions were presented on the right panel of figure 4 (CDM), as well as on the left (SIDM) and right (CDM+SNF) panels of figure 5 in Burger and Zavala (2019). In these toy model simulations, SNF was implemented as a simplified analytical model and $\sigma_T/m_\chi = 2\,\mathrm{cm^2 g^{-1}}$ was adopted as the fiducial SIDM cross section. On the other panels, we show the distributions of energy per unit mass of stars formed in three of the hydrodynamic simulations introduced above. On the middle panel, we show the energy distribution of the 150 most metal-rich stars which are $0.8 - 0.9\,Gyr$ old in the smooth SF ($n_{\rm th} = 0.1\,\mathrm{cm^{-3}}$) CDM simulation (blue line) and the smooth SF SIDM simulation ($\sigma_T/m_\chi = 1\,\mathrm{cm^2 g^{-1}}$, red line) simulation. On the right panel, we show the energy distribution of the 150 most metal-rich stars which are $0.8 - 0.9\,Gyr$ old in the CDM simulations with smooth and bursty SF with blue (the same blue line as in the middle panel) and green lines, respectively. Qualitatively, we find that the expectations of Burger and Zavala (2019) are verified: adiabatic energy injection (smooth SNF) or redistribution (SIDM) results in a narrower energy distribution of kinematic tracers/stars compared to impulsive energy injection (bursty SNF).*

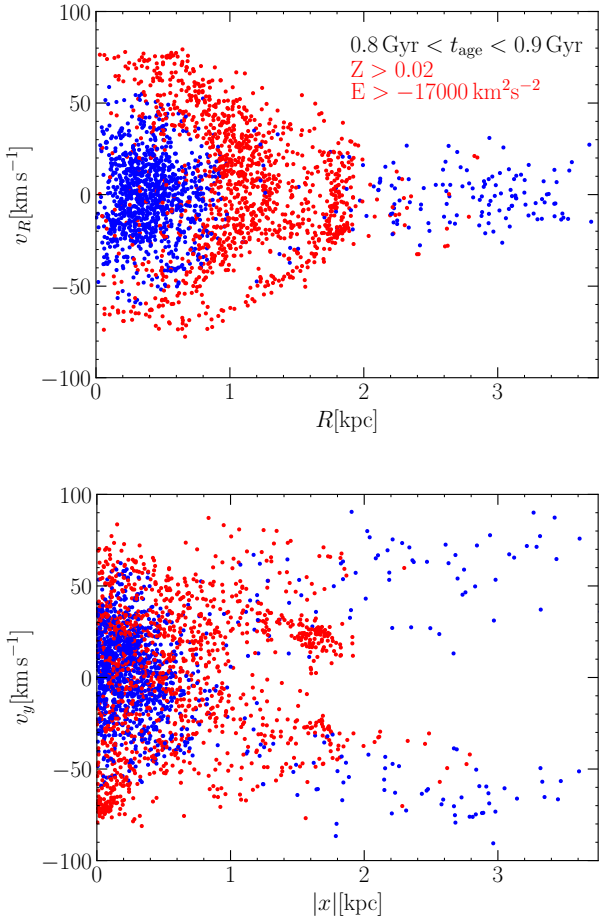

Extended Data Figure 9.4. **Shell-like signatures of early stage phase mixing caused by impulsive SN feedback in two different projections.** *We show scatter plots of the phase space coordinates of stars with ages in the range* $0.8-0.9\,\mathrm{Gyr}$ *after* $3\,\mathrm{Gyr}$ *of simulation time in the CDM run with* $n_{\mathrm{th}} = 100\,\mathrm{cm}^{-3}$. *We show two different projections of the phase space coordinates for all stars:* (R, v_R) *(top panel) and* $(|x|, v_y)$ *(right panel). The star particles selected by the cuts indicated in the left panel are marked red, other star particles blue. We chose the indicated metallicity and specific energy cuts to select the stars that make up the shell-like feature in radial phase space. Notice that the clumps seen in the edge-on projection consist of those same star particles.*

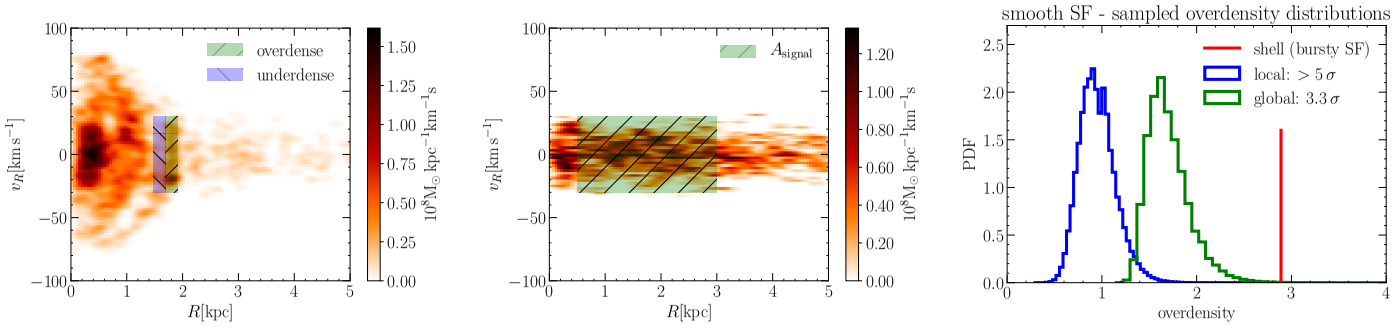

Extended Data Figure 9.5. **Significance of the overdense shell in the CDM simulation with bursty SF.** *The left panel illustrates how the overdensity is defined. The green shaded rectangle defines the region in phase space in which the shell seen in the left panel intersects the phase space region which is occupied in the smooth SF case (see middle panel; essentially defined by the range of radial velocities in the smooth SF case). Dividing the mass contained within this rectangular area in phase space by the mass contained within the adjacent blue shaded rectangle defines the shell overdensity. The middle panel defines* $A_{\rm signal}$*, i.e. the area in phase space in which we expect shell-like overdensities. In the right panel we derive the significance of the shell-like overdensity as follows. Assuming that the underlying mass distribution is given by the smooth SF case, we use the mass distribution of the middle panel as a target distribution for random sampling. We resample the full distribution* 10^7 *times, using the same number of particles as in the original distribution (top right panel). The resulting distribution of the local (global) overdensity is shown as a blue (green) curve. Here, local overdensity is defined using exactly the same phase space areas as in the left panel, while global overdensity refers to the maximum overdensity over* $A_{\rm signal}$ *as defined in the middle panel. The red line is the measured shell overdensity. The estimated global (local) significance of the shell overdensity is* $3.3\,\sigma$ *(*$> 5\,\sigma$*).*

Chapter 10

Summary and Outlook

In this thesis, I have investigated what the kinematic properties of stars that orbit in a virialized system – typically a galaxy contained within a DM halo – can tell us about the evolutionary history of the system's gravitational potential. The bulk part of the work presented in this thesis has been focused on the cusp-core problem in dwarf galaxies. We have compared how the kinematic properties of stellar tracers are affected by two different core formation mechanisms – impulsive feedback from supernovae and adiabatic redistribution of energy in the DM halo due to DM. Moreover, we have taken a close look on how actions of particles evolve in potentials whose time evolution is neither impulsive nor fully adiabatic. We have furthermore shown that the evolution of the radial action distribution of a distribution of tracers orbiting within a time-dependent spherical potential can be derived from the evolution of the action of an individual tracer particle, by treating the tracer distribution as a microcanonical ensemble. There are several interesting avenues for future research starting from the results presented here. To provide an outline of potential follow-up work, we will here briefly review the key results of the five articles presented in this thesis, along with ideas for future projects that constitute a natural continuation of this work.

In Article 1 (chapter 5) we focused on the kinematic properties of stellar tracers in an isolated, core forming, dwarf-sized DM halo. We modeled both adiabatic core formation (through elastic self-scattering between the DM particles) and impulsive core formation (through a sudden removal of an external potential, approximating the effect of SN feedback in DMO simulations). We ran our simulations using a spherically symmetric DM halo and demonstrated that radial actions of tracer particles are conserved only if core formation is adiabatic. The most striking result concerned orbital families, sets of kinematic tracers (stars) with similar integrals of motion, i.e. energy and angular momentum. We found that orbital families remain united in potentials that evolve adiabatically, while they split up into several distinct shells in potentials that evolve impulsively. These shells appear even if the host potential is not spherically symmetric, despite their lifetime being reduced to ~ 1 dynamical time. Eventually, these shell-like features are diluted / erased by phase mixing.

Two of the further articles presented in this thesis explicitly follow up on the work presented in Article 1. In Article 3 (chapter 7), we took a closer look at the impulsive,

SN feedback driven mechanism of core formation. We qualitatively confirmed an earlier (analytic) result by Peñarrubia et al. (2012) that relates the difference between a DM halo's initial and its final density profile to the energy that has been injected by supernovae. To that aim, we developed an effective model for SN feedback in DMO simulations, in which both the gravitational pull of an external "baryonic" potential, and stochastic SN feedback – with a distribution of "superbubbles" that is sampled from the external distribution of baryons – are taken into account. Apart from confirming that the energy injected by SN feedback is the deciding factor when it comes to whether a core can form or not, we obtained two further, general results. Specifically, we demonstrated that SN feedback is a more efficient cusp-core transformation mechanism if feedback is more concentrated, provided the feedback energy injected into the interstellar medium is sufficiently large to drive galactic scale winds. Finally, we showed that SN feedback can be an effective core formation mechanism only if the associated change in the total gravitational potential is impulsive. As a consequence, the shell-like features first identified in Article 1 appear in the radially projected phase space distribution of tracers in all simulations in which a core is formed and the external potential is spherically symmetric. In simulations with an external disk potential, no such features are found at the end of the simulation[1]. Finally, we found in Article 5 (chapter 9) that the shell-like features discovered in Article 1 are recovered in hydrodynamic simulations of an isolated, SMC-like galaxy in a live DM halo. Based on the premise that stars with similar formation times and metallicities constitute orbital families (see also Price-Whelan et al. 2021), we demonstrated that transitory shell-like features appear in the radial projection of the phase space distribution of stars with similar ages, despite the strong deviation from spherical symmetry induced by the SMC-like disk galaxy. Consistent with our findings in Article 3, whether or not such features appeared depended mainly on the burstiness of the simulated star formation history, and in turn on the impulsiveness of SN feedback. The (non-) observation of such shell-like features in the kinematic data of stars in nearby dwarf galaxies would therefore strengthen (disfavour) SN feedback as a core formation mechanism. Finally, we outlined the theoretical mechanism that is behind the formation of these shells. According to Pontzen and Governato (2012), impulsive changes in the gravitational potential change the energy of kinematic tracers by an amount the depends on the tracers' orbital phase. Thus, impulsive SN feedback creates an energy spread in orbital families – and the phase space shells are simply signatures of early stage phase mixing.

