Dissertation submitted to the
Combined Faculties of the Natural Sciences and Mathematics
of the Ruperto-Carola-University of Heidelberg. Germany
for the degree of
Doctor of Natural Sciences

Construction and stability of disk galaxies, and the radial migration of their stars

Put forward by:

Denis Yurin

born in: Kazakhstan

Oral examination: 17.12.2014

Referees:

Prof. Volker Springel

Prof. Bjoern Malte Schaefer

To Julia

Contents

Abstract

This thesis studies the stability of disk galaxies and the radial migration rate of their stars in self-consistent cosmological models of the formation of Milky Way-sized galaxies. In order to carry out appropriate numerical experiments, we first develop a new method for creating multi-component N-body galaxy models in a stationary state. Unlike previous techniques, this approach can flexibly cope with nearly arbitrary axisymmetric density distributions, and allows the construction of disk galaxy models with distribution functions that have three integrals of motion. To demonstrate the capability and accuracy of our parallel code GALIC in which we implemented the method, we examine 20 different galaxy models and study their stability when evolved as a live N-body system, finding very good results. We then apply the method to study the evolution of thin disk galaxies inserted in high-resolution dark matter halos drawn from the Aquarius simulation suite. The galaxy models are constructed with GALIC and are adiabatically grown in the evolving dark matter halo from redshifts $z = 1.3$ to $z = 1.0$, and then evolved live for a period of about 6 Gyrs to the present epoch. Our analysis of the simulations explores to what extent the galaxies are affected by the dark matter halo's triaxiality and the large number of dark matter subhalos orbiting in it, and by how much the disk orientation is tumbling during this evolution. Finally, we study the radial migration of stars in hydrodynamical simulations of the same Milky Way-sized galaxies, carried out with the novel moving-mesh code AREPO. We are especially interested in the question whether radial migration can strongly modify metallicity gradients and the age-metallicity relation in such galaxies, and whether it can potentially contribute to the formation of a thick disk component.

Zusammenfassung

In der vorliegenden Arbeit wird die Stabilität von Milchstraßen und die radiale Migrationsrate von Sternen in großen Scheibengalaxien mit Hilfe von selbstkonsistenten kosmologischen Galaxienentstehungsmodellen untersucht. Um hierfür geeignete numerische Experimente durchzuführen, entwickeln wir in einem ersten Schritt eine neue Methode zur Bestimmung des stationären Zustands eines N-Teilchen Galaxienmodells. Im Gegensatz zu bisherigen Verfahren kann unsere neue Methode auf beliebige achsensymmetrische Dichteverteilungen angewandt werden. Dies erlaubt die Konstruktion von Galaxienmodellen mit Verteilungsfunktionen welche drei Erhaltungsgrößen beinhalten. Um die Leistungsfähigkeit und Genauigkeit unseres neuen parallelen GALIC Codes zu testen, führen wir N-Teilchen Simulationen von 20 verschiedenen Scheibengalaxien durch.

Wir wenden unsere Methode auf die Entwicklung von dünnen Scheibengalaxien in hochaufgelösten Halos aus dunkler Materie an. Die Galaxienmodelle werden mit dem GALIC Code erstellt und wachsen anschließend zwischen Rotverschiebung $z = 1.3$ und $z = 1.0$ adiabatisch im sich entwickelnden Halo. Im Anschluss wird das gesamte System selbstkonsistent für sechs Milliarden Jahre bis zum heutigen Zeitpunkt entwickelt. In der Analyse dieser Simulationen untersuchen wir die Auswirkungen der Triaxialität der Halos auf die Galaxien, sowie den Einfluss der großen Zahl von Satellitenhalos, welche im Haupthalo kreisen. Insbesondere analysieren wir die Veränderung der Orientierung der Galaxienscheiben während der Simulationen.

Schließlich studieren wir die radiale Sternmigration in hydrodynamischen Simulationen in Galaxien von der Größe der Milchstraße. Diese wurden mit dem neuartigen AREPO Code ausgeführt, welcher die hydrodynamischen Gleichungen auf einem sich bewegenden Voronoi-Gitter löst. Insbesondere interessiert uns die Frage, ob die radiale Migration die Metallizitätsgradienten und die Alter-Metallizität Relation signifikant beeinflusst und einen Beitrag zur Bildung einer dünnen Scheibenkomponente liefert.

1

Introduction

*"What is real? How do you define 'real'?
If you're talking about what you can feel,
what you can smell, what you can taste
and see, then 'real' is simply electrical
signals interpreted by your brain."*
— Morpheus, *The Matrix*

Galaxy formation is a phenomenon of great complexity that incorporates an enormous variety of intertwined physical processes acting on an extremely wide range of scales. The geometrical complexity and lack of simple symmetries make it basically impossible to establish an accurate analytical model of this phenomenon. The only known way to avoid these restrictions and obtain "predictions of the sky" that resemble the level of detail in available observational data is through numerical modeling (see Frenk & White 2012, for a recent review).

The best numerical models of the Universe available today are based on the Lambda Cold Dark Model (ΛCDM). They take initial perturbations observed in the cosmic microwave background (CMB) as an input, and compute the large-scale structure of the Universe at the present-day as output. ΛCDM is not just successful in reproducing many observational facts, it also has substantial predictive power and confronts us with deep puzzles about the nature of the energy and matter content of the Universe. To be the best among potential alternative models, ΛCDM demands that the Universe contains only about 30% of the matter required for closure, to have nevertheless a flat spatial geometry, and to have acoustic peaks of specific lengths and amplitudes in CMB power spectrum that suggest the existence of a dark energy component.

Unfortunately, despite of all the cosmological successes, ΛCDM still poses us with crucial questions on galactic scales. One of these issues is that theoretical models of galaxy formation in ΛCDM typically greatly underestimate the number of thin disk galaxies in the Universe. The reason is that ΛCDM predicts a large number of mergers (during the process of hierarchical assembly of structures). If disk galaxies merge with other galaxies the merger will likely destroy, or at least greatly disrupt the disk, implying that the resulting galaxies are not expected to be thin disk galaxies with small or absent stellar bulges. However, the majority of today's galaxies (65%) are observed to be disk galaxies (Lintott et al. 2008), many of them with small central bulges. This problem does not necessarily mean that the ΛCDM model is potentially wrong, it could rather also mean that our understanding of galaxy formation is still inadequate, and that it requires further refinement to accurately predict the population of galaxies expected to form in the Universe. Currently, a lot of work is therefore being done that try to improve theoretical galaxy formation models. In simulation models, these efforts aim to reduce the freedom in parameterizations of star formation and feedback physics, to improve the numerical integration accuracy, to incorporate additional physics that is potentially important, and to refine the comparison to observational data.

Numerical models have rapidly become more powerful in the recent past. Using them, we have now entered a new era of astronomical research, where computer simulations have become a new tool that can mimic reality in an increasingly realistic fashion. We can carry out controlled experiments of synthetic astrophysical processes, and take their results as an approximate reflection of real experiments. This is particularly fascinating in astronomy, were the corresponding real experiments can actually not be carried out. Controlled experimentation allows us to decompose complex physical processes into their components, and in this way to derive their properties and their roles in the system as a whole. Moreover, the results of computer simulations can be observed in great detail over the entire phase-

9

space at any angle and as close as we like. This offers an unprecedented opportunity to see, for example, how exactly the large-scale structure of the Universe emerged from the primordial noise, how razor thin galaxies are able to survive despite the hazardous merging, or how processes like the radial migration of stars in disk galaxies affect the properties of individual galaxies.

In this thesis, we aim to exploit numerical simulation techniques to study different aspects of the structure, stability and evolution of disk galaxies. In essence, the thesis addresses three different, interrelated topics which consist of (1) a new method to create equilibrium galaxy systems for controlled experimentation and its implementation in the GALIC code, (2) a set of controlled experiments to study the stability of fully formed disk galaxies in a realistic ΛCDM cosmological environment, and (3) an analysis of radial migration of stars in hydrodynamical simulations of Milky Way-sized galaxies carried out with the moving-mesh code AREPO. In this introductory chapter, we in addition provide a short exposition of the cosmological context and the employed numerical techniques.

1.1 Cosmic initial conditions for structure formation

The initial conditions for ΛCDM consist of a near-uniform, slightly perturbed distribution of matter, where the statistics of the fluctuations are given by a nearly scale-invariant Gaussian random field with two free parameters, the *amplitude* of perturbations and the *spectral index*, which measures the slight deviation from exact scale invariance, known as the Harrison-Zel'dovich spectrum (Harrison 1970; Zeldovich 1972). The density fluctuations are of adiabatic nature, where all matter components are coupled and perturbed in the same way, so that the total density varies but the ratio of baryons to dark matter remains spatially invariant. According to recent observational constraints (Planck Collaboration et al. 2013), the

dark matter is assumed to constitute 84.5% of the total matter content. The ordinary matter (15.5%), in turn, is inferred to consist mostly of hydrogen ($\sim 75\%$) and helium ($\sim 25\%$), with a very small admixture of heavier elements ('metals' in astronomy) dominated by lithium. These conditions are thought to exit when the Universe was much smaller than it is now.

After the expansion of the Universe had been discovered (Lemaître 1927; Hubble 1929), it was not difficult to infer that when the evolution is traced backward in time, the Universe has to contract, get denser and become hotter (Lemaître 1931). Extrapolating the known laws of physics to the conditions of the increasing density and temperature, right up to the known limit of their applicability (roughly expected at the Planck scale, $10^{93}\mathrm{g/cm}^3$, $10^{32}\,\mathrm{K}$), it was possible to form a hypothetical picture of the earliest processes taking place in the Universe, known as the Big Bang Theory.

According to this theory, the Universe has gone through a series of epochs and phase transitions during its early expansion. The very first period accessible for substantive theoretical description is the Grand Unification Epoch, when three of the four known interactions (strong, weak and electromagnetic) were merged into one gauge force. As the temperature fell further ($< 10^{28}\,\mathrm{K}$) the unified gauge force split into strong and electroweak forces, and matter took the form of a quark-gluon plasma. The temperature was so high and the random motions of particles were so energetic that particle-antiparticle pairs of all kinds were continuously created and destroyed in collisions. At some point a still poorly understood reaction called baryogenesis broke the symmetry of particle/antiparticle creation and destruction, leading to a very small excess of leptons and quarks over antileptons and antiquarks, of order of one normal extra particle for every 30 million particle-antiparticle pairs. This was enough however to cause matter to dominate over antimatter.

As the density and temperature of the Universe continued to decrease, the typical energy of each particle was decreasing as well, until electroweak symmetry spontaneously broke and the electroweak force split into weak

and electromagnetic forces. At this point, the theoretical picture becomes much clearer and robust, since particle energies dropped to values that can be achieved in particle physics experiments. About 10^{-6} seconds later, quarks and gluons had cooled sufficiently to combine and form baryons such as protons and neutrons, as well as their antimatter partner. Along with a further drop in temperature and density, the balance between creation and annihilation of proton-antiproton and neutron-antineutron pairs became shifted toward annihilation, and the number of surviving baryons quickly approached zero. But since the baryons were born in slightly greater number than antibaryons, thanks to a small excess of quarks over antiquarks, a small fraction of them, approximately 1 per 10 billion, was left over when the creation of new baryon-antibaryon pairs was no longer possible and all existing pairs had annihilated. Something similar happened for electrons and positrons about one second later. Eventually, the remaining protons, neutrons and electrons were no longer moving relativistically and the energy density of the universe became dominated by photons instead (with a minor contribution from neutrinos). After a few minutes, the temperature was about a billion Kelvin and the density was about the density of air. During this time, neutrons combined with protons and formed the universe's deuterium and helium nuclei with a small admixture of light isotopes in a process called Big Bang nucleosynthesis (Alpher et al. 1948).

Once about 380 thousand years had passed, the temperature had dropped so much that the ionization and recombination of protons and electrons went out of equilibrium, and it became possible to form hydrogen atoms. Shortly after, the photons were released from Thomson scattering on free electrons and now could travel freely through the Universe, largely unimpeded. As a consequence of this thermal history, this radiation still has to be present everywhere in the Universe, even today (Gamow 1948; Alpher & Herman 1948). And, indeed, this primordial relic light was detected as cosmic microwave background radiation (CMB), coming uniformly to us from all directions (Penzias & Wilson 1965). This confirms a remarkable
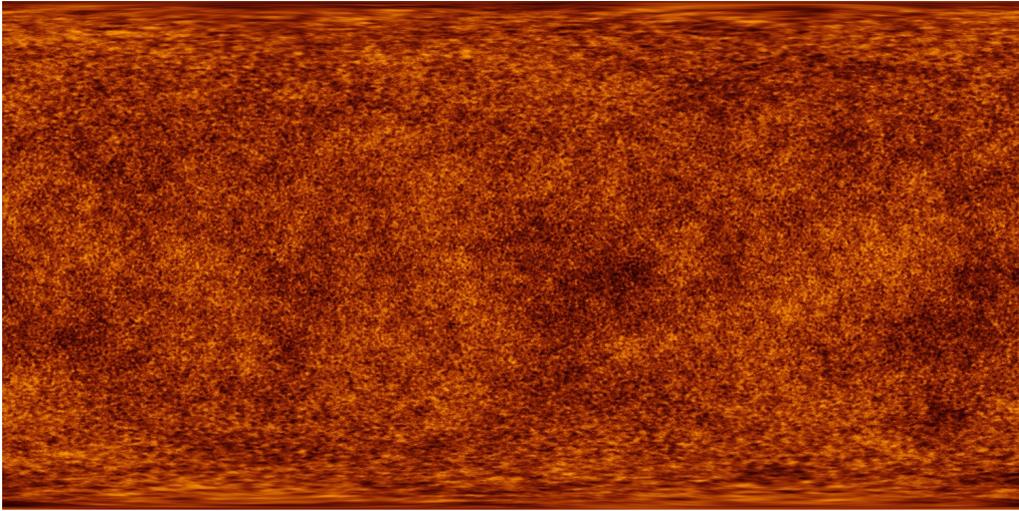
Figure 1.1: All-sky map of CMB anisotropies measured by ESA's Planck space telescope (Planck Collaboration et al. 2013). It shows tiny temperature fluctuations that correspond to regions of slightly different densities at very early times, representing the seeds of all future structure such as today's galaxies.

prediction of the Big Bang Theory.

The detailed study of the CMB, first with the Cosmic Background Explorer (COBE) satellite (Smoot et al. 1992), revealed that in addition to forming a highly isotropic nearly perfect black body signal, there are extremely faint fluctuations of relative amplitude 4×10^{-5} imprinted in the temperature of the radiation (see Fig. 1.1), after the dipole component due to our own motion relative to the CMB is subtracted. Since the radiation and matter were coupled before recombination these fluctuations have to reflect inhomogeneities in the matter distribution at the distant epoch where the radiation was released. The Big Bang theory itself provided no explanation for these inhomogeneities and their origin. Moreover, the statistical properties of CMB anisotropies were found to be the same at angles exceeding the cosmological horizon, i.e. for separations where a causal contact seems impossible. This marks the *horizon problem*, describing the puzzle that different parts of the Universe should have not been able to communicate with each other yet they show highly similar properties. The

latter can only be naturally understood if they have been in casual contact in the past after all (Misner 1969). In addition, the Big Bang theory alone fails to explain the *flatness problem*, describing the observational fact that the Universe's matter and energy density appear to be fine-tuned to critical density (Dicke & Peebles 1979).

To solve these problems, a modification of the simple Big Bang scenario, known as *inflation theory*, has been proposed (Guth 1981; Linde 1982). According to this theory the initially hyper-dense and hot Universe has gone through a phase of rapid, exponential expansion (called inflation) driven by the vacuum energy of a hypothetical quantum field (inflation). This rapid expansion smoothed out the matter beyond the cosmological horizon, providing the isotropy and homogeneity of its properties in areas that have later not been able to communicate any more. Another generic consequence of this proposal is that density perturbations generated by quantum fluctuations at early times show up on macroscopic scales in the density distribution of the Universe and become imprinted on it. Also according to the inflation scenario, the Universe expanded so much beyond the cosmological horizon that its curvature radius grew to be so large that the geometry of the observable Universe should appear basically flat. While the Big Bang theory combined with inflation tells nothing about the true nature of Dark Matter (Zwicky 1933; Rubin & Ford 1970) and Dark Energy (Riess et al. 1998; Perlmutter et al. 1999), it provides a clear and very successful theoretical framework for the growth of cosmic structure and the origin of the ΛCDM initial conditions.

1.2   Galaxy formation processes

In studying galaxy formation, we are faced with a wide mix of object sizes and a diverse set of physical processes whose properties and detailed laws are not all well understood. This complexity hampers a formalization of galaxy formation theory and makes even a qualitative understanding of this natural

phenomenon a non-trivial goal. However, significant strides in unraveling galaxy formation have been achieved in recent times. The purpose of this subsection is to briefly outline a qualitative picture of galaxy formation, corresponding to the current understanding in the field.

Once the baryonic matter decouples from radiation and stops being smoothed by its pressure, it quickly catches up with the growing roughness of dark matter spreading over its potential wells. By the typical starting redshifts of cosmological simulations, $z \sim 100$, the spatial distribution and the level of inhomogeneities in baryons and dark matter become almost indistinguishable (Zel'dovich 1970).

The reason for the heterogeneity in the early structure growth lies in the interplay of cosmic expansion and gravity. Overdense regions expand slower than rarefied areas because there gravity is working more strongly against the expansion. As a result, a dense region tends to become denser and smaller, while a low density region becomes more diffuse and bigger. The expanding regions of lower density, so-called voids, acquire the topology of inflating bubbles (Dubinski et al. 1993). Matter appears "pushed" to their periphery/outskirts and becomes more concentrated in the spaces between the bubbles. As the expansion continues, the bubbles contact each other and the matter is further "displaced" to the contact interfaces, forming so-called *filaments* that connect regions of maximum matter concentration in the inter-bubble space. Thus the matter distribution starts to resemble a foam-like Voronoi mesh (Aragón-Calvo et al. 2007), known as the *Cosmic Web*, as depicted in Fig. 1.2.

At a certain stage, the most overdense regions become so dense that the self-gravity completely stops their expansion, they turn around and begin to collapse. These regions decouple from the background expansion. Because the dark matter is believed to be effectively collisionless, it collapses in a dissipation-less way and shrinks to half the size reached at maximum expansion, forming quasi-virialized objects called *dark matter halos*. In contrast to the dark matter, the baryonic gas is collisional and subject to electro-
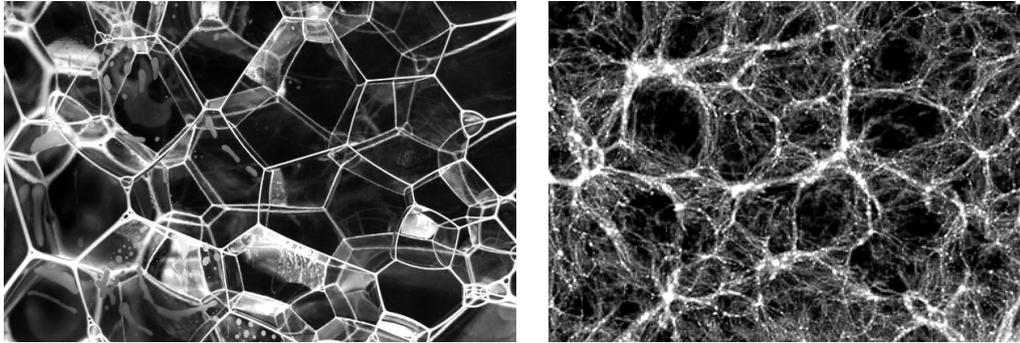
15

Figure 1.2: Left image: Soap bubbles (adopted from Shutterstock.com). Each bubble represents approximately a specific region that is closer to its center than to the center of any other bubble. Such regions are called Voronoi cells. Right image: The cosmic web (adopted from Springel et al. 2005b), characterized by a foam-like distribution of matter. The void intersections are identified with walls in matter the density. The edges forming at the rim of each wall are identified with filaments in the matter distribution. The densest regions with a maximum concentration of matter are located at the nodes of this web.

magnetic forces. Therefore it can cool radiatively (via Bremsstrahlung radiation, atomic excitation and de-excitation mechanisms, as well as inverse Compton scattering of CMB photons). Because of its dissipative nature, gas can loose its pressure support and contract to much smaller scales, reaching the kind of densities needed to trigger star formation.

The first stars in the Universe were likely the most massive and the most powerful in cosmic history (Bromm et al. 2002). This is because for later generations of stars the temperature of the gas from which stars are formed is expected to be lower due to the presence of metals, and the lower the temperature, the lower the gas pressure and therefore the less mass is needed before the final runaway collapse to make a star sets in (Larson 2000). The powerful radiation of the first stars ionized and heated the surrounding gas, possibly preventing other stars from forming out of the remaining gas in the same cloud (Abel et al. 2002). These single, massive stars are hence expected to sit in the centers of virialized dark matter clumps filled with hot gas. Later, these sites will become the seeds of the first galaxies.

The life of the first stars, known as Population-III stars, was very short

and ended with giant explosions that swept their material into the surrounding gas and enriched it with first heavy elements. Enriched gas cools much more quickly cool and contracts efficiently. It is believed that the formation rate of the second generation of stars (Population II) defines the original morphology of the galaxies (*protogalaxies*). If most of the gas turns into stars as it falls in, the collapse becomes dissipation-less, and infall motions turn into random motions of stars, resulting in an *elliptical galaxy*. On the other hand, if the cloud remains gaseous during collapse, the gravitational energy can be effectively dissipated via shocks and radiative cooling. In this case, the cloud will shrink until it is supported by angular momentum, leading to the formation of a rotationally-supported disk, and eventually to the formation of a *spiral galaxy* (White & Rees 1978).

As the cosmic expansion continues, increasingly large areas of higher than average density are being decoupled from the background expansion, turn around and start to collapse. This leads to the hierarchical merging of protogalaxies in a bottom-up way, in which ever larger regions decouple from the expansion and their embedded smaller structures merge to form bigger structures (Searle & Zinn 1978). The morphology of each galaxy depends sensitively on the type and amount of mergers that it has experienced during its formation. If a galaxy that is originally a spiral undergoes a number of major mergers this unavoidably destroys its disk in a process of violent relaxation that transforms infall energy of the progenitors into internal energy of the virialized remnant. Effectively, this converts a spiral galaxy with ordered stellar motion into a fuzzy ball of unordered stellar orbits, looking like an elliptical galaxy. If there is still remaining gas that cools onto the elliptical galaxy, a new disk may form, resulting in a disk-bulge system. Potentially, a spiral galaxy can also protect its morphology during a merging process if only mergers with small objects occur, i.e. if its merging history is quiet enough. This leads to the expectation that spiral galaxies are preferentially found in relative isolation, whereas ellipticals should mostly populate crowded environments like a galaxy cluster

(Dressler 1980).

By $z \sim 1$ major mergers become rare in almost all parts of the Universe (Ryan et al. 2008) and internal dynamical effects come to the forefront. For example, a successfully formed spiral galaxy may be intrinsically unstable and thereby significantly change its morphology at later time. This mainly affects the evolution of massive spiral galaxies (Melvin et al. 2014). A thin disks with too high a surface density are often subject to non-axisymmetric instabilities (Sellwood 2013), which produce cigar-like structure in their centers, known as *bar*.

The evolution of stars also affects the evolution of galaxies via so-called feedback processes. Radiation and blastwaves from supernovae explosions heat the surrounding gas and potentially vent it out of galaxies, significantly reducing the rate of star formation (Larson 1974). On the other hand, supernova explosions can also compress the surrounding gas and thus enhance the star formation rate locally. Another important source of feedback is provided by *Active Galactic Nuclei* (AGN) (Silk & Rees 1998), i.e. accreting supermassive black holes at the centers of galaxies. If they manage to transfer part of the enormous energy liberated during black hole growth into the surrounding gas, they may play an important role in shutting off star formation in very large galaxies, especially those at the centers of galaxy clusters. We note however that this picture of galaxy formation, especially the role played by feedback processes, is still very sketchy and contains aspects that are poorly understood. Many ongoing theoretical studies and observational campaigns therefore aim to refine it further.

1.3   Numerical modeling techniques

In contrast to the qualitative view described above, modern numerical models allow a description of the hierarchical merging of galaxies in its full geometrical complexity. They operate by discretizing the underlying physical objects and processes, allowing a description in terms of the processes that

act on small spatial elements. For example, computer models describe the motion of individual dark matter particles (actually, they use more massive "macro" particles) under the action of gravity, and the motion of gas elements under the action of simple hyperbolic conservation laws describing ideal hydrodynamics. Nonlinearities in the interactions are not a problem for the numerical models, and complex behaviour of systems can arise from the collective behaviour of all the individual elements. Typical simulations are computed as time evolution models, where the dynamics over a small time interval is approximated with linear solutions to the dynamical equations, based on an initial state of the system. Then the new predicted states of every element are considered as the new initial state, and the procedure is repeated, until the evolution of the system is described over the desired finite time interval.

There are two main concerns with numerical methods: *computational cost* and *accuracy*. They both depend sensitively on the *resolution*, i.e. the number of elements on which we discretize the system. This number is sometimes also limited by how well we understand the properties of the small discrete elements that are used; if very small scale processes are not understood, then it may make no sense to try to resolve them. The higher the resolution, the smaller the discrete elements become. Also, they then become more linear in their properties, and consequently the description of their behavior becomes more precise. However, a larger number of elements implies more calculations and hence a larger computational cost. At the same time, the better we understand the properties and laws of individual elements, the more precisely we can predict their behavior, or alternatively, we can reduce computational cost while keeping the same accuracy. Finding an optimum compromise between accuracy and computational cost for given computational resources is a difficult task. In practice, this optimum is often determined heuristically, and is in any case dictated by the limited available computing resources.

For the numerical treatment of gas, several different discretization meth-

ods have been invented, among them splitting gas into individual particles (so-called smoothed particle hydrodynamics, SPH) (Lucy 1977; Gingold & Monaghan 1977; Monaghan 1992), into Cartesian cells (Stone & Norman 1992; Cen 1992; Teyssier 2002) or into unstructured and moving Voronoi cells (Springel 2010; Duffell & MacFadyen 2011). Interestingly, the numerical properties and errors of these gas elements are different, although the underlying physical equations are the same. In the simplest versions of these methods, the gas is treated on the assumption that its properties at the molecular level do not change. However, the evolution of galaxies is also sensitive to processes associated with a change of internal properties of atoms, such as *radiative cooling*, *photo-ionization* and *chemical enrichment* (Katz et al. 1992; Steinmetz & Mueller 1994). Such processes are part of the subgrid physics and they are added to a numerical model of an ideal gas as a set of additional properties and laws. Another critical, but less well understood subgrid process is the transformation of gas into stars (Springel & Hernquist 2003). In addition to changing the state of the gas elements, this leads to the creation or removal of star particles. Often, the prescriptions used by numerical models for these processes are not based on fundamental physical laws, but consist of ad hoc equations that are calibrated against observational data, or in some cases on smaller-scale simulations with high resolution. In either case, such sub-grid prescriptions may not be valid in all physical regimes (see Vogelsberger et al. 2013, for details).

Gravity, which determines the dynamics of the star and dark matter particles and also plays an important role in the dynamics of gas elements is in principle a fairly simple physical law. However, its computational cost is very high because it is a long range force. In order to calculate the gravitational force acting on one particle one needs to calculate the interaction with all the other particles, resulting in a N^2 scaling of the computational cost to calculate the gravitational force for all particles, where N is the total number of particles. A variety of methods had been developed to reduce this cost, among them Particle-Mesh (Hockney & Eastwood 1981),

Multi-Grid (Fedorenko 1962; Brandt 1977) and Tree-based (Barnes & Hut 1986) methods. Involving approximations to the force they can drop the cost to $N \log N$, while keeping the force accuracy good enough for the purposes of collisionless systems. One can also gain computational efficiency by adaptive time integration. In essence, by analyzing the behavior of a particle orbit in a self-gravitating N-body system one can detect the level of nonlinearity in the orbit and compute the force less frequently when the motion is close to linear. This idea is the basis of hierarchical time stepping techniques that are commonly employed in cosmological simulations.

Accounting for stellar evolution is necessary to properly describe the changing properties of star particles in simulations, which usually represent full stellar populations of at least tens of thousands of stars. Fortunately, these stellar populations can be viewed as macro-particules whose evolution is clearly defined based on spectro-photometric and stellar population synthesis models that have been developed for decades. Their behavior can thus be predicted with a high degree of accuracy, just based on the initial stellar mass function and chemical composition of the star-forming gas (Springel & Hernquist 2003).

The expansion of the Universe needs to be reflected in numerical models of galaxy formation as well. There are two ways to accomplish this. First, a simple model of cosmic structure formation can be constructed by using a sphere of matter that is initially put into homologous expansion, i.e. a velocity proportional to the distance from the center is added to dark matter particles and gas elements. Due to Birkhoff's theorem, the expansion at a given radius is slowed down by the gravity of the enclosed mass, and the resulting equation of motion agrees with the one derived through the Friedman-Lemaitre solutions of general relativity. Alternatively, one can subtract the expansion from the proper motion, and describe the motion in terms of a peculiar velocity relative to an expanding homogenous background. This practice, known as integration in comoving coordinates, is usually adopted. It also allows the treatment of effectively infinite spaces

in a Newtonian approximation by applying the trick of periodic replication to a fundamental simulation box.

Today, the main efforts in developing better numerical models is mainly aimed towards a better description of the sub-resolution physics, and to improve the balance between computational cost and accuracy. The resolution that can ultimately be used for simulations of galaxy formation is determined both by technological advances and our capability to utilize it. It is in fact highly non-trivial to leverage the full power of modern parallel supercomputers, because not all calculations can be easily parallelized in a perfectly efficient fashion.

1.4 Important codes and simulation projects

To bring a numerical model to a life and use it scientifically, one needs to write computer code that realizes all operations prescribed by the numerical model. Nowadays, a computational astrophysicist is as likely to be found puzzling over the engineering of a complex computer code and struggling to make it run on a computer cluster than to be found working with mathematical equations on paper (or chalkboard). Making good use of the rapidly growing power of modern computer hardware is a great challenge for astrophysics, requiring also some knowledge in computer science.

Over the last decades several advanced codes for cosmological simulations have been developed, among them ENZO, Flash, RAMSES, GADGET, Gasoline, HYDRA, and AREPO. A special mentioning should perhaps be made of the GADGET code (Springel et al. 2001b), which is probably the most widely employed code in the field – and we will also make use of it in this work. It is a general purpose, publicly available code that computes gravitational forces with a hierarchical tree algorithm and represents gas by means of smoothed particle hydrodynamics (SPH). It features hierarchical block time stepping that significantly reduces the total calculation cost when a wide dynamic range in time is present. The code can integrate

isolated systems in proper or in comoving coordinates both with or without periodic boundary conditions, and is parallelized using the message passing interface (MPI), based on work-load balanced domain decomposition. This allows the efficient use of massively parallel computers with distributed memory.

The GADGET code was used in 2005 to carry out the biggest cosmological simulation of the growth of dark matter structure of its time, known as *Millenium Simulation* (Springel et al. 2005b). This calculation has subsequently seen very intense scientific analysis. To follow the dark matter distribution it used $\approx 10^{10}$ particles in a cubic region $500\,h^{-1}\mathrm{Mpc}$ on a side. A supercomputer located in Garching, Germany, executed the simulation, for more than a month. The output of the simulation needed about 25 terabytes of storage. Application of semi-analytic techniques to the stored output of this calculation predicted the distribution of $\sim 10^7$ galaxies.

In 2008, the most comprehensive study for individual dark matter halos thus far was done, known as *Aquarius Project* (Springel et al. 2008a). This was also carried out with the GADGET code. Aquarius simulated a sample of 6 ultra-highly resolved Milky-way sized halos using a so-called zoom-in simulation technique. In our thesis work, we will also study the Aquarius halos and use them for our disk insertion experiments. The halos of the Aquarius simulation project are especially interesting due to the high quality of the initial conditions, which in particular allow accurate convergence studies.

We will also analyze hydrodynamical simulations of the Aquarius halos, which self-consistently form stellar disks. In fact, Marinacci et al. (2014a) achieved for the first time disk galaxies with realistic properties in such fully self-consistent cosmological hydrodynamic simulations. This was accomplished with the help of the AREPO code, which offers a much more accurate method for treating gas dynamics than the older SPH technique. In addition, these simulations use a comprehensive model for galaxy formation physics that is more sophisticated than in previous work. This model

23

was also employed in the Illustris simulation project (Vogelsberger et al. 2014b), the thus far largest and most ambitious attempt to model galaxy formation with cosmological hydrodynamical simulations in a representative part of the Universe.

1.5   Simulations as virtual laboratories

Unlike most other disciplines in natural science, astronomy is not a 'laboratory science', because nobody can directly experiment with objects in the sky, no one can rearrange them. However, now that we have entered the era of advanced computer simulations, one can create a synthetic universe in a computer and "move" stars and galaxies in it. Numerical experiments are similarly affected by assumptions, caveats, and limitations associated with any traditional, laboratory-based experiments. And indeed, in common terminology in the field, simulated results are often described as being "empirical" data, a term usually reserved for natural phenomena rather than for numerical models of nature. Or as Kevin Heng (2014) puts it, simulated data are referred to as "data sets", seemingly placing them on an equal footing with observed natural phenomena.

The danger of blurring this distinction is probably small in reality, even though virtual laboratories offer opportunities for experimentation and observation that in some sense even go beyond what is possible in any real physical laboratory. In a virtual laboratory on the computer, we can turn off or turn on any processes we want, analyzing how they affect each other. Or we can change their properties (even into unphysical regimes) in order to understand what role they play in detail. We can trace back in time the evolution of any object to see how its properties emerged. We can observe a process under any spatial and time angle, unrestricted by sensitive limitations of instruments.

Applied to astrophysics in the form of this thesis, the primary objective of such experimentation is to improve our understanding of galaxy forma-
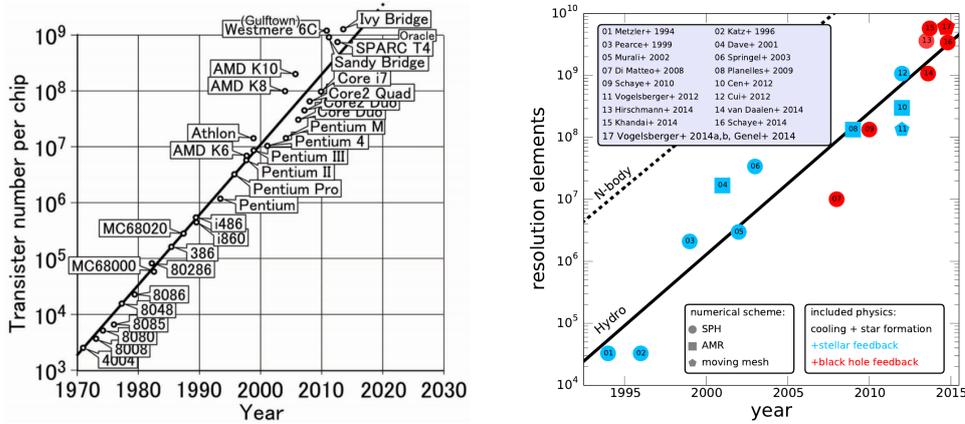
Figure 1.3: Left image: Moore's Law for microprocessors, showing that the number of transistors in microchips doubles approximately every two years (adopted from WikimediaFoundation.org). Right image: Growth of resolution elements in cosmological simulations as a function of publication year (adopted from Genel et al. 2014). This shows the same tendency to double roughly every two years.

tion. Also, as numerical modeling has been demonstrated to have predictive power, it is in principle possible to discover physical phenomena that have not been detected yet in observational data. There is no doubt however that when interpreting numerical experiments we have to be careful. To use a simulation as a laboratory, one has to understand its limitations – otherwise, one may wrongly identify an artifact as a result. For example, numerical oscillations or enhanced viscosity that arise from errors in numerical schemes could be easily be mistaken as physical effects.

Since Moore's law (Moore 1965) still holds true (see Fig. 1.3), continuing the trend over the last 40 years, we can hope that the accuracy and level of detail reached in numerical simulations will still continue to grow (Genel et al. 2014). And consequently, the reliability of virtual experimentation should continue to grow as well. If one day all of the relevant physical laws of galaxy formation will be captured and incorporated in a numerical simulation, then one may end up with an ab-initio prediction of galaxy formation that nearly perfectly reproduces the physical reality. At the moment, this is still a distant goal.

1.6   Structure of this thesis

The key properties of spiral galaxies, which are also commonly called disk galaxies, are that they are thin stellar systems which rotate rapidly and often show spiral structure. One of the main challenges to our qualitative understanding of galaxy formation in the ΛCDM cosmogony is the great number of thin disk galaxies observed in the local Universe. Many of them are also found to contain old stars in their disks. But the problem is that disks are very fragile, and mergers with other galaxies can easily destroy thin disks. This conceptually challenges the hierarchical "bottom-up" scenario, where basically all galaxies form as the end result of a number of merging processes.

This concern becomes a real problem once one considers that our best theoretical ΛCDM simulation models are thus far all failing to form a sufficient number of disk galaxies with low bulge-to-disk ratio (Vogelsberger et al. 2014b; Genel et al. 2014). There is a hope that this may simply be due to a lack of resolution and inadequate modeling of the relevant physics, preventing the formation of disk galaxies. In this case, we may just need to wait until numerical simulations reach a level where numerical resolution and physics modeling is not a limitation any more.

Unfortunately there is however also the nagging possibility that something is conceptually wrong, and future simulations may then not be able to resolve this issue. To address this question already today one can for example consider controlled experiments that place fully formed thin disk galaxies with realistic properties inside a live dark matter halo as it assembles and evolves cosmologically, and see whether this galaxy can survive in the tumultuous cosmological environment.

Unfortunately there are obstacles on the way to such an experiment. In particular, one evidently needs a disk galaxy that is in equilibrium with itself and with the dark matter halo it is going to be placed in to carry out such an experiment. We have to make sure that the disk model will

not be destroyed or deformed because of intrinsic instabilities, but rather is only affected by the host's halo evolution, otherwise any conclusion about whether thin disks can survive or not would be misleading. However, it is actually quite challenging to realize particle disks with prescribed realistic properties that are in near perfect dynamical equilibrium. Several such methods were invented that can in principle be tried to achieve this goal, among them are distribution function based and moment-based methods, Schwarzschild's method, and the made-to-measure method, but none of them is actually suitable and accurate enough for the controlled experiment we have described.

In order to overcome this problem, we have developed a new method that is powerful enough to construct high-quality steady-state compound disk galaxies for such experiments. Our method will be described in full in the next chapter of this thesis and originates in principle from a simple idea (see Fig. 1.4), namely to iteratively adjust the initial velocities of the particles until their time-averaged orbits reproduce a prescribed mass distribution. In order for this idea to work the velocity adjustments need to be well balanced and not biased to specific particles or specific directions, which is one reason why a practical realisation of the idea is difficult. To arrive at a robust scheme, we combine our method with analytical solutions of the Jeans equations to obtain second velocity moments that are imposed as additional velocity constrains. This allowed us to develop a general purpose code which we dubbed GALIC that can create composite halo-disk-bulge galaxy models with prescribed density distribution and velocity anisotropy.

Using this newly developed code we initialized well equilibrated disks in the full set of eight Aquarius halos and re-simulated their joint evolution, as described in Chapter 3 of this thesis. We analyzed in particular the role played by the triaxiality of dark matter halos, by infalling dark matter subhalos and by the initial structural properties of the disk and bulge. Interestingly, we find that disks tumble substantially in space as they are torqued by the triaxial dark halo potential, but their survival is actually
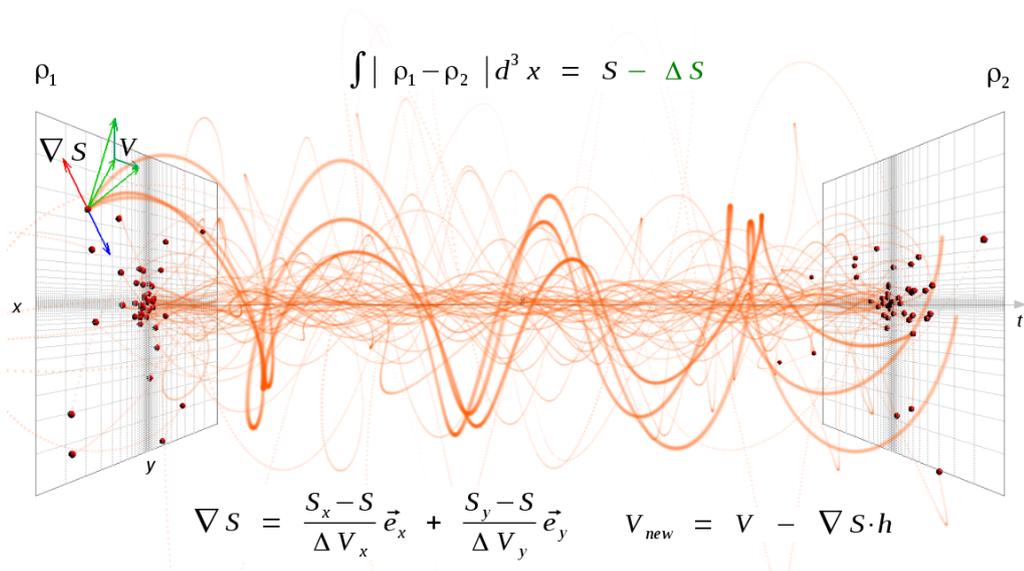
$$\int |\,\rho_1 - \rho_2\,| \, d^3x \;=\; S \,-\, \Delta S$$

$$\nabla S \;=\; \frac{S_x - S}{\Delta V_x}\,\vec{e}_x \;+\; \frac{S_y - S}{\Delta V_y}\,\vec{e}_y \qquad V_{new} \;=\; V \,-\, \nabla S \cdot h$$

Figure 1.4: This plot illustrates our original idea that led us to develop a new method for the construction of N-body systems in collisionless equilibrium, as described in full in Chapter 2 of this thesis.

less problematic than it may seem, provided the disks are light enough and the dark matter and/or stellar bulge is concentrated enough to avoid the formation of a stellar bar.

Finally, in Chapter 4 of this thesis we consider another important science topic related to stellar disks, namely the radial migration of stars. If a disk galaxy is assumed to be a largely unperturbed system then one would expect to have a nice correlation between age and metallicities of stars. However, observationally it is found that the stars of a given age and at a given radius have a broad spread of metallicity (Edvardsson et al. 1993). This poses the question where this spread of metallicities comes from. One proposal is that it may result from a radial redistribution of stars due to scattering of circular orbits on transient spiral waves, which may alter the angular moment of stars without increasing their random motions much (Sellwood & Binney 2002). Because there is a metallicity gradient in disks, stars that formed in the inner disk are expected to be more metal rich than those formed in outer parts. So as a result of radial migration, the

scatter in the age-metallicity relation may be substantially affected. In addition, it has been speculated that radial migration may form a so-called "thick-disk" component in the Milky Way. We study the importance of radial migration in a suite of hydrodynamical cosmological simulations of disk galaxy formation carried out with the AREPO code. All previous work on the subject has in contrast relied on idealized isolated galactic disc, an approach that suffers from a significant dependence on arbitrary assumptions made for the initial conditions, something that is eliminated in our approach.

In brief, this thesis is organized as follows. In Chapter 2, we describe our new method to create equilibrium galactic systems and our new code GALIC that implements the approach. In Chapter 3, we describe a set of controlled numerical experiments aimed at studying the influence of dark matter halo assembly on the stability of fully formed disk galaxies that we insert into the Aquarius halos with the help of the GALIC code. In Chapter 4 we analyze the radial migration of stars in state-of-the-art hydrodynamical simulations of Milky Way-sized galaxies carried out with the moving-mesh code AREPO. Finally, we conclude with a brief summary of our findings in Chapter 5.

The contents of Chapter 2 have been published in slightly modified form in Yurin & Springel, Monthly Notves of the Royal Astronomical Society, 444, 62 (2014). Submission of Chapters 3 and 4 as journal publications is in preparation.

A new iterative method for the construction of N-body galaxy models in collisionless equilibrium

"Think beyond the limits"
— HITS' slogan

Abstract

We describe a new iterative approach for the realization of equilibrium N-body systems for given density distributions. Our method uses elements of Schwarzschild's technique and of the made-to-measure method, but is based on a different principle. Starting with some initial assignment of particle velocities, the difference of the time-averaged density response produced by the particle orbits with respect to the initial density configuration is characterized through a merit function, and a stationary solution of the collisionless Boltzmann equation is found by minimizing this merit function directly by iteratively adjusting the initial velocities. Because the distribution function is in general not unique for a given density structure, we augment the merit function with additional constraints that single out a desired target solution. The velocity adjustment is carried out with a stochastic process in which new velocities are randomly drawn from an approximate solution of the distribution function, but are kept only when they improve the fit. Our method converges rapidly and is flexible enough to allow the construction of solutions with third integrals of motion, including disk galaxies in which radial and vertical dispersions are different. A parallel code for the calculation of compound galaxy models with this new method is made publicly available.

2.1   Introduction

The large number of stars and dark matter particles in galaxies and galaxy clusters makes them essentially perfect collisionless systems. Their dynamics is hence described by the collisionless Boltzmann equation, coupled to self-gravity through Poisson's equation. Relaxed systems correspond to stationary solutions of these equations, and much of the field of galactic dynamics is concerned with understanding different aspects of these solutions (see Binney & Tremaine 2008, for an excellent exposition). This is particularly important for using observational probes of kinematics to infer, for example, something about the underlying density distribution.

Numerical N-body simulations have become a primary work-horse to study collisionless systems, both in stationary and dynamic situations. Prominent examples include the study of bar instabilities (e.g. Athanassoula 2002), the formation of spiral waves (e.g. D'Onghia et al. 2013), or major and minor mergers of galaxies (e.g. Barnes & Hernquist 1992; Hernquist & Mihos 1995). They are also actively used to study the response of disks to the bombardment by dark matter clumps (e.g. Kazantzidis et al. 2008; D'Onghia et al. 2010), or the radial migration of stars caused by resonance scattering (e.g. Sellwood & Binney 2002), and many more.

In carrying out numerical experiments targeting these questions, a recurrent challenge is to construct suitable initial conditions. One usually requires them to be in a reasonably stable, approximate equilibrium in the beginning, otherwise any subsequent dynamics may be dominated or heavily contaminated by the specific out-of-equilibrium state one started out with. Often, one has a relatively clear notion of the density structure one wants to realize, but initializing the particle velocities appropriately is a quite non-trivial problem. This is because doing this *perfectly* requires knowledge of the full distribution function (DF) of the system, or in other words, availability of a stationary solution of the collisionless Boltzmann equation. However, such solutions are analytically known only for a very

limited number of density distributions.

There is hence significant demand to construct equilibrium solutions numerically, not only for realizing N-body initial conditions but also in the context of modelling observational data sets. In the latter case, finding such models is a main component of the reverse-engineering process aimed at constructing self-consistent three-dimensional systems that reproduce the observations. They can then be examined in great detail, allowing insights into properties that are not directly observable (Cretton et al. 1999; van den Bosch et al. 2012).

To our knowledge, there are presently mainly five different methods in use for constructing such equilibrium models:

1. DF-based: For certain mass distributions, the distribution function (DF) can be analytically calculated or accurately approximated. Unfortunately, this ideal case is not generally available for arbitrary density distributions. The main problem is that we do not know the analytical form of the third integral of motion. In some cases it may be reasonably approximated, but this leads at best to nearly self-consistent solutions (Kuijken & Dubinski 1995; Widrow & Dubinski 2005). Nevertheless, there are some useful classes of solutions known, for example for spherical galaxies (Osipkov 1979; Jaffe 1983; Merritt 1985; Hernquist 1990). However, because many real systems are not particularly close to any of these parameterized classes of systems, the approach is rather restrictive in practice.

2. Moment-based: Moments of the velocity distribution can be calculated or estimated with the hierarchy of Jeans equations. If one neglects higher order moments and assumes a functional form for the velocity distribution (often taken to be Gaussian, e.g. Hernquist 1993; Springel & White 1999) that reproduces the estimated moments, one obtains an approximate distribution function. This method is quite general and can be applied to all mass distributions. Since the true

velocity distribution function is usually close to a triaxial Gaussian for much of the mass of a system, the method typically produces systems that are roughly in equilibrium. But the crux is that this equilibrium is by no means perfect, and that it is hard to overcome this limitation within this method. Especially difficult are the central regions of galaxies; when the constructed ICs are evolved in time, one here typically finds density ripples propagating through the system while it relaxes to a true equilibrium state. This can interfere with the interpretation of numerical experiments, especially when they require particularly quiet ICs.

3. Orbit-based method: Schwarzschild (1979) introduced a radically different approach to solve the problem. He suggested to integrate a wide variety of orbits in a given potential, and then to distribute the mass of the system over this orbit library such that the time-averaged density of the system becomes as close as possible to the one corresponding to the potential. Finding the weights of each of these orbits defines a linear optimization problem with positive coefficients, which can be solved iteratively. The resulting weights then effectively define the velocity distribution function. A practical problem with this method is that the size of the orbit library is severely constrained by the available memory. Moreover, the method is ill-conditioned in its basic form, something that needs to be cured by adding ad-hoc assumptions such as smoothness constraints or maximum entropy measures for the weights. Also, the velocity distribution functions constructed with this method are typically very noisy and may feature large jumps. One needs to smooth them, but the required level of smoothing is hard to define. Many attempts have been made to overcome these difficulties (e.g. Vandervoort 1984; Jalali & Tremaine 2011).

4. Made-to-measure: Attempts to improve on Schwarzschild's method

have resulted in a new technique where the orbit integration process and the mass/weight redistribution are combined. This 'made-to-measure' technique makes the storage of a full orbit library unnecessary and therefore removes the memory barrier. But it still requires a smoothing procedure for the velocity distributions (Syer & Tremaine 1996; Dehnen 2009).

5. Guided-relaxation: Another class of methods exploits the fact that any isolated system left to itself tends to an equilibrium state. Knowing the target mass distribution we may try to directly relax to it by steering a system appropriately in the process. This guiding can be done by introducing an additional force, e.g. adiabatic drag on the vertical components of the particle velocities in order to squeeze the system (Holley-Bockelmann et al. 2001), or we may restrict particle mobility such that the target density distribution is maintained and the systems evolves towards a self-consistent equilibrium state (Rodionov et al. 2009). A disadvantage of this approach is that it involves one of the other methods to create an initial state for the further relaxation. Also, there is only limited control on the outcome, making it, e.g., difficult to construct systems with a prescribed velocity anisotropy.

The purpose of this work is to introduce a new, flexible approach for the construction of compound N-body models of axisymmetric galaxies in an essentially perfect equilibrium state. The method only requires the specification of the density profiles of the different components and a selection of the desired bulk properties of the velocity structure, such as the degree of rotational support or the ratio between radial and vertical velocity dispersion in the disk plane. Our code then constructs an N-body system that is in equilibrium and fulfills the imposed constraints on the velocity structure. Implicitly, it hence also provides a solution for the full 3D distribution function. This is achieved for essentially arbitrary axisymmetric density

structure and by taking the mutual influence of different mass components (if present) fully into account. We argue that the resulting flexibility and accuracy makes our approach an attractive alternative compared with other IC generation methods in the literature.

This chapter is structured as follows. In Section 2.2, we describe the basic methodology adopted in our method, which consists of an iterative procedure to adjust the velocities of an N-body realization of a galaxy model until the prescribed density structure is maintained self-consistently under time evolution, and the imposed velocity constraints are fulfilled. In Section 2.3, we highlight how we specify velocity constraints for different structural choices. They take the form of second velocity moments which we determine through solutions of the Jeans equations. We then specify in Section 2.4 various implementation details of our numerical methods as realized in the GALIC code introduced here. Section 2.5 is concerned with a brief description of the specific density profile models currently implemented in this code; these are employed for a suite of tests presented in Section 2.6. Finally, we conclude with a summary of our findings in Section 2.7.

2.2   Methodology

If density profiles for all collisionless mass components of a galaxy model are given, a random N-body realization of particle positions can be easily created by interpreting the density distribution as a probability field for a Poisson point process. But assigning suitable velocities to the particles is the difficult step. Our idea is to do this iteratively: Starting from some guess for the particle velocities, we try to correct them such that the system becomes closer to the desired equilibrium state. This is in some sense similar to Schwarzschild's method, where one models the distribution function from which the velocities are drawn through a set of weights, which are then iteratively adjusted until a global merit function is extremized. Differently from this technique, we however optimize the velocities of each particle

directly. This eliminates the explicit orbit library of the Schwarzschild method, and all the restrictions that come with it. Instead, the particles of our N-body model themselves define the orbit set that is optimized. Importantly, this set is free of any discreteness restrictions or potential biases due to incompleteness of the Schwarzschild orbit library.

As basic merit function that is optimized we consider the difference between the target density field and the actual density response created by our N-body realization with the currently assigned initial velocities. The density response is here defined as the time-averaged density field of the N-body orbits, calculated in the static potential of the target density distribution. For a steady state system, this density response is supposed to be time invariant, and equal to the initial density field. We can readily imagine several different optimization schemes that adjust individual particle velocities iteratively such that this difference becomes as small as possible, for example multi-dimensional steepest decent.

Before discussing the details of our specific solution for this, it is however prudent to consider two apparent conceptual problems with the basic approach as outlined thus far. One is that the density structure does not uniquely specify the velocity structure of an equilibrium model, or in other words, there can be more than one steady-state solution of the collisionless Boltzmann equation for a given density structure. For example, for a spherically symmetric mass distribution, one can have solutions where the velocity distribution is isotropic everywhere, and the distribution function depends only on energy (the 'ergodic' case). But there are also solutions with an anisotropy between radial and tangential motions. Furthermore, one can also have many different axisymmetric solutions that feature different degrees of net rotation.

It is hence not clear to which equilibrium solution our adjustment scheme would converge when only the density response is optimized. This ambiguity can be lifted by making a selection for the desired type of solution one wants to obtain, and to suitably incorporate this constraint in the merit

function. For example, one may request to obtain an anisotropic solution with a certain prescribed ratio of radial and tangential velocity dispersions. We can then augment our density based merit function with further conditions that enforce this velocity structure.

A second problem, of perhaps somewhat lesser importance, is the possibility of overfitting individual particle velocities. In the continuum limit of a collisionless system, individual particles are completed uncorrelated from each other. An iterative optimization approach will however always adjust a particle's velocity given the current realization for positions and velocities of all other particles. This can in principle introduce undesired correlations between particles. Related to this, one may easily end up in an unfavorable local minimum of the merit function. We largely eliminate this effect by using a special optimization strategy in which new values for the velocity of a given particle are not searched in the vicinity of the current velocity, but rather globally in a random fashion, completely independent of the particle's current velocity.

In the following, we first discuss our formalism for determining the density response of a particular realization and for measuring its goodness of fit through a merit function. We then extend the discussion to merit functions for the velocity moments, and present our approach for optimizing both of them concurrently.

2.2.1 Density merit function

Consider a collisionless N-body system with N particles, initial coordinates $\hat{\mathbf{x}}_i$, and initial velocities $\hat{\mathbf{v}}_i$. We assume that an initial density distribution $\rho_0(\mathbf{x})$ is given, which can be used to create a realization of the coordinates $\hat{\mathbf{x}}_i$ by random sampling. Determining the $\hat{\mathbf{v}}_i$ is more complicated, however; we want to initialize them such that a stationary solution of the collisionless Boltzmann equation is obtained where the particles move in a given, stationary gravitational potential $\Phi(\mathbf{x})$. In other words, the collection of particles should move such that the density field they create is time invariant and

identical to the initial density distribution. If this is achieved, the gravitational field can then also be chosen self-consistently as the one created by the mass distribution itself (plus additional contributions by other mass distributions, if desired), such that one obtains a stationary self-gravitating solution of the Poisson-Vlasov system.

The density field created by the particles of our system can be formally expressed through a superposition of Dirac delta functions:

$$\rho(\mathbf{x}, t, \hat{\mathbf{v}}_1, \ldots, \hat{\mathbf{v}}_N) = \sum_{i=1}^{N} m_p \, \delta(\mathbf{x}_{\mathrm{orbit}}(\hat{\mathbf{x}}_i, \hat{\mathbf{v}}_i, t) - \mathbf{x}), \qquad (2.1)$$

where the function $\mathbf{x}_{\mathrm{orbit}}(\hat{\mathbf{x}}', \hat{\mathbf{v}}', t)$ describes the time-dependent orbit of a particle starting in the phase-space point $(\hat{\mathbf{x}}', \hat{\mathbf{v}}')$. Note that in the expression for the density field we have explicitly retained the dependence on the initial values of the particles velocities (which we have yet to determine), whereas the initial positions can be viewed as fixed parameters.

Next, we define the time-averaged density response for the chosen initial velocities as

$$\bar{\rho}(\mathbf{x}, \hat{\mathbf{v}}_1, \ldots, \hat{\mathbf{v}}_N) = \lim_{T \to \infty} \frac{1}{T} \int_0^T \rho(\mathbf{x}, t, \hat{\mathbf{v}}_1, \ldots, \hat{\mathbf{v}}_N) \, \mathrm{d}t. \qquad (2.2)$$

The best steady-state for the system can be defined as the smallest possible difference between the time-averaged density and the initial density field. To this end, we introduce an objective function

$$S(\hat{\mathbf{v}}_1, \ldots, \hat{\mathbf{v}}_N) = \int | \bar{\rho}(\mathbf{x}, \hat{\mathbf{v}}_1, \ldots, \hat{\mathbf{v}}_N) - \rho_0(\mathbf{x}) | \, \mathrm{d}\mathbf{x}, \qquad (2.3)$$

which measures the L_1-norm of the difference between the two fields. The linear weighting of the mass difference at a given location is motivated by the source term in Poisson's equation, which is ultimately what we want to keep constant as much as possible in a steady state to avoid potential fluctuations that can modify particle energies.

Therefore, the task to construct a best possible steady-state that has a given density distribution is primarily about finding the $\hat{\mathbf{v}}_i$ such that the

difference S defined by equation (2.3) reaches a minimum. Note that this can in principle be viewed as a high-dimensional minimization problem with respect to the initial velocities. This could, for example, be tackled with the method of steepest decent. A direct adjustment of the velocities to minimize the function S is indeed the central idea we pursue in this chapter, yielding a novel scheme for constructing equilibrium solutions. There are however a number of obstacles that make such a minimization non-trivial.

First of all, the problem needs to be somehow discretized, otherwise we cannot meaningfully define a density field for a finite number of particles. We will deal with this aspect in the remainder of this subsection. A further conceptual problem, namely the non-uniqueness of the obtained solutions, needs to be addressed as well.

Let's assume we discretize the space covered by our system in terms of cells of volume V_j, indexed by j. The cells cover the volume but they do not need to be of the same size (e.g. we may choose to use adaptive logarithmic grids, as we do in practice). The merit function (2.3) can now be written as

$$S(\hat{\mathbf{v}}_1, \ldots, \hat{\mathbf{v}}_N) = \sum_j \left| \overline{M}_j(\hat{\mathbf{v}}_1, \ldots, \hat{\mathbf{v}}_N) - M_j^0 \right|, \tag{2.4}$$

where $\overline{M}_j$ and M_j^0 give the masses of the time-averaged and the initial density field that fall into cell j, respectively. To determine $\overline{M}_j$, we add the time-averaged contributions of the orbits of all particles to this spatial cell. Since the trajectories of different points require different times to saturate their impact on the common averaged density, it is computationally more efficient to follow their orbits over individually chosen time scales T_i. We can then write

$$\overline{M}_j = \sum_{i=1}^N \int_{\text{cell } j} d\mathbf{x} \int \frac{dt}{T_i} \, m_p \, \delta(\mathbf{x}_{\text{orbit}}(\hat{\mathbf{x}}_i, \hat{\mathbf{v}}_i, t) - \mathbf{x}). \tag{2.5}$$

This reduces to

$$\overline{M}_j = \sum_{i=1}^N m_p \, \tau_j^{\text{orbit}}(\hat{\mathbf{x}}_i, \hat{\mathbf{v}}_i), \tag{2.6}$$

where $\tau_j^{\mathrm{orbit}}(\hat{\mathbf{x}}_i, \hat{\mathbf{v}}_i)$ gives the fraction of time an orbit starting in the given phase-space point spends in cell j. The expected mass in the cell, M_j^0, is simply given by

$$M_j^0 = \int_{\mathrm{cell}\ j} \rho_0(\mathbf{x})\mathrm{d}\mathbf{x}. \tag{2.7}$$

We note that the above equations correspond to so-called nearest grid point assignment of the current position of a particle to the mesh. One can replace this with a higher-order assignment scheme if desired, with the simplest possibility being clouds-in-cell assignment[1].

A more important question concerns the choice of the spatial binning scheme. There should be enough bins to resolve all relevant detail of the density distribution, but the Poisson noise affecting $\overline{M}_j$ due to the finite number of particles we use clearly limits the minimum size of a bin that is reasonable. In order to make the noise in each bin comparable, it is advantageous to choose the bins sizes such that they contain roughly equal mass. We follow this strategy by adopting a hierarchical adaptive binning scheme combined with a logarithmic grid. We will describe this approach in full detail in Section 2.4.1.

2.2.2 Velocity dispersion merit functions

As we discussed earlier, the requirement of a stationary density field does not in general imply a unique solution for the distribution function. For example, in an axisymmetric system, it would always be possible to flip the signs of the azimuthal velocities to generate, e.g., a system where all particles orbit around the z-axis with positive L_z, or with negative L_z, or with any desired mixture of the two. It is hence unclear in which minimum one ends up when one would try to directly minimize S with respect to the $\hat{\mathbf{v}}_i$.

In order to lift this ambiguity and make the solution more well defined, we need to add additional constraints that drastically reduce the accept-

---

[1] We actually use the latter in our implementation, even though the improvement relative to nearest grid point assignment is here minor.

able solution space. We do this by invoking symmetry assumptions about the velocity structure of the system. This then allows solving the Jeans equations for the second velocity moments, which we impose as a further optimization constraint.

In general, there are three first moments and three second (reduced) moments of the velocity distribution function at every point. We will here focus on axisymmetric systems and employ cylindrical coordinates (R, ϕ, z). In a stationary system, we always have $\langle v_R \rangle = 0$ and $\langle v_z \rangle = 0$. We shall now assume that as part of specifying our desired target system, we provide enough assumptions such that the three dispersions and the azimuthal streaming can be calculated everywhere, i.e. $\sigma_R^2 = \langle v_R^2 \rangle$, $\sigma_\phi^2 = \langle (v_\phi - \overline{v_\phi})^2 \rangle$, $\sigma_z^2 = \langle v_z^2 \rangle$, and the mean azimuthal streaming $\overline{v_\phi} = \langle v_\phi \rangle$ can be considered to be known as a function of (R, z). How we compute these quantities in practice for different cases will be discussed in Section 2.3.

In order to impose these velocity moments as additional constraints on the initial velocities $\hat{\mathbf{v}}_i$, we again consider spatial bins indexed by j, allowing us to estimate the actual (initial) velocity dispersions of our particular realization. For example, the average radial dispersion in bin j is given by

$$\left[\sigma_R^2\right]_j^{\text{actual}} = \frac{1}{M_j} \sum_{\hat{\mathbf{x}}_i \text{ in cell } j} m_p (\hat{\mathbf{v}}_i \cdot \mathbf{e}_R^{(i)})^2. \tag{2.8}$$

The normalization factor

$$M_j = \sum_{\hat{\mathbf{x}}_i \text{ in cell } j} m_p \tag{2.9}$$

is simply equal to the mass of the initial realization that falls into the spatial bin. The vector $\mathbf{e}_R^{(i)}$ is the radial unit vector at the position of particle i. The expected target velocity dispersion in the bin is given by

$$\left[\sigma_R^2\right]_j^{\text{target}} = \frac{1}{M_j^0} \int_{\text{cell } j} \rho_0(\mathbf{x})\, \sigma^2(\mathbf{x})\, \mathrm{d}\mathbf{x}, \tag{2.10}$$

where M_j^0 is the mass expected in the continuum in the cell. We may then define a merit function that measures the deviation of the actually realized

velocity dispersion relative to the target value. To this end we adopt

$$Q_R = \sum_j \frac{|\,[\sigma_R^2]_j^{\text{actual}} - [\sigma_R^2]_j^{\text{target}}\,|}{[\sigma_R^2]_j^{\text{target}}}. \tag{2.11}$$

Similarly, we define merit functions Q_z and Q_ϕ for the dispersions in the z-direction and the azimuthal ϕ-direction, respectively. In the azimuthal direction, we actually measure and constrain both the full second moment $\langle v_\phi^2 \rangle$ and the dispersion relative to the mean streaming σ_ϕ^2. Combined, this then also constrains the mean streaming itself.

As for the density case, we need to adopt a suitable discretization for the spatial bins. In order to avoid biases due to different noise levels in the bins, we adopt bins in which the mass per bin (or, equivalently, the number of particles) is roughly constant, a situation we realize with the help of a hierarchical adaptive binning scheme. We note that it is possible to employ the same spatial bins as used for the density merit function, but if desired the corresponding target value for the mass per bin can also be chosen differently.

2.2.3 Optimization procedure

Our goal is to iteratively adjust the $\hat{\mathbf{v}}_i$ such that S, Q_R, Q_ϕ and Q_z are simultaneously minimized. We do this by combining these quantities into a single goodness-of-fit parameter,

$$S_{\text{global}} = S + \chi(Q_R + Q_\phi + Q_z), \tag{2.12}$$

where the constant χ is adjusted such that S and $Q_R + Q_\phi + Q_z$ are of the same magnitude and have the same units. In other words, we give equal weight to the density and velocity constraints.

The function $S_{\text{global}}(\hat{\mathbf{v}}_1, \ldots, \hat{\mathbf{v}}_N)$ depends only on the initial particle velocities. Hence we are formally charged with the task to find its minimum in the high-dimensional space of all the $3N$ velocity components. Trying to find this minimum is a computationally rather tricky problem, because

the function will feature a large number of local minima in which a direct search may easily get stuck. Also, the function is non-linear and expensive to evaluate – calculating S involves orbit integrations of a large number of particles over a long time interval. Even if a single force calculation is comparatively cheap due to the static potential, the cumulative CPU cost can become demanding, especially since we are not dealing just with a single particle but rather with a (potentially quite large) particle collection of size N.

Nevertheless, it is still possible to estimate the local gradient of S_{global} with respect to the $\hat{\mathbf{v}}_i$ and then to move in the direction of steepest decent by simultaneously modifying all velocities in the direction opposite to the gradient. But finding a local minimum in this way will still be very hard (we have tried); typically, one will instead overshoot in at least one of the many dimensions of the problem.

Another consideration also argues against this brute force approach. Physically, we expect that the particles should be completely uncorrelated in proper collisionless initial conditions. Directly minimizing S_{global} simultaneously with respect to all velocities invokes the danger of 'overfitting', where a low value of the merit function is obtained through the introduction of velocity correlations in the specific N-body realization of the system.

Our solution to these problems involves two components. First, we serialize the minimization procedure, i.e. we always pick only one particle randomly, and then optimize its velocity such that S_{global} is reduced. Second, we do not actually try to adjust the velocity of the single particle such that S_{global} is necessarily minimized, as may be done by using the result of a line search along a single parameter. Rather, we simply randomly pick a new guess for the particle's velocity and (re)evaluate the merit function for this choice. If the proposed velocity improves the fit, we retain it as the new velocity of the particle, otherwise we simply keep the particle's old velocity and proceed with the next particle. This is simply repeated until the fit cannot be improved significantly any more. We note that this ap-

proach bears some resemblance to Monte Carlo Markov Chain techniques, except that we are here trying to find a global optimum rather than exploring a likelihood surface where one also moves occasionally away from the optimum with a certain probability.

The distribution from which one draws the trial velocities is in principle arbitrary, provided it is broad enough to sample all allowed velocities. However, it is highly advantageous to make it close to the target distribution function, because in this case the convergence speed can be expected to be particularly rapid (just as in MCMC). In our case, we can simply use Gaussians for that, as we already have the second moments in hand based on our Jeans solutions and the distribution function will in most cases resemble a Gaussian locally, so this should facilitate rapid convergence. Note that every new trial velocity we pick is completely independent of the previous value, as well as of the velocities of all other particles. This helps to minimize correlations between different particles in the created initial conditions, and it prevents to get easily stuck in a local minimum. Nevertheless, velocity correlations are not completely absent, because the acceptance decision for the velocity of a particular particle still depends on the discrete spectrum of velocities realized at this instant for all the other particles. But as our results show, any present residual correlations do not seem to negatively impact the quality of the created initial conditions.

In practice, we choose to process all particles in a random order. In each pass over the particles, we pick for a given particle one of its three principal coordinate directions and draw a random trial value for the corresponding velocity component. We note that the evaluation of S_{global} can be significantly accelerated if only one particle is varied. In this case, only the summed orbital response of all particles needs to stored, without requiring storage of all the responses individually. Evaluating S_{global} for a changed velocity of one particle then boils down to calculating the orbit response for this particle twice, both for the old and new velocities. The differential between the two results can then be appropriately added to the global

response to assess the change in S_{global}.

2.3   Velocity constraints

As discussed above, a problematic aspect of optimizing only a density merit function is that it is ambiguous to which solution this will converge. Recall that for a given density distribution there will in general be a vast number of possible distribution functions. The iterative optimization will yield a particular realization of one of these distribution functions, and this solution might depend on the initial velocity guesses one has used at the beginning. In order to make the solution well defined, we need to impose additional constraints that reflect the desired properties of the specific solution one is looking for. We do this in terms of second moments of the velocity distribution and by forcing the system to converge to a solution that features these moments. The moments themselves are calculated from the Jeans equations. Different possibilities for a specification of the desired properties of the target system exist.

2.3.1   Spherically symmetric distribution functions

If the density structure is spherically symmetric and the velocity distribution function depends at most on the magnitude of the angular momentum, we can make use of the spherically symmetric Jeans equation for the second radial velocity moment,

$$\frac{\partial(\rho\sigma_r^2)}{\partial r} + 2\frac{\beta\rho\sigma_r^2}{r} + \rho\frac{\partial\Phi}{\partial r} = 0. \tag{2.13}$$

Here $\sigma_r^2 = \langle v_r^2 \rangle$ is the radial dispersion. The velocity distribution functions in the transverse directions at any given position need not be equal to that in the radial direction, but we have $\sigma_\theta = \sigma_\phi$ due to the assumed symmetry. The degree of radial–tangential anisotropy is usually measured in terms of

$$\beta = 1 - \frac{\sigma_t^2}{2\sigma_r^2} = 1 - \frac{\sigma_\phi^2}{\sigma_r^2}, \tag{2.14}$$

where $\sigma_t^2 = \sigma_\theta^2 + \sigma_\phi^2 = 2\sigma_\theta^2$ measures the total tangential dispersion, and due to spherical symmetry we have $\sigma_\theta = \sigma_\phi$. If the distribution is isotropic, we have $\beta = 0$. If the orbits are biased towards radial motions we have $\beta > 0$, while for $\beta < 0$ they are preferentially tangential.

For given $\rho(r)$ and prescribed $\beta(r)$, and thanks to the purely radial dependence, equation (2.13) becomes an ordinary differential equation for $\rho\sigma_r^2$ which can be readily integrated using the boundary condition $\rho\sigma_r^2 = 0$ for large radii. Dividing the solution by the density then yields the dispersion $\sigma_r^2(r)$ as a function of radius, and from it we also obtain $\sigma_\theta^2(r) = (1-\beta)\sigma_r^2(r)$.

We note that we may choose β to be a function of radius, as suggested by the structure measured for cosmological dark matter halos. Hansen & Moore (2006) found that the local anisotropy of dark matter halos correlates well with the logarithmic slope

$$\alpha = \frac{d\ln\rho}{d\ln r} \tag{2.15}$$

of the density profile. Their numerical results are well fit by the relation

$$\beta(r) = -0.15 - 0.2\alpha, \tag{2.16}$$

which we adopt as an additional option in our IC code. This implies nearly isotropic orbits in the center of a Hernquist or NFW halo, and a growing preference for more radial dispersion as a function of distance.