The work presented in Articles 1,3, and 5 is purely theoretical and based on numerical simulations of isolated systems. Therefore, there are two obvious directions that future research can take based on the results presented in these articles. Since we predict transitory shell-like features in phase space as a universal signature of impulsive SN feedback in dwarf galaxies, searching for those signatures in nearby, star-forming field dwarfs will be a worthwhile effort. The non-detection of such features in star-forming field dwarfs with a confirmed core would be a serious blow to the ΛCDM paradigm, and therefore future efforts will be dedicated to searching those features. At the current time, data on the kinematics, metallicities, and ages of individual stars in isolated dwarf

[1] They may appear as transitory features. However, as the discussion in this Article focused on the effect of symmetry, we did not check this explicitly at that point, and instead presented a more comprehensive discussion of the phase space shells – and how they may be observed – in Article 5.

galaxies is still sparse. As an exercise, we looked for features associated with impulsive SN feedback in the stellar phase space distribution of the MW dwarf spheroidal galaxy Fornax. Fornax has been claimed to have a core (Peñarrubia et al. 2012) and the required data is available for some of its member stars (Walker et al. 2009). However, Fornax is far from being an ideal target for this analysis, since star formation was quenched $\sim$ 1Gyr ago Rusakov et al. (2021), and, being a MW satellite, it is subject to tidal forces that can dilute signatures of impulsive SN feedback on rather short timescales, as we demonstrated in Article 1. Fortunately, better data is around the corner. The Roman Space telescope (Khan et al. 2018) will provide detailed photometric data for stars within dwarf galaxies in the local volume, including field dwarfs as well as satellites of the MW and M31. Combined with ground based telescopes, this will enable the precise determination of the star formation history in these dwarf galaxies, as well as the approximate ages of individual stars (de Boer et al. 2012, Rusakov et al. 2021). Obtaining furthermore the kinematics of those stars with spectroscopic measurements from ground based telescopes, we will be able to look for phase space shells in all nearby dwarfs. Ideal targets are star-forming, isolated dwarf galaxies with claimed – or confirmed – cores. Until these data become available, more theoretical work can be done in the interest of a more comprehensive analysis. The results presented in Articles 1, 3, and 5 are derived from simulations of isolated systems. Since we are expecting data from dwarf galaxies in the Local Group, an obvious step forward is to explore different DM models and different implementations of SN feedback in a series of hydrodynamic zoom simulations of a similar volume. The zoom volume will be chosen in such a way that its global properties are similar to those of the Local Group. The DM models used in such a suite of zoom simulations should be within current constraints. It is therefore advisable to chose velocity-dependent cross sections when modeling SIDM, as is done in several ETHOS models(Cyr-Racine et al. 2016). Such models can also include deviations from the linear predictions of CDM, i.e. reduction of small scale structure due to collisional – or free streaming – damping. The results of such a suite of zoom simulations will allow for a better modeling – and quantification – of the shell-like features of impulsive SN feedback and their appearance in different cosmologies.

Such a simulation suite will also be helpful to further investigate the results presented in Article 4 (chapter 8). Therein, we looked at different properties of the gas and the stars in 16 hydrodynamical simulations of an isolated, SMC-like system, each of them with a different combination of the SIDM momentum transfer cross section and the gas density threshold for star formation. Two results were particularly striking. We observed an adiabatic expansion of galaxies in simulations with smooth star formation and SIDM induced core formation. Moreover, we showed that impulsive SN feedback creates steep stellar age gradients, both in CDM and in SIDM simulations. While these results are conceptually very interesting, they do not constitute "smoking gun" signatures, such as the phase space shells reported in Article 5. An extended study – preferably using the above-mentioned suite of zoom simulations – is key in order to identify clear differences between different DM models / SN feedback realizations within galaxies that resemble dwarfs in the Local Group. Such an improved classification of, for instance, the age gradients and the relative size of galaxies can then be used to differentiate between models, using current and future observational data sets.

In Article 2 (chapter 6), we shifted the focus away from the study of dwarf galaxies.

Instead, we here focused on the differences between adiabatic and impulsive evolution of gravitational potentials in general. Radial actions are adiabatic invariants in evolving spherical potentials, but they change discontinuously in impulsively evolving (spherical) potentials. In Article 2, we demonstrated that the evolution of radial actions in mildly time-dependent potentials is, to linear order, given by an oscillation around a constant value, the amplitude of which is determined by the shape of the potential and the rate at which it changes. For ensembles of kinematic tracers, we showed that the regime in which the evolution of the actions of individual particles is oscillatory corresponds to a "diffusive" regime; and we developed a statistical diffusion theory for the evolution of radial action distributions of kinematic tracers in time-dependent spherical potentials. Moreover, we demonstrated that an asymmetric drift towards lower radial actions occurs in the impulsive, non-diffusive regime. We hypothesized that this asymmetric drift may be related to the formation of cusps in collapsing CDM haloes, and argued that the analysis of tidal streams in the MW (e.g. Buist and Helmi 2015) can be improved by including the diffusion in radial action space that is related to the time-dependence of the MW potential as a baseline systematic error. Currently, our diffusion theory is limited to the analysis of tracers in spherical potentials. In the future, we will improve on that by developing a theory for massive particles and extending the formalism into two dimensions in integral-of-motion space. In this way, we will be able to model resonant diffusion and quantify the importance of diffusion effects for the analysis of tidal streams. Moreover, developing a theory for massive particles is an important step in order to use our formalism to make stronger statements about the (presumed) interconnection between the fast evolving potentials of recently collapsed DM haloes and the formation of primordial density cusps.

Chapter 11

Appendix - TBTF and alternatives to CDM

In Zavala et al. (2019)[1], we took a closer look at how the modified version of TBTF (right panel of figure 3.1) may be used to test predictions from alternative DM models. Using DMO high resolution zoom simulations of a MW-like halo, we investigated how well the circular velocity profiles of the most massive simulated subhaloes in different cosmologies match the measured circular velocities of the MW's dwarf satellites. We find that:

- SIDM with $\sigma_T/m_\chi = 1\,\mathrm{cm}^2\mathrm{g}^{-1}$ is disfavoured by the data.
- Thermal relic WDM with $m_{\mathrm{WDM}} = 2.3\,\mathrm{keV}$ improves the number of matches considerably. The measured circular velocities of 21/24 satellites are consistent with the density profiles of the heaviest subhaloes
- ETHOS-4 (Cyr-Racine et al. 2016, Vogelsberger et al. 2016), a model with both collisional damping at early times and self-scattering at late times, also increases the number of matches between theory and data (19/24).
- In contrast to SIDM with $\sigma_T/m_\chi = 1\,\mathrm{cm}^2\mathrm{g}^{-1}$, SIDM with a strongly velocity-dependent cross section can explain the circular velocities of the ultra-faints. The velocity dependence of the cross section must be such that the heaviest subhaloes undergo gravothermal collapse (Balberg et al. 2002, Colín et al. 2002, Koda and Shapiro 2011, Pollack et al. 2015, Nishikawa et al. 2020), significantly increasing their central densities, while lighter subhaloes form extended constant density cores. This naturally creates a more diverse distribution of the subhaloes' circular velocity profiles and increases the number of matched data points compared to CDM (19/24).

As all simulations are DMO, the effect of baryons is not taken into account. As we argue in Zavala et al. (2019), baryons are likely to increase the diversity of the subhaloes'

[1] 1904.09998

circular velocity profiles and may thus universally lead to a better match between theory and data. However, baryonic effects are unlikely to increase the central densities of subhaloes, and thus the statement that $1\,\mathrm{cm}^2\mathrm{g}^{-1}$-SIDM is disfavoured by the data is likely solid.

Bibliography

Aarseth, S. J. (2003). *Gravitational N-Body Simulations: Tools and Algorithms*. Cambridge Monographs on Mathematical Physics. Cambridge University Press.

Afshordi, N., Mohayaee, R., and Bertschinger, E. (2009). Hierarchical phase space structure of dark matter haloes: Tidal debris, caustics, and dark matter annihilation. Phys. Rev. D, 79(8):083526.

Agertz, O., Kravtsov, A. V., Leitner, S. N., and Gnedin, N. Y. (2013). Toward a Complete Accounting of Energy and Momentum from Stellar Feedback in Galaxy Formation Simulations. ApJ, 770(1):25.

Agertz, O., Teyssier, R., and Moore, B. (2011). The formation of disc galaxies in a ΛCDM universe. MNRAS, 410(2):1391–1408.

Aguilar, L. A. and White, S. D. M. (1985). Tidal interactions between spherical galaxies. ApJ, 295:374.

Aguilar, L. A. and White, S. D. M. (1986). The Density Profiles of Tidally Stripped Galaxies. ApJ, 307:97.

Alam, S. M. K., Bullock, J. S., and Weinberg, D. H. (2002). Dark Matter Properties and Halo Central Densities. ApJ, 572(1):34–40.

Amorisco, N. C., Zavala, J., and de Boer, T. J. L. (2014). Dark Matter Cores in the Fornax and Sculptor Dwarf Galaxies: Joining Halo Assembly and Detailed Star Formation Histories. ApJ, 782(2):L39.

Anderhalden, D. and Diemand, J. (2013). Density profiles of CDM microhalos and their implications for annihilation boost factors. J. Cosmology Astropart. Phys., 2013(4):009.

Angulo, R. E., Hahn, O., and Abel, T. (2013). The warm dark matter halo mass function below the cut-off scale. MNRAS, 434(4):3337–3347.

Angulo, R. E., Lacey, C. G., Baugh, C. M., and Frenk, C. S. (2009). The fate of substructures in cold dark matter haloes. MNRAS, 399(2):983–995.

Antonov, V. A. (1961). Remarks on the Problem of Stability in Stellar Dynamics. Soviet Ast., 4:859.

Arad, I. and Johansson, P. H. (2005). A numerical comparison of theories of violent relaxation. MNRAS, 362(1):252–262.

Arraki, K. S., Klypin, A., More, S., and Trujillo-Gomez, S. (2014). Effects of baryon removal on the structure of dwarf spheroidal galaxies. MNRAS, 438(2):1466–1482.

Asplund, M., Grevesse, N., Sauval, A. J., and Scott, P. (2009). The Chemical Composition of the Sun. *Annual Review of Astronomy and Astrophysics*, 47(1):481–522.

Bahl, H. and Baumgardt, H. (2014). A comparison of the distribution of satellite galaxies around Andromeda and the results of ΛCDM simulations. MNRAS, 438(4):2916–2923.

Balberg, S., Shapiro, S. L., and Inagaki, S. (2002). Self-Interacting Dark Matter Halos and the Gravothermal Catastrophe. ApJ, 568(2):475–487.

Balescu, R. (1960). Irreversible Processes in Ionized Gases. *Physics of Fluids*, 3(1):52–63.

Barnes, J. and Hut, P. (1986). A hierarchical O(N log N) force-calculation algorithm. Nature, 324(6096):446–449.

Baugh, C. M., Gaztanaga, E., and Efstathiou, G. (1995). A comparison of the evolution of density fields in perturbation theory and numerical simulations - II. Counts-in-cells analysis. MNRAS, 274(4):1049–1070.

Baumann, D., Nicolis, A., Senatore, L., and Zaldarriaga, M. (2012). Cosmological non-linearities as an effective fluid. J. Cosmology Astropart. Phys., 2012(7):051.

Behroozi, P. S., Wechsler, R. H., and Conroy, C. (2013). The Average Star Formation Histories of Galaxies in Dark Matter Halos from z = 0-8. ApJ, 770(1):57.

Benítez-Llambay, A., Frenk, C. S., Ludlow, A. D., and Navarro, J. F. (2019). Baryon-induced dark matter cores in the EAGLE simulations. MNRAS, 488(2):2387–2404.

Benson, A. J., Cole, S., Frenk, C. S., Baugh, C. M., and Lacey, C. G. (2000). The nature of galaxy bias and clustering. MNRAS, 311(4):793–808.

Benson, A. J., Lacey, C. G., Baugh, C. M., Cole, S., and Frenk, C. S. (2002). The effects of photoionization on galaxy formation - I. Model and results at z=0. MNRAS, 333(1):156–176.

Bertone, G. and Hooper, D. (2018). History of dark matter. *Reviews of Modern Physics*, 90(4):045002.

Bertschinger, E. (1985). Self-similar secondary infall and accretion in an Einstein-de Sitter universe. ApJS, 58:39–65.

Binney, J. and Tremaine, S. (1987). *Galactic dynamics*.

Binney, J. and Tremaine, S. (2008). *Galactic Dynamics: Second Edition*.

Blumenthal, G. R., Faber, S. M., Primack, J. R., and Rees, M. J. (1984). Formation of galaxies and large-scale structure with cold dark matter. Nature, 311:517–525.

Bode, P., Ostriker, J. P., and Turok, N. (2001). Halo Formation in Warm Dark Matter Models. ApJ, 556(1):93–107.

Boehm, C., Riazuelo, A., Hansen, S. H., and Schaeffer, R. (2002). Interacting dark matter disguised as warm dark matter. *Phys. Rev. D*, 66:083505.

Boehm, C., Schewtschenko, J. A., Wilkinson, R. J., Baugh, C. M., and Pascoli, S. (2014). Using the Milky Way satellites to study interactions between cold dark matter and radiation. MNRAS, 445:L31–L35.

Bohr, S., Zavala, J., Cyr-Racine, F.-Y., Vogelsberger, M., Bringmann, T., and Pfrommer, C. (2020). ETHOS - an effective parametrization and classification for structure formation: the non-linear regime at $z \gtrsim 5$. MNRAS, 498(3):3403–3419.

Bonamigo, M., Despali, G., Limousin, M., Angulo, R., Giocoli, C., and Soucail, G. (2015). Universality of dark matter haloes shape over six decades in mass: insights from the Millennium XXL and SBARBINE simulations. MNRAS, 449(3):3171–3182.