A particularly simple choice for $\beta(r)$ is the isotropic case, $\beta = 0$, where the velocity distribution function is independent of direction at every point. In this ergodic case, the distribution function depends only on energy. Hernquist (1990) constructed such a solution for a density profile of the form $\rho(r) \propto r^{-1}(r+a)^{-3}$, which is cosmologically particularly relevant as it has a shape similar to the NFW density profile (Navarro et al. 1997) measured for relaxed halos in cold dark matter structure formation simulations. This makes the isotropic Hernquist model a particularly useful analytic distribution, and we will also use it here to verify our procedures. We note however that a yet more realistic model would be one with a radially vary-

ing anisotropy $\beta(r)$. No analytic distribution functions are known for this case, but such models can be readily constructed with our new method.

2.3.2 Axisymmetric systems with two integrals of motion

For axisymmetric systems, the angular momentum L_z around the z-axis is a conserved quantity for all orbits, hence we expect the distribution function to depend on L_z besides energy E. In general, there can be a third integral of motion, I_3, which is however often not easy to identify and therefore considered "non-classical". If one disregards I_3 and assumes that the distribution function is only a function of E and L_z, then the situation simplifies considerably, as one can then infer that all mixed moments of the velocity distribution vanish (i.e. $\langle \sigma_R \sigma_z \rangle = 0$). In this case the axisymmetric Jeans equations simplify considerably and can be comparatively easily solved.

With two integrals of motion, the non-trivial axisymmetric Jeans equations become:

$$\frac{\partial(\rho\sigma_z^2)}{\partial z} + \rho\frac{\partial\Phi}{\partial z} = 0, \tag{2.17}$$

and

$$\langle v_\phi^2 \rangle = \sigma_R^2 + \frac{R}{\rho}\frac{\partial(\rho\sigma_R^2)}{\partial R} + R\frac{\partial\Phi}{\partial R}. \tag{2.18}$$

The mean streaming motions in the radial and vertical directions vanish, $\langle v_R \rangle = \langle v_z \rangle = 0$ (but not necessarily in the azimuthal direction), and importantly, the radial and vertical dispersions are equal everywhere, $\sigma_R^2 = \sigma_z^2$.

This in particular means that the density distribution fully specifies the vertical and radial dispersions in the meridional plane (R, z). They can be explicitly calculated as

$$\sigma_R^2 = \sigma_z^2(R, z) = \frac{1}{\rho(R, z)}\int_z^\infty \rho(z', R)\frac{\partial\Phi}{\partial z}(R, z')\,\mathrm{d}z'. \tag{2.19}$$

Once these dispersions are known, we can now determine the second moment $\langle v_\phi^2 \rangle$ of the azimuthal motion from the radial Jeans equation (2.18). However the mean streaming $\langle v_\phi \rangle$ in the azimuthal direction is not specified

by the Jeans equations. Indeed, $\langle v_\phi \rangle$ does not have to be zero if there is net rotation. For any given solution with non-zero $\langle v_\phi \rangle$, one can readily construct new, equally valid equilibrium solutions, for example by reversing all or a fraction of the particles' ϕ-motions. It is hence clear that the requirement of axisymmetry does not specify $\langle v_\phi \rangle$. In fact, we are (within limits) free to set this.

We adopt the parameterization

$$\langle v_\phi \rangle^2 = k^2 \left[\langle v_\phi^2 \rangle - \sigma_R^2 \right] \tag{2.20}$$

suggested by Satoh (1980) to specify the mean streaming. For the interesting choice $k = 1$, we obtain for the azimuthal dispersion

$$\sigma_\phi^2 \equiv \langle v_\phi^2 \rangle - \langle v_\phi \rangle^2 = \sigma_R^2 = \sigma_z^2, \tag{2.21}$$

i.e. σ_ϕ^2 is then equal to the radial and vertical dispersions. This defines the case of an isotropic rotator. But we may also adopt a lower or higher value for k, or even one with a spatial dependence, up to the maximum allowed local value of

$$k_{\max}^2 = \frac{\langle v_\phi^2 \rangle}{\langle v_\phi^2 \rangle - \sigma_R^2}. \tag{2.22}$$

In k climbs up to this value, the azimuthal dispersion vanishes and we have $\langle v_\phi \rangle^2 = \langle v_\phi^2 \rangle$, corresponding to a system with the maximum possible angular momentum for a given density structure. In our GALIC code, we either choose a constant k or specify k in units of $k_{\max}$ when the case of a $f(E, L_z)$ distribution function is selected.

2.3.3   General systems with three integrals of motion

While simple disk models can be constructed as isotropic rotators, observations in the Milky Way at the Solar circle suggest that σ_R is not equal to σ_z. Rather, the two dispersions are related approximately by $\sigma_z \simeq 0.5 \, \sigma_R$ (Binney & Merrifield 1998). Even if the Milky Way can still be described well as an axisymmetric system, this already means that the distribution function
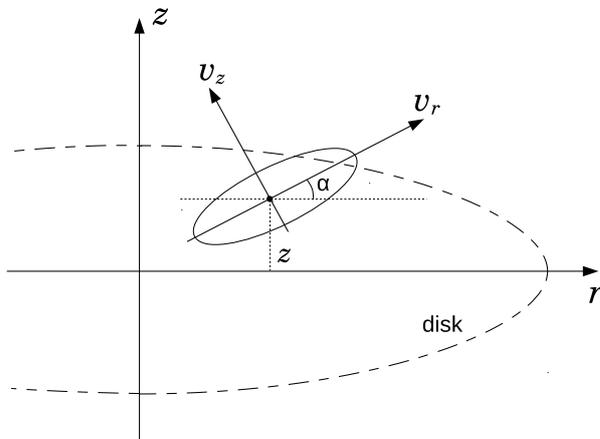
Figure 2.1: Sketch of the geometry adopted to describe the assumed tilt of the velocity ellipsoid in the $f(E, L_z, I_3)$ case.

is not only dependent on (E, L_z); instead, a third integral of motion must play an important role. In the outer parts of the disk, this is approximately given by the energy of the vertical motion in the disks potential, but in the inner parts of the disk this identification presumably becomes a poor approximation.

Another interesting observational fact is that the velocity ellipsoid above the disk mid plane is not aligned with the coordinate plane; instead, it appears tilted (i.e. $\langle v_R v_z \rangle \neq 0$) towards the centre of the system. Using RAVE velocity data, Siebert et al. (2008a) and Binney et al. (2014) quantified the tilt at $1\,\mathrm{kpc}$ above the disk to be around $\alpha = 7^o$ (see sketch of Fig. 2.1). We are hence forced to apply the general axisymmetric Jeans equations, which take the form:

$$\frac{\partial(\rho\sigma_z^2)}{\partial z} + \rho\frac{\partial\Phi}{\partial z} + \frac{1}{R}\frac{\partial(R\rho\,\langle v_R v_z\rangle)}{\partial R} = 0, \tag{2.23}$$

$$\langle v_\phi^2 \rangle = \sigma_R^2 + \frac{R}{\rho}\frac{\partial(\rho\sigma_R^2)}{\partial R} + R\frac{\partial\Phi}{\partial R} + \frac{R}{\rho}\frac{\partial(\rho\,\langle v_R v_z\rangle)}{\partial R}, \tag{2.24}$$

$$\frac{\partial(\rho\,\langle v_z v_\phi\rangle)}{\partial z} + \frac{1}{R^2}\frac{\partial(R^2\rho\,\langle v_R v_\phi\rangle)}{\partial R} = 0. \tag{2.25}$$

Velocity constraints

This system of equations is significantly under-specified and additional assumptions are needed for closure. We shall assume that the velocity ellipsoid is not tilted in the ϕ-direction, hence $\langle v_z v_\phi \rangle = \langle v_R v_\phi \rangle = 0$. This eliminates the third equation. However, we need to retain a tilt in the meriodonal plane, as encoded by $\langle v_R v_z \rangle$. If α is the local angle between the velocity ellipsoid and the R-axis, this mixed moment can be expressed in terms of the radial and vertical moments, i.e. we have

$$\langle v_R v_z \rangle = \frac{1}{2} \tan(2\alpha) \left[\sigma_R^2 - \sigma_z^2 \right].\tag{2.26}$$

For reference, the dispersions in the rotated coordinate frame (R', z') are given by

$$\left\langle v_{R'}^2 \right\rangle = \sigma_R^2 \cos^2(\alpha) + \langle v_R v_z \rangle \sin(2\alpha) + \sigma_z^2 \sin^2(\alpha),\tag{2.27}$$

$$\left\langle v_{z'}^2 \right\rangle = \sigma_R^2 \sin^2(\alpha) - \langle v_R v_z \rangle \sin(2\alpha) + \sigma_z^2 \cos^2(\alpha).\tag{2.28}$$

The tilt angle is the one for which $\langle v_{R'} v_{z'} \rangle = 0$, by construction.

Interestingly, the tilt observed for the Galaxy at the Solar circle is consistent with the velocity ellipsoid pointing approximately to the center of the Galaxy; the most recent determination by Binney et al. (2014) gives $\alpha \sim 0.8 \arctan(z/R)$. We here assume for definiteness that this alignment is perfect and holds throughout the system, in which case the angle α is simply given by

$$\tan \alpha = \frac{z}{R}.\tag{2.29}$$

Specifying the orientation of the velocity ellipsoid in this way has the nice property of naturally producing a spherically symmetric orientation close to the galactic centre, i.e. the 'disk regime' seamlessly transitions to a 'bulge regime'. Far out in a thin disk, the velocity ellipsoid will align with the coordinate axes, while near to the centre the situation becomes closer to that in a spherically symmetric case with a radial alignment, which seems plausible.

Prescribing the tilt angle is not yet enough to solve equations (2.23) and (2.24), because they involve four unknowns. An additional assumption is

required. To this end, we adopt a prescribed relation between the radial and vertical dispersions in the tilted velocity ellipsoids, namely

$$\left\langle v_{R'}^2 \right\rangle = f_R \left\langle v_{z'}^2 \right\rangle, \tag{2.30}$$

where f_R is a factor specifying the anisotropy between radial and transverse motions. For the disk of a Milky Way like galaxy, we would expect $f_R \simeq 2$ at the Solar circle, but little is own about a potential radial variation of this value. We also note in passing that the Toomre stability criterion depends sensitively on σ_R, so invoking values $f_R > 1$ is one way of stabilizing a stellar disk of given thickness against axisymmetric perturbations. For simplicity, we shall assume a spatially constant value for f_R in the disk, but note that our techniques could be easily generalized to include a radial or vertical variation of this factor.

Given the above model for the dispersions, we can now express the mixed moment $\langle v_R v_z \rangle$ through the vertical dispersion, namely

$$\langle v_R v_z \rangle = h\,\sigma_z^2, \tag{2.31}$$

where the function $h = h(R, z)$ is given by

$$h = \frac{(f-1)\tan(2\alpha)}{2\cos^2(\alpha) - 2f\sin^2(\alpha) + (1+f)\sin(2\alpha)\tan(2\alpha)}, \tag{2.32}$$

and the shortcut $f = f_R$ is understood. The Jeans equation (2.23) now becomes an inhomogenous first order partial differential equation (PDE) for σ_z^2. Defining $q \equiv \rho\sigma_z^2$, the relevant equation takes the form

$$\frac{\partial q}{\partial z} + \frac{\partial(hq)}{\partial R} + \frac{hq}{R} + \rho\frac{\partial \Phi}{\partial z} = 0. \tag{2.33}$$

We can solve this PDE numerically with the methods of lines by discretizing in R and replacing the spatial R-derivative with a finite difference approximation. We can then integrate the resulting system of coupled ordinary differential equations along the z-direction, starting at $z \simeq \infty$ and ending up at $z = 0$. The initial condition is $q(R, z = \infty) = 0$, augmented with the

boundary condition $q(R = \infty, z) = 0$. For numerical stability, one needs to take care that an upwind finite difference estimate for the R-derivative is used. Note also that the hq/R term is not singular for $R = 0$, because $h/R \rightarrow (f - 1)/f$ for $R \rightarrow 0$.

Having obtained a solution for $q(R, z)$, we then readily have σ_z, σ_R^2, $\langle v_R v_z \rangle$, as well as $\langle v_{R'}^2 \rangle$ and $\langle v_{z'}^2 \rangle$ throughout the meridional plane. Similar as with the axisymmetric $f(E, L_z)$ case, we still have the freedom to choose a streaming velocity in the ϕ-direction, except that now we have to use equation (2.24) to infer the corresponding dispersion available in the azimuthal direction. We continue to use the parametrization of equation (2.21) for the azimuthal streaming. For the case $k = 1$, we then get $\sigma_\phi^2 = \sigma_R^2$ in the mid-plane.

We note that one can also obtain from epicycle theory a statement about the relation between $\langle (v_\phi - v_c)^2 \rangle$ and σ_R^2, valid for small radial dispersions σ_R, namely

$$\frac{\langle (v_\phi - v_c)^2 \rangle}{\sigma_R^2} \simeq \frac{1}{\gamma^2} \tag{2.34}$$

where $\gamma = 2\Omega/\kappa$. $\Omega^2 = \frac{1}{R}\frac{\partial \Phi}{\partial R}$ is the circular orbit frequency, and

$$\kappa^2 = R\frac{\mathrm{d}\Omega^2}{\mathrm{d}R} + 4\Omega^2 \tag{2.35}$$

is the epicycle frequency. Typically we have $1/\gamma^2 \simeq 0.5$. We note that equation (2.34) is only reliable for very cold thin disks, with $\sigma_R \ll v_c$ (see Binney & Tremaine 2008). Interestingly, combined with equation (2.24), the epicycle approximation gives the azimuthal streaming (and hence also the axisymmetric drift) in the equatorial plane as

$$\langle v_\phi \rangle = v_c + \frac{\sigma_R^2}{2v_c}\left(\frac{\partial \ln(\rho \sigma_R^2)}{\partial \ln R} + \frac{\gamma^2 - 1}{\gamma^2}\right). \tag{2.36}$$

On the other hand, we obtain from equation (2.20) the following expression for the streaming velocity to leading order in σ_R/v_c:

$$\langle v_\phi \rangle = k\, v_c + \frac{k\, \sigma_R^2}{2v_c}\left(\frac{\partial \ln(\rho \sigma_R^2)}{\partial \ln R}\right). \tag{2.37}$$

Consistency with the epicycle approximation hence requires $k = 1$ for thin cold disks. The residual difference grows for large σ_R/v_c, but note that in this limit the epicycle approximation becomes inaccurate anyway.

2.4 Implementation details

2.4.1 Adaptive logarithmic binning

To account for the typical power-law growth of the density towards the center in self-gravitating systems, we generally employ logarithmic grids. For the sake of simplicity, we restrict ourselves to axisymmetric systems in this chapter, and also assume mirror symmetry with respect to the $z = 0$ plane. Adopting cylindrical coordinates, this means we only have to cover the positive quadrant in the (R, z)-plane. We assume that the mass distribution is fully contained inside a cube of side-length $2d_{\mathrm{max}}$, i.e. our mesh needs to cover the region $0 \leq R < d_{\mathrm{max}}$ and $0 \leq z < d_{\mathrm{max}}$. If we use $N_{\mathrm{bin}} = 2^l$ bins per dimension, and require that the width of the bins grows by a constant factor f from bin to bin, the borders of the bins can be written as

$$R_i = d_{\mathrm{base}}(f^i - 1), \tag{2.38}$$

$$z_j = d_{\mathrm{base}}(f^j - 1), \tag{2.39}$$

with $i, j \in [0, 1, \ldots, N_{\mathrm{bin}}]$. The bin (i, j) (with $0 \leq i, j < N_{\mathrm{bin}}$) then covers $[R_i, R_{i+1}] \times [z_j, z_{j+1}]$ in the (R, z)-plane and has volume

$$V_{ij} = 2\pi(R_{i+1}^2 - R_i^2)(z_{j+1} - z_j). \tag{2.40}$$

To cover the full volume, d_{base} and f need to be chosen such that

$$d_{\mathrm{max}} = d_{\mathrm{base}}(f^{N_{\mathrm{bin}}} - 1). \tag{2.41}$$

This still leaves room for one additional constraint to fully specify the quantities d_{base} and f. We typically address this by requiring that the first bin, bounded by $R_1 = d_{\mathrm{base}}(f - 1)$, encloses a small prescribed fraction of the
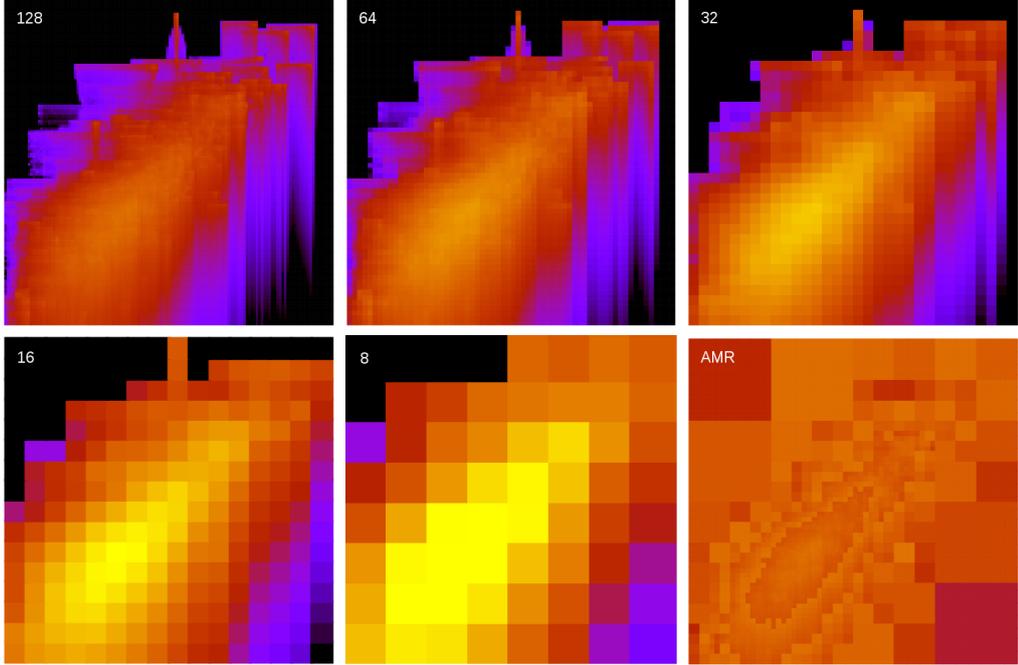
Figure 2.2: Density response of our hierarchical binning scheme. The top left panel shows the average orbit response as recorded on the finest grid used (which is a logarithmic grid with 256^2 pixels). The next four panels show coarsened representations of this field, with resolutions of 64^2 to 8^2 pixels (this continues to even coarser meshes that are not shown). Finally, the bottom right panel shows the variable resolution response that is actually used to compare with the target density distribution, based on bins containing roughly equal mass.

total mass of the system (e.g. 10^{-6}), such that the central region is still well resolved by the grid.

As we discussed earlier, the objective functions assessing the density response and the initial velocity distribution work best if the spatial bins are chosen such that they contain approximately constant mass. We realize such a scheme by first constructing the mass response on a relatively fine grid, given by the level l_{max}. We then recursively construct a set of coarsened meshes on levels $l_{max} - 1$, $l_{max} - 2$, ..., 1, 0, until there is only one cell left covering the whole quadrant. Computing the mass response of one of the grid cells of a coarsened mesh is done recursively by summing over the

corresponding four cells in the finer mesh one level higher. Similarly for the velocity dispersion fields.

Evaluating the objective functions then proceeds with a recursive algorithm that walks the tree of nested mesh cells. Beginning at the 'root node', a mesh cell is included in the sum if it contains less than a certain threshold mass or if it is already a cell of the finest level. Otherwise, the mesh cell is 'opened', and its four daughter cells are considered in turn as candidates for being included in the sum. This procedure automatically selects a close to optimum set of cells of different sizes. Note that the union of the cells that enter the sum form a space-covering tessellation, i.e. each point in the (R, z)-plane is accounted for exactly once.

In Figure 2.2 we show an example for the mass response grid of a set of orbits in the (R, z)-plane for a mesh with $N_{\mathrm{bin}} = 2^8$ cells per dimension, together with the hierarchy of the next four coarsened representations at higher levels. The final panel on the bottom right shows a mixed image of variable resolution, indicating what is effectively used in the adaptively calculated sum that defines the merit function.

2.4.2 Orbit integration

In order to efficiently and accurately compute orbits of particles for arbitrary mass distributions, we produce a look-up table of the gravitational potential and its derivatives $\partial \Phi / \partial R$ and $\partial \Phi / \partial z$ in the positive quadrant of the (R, z)-plane. Due to the axisymmetry we assume, this is sufficient to obtain the forces and the potential everywhere through a table look-up. We use a fine logarithmic grid in R and z and bilinear interpolation for the look-up table.

In order to allow a computation of the forces for arbitrary density distributions without analytic solutions of Poisson's equation, we numerically evaluate the potential and forces on the fine grid based on randomly sampling the density distribution with a very large number of fiducial particles combined with a calculation of the forces and potentials with a gravitational tree algorithm. This procedure is very flexible and accurate. In order to

help reducing noise effects from the sampling to a negligible level, a large number of fiducial points is used, and additionally, we evaluate and average the field at a set of different azimuthal angles.

For the orbit integration of particles, which gives us $\mathbf{x}_{\mathrm{orbit}}(\hat{\mathbf{x}}_i, \hat{\mathbf{v}}_i, t)$, we use the leapfrog scheme with adaptive timestep based on the kick-drift-kick formulation. If $(\mathbf{x}^{(n)}, \mathbf{v}^{(n)})$ denote position and velocity after step n, then the update to the next step is obtained through

$$\mathbf{v}^{(n+1/2)} = \mathbf{v}^{(n)} + \mathbf{a}^{(n)} \Delta t_n / 2, \tag{2.42}$$

$$\mathbf{x}^{(n+1)} = \mathbf{x}^{(n)} + \mathbf{v}^{(n+1/2)} \Delta t_n, \tag{2.43}$$

$$\mathbf{v}^{(n+1)} = \mathbf{v}^{(n+1/2)} + \mathbf{a}^{(n+1)} \Delta t_n / 2. \tag{2.44}$$

We set the size Δt_n of the timestep of step n as

$$\Delta t_n = \min \left(\eta_{\mathrm{orbit}} \frac{V_{200}}{|\mathbf{a}^{(n)}|}, \eta_{\mathrm{mesh}} \frac{d_{\mathrm{cell}}^{(n)}}{|\mathbf{v}^{(n)}|} \right), \tag{2.45}$$

where V_{200} is the circular velocity of the halo of the constructed galaxy and $d_{\mathrm{cell}}^{(n)}$ is the dimension of the mesh cell at the particle's current location. The dimensionless coefficients η_{orbit} and η_{mesh} are meant to ensure an accurate integration of the orbit and a precise accounting of the time spent by the orbit in each of the bins used for recording the density response.

We select the integrated timespan T_i for each particle individually. To this end we use the circular velocity at the particle's initial position, and introduce a dimensionless factor η_{timespan} for scaling the circular orbital time at the local distance. Explicitly, we set

$$T_i = \eta_{\mathrm{timespan}} \frac{2\pi |\hat{\mathbf{x}}_i|}{v_{\mathrm{circ}}(\hat{\mathbf{x}}_i)}, \tag{2.46}$$

where $v_{\mathrm{circ}}(\hat{\mathbf{x}}_i) \equiv (|\hat{\mathbf{x}}_i| \, |\hat{\mathbf{a}}_i(\hat{\mathbf{x}}_i)|)^{1/2}$. We typically found $\eta_{\mathrm{timespan}} = 10.0$ to be sufficient, yielding an average number of about 15 orbits for the particles of a typical halo.

2.4.3 Optimization procedure

As discussed earlier, we in principle would like to optimize the particles sequentially. Unfortunately, this immediately poses a serious problem for any efficient parallelization. If we enforce strictly sequential iterative adjustments of the particle velocities (such that a subsequent evaluation of the merit function already takes the effects of a potential change of the previous particle's velocity fully into account), then the optimization can evidently not be done concurrently for several different particles.

However, we have found that in practice we still obtain good results if we allow a small fraction of all particles to be treated simultaneously, each remaining unaware of the changes in the other particles until these are 'committed' at the end of the concurrent phase. With this approach, we can exploit massive parallelism in the optimization procedure (as implemented in our GALIC code).

For definiteness, this practical aspect of our optimization scheme is controlled by a parameter f_{opt} which gives the fraction of particle orbits that can be set to new starting velocities without taking note of each other. Our default values for this parameter is $f_{\mathrm{opt}} = 0.001$, meaning that our code will process the particles in batches of size $f_{\mathrm{opt}} N_{\mathrm{part}}$ particles from a randomly shuffled list of all particles. In each batch, all the trial velocities are drawn and evaluated independently (hence this can be done in parallel), and only at the end the velocity updates are committed to the new global response of the system, affecting the next batch.

When a particle is selected for optimization, we first randomly select one of the three primary coordinate directions, and then replace the corresponding velocity component with one drawn from the corresponding Gaussian distribution. In this way, each of the optimizations effectively couples only to one of the velocity dispersion measures. We found this advantageous also for the following reason. To exclude any possibility that systematic binning effects might prefer orbits that start, for example, with positive v_R as opposed to negative v_R, we actually assess orbits by averaging the merit
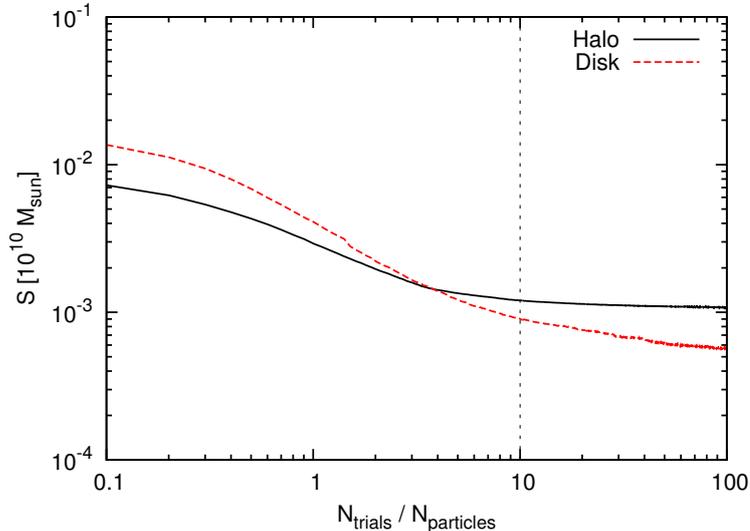
Figure 2.3: Decline of the merit function in a typical optimization run (actually the model D1 from Table 1) as a function of the number of attempted velocity adjustments in units of the particle number of the corresponding component. The solid line shows the result for the dark matter particles, while the dashed lines is for the disk particles.

functions for orbits both with v_R and $-v_R$ velocities, and likewise for the v_z velocities. This guarantees symmetry of the resulting velocity distribution functions in these two directions, and in particular, $\langle v_R \rangle = \langle v_z \rangle = 0$. However, in the ϕ-direction, this reversal trick is not indicated, both because here the symmetry of the binning procedure excludes the possibility of any such effects by construction (unlike for the R- and z-directions), and because in the ϕ-direction orbits with a reversed ϕ-velocity are not necessarily equally probable. We note that to ensure that all particles remain bound, we reject any trial velocity that is larger than $\eta_{max} v_{esc}$, where v_{esc} is the local escape velocity and $\eta_{max} = 0.9999$ is a parameter very close to 1.

In Figure 2.3, we show the decline of the value of the merit function as a function of the number of velocity optimizations that have been attempted by the code, in units of the total particle number, for a typical initial conditions model where the initial velocity guess were computed with

the moment-based method. We see that after ~ 3 optimization attempts for each particle, the initial convergence speed slows down significantly, and a stationary state in which no further improvement appears possible is reached after approximately ~ 5 optimizations. We find this is a quite typical behaviour in all of our models. Conservatively, we usually run our models to ~ 10 optimizations per particle.

2.4.4   Determination of the initial realization

There are only two functions that need to be provided for any desired density distribution that should be treated with our scheme. For each component of the system (i.e. halo, disk and/or bulge), one function needs to return the density of the component at a given point, the other must return a randomly sampled coordinate from the density field, i.e. the probability density of the corresponding point process must be proportional to the density field. Having these functions in hand, we can create the $\hat{\mathbf{x}}_i$ simply by randomly sampling each density component present in the target system. Also, we can create a (large) fiducial set of points for evaluating the force field to arbitrary precision with a tree algorithm. Finally, we can make use of the function returning the continuum density in solving the Jeans equations.

Since in our approach we anyway compute the second moments with the Jeans equations, we may as well initialize initial guesses for the particle velocities $\hat{\mathbf{v}}_i$ by drawing randomly from Gaussian distributions with the correct local dispersions. This corresponds to the frequently invoked approximation of adopting triaxial Gaussians for the local velocity distribution function, and since this is in most cases reasonably close to the correct distribution, it accelerates convergence. The iteration method is then in essence only responsible for determining the higher-order moments of the velocity distribution function.

2.4.5   Parallelization approach

Our C-code for creating stable initial conditions with the scheme described here, GALIC, has been fully parallelized for distributed memory machines using the message passing interface (MPI). For calculating the gravitational field in the (R, z)-plane, we let each MPI-task sample particles independently. The resulting particle set is then subjected to a domain decomposition, and a parallel distributed tree algorithm derived from the well-known GADGET simulation code (Springel et al. 2001b; Springel 2005a) is invoked to compute the force field.

For creating the particles of the actual initial conditions, we again let each MPI-task create a random, disjoint subset of the target particle set for each mass component. Then we let each MPI-task work independently and in parallel on the orbit optimizations associated with one batch of size $f_{\mathrm{opt}}N_{\mathrm{part}}$. The results are then interchanged and the sums over the orbit responses are updated accordingly, allowing the next cycle of optimizations to proceed. As a result, the scalability of our code is essentially perfect provided $f_{\mathrm{opt}}N_{\mathrm{part}}$ is substantially larger than N_{CPU}, otherwise work-load imbalances may become substantial as not all tasks could then be expected to have roughly equal amounts of work in each batch.

We have also made use of some of the I/O code from GADGET when writing the final initial conditions to disk. They can be stored in any of the three file formats supported by GADGET (including one in HDF5-format), thereby simplifying the subsequent use of the ICs with this simulation code, or the application of existing file format conversion tools from GADGET's format to other simulation codes. Finally, our use of parallel I/O routines also facilitates the creation of extremely large galaxy models, if desired.

2.5   Galaxy models

The approach outlined above is quite general and can be used with nearly arbitrary axisymmetric density profiles. For definiteness, we describe in

this section a specific set of parameterizations for dark matter halos, stellar disks and stellar bulges, which we shall use in our test galaxy models. These parameterizations follow models widely employed in the literature.

We usually model the dark matter density profile as a spherically symmetric halo with density

$$\rho_{\mathrm{dm}}(r) = \frac{M_{\mathrm{dm}}}{2\pi} \frac{a}{r(r+a)^3},$$

(2.47)

where a is the scale factor. Following Springel et al. (2005a), we can relate a to the concentration c of a corresponding NFW halo of mass $M_{200} = M_{\mathrm{dm}}$ such that the shape of the density profile in the inner regions is identical. The relation between a and c is then given by

$$a = \frac{r_{200}}{c} \sqrt{2[\ln(1+c) - c/(1+c)]},$$

(2.48)

where r_{200} and M_{200} are the virial radius and virial mass of the NFW halo, respectively.

We may also consider axisymmetric dark matter halos with either prolate or oblate distortions. For simplicity, we assume that the isodensity contours of the distorted shape are ellipses, effectively created by linearly distorting the spherical shape along the symmetry axis. If $s = a/c$ is the (radially constant) stretch factor, and $a = b$ and c are the axes of some isodensity ellipsoid, then a prolate halo has $c/a > 1$ (and hence $s < 1$), while an oblate halo has $c/a < 1$ (with $s > 1$). We can then define the density profile of the distorted halo as

$$\tilde{\rho}_{\mathrm{dm}}(R, z) \equiv s\,\rho_{\mathrm{dm}}(\sqrt{R^2 + s^2 z^2}),$$

(2.49)

which leaves the total mass invariant.

For the disk, we adopt in general a model with exponential radial scale length, and a sech^2-profile in the vertical direction. More specifically, the

3D disk density profile[2] is described by

$$\rho_\star(R, z) = \frac{M_\star}{4\pi z_0 \, h^2} \, \text{sech}^2 \left(\frac{z}{z_0} \right) \exp \left(-\frac{R}{h} \right).$$ (2.50)

The disk scale length h can either be set to a prescribed value, or calculated by assuming that the disk contains a certain fraction of the specific angular momentum of the halo (see, e.g., Mo et al. 1998). We assume a radially constant scale height z_0 of the disk, but this could be easily modified if desired. Usually, we parameterize z_0 in terms of the disk scale length, with typical disks lying in the range $z_0/h \sim 0.1 - 0.3$.

Finally, we model a stellar bulge (if present) with a Hernquist halo as well, using the profile

$$\rho_\text{b}(r) = \frac{M_\text{b}}{2\pi} \, \frac{b}{r(r + b)^3}.$$ (2.51)

The bulge scale length b is prescribed through a parameter that gives its size in units of the halo's scale length.

We specify both the bulge and disk masses as fractions m_d and m_b of the total mass, i.e. $M_\text{d} = m_\text{d} M_\text{tot}$ and $M_\text{b} = m_\text{b} M_\text{tot}$. This parameterization has previously been adopted in the study of Mo et al. (1998) on disc structure, as well as in some earlier work on compound disc galaxy models (e.g. Springel & White 1999; Springel et al. 2005a).

All the many reasonable combinations of the above components, together with the various velocity structures possible for them, produce a fairly large number of possibilities our code GALIC has to deal with. In particular, requiring that a galaxy model always needs to have a dark matter halo (either of spherical or oblate/prolate shape), that a disk can either be present or absent, and that a bulge is optional but may have different shapes if present (spherical, oblate, or prolate), we already arrive at 12 possible combinations of these three components. Of the corresponding models, only 2 have spherically symmetric potentials (namely either the model

---

[2]We note that equation 10 of Springel et al. (2005a) contains a typo in the form of an extraneous factor $1/2$ in the argument of the sech-function. All model calculations in that paper have however been done correctly, based on equation 28 of Springel & White (1999), which is what we adopt here too.

with just a spherical halo, or the model with spherical halo and a spherical bulge), allowing ergodic $f(E)$ distribution functions or $f(E, |\mathbf{L}|)$ models for them (besides axisymmetric $f(E, L_z)$ or $f(E, L_z, I_3)$ distribution functions possible for all the models). Allowing just for different combinations of these extra velocity structures, this means that the 2 density models really correspond to 6 possible variants. Similarly, the other 10 possible density models give rise to 54 possible velocity variants, so that we have of order 60 valid combinations of density model and associated velocity structures. Of course, many of these models really correspond to a continuum of further possibilities once the additional free parameters describing, for example, the degree of net rotation or the radial velocity anisotropy are used.

It is clear that we cannot present exhaustive tests of all these possibilities in this work. Rather, we instead focus on a representative selection of models which we list in Table 2.1. This sample of models covers a good fraction of the space of possible model variants of interest, hence our tests should give a good assessment of how well our techniques work in practice. We consider, in particular, models that contain only a dark matter halo (denoted as 'H1', 'H2', etc.), that feature a pure disk embedded in a halo (labeled 'D1', 'D2', etc.), that contain a bulge but no disk inside a halo ('B1', etc.), and finally, models that feature both a disk and a bulge ('M1', 'M2', etc.). In each of these four groups, we consider models with a variety of velocity structures, and/or different halo or bulge shapes. Detailed test results for the produced initial conditions will be discussed in the next section.