Bond, J. R., Cole, S., Efstathiou, G., and Kaiser, N. (1991). Excursion Set Mass

Functions for Hierarchical Gaussian Fluctuations. ApJ, 379:440.

Bose, S., Frenk, C. S., Jenkins, A., Fattahi, A., Gómez, F. A., Grand, R. J. J., Marinacci, F., Navarro, J. F., Oman, K. A., Pakmor, R., Schaye, J., Simpson, C. M., and Springel, V. (2019a). No cores in dark matter-dominated dwarf galaxies with bursty star formation histories. MNRAS, 486(4):4790–4804.

Bose, S., Hellwing, W. A., Frenk, C. S., Jenkins, A., Lovell, M. R., Helly, J. C., and Li, B. (2016). The Copernicus Complexio: statistical properties of warm dark matter haloes. MNRAS, 455(1):318–333.

Bose, S., Vogelsberger, M., Zavala, J., Pfrommer, C., Cyr-Racine, F.-Y., Bohr, S., and Bringmann, T. (2019b). ETHOS - an Effective Theory of Structure Formation: detecting dark matter interactions through the Lyman-α forest. MNRAS, 487(1):522–536.

Bosma, A. (1981a). 21-cm line studies of spiral galaxies. I. Observations of the galaxies NGC 5033, 3198, 5055, 2841, and 7331. AJ, 86:1791–1824.

Bosma, A. (1981b). 21-cm line studies of spiral galaxies. II. The distribution and kinematics of neutral hydrogen in spiral galaxies of various morphological types. AJ, 86:1825–1846.

Bovill, M. S. and Ricotti, M. (2009). Pre-Reionization Fossils, Ultra-Faint Dwarfs, and the Missing Galactic Satellite Problem. ApJ, 693(2):1859–1870.

Bower, R. G. (1991). The evolution of groups of galaxies in the Press-Schechter formalism. MNRAS, 248:332–352.

Boylan-Kolchin, M., Bullock, J. S., and Kaplinghat, M. (2011). Too big to fail? The puzzling darkness of massive Milky Way subhaloes. MNRAS, 415(1):L40–L44.

Boylan-Kolchin, M., Bullock, J. S., and Kaplinghat, M. (2012). The Milky Way's bright satellites as an apparent failure of ΛCDM. MNRAS, 422(2):1203–1218.

Boylan-Kolchin, M., Ma, C.-P., and Quataert, E. (2008). Dynamical friction and galaxy merging time-scales. MNRAS, 383(1):93–101.

Boylan-Kolchin, M., Springel, V., White, S. D. M., and Jenkins, A. (2010). There's no place like home? Statistics of Milky Way-mass dark matter haloes. MNRAS, 406(2):896–912.

Boylan-Kolchin, M., Springel, V., White, S. D. M., Jenkins, A., and Lemson, G. (2009). Resolving cosmic structure formation with the Millennium-II Simulation. MNRAS, 398(3):1150–1164.

Bringmann, T. (2009). Particle models and the small-scale structure of dark matter. *New Journal of Physics*, 11(10):105027.

Brook, C. B. and Di Cintio, A. (2015). Expanded haloes, abundance matching and too-big-to-fail in the Local Group. MNRAS, 450(4):3920–3934.

Brooks, A. M., Papastergis, E., Christensen, C. R., Governato, F., Stilp, A., Quinn, T. R., and Wadsley, J. (2017). How to Reconcile the Observed Velocity Function of Galaxies with Theory. ApJ, 850(1):97.

Brooks, A. M. and Zolotov, A. (2014). Why Baryons Matter: The Kinematics of Dwarf Spheroidal Satellites. ApJ, 786(2):87.

Bryan, G. L. and Norman, M. L. (1998). Statistical Properties of X-Ray Clusters: Analytic and Numerical Comparisons. ApJ, 495(1):80–99.

Buck, T., Dutton, A. A., and Macciò, A. V. (2016). Simulated ΛCDM analogues of the thin plane of satellites around the Andromeda galaxy are not kinematically coherent

structures. MNRAS, 460(4):4348–4365.
Buckley, M. R., Zavala, J., Cyr-Racine, F.-Y., Sigurdson, K., and Vogelsberger, M. (2014). Scattering, damping, and acoustic oscillations: Simulating the structure of dark matter halos with relativistic force carriers. Phys. Rev. D, 90(4):043524.
Buist, H. J. T. and Helmi, A. (2015). The evolution of streams in a time-dependent potential. A&A, 584:A120.
Bullock, J. S. and Boylan-Kolchin, M. (2017). Small-Scale Challenges to the ΛCDM Paradigm. ARA&A, 55(1):343–387.
Bullock, J. S., Kolatt, T. S., Sigad, Y., Somerville, R. S., Kravtsov, A. V., Klypin, A. A., Primack, J. R., and Dekel, A. (2001). Profiles of dark haloes: evolution, scatter and environment. MNRAS, 321(3):559–575.
Bullock, J. S., Kravtsov, A. V., and Weinberg, D. H. (2000). Reionization and the Abundance of Galactic Satellites. ApJ, 539(2):517–521.
Burger, J. D., Peñarrubia, J., and Zavala, J. (2020). Conservation of radial actions in time-dependent spherical potentials. *arXiv e-prints*, page arXiv:2012.00737.
Burger, J. D. and Zavala, J. (2019). The nature of core formation in dark matter haloes: adiabatic or impulsive? MNRAS, 485(1):1008–1028.
Burger, J. D. and Zavala, J. (2021). SN-driven mechanism of cusp-core transformation: an appraisal. *arXiv e-prints*, page arXiv:2103.01231.
Burkert, A. (1995). The Structure of Dark Matter Halos in Dwarf Galaxies. ApJ, 447:L25–L28.
Carlson, E. D., Machacek, M. E., and Hall, L. J. (1992). Self-interacting Dark Matter. ApJ, 398:43.
Carrasco, J. J. M., Hertzberg, M. P., and Senatore, L. (2012). The effective field theory of cosmological large scale structures. *Journal of High Energy Physics*, 2012:82.
Chabrier, G. (2001). The Galactic Disk Mass Budget. I. Stellar Mass Function and Density. ApJ, 554(2):1274–1281.
Chan, T. K., Kereš, D., Oñorbe, J., Hopkins, P. F., Muratov, A. L., Faucher-Giguère, C. A., and Quataert, E. (2015). The impact of baryonic physics on the structure of dark matter haloes: the view from the FIRE cosmological simulations. MNRAS, 454(3):2981–3001.
Chandrasekhar, S. (1943). Dynamical Friction. I. General Considerations: the Coefficient of Dynamical Friction. ApJ, 97:255.
Cole, S., Aragon-Salamanca, A., Frenk, C. S., Navarro, J. F., and Zepf, S. E. (1994). A recipe for galaxy formation. MNRAS, 271:781–806.
Cole, S. and Kaiser, N. (1989). Biased clustering in the cold dark matter cosmogony. MNRAS, 237:1127–1146.
Cole, S. and Lacey, C. (1996). The structure of dark matter haloes in hierarchical clustering models. MNRAS, 281:716.
Colín, P., Avila-Reese, V., Valenzuela, O., and Firmani, C. (2002). Structure and Subhalo Population of Halos in a Self-interacting Dark Matter Cosmology. ApJ, 581(2):777–793.
Colless, M., Dalton, G., Maddox, S., Sutherland, W., Norberg, P., Cole, S., Bland-Hawthorn, J., Bridges, T., Cannon, R., Collins, C., Couch, W., Cross, N., Deeley, K., De Propris, R., Driver, S. P., Efstathiou, G., Ellis, R. S., Frenk, C. S., Glazebrook, K., Jackson, C., Lahav, O., Lewis, I., Lumsden, S., Madgwick, D., Peacock, J. A.,

Peterson, B. A., Price, I., Seaborne, M., and Taylor, K. (2001). The 2dF Galaxy Redshift Survey: spectra and redshifts. MNRAS, 328(4):1039–1063.

Colombi, S., Dodelson, S., and Widrow, L. M. (1996). Large-Scale Structure Tests of Warm Dark Matter. ApJ, 458:1.

Conn, A. R., Lewis, G. F., Ibata, R. A., Parker, Q. A., Zucker, D. B., McConnachie, A. W., Martin, N. F., Valls-Gabaud, D., Tanvir, N., Irwin, M. J., Ferguson, A. M. N., and Chapman, S. C. (2013). The Three-dimensional Structure of the M31 Satellite System; Strong Evidence for an Inhomogeneous Distribution of Satellites. ApJ, 766(2):120.

Conroy, C., Wechsler, R. H., and Kravtsov, A. V. (2006). Modeling Luminosity-dependent Galaxy Clustering through Cosmic Time. ApJ, 647(1):201–214.

Cooray, A. and Sheth, R. (2002). Halo models of large scale structure. Phys. Rep., 372(1):1–129.

Correa, C. A. (2021). Constraining velocity-dependent self-interacting dark matter with the Milky Way's dwarf spheroidal galaxies. MNRAS, 503(1):920–937.

Cowsik, R. and McClelland, J. (1973). Gravity of Neutrinos of Nonzero Mass in Astrophysics. ApJ, 180:7–10.

Crain, R. A., Theuns, T., Dalla Vecchia, C., Eke, V. R., Frenk, C. S., Jenkins, A., Kay, S. T., Peacock, J. A., Pearce, F. R., Schaye, J., Springel, V., Thomas, P. A., White, S. D. M., and Wiersma, R. P. C. (2009). Galaxies-intergalactic medium interaction calculation - I. Galaxy formation as a function of large-scale environment. MNRAS, 399(4):1773–1794.

Cuesta, A. J., Prada, F., Klypin, A., and Moles, M. (2008). The virialized mass of dark matter haloes. MNRAS, 389(1):385–397.

Cyr-Racine, F.-Y., Sigurdson, K., Zavala, J., Bringmann, T., Vogelsberger, M., and Pfrommer, C. (2016). ETHOS—an effective theory of structure formation: From dark particle physics to the matter distribution of the Universe. Phys. Rev. D, 93(12):123527.

Dalla Vecchia, C. and Schaye, J. (2012). Simulating galactic outflows with thermal supernova feedback. MNRAS, 426(1):140–158.

Davé, R., Spergel, D. N., Steinhardt, P. J., and Wandelt, B. D. (2001). Halo Properties in Cosmological Simulations of Self-interacting Cold Dark Matter. ApJ, 547(2):574–589.

Davis, M., Efstathiou, G., Frenk, C. S., and White, S. D. M. (1985). The evolution of large-scale structure in a universe dominated by cold dark matter. ApJ, 292:371–394.

de Bernardis, P., Ade, P. A. R., Bock, J. J., Bond, J. R., Borrill, J., Boscaleri, A., Coble, K., Crill, B. P., De Gasperis, G., Farese, P. C., Ferreira, P. G., Ganga, K., Giacometti, M., Hivon, E., Hristov, V. V., Iacoangeli, A., Jaffe, A. H., Lange, A. E., Martinis, L., Masi, S., Mason, P. V., Mauskopf, P. D., Melchiorri, A., Miglio, L., Montroy, T., Netterfield, C. B., Pascale, E., Piacentini, F., Pogosyan, D., Prunet, S., Rao, S., Romeo, G., Ruhl, J. E., Scaramuzzi, F., Sforna, D., and Vittorio, N. (2000). A flat Universe from high-resolution maps of the cosmic microwave background radiation. Nature, 404(6781):955–959.

de Blok, W. J. G., Walter, F., Brinks, E., Trachternach, C., Oh, S. H., and Kennicutt, R. C., J. (2008). High-Resolution Rotation Curves and Galaxy Mass Models from THINGS. AJ, 136(6):2648–2719.

de Boer, T. J. L., Tolstoy, E., Hill, V., Saha, A., Olszewski, E. W., Mateo, M., Starkenburg, E., Battaglia, G., and Walker, M. G. (2012). The star formation and chemical evolution history of the Fornax dwarf spheroidal galaxy. *Astronomy & Astrophysics*, 544:A73.

De Rijcke, S., Fouvry, J.-B., and Pichon, C. (2019). Instabilities in disc galaxies: from noise to grooves to spirals. MNRAS, 484(3):3198–3208.

Dehnen, W. (1993). A Family of Potential-Density Pairs for Spherical Galaxies and Bulges. MNRAS, 265:250.

Dehnen, W. (2000). A Very Fast and Momentum-conserving Tree Code. ApJ, 536(1):L39–L42.

Dehnen, W. (2002). A hierarchical o(n) force calculation algorithm. *Journal of Computational Physics*, 179(1):27–42.

Dehnen, W. and Read, J. I. (2011). N-body simulations of gravitational dynamics. *European Physical Journal Plus*, 126:55.

Delos, M. S., Erickcek, A. L., Bailey, A. P., and Alvarez, M. A. (2018). Density profiles of ultracompact minihalos: Implications for constraining the primordial power spectrum. Phys. Rev. D, 98(6):063527.

Despali, G., Giocoli, C., and Tormen, G. (2014). Some like it triaxial: the universality of dark matter halo shapes and their evolution along the cosmic time. MNRAS, 443(4):3208–3217.

Di Cintio, A., Brook, C. B., Macciò, A. V., Stinson, G. S., Knebe, A., Dutton, A. A., and Wadsley, J. (2014). The dependence of dark matter profiles on the stellar-to-halo mass ratio: a prediction for cusps versus cores. MNRAS, 437(1):415–423.

Di Cintio, A. and Lelli, F. (2016). The mass discrepancy acceleration relation in a ΛCDM context. MNRAS, 456(1):L127–L131.

Diemand, J., Kuhlen, M., and Madau, P. (2007). Formation and Evolution of Galaxy Dark Matter Halos and Their Substructure. ApJ, 667(2):859–877.

Diemer, B. and Kravtsov, A. V. (2015). A Universal Model for Halo Concentrations. ApJ, 799(1):108.

D'Onghia, E., Springel, V., Hernquist, L., and Keres, D. (2010). Substructure Depletion in the Milky Way Halo by the Disk. ApJ, 709(2):1138–1147.

Drlica-Wagner, A., Bechtol, K., Rykoff, E. S., Luque, E., Queiroz, A., Mao, Y. Y., Wechsler, R. H., Simon, J. D., Santiago, B., Yanny, B., Balbinot, E., Dodelson, S., Fausti Neto, A., James, D. J., Li, T. S., Maia, M. A. G., Marshall, J. L., Pieres, A., Stringer, K., Walker, A. R., Abbott, T. M. C., Abdalla, F. B., Allam, S., Benoit-Lévy, A., Bernstein, G. M., Bertin, E., Brooks, D., Buckley-Geer, E., Burke, D. L., Carnero Rosell, A., Carrasco Kind, M., Carretero, J., Crocce, M., da Costa, L. N., Desai, S., Diehl, H. T., Dietrich, J. P., Doel, P., Eifler, T. F., Evrard, A. E., Finley, D. A., Flaugher, B., Fosalba, P., Frieman, J., Gaztanaga, E., Gerdes, D. W., Gruen, D., Gruendl, R. A., Gutierrez, G., Honscheid, K., Kuehn, K., Kuropatkin, N., Lahav, O., Martini, P., Miquel, R., Nord, B., Ogando, R., Plazas, A. A., Reil, K., Roodman, A., Sako, M., Sanchez, E., Scarpine, V., Schubnell, M., Sevilla-Noarbe, I., Smith, R. C., Soares-Santos, M., Sobreira, F., Suchyta, E., Swanson, M. E. C., Tarle, G., Tucker, D., Vikram, V., Wester, W., Zhang, Y., Zuntz, J., and DES Collaboration (2015). Eight Ultra-faint Galaxy Candidates Discovered in Year Two of the Dark Energy Survey. ApJ, 813(2):109.

Dutton, A. A., Buck, T., Macciò, A. V., Dixon, K. L., Blank, M., and Obreja, A. (2020). NIHAO - XXV. Convergence in the cusp-core transformation of cold dark matter haloes at high star formation thresholds. MNRAS, 499(2):2648–2661.

Dutton, A. A., Macciò, A. V., Frings, J., Wang, L., Stinson, G. S., Penzo, C., and Kang, X. (2016). NIHAO V: too big does not fail - reconciling the conflict between ΛCDM predictions and the circular velocities of nearby field galaxies. MNRAS, 457(1):L74–L78.

Eddington, A. S. (1916). The distribution of stars in globular clusters. MNRAS, 76:572–585.

Efstathiou, G. (1992). Suppressing the formation of dwarf galaxies via photoionization. MNRAS, 256(2):43P–47P.

Efstathiou, G., Davis, M., White, S. D. M., and Frenk, C. S. (1985). Numerical techniques for large cosmological N-body simulations. ApJS, 57:241–260.

Einasto, J., Saar, E., Kaasik, A., and Chernin, A. D. (1974). Missing mass around galaxies - Morphological evidence. Nature, 252:111–113.

Einstein, A. (1905). Über die von der molekularkinetischen theorie der wärme geforderte bewegung von in ruhenden flüssigkeiten suspendierten teilchen. *Annalen der Physik*, 322(8):549–560.

Eke, V. R., Cole, S., and Frenk, C. S. (1996). Cluster evolution as a diagnostic for Omega. MNRAS, 282:263–280.

Eke, V. R., Navarro, J. F., and Steinmetz, M. (2001). The Power Spectrum Dependence of Dark Matter Halo Concentrations. ApJ, 554(1):114–125.

Elbert, O. D., Bullock, J. S., Garrison-Kimmel, S., Rocha, M., Oñorbe, J., and Peter, A. H. G. (2015). Core formation in dwarf haloes with self-interacting dark matter: no fine-tuning necessary. MNRAS, 453:29–37.

Ellis, J., Hagelin, J. S., Nanopoulos, D. V., Olive, K., and Srednicki, M. (1984). Supersymmetric relics from the big bang. *Nuclear Physics B*, 238(2):453–476.

Emami, N., Siana, B., Weisz, D. R., Johnson, B. D., Ma, X., and El-Badry, K. (2019). A Closer Look at Bursty Star Formation with L $_{H\alpha}$ and L $_{UV}$ Distributions. *Astrophysical Journal*, 881(1):71.

Errani, R., Peñarrubia, J., and Walker, M. G. (2018). Systematics in virial mass estimators for pressure-supported systems. MNRAS, 481(4):5073–5090.

Fattahi, A., Navarro, J. F., Frenk, C. S., Oman, K. A., Sawala, T., and Schaller, M. (2018). Tidal stripping and the structure of dwarf galaxies in the Local Group. MNRAS, 476(3):3816–3836.

Ferrero, I., Abadi, M. G., Navarro, J. F., Sales, L. V., and Gurovich, S. (2012). The dark matter haloes of dwarf galaxies: a challenge for the Λ cold dark matter paradigm? MNRAS, 425(4):2817–2823.

Fitts, A., Boylan-Kolchin, M., Elbert, O. D., Bullock, J. S., Hopkins, P. F., Oñorbe, J., Wetzel, A., Wheeler, C., Faucher-Giguère, C.-A., Kereš, D., Skillman, E. D., and Weisz, D. R. (2017). fire in the field: simulating the threshold of galaxy formation. MNRAS, 471(3):3547–3562.

Fouvry, J. B., Pichon, C., Magorrian, J., and Chavanis, P. H. (2015). Secular diffusion in discrete self-gravitating tepid discs II. Accounting for swing amplification via the matrix method. A&A, 584:A129.

Freeman, K. C. (1970). On the Disks of Spiral and S0 Galaxies. ApJ, 160:811.

Frenk, C. S. and White, S. D. M. (2012). Dark matter and cosmic structure. *Annalen der Physik*, 524(9-10):507–534.

Frenk, C. S., White, S. D. M., and Davis, M. (1983). Nonlinear evolution of large-scale structure in the universe. ApJ, 271:417–430.

Frenk, C. S., White, S. D. M., Davis, M., and Efstathiou, G. (1988). The Formation of Dark Halos in a Universe Dominated by Cold Dark Matter. ApJ, 327:507.

Fry, A. B., Governato, F., Pontzen, A., Quinn, T., Tremmel, M., Anderson, L., Menon, H., Brooks, A. M., and Wadsley, J. (2015). All about baryons: revisiting SIDM predictions at small halo masses. MNRAS, 452:1468–1479.

Ganeshaiah Veena, P., Cautun, M., van de Weygaert, R., Tempel, E., Jones, B. J. T., Rieder, S., and Frenk, C. S. (2018). The Cosmic Ballet: spin and shape alignments of haloes in the cosmic web. MNRAS, 481(1):414–438.

Gao, L., Frenk, C. S., Boylan-Kolchin, M., Jenkins, A., Springel, V., and White, S. D. M. (2011). The statistics of the subhalo abundance of dark matter haloes. MNRAS, 410(4):2309–2314.

Gao, L., Navarro, J. F., Frenk, C. S., Jenkins, A., Springel, V., and White, S. D. M. (2012). The Phoenix Project: the dark side of rich Galaxy clusters. MNRAS, 425(3):2169–2186.

Gao, L., White, S. D. M., Jenkins, A., Stoehr, F., and Springel, V. (2004). The subhalo populations of ΛCDM dark haloes. MNRAS, 355(3):819–834.

Garrison-Kimmel, S., Boylan-Kolchin, M., Bullock, J. S., and Lee, K. (2014). ELVIS: Exploring the Local Volume in Simulations. MNRAS, 438(3):2578–2596.

Garrison-Kimmel, S., Rocha, M., Boylan-Kolchin, M., Bullock, J. S., and Lally, J. (2013). Can feedback solve the too-big-to-fail problem? MNRAS, 433(4):3539–3546.

Geller, M. J. and Huchra, J. P. (1989). Mapping the Universe. *Science*, 246(4932):897–903.

Genina, A., Benítez-Llambay, A., Frenk, C. S., Cole, S., Fattahi, A., Navarro, J. F., Oman, K. A., Sawala, T., and Theuns, T. (2018). The core-cusp problem: a matter of perspective. MNRAS, 474(1):1398–1411.

Genina, A., Read, J. I., Fattahi, A., and Frenk, C. S. (2020). Can tides explain the low dark matter density in Fornax? *arXiv e-prints*, page arXiv:2011.09482.

Gillet, N., Ocvirk, P., Aubert, D., Knebe, A., Libeskind, N., Yepes, G., Gottlöber, S., and Hoffman, Y. (2015). Vast Planes of Satellites in a High-resolution Simulation of the Local Group: Comparison to Andromeda. ApJ, 800(1):34.

Gnedin, O. Y., Hernquist, L., and Ostriker, J. P. (1999). Tidal Shocking by Extended Mass Distributions. ApJ, 514(1):109–118.

Gnedin, O. Y. and Zhao, H. (2002). Maximum feedback and dark matter profiles of dwarf galaxies. MNRAS, 333(2):299–306.

Goodman, J. and Binney, J. (1984). Adding a point mass to a spherical stellar system. MNRAS, 207:511–515.

Gott, J. Richard, I., Jurić, M., Schlegel, D., Hoyle, F., Vogeley, M., Tegmark, M., Bahcall, N., and Brinkmann, J. (2005). A Map of the Universe. ApJ, 624(2):463–484.