2.6 Test results

2.6.1 Models with a single halo component

The simplest of our models is 'H1', featuring a Hernquist dark matter halo with an isotropic velocity distribution. For definiteness, we adopt $v_{200} = 200\,\mathrm{km\,s^{-1}}$ and $c = 10$ to set the total mass and concentration of the halo,

Model	Components and their shape	Imposed velocity structure
H1	spherical dark halo	ergodic (i.e. isotropic Hernquist model)
H2	spherical dark halo	spherical anisotropy, radial orbits dominating, $\beta = 0.5$
H3	spherical dark halo	spherical anisotropy, tangential orbits dominating, $\beta = -1.0$
H4	spherical dark halo	radially varying anisotropy, $\beta(r) = -0.15 - 0.20\frac{d\log\rho}{d\log r}$
H5	spherical dark halo	axisymmetric velocity structure, isotropic rotator with $k = 1$
H6	prolate dark halo with $s = 0.85$	axisymmetric velocity structure, no net rotation
H7	oblate dark halo with $s = 1.15$	axisymmetric velocity structure, no net rotation
B1	spherical dark halo, spherical bulge	ergodic
B2	spherical dark halo, spherical bulge	different anisotropies for bulge and halo, $\beta_{\mathrm{halo}} = 0.5$, $\beta_{\mathrm{bulge}} = -1.0$
B3	prolate dark halo $s = 0.85$, spherical bulge	axisymmetric velocity structure, no net rotation
B4	oblate dark halo $s = 1.15$, prolate bulge $s = 0.85$	axisymmetric velocity structure, no net rotation
D1	spherical dark halo, thin disk	axisymmetric velocity structure for halo and disk, disk isotropic rotator
D2	prolate dark halo with $s = 0.85$, thin disk	axisymmetric velocity structure for halo and disk, disk isotropic rotator
D3	spherical dark halo, thin disk	disk with $f(E, L_z, I_3)$ structure and $f_R = 2.0$, halo axisymmetric with $k = 0$
D4	spherical dark halo, thin disk	disk with $f(E, L_z, I_3)$ and $f_{R,\mathrm{disk}} = 4.0$, halo axisymmetric with $k = 0.5$
D5	prolate dark halo with $s = 0.85$, thin disk	disk with $f(E, L_z, I_3)$ and $f_{R,\mathrm{disk}} = 2.0$, halo axisymmetric isotropic rotator
M1	spherical dark halo, spherical bulge, thin disk	axisymmetric structure for halo and bulge (no rotation), disk isotropic rotator
M2	spherical dark halo, spherical bulge, thin disk	axisymmetric velocities for halo/bulge, disk with $f(E, L_z, I_3)$, $f_{R,\mathrm{disk}} = 2.0$
M3	spherical dark halo, spherical bulge, thin disk	disk with $f(E, L_z, I_3)$ and $f_R = 4.0$, bulge no rotation, halo with $k = 0.1$
M4	prolate dark halo $s = 0.85$, oblate bulge $s = 1.15$, thin disk	disk with $f(E, L_z, I_3)$, $f_R = 2.0$, halo and bulge both isotropic rotators

Table 2.1: Set of basic galaxy models constructed for testing purposes with the methods outlined in this paper. Unless stated otherwise, we have used 10^6 particles for each model component. The models labeled 'H1', 'H2', etc., contain only a dark matter halo but differ in the halo shape or the assumptions made for the velocity structure. The models denoted 'B1', 'B2', and so on, contain a bulge in addition to the halo, and the modes with 'D1', 'D2', etc., feature a disk in addition to the halo. Finally the models 'M1', 'M2', etc., contain both a stellar bulge and a stellar disk, next to a dark matter halo.
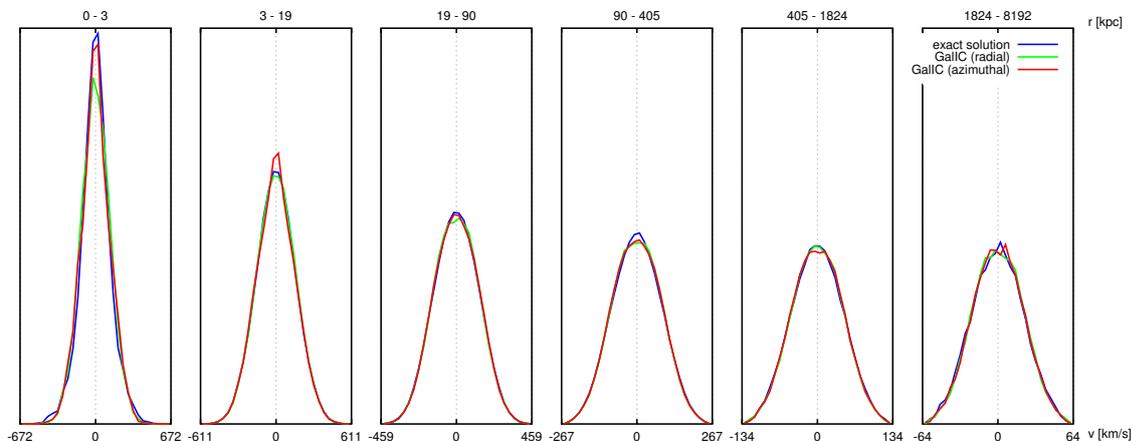
Figure 2.4: Radial and azimuthal velocity distribution functions in different radial shells for the isotropic Hernquist sphere. The blue line show the exact analytic solution. The red and green lines show the constructed solution with our method, separately for azimuthal and radial velocity components. The numbers at the top of each panel indicate the radial range of the measurement.

and we use $N = 10^6$ particles in order to have enough sampling points for a reliable measurement of the produced velocity distribution function.

In Figure 2.4, we show radial and azimuthal velocity distribution functions measured from the ICs produced by our code for this classic Hernquist model, where the analytic distribution function is known analytically. We measure the distribution function in a set of 6 radial shells, as labelled in the different panels. In each panel, we show the analytic distribution function in blue, and the one produced by the GALIC code in red (azimuthal direction) and green (radial direction), respectively. We can nicely see from the figure that the model calculated by GALIC reproduces the expected distribution function rather well in all radial shells, without any significant difference. In particular, note that the model produces the platykurtic nature of the velocity distribution of the Hernquist sphere, which directly shows the presence of higher-order moments that are missed by simpler moment-based methods but are capture by our new approach. This is seen explicitly in Figure 2.5, where we compare the shape of the produced radial velocity distribution function to a Gaussian with the same dispersion.

r: 100.0 - 200.0 [kpc]
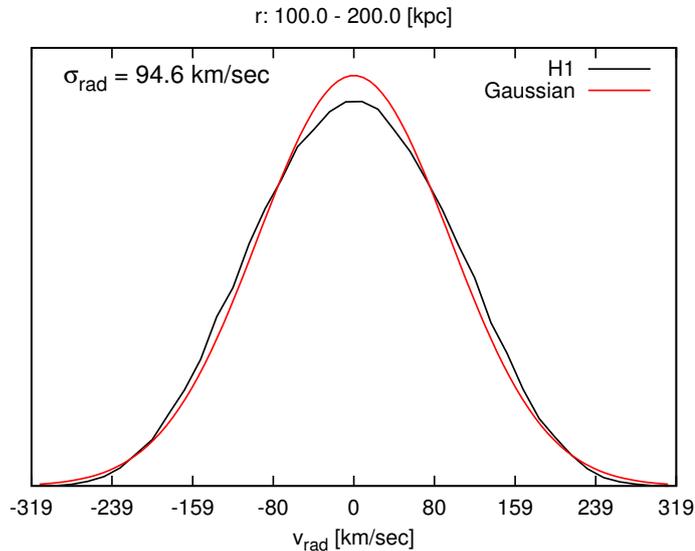
Figure 2.5: Distribution function of the radial velocities in a Hernquist model within a thick radial shell. The black line gives the result of our code for the H1 model, while the red curve is a normal distribution with the same dispersion. The correct platykurtic shape of the distribution function (which is missed in moment-based approaches) is reproduced by our method.

The most important critical test of a method's ability to create initial conditions in equilibrium is however to check the stability of the ICs in a self-consistent simulation under its own self-gravity. To this end, we use the GADGET N-body code, with force accuracy and time integration parameters set conservatively such that energy conservation is excellent. To control discreteness effects in the potential we set the gravitational softening length to a value of 0.05 kpc. In this way we make sure that any secular evolution that is seen really reflects imperfections of the ICs rather than being influenced also by N-body integration errors or two-body relaxation.

Figure 2.6 shows the relative deviation of the spherically averaged density profile from the initial values at different radii and different simulation times, for our H-models. Different line colors mark the different times, as labelled. We have here restricted the simulation time to 1 Gyr, but note that nothing qualitatively changes if this is expanded to 10 Gyrs, significantly
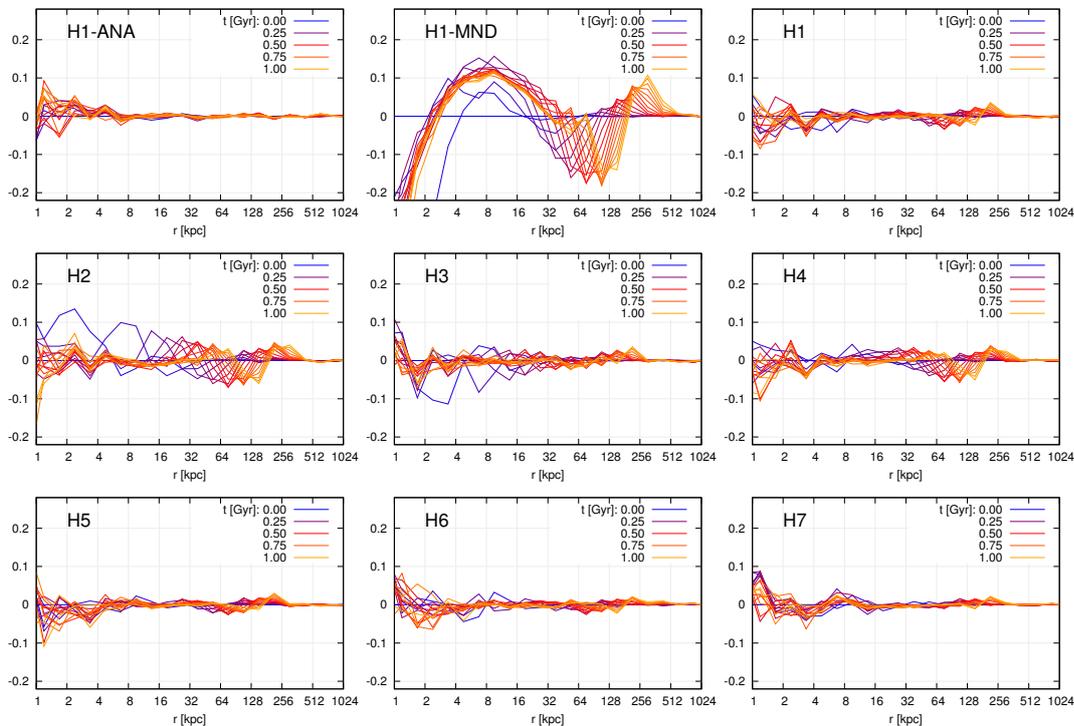
Figure 2.6: Density profile changes of different realizations of our halo-only H-models when evolved in time. The panels show the relative deviation of the spherically averaged density profile from the initial values at different radii and different simulation times. Different line colors mark the different times, as labelled. The top left panel corresponds to initial conditions for the isotropic H1-model realized with the analytic distribution function, while the top middle panel is for the moment-based approach, for comparison. All other results (H1 to H7) are for our new method as implemented in the GALIC code.

longer than the dynamical time of the galaxy model. Let us first focus on a comparison of runs for three different initial conditions constructed for H1, the isotropic Hernquist model. The simulation starting from ICs created with the analytic distribution function is shown in the top left panel, the top right panel shows our GALIC technique, and the top middle panels gives the moment-based method (here realized with the MAKENEWDISK code described in Springel et al. 2005a). As can be seen, our result (top right panel) is nearly indistinguishable from the analytic initial conditions. There is a hint of some small deviations standing out of the noise compared
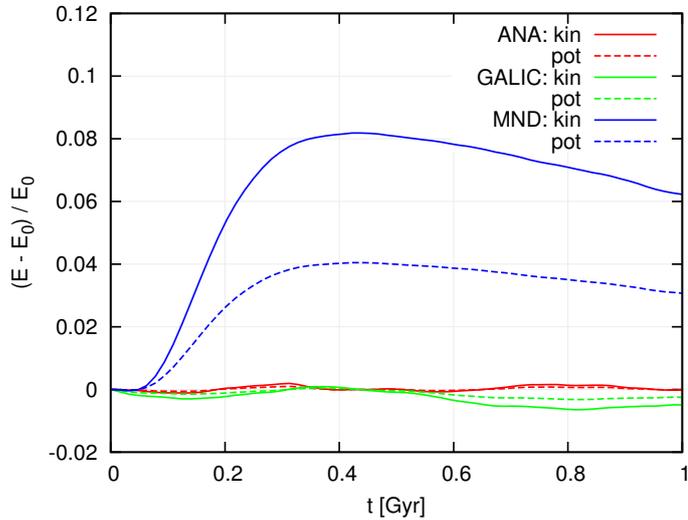
Figure 2.7: Time evolution of the relative change of kinetic and potential energies in the H1 model, for initial conditions realized either through the analytic distribution function (red), through the GALIC code (green), or with a moments based approach (blue).

with the analytic solution, but this is very small if real at all. In contrast, for the moment-based method we see a prominent perturbation propagating outwards, irreversibly changing the mass distribution of the system as it relaxes to a new equilibrium state.

Another view of this difference in the dynamical evolutions of these three simulations is given in Figure 2.7, where we compare the relative changes of the kinetic and potential energies of the three runs as a function of time. We can see that both the N-body realization drawn from the analytic Hernquist distribution function and the GALIC result show essentially stationary energies over the simulation, as expected from a virialized system in equilibrium. In contrast, the moment-based approach shows a rapid evolution in the energies in the first $\sim 300\,\mathrm{Myr}$, and only then settles into a stationary state. Note that in this initial phase, the central potential fluctuates, allowing individual particles to change their energies and the system to relax to a new equilibrium.

We now turn to considerably more demanding models that have an
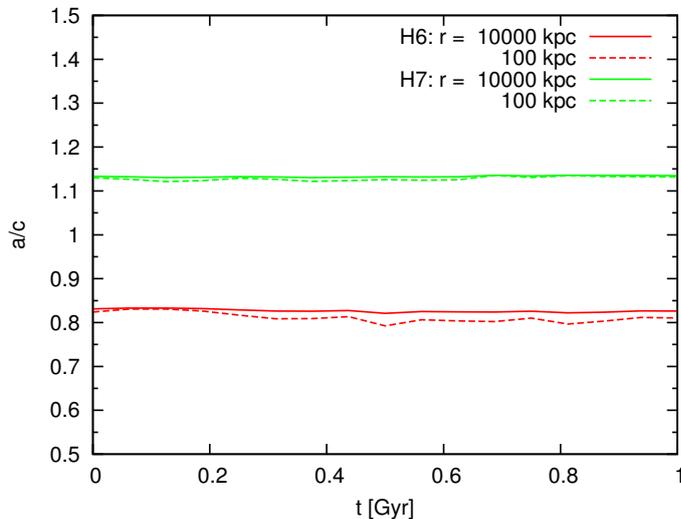
Figure 2.8: Halo shapes of models H6 and H7 as a function of evolution time. We show the ratios of the principal eigenvalues of the moment-of-inertia tensor at two different radii as a function of time. The values reproduce the intended shapes according to Table 1, and are constant in time.

anisotropic, but still spherically symmetric velocity structure. These are models H2, H3, and H4, characterized by asymmetry parameters $\beta = -1$ (for H2) and $\beta = 0.5$ (for H3), corresponding to the cases $\sigma_r^2 = \sigma_\theta^2/2$ and $\sigma_r^2 = 2\sigma_\theta^2$, respectively. The model H4 adopts a radially varying profile $\beta(r)$ as suggested by cosmological simulations.

In Figure 2.9, we show radial profiles of the radial and azimuthal velocity dispersion profiles for initial conditions produced by GALIC for these 4 cases (and for completeness also for all other H-models), both at the initial time and after different times of evolution. For reference, we also include in the figure panels the result for H1 (top left panel) as a grey line, which is the isotropic $\beta = 0$ case. We see that the initial conditions code manages to accurately impose the desired velocity anisotropy at the initial time. Upon time evolution, these velocity dispersion profiles are quite well maintained, but not perfectly in the very central regions for models H2, and to lesser extent, H3. While these two models still manage to maintain a directional difference in the velocity dispersions in the innermost halo even after 1 Gyr
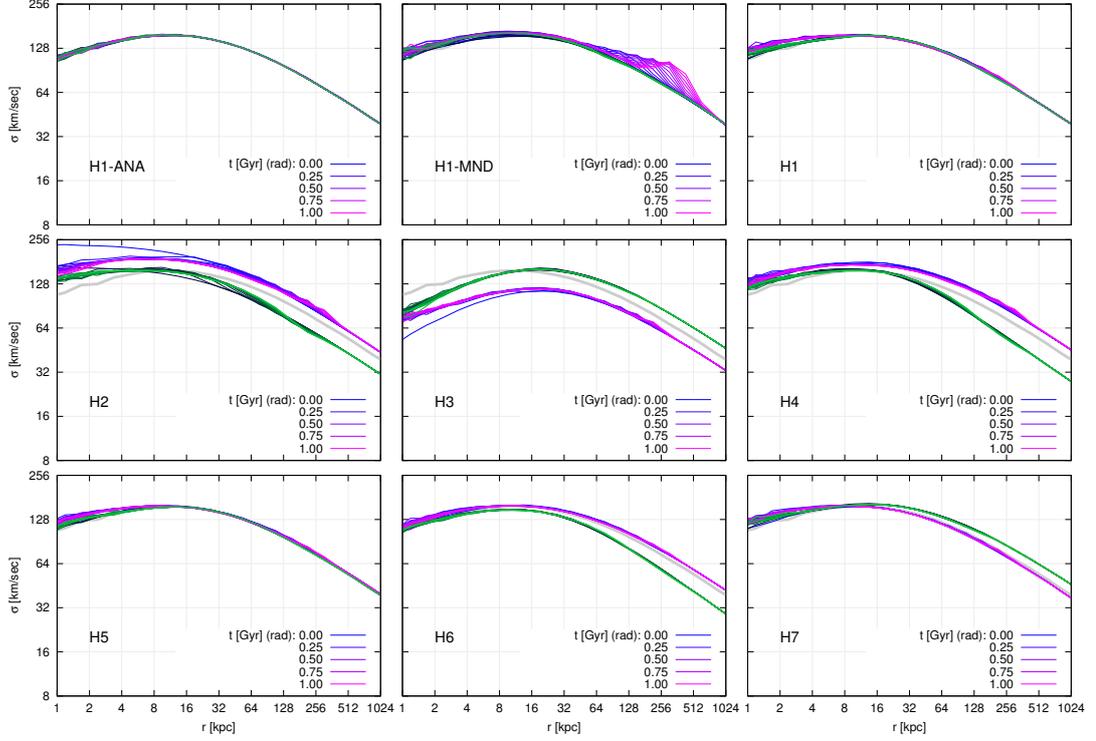
Figure 2.9: Radial and azimuthal velocity dispersion profiles for our models H1 to H7. In each panel, we show the initial conditions produced by GALIC, and the evolved states after different times. The expected profiles based on the Jeans equations are shown underneath. The grey line reproduced the dispersion profile expected for the H1 model, for comparison. In the top row, we also show results for H1 obtained with the analytic distribution function (top left), and with the moment-based method (top, middle).

of evolution, the initial profile is not fully retained in the very central region. Overall, we consider these results however still to be quite good.

We note however that the density profiles of these anisotropic models are accurately retained even in the centres in these anisotropic cases. Some of the panels of Figure 2.6 report the density variations of the anisotropic models H2-H4 upon time evolving their ICs. The relative fluctuations of the density profiles are very small, consistent with the findings for the simple H1 model. Note that in Fig. 2.6 we also include results for the models H5, H6, and H7. These latter three models now feature an axisymmetric assumption

for their velocity structure. H5 is actually slowly rotating, whereas H6 and H7 have prolate or oblate shape distortions, respectively. The absence of any significant time evolution in the spherically averaged density profiles shown in Fig. 2.6 indicates that these models are also rather robust and in good equilibrium.

This is also confirmed by a look at their velocity dispersion profiles shown in Figure 2.9, and a direct analysis of the halos shapes of models H6 and H7 at the final times. Simple measurements of the eigenvalues of their moment of inertia tensors as a function of time (see Figure 2.8) confirm that the imposed halo shapes are accurately retained over time. We also note that the time evolution of the kinetic and potential energies (not shown) confirms that the models are in good equilibrium.

2.6.2 Systems with a bulge and a halo

Next, we consider models that are slightly more complicated and feature two different mass components of very different spatial extent, a Hernquist halo with an embedded, much small stellar bulge, also modelled with a Hernquist profile. Our model B1 simply consists of two ergodic systems nested into each other. B2 varies that by invoking different velocity anisotropies for halo and bulge, with a preference for radial orbits in the halo and tangential ones in the bulge. Finally, B3 and B4 test different shape distortions for halo and bulge, under the assumption of an axisymmetric velocity structure and no net rotation.

Figure 2.10 shows the changes in the radially averaged density profiles of these systems when they are evolved in time, separately for the stellar bulge and the dark matter components. The stability appears to be excellent in most cases, something that is also confirmed by other measures, such as the time evolution of kinetic and potential energies. Only the B2 model performs slightly worse, an outcome that we blame on the dominance of radial orbits in the dark matter component of this model even in the very centre (similar as in H2), and some of these orbits can be affected by the
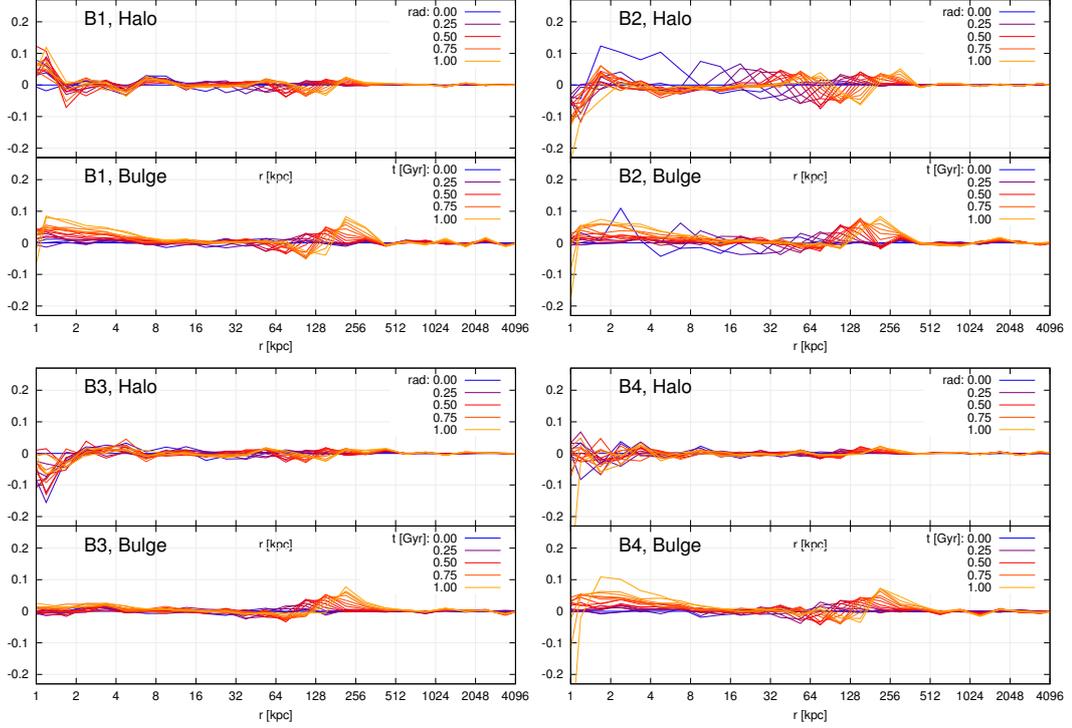
Figure 2.10: Density profiles changes of different bulge models (B-models in Table 1) when evolved in time. For each of the four models, two panels are given that show the relative deviation of the spherically averaged density profile of halo and bulge components relative to the initial values, as a function of radius and for different simulation times. Different line colors mark the different times, as labelled.

radial orbit instability (Buyle et al. 2007).

2.6.3 Systems with a halo, a disk, and an (optional) bulge

We now turn to the much more challenging case of models containing a thin stellar disk. In Figure 2.11, we show the rotation curves of our models containing just a halo and a disk (D-models, top panel), and those of our models containing in addition a bulge as well (M-models, bottom panel). In the innermost regions, the disk dominates slightly over the spheroidal halo in the D-models. The specific parameters chosen for the D-models are
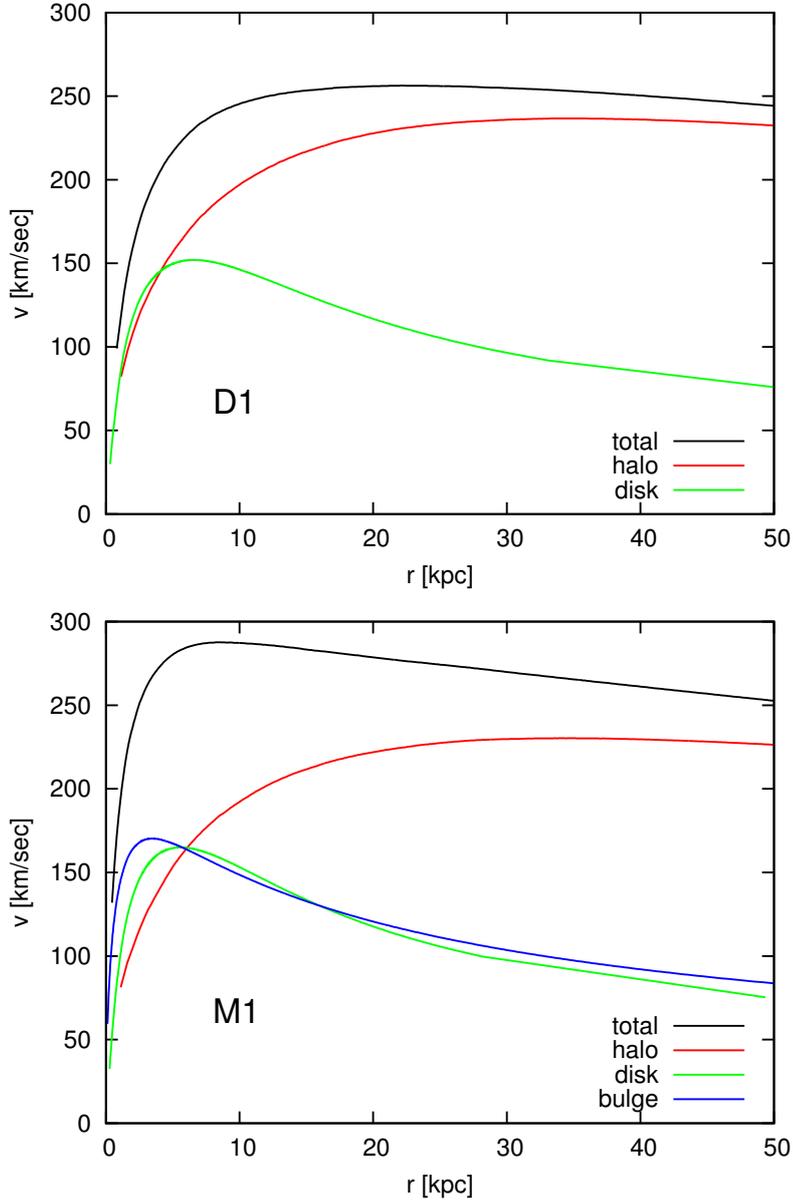
Figure 2.11: Rotation curves of the different mass components in our D-models (top panel) which contain only a dark matter halo and a stellar disk, and our M-models (bottom panel) which in addition contain a central bulge.
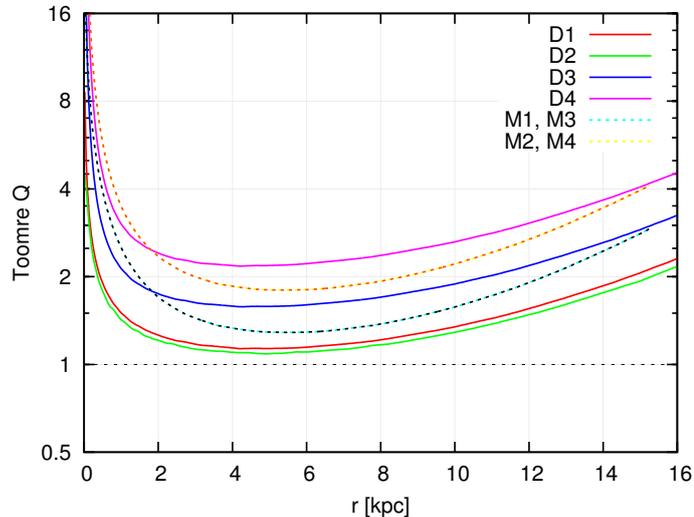
Figure 2.12: Stability against axisymmetric perturbations as expressed through the Toomre Q-parameter, for different D- and M-models.

$m_d = 0.035$ and $\lambda = 0.035$, and the M-models contain additionally a bulge with $m_b = 0.05$ and a scale length set to a tenth of that of the halo. We note that these choices are somewhat arbitrary and not meant to represent a specific system such as the Milky Way; our methods work with similar quality when the parameters are varied over a plausible range.

Interestingly, depending on what assumptions we make about the velocity structure of the disk systems, the expected stability with respect to axisymmetric perturbations can be quite different. In Figure 2.12, we show Toomre's Q-parameter for the disk models D1 to D4, as well as for our M-models. As we see, D1 actually nearly straddles the stability boundary at $Q = 1$, and can hence be expected to be somewhat more prone to axisymmetric perturbations than D3, were Q is boosted thanks to a higher radial velocity dispersion.

In Figure 2.13, we show the time evolution of the azimuthally averaged projected disk surface density profiles, for models D1-D4, and for M1-M4. We can see that all models are reassuringly stable. The improvement compared with moment-based methods such as that implemented the MAK-
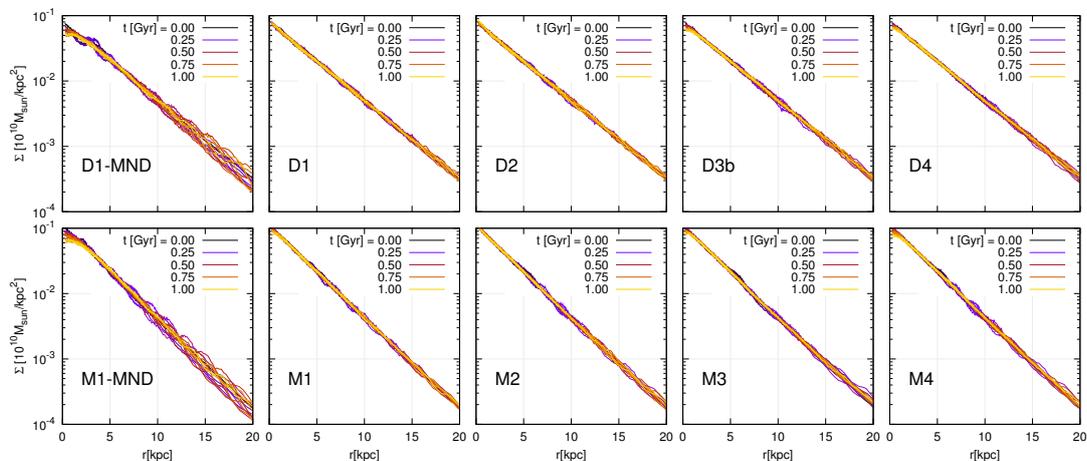
Figure 2.13: Radial surface density profiles of the stellar disk component under time evolution, for D1-D4 (top row), and M1-M4 (bottom row). The left column shows the D1 and M1 models again, but this time for initial conditions produced with the moment-based approach implemented in MAKENEWDISK.

ENEWDISK code can perhaps be best appreciated by comparing to the results for this method, which are give in the leftmost panels of Fig. 2.13 for models D1 and M1.

Finally, a complementary view of the disk stability is obtained by considering the time evolution of the vertical density structure of the disks, wich is shown in Figure 2.14. Again, the models D1-D4 and M1-M4 are seen to retain their disk density structure accurately, relatively independent of the different variants of halo and bulge shapes, and the different degrees of rotation that we tried. Only M3 performs noticeably worse than the other models in the outer disk. When we compare the D1 and M1 models to corresponding realizations obtained with the moment based approach (left most panels in the figure), there is a clear improvement.

2.6.4 Dependence on nuisance parameters

Our iterative method for finding equilibrium galaxy models with the GALIC code involves several free parameters, for example the fraction of particles that is allowed to be concurrently optimized, the number of optimization
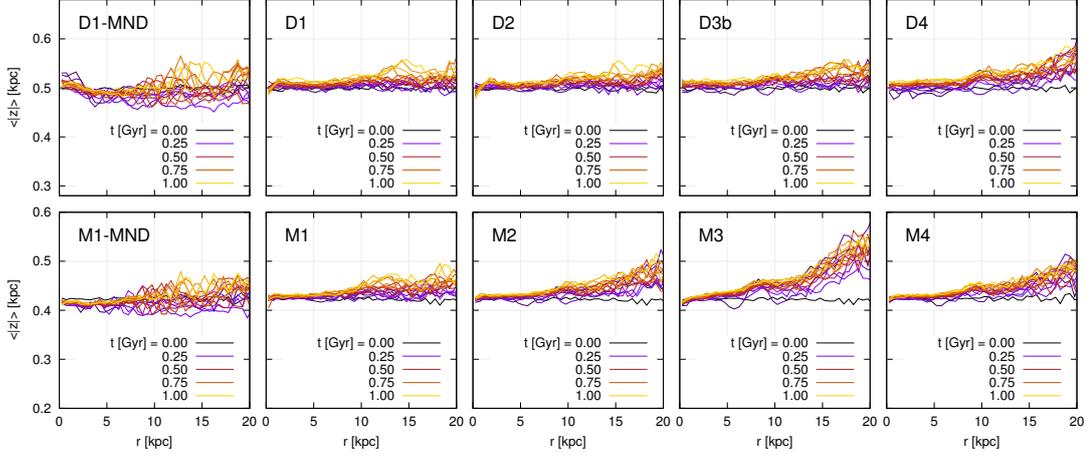
Figure 2.14: Vertical disk height as a function of radius at different times when our disk models are evolved in time, for the D1-D4 models (top row), and the M1-M4 models (bottom row). In the left column, we show for comparison the corresponding results when the initial conditions for D1 and M1 are constructed with a moment-based approach.

cycles before a randomization is carried out, the resolution of the density response grid, the length of time over which orbits are integrated, and a few more minor ones.

We have carefully tested whether our results depend significantly on the settings of any of these parameters. This is fortunately not the case. We find that our results are rather robust when any of these nuisance parameters is changed around our default settings. As a case in point, we show in Figure 2.15 an explicit test for the number of orbits that are integrated, comparing results produced for the M1 model where this parameter has been lowered by a factor of 2, or increased by a factor of 2, compared with our default choice. Reassuringly, we see that the density deviations occurring in time evolutions of the produced ICs are of very similar magnitude, i.e. their quality appears indistinguishable.

We also find that that the grid resolution used for recording density and velocity dispersion responses plays only a negligible role for the results, provided the finest possible level is not overly coarse. This can be understood
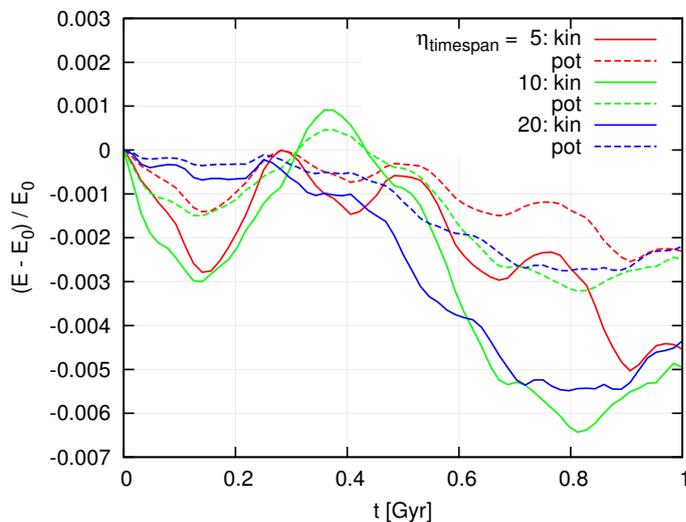
Figure 2.15: Test of the dependence of our results on the integration time of particle orbits. We show the relative changes in potential and kinetic energies when evolving initial conditions constructed with integration times lowered or increased by factors of two relative to our default value. The particular system used here is M1, but similar results are found for other models and other changes of numerical parameters in our code.

as a result of our adaptive binning prescription. Potentially more important may be the value of the number of particles required in a cell before it may be split up in finer cells. But we have also found here that varying this parameter over a significant range does not change the results appreciably. Finally, the last free parameter that we have extensively tested is the integration accuracy of the orbits. Here our typical relative energy errors for the integration of individual orbits are below 10^{-3}, already suggesting that this should be good enough and not introduce any significant errors into the results. Indeed, investing more computational effort and lowering the integration errors through finder timestepping does not change our results in any significant way.