Griffen, B. F., Ji, A. P., Dooley, G. A., Gómez, F. A., Vogelsberger, M., O'Shea, B. W., and Frebel, A. (2016). The Caterpillar Project: A Large Suite of Milky Way Sized

Halos. ApJ, 818(1):10.
Gunn, J. E. and Gott, J. Richard, I. (1972). On the Infall of Matter Into Clusters of Galaxies and Some Effects on Their Evolution. ApJ, 176:1.
Guth, A. H. (1981). Inflationary universe: A possible solution to the horizon and flatness problems. Phys. Rev. D, 23(2):347–356.
Guth, A. H. and Pi, S. Y. (1982). Fluctuations in the New Inflationary Universe. Phys. Rev. Lett., 49(15):1110–1113.
Hanany, S., Ade, P., Balbi, A., Bock, J., Borrill, J., Boscaleri, A., de Bernardis, P., Ferreira, P. G., Hristov, V. V., Jaffe, A. H., Lange, A. E., Lee, A. T., Mauskopf, P. D., Netterfield, C. B., Oh, S., Pascale, E., Rabii, B., Richards, P. L., Smoot, G. F., Stompor, R., Winant, C. D., and Wu, J. H. P. (2000). MAXIMA-1: A Measurement of the Cosmic Microwave Background Anisotropy on Angular Scales of 10'-5°. ApJ, 545(1):L5–L9.
Hansen, S. H. and Moore, B. (2006). A universal density slope Velocity anisotropy relation for relaxed structures. New A, 11(5):333–338.
Harvey, D., Revaz, Y., Robertson, A., and Hausammann, L. (2018). The impact of cored density profiles on the observable quantities of dwarf spheroidal galaxies. MNRAS, 481:L89–L93.
Hawking, S. W. (1982). The development of irregularities in a single bubble inflationary universe. *Physics Letters B*, 115(4):295–297.
Hayashi, E., Navarro, J. F., and Springel, V. (2007). The shape of the gravitational potential in cold dark matter haloes. MNRAS, 377(1):50–62.
Hayashi, E., Navarro, J. F., Taylor, J. E., Stadel, J., and Quinn, T. (2003). The Structural Evolution of Substructure. ApJ, 584(2):541–558.
Hellwing, W. A., Frenk, C. S., Cautun, M., Bose, S., Helly, J., Jenkins, A., Sawala, T., and Cytowski, M. (2016). The Copernicus Complexio: a high-resolution view of the small-scale Universe. MNRAS, 457(4):3492–3509.
Hernquist, L. (1990). An Analytical Model for Spherical Galaxies and Bulges. *Astrophys. J.*, 356:359.
Hernquist, L. (1993). N-Body Realizations of Compound Galaxies. *Astrophysical Journal Supplement*, 86:389.
Heyvaerts, J. (2010). A Balescu-Lenard-type kinetic equation for the collisional evolution of stable self-gravitating systems. MNRAS, 407(1):355–372.
Hopkins, P. F., Quataert, E., and Murray, N. (2012). The structure of the interstellar medium of star-forming galaxies. MNRAS, 421(4):3488–3521.
Hopkins, P. F., Wetzel, A., Kereš, D., Faucher-Giguère, C.-A., Quataert, E., Boylan-Kolchin, M., Murray, N., Hayward, C. C., Garrison-Kimmel, S., Hummels, C., Feldmann, R., Torrey, P., Ma, X., Anglés-Alcázar, D., Su, K.-Y., Orr, M., Schmitz, D., Escala, I., Sanderson, R., Grudić, M. Y., Hafen, Z., Kim, J.-H., Fitts, A., Bullock, J. S., Wheeler, C., Chan, T. K., Elbert, O. D., and Narayanan, D. (2018). FIRE-2 simulations: physics versus numerics in galaxy formation. MNRAS, 480(1):800–863.
Hui, L., Ostriker, J. P., Tremaine, S., and Witten, E. (2017). Ultralight scalars as cosmological dark matter. Phys. Rev. D, 95(4):043541.
Ibata, R. A., Lewis, G. F., Conn, A. R., Irwin, M. J., McConnachie, A. W., Chapman, S. C., Collins, M. L., Fardal, M., Ferguson, A. M. N., Ibata, N. G., Mackey, A. D., Martin, N. F., Navarro, J., Rich, R. M., Valls-Gabaud, D., and Widrow, L. M. (2013).

A vast, thin plane of corotating dwarf galaxies orbiting the Andromeda galaxy. Nature, 493(7430):62–65.

Ishiyama, T. (2014). Hierarchical Formation of Dark Matter Halos and the Free Streaming Scale. ApJ, 788(1):27.

Ishiyama, T., Prada, F., Klypin, A. A., Sinha, M., Metcalf, R. B., Jullo, E., Altieri, B., Cora, S. A., Croton, D., de la Torre, S., Millán-Calero, D. E., Oogi, T., Ruedas, J., and Vega-Martínez, C. A. (2020). The Uchuu Simulations: Data Release 1 and Dark Matter Halo Concentrations. *arXiv e-prints*, page arXiv:2007.14720.

Jaynes, E. T. (1957). Information Theory and Statistical Mechanics. *Physical Review*, 106(4):620–630.

Jenkins, A. (2010). Second-order Lagrangian perturbation theory initial conditions for resimulations. MNRAS, 403(4):1859–1872.

Jenkins, A., Frenk, C. S., White, S. D. M., Colberg, J. M., Cole, S., Evrard, A. E., Couchman, H. M. P., and Yoshida, N. (2001). The mass function of dark matter haloes. MNRAS, 321(2):372–384.

Jing, Y. P. and Suto, Y. (2002). Triaxial Modeling of Halo Density Profiles with High-Resolution N-Body Simulations. ApJ, 574(2):538–553.

Kahlhoefer, F., Kaplinghat, M., Slatyer, T. R., and Wu, C.-L. (2019). Diversity in density profiles of self-interacting dark matter satellite halos. J. Cosmology Astropart. Phys., 2019(12):010.

Kahlhoefer, F., Schmidt-Hoberg, K., Frandsen, M. T., and Sarkar, S. (2014). Colliding clusters and dark matter self-interactions. MNRAS, 437(3):2865–2881.

Kaplinghat, M., Tulin, S., and Yu, H.-B. (2016). Dark Matter Halos as Particle Colliders: Unified Solution to Small-Scale Structure Puzzles from Dwarfs to Clusters. Phys. Rev. Lett., 116(4):041302.

Kauffmann, G. (2014). Quantitative constraints on starburst cycles in galaxies with stellar masses in the range 10^8-10^{10} $M_\odot$. MNRAS, 441(3):2717–2724.

Kauffmann, G., Colberg, J. M., Diaferio, A., and White, S. D. M. (1999). Clustering of galaxies in a hierarchical universe - I. Methods and results at z=0. MNRAS, 303(1):188–206.

Kauffmann, G., Nusser, A., and Steinmetz, M. (1997). Galaxy formation and large-scale bias. MNRAS, 286(4):795–811.

Kauffmann, G. and White, S. D. M. (1993). The merging history of dark matter haloes in a hierarchical universe. MNRAS, 261:921–928.

Kauffmann, G., White, S. D. M., and Guiderdoni, B. (1993). The formation and evolution of galaxies within merging dark matter haloes. MNRAS, 264:201–218.

Kazantzidis, S., Magorrian, J., and Moore, B. (2004). Generating equilibrium dark matter halos: Inadequacies of the local Maxwellian approximation. *Astrophys. J.*, 601:37–46.

Kazantzidis, S., Mayer, L., Mastropietro, C., Diemand, J., Stadel, J., and Moore, B. (2004). Density Profiles of Cold Dark Matter Substructure: Implications for the Missing-Satellites Problem. ApJ, 608(2):663–679.

Keller, B. W. and Wadsley, J. W. (2017). ΛCDM is Consistent with SPARC Radial Acceleration Relation. ApJ, 835(1):L17.

Khan, R., Williams, B. F., Dalcanton, J., and WFIRST Infrared Nearby Galaxies Survey (WINGS) (2018). WFIRST: Simulating and Analyzing Wide Field, High-Resolution

Images of Nearby Galaxies. In *American Astronomical Society Meeting Abstracts #231*, volume 231 of *American Astronomical Society Meeting Abstracts*, page 150.40.

Kirby, E. N., Bullock, J. S., Boylan-Kolchin, M., Kaplinghat, M., and Cohen, J. G. (2014). The dynamics of isolated Local Group galaxies. MNRAS, 439(1):1015–1027.

Klypin, A., Yepes, G., Gottlöber, S., Prada, F., and Heß, S. (2016). MultiDark simulations: the story of dark matter halo concentrations and density profiles. MNRAS, 457(4):4340–4359.

Klypin, A. A. and Shandarin, S. F. (1983). Three-dimensional numerical model of the formation of large-scale structure in the Universe. MNRAS, 204:891–907.

Klypin, A. A., Trujillo-Gomez, S., and Primack, J. (2011). Dark Matter Halos in the Standard Cosmological Model: Results from the Bolshoi Simulation. ApJ, 740(2):102.

Knollmann, S. R. and Knebe, A. (2009). AHF: Amiga's Halo Finder. ApJS, 182(2):608–624.

Koda, J. and Shapiro, P. R. (2011). Gravothermal collapse of isolated self-interacting dark matter haloes: N-body simulation versus the fluid model. MNRAS, 415(2):1125–1137.

Kroupa, P. (2002). The Initial Mass Function of Stars: Evidence for Uniformity in Variable Systems. *Science*, 295(5552):82–91.

Krumholz, M. R. and Tan, J. C. (2007). Slow Star Formation in Dense Gas: Evidence and Implications. *Astrophysical Journal*, 654(1):304–315.

Kunkel, W. E. and Demers, S. (1976). The Magellanic Plane. In *The Galaxy and the Local Group*, volume 182, page 241.

Kuzio de Naray, R., McGaugh, S. S., and de Blok, W. J. G. (2008). Mass Models for Low Surface Brightness Galaxies with High-Resolution Optical Velocity Fields. ApJ, 676(2):920–943.

Lacey, C. and Cole, S. (1993). Merger rates in hierarchical models of galaxy formation. MNRAS, 262(3):627–649.

Laporte, C. F. P. and Peñarrubia, J. (2015). Under the sword of Damocles: plausible regeneration of dark matter cusps at the smallest galactic scales. MNRAS, 449:L90–L94.

Lazar, A., Bullock, J. S., Boylan-Kolchin, M., Chan, T. K., Hopkins, P. F., Graus, A. S., Wetzel, A., El-Badry, K., Wheeler, C., Straight, M. C., Kereš, D., Faucher-Giguère, C.-A., Fitts, A., and Garrison-Kimmel, S. (2020). A dark matter profile to model diverse feedback-induced core sizes of ΛCDM haloes. MNRAS, 497(2):2393–2417.

Lehmann, E. L. and Casella, G. (1998). *Theory of Point Estimation*. Springer-Verlag, New York, NY, USA, second edition.

Lenard, A. (1960). On Bogoliubov's kinetic equation for a spatially homogeneous plasma. *Annals of Physics*, 10(3):390–400.

Levin, Y., Pakter, R., and Rizzato, F. B. (2008). Collisionless relaxation in gravitational systems: From violent relaxation to gravothermal collapse. Phys. Rev. E, 78(2):021130.

Levin, Y., Pakter, R., Rizzato, F. B., Teles, T. N., and Benetti, F. P. C. (2014). Nonequilibrium statistical mechanics of systems with long-range interactions. Phys. Rep., 535(1):1–60.

Linde, A. D. (1982). A new inflationary universe scenario: A possible solution of the

horizon, flatness, homogeneity, isotropy and primordial monopole problems. *Physics Letters B*, 108(6):389–393.

Lovell, M. R., Frenk, C. S., Eke, V. R., Jenkins, A., Gao, L., and Theuns, T. (2014). The properties of warm dark matter haloes. MNRAS, 439(1):300–317.

Lubimov, V. A., Novikov, E. G., Nozik, V. Z., Tretyakov, E. F., and Kosik, V. S. (1980). An estimate of the ν_e mass from the β-spectrum of tritium in the valine molecule. *Physics Letters B*, 94(2):266–268.

Ludlow, A. D., Navarro, J. F., Angulo, R. E., Boylan-Kolchin, M., Springel, V., Frenk, C., and White, S. D. M. (2014). The mass-concentration-redshift relation of cold dark matter haloes. MNRAS, 441(1):378–388.

Ludlow, A. D., Navarro, J. F., White, S. D. M., Boylan-Kolchin, M., Springel, V., Jenkins, A., and Frenk, C. S. (2011). The density and pseudo-phase-space density profiles of cold dark matter haloes. MNRAS, 415(4):3895–3902.

Lukić, Z., Heitmann, K., Habib, S., Bashinsky, S., and Ricker, P. M. (2007). The Halo Mass Function: High-Redshift Evolution and Universality. ApJ, 671(2):1160–1181.

Lynden-Bell, D. (1967). Statistical mechanics of violent relaxation in stellar systems. MNRAS, 136:101.

Lynden-Bell, D. (1976). Dwarf galaxies and globular clusters in high velocity hydrogen streams. MNRAS, 174:695–710.

Lynden-Bell, D. (1982). On the N-body problem in Dirac's cosmology. *The Observatory*, 102:86–87.

Lynden-Bell, D. (1999). Negative Specific Heat in Astronomy, Physics and Chemistry. *Physica A Statistical Mechanics and its Applications*, 263(1-4):293–304.

Lynden-Bell, D. and Kalnajs, A. J. (1972). On the generating mechanism of spiral structure. MNRAS, 157:1.

Lynden-Bell, D. and Lynden-Bell, R. M. (1977). On the negative specific heat paradox. MNRAS, 181:405–419.

Lynden-Bell, D. and Wood, R. (1968). The gravo-thermal catastrophe in isothermal spheres and the onset of red-giant structure for stellar systems. MNRAS, 138:495.

Macciò, A. V., Udrescu, S. M., Dutton, A. A., Obreja, A., Wang, L., Stinson, G. R., and Kang, X. (2016). NIHAO X: reconciling the local galaxy velocity function with cold dark matter via mock H I observations. MNRAS, 463(1):L69–L73.

Makino, J. (1991). Optimal Order and Time-Step Criterion for Aarseth-Type N-Body Integrators. ApJ, 369:200.

Marinacci, F., Sales, L. V., Vogelsberger, M., Torrey, P., and Springel, V. (2019). Simulating the interstellar medium and stellar feedback on a moving mesh: implementation and isolated galaxies. MNRAS, 489(3):4233–4260.

Martin, C. L., Shapley, A. E., Coil, A. L., Kornei, K. A., Bundy, K., Weiner, B. J., Noeske, K. G., and Schiminovich, D. (2012). Demographics and Physical Properties of Gas Outflows/Inflows at $0.4 < z < 1.4$. ApJ, 760:127.

McConnachie, A. W. (2012). The Observed Properties of Dwarf Galaxies in and around the Local Group. AJ, 144(1):4.

McGaugh, S. S. (2015). A tale of two paradigms: the mutual incommensurability of ΛCDM and MOND. *Canadian Journal of Physics*, 93(2):250–259.

McGaugh, S. S., Lelli, F., and Schombert, J. M. (2016). Radial Acceleration Relation in Rotationally Supported Galaxies. Phys. Rev. Lett., 117(20):201101.

Meiron, Y. and Kocsis, B. (2018). Diffusion and Mixing in Globular Clusters. ApJ, 855:87.

Melott, A. L. (1983). Massive neutrinos in large-scale gravitational clustering. ApJ, 264:59–78.

Meszaros, P. (1974). The behaviour of point masses in an expanding cosmological substratum. A&A, 37(2):225–228.

Monaghan, J. J. and Lattanzio, J. C. (1985). A refined particle method for astrophysical problems. A&A, 149(1):135–143.

Moore, B. (1994). Evidence against dissipation-less dark matter from observations of galaxy haloes. Nature, 370(6491):629–631.

Moster, B. P., Somerville, R. S., Maulbetsch, C., van den Bosch, F. C., Macciò, A. V., Naab, T., and Oser, L. (2010). Constraints on the Relationship between Stellar Mass and Halo Mass at Low and High Redshift. ApJ, 710(2):903–923.

Nakamura, T. K. (2000). Statistical Mechanics of a Collisionless System Based on the Maximum Entropy Principle. ApJ, 531(2):739–743.

Navarro, J. F., Benítez-Llambay, A., Fattahi, A., Frenk, C. S., Ludlow, A. D., Oman, K. A., Schaller, M., and Theuns, T. (2017). The origin of the mass discrepancy-acceleration relation in ΛCDM. MNRAS, 471(2):1841–1848.

Navarro, J. F., Eke, V. R., and Frenk, C. S. (1996a). The cores of dwarf galaxy haloes. MNRAS, 283(3):L72–L78.

Navarro, J. F., Frenk, C. S., and White, S. D. M. (1996b). The Structure of Cold Dark Matter Halos. ApJ, 462:563.

Navarro, J. F., Frenk, C. S., and White, S. D. M. (1997). A Universal Density Profile from Hierarchical Clustering. ApJ, 490(2):493–508.

Navarro, J. F., Ludlow, A., Springel, V., Wang, J., Vogelsberger, M., White, S. D. M., Jenkins, A., Frenk, C. S., and Helmi, A. (2010). The diversity and similarity of simulated cold dark matter haloes. MNRAS, 402(1):21–34.

Neto, A. F., Gao, L., Bett, P., Cole, S., Navarro, J. F., Frenk, C. S., White, S. D. M., Springel, V., and Jenkins, A. (2007). The statistics of Λ CDM halo concentrations. MNRAS, 381(4):1450–1462.

Nishikawa, H., Boddy, K. K., and Kaplinghat, M. (2020). Accelerated core collapse in tidally stripped self-interacting dark matter halos. Phys. Rev. D, 101(6):063009.

Oh, S.-H., de Blok, W. J. G., Brinks, E., Walter, F., and Kennicutt, Robert C., J. (2011). Dark and Luminous Matter in THINGS Dwarf Galaxies. *Astronomical Journal*, 141(6):193.

Oh, S.-H., Hunter, D. A., Brinks, E., Elmegreen, B. G., Schruba, A., Walter, F., Rupen, M. P., Young, L. M., Simpson, C. E., Johnson, M. C., Herrmann, K. A., Ficut-Vicas, D., Cigan, P., Heesen, V., Ashley, T., and Zhang, H.-X. (2015). High-resolution Mass Models of Dwarf Galaxies from LITTLE THINGS. AJ, 149(6):180.

Oman, K. A., Marasco, A., Navarro, J. F., Frenk, C. S., Schaye, J., and Benítez-Llambay, A. (2019). Non-circular motions and the diversity of dwarf galaxy rotation curves. MNRAS, 482(1):821–847.

Oman, K. A., Navarro, J. F., Fattahi, A., Frenk, C. S., Sawala, T., White, S. D. M., Bower, R., Crain, R. A., Furlong, M., Schaller, M., Schaye, J., and Theuns, T. (2015). The unexpected diversity of dwarf galaxy rotation curves. MNRAS, 452(4):3650–3665.

Oman, K. A., Navarro, J. F., Sales, L. V., Fattahi, A., Frenk, C. S., Sawala, T., Schaller,

M., and White, S. D. M. (2016a). Missing dark matter in dwarf galaxies? MNRAS, 460(4):3610–3623.

Oman, K. A., Navarro, J. F., Sales, L. V., Fattahi, A., Frenk, C. S., Sawala, T., Schaller, M., and White, S. D. M. (2016b). Missing dark matter in dwarf galaxies? MNRAS, 460(4):3610–3623.

Orkney, M. D. A., Read, J. I., Rey, M. P., Nasim, I., Pontzen, A., Agertz, O., Kim, S. Y., Delorme, M., and Dehnen, W. (2021). EDGE: Two routes to dark matter core formation in ultra-faint dwarfs. *arXiv e-prints*, page arXiv:2101.02688.

Ostriker, J. P. and Peebles, P. J. E. (1973). A Numerical Study of the Stability of Flattened Galaxies: or, can Cold Galaxies Survive? ApJ, 186:467–480.

Ostriker, J. P., Peebles, P. J. E., and Yahil, A. (1974). The Size and Mass of Galaxies, and the Mass of the Universe. ApJ, 193:L1.

Padmanabhan, T. (1989). A Note on the Volume Occupied in Phase Space by a System of Gravitationally Interacting Particles. ApJ, 344:848.

Padmanabhan, T. (1990). Statistical mechanics of gravitating systems. Phys. Rep., 188(5):285–362.

Papastergis, E., Giovanelli, R., Haynes, M. P., and Shankar, F. (2015). Is there a "too big to fail" problem in the field? A&A, 574:A113.

Papastergis, E. and Shankar, F. (2016). An assessment of the "too big to fail" problem for field dwarf galaxies in view of baryonic feedback effects. A&A, 591:A58.

Pawlowski, M. S., Famaey, B., Merritt, D., and Kroupa, P. (2015). On the Persistence of Two Small-scale Problems in ΛCDM. ApJ, 815(1):19.

Pawlowski, M. S. and Kroupa, P. (2013). The rotationally stabilized VPOS and predicted proper motions of the Milky Way satellite galaxies. MNRAS, 435(3):2116–2131.

Pawlowski, M. S., Pflamm-Altenburg, J., and Kroupa, P. (2012). The VPOS: a vast polar structure of satellite galaxies, globular clusters and streams around the Milky Way. MNRAS, 423(2):1109–1126.

Peñarrubia, J. (2013). Dynamical invariants and diffusion of merger substructures in time-dependent gravitational potentials. MNRAS, 433(3):2576–2587.

Peñarrubia, J. (2015). A probability theory for non-equilibrium gravitational systems. MNRAS, 451(4):3537–3550.

Peñarrubia, J. (2019). Stochastic tidal heating by random interactions with extended substructures. MNRAS, 484(4):5409–5436.

Peñarrubia, J., Pontzen, A., Walker, M. G., and Koposov, S. E. (2012). The Coupling between the Core/Cusp and Missing Satellite Problems. ApJ, 759(2):L42.

Peacock, J. A. and Smith, R. E. (2000). Halo occupation numbers and galaxy bias. MNRAS, 318(4):1144–1156.

Peebles, P. J. E. (1993). *Principles of Physical Cosmology*.

Perlmutter, S., Aldering, G., Goldhaber, G., Knop, R. A., Nugent, P., Castro, P. G., Deustua, S., Fabbro, S., Goobar, A., Groom, D. E., Hook, I. M., Kim, A. G., Kim, M. Y., Lee, J. C., Nunes, N. J., Pain, R., Pennypacker, C. R., Quimby, R., Lidman, C., Ellis, R. S., Irwin, M., McMahon, R. G., Ruiz-Lapuente, P., Walton, N., Schaefer, B., Boyle, B. J., Filippenko, A. V., Matheson, T., Fruchter, A. S., Panagia, N., Newberg, H. J. M., Couch, W. J., and Project, T. S. C. (1999). Measurements of Ω and Λ from 42 High-Redshift Supernovae. ApJ, 517(2):565–586.

Peter, A. H. G., Rocha, M., Bullock, J. S., and Kaplinghat, M. (2013). Cosmological

simulations with self-interacting dark matter - II. Halo shapes versus observations. MNRAS, 430(1):105–120.

Phillips, J. I., Cooper, M. C., Bullock, J. S., and Boylan-Kolchin, M. (2015). Are rotating planes of satellite galaxies ubiquitous? MNRAS, 453(4):3839–3847.

Pilipenko, S. V., Sánchez-Conde, M. A., Prada, F., and Yepes, G. (2017). Pushing down the low-mass halo concentration frontier with the Lomonosov cosmological simulations. MNRAS, 472(4):4918–4927.

Planck Collaboration, Aghanim, N., Akrami, Y., Ashdown, M., Aumont, J., Baccigalupi, C., Ballardini, M., Banday, A. J., Barreiro, R. B., Bartolo, N., Basak, S., Battye, R., Benabed, K., Bernard, J. P., Bersanelli, M., Bielewicz, P., Bock, J. J., Bond, J. R., Borrill, J., Bouchet, F. R., Boulanger, F., Bucher, M., Burigana, C., Butler, R. C., Calabrese, E., Cardoso, J. F., Carron, J., Challinor, A., Chiang, H. C., Chluba, J., Colombo, L. P. L., Combet, C., Contreras, D., Crill, B. P., Cuttaia, F., de Bernardis, P., de Zotti, G., Delabrouille, J., Delouis, J. M., Di Valentino, E., Diego, J. M., Doré, O., Douspis, M., Ducout, A., Dupac, X., Dusini, S., Efstathiou, G., Elsner, F., Enßlin, T. A., Eriksen, H. K., Fantaye, Y., Farhang, M., Fergusson, J., Fernandez-Cobos, R., Finelli, F., Forastieri, F., Frailis, M., Fraisse, A. A., Franceschi, E., Frolov, A., Galeotta, S., Galli, S., Ganga, K., Génova-Santos, R. T., Gerbino, M., Ghosh, T., González-Nuevo, J., Górski, K. M., Gratton, S., Gruppuso, A., Gudmundsson, J. E., Hamann, J., Handley, W., Hansen, F. K., Herranz, D., Hildebrandt, S. R., Hivon, E., Huang, Z., Jaffe, A. H., Jones, W. C., Karakci, A., Keihänen, E., Keskitalo, R., Kiiveri, K., Kim, J., Kisner, T. S., Knox, L., Krachmalnicoff, N., Kunz, M., Kurki-Suonio, H., Lagache, G., Lamarre, J. M., Lasenby, A., Lattanzi, M., Lawrence, C. R., Le Jeune, M., Lemos, P., Lesgourgues, J., Levrier, F., Lewis, A., Liguori, M., Lilje, P. B., Lilley, M., Lindholm, V., López-Caniego, M., Lubin, P. M., Ma, Y. Z., Macías-Pérez, J. F., Maggio, G., Maino, D., Mandolesi, N., Mangilli, A., Marcos-Caballero, A., Maris, M., Martin, P. G., Martinelli, M., Martínez-González, E., Matarrese, S., Mauri, N., McEwen, J. D., Meinhold, P. R., Melchiorri, A., Mennella, A., Migliaccio, M., Millea, M., Mitra, S., Miville-Deschênes, M. A., Molinari, D., Montier, L., Morgante, G., Moss, A., Natoli, P., Nørgaard-Nielsen, H. U., Pagano, L., Paoletti, D., Partridge, B., Patanchon, G., Peiris, H. V., Perrotta, F., Pettorino, V., Piacentini, F., Polastri, L., Polenta, G., Puget, J. L., Rachen, J. P., Reinecke, M., Remazeilles, M., Renzi, A., Rocha, G., Rosset, C., Roudier, G., Rubiño-Martín, J. A., Ruiz-Granados, B., Salvati, L., Sandri, M., Savelainen, M., Scott, D., Shellard, E. P. S., Sirignano, C., Sirri, G., Spencer, L. D., Sunyaev, R., Suur-Uski, A. S., Tauber, J. A., Tavagnacco, D., Tenti, M., Toffolatti, L., Tomasi, M., Trombetti, T., Valenziano, L., Valiviita, J., Van Tent, B., Vibert, L., Vielva, P., Villa, F., Vittorio, N., Wandelt, B. D., Wehus, I. K., White, M., White, S. D. M., Zacchei, A., and Zonca, A. (2020). Planck 2018 results. VI. Cosmological parameters. A&A, 641:A6.