2.7   Discussion and conclusions

In this chapter, we introduced a new iterative method for constructing equilibrium N-body galaxy models. The principle advantage of our method is that it can produce N-body systems that are essentially in exact equilibrium for rather general density distributions, making the method ideal for studies of galaxy dynamics and numerical experiments with isolated or colliding galaxies. Compared to alternative schemes like the Schwarzschild method, our approach eliminates restrictions arising from a finite orbit library or from required regularization schemes. Also, our method allows a natural inclusion of simulation aspects like the need for a gravitational softening.

The test results we have analyzed show a considerable improvement of the quality of the created initial conditions compared to existing codes such as the moment-based MAKENEWDISK, a technique that has been used in numerous studies over recent years. This is possible thanks to the absence of any assumptions in our approach with respect to the importance of higher-order moments of the velocity distribution function. The main disadvantage of our method lies in its higher computational cost compared to moment-based approaches. However, thanks to the scalable parallelization implemented of our code, this should not be a serious restriction in practice. For example it took about 5 hours on 96 AMD-6174 cores (2.2 GHz) to compute high-quality solutions for our most compicated models M1-4, while for the one component models it took only 1 hour. And since poor ICs may also impact any further scientific investigation, the additional CPU effort invested for better ICs should in many cases be well worth the effort.

Finally, we also note that numerous optimizations in our code could well be made to reduce its CPU time consumption. For example, the time to full convergence for a system with large N may be reduced considerably by first treating a smaller subsample of the particles with correspondingly higher mass. Once this system has fully converged, one could then create

the large target realization from it through bootstrap sampling, followed by briefly relaxing the big system to the final equilibrium.

In future work, it might be interesting to extend our approach to genuinely triaxial systems, which are of course considerably more challenging than the axisymmetric case considered here. One could even include additional phenomena such as figure rotation. An important challenge is here to suitably store the density response grid. Here, our approach, which only requires essentially one such response grid should be considerably less restrictive than Schwarzschild's method because the requirement to store a huge orbit library is avoided. In the meantime, we publicly release our GALIC code (http://www.h-its.org/tap/galic) , hoping that it proves useful for future N-body studies in galactic dynamics.

The stability of stellar disks in Milky-Way sized dark matter halos

FAUST:
Nun gut, wer bist du denn?

MEPHISTOPHELES:
Ein Teil von jener Kraft,
Die stets das Böse will
Und stets das Gute schafft.

— Faust, *Goethe*

Abstract

We employ an improved methodology to insert live stellar disks into high-resolution dark matter simulations of Milky Way sized halos, allowing us to investigate the fate of thin stellar disks in the tumultuous environment of cold dark matter structures. We study a set of eight different halos, drawn from the Aquarius simulation project, in which stellar disks are adiabatically grown with a prescribed structure, and then allowed to self-consistently evolve. The initial velocity distribution is set-up in very good equilibrium with the help of the GALIC code. We find that the residual triaxiality of the halos leads to significant disk tumbling, qualitatively confirming earlier work. We show that the disk turning motion is unaffected by structural properties of the galaxies such as the presence or absence of a bulge or bar. In typical Milky Way sized dark matter halos, we expect an average turning of the disks by about 40 degrees between $z = 1$ and $z = 0$, over the coarse of 6 Gyr. We also investigate the impact of the disks on substructures, and conversely, the disk heating rate caused by the dark matter halo substructures. The presence of disks reduces the central subhalo abundance by a about a factor of two, caused by an increased evaporation rate due to gravitational shocks from disk passages. We find that substructures are important for heating the outer parts of stellar disks but do not appear to significantly affect their inner parts.

3.1 Introduction

The ΛCDM cosmogony represents the leading theoretical model for structure formation and predicts a hierarchical growth of galaxies in which dark matter halos have prominent triaxial shapes and are full of substructures. It is not well understood how grand design spiral galaxies can form and survive in this violent environment, which seems at the outset quite hostile towards the long-term survival of thin, cold stellar disks. Such disks may easily become distorted and tilted by triaxial potentials, they may suffer from bar instabilities or excessive heating from dark matter substructures and merger events.

Until recently, full hydrodynamical cosmological simulations have not been particularly successful in making realistic thin disk galaxies without dominating bulges, even though some remarkable progress has recently been achieved on this long-standing problem (Governato et al. 2010; Agertz et al. 2011; Guedes et al. 2011; Aumer et al. 2013; Stinson et al. 2013; Marinacci et al. 2014a). But even in the most recent simulation works, the disks have generally been found to be too thick (e.g. Marinacci et al. 2014a). This could plausibly be related to inaccurate modelling of the disk formation physics, to a lack of numerical resolution, or to a combination thereof. However, another possibility is that the hierarchical nature of dark matter halo growth in cold dark matter models, combined with the strong triaxiality of CDM halos, quite generally causes excessive disk heating. In this case, CDM scenarios may pose restrictive intrinsic limits on the possible abundance of thin stellar disks. It is therefore important to shed more light on the question under which conditions thin stellar disks can comfortably survive in CDM halos.

It is difficult to study this issue systematically with full cosmological hydrodynamical simulations, due to their very high computational cost and the lack of freedom to prescribe the structural disk properties of the galaxies. An alternative approach is to study the stability of well-resolved col-

lisionless stellar disks that are inserted into dark matter halos in a suitable fashion. This has been explored in numerous previous works, but most of these studies employed isolated toy models of halos and disks, and not the full cosmological context (e.g. Walker et al. 1996; Sellwood et al. 1998; Velazquez & White 1999; Athanassoula & Misiriotis 2002; Debattista et al. 2006; Gauthier et al. 2006; Kazantzidis et al. 2008; Read et al. 2008; Machado & Athanassoula 2010). Only a few works have tried to achieve a consistent cosmological embedding, where the disk is somehow inserted by hand in a suitable fashion in a growing dark matter (Berentzen & Shlosman 2006; D'Onghia et al. 2010; DeBuhr et al. 2012). If this is done carefully, one can hope to reach much higher resolution than accessible in full hydrodynamical simulations, while at the same time having full control about the structural parameters of the inserted disk galaxy. The latter allows in principle an easy exploration of the full space of observationally realistic disk parameters.

The most sophisticated variant of this approach has recently been presented by DeBuhr et al. (2012), who introduced disks in a subset of the halos studied in the "Aquarius" project (Springel et al. 2008a), which consists of high-resolution simulations of Milky Way-sized halos in a ΛCDM universe. The Aquarius halos have been particularly well studied, with a subset of them also being followed up hydrodynamically (Scannapieco et al. 2009, 2012; Aumer et al. 2013, 2014; Marinacci et al. 2014a,b; Pakmor et al. 2014; Okamoto et al. 2014). The corresponding initial conditions exist at different resolutions and are of high quality, allowing converged results even including the density profiles of individual dark matter subhalos (Springel et al. 2008a).

The study of DeBuhr et al. (2012) has examined halos A, B, C, and D of the original Aquarius set. A general result of their work was that disks in CDM halos of Milky-Way sized halos appear to be rather brittle, and can be expected to substantially change over the coarse of a few Gyrs. In particular, DeBuhr et al. (2012) have found that their disks tumble

substantially, and universally grow bars, unless very light disks are used. A substantial fraction of the initial disk material reached large heights above and below the disk plane, and significant warps in the disks where detected as well.

However, DeBuhr et al. (2012) used a comparatively simple method to initialize the disk velocity distribution function, and only one family of initial galaxy structures was considered. In the present work we try to improve on this earlier work in several respects, in particular by extending the study to a larger halo sample that encompasses eight Aquarius halos, by using a more sophisticated and flexible method to initialize the initial disk models based on our GALIC code (Yurin & Springel 2014), by including also models with stellar bulges, and by checking the robustness of our results with a convergence study at much higher resolution than used in previous work. We also consider additional lines of analysis, for example by examining the mutual impact of the disk and the substructures onto each other. Substructures may induce heating of the disk, but they may also themselves be depleted through gravitational shocks experienced during disk transition or pericentre passage.

In this work we are especially interested in the question how universal the disk tumbling phenomenon is, and to what extent it is affected by the initial orientation and structure of the disk galaxy relative to its hosting dark matter halo. We would also like to better understand under which conditions strong bars can be avoided in live dark matter halos, and whether the large number of dark matter substructures poses a significant problem for disk stability.

This chapter is structured as follows. In Section 3.2 we review the methodology we apply, while in Section 3.3 we describe the simulation set we have carried out. Section 3.4 is then devoted to an analysis of our results for pure disk models, whereas in Section 3.5 we turn to models that also include a central bulge. In Section 3.6 we briefly analyze the impact of disks on substructures and vice versa, and in Section 3.7 we examine the

robustness of our results with respect to numerical resolution. Finally, we conclude with a discussion and summary in Section 3.8.

3.2   Methodology

In this study, we carry out resimulations of the "Aquarius" set of initial conditions of Milky Way-sized dark matter halos. They have previously been analyzed in a number of studies (Springel et al. 2008a,b; Navarro et al. 2010; Xu et al. 2009), where in particular the high quality of the zoom initial conditions (created by Adrian Jenkins) was demonstrated, allowing good convergence of all dark matter properties of the halos. We shall mainly work with 'level 5' in the nomenclature of Springel et al. (2008a), were the dark matter halos have close to 1 million particles in the final virial radius, corresponding to a dark matter particle mass of about $\sim 3 \times 10^6 \, \mathrm{M_\odot}$. For selected models, we also carry out simulations at 8 times ('level 4') and 64 times ('level 3') better resolution, reaching up to ~ 50 million particles dark matter and ~ 19 million star particles in the virial radius.

In order to set-up a live stellar disk in the evolving dark matter (re-)simulations of the Aquarius halos, we proceed similarly to DeBuhr et al. (2012). At a certain redshift z_{insert} we place a number N_{disk} of massless star particles into the dark matter halo, sampling a prescribed density distribution placed at the centre of the potential minimum and oriented along one of the principal axes of the halo at that time. The disk mass is then grown linearly in time to a final mass M_d reached at a redshift z_{live}. During the growth phase, the relative distances of the disk particles with respect to each other are kept fixed in physical coordinates, and the whole set of disk particles is coherently moved as a solid body under the total gravitational force experienced by all the particles of the disk. The initial velocity of all the disk particles is set equal to the bulk velocity of the inner dark matter halo (defined as $r < R_{200}/4$, where R_{200} is the radius enclosing a mean overdensity 200 times the critical density). While the disk mass is

ramped up, the dark matter particles start to experience the additional gravitational force due to the disk particles, so that the dark matter halo reacts adiabatically to the growing disk. This process of inserting the disk ensures that the disk stays well centred in the moving dark matter halo, and the realization of the disk potential through particles avoids the need to make any approximation in representing the disk potential; in fact, the density structure of the inserted stellar system can be chosen freely as long as it is physical.

When the disk has reached its final mass, we continue the simulation by treating the disk particles as live, which simply means that from this point on they are treated as ordinary collisionless particles with independent orbits in an evolving gravitational potential. This requires an initialization of the initial velocities of the star particles at z_{live}, which should be done such that the disk is in a self-consistent dynamical equilibrium at this instant. Formally, this corresponds to finding a stationary solution of the Poisson-Vlasov system, which is a quite non-trivial problem for general mass distributions (Binney & Tremaine 2008). In DeBuhr et al. (2012), an approximate solution of the Jeans equation was used to initialize the disk velocities, but the details of this procedure were not described. Here we adopt the method of Yurin & Springel (2014), which iteratively derives a high quality distribution function for the disk particles. The method as implemented in the publicly available GALIC code integrates a large set of test particles representing the target mass distribution in the self-consistent gravitational potential of the system, and adjusts the velocities until the deviation between the time-averaged density response of the particle orbits and the target density distribution is minimized.

In order to take the actual shape of the dark matter halo into account at the moment the disk goes live, we directly use the dark matter particles in the cosmological simulation to compute the halo forces in GALIC. However, as the stable version of the code is at present restricted to axisymmetric disk models, the dark matter halo forces are averaged in the azimuthal di-

rection, i.e. the dark matter force field is "axisymmetrized" without actually changing the dark matter halo. We note that this approximation could be avoided in principle in future refinements of the method, in which case it would then become possible to insert ellipsoidal disks. For the moment, we stick however with inserting axisymmetric disk models with a correspondingly axisymmetric velocity structure. Also, we note that GALIC uses additional constraints on the velocity dispersions in order to single out a desired target solution among the many different degenerate solutions that are in principle possible for reproducing the same density distribution. For example, GALIC can realize disk models with three integrals of motion and tilted velocity ellipsoids, similar to what is observed for the Milky Way (Siebert et al. 2008b; Binney et al. 2014; Büdenbender et al. 2014). However, for simplicity, we have restricted ourselves to simpler models with two integrals of motion (E, L_z), implying that the radial and vertical velocity dispersions are equal.

In practice, we have used the moving-mesh code AREPO (Springel 2010) for evolving our N-body systems in time. While the hydrodynamical features of this code are not exercised in this work, its N-body solver represents an improved and more efficient realization of the algorithmic methods of the GADGET code (Springel 2005b), which we found convenient to make use of. In particular, the built-in parallel version of the SUBFIND algorithm (Springel et al. 2001a) used to identify dark matter substructures is more efficient in AREPO than in GADGET. We run SUBFIND regularly on the fly while the N-body system is involved in order to track the masses and positions of all halos and subhalos, as well as to measure basic dark matter halo properties such as their shape orientations.

To facilitate the realization of a large number of disk insertion simulations, we have largely automized the process of putting in live disks into the Aquarius halos. To this end, the cosmological simulation automatically measures the dark matter halo orientation at z_{insert}, inserts the initially rigid stellar system, and continues with the disk growth phase until z_{live},

at which point a snapshot file is written. A special version of GALIC is then started by a script that replaces the disk velocities in this snapshot file with a self-consistent stationary solution, taking the halo motion, the disk orientation, the Hubble flow at that epoch, etc., into account. Then, the AREPO simulation of the N-body system is continued from z_{live} to $z = 0$.

To determine the orientation of the dark matter halo, we calculate the principal axes of the moment-of-inertia tensor of the dark matter particles in a spherical region of size $R_{\text{vir}}/4$. By restricting ourselves to the inner region, we avoid strong influences from the less well relaxed outer parts of halos. In general, analysis of the halo shapes of CDM halos have found only mild variations of the axis ratios with radius (Allgood et al. 2006), and quite stable directions of the ellipsoidal axis (Hayashi et al. 2007). By measuring the moment-of-inertia tensor in a spherical aperture the axis ratios are biased low, but the directions of the principal axes, which is all that matters for our purposes should be unaffected and line up well with the principal axes of the halo potential. In our simulations we have aligned the spin axis of the disk either with the minor or the major axis identified in this way. The direction of the spin axis was chosen such that the angle with the dark matter angular momentum of the inner halo (again using $r < R_{200}/4$ for selecting this) was minimized. To take account of the fact that we expect the formed galaxy to be stationary in physical coordinates, we keep the gravitational softening fixed in our runs in comoving coordinates until $z = 1.5$, and fixed in physical coordinates thereafter. The same softening length is used for the stellar particles and the high-resolution dark matter particles.

The above methodology can be readily generalized to stellar systems that include a bulge component besides a disk. To this end an additional set of particles is inserted and grown in parallel to the disk component, and GALIC then calculates initial velocities for the bulge component as well.

For test simulations, we have also implemented a "roundening" procedure for the dark matter halo. To this end, all dark matter particles

contained in the friends-of-friends (FOF) group of the target halo at z_{insert} are rotated by a random angle around the halo centre, and their velocity vectors are randomized in direction in the rest frame of the halo. By construction, this procedure makes the halo spherically symmetric and smooth with isotropic distribution function while keeping the spherically averaged density profile and kinetic energy in random motions unchanged. All substructure in the FOF halo is eliminated as well. However, as a result of this procedure the halo will be slightly out of equilibrium initially. It has however enough time to relax again during the disk growth phase, so that when the disk goes live it does so in a stationary halo that is spherical apart from asymmetries induced by the disk growth itself. In a variant of this procedure, we restrict the rounding operation to just those particles bound in substructures (which amount to a few percent of the total mass of the halo). This allows the creation of dark matter host halos that retain the cosmological triaxiality but are largely pruned of dark matter substructures, at least at times close to z_{insert}. Later, new substructures will fall in due to halo growth.

3.3   Simulation set

For definiteness and ease of comparison with DeBuhr et al. (2012), we adopt for our default disk insertion runs their choice of disk mass, $M_d = 5 \times 10^{10} \, \text{M}_{\odot}$, and disk scale length, $R_d = 3 \, \text{kpc}$. Also, we use in these default models their choices of $z_{\text{insert}} = 1.3$, $z_{\text{live}} = 1.0$, and $N_{\text{disk}} = 200,000$, combined with a canonical thickness of 0.2 times the scale length. For our adopted cosmology[1], the growth period of the disk then lasts $T_{\text{growth}} \simeq 1 \, \text{Gyr}$, and the live evolution of the disk system proceeds for $T_{\text{live}} \simeq 6 \, \text{Gyr}$.

---

[1]The cosmology adopted in the Aquarius project is the same one used in the Millennium simulation, and is characterized by $\Omega_0 = 0.25$, $\Omega_\Lambda = 0.75$, $\sigma_8 = 0.9$, $n_s = 1.0$, and a Hubble constant of $H_0 = 73 \, \text{km} \, \text{s}^{-1} \text{Mpc}^{-1}$, consistent with with the WMAP-1 and WMAP-5 cosmological constraints. The small offset of the cosmological parameters with the most recent determinations by Planck does not matter for the purposes of this study.

We have carried out simulations for eight different Aquarius dark matter halos, labelled A to H, following the notation of Springel et al. (2008a) and Scannapieco et al. (2009). In all of these, the pure disk models were run both with minor and major axes orientations, forming a set of 16 default models. The simulations are so-called 'zoom-simulations' where the mass resolution of the initial particle load has a strong spatial resolution; a high-resolution sampling of the Lagrangian region of the target halo is surrounded by shells of progressively more massive particles, so that the target halo feels the same gravitational tidal fields as if it was forming in a simulation where the whole periodic box of side-length 137 Mpc was uniformly followed at the high resolution.

In additional simulation sets, either the structural properties of the inserted galaxy models were modified, the time of the disk insertion was varied, or additional experiments like a rounding of the dark matter halo or a (partial) elimination of substructure was carried out. Most of these additional runs were only done for the minor orientation, as we generally find only small systematic differences between the minor and major orientations, with a small preference for a higher stability of the minor orientation. Finally, we have done runs at higher resolution for selected models of the A-halo, the pure disk model with the minor rotation, and also the default disk plus bulge with minor axis orientation.

Besides the default disk series, our other main series of runs consists of disk plus bulge models where the total stellar mass was kept fixed at $M_\star = 5 \times 10^{10} \, \mathrm{M_\odot}$, but one third of the stellar mass was moved to a spherical stellar bulge, modelled with a Hernquist profile with scale length $a = 2 \, \mathrm{kpc}$, with the rest staying in a disk with exponential surface density profile. Note that these systems still have roughly the right stellar mass expected based on abundance matching arguments for halos of this size, and the disk-to-bulge mass ratio of 2:1 is still reasonably large. In another series we have made this ratio more extreme, by exchanging the masses of disk and bulge, yielding a disk-to-bulge mass ratio of 1:2. In addition, we have considered

#	Simulation names	Disk parameters		Alignment	Bulge parameters		ϵ_{grav} [kpc]	Notes
		M_d [10^{10}M$_\odot$]	N_{disk}		M_b [10^{10}M$_\odot$]	N_{bulge}		
1	A5..H5-minor	5.0	2×10^5	minor			0.68	
	A4-minor	5.0	1.6×10^6	minor			0.34	
	A3-minor	5.0	1.28×10^7	minor			0.17	
2	A5..H5-major	5.0	2×10^5	major			0.68	
3	A5..H5-with-bulge-minor	3.33	2×10^5	minor	1.67	1×10^5	0.68	
	A4-with-bulge-minor	3.33	1.6×10^6	minor	1.67	8×10^5	0.34	
	A3-with-bulge-minor	3.33	1.28×10^7	minor	1.67	3.2×10^6	0.17	
4	A5..H5-lighter-disk-minor	3.33	2×10^5	minor			0.68	
5	A5..H5-massive-bulge-minor	1.67	2×10^5	minor	3.33	2×10^5	0.68	
6	A5..H5-rounded-minor	5.0	2×10^5	minor			0.68	rounded dark halo (like #1)
7	A5..H5-with-bulge-rounded	3.33	2×10^5	minor	1.67	1×10^5	0.68	rounded dark halo (like #3)
8	A5..H5-subs-wiped-minor	5.0	2×10^5	minor			0.68	rounded dark halo (like #1)
9	A5..H5-with-bulge-subs-wiped	3.33	2×10^5	minor	1.67	1×10^5	0.68	rounded dark halo (like #3)
10	A5..H5-reorient-minor	5.0	2×10^5	minor			0.68	growth phase reorientation
11	A5..H5-reorient-major	5.0	2×10^5	major			0.68	growth phase reorientation
12	A5..H5-late-insert-minor	5.0	2×10^5	minor			0.68	$z_{\mathrm{ins}} = 0.5$, $z_{\mathrm{live}} = 0.364$

Table 3.1: Overview of our simulation sets and their basic numerical parameters. We have organized the runs in different series, as illustrated in the table. The first series consists of our default pure disk runs, inserted along the minor axis. Here we also carried out runs for the A-halo at 8 times and 64 times higher resolution, respectively. Series #2 repeats the level-5 runs with a major orientation of the disk. In series #3 we replace the disk with a bulge+disk system in where one third of the mass is moved to a disk. In series #4 this lighter disk is kept but the bulge is omitted, while in series #5 we swap the masses of disk and bulge wis that we end up with a relatively massive bulge and a disk of half the mass of the bulge. The remaining series represent special simulations to test various aspects of our procedures. Series #6 and #7 repeat the runs of series #1 and #3, respectively, but this time the dark matter halo is "rounded"when the disk is inserted, as described in the text. Series #8 and #9 restrict the rounding to particles bound in substructures, so that smooth dark matter halos are produced. In series #10 and #11 we have tested for minor and major orientations whether a continuous reorientation of the disk during the growth phase between $z_{\mathrm{insert}} = 1.0$ and $z_{\mathrm{live}} = 1.3$ helps in reducing disk tumbling. Finally, series #12 delays the disk insertion to much later time, where the dark matter halos have relaxed more and therefore may potentially make it easier for disks to survive unaffected.

a series of runs where only the disk mass was reduced by one third relative to our default disk model ("light disk" models) and the bulge was omitted.

Table 3.1 gives an overview of these different simulation sets and lists some of their most important numerical parameters. In all the runs, we have used conservative integration settings for the tree force accuracy and time integration timestep in order to ensure that all simulations are unaffected by orbit integration errors.

3.4 Results for pure disk models

3.4.1 Disk orientation and visual morphology

In Figure 3.1 we show projected images of the time evolution of the stellar disk material in our eight Aquarius halos, where a pure disk of mass $5 \times 10^{10}\,\mathrm{M_\odot}$ and scale length $R_d = 3\,\mathrm{kpc}$ is inserted along the minor axis of the halos. In the different panels of the figure, the disk stars have been turned into a face-on orientation, taking the spin angular momentum of the stars in the central region of the disk (within $5\,\mathrm{kpc}$) to define the disk normal. All the models almost immediately form very strong bars, consistent with the findings of DeBuhr et al. (2012) for halos A-D.

When viewed in an edge-orientation, as shown in Figure 3.2, a planar disk-like distribution of the majority of stars is maintained in all the cases, but the presence of the strong bars is clearly revealed by pronounced X-shaped features in the centre of the galaxies. It is also evident that the disks are thickened to different degrees. Halo E sports a particularly thick disk at the end, and some systems, notably halos B and F, show substantial bending in the periphery of the disks. Nevertheless, the amount of stellar material significantly outside the disk plane seems to be rather limited, somewhat different from the findings of DeBuhr et al. (2012) for halos A-D.

The equivalent simulations for inserting the disks along the major orientation of the halos (series #2) yield qualitatively very similar results, and we therefore refrain from showing the corresponding images. Figure 3.3

Results for pure disk models
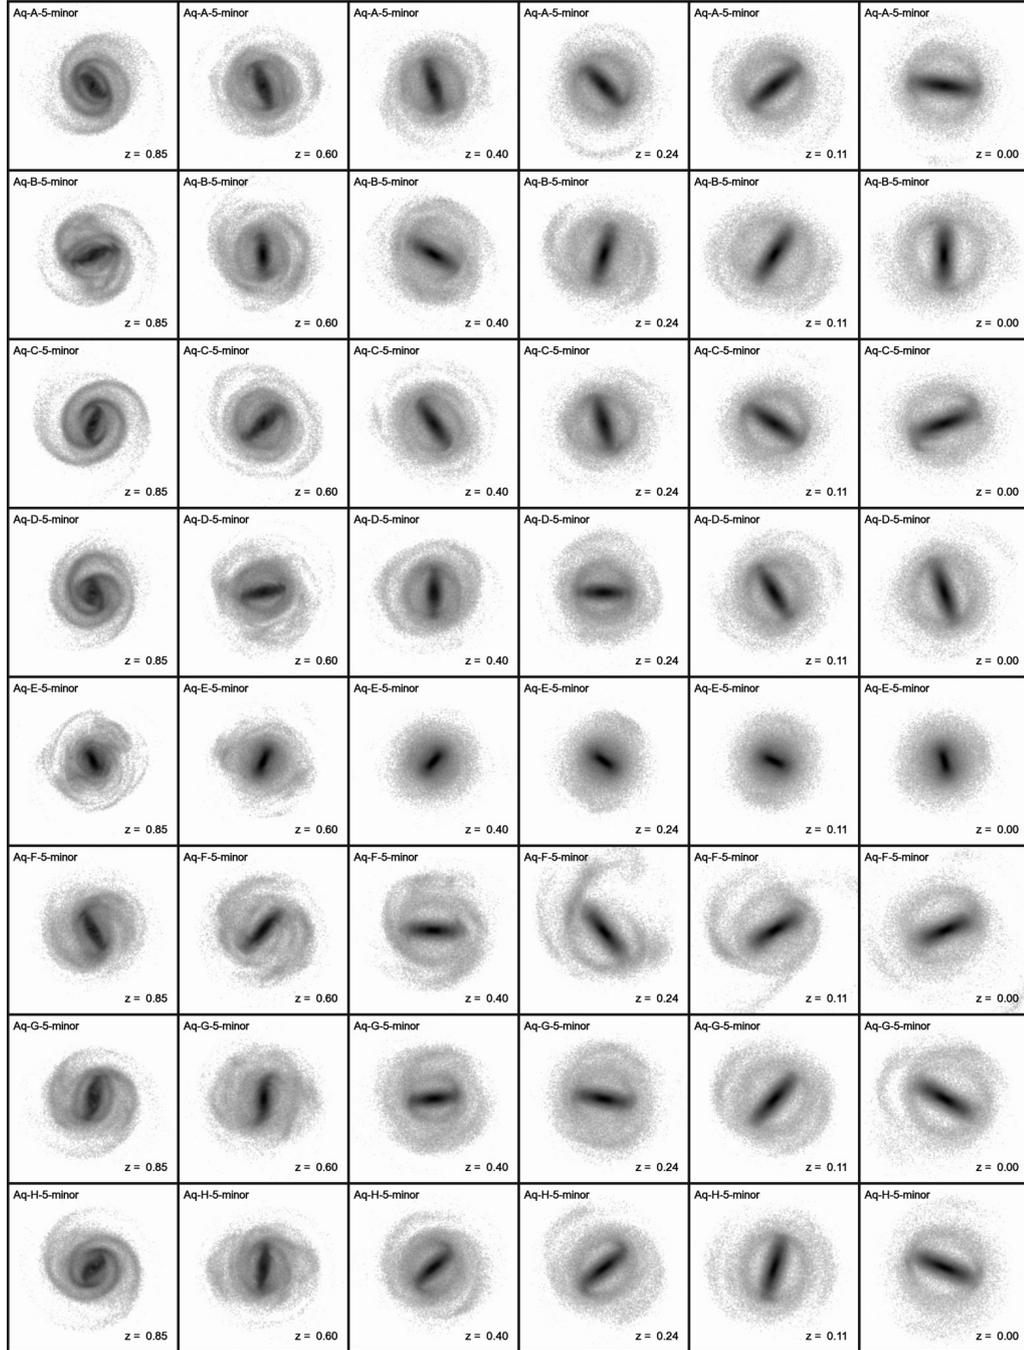
Figure 3.1: Face-on projections of the stellar mass of the disks in our default runs (series #1). Each row shows a different Aquarius halo, as labeled, with the columns left to right showing different times ranging from $z = 0.85$ to $z = 0$. Each panel has a fixed physical size of 54 kpc on a side, and uses the same logarithmic grey scale (covering a dynamic range of 1000 in surface density) for visualizing the adaptively smoothed surface density of star particles. In each panel, the stellar particles have been turned independently into a face orientation as determined by the angular momentum vector of the stars within the central 5 kpc of the disk.
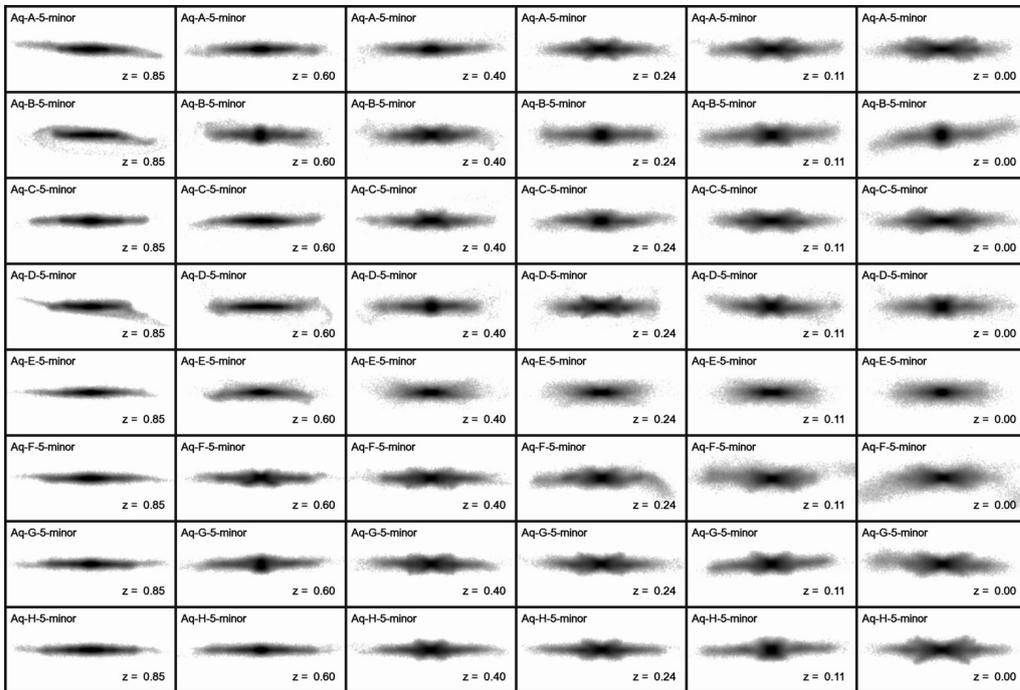
Figure 3.2: Edge-on projections of the stellar disks in our default runs with pure disks. This images correspond to side-views of the corresponding images shown in Figure 3.1, with an unchanged physical size of each panel in the horizontal direction (54 kpc), and an identical color-scale.

gives instead an overview of the time evolution of the orientation of the halos' minor, major and intermediate axes, as defined based on the moment of inertia tensor of the central dark matter distribution in a spherical aperture of size $R_{200}/4$. We also include in the plots the evolving orientation of the disk spin axis as a function of time (based on the minor orientation runs of series #1), as well as the dark matter spin direction of the central region of the halo. The individual panels are Mollweide projections of the corresponding direction angles, one for each Aquarius halo. The origin of the projection has been shifted such that the initial position of the disk orientation lies at the center of the corresponding Mollweide map.

Interestingly, in several of the systems, the dark matter halo spin is quite well aligned with the minor axis. This is in particular the case for systems A, E, G, and H, and to a lesser degree in B. In halos C and D,
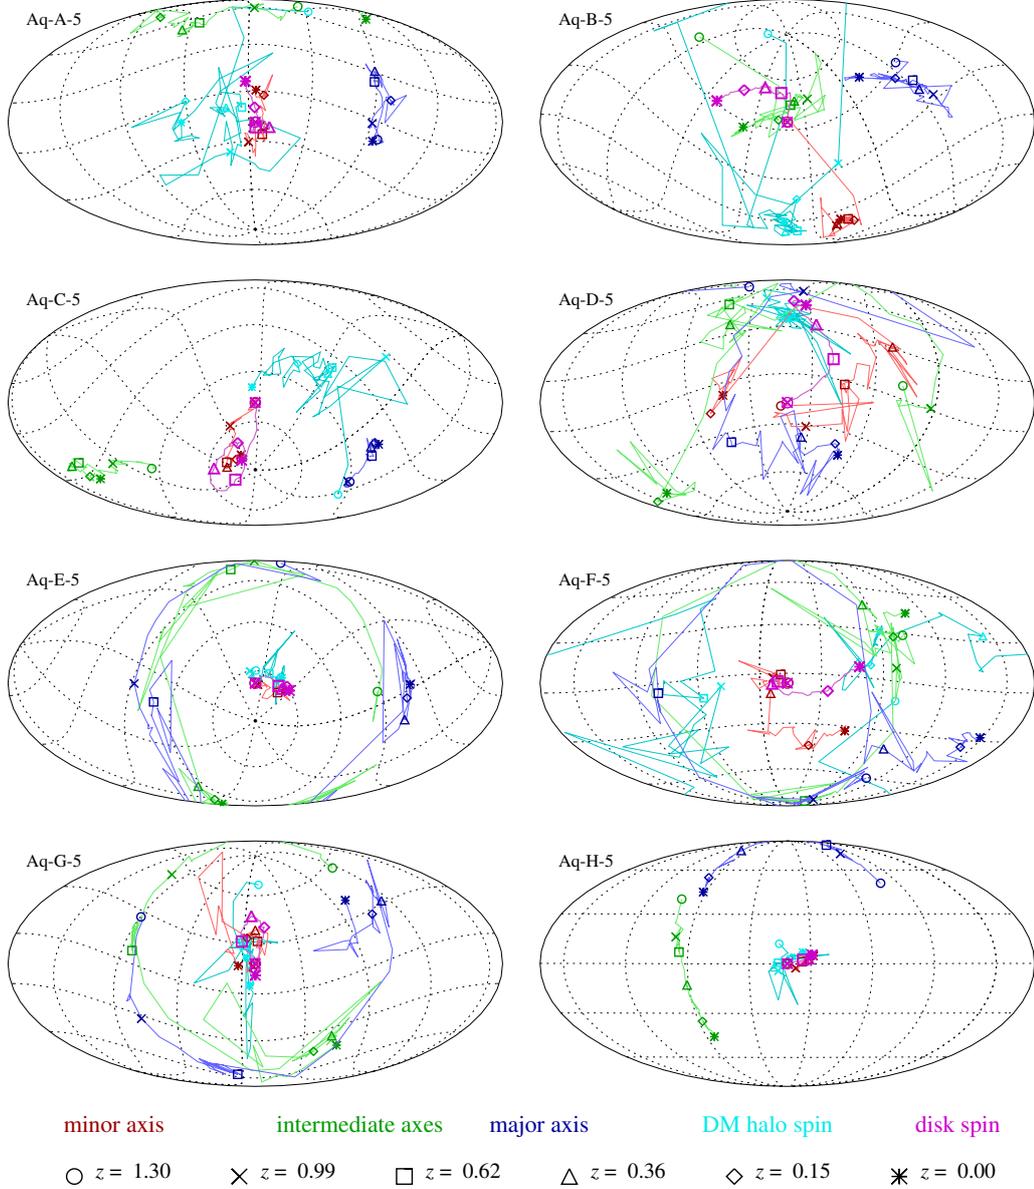
Figure 3.3: Orientations of the principal dark matter halo axes, the stellar disk spin vector, and the central dark matter halo angular momentum as a function of time between $z = 1.3$ and $z = 0$ in our eight Aquarius halos. In each panel, a different simulation corresponding to our default series #1 runs is shown, as labeled. In each of the displayed runs, a pure disk is inserted at $z = 1.3$ and grown to its final mass at $z = 1.0$ keeping its shape and orientation fixed in time during the growth phase. From $z = 1.0$ to $z = 0$, the disk is evolved live. Different redshifts are singled out with symbols, as labelled. In each panel, the vector orientations are shown in a Mollweide projection of the unit sphere, with the initial orientation of the disk aligned with the centre of the map.
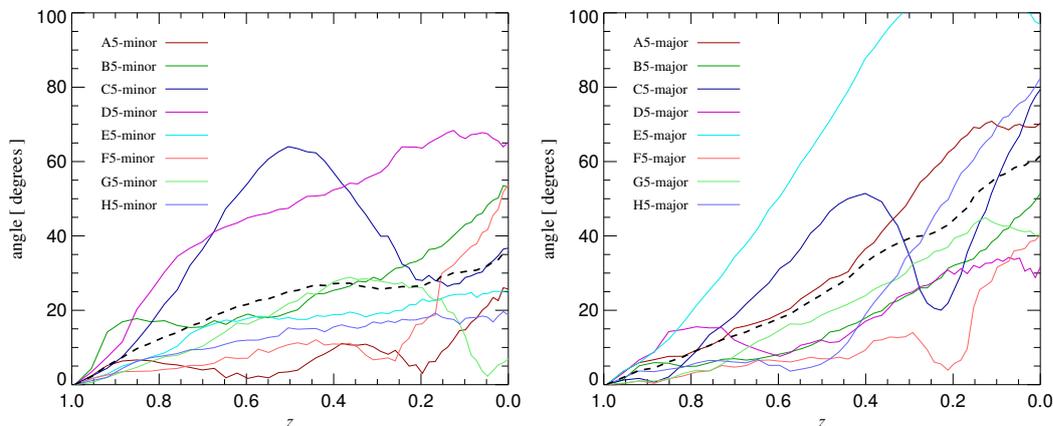
Figure 3.4: Angle between the current disk spin axis and the initial disk orientation when the disk goes live (at $z = 1.0$), in our default disk models of simulation series #1, as a function of time. The left panel shows our results if the disks are aligned along the minor axes of their hosting dark matter halos, the right panel is for the major axes. Most of the disk galaxies experience substantial tilting during the 6 Gyr of evolution from $z = 1$ to $z = 0$, independent of whether they are initially oriented along the minor or the major axis of the dark matter halo.

the dark matter spin is reasonable stable in orientation but offset from the minor axis, while in F it wanders all over the place. Studies of cosmological halos have long found a preference of the dark matter halo spin to line up with the minor axis (e.g. Hayashi et al. 2007), a trend also seen here. Interestingly, there is growing observational evidence from alignment studies of SDSS galaxies (Zhang et al. 2014) that favor a picture where disk spins line up with the dark matter halo angular momentum of the inner regions of halos, suggesting that the minor orientation arises naturally and is actually preferred.