Plummer, H. C. (1911). On the problem of distribution in globular star clusters. MNRAS, 71:460–470.

Polisensky, E. and Ricotti, M. (2011). Constraints on the dark matter particle mass from the number of Milky Way satellites. Phys. Rev. D, 83(4):043506.

Pollack, J., Spergel, D. N., and Steinhardt, P. J. (2015). Supermassive Black Holes from Ultra-strongly Self-interacting Dark Matter. ApJ, 804(2):131.

Pontzen, A. and Governato, F. (2012). How supernova feedback turns dark matter cusps

into cores. MNRAS, 421(4):3464–3471.
Pontzen, A. and Governato, F. (2013). Conserved actions, maximum entropy and dark matter haloes. MNRAS, 430(1):121–133.
Pontzen, A., Read, J. I., Teyssier, R., Governato, F., Gualandris, A., Roth, N., and Devriendt, J. (2015). Milking the spherical cow - on aspherical dynamics in spherical coordinates. MNRAS, 451(2):1366–1379.
Potter, D., Stadel, J., and Teyssier, R. (2017). PKDGRAV3: beyond trillion particle cosmological simulations for the next era of galaxy surveys. *Computational Astrophysics and Cosmology*, 4(1):2.
Power, C., Navarro, J. F., Jenkins, A., Frenk, C. S., White, S. D. M., Springel, V., Stadel, J., and Quinn, T. (2003). The inner structure of ΛCDM haloes - I. A numerical convergence study. MNRAS, 338(1):14–34.
Prada, F., Klypin, A. A., Cuesta, A. J., Betancort-Rijo, J. E., and Primack, J. (2012). Halo concentrations in the standard Λ cold dark matter cosmology. MNRAS, 423(4):3018–3030.
Preskill, J., Wise, M. B., and Wilczek, F. (1983). Cosmology of the invisible axion. *Physics Letters B*, 120(1-3):127–132.
Press, W. H. and Schechter, P. (1974). Formation of Galaxies and Clusters of Galaxies by Self-Similar Gravitational Condensation. ApJ, 187:425–438.
Price-Whelan, A. M., Hogg, D. W., Johnston, K. V., Ness, M. K., Rix, H.-W., Beaton, R. L., Brownstein, J. R., García-Hernández, D. A., Hasselquist, S., Hayes, C. R., Lane, R. R., Shetrone, M., Sobeck, J., and Zasowski, G. (2021). Orbital Torus Imaging: Using Element Abundances to Map Orbits and Mass in the Milky Way. *Astrophysical Journal*, 910(1):17.
Quinlan, G. D. and Tremaine, S. (1990). Symmetric Multistep Methods for the Numerical Integration of Planetary Orbits. AJ, 100:1694.
Randall, S. W., Markevitch, M., Clowe, D., Gonzalez, A. H., and Bradač, M. (2008). Constraints on the Self-Interaction Cross Section of Dark Matter from Numerical Simulations of the Merging Galaxy Cluster 1E 0657-56. ApJ, 679(2):1173–1180.
Read, J. I., Agertz, O., and Collins, M. L. M. (2016). Dark matter cores all the way down. MNRAS, 459(3):2573–2590.
Read, J. I., Goerdt, T., Moore, B., Pontzen, A. P., Stadel, J., and Lake, G. (2006). Dynamical friction in constant density cores: a failure of the Chandrasekhar formula. MNRAS, 373(4):1451–1460.
Read, J. I., Walker, M. G., and Steger, P. (2018). The case for a cold dark matter cusp in Draco. MNRAS, 481(1):860–877.
Read, J. I., Walker, M. G., and Steger, P. (2019). Dark matter heats up in dwarf galaxies. MNRAS, 484(1):1401–1420.
Rees, M. J. and Ostriker, J. P. (1977). Cooling, dynamics and fragmentation of massive gas clouds: clues to the masses and radii of galaxies and clusters. MNRAS, 179:541–559.
Riess, A. G., Filippenko, A. V., Challis, P., Clocchiatti, A., Diercks, A., Garnavich, P. M., Gilliland, R. L., Hogan, C. J., Jha, S., Kirshner, R. P., Leibundgut, B., Phillips, M. M., Reiss, D., Schmidt, B. P., Schommer, R. A., Smith, R. C., Spyromilio, J., Stubbs, C., Suntzeff, N. B., and Tonry, J. (1998). Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant. AJ, 116(3):1009–1038.

Robertson, A., Harvey, D., Massey, R., Eke, V., McCarthy, I. G., Jauzac, M., Li, B., and Schaye, J. (2019). Observable tests of self-interacting dark matter in galaxy clusters: cosmological simulations with SIDM and baryons. MNRAS, 488(3):3646–3662.

Robertson, A., Massey, R., and Eke, V. (2017). What does the Bullet Cluster tell us about self-interacting dark matter? MNRAS, 465(1):569–587.

Robles, V. H., Bullock, J. S., Elbert, O. D., Fitts, A., González-Samaniego, A., Boylan-Kolchin, M., Hopkins, P. F., Faucher-Giguère, C.-A., Kereš, D., and Hayward, C. C. (2017). SIDM on FIRE: hydrodynamical self-interacting dark matter simulations of low-mass dwarf galaxies. MNRAS, 472(3):2945–2954.

Rocha, M., Peter, A. H. G., Bullock, J. S., Kaplinghat, M., Garrison-Kimmel, S., Oñorbe, J., and Moustakas, L. A. (2013). Cosmological simulations with self-interacting dark matter - I. Constant-density cores and substructure. MNRAS, 430(1):81–104.

Rodríguez-Puebla, A., Behroozi, P., Primack, J., Klypin, A., Lee, C., and Hellinger, D. (2016). Halo and subhalo demographics with Planck cosmological parameters: Bolshoi-Planck and MultiDark-Planck simulations. MNRAS, 462(1):893–916.

Rosdahl, J., Schaye, J., Teyssier, R., and Agertz, O. (2015). Galaxies that shine: radiation-hydrodynamical simulations of disc galaxies. MNRAS, 451(1):34–58.

Rubin, V. C., Burstein, D., Ford, W. K., J., and Thonnard, N. (1985). Rotation velocities of 16 SA galaxies and a comparison of Sa, SB and SC rotation properties. ApJ, 289:81–104.

Rubin, V. C., Ford, W. K., J., and Thonnard, N. (1980). Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 (R=4kpc) to UGC 2885 (R=122kpc). ApJ, 238:471–487.

Rusakov, V., Monelli, M., Gallart, C., Fritz, T. K., Ruiz-Lara, T., Bernard, E. J., and Cassisi, S. (2021). The bursty star formation history of the Fornax dwarf spheroidal galaxy revealed with the HST. *Monthly Notices of the RAS*, 502(1):642–661.

Saha, P. and Tremaine, S. (1992). Symplectic Integrators for Solar System Dynamics. AJ, 104:1633.

Sameie, O., Yu, H.-B., Sales, L. V., Vogelsberger, M., and Zavala, J. (2020). Self-Interacting Dark Matter Subhalos in the Milky Way's Tides. Phys. Rev. Lett., 124(14):141102.

Sánchez-Conde, M. A. and Prada, F. (2014). The flattening of the concentration-mass relation towards low halo masses and its implications for the annihilation signal boost. MNRAS, 442(3):2271–2277.

Sanderson, R. E., Hartke, J., and Helmi, A. (2017). Modeling the Gravitational Potential of a Cosmological Dark Matter Halo with Stellar Streams. ApJ, 836(2):234.

Sanderson, R. E., Helmi, A., and Hogg, D. W. (2015). Action-space Clustering of Tidal Streams to Infer the Galactic Potential. ApJ, 801(2):98.

Santos-Santos, I. M. E., Navarro, J. F., Robertson, A., Benítez-Llambay, A., Oman, K. A., Lovell, M. R., Frenk, C. S., Ludlow, A. D., Fattahi, A., and Ritz, A. (2020). Baryonic clues to the puzzling diversity of dwarf galaxy rotation curves. MNRAS, 495(1):58–77.

Sawala, T., Frenk, C. S., Fattahi, A., Navarro, J. F., Bower, R. G., Crain, R. A., Dalla Vecchia, C., Furlong, M., Helly, J. C., Jenkins, A., Oman, K. A., Schaller, M., Schaye, J., Theuns, T., Trayford, J., and White, S. D. M. (2016). The APOSTLE simulations: solutions to the Local Group's cosmic puzzles. MNRAS, 457(2):1931–1943.

Schneider, A., Smith, R. E., Macciò, A. V., and Moore, B. (2012). Non-linear evolution of cosmological structures in warm dark matter models. MNRAS, 424(1):684–698.

Schneider, A., Smith, R. E., and Reed, D. (2013). Halo mass function and the free streaming scale. MNRAS, 433(2):1573–1587.

Seljak, U. (2000). Analytic model for galaxy and dark matter clustering. MNRAS, 318(1):203–213.

Sheth, R. K., Mo, H. J., and Tormen, G. (2001). Ellipsoidal collapse and an improved model for the number and spatial distribution of dark matter haloes. MNRAS, 323(1):1–12.

Silk, J. (1968). Cosmic Black-Body Radiation and Galaxy Formation. ApJ, 151:459.

Sirko, E. (2005). Initial Conditions to Cosmological N-Body Simulations, or, How to Run an Ensemble of Simulations. ApJ, 634(2):728–743.

Smith, M. C., Sijacki, D., and Shen, S. (2018). Supernova feedback in numerical simulations of galaxy formation: separating physics from numerics. MNRAS, 478(1):302–331.

Smith, R. E., Peacock, J. A., Jenkins, A., White, S. D. M., Frenk, C. S., Pearce, F. R., Thomas, P. A., Efstathiou, G., and Couchman, H. M. P. (2003). Stable clustering, the halo model and non-linear cosmological power spectra. MNRAS, 341(4):1311–1332.

Spergel, D. N. and Steinhardt, P. J. (2000). Observational Evidence for Self-Interacting Cold Dark Matter. Phys. Rev. Lett., 84(17):3760–3763.

Spitzer, Lyman, J. (1958). Disruption of Galactic Clusters. ApJ, 127:17.

Springel, V. (2005). The cosmological simulation code GADGET-2. MNRAS, 364(4):1105–1134.

Springel, V. (2010). E pur si muove: Galilean-invariant cosmological hydrodynamical simulations on a moving mesh. MNRAS, 401(2):791–851.

Springel, V. (2016). High Performance Computing and Numerical Modelling. *Saas-Fee Advanced Course*, 43:251.

Springel, V., Di Matteo, T., and Hernquist, L. (2005a). Modelling feedback from stars and black holes in galaxy mergers. MNRAS, 361(3):776–794.

Springel, V., Frenk, C. S., and White, S. D. M. (2006). The large-scale structure of the Universe. Nature, 440(7088):1137–1144.

Springel, V. and Hernquist, L. (2003). Cosmological smoothed particle hydrodynamics simulations: a hybrid multiphase model for star formation. MNRAS, 339(2):289–311.

Springel, V., Wang, J., Vogelsberger, M., Ludlow, A., Jenkins, A., Helmi, A., Navarro, J. F., Frenk, C. S., and White, S. D. M. (2008). The Aquarius Project: the subhaloes of galactic haloes. MNRAS, 391(4):1685–1711.

Springel, V., White, S. D. M., Jenkins, A., Frenk, C. S., Yoshida, N., Gao, L., Navarro, J., Thacker, R., Croton, D., Helly, J., Peacock, J. A., Cole, S., Thomas, P., Couchman, H., Evrard, A., Colberg, J., and Pearce, F. (2005b). Simulations of the formation, evolution and clustering of galaxies and quasars. Nature, 435(7042):629–636.

Springel, V., White, S. D. M., Tormen, G., and Kauffmann, G. (2001). Populating a cluster of galaxies - I. Results at [formmu2]z=0. MNRAS, 328(3):726–750.

Starobinsky, A. A. (1982). Dynamics of phase transition in the new inflationary universe scenario and generation of perturbations. *Physics Letters B*, 117(3-4):175–178.

Strigari, L. E., Frenk, C. S., and White, S. D. M. (2017). Dynamical Models for the Sculptor Dwarf Spheroidal in a ΛCDM Universe. ApJ, 838(2):123.

Szalay, A. S. and Marx, G. (1976). Neutrino rest mass from cosmology. A&A, 49(3):437–441.

Taylor, J. E. and Babul, A. (2001). The Dynamics of Sinking Satellites around Disk Galaxies: A Poor Man's Alternative to High-Resolution Numerical Simulations. ApJ, 559(2):716–735.

Taylor, J. E. and Navarro, J. F. (2001). The Phase-Space Density Profiles of Cold Dark Matter Halos. ApJ, 563(2):483–488.

Tinker, J., Kravtsov, A. V., Klypin, A., Abazajian, K., Warren, M., Yepes, G., Gottlöber, S., and Holz, D. E. (2008). Toward a Halo Mass Function for Precision Cosmology: The Limits of Universality. ApJ, 688(2):709–728.

Tollerud, E. J., Boylan-Kolchin, M., and Bullock, J. S. (2014). M31 satellite masses compared to ΛCDM subhaloes. MNRAS, 440(4):3511–3519.

Tollet, E., Macciò, A. V., Dutton, A. A., Stinson, G. S., Wang, L., Penzo, C., Gutcke, T. A., Buck, T., Kang, X., Brook, C., Di Cintio, A., Keller, B. W., and Wadsley, J. (2016). NIHAO - IV: core creation and destruction in dark matter density profiles across cosmic time. MNRAS, 456(4):3542–3552.

Tomozeiu, M., Mayer, L., and Quinn, T. (2016). Tidal Stirring of Satellites with Shallow Density Profiles Prevents Them from Being Too Big to Fail. ApJ, 827(1):L15.

Tormen, G., Diaferio, A., and Syer, D. (1998). Survival of substructure within dark matter haloes. MNRAS, 299(3):728–742.

Toro, E. F. (1997). *Riemann Solvers and Numerical Methods for Fluid Dynamics*. Springer, Berlin.

Tremaine, S. and Weinberg, M. D. (1984). Dynamical friction in spherical systems. MNRAS, 209:729–757.

Trujillo-Gomez, S., Klypin, A., Primack, J., and Romanowsky, A. J. (2011). Galaxies in ΛCDM with Halo Abundance Matching: Luminosity-Velocity Relation, Baryonic Mass-Velocity Relation, Velocity Function, and Clustering. ApJ, 742(1):16.

Tulin, S. and Yu, H.-B. (2018). Dark matter self-interactions and small scale structure. Phys. Rep., 730:1–57.

Turner, H. C., Lovell, M. R., Zavala, J., and Vogelsberger, M. (2020). The Onset of Gravothermal Core Collapse in Velocity Dependent Self-Interacting Dark Matter Subhaloes. *arXiv e-prints*, page arXiv:2010.02924.

Tyson, J. A., Kochanski, G. P., and Dell'Antonio, I. P. (1998). Detailed Mass Map of CL 0024+1654 from Strong Lensing. ApJ, 498(2):L107–L110.

Utrobin, V. P. and Chugai, N. N. (2011). Supernova 2000cb: high-energy version of sn 1987a. *A&A*, 532:A100.

van Albada, T. S., Bahcall, J. N., Begeman, K., and Sancisi, R. (1985). Distribution of dark matter in the spiral galaxy NGC 3198. ApJ, 295:305–313.

van den Bosch, F. C. and Ogiya, G. (2018). Dark matter substructure in numerical simulations: a tale of discreteness noise, runaway instabilities, and artificial disruption. MNRAS, 475(3):4066–4087.

van den Bosch, F. C., Ogiya, G., Hahn, O., and Burkert, A. (2018). Disruption of dark matter substructure: fact or fiction? MNRAS, 474(3):3043–3066.

Vega-Ferrero, J., Yepes, G., and Gottlöber, S. (2017). On the shape of dark matter haloes from MultiDark Planck simulations. MNRAS, 467(3):3226–3238.

Vera-Ciro, C. A., Sales, L. V., Helmi, A., Frenk, C. S., Navarro, J. F., Springel, V.,

Vogelsberger, M., and White, S. D. M. (2011). The shape of dark matter haloes in the Aquarius simulations: evolution and memory. MNRAS, 416(2):1377–1391.

Viel, M., Becker, G. D., Bolton, J. S., and Haehnelt, M. G. (2013). Warm dark matter as a solution to the small scale crisis: New constraints from high redshift Lyman-α forest data. Phys. Rev. D, 88(4):043502.

Vogelsberger, M., Helmi, A., Springel, V., White, S. D. M., Wang, J., Frenk, C. S., Jenkins, A., Ludlow, A., and Navarro, J. F. (2009). Phase-space structure in the local dark matter distribution and its signature in direct detection experiments. MNRAS, 395(2):797–811.

Vogelsberger, M., Marinacci, F., Torrey, P., and Puchwein, E. (2020). Cosmological simulations of galaxy formation. *Nature Reviews Physics*, 2(1):42–66.

Vogelsberger, M., Zavala, J., Cyr-Racine, F.-Y., Pfrommer, C., Bringmann, T., and Sigurdson, K. (2016). ETHOS - an effective theory of structure formation: dark matter physics as a possible explanation of the small-scale CDM problems. MNRAS, 460(2):1399–1416.

Vogelsberger, M., Zavala, J., and Loeb, A. (2012). Subhaloes in self-interacting galactic dark matter haloes. MNRAS, 423(4):3740–3752.

Vogelsberger, M., Zavala, J., Simpson, C., and Jenkins, A. (2014). Dwarf galaxies in CDM and SIDM with baryons: observational probes of the nature of dark matter. MNRAS, 444:3684–3698.

von Neumann, J. (1951). Various techniques used in connection with random digits. *National Bureau of Standards Applied Mathematics Series*, 12:36–38.

Walker, M. G., Mateo, M., and Olszewski, E. W. (2009). Stellar Velocities in the Carina, Fornax, Sculptor, and Sextans dSph Galaxies: Data From the Magellan/MMFS Survey. AJ, 137(2):3100–3108.

Walker, M. G. and Peñarrubia, J. (2011). A Method for Measuring (Slopes of) the Mass Profiles of Dwarf Spheroidal Galaxies. ApJ, 742(1):20.

Wang, J., Bose, S., Frenk, C. S., Gao, L., Jenkins, A., Springel, V., and White, S. D. M. (2020). Universal structure of dark matter haloes over a mass range of 20 orders of magnitude. Nature, 585(7823):39–42.

Wang, J., Navarro, J. F., Frenk, C. S., White, S. D. M., Springel, V., Jenkins, A., Helmi, A., Ludlow, A., and Vogelsberger, M. (2011). Assembly history and structure of galactic cold dark matter haloes. MNRAS, 413(2):1373–1382.

Warren, M. S., Abazajian, K., Holz, D. E., and Teodoro, L. (2006). Precision Determination of the Mass Function of Dark Matter Halos. ApJ, 646(2):881–885.

Wechsler, R. H., Bullock, J. S., Primack, J. R., Kravtsov, A. V., and Dekel, A. (2002). Concentrations of Dark Halos from Their Assembly Histories. ApJ, 568(1):52–70.

Weisz, D. R., Dolphin, A. E., Skillman, E. D., Holtzman, J., Gilbert, K. M., Dalcanton, J. J., and Williams, B. F. (2014). The Star Formation Histories of Local Group Dwarf Galaxies. I. Hubble Space Telescope/Wide Field Planetary Camera 2 Observations. ApJ, 789(2):147.

Wetzel, A. R., Hopkins, P. F., Kim, J.-h., Faucher-Giguère, C.-A., Kereš, D., and Quataert, E. (2016). Reconciling Dwarf Galaxies with ΛCDM Cosmology: Simulating a Realistic Population of Satellites around a Milky Way-mass Galaxy. ApJ, 827(2):L23.

Wheeler, C., Hopkins, P. F., Pace, A. B., Garrison-Kimmel, S., Boylan-Kolchin, M.,

Wetzel, A., Bullock, J. S., Kereš, D., Faucher-Giguère, C.-A., and Quataert, E. (2019). Be it therefore resolved: cosmological simulations of dwarf galaxies with 30 solar mass resolution. MNRAS, 490(3):4447–4463.

White, M. (2001). The mass of a halo. A&A, 367:27–32.

White, S. D. M. (1994). Formation and Evolution of Galaxies: Les Houches Lectures. *arXiv e-prints*, pages astro–ph/9410043.

White, S. D. M. and Frenk, C. S. (1991). Galaxy Formation through Hierarchical Clustering. ApJ, 379:52.

White, S. D. M., Frenk, C. S., and Davis, M. (1983). Clustering in a neutrino-dominated universe. ApJ, 274:L1–L5.

Yoshida, N., Springel, V., White, S. D. M., and Tormen, G. (2000). Weakly Self-interacting Dark Matter and the Structure of Dark Halos. ApJ, 544(2):L87–L90.

Yuan, Z., Qian, Y. Z., and Jing, Y. P. (2016). Estimating the evolution of gas in the Fornax dwarf spheroidal galaxy from its star formation history: an illustrative example. MNRAS, 456(3):3253–3264.

Zavala, J. and Frenk, C. S. (2019). Dark Matter Haloes and Subhaloes. *Galaxies*, 7(4):81.

Zavala, J., Lovell, M. R., Vogelsberger, M., and Burger, J. D. (2019). Diverse dark matter density at sub-kiloparsec scales in Milky Way satellites: Implications for the nature of dark matter. Phys. Rev. D, 100(6):063007.

Zavala, J., Vogelsberger, M., and Walker, M. G. (2013). Constraining self-interacting dark matter with the Milky way's dwarf spheroidals. MNRAS, 431:L20–L24.

Zel'Dovich, Y. B. (1970). Reprint of 1970A&A.....5...84Z. Gravitational instability: an approximate theory for large density perturbations. A&A, 500:13–18.

Zentner, A. R., Berlind, A. A., Bullock, J. S., Kravtsov, A. V., and Wechsler, R. H. (2005). The Physics of Galaxy Clustering. I. A Model for Subhalo Populations. ApJ, 624(2):505–525.

Zentner, A. R. and Bullock, J. S. (2003). Halo Substructure and the Power Spectrum. ApJ, 598(1):49–72.

Zhao, D. H., Jing, Y. P., Mo, H. J., and Börner, G. (2009). Accurate Universal Models for the Mass Accretion Histories and Concentrations of Dark Matter Halos. ApJ, 707(1):354–369.

Zolotov, A., Brooks, A. M., Willman, B., Governato, F., Pontzen, A., Christensen, C., Dekel, A., Quinn, T., Shen, S., and Wadsley, J. (2012). Baryons Matter: Why Luminous Satellite Galaxies have Reduced Central Masses. ApJ, 761(1):71.

Zwicky, F. (1933). Die Rotverschiebung von extragalaktischen Nebeln. *Helvetica Physica Acta*, 6:110–127.