What is also evident from Figure 3.3 is that the disk spins of several of the models wander away significantly from their original orientation, especially in those cases where the minor axis shows a change of orientation as well. In cases where it moves little, such as E and H, the disk spin exhibits a fairly stable orientation. A clearer view of the size of the changes in orientation is given in Figure 3.4, where the angle between the current disk

orientation and its initial orientation is shown as a function of time, from redshift $z = 1$ to $z = 0$. The left panel shows our results when the disk is initially aligned with the minor orientation, the right panel is for the major orientation. Some of the models show rather substantial reorientations of the disks reaching up to 65 degrees in the minor cases, and even larger angles beyond 110 degrees for the major orientation. The average tumbling angle (shown as a dashed line) for the eight systems is 35 degrees for the minor orientation, and 60 degrees for the major orientation. The substantially smaller average tilt angle for the minor orientation suggests that this orientation typically offers better long-term directional stability than the major orientation. We note that our results for A-D show a large resemblance to those of DeBuhr et al. (2012), for example, we find the same characteristic evolution pattern for halos a A and C. This is reassuring, given the independent and at a technical level quite different methodology to introduce and simulate the live disks. However, there are also some quantitative differences, and for a subset of the systems we tend to find somewhat smaller angles than DeBuhr et al. (2012).

3.4.2 Radial and vertical structure, and its evolution

As pointed out by DeBuhr et al. (2012), it is perhaps not too surprising that these systems show such strong tendencies to form bars. Whereas their rotation curve structure, shown in Figure 3.5, in principle suggests that the disks are not exceeding the rotation curve contribution of the dark matter anywhere (apart from halo G for a small region) and are thus far away from being maximal disks, the simple criterion of Efstathiou et al. (1982) for stability against bar formation,

$$Q_{\mathrm{bar}} \equiv \frac{v_{\max}}{(GM_d/R_d)^{1.2}} > 1.1 \tag{3.1}$$

is violated for all the models. Here $v_{\max}$ is the maximum rotation curve velocity, and M_d and R_d refer to mass and scale-length of the exponential
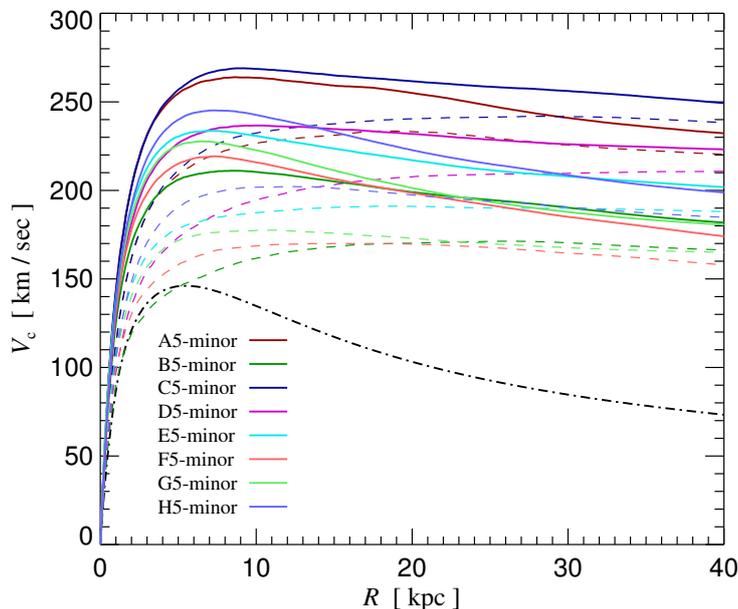
Figure 3.5: Rotation curves of our default disk models (series #1) at redshift $z = 1$, when the disk goes live. The rotation curve velocity is here defined based on the enclosed mass at a given radius. Note that the disk is kept the same in all these eight simulations, hence its contribution to the rotation curve (dot-dashed line) is always the same. The dashed lines show the contributions of the different dark matter halos, while the solid curves give the total rotation curves of the eight systems, with different colours as specified in the legend.

stellar disk. In fact, the values for $Q_{\rm bar}$ after the disks have been inserted are 0.99, 0.79, 1.0, 0.88, 0.87, 0.82, 0.85, and 0.92, for A to H, respectively.

In this context it is also interesting to look at Toomre's stability parameter for axisymmetric stability of stellar disks (Toomre 1964),

$$Q_{\rm Toomre} \equiv \frac{\sigma_R \, \kappa}{3.36 \, G \Sigma} > 1, \tag{3.2}$$

where σ_R is the radial velocity dispersion, Σ the surface density, and

$$\kappa^2 = \frac{3}{R} \frac{\partial \Phi}{\partial R} + \frac{\partial^2 \Phi}{\partial R^2} \tag{3.3}$$

is the epicycle frequency. The value of $Q_{\rm Toomre}$ is shown as a function of radius in Figure 3.6. Interestingly, most of the models are Toomre stable,
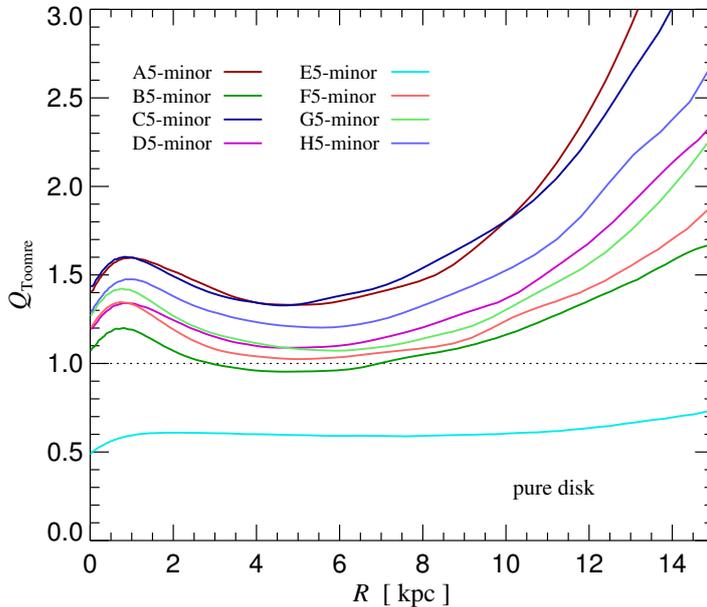
Figure 3.6: Radial run of Toomre's Q-parameter (Toomre 1964) for axisymmetric stability for our disk simulations of series #1, at $z_{\text{live}} = 1$. Model B5 is expected to be marginally unstable against axisymmetric instabilities, whereas E5 should be strongly unstable. Evidence for this is in fact seen in Figure 3.1, where E5 shows residual axisymmetric ring-like features in its disk at $z = 0.85$ that are not present in this form in the other runs.

with B and F being marginal cases, but the light halo E is clearly predicted to be unstable against axisymmetric instabilities. And indeed, inspecting the stellar images at $z = 0.85$ in Fig. 3.1 one can clearly see ring-like spherical features that are absent in this form in the other models, providing evidence that such instabilities have occurred in the early evolution of the system.

We find further signs for this special evolution of halo E in the evolution of the structural properties of the systems, which we examine next. In Figure 3.7 we show the radial and vertical density profiles, at a set of different times. The exponential surface density profile measured for the face-on orientation of the disks is quite robust and more or less retains its
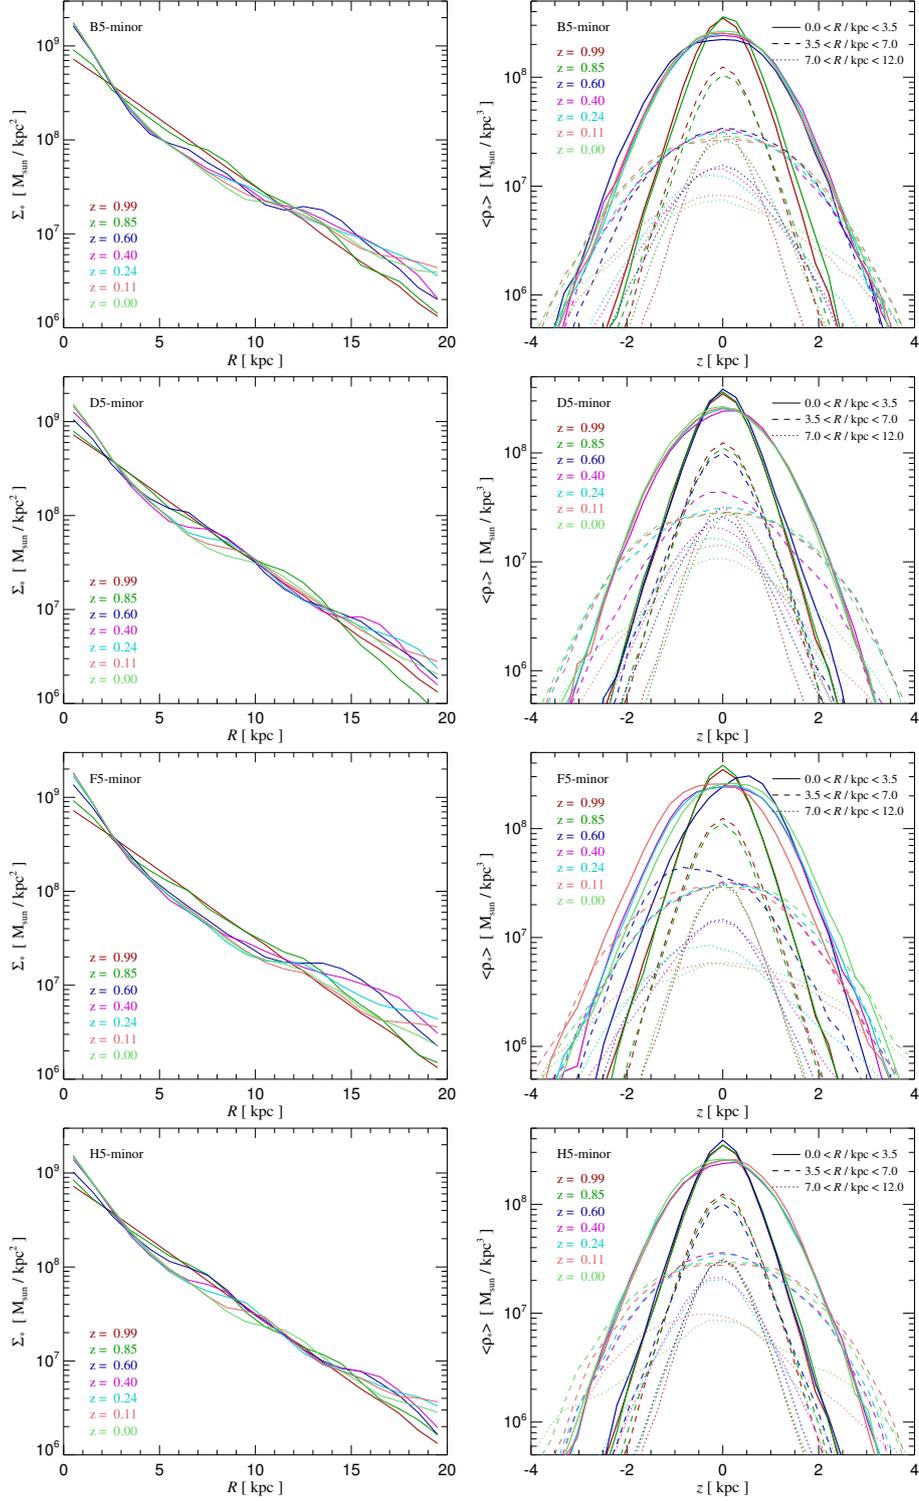
Figure 3.7: Radial and vertical disk density profiles for our default disk simulations (series #1) as a function of time. We show results for all of our eight Aquarius halos, in each case with a pair of panels where the surface density profile is shown on the left, and the vertical density profiles on the right. For the latter, three families of curves are shown, corresponding to averages over different radial ranges ($R < 3.5 \, \mathrm{kpc}$, $3.5 \, \mathrm{kpc} < R < 7 \, \mathrm{kpc}$, and $7 \, \mathrm{kpc} < R < 12 \, \mathrm{kpc}$, as labelled).
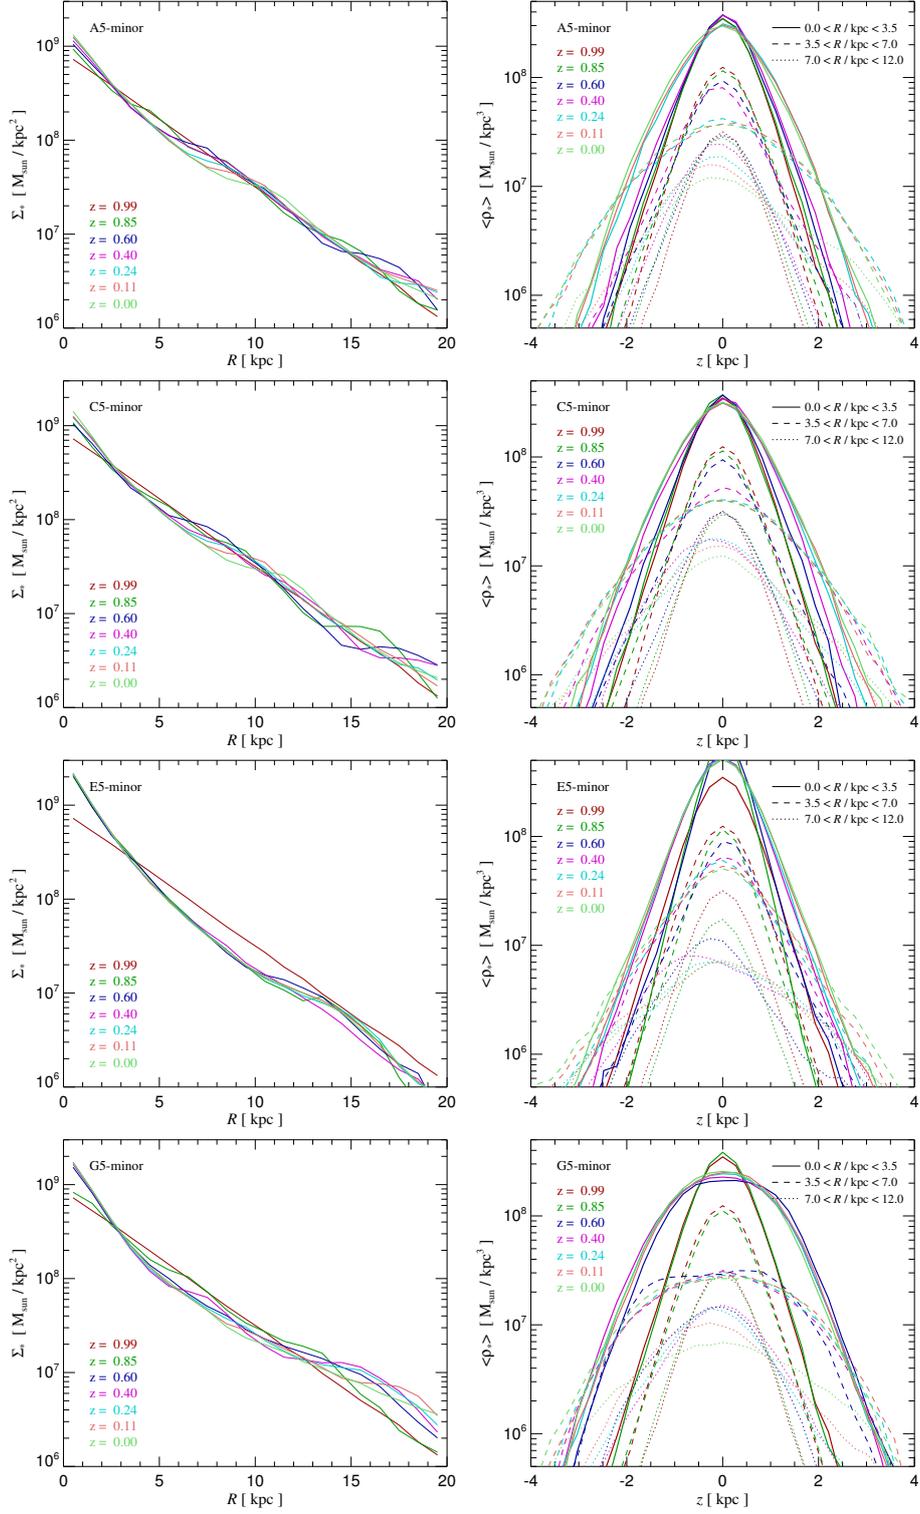
Figure 3.7 (continue): In all panels, the initial density distribution when the disk goes live is shown together with 6 subsequent times down to $z = 0$, and at each time, the disk plane has been defined based on the angular momentum of the stars in the inner 5 kpc.
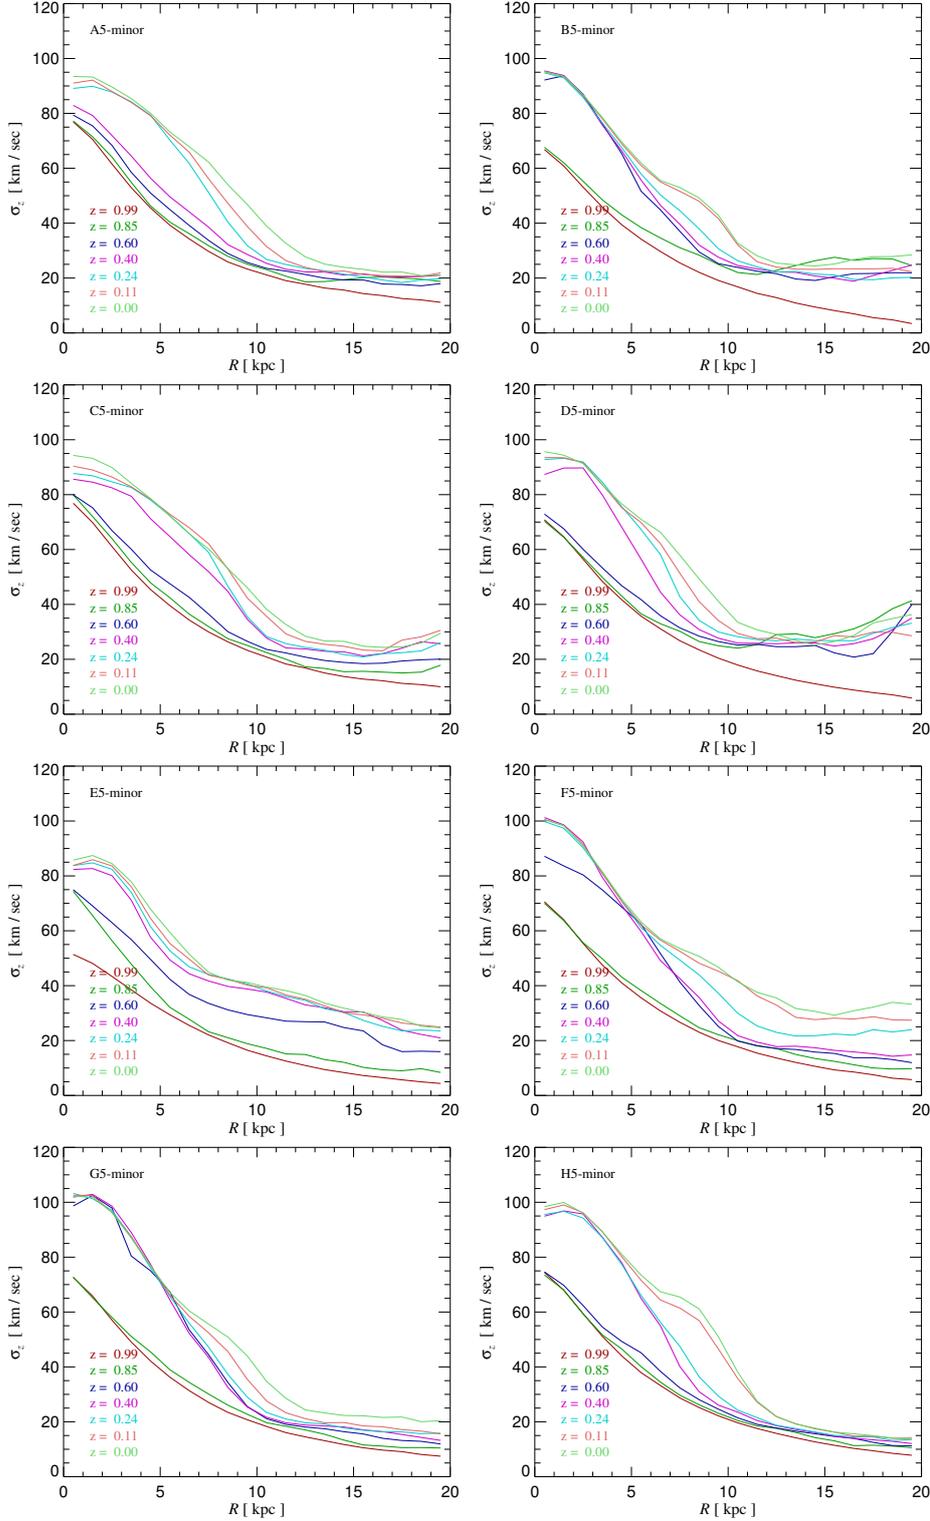
Figure 3.8: Evolution of the vertical velocity dispersion profile of the disk particles in our default disk insertion simulations (series #1). The panels give the expectation value of the vertical stellar velocity $\sigma_z = \langle v_z^2 \rangle^{1/2}$ measured in in different cylindrical shells aligned with the stellar spin axis of the stars at the corresponding times.
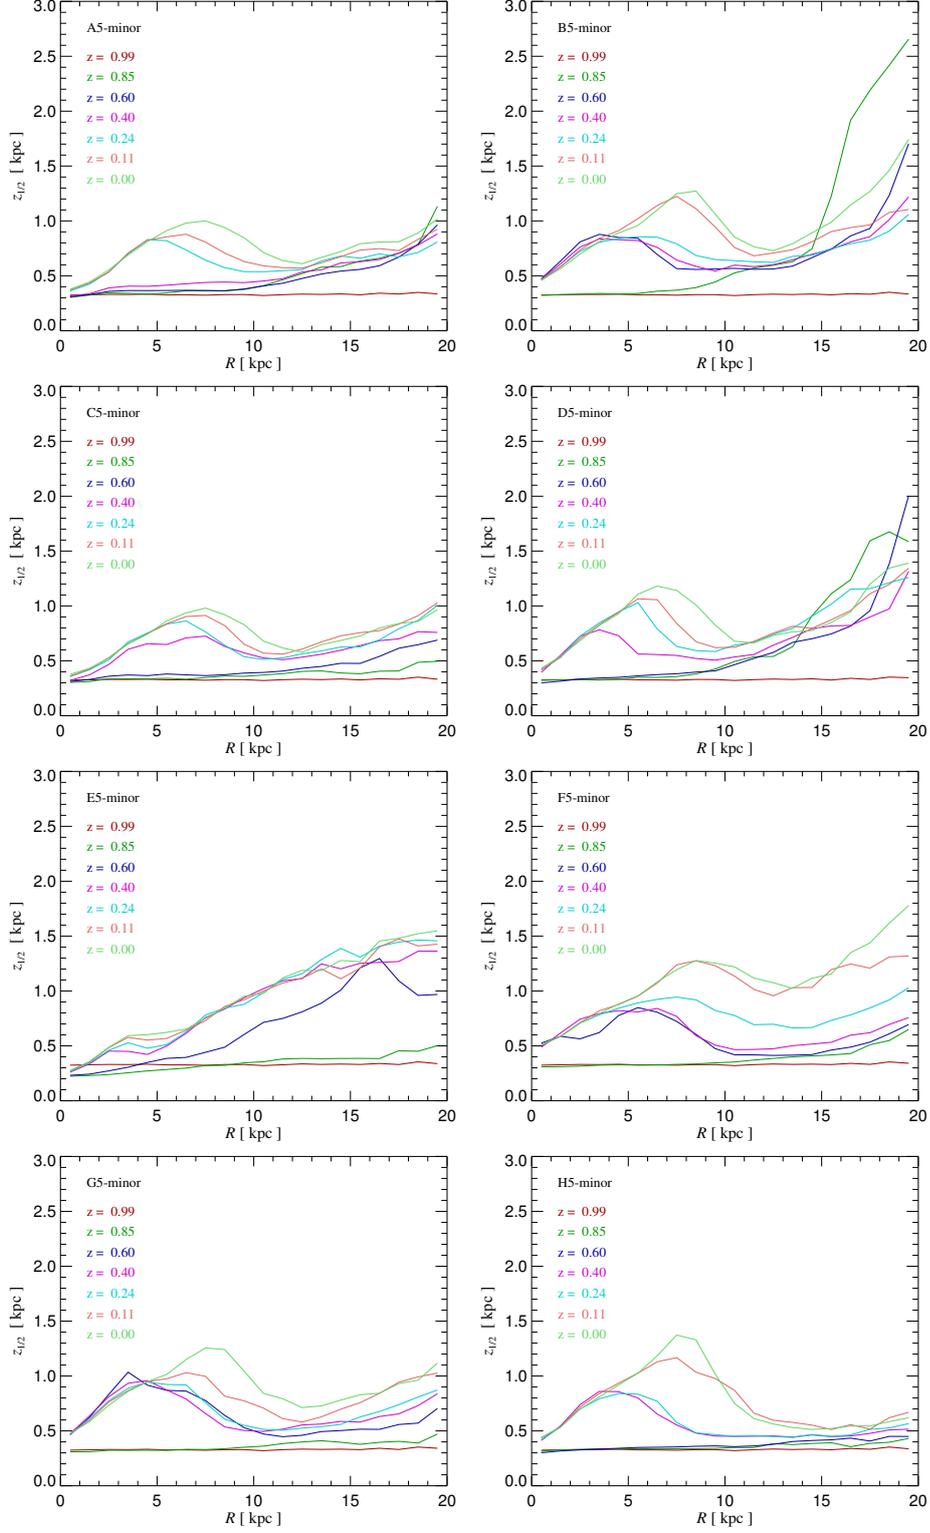
Figure 3.9: Evolution of the half-mass height profile of the disk particles in our default disk insertion simulations (series #1). The height profiles give the median of $|z|$ relative to the disk plane, i.e. half the stars have distances below/above $z_{1/2}$ from the disk plane.

initial shape, despite the rather dramatic bar formation events occurring in these simulations. In contrast, the vertical density profiles (three families of curves are shown, corresponding to different radial ranges, as labeled) show the damaging impact of the forming bars more clearly. In particular, a relatively sudden transition to a new vertical equilibrium with a thicker profile is apparent in most of the models. One interesting difference with DeBuhr et al. (2012) is that our models A-D show substantially less broadening in the outer parts of the disks. This is presumably a reflection of our more accurate approach to initialize the velocities of the initial disk models.

Further support for this is provided by the evolution of kinematic quantities, for example those shown in Figure 3.8 and Figure 3.9. The eight panels in Fig. 3.8 give the evolution of the profiles of the vertical velocity dispersion σ_z^2, while the eight panels of Fig. 3.9 measure the disk height in terms of the median $z_{1/2}$ of the absolute value of the vertical z-coordinates of the star particles relative to the disk plane. Or in other words, half of the stellar mass has a height above the central disk plane less than $z_{1/2}$. Comparing again to DeBuhr et al. (2012), we see that our $\sigma_z(r)$ profile at $z = 0.85$ is much closer to the initial profile than in their case, suggesting that our disk models are in better dynamical equilibrium initially. This equilibrium is however anyway destroyed relatively quickly by the onset of bar formation. Another notable difference are the much smaller values for $z_{1/2}$ we find in the outer parts of the disks, corroborating the observation that our disk models appear to be in better equilibrium.

3.4.3 Bar strength and vertical heating

The formation of the bar can also be studied more quantitatively, for example through measuring a bar strength indicator, or by looking at the vertical heating of the stellar disk. As a simple global measure of the total amount of vertical heating we can use the quantity

$$\zeta = \frac{\langle v_z^2 \rangle}{\langle v_z^2 \rangle_0},\tag{3.4}$$
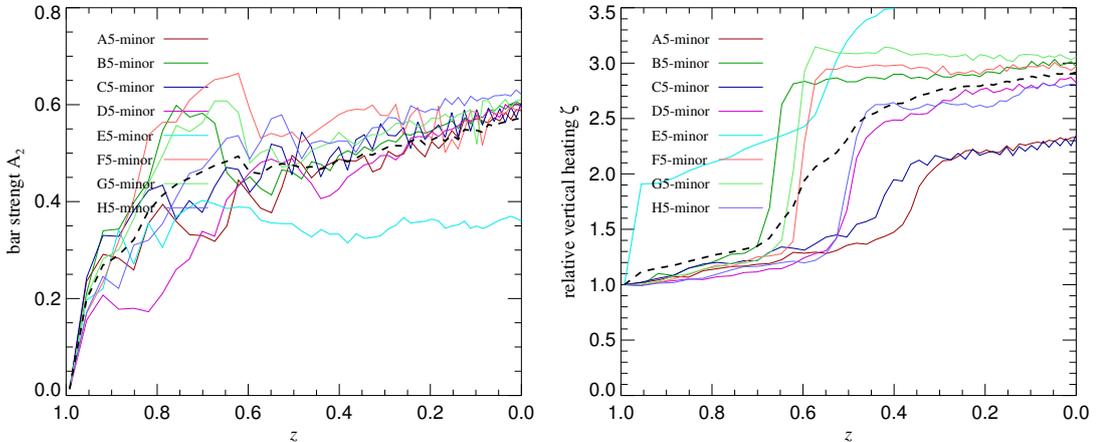
Figure 3.10: Bar strength parameter A_2 (left panel) and relative vertical heating (right panel) in the simulations of our default pure disk insertion simulations (series #1), as a function of time. It is clearly seen that all the models develop a strong bar characterized by $A_2 \simeq 0.6$, except for model E, which yields $A_2 \sim 0.3$ at the end. The latter model is special as it shows substantial vertical heating right after the disk becomes live. This is because this system is instable against axisymmetric instabilities (see Fig. 3.6).

which simply is equal to the total kinetic energy in vertical motion relativ to the initial value of this quantity at the time the disk starts to evolve live.

As a characterization of the bar strength, we adopt a simple measure for the $m = 2$ Fourier mode of the disk, as is often done to quantify bars. For definiteness and ease of comparison, we measure the bar strength similarly as DeBuhr et al. (2012), by first determining the quantities

$$a_2^{(b)} = \sum_{i \in b} m_i \cos(2\phi_i) \tag{3.5}$$

$$b_2^{(b)} = \sum_{i \in b} m_i \sin(2\phi_i) \tag{3.6}$$

for a set of 30 radial bins between $R = 0$ and $R = 2\,R_d = 6\,\mathrm{kpc}$. Here ϕ_i refers to the azimuthal angle of each disk star, and the sums extend only over the particles in a radial bin b. Defining $c_2 = \sqrt{a_2^2 + b_2^2}$, we then

104

calculate a bar strength parameter as

$$A_2 = \frac{\sum_b R_b \, c_2^{(b)}}{\sum_b' R_b' M_b'} \tag{3.7}$$

where M_b is the mass falling into bin b, and R_b is the bin radius.

In Figure 3.10, we show our results for the time evolution of the A_2 parameter in the left panel, and in the right panel we give the time evolution of the vertical heating parameter ζ. Consistent with the evolution of the visual morphology, all the systems show a rapidly growing bar signal in their early evolution, with all the models except E converging to a similarly high bar strength of $A_2 \simeq 0.6$ at the end.

Interestingly, the relative vertical heating parameter shows only a small and slow growth in the beginning, but then jumps up rapidly by a large factor, followed by a stabilization at a new high level. In contrast, the bar strength indicator becomes high already significantly earlier. Presumably this is a combination of two different effects. One is that some disk asymmetries quickly develop in the early disk evolution simply because of the non-sphericity of the halo potential, which for example manifests itself in pronounced spiral patters in the disks. The other is that during the initial phase of bar formation the density contrast of the bar grows without yet leading to a notable change in the vertical structure. Only later, once the bar "collapses" or buckles, a sudden transition to a new equilibrium occurs, and this is associated with some degree of violent relaxation and significant vertical heating.

3.5   Models with bulges and lighter disks

The strong bars formed in the models considered in the previous section raise the question under which conditions disk galaxies could survive in the Aquarius halos and maintain a nice, disky morphology all the way to $z = 0$. It appears clear from the preceding results that the bar criterion of Efstathiou et al. (1982) should be taken as an important first guide. In-

creasing the value of Q_{bar} in a given dark matter halo can be achieved first of all by making the disk lighter, and/or by making it larger. However, we note that there is only limited room for such changes if one wants to keep the disk parameters close to observational inferences, such as the stellar mass–halo mass relationship derived from abundance matching (Guo et al. 2010; Moster et al. 2010), or direct constraints on the size–stellar mass relationship (Shen et al. 2003). Another problematic aspect of simply adopting a very light disk is that this will greatly reduce the roundening effects of the dark matter halo due to the growth of a baryonic mass distribution in the centre. Such a rounding is however essential to keep the destructive effects of a highly triaxial dark matter potential on disk stability and morphology at bay.

Another approach to improve the stability of the disk against bar formation is to add a central stellar bulge in addition to the disk. Such a bulge increases the circular velocity of the spheroidal component of the system, i.e. it contributes to the numerator of the bar stability criterion. At the same time, growing a central spherical bulge in a dark matter halo is an effective way to rounden the dark matter potential. Of course, on the other hand, adding a central bulge is limited by the observed bulge-to-disk ratios, and is not a promising option for explaining pure disk galaxies.

To test these options, we have considered a few additional sets of runs. Our default standard bulge models (series #3) have the same stellar mass as our default disk models, but one third of the stellar mass is put into a bulge, with two thirds remaining in the disk. Such bulge-to-disk ratios are about the smallest ones that the present generation of hydrodynamical simulations of disk formation can achieve (e.g. Marinacci et al. 2014a).

In addition, we have also run models were we omitted this bulge, which is exploring the lighter disk option at some level. Here the disk mass was reduced to 2/3 of the value in our default models (series #4). This is complemented by a further set of runs in which we swapped the masses of disk and bulge of our standard bulge models, ending up with 1/3 in the

disk and 2/3 in the bulge (series #5). We would expect these models to be extraordinarily stable against bar formation.

Finally, to test the influence of the residual triaxiality of the dark matter halos on the disk stability, we have also run two sets of simulations where we artificially sphericalized the dark matter halos at the instant of disk insertion, i.e. at $z = 1.3$ (series #6 and #7). To this end, we simply took all dark matter particles in the FOF halo at that time and rotated them randomly around the halo center. Also, their velocities were isotropized by turning them randomly in the rest frame of the halo. In this way, the density structure and potential energy of the halo was approximately maintained, with any deviation from non-equilibrium decaying away during the disk growth phase. When the disk goes live at $z = 1.0$ it then does so in an essentially spherical halo, but the mass growth of the halo, including the accretion of newly infalling substructures onto the halo, stays unaffected.

We begin our review of the results of these runs with images of the disks in the runs with the small bulges. These are shown in face-on and edge-on orientations in Figures 3.11 and 3.12, respectively. Note that the bulge stars are not included in these images. Compared to the corresponding images for the simulations with pure disk, the disks are clearly much less effected by bulge formation, although most systems still do form bars, albeit of weaker strength and at later times. The edge-on projects of the disk stars correspondingly show much less evidence for central bars.

The rotation curves of the models with bulges, as well as the corresponding profiles of the Toomre stability parameter are shown in Figure 3.13. None of the models is any more dominated by baryon at any radius, and as expected, the resistance against axisymmetric instabilities is increased substantially. Note however that halo E is still predicted to be Toomre unstable. This is confirmed by visual inspection of simulation E at $z = 0.85$ in Fig. 3.11, which reveals conspicuous axisymmetric rings of stars.

The weaker bar strengths in the bulge-runs also become evident in Figure 3.14, which shows the time evolution of the bar strength parameter A_2
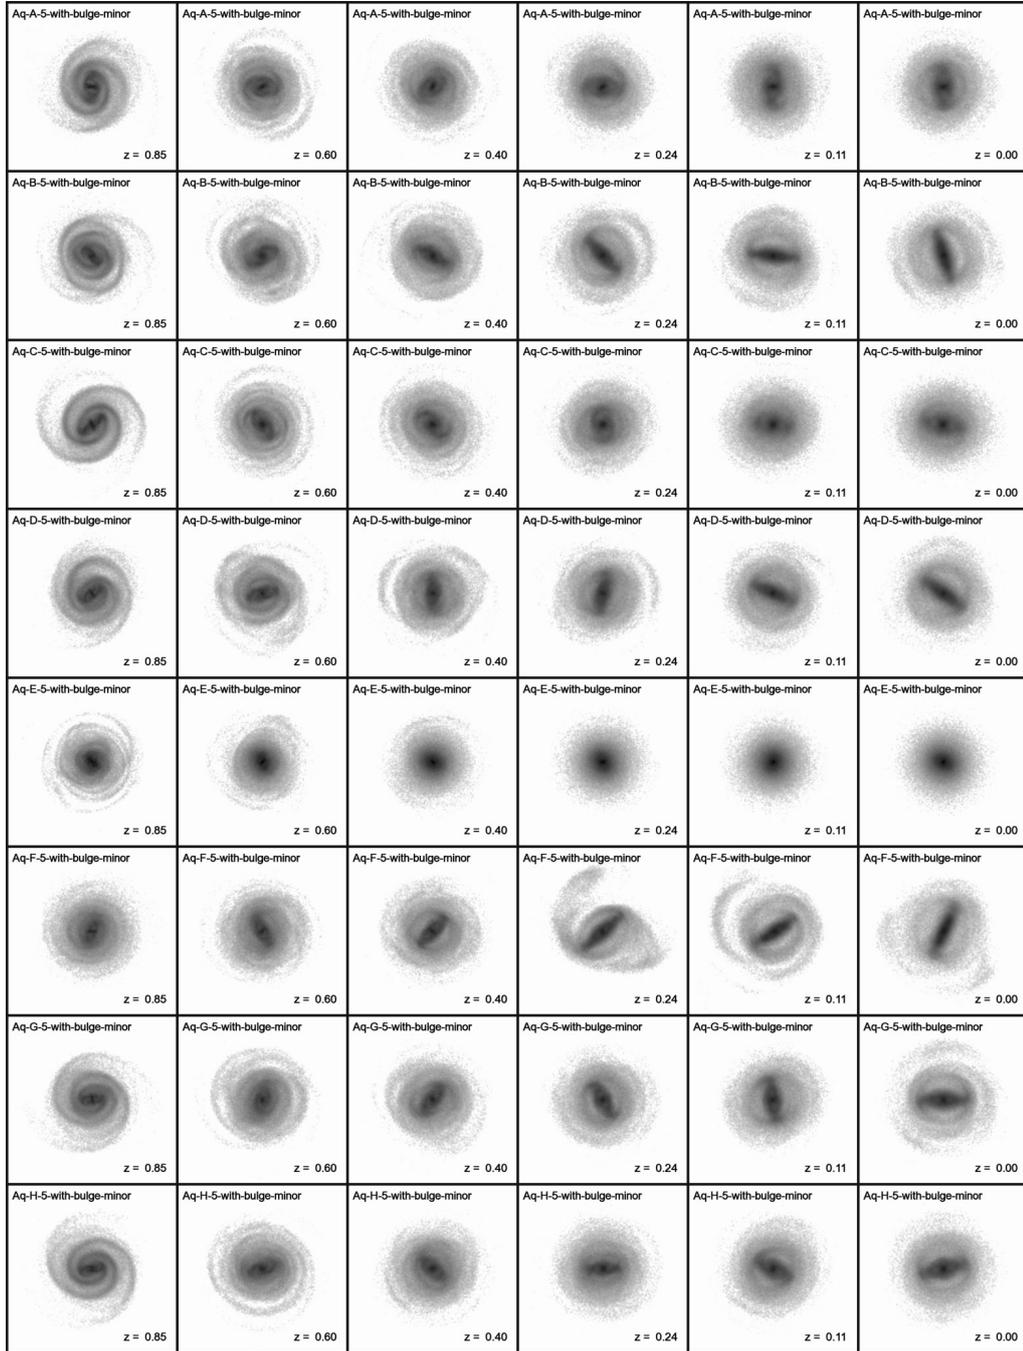
Figure 3.11: Face-on projections of the stellar mass of the disks in our runs with bulges (series #3), where one third of the stellar mass is moved to a central spheroidal bulge and two thirds are kept in the disk. Otherwise, the outline of the figures and the image generation method corresponds exactly to that of Fig. 3.1.
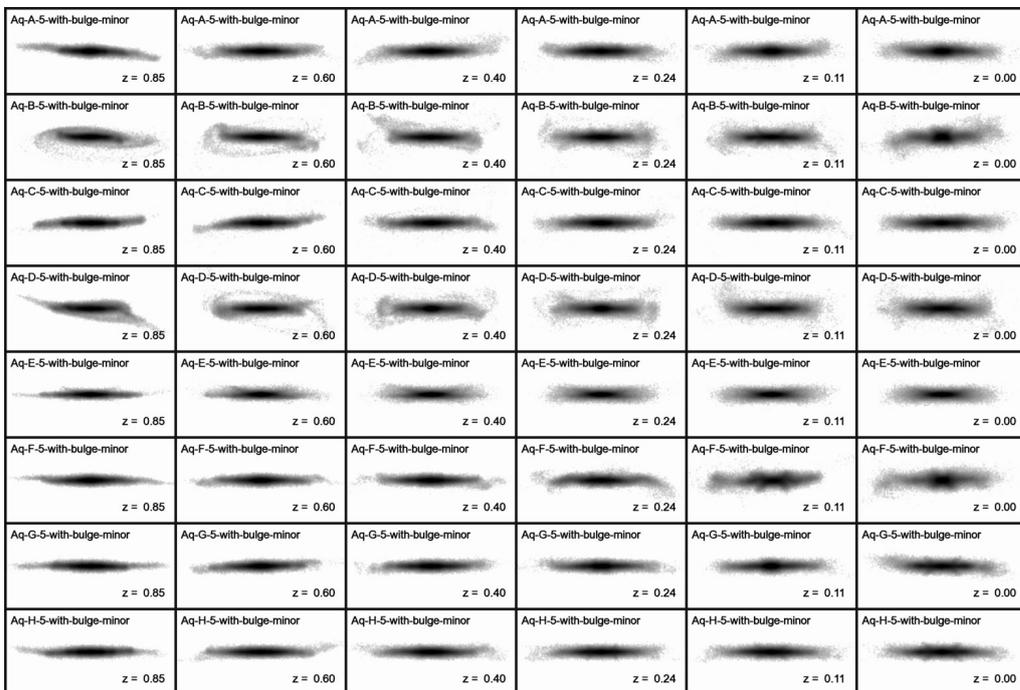
Figure 3.12: Edge-on projections of the stellar disks in our runs with bulges, giving the side-views of the corresponding images shown in Figure 3.11. As in the corresponding images of Fig. 3.2, each panel has an unchanged physical size in the horizontal direction (54 kpc) and uses an identical color-scale as in the face-on images.

and the vertical heating parameter ξ of the disks of our eight standard bulge models. Comparing to the corresponding results for the pure disk models (Fig. 3.10) clearly shows a much weaker trend towards making strong bars. Most notably the early evolution is very different, where all the models show a small value of A_2 as a result of deviations from axisymmetry due to the aspherical dark halo potential, but no strong bulge signal is present yet. Only after a few Gyr, a subset of the models starts to grow a significant bar.

It is also interesting to examine differences in the evolution of the structural properties of the models with bulges. This is shown in Figure 3.15 in terms of the radial and vertical density profiles, and in Figure 3.16 and Figure 3.17 in terms of the vertical velocity dispersion and half mass height
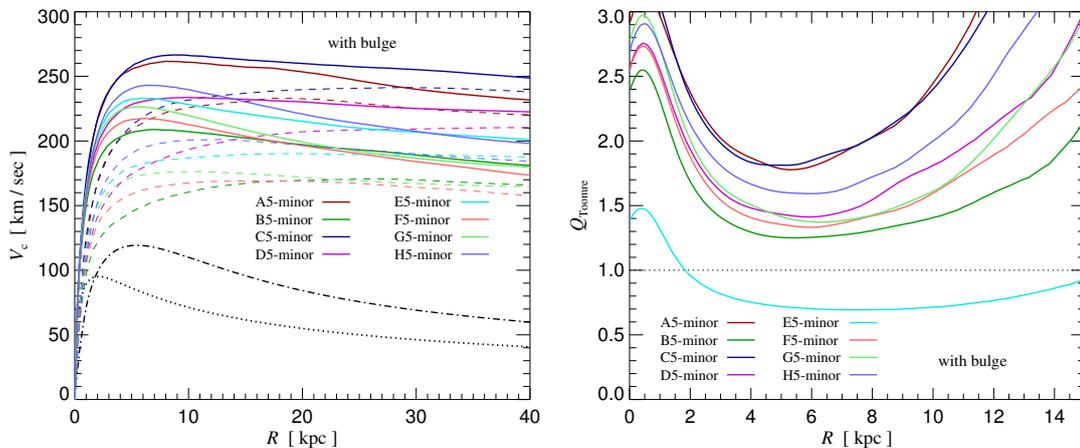
Figure 3.13: Rotation curves and Toomre-Q stability parameter of our default simulations with bulges (series #3), as a function of radius. The left panel shows the rotation curves, where the disk (dot-dashed) and bulge (dotted) always make the same contribution due to their constant parameters. The different dark matter halo contributions are shown by dashed lines, and the total rotation curves by solid lines. The panel on the run illustrates the expected stability against axisymmetric instabilities. While the bulge has a substantial stabilizing influence (compare to Fig. 3.6), model E5 is still found to be unstable for the adopted stellar parameters of disk and bulge.

as a function of radius. We see a very substantial improvement in the structural stability of the models with respect to the runs with pure disks. This is reflecting both, the stronger stability against bar formation, and the slightly rounder dark matter halo potential due to the more concentrated distribution of the baryons.

Getting completely rid-off the bars is achieved in our models that adopt a massive bulge and a light disk (series #5). Here the disk images show a remarkable degree of stability. For conciseness, we refrain from showing their uneventful time evolution. Instead, we collect all our simulations in a single plot of the bar strength versus the initial Q_{bar}-parameter of Efstathiou et al. (1982), allowing us to assess how well this venerable criterion works in the context of full cosmological CDM models. Figure 3.18 compiles our different runs, with the bar strength A_2 measured at $z = 0$, and the Q_{bar}-
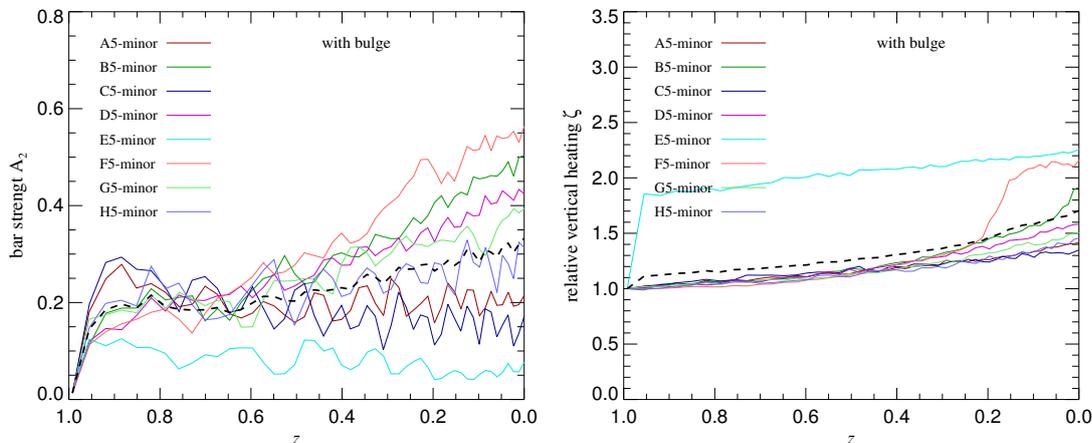
Figure 3.14: Bar strength parameter A_2 (left panel) and relative vertical heating (right panel) in our default disk+bulge simulations (series #3), as a function of time. These evolutions can be directly compared to the corresponding results for the pure disk case shown in Fig. 3.10. Now only a subset of the systems develops bars, and even if this happens, the bars are weaker and form later. Model E still appears as an outlier, caused by its instability against axisymmetric perturbations. We note that the early growth of the A_2 indicator to values of around ~ 0.2 does not really measure a bar; it is presumably caused by a quick distortion of the spherical disk into an ellipsoidal disk due to the the residual asymmetry of the dark matter halo potential.

parameter evaluated at $z = 1.0$ when the disk goes live. Remarkably, the threshold value 1.1 originally introduced by Efstathiou et al. (1982) for the dividing line between bar-unstable and stable models still serves as a surprisingly robust indicator, even in the light of all sorts of other complicating factors. In particular, we note that the different symbols in the figure show a broad range of simulation models, including runs with and without bulge, with minor or major axis orientation, with ordinary or artificially rounded dark matter halos, etc. Irrespective of these factors, it appears that the strength of the disk self-gravity relative to the supporting spheroidal potential is by far the most decisive parameter for governing stability against the formation of strong bars. We note that this therefore cannot be ignored in the interpretation of mass models derived for the Milky
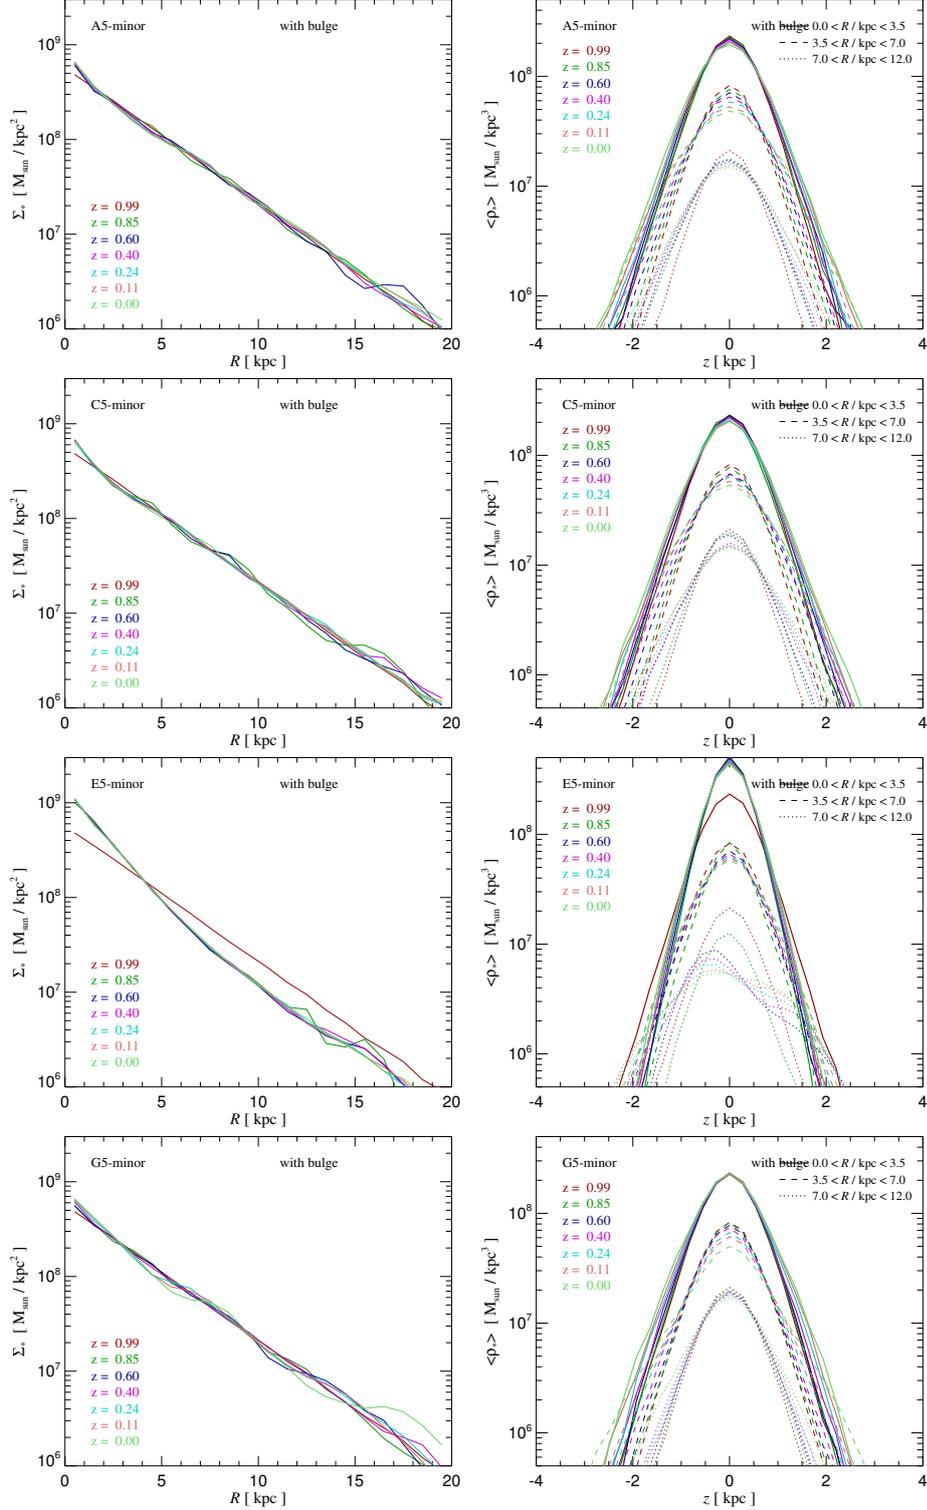
Figure 3.15: Radial and vertical disk density profiles for our default disk+bulge simulations (series #3) as a function of time. As in the corresponding Fig. 3.7 for the pure disk runs, we show results for all of our eight Aquarius halos, in each case with a pair of panels where the surface density profile is shown on the left, and the vertical density profiles on the right.
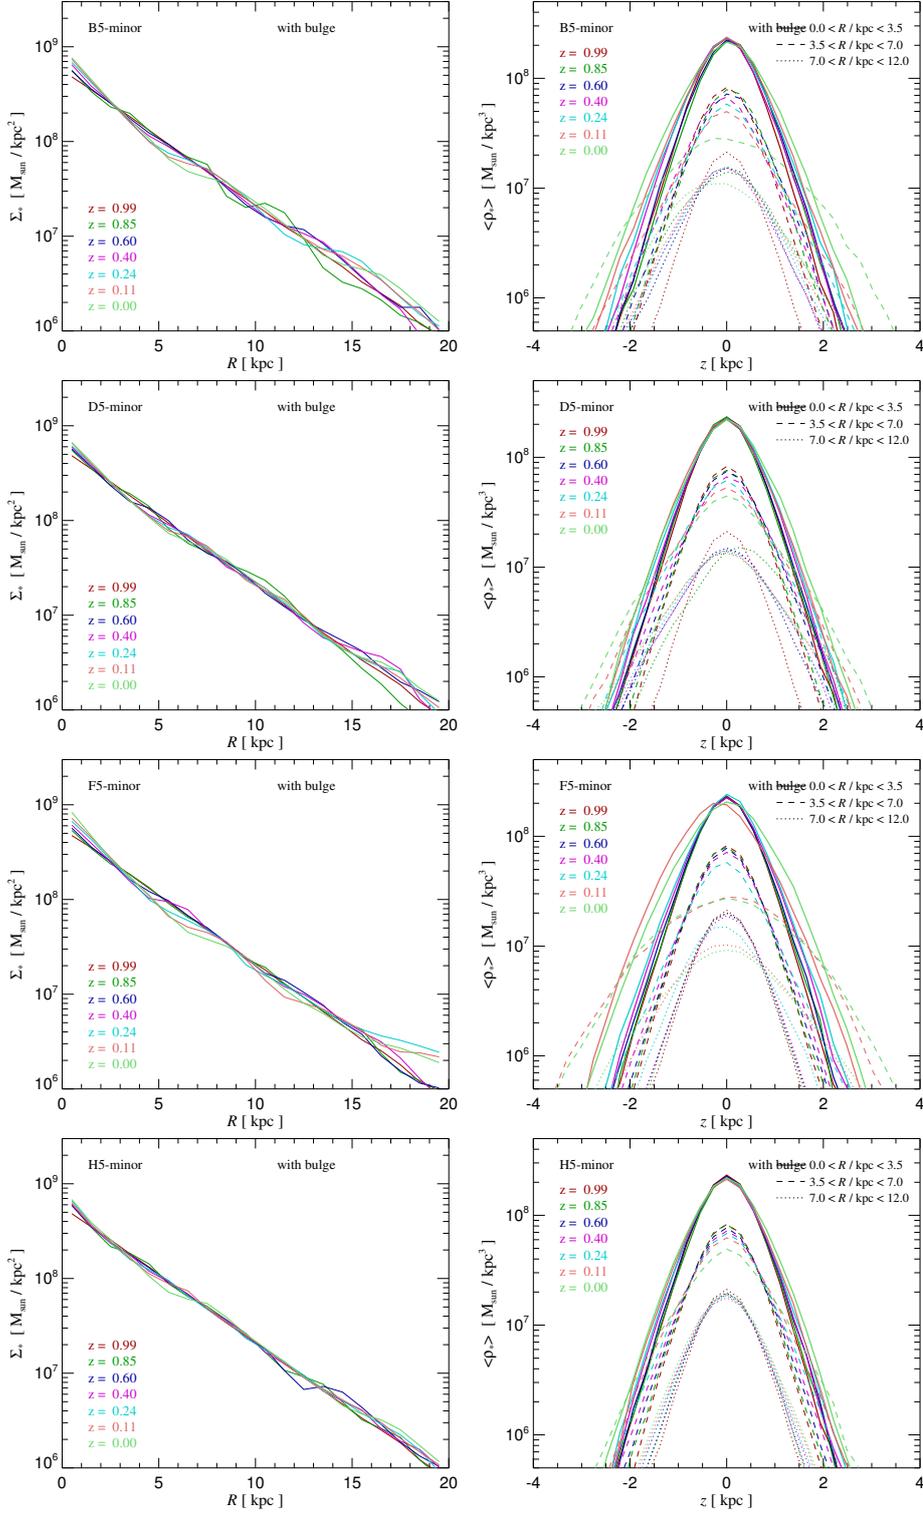
Figure 3.15 (continue): Especially in the much more modest evolution of the vertical structure it is evident that these disks evolve comparatively little; i.e. the bulge has largely stabilized the disks against strong bar formation.
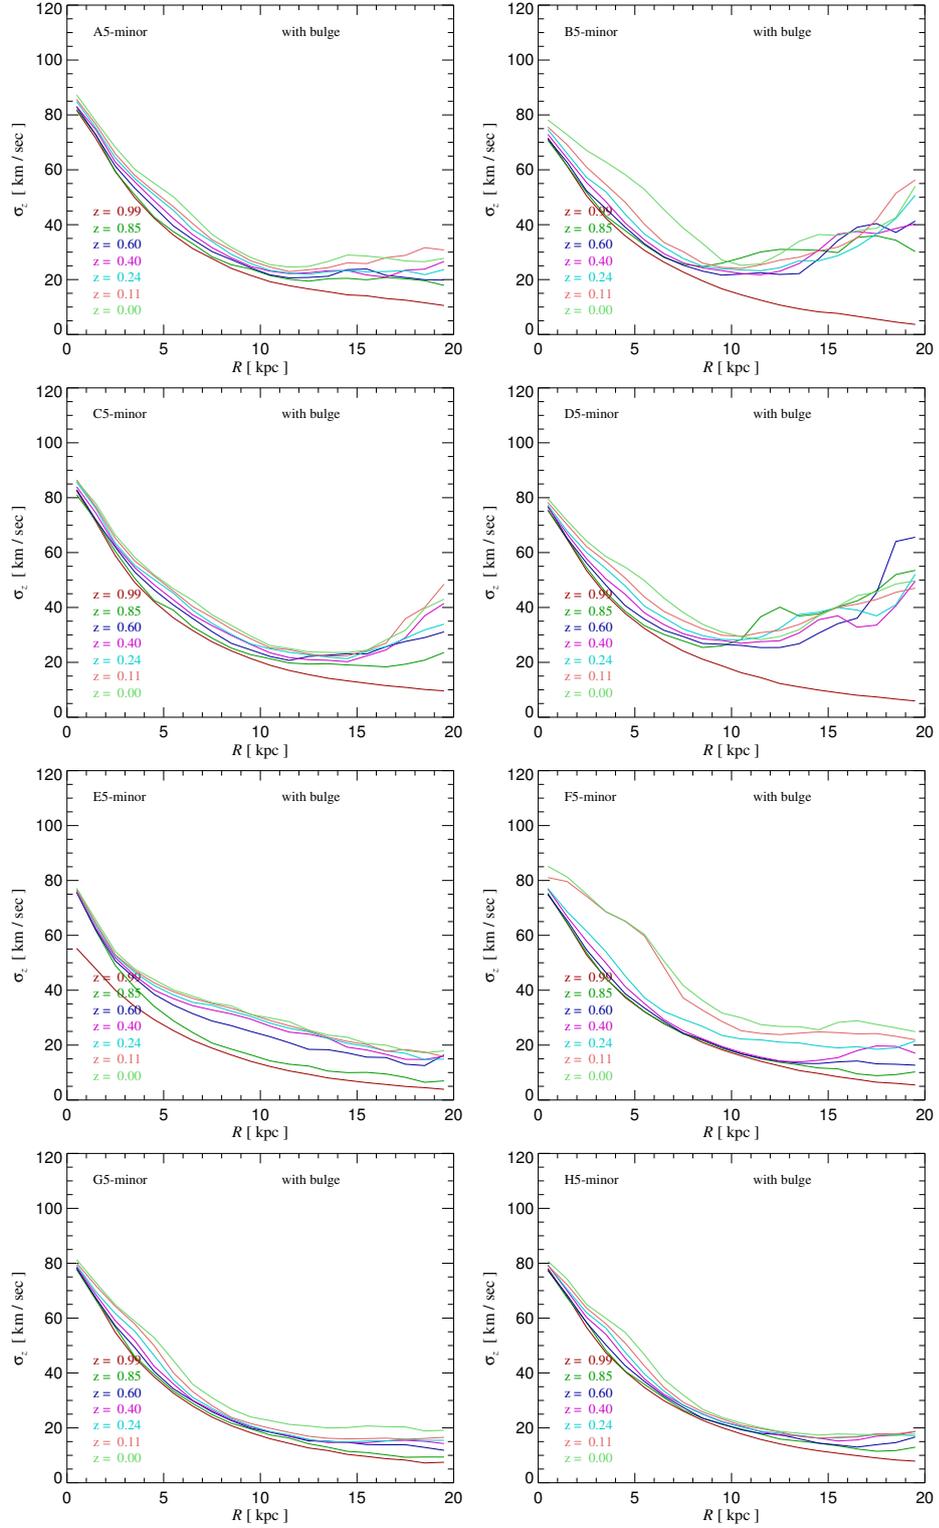
Figure 3.16: Time evolution of the vertical velocity dispersion profile and of the disk particles in our default disk+bulge simulations (series #3). These results correspond directly to the simulations for our pure disk simulations shown in Fig. 3.9. Note that only disk star particles are included in the measurements.
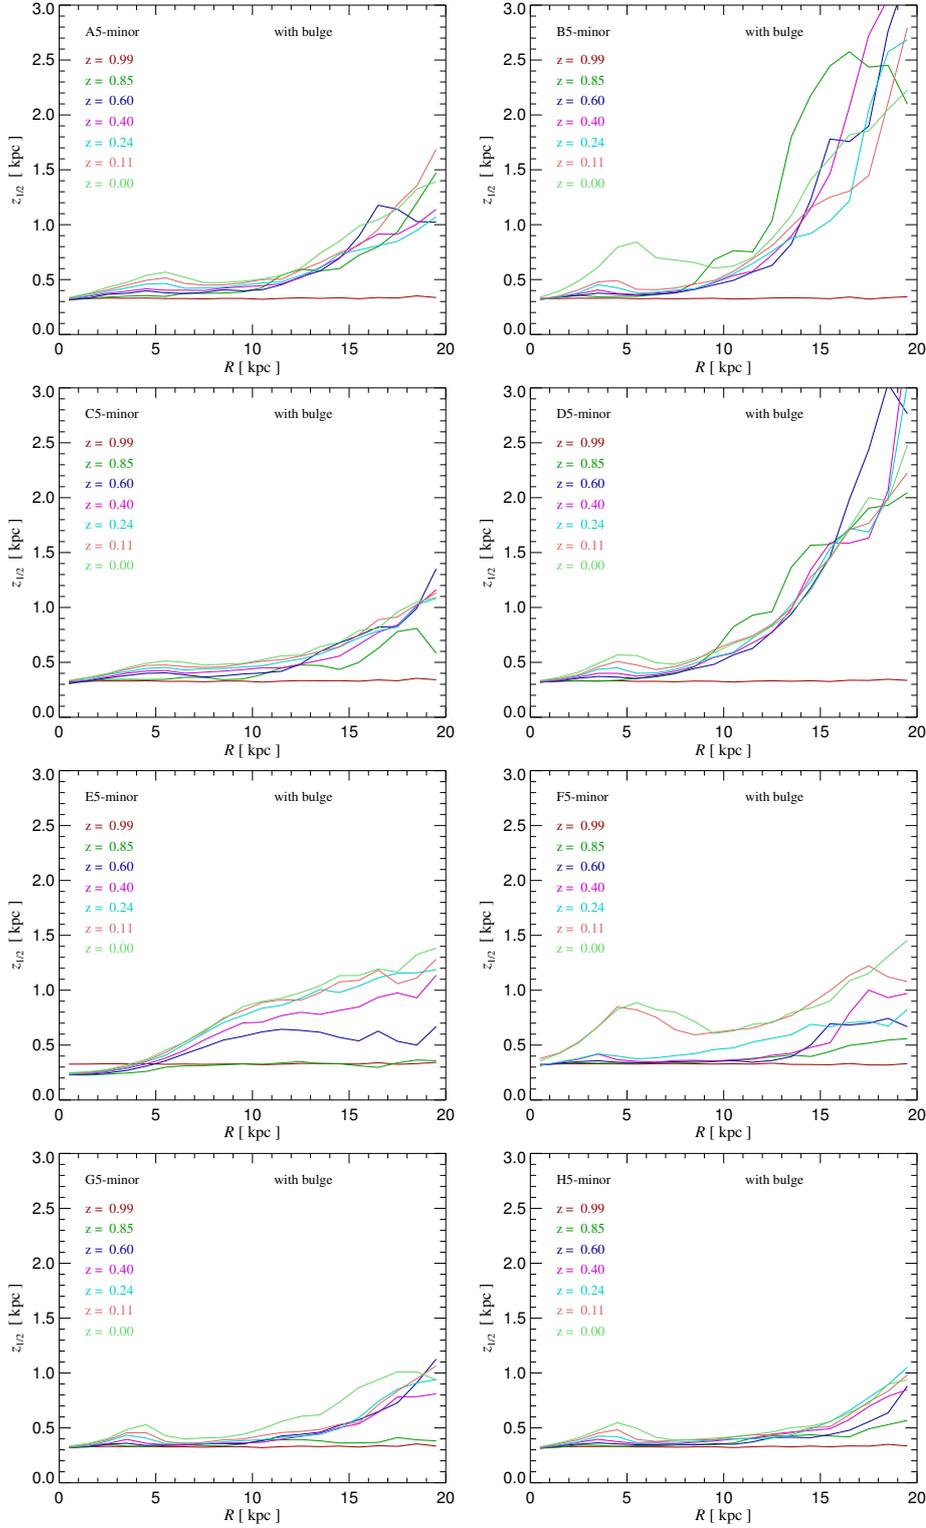
Figure 3.17: Time evolution of the half-mass height profile of the disk particles in our default disk+bulge simulations (series #3). These results correspond directly to the simulations for our pure disk simulations shown in Fig. 3.9. Note that only disk star particles are included in the measurements.

115
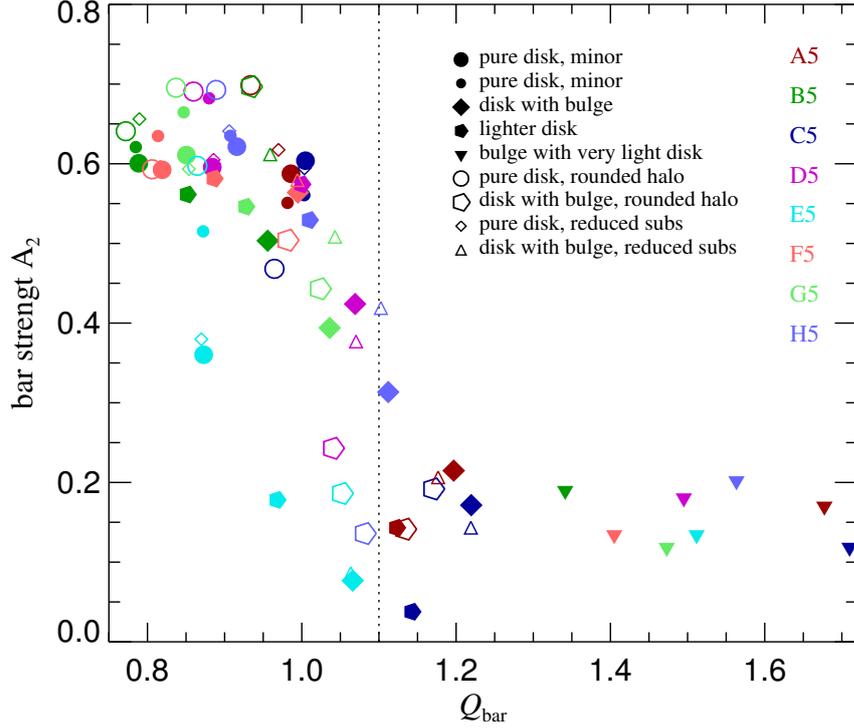
Figure 3.18: Compilation of the measured bar strength parameter A_2 at $z = 0$ in our simulation set versus the Q_{bar} parameter of Efstathiou et al. (1982). We here collect results for all of our runs (except for the late insertion and reorientation ones). Different colours are used for the different Aquarius halos, and different symbols for the different simulation series, as indicated in the legend (the number in the symbol key refers to series # in Table 3.1). It is clear that the criterion $Q_{\text{bar}} \geq 1.1$ formulated by Efstathiou et al. (1982) for indicating stability against bar formation works pretty well for our simulations. Other parameters such as minor or major disk orientation, or the residual a-sphericity of the dark matter halo, appear to be only of secondary importance.
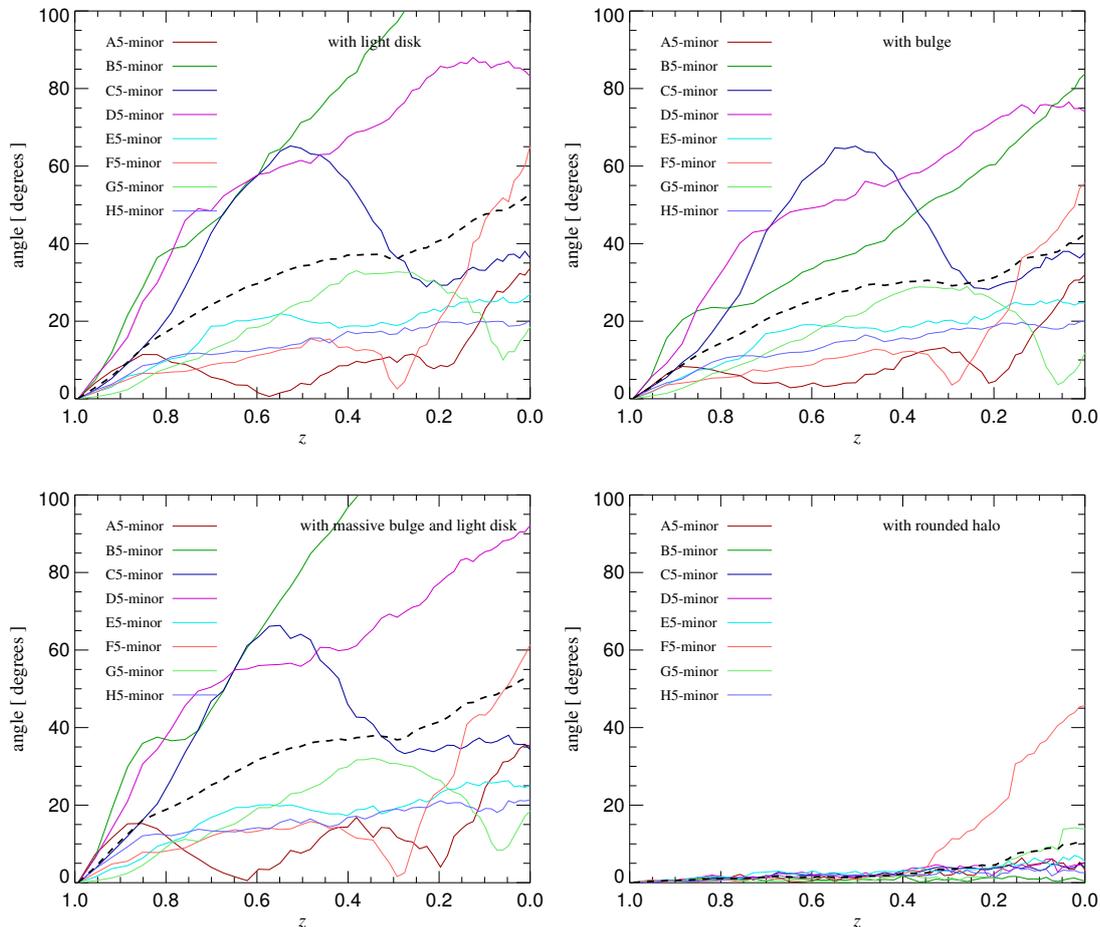
Figure 3.19: Angle between the current disk spin axis and the initial disk orientation as a function of time, for different types of simulations. The top left panel shows results for our simulation series #4, where a lighter disk is used compared with our default disk runs. The top right panel shows our default disk+bulge models instead (series #3), whereas the bottom left gives the tilt angle evolution for our simulations with massive bulges and a light disk (series #5). Finally the bottom right is for our default pure disk simulations, but this time the dark matter halo has been artificially rounded at $z = 1.3$ when the disk was inserted (series #6). Interestingly, the disk structural properties appear to play only a minor role for the disk tumbling (compare also to Fig. 3.4). The latter is clearly governed by the dark matter halo and could only be eliminated by substantial roundening that is far stronger than the effects achieved by the halo's response to the growing baryonic component.

117

Way. For example, the rotation curve decomposition derived by Bovy & Rix (2013) suggests that the corresponding live galaxy model should be violently bar unstable according to our results.

In light of these differences, it is now interesting to consider the stability of these models with respect to the spatial orientation of the disks. In Figure 3.19 we show results for several of our modified runs, including the ones with a lighter disks, the ones with a bulge, the ones with a very massive bulge, and the ones for a rounded dark matter halo. The latter are for the pure disk case (series #6), but the results for the bulge case (series #7) look essentially identical.

Comparing with the corresponding results in Fig. 3.4 for the pure disk case, it is evident that the amount of disk tumbling is fairly independent of the structural properties of the galaxies. In particular, it does not matter much whether a bar is present or not. Apparently, the reorientation of the disk is primarily controlled by the tumbling of the dark matter halo and the torques it exerts on the disk, and this is only marginally affected by the growth of the baryonic disk/bulge system. Only when the halos are artificially rounded and any figure rotation of the inner dark matter halo is stopped by construction, the disk orientation remains stationary. Except for halo F – its disk turns even in this case by a substantial angle, starting at $z \simeq 0.35$. The same characteristic turning motion of F-disks is also seen in the standard runs at this time, suggesting that this is caused by the fly-by of a massive substructure that interacts with the disk at this time and torques it substantially.

The results above suggest that disk tumbling of significant size is virtually inevitable in CDM halos. We expect typical tumbling rates of about 40 degrees from $z = 1$ to $z = 0$, or about $\sim 6 - 7$ degrees per Gigayear on average. Some systems may have up to 2-3 times that, while others stay below it by a similar factor. Occasionally, disks may also brought into a turning motion by a close encounters with a substructure. Importantly, our results show that disk can survive such reorientations largely unaffected,
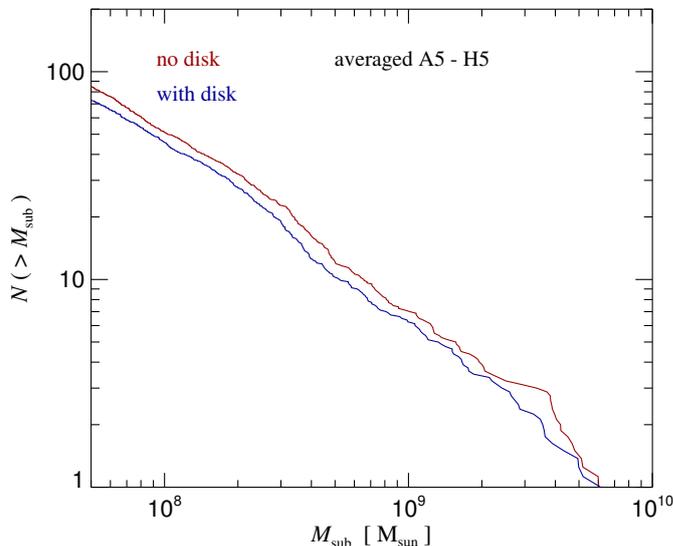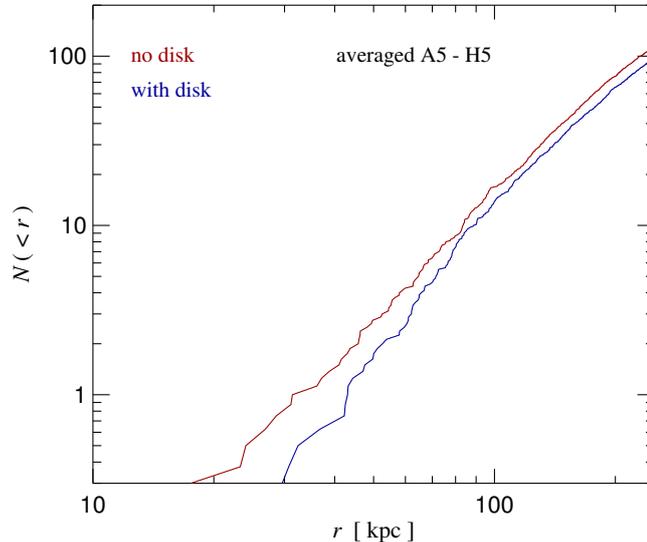
Figure 3.20: Change of the cumulative abundance of substructure as a function of mass due to presence of a stellar disk. We count all substructures identified by SUBFIND within 200 kpc of the halo centers, and averaged over our 8 Aquarius halos. The "no disk" runs refer to results at $z = 0$ of pure dark matter only runs of the halos where no stellar component is inserted, while the "with disk" results are for our default disk systems of simulation series #1.

i.e. they are not in apparent conflict with the observed abundance of thin stellar disks.

3.6   Impact of substructures

Cold dark matter subhalos contain a large amount of substructures, raising the question whether they may interfere with the stability of cold stellar disks and induce substantial heating. We note however that it is well established that substructures populate primarily the outer parts of dark matter halos (e.g. Ghigna et al. 1998; Diemand et al. 2004; Springel et al. 2008a), leaving the inner halo relatively smooth. Also, the subhalo mass function is skewed very slightly to being dominated by the most massive subhalo systems. Those are expected to dominate the heating (Springel et al. 2008a),

Figure 3.21: Change of the cumulative abundance of substructure as a function of radius due to presence of a stellar disk. We count all substructures identified by SUBFIND with a mass larger than $6 \times 10^7 \, \mathrm{M_\odot}$, and average over our 8 primary galaxy models. The comparison corresponds to that shown in Fig. 3.20, but considers the radial distribution instead of the mass distribution.

but their number is small.

We here use our models to check whether subhalos contribute significantly to the disk heating, and whether the disk in turn plays a significant role in reducing substructure abundance when they pass through pericenter and experience gravitational tidal shocks from the disk or the enhanced central cusp. There is a body of previous work on this subject, largely based on much simpler toy simulations than studied here. Our analysis is so far the most elaborate attempt to study this in the correct cosmological setting, and in particular takes the expected system-to-system variation into account.

We consider first the subhalo abundance in runs without any disk (i.e. these are *dark-matter only* runs of A to H at $z = 0$), and compare it to the one found at $z = 0$ in our default runs with disk. In Figure 3.20, we show the cumulative abundance of substructures as a function of mass in both

types of simulations. To emphasize the mean difference in a clear way we show the averaged abundance over all eight systems we simulated; we note that orienting the disk along minor or major axis makes no difference here. There is a $\sim 30\%$ reduction of substructure abundance across all mass scales in the runs with the disk. This can be understood as an effect of accelerated substructure depletion due to the gravitational shocks the substructures experience as they pass through the disk. This enhanced destruction rate shows up particularly strongly in the halo center, as evidenced by Figure 3.21, which gives the cumulative average abundance of substructures with mass larger than $6 \times 10^7 \, \mathrm{M_\odot}$ as a function of radius. In the inner parts of the halo, there is about a factor of two reduction of the subhalo abundance. These results are in good agreement with the analysis of D'Onghia et al. (2010), who carried out orbit integrations for the subhalos found in the dark matter only simulations of Aquarius and estimated their evaporation rate analytically by summing up the impact of gravitational shocks (Ostriker et al. 1972) experienced during disk passages.

In order to explicitly test for the influence of substructures on disk heating and disk stability, we have carried out a series of runs where we varied our "rounding halo" experiment above. Instead of sphericalizing all dark matter particles in the halo at $z = 1.3$, we have done so only for the dark matter particles bound in substructures. Combined, these subhalos amount to a few percent of the mass of the halo, so by redistributing the substructure particles in a spherical fashion, the halo is made smooth without affecting its dynamical equilibrium much. We note however that a large fraction of the substructures found in a halo at low redshift will be accreted at $z < 1.3$; they are unaffected, and so this cleaning of substructure will only temporarily make the halo smooth. This is seen explicitly in Figure 3.22, where we show the abundance of substructures as a function or radius at different times, comparing an ordinary disk run with a run where the substructure cleaning has been done at $z = 1.3$. While at $z = 1.25$ the inner halo is still largely devoid of subhalos in the cleaned run, in the redshift range
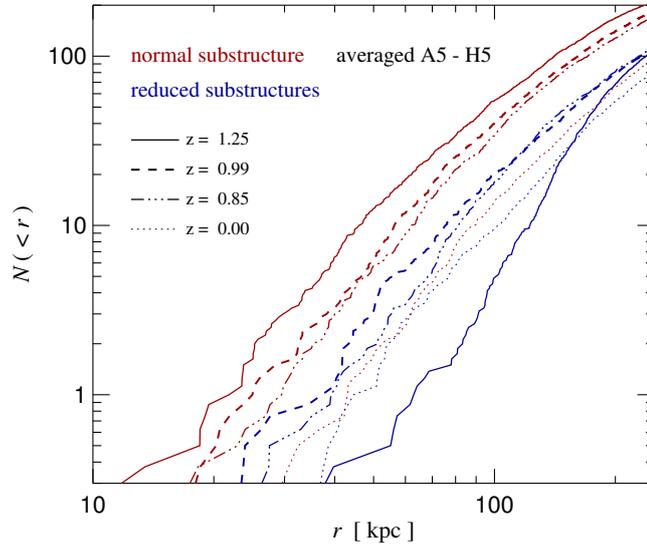
121

Figure 3.22: Illustration of the impact of our "substructure cleaning" procedure where at the disk insertion redshift of $z = 1.3$ all particles in bound substructures of the central FOF group are spherically redistributed. The plot shows the cumulative radial abundance of substructures with mass above $6 \times 10^7 \, \mathrm{M_\odot}$ at different times, comparing runs where the substructure reduction was carried out with normal simulations where this was not done. At $z = 1.25$, the reduction of subhalos in the inner parts of the target halo is still very strong, but soon the subhalo population is replenished by the accretion of additional mass and substructures. This reduces the substructure reduction effect at late times substantially, but in the range $1.0 < z < 0.85$ it is more or less constant as a function of radius and amounts to a suppression of around a factor of 2.

$1 < z < 0.85$, the abundance is suppressed on average by a factor of 2, whereas towards $z = 0$ it is down by only $\sim 20\%$.

We now use the subhalo cleaning run to look at differences of the disk heating rates between $z = 1.0$ and $z = 0.85$. At this time, the disks are still largely intact even in the runs that form strong bulges, and here the substructure suppression in the runs with the subhalo cleaning is still substantial and fairly uniform across radius and in time. In Figure 3.23 we show the difference in the vertical velocity dispersion in the standard runs and the runs with subhalo cleaning. We clearly see evidence for an enhanced
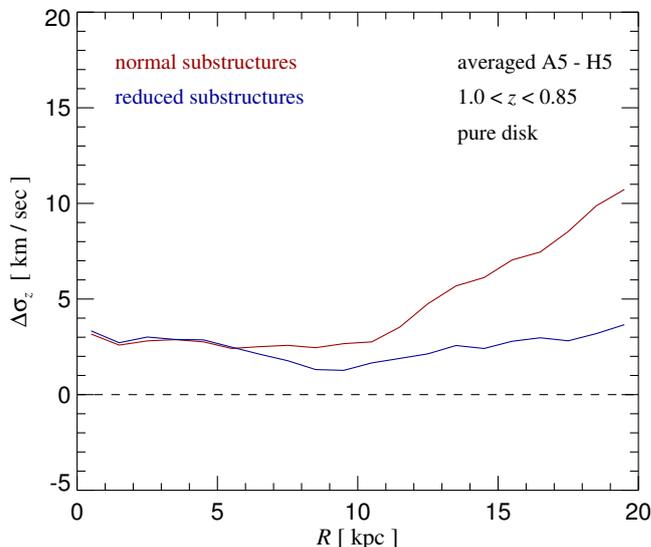
Figure 3.23: Differential heating of the stellar disk in runs with normal or reduced substructure abundance. We here show radial profiles of the increase of the vertical velocity dispersion of disk stars between times $z = 1$ and $z = 0.85$ in our runs with pure disks. The reference simulations correspond to series #1, while in the comparison runs of series #8 the substructure has been smoothed out at the disk insertion time ($z_{\text{insert}} = 1.3$). In the runs with reduced substructure abundance, we find a substantially reduced hearing rate in the outer parts of the disk whereas in the inner parts, within two disk scale length no difference is detected.

heating rate of the disks in their outer parts, amounting to several $\text{km}\,\text{s}^{-1}$ over the coarse of 1 Gyr. However, within two disk scale lengths, there is virtually no detectable difference in the disk heating rate, suggesting that substructure heating is negligible for the bulk of the disk's stellar mass. It may however play an important role in contributing to flaring of the stellar disk in the outer parts.

3.7 Resolution dependence

So far, all our results have been based on 'level-5' Aquarius simulations, with a resolution of 2×10^5 particles in the disk, a dark matter halo resolved by
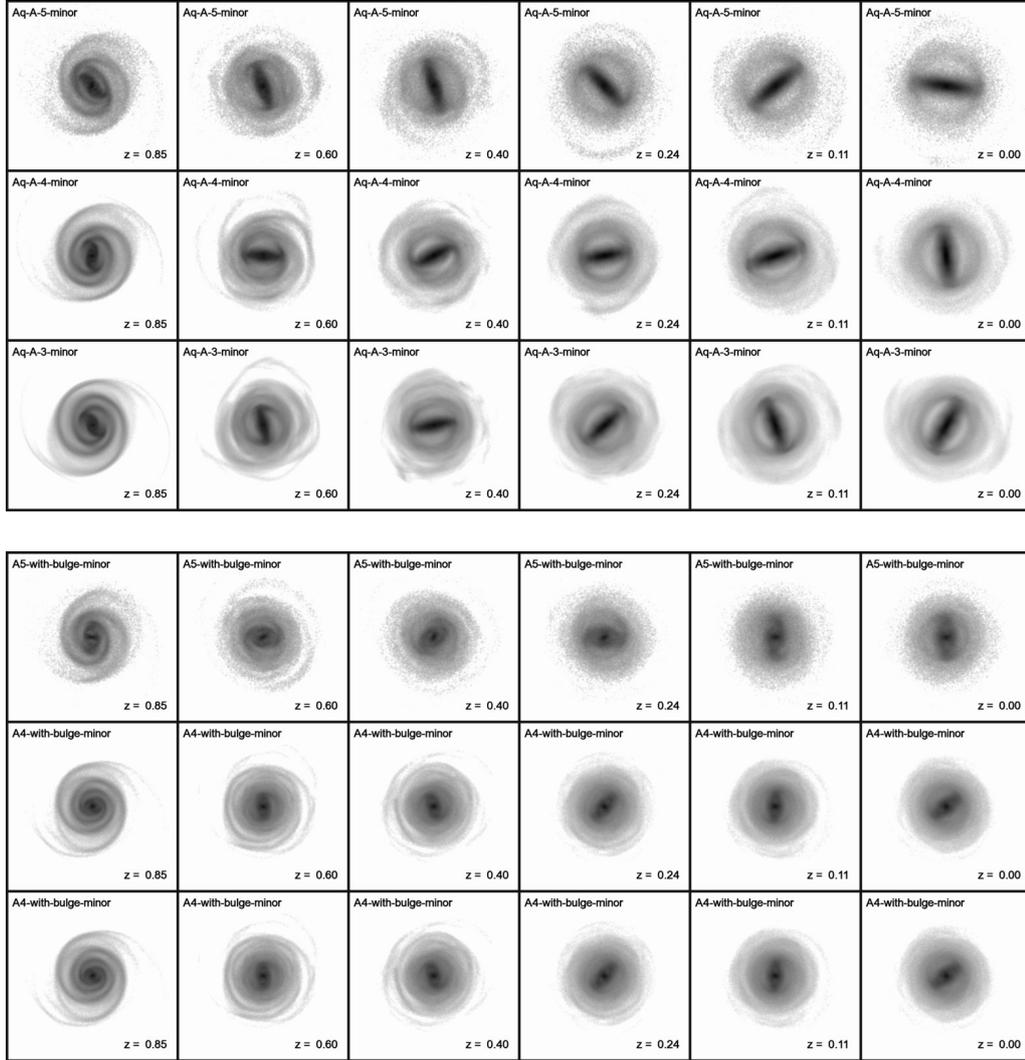
Figure 3.24: Resolution dependence of disk morphology in our default runs with pure disks (series #1) and disks+bulges (series #3). The top three rows of images compare the face-on projections of the stellar disks at different times of our simulations of $A5$, $A4$, and $A3$ in the runs with pure disks. The corresponding simulations when a third of the stellar mass is moved to a central bulge are shown in the bottom three rows. It is reassuring that the runs agree rather well, apart from the phase angle of the bar.

about 10^6 particles, and a gravitational softening length of 680 pc. Some effects of galactic dynamics can depend strongly on numerical resolution (e.g. Weinberg & Katz 2002), so it is advisable to check the robustness and numerical convergence of our primary results in the relevant regime.

To this end we have repeated our primary A-5 simulations at two higher resolution levels, where the particle number is increased by factors of 8 and 64, respectively, in both the dark matter and stellar components. The gravitational softening lengths are reduced by factors of 2 and 4, respectively. In our highest resolution simulation A-3, this means that the stellar disk is represented with 12.8 million particles, the stellar bulge (if present) with 6.4 million, and the dark matter halo with about 50 million particles.

In Figure 3.24 we show a visual comparison of the stellar disk evolution seen in the three resolution levels. The top three rows compare pure disk models at the resolution levels 5, 4, and 3, whereas the bottom three rows give the same comparison for our standard models with bulges. Overall, the disk morphologies are very similar in both resolution sequences. The position angles of the bar motion do not line up exactly, but such phase differences are to be expected. However, the overall morphological evolution is clearly very similar, which suggests good convergence of the structural evolution of the galaxies.

More quantitatively, we show in Figure 3.25 a convergence study of the tilt angle evolution in the two resolution sequences. Especially the two high resolution runs line up remarkably well. The lower resolution run shows a small offset in its evolution in comparison, but since there is no systematic trend with resolution we interpret this small difference in angle as a chance effect. Other quantities we examined are similarly stable with respect to resolution. We hence believe that already the level-5 resolution provides robust results for the quantities studied in this work.
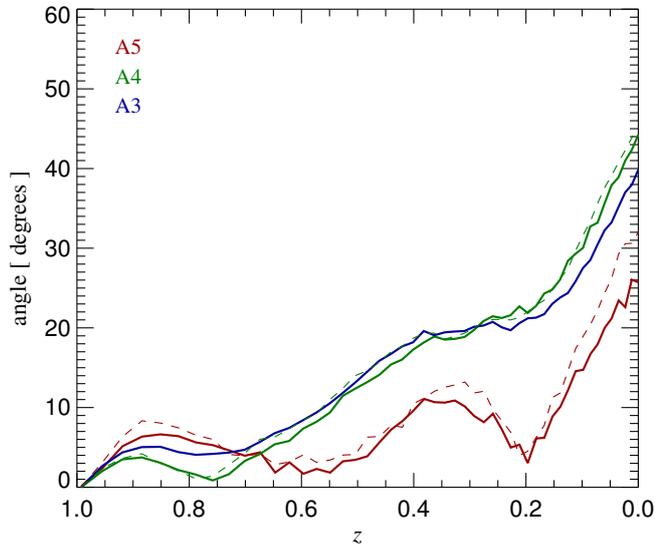
125

Figure 3.25: Disk tilt angle measured for simulations at different numerical resolutions, ranging from Aquarius level-5 to level-3. Two types of runs are shown, the A-halo with a pure disk (from series #1), and the A-halo with a disk+bulge (from series #3). The tilt angles of the disks do not line up quantitatively in detail for the different resolutions, but the evolution qualitatively still agrees reasonably well, without any indication of systematic trends with resolution. The tilt angle evolution for a given resolution is almost identical for runs with or without a bulge, showing that it is governed almost exclusively by the dark matter halo, so that the difference we find between A3, A4, and A5 likely originate in small differences of the dynamical states of their dark matter halos when the disk is inserted.

3.8   Discussion and conclusions

In this study, we have analyzed the stability of disk galaxies inserted into high-resolution zoom-simulations of the formation of Milky Way-sized dark matter halos, with initial conditions taken from the Aquarius project. We have refined a methodology previously used by DeBuhr et al. (2012), most notably by using a more sophisticated approach to determine the initial velocity distribution of the star particles. For the latter we employed the iterative method realized in the GALIC code (Yurin & Springel 2014), which is capable to compute high-quality stationary solutions in general dark mat-

ter halos. For the dark matter potential, we directly used the distribution of dark matter particles found in the Aquarius halos, without using any approximation besides imposing axisymmetry on the force field.

Using this improved methodology, we have extended the analysis of the Aquarius halos to a larger halo sample (using eight systems, A to H), and to structural variants that also include systems with central bulges. We have also investigated a number of toy simulations where the dark matter halos were artificially rounded or dark matter substructures were erased, in order to highlight the impact of residual triaxiality or of substructures on the dynamical evolution of the disks.

Our main findings can be summarized as follows:

1. The presence/absence of a stellar bulge, as well as the presence/absence of a stellar bar, do not affect the tumbling of disk galaxies. The turning motion of disks appears to be primarily driven by the triaxiality of the halo and its figure rotation; in rare cases encounters with massive substructures can also initiate substantial disk tilt.

2. Disks initially oriented along the dark matter halo's minor axis show marginally better directional stability than disks oriented along the major axis.

3. We predict that an average tumbling angle of about 40 degrees over 6 Gyr between $z = 1$ and $z = 0$ should be quite typical for disk galaxies, corresponding to $6 - 7$ degrees per Gyr. Importantly, thin disks can survive such tumbling rates in a largely unaffected way.

4. We confirm that the simple criterion of Efstathiou et al. (1982) is a good predictor for bar instability also in the full cosmological context. Other parameters appear to be at most of secondary importance, in particular, we do not find strong evidence that disk systems in round halos are less prone to bar formation than the disks in the mildly triaxial CDM halos studies here.

5. There is a significant depletion of dark matter substructures due to the presence of a massive disk. If part of the disk mass is put into a bulge instead, the effect is slightly reduced (**t**o be confirmed), suggesting that gravitational shocking at the disk is indeed more important than the enhancement of the central core density and the associated increase of pericenter at the halo cusp.

6. Dark matter substructure appears to significantly contribute to disk heating in the outer parts of disks, while this appears negligible in the inner regions within $\simeq 2$ disk scale lengths.

7. Our quantitative numerical results appear unaffected by numerical resolution, as evidence by our resolution tests that cover a factor of 64 in mass resolution, and a factor of 4 in gravitational softening length.

Overall, our results suggests that the survival of thin stellar disks is in principle not a problem in CDM halos. In particular, the triaxiality and high substructure abundance in CDM halos do not preclude the survival of thin stellar disks, even though we would expect them to tumble slowly with time. What is arguably more difficult to understand is how massive, cold stellar disks can survive strong bar formation. Cuspy cold dark matter actually help here, but the low central dark matter densities often inferred observationally for the inner parts of galaxies (even for the Milky Way Bovy & Rix 2013) make it unclear how such observed systems manage to sport only a small or no bar. This remains an interesting topic for further study.

The methods used here demonstrate the power of our technique to introduce stellar systems into growing dark matter halos. This allows one to carry out traditional galactic dynamics work at very high resolution within a realistic and complex cosmological environment, something that promises to be a worthwhile avenue for future work.

4

The radial migration of stars in disc galaxies formed in moving-mesh cosmological simulations

Abstract

Radial migration of stars has been suggested as an important mechanism for shaping the metallicity structure in disk galaxies, and for possibly contributing significantly to the formation of a thick disk component. Here we use for the first time fully cosmological hydrodynamical simulations of disk galaxy formation to quantitatively assess the level of radial redistribution of stars. We consider simulation models of six different halos carried out with a moving-mesh hydrodynamical code which forms galaxies that are broadly consistent with the expected structure of Milky Way sized galaxies. We find that the radial redistribution can be described to good accuracy as a random walk process with a typical diffusion constant of $0.8\,\mathrm{kpc}^2\,\mathrm{Gyr}^{-1}$. At this level, the strength of radial migration present in the simulations affects the metallicity gradients only by an insignificant amount. We confirm previous findings that disk thickness and vertical velocity dispersion are both a function of stellar age and of the amount of radial migration. Outwardly migrating stellar populations reduce their velocity dispersion but keep their thickness approximately constant, suggesting that the impact of radial migration on disk thickening is at most of limited importance. Our resolution tests indicate that our results for radial migration are numerically well converged, but a possible dependence of the stellar disk structure on feedback physics constitutes an important systematic uncertainty.

4.1 Introduction

Understanding the build-up of the stellar disk of our own Milky Way and of galaxies of similar size remains one of the most interesting problems in galaxy formation. In hierarchical models based on the prevailing ΛCDM cosmological model, we expect a 'bottom-up' formation of galaxies (White & Rees 1978) and a substantial content of dark matter even in the innermost regions of late-time galaxies (Navarro et al. 1996). It is still unclear whether this can be consistent in detail with the observed structural properties of galaxies, such as their rotations curves or the claimed think/thick structure of the stellar disk. In addition, their hosting dark matter halos should be full of smaller subhalos (e.g. Moore et al. 1999), likely providing a strong source of perturbations in the gravitational potential, making it difficult to understand how thin stellar disks and stellar streams can survive. It is therefore important to confront detailed numerical simulations of galaxy formation and evolution with the increasingly reach data on galactic structure.

For example, recent observations such as the Gaia-ESO survey (Bergemann et al. 2014) show that there is a large scatter in metallicity at given stellar age, significantly degrading the utility of the age-metallicity relation for constructing reliable models of the build up of the disk. This also challenges the idea that all the stars present at a local region of the disk formed more or less in close proximity to each other, i.e. from the same patch of gas. More sophisticated chemo-dynamical models that either model the complex gas flows in galaxies faithfully or take into account dynamical redistribution of stars in galaxies appear therefore needed.

A large spread in metallicity at a given stellar age, even in an inside-out disk formation scenario, could for example arise if stars are not restricted to their birth radius, but rather that any local sample of stars really comprises birthplaces at very different radii. This idea of significant radial migration has a long history (Wielen 1977; Wielen et al. 1996; Sellwood & Binney

2002), and was recently revived with a number of galaxy simulations in isolation (Roškar et al. 2012; Solway et al. 2012; Roškar & Debattista 2014). The simulations largely confirmed an important role of the corotation resonance mechanism proposed by Sellwood & Binney (2002) for facilitating radial migration. Scattering at this resonance is able to exchange angular momentum and energy without causing significant heating of the orbits. On average, the vertical action is expected to be conserved during radial migration (Solway et al. 2012), but alternatively, some theoretical works have also conjectured that the vertical energy is conserved instead (Schönrich & Binney 2009).

In full cosmological simulations, additional mechanisms can potentially redistribute material in the disks, such as satellite mergers (Quillen et al. 2009; Kazantzidis et al. 2009) or growing bars (Debattista et al. 2006; Minchev & Famaey 2010; Di Matteo et al. 2013). In general, it may be difficult to cleanly separate these processes in cosmological simulations, even though they are clearly physically distinct. But the advantage of such self-consistent models is that they do not rely on ad-hoc choices of initial conditions which may potentially distort the results. Despite this advantage, so far no analysis of the radial migration rates in full cosmological hydrodynamical simulations of disk galaxy formation has been attempted.

A primary reason of this is that such calculations have simply not been available until very recently. Only the newest generation of such hydrodynamical cosmological simulations finally succeeded in making reasonable disk galaxies, with roughly the right stellar mass and right stellar size (e.g. Governato et al. 2010; Agertz et al. 2011; Guedes et al. 2011; Aumer et al. 2013; Stinson et al. 2013; Marinacci et al. 2014a). In the present study we want to take advantage of this important advance and carry out a first analysis of radial migration in such fully cosmological disk formation models, using the set of simulations first presented in Marinacci et al. (2014a).

This chapter is structured as follows. In Section 4.2 we briefly review the methodology of our hydrodynamic simulations. We then present mea-

surements of the extent of radial migration in Section 4.3, implications for the metallicity structure in Section 4.4, and results for changes in the vertical structure of migrating stellar population in Section 4.5. Finally, we briefly examine the convergence and numerical robustness of our results in Section 4.6. Conclusions and a summary are given in Section 4.7.

4.2 Methodology

In this study, we analyse the hydrodynamical simulations of disk galaxy formation previously analysed in Marinacci et al. (2014a,b). In brief, they are hydrodynamical versions of the 'Aquarius' initial conditions of Milky Way-sized galaxies (Springel et al. 2008a), carried out with the moving-mesh code AREPO (Springel 2010) and a sophisticated model to treat star formation and associated feedback processes. The Aquarius halos were chosen to be mildly isolated and to have a quiet merger history, which favours the formation of disk galaxies, but otherwise they were purely mass-selected to have a mass around $10^{12}\,\mathrm{M_\odot}$ at $z = 0$ (see Springel et al. 2008a, for details on the selection procedure and the technical set-up of the Aquarius "zoom" simulations). We here focus on a subset of six of the eight Aquarius systems simulated by Marinacci et al. (2014a), disregarding halos D and E as they produced only poorly defined disks at low redshift.

The hydrodynamical simulations examined here used a comprehensive model for galaxy formation physics described in full in Vogelsberger et al. (2013). It implements the baryonic processes considered to be most important for regulating star formation and galaxy growth. Specifically, the hydrodynamic simulations include metal-dependent radiative cooling, a sub-resolution model for regulating star formation in the interstellar medium, a self-consistent treatment of stellar evolution and metal enrichment (following 9 elements, H, He, C, N, O, Ne, Mg, Si and Fe), a galactic wind model, growth of supermassive black holes and associated energy feedback, and a UV photoionizing background. The model has been specifically developed

for the AREPO moving-mesh code and has been shown to successfully re-produce a number of key properties of the observed galaxy distribution. In fact, the simulations analysed here are arguably among the most successful simulation models of Milky Way sized galaxies available thus far (Marinacci et al. 2014a; Pakmor et al. 2014). The disk galaxies formed in the simulations have the right stellar mass expected for halos of this mass based on abundance matching arguments, and they also have approximately the size expected observationally for this stellar mass. Furthermore, in applications of the same galaxy formation model to large cosmological volumes in the 'Illustris' simulation project (Vogelsberger et al. 2014a,c) it has been shown that the model is quite successful in reproducing a broad range of observations both at high redshift and at the present epoch (see also Genel et al. 2014; Bird et al. 2014).

In this work, we mainly focus on the 'level-5' resolution in the notation of Springel et al. (2008a) and Marinacci et al. (2014a), which corresponds to a baryonic mass resolution of $\simeq 4.1 \times 10^5 \, \mathrm{M}_\odot$, a dark matter mass resolution of $\simeq 2.2 \times 10^6 \, \mathrm{M}_\odot$, and a gravitational softening length of $680 \, \mathrm{pc}$ for dark matter and baryons in the high-resolution region. However, we will also carry out a resolution study for halo C, for which an 8 times better ('level-4') and an 8 times worse resolution ('level-6') is available.

4.3 Extent of radial migration

In Figure 4.1, we give a visual impression of the degree of radial mixing in our different disk systems. In the left-most column, the identified "clean" disk stars are shown in different radial annuli of width $3 \, \mathrm{kpc}$, using different colours. The same stars are then shown again for subsequent simulation output times, over a timespan of $2.5 \, \mathrm{Gyrs}$.

It is evident that the models show a range of different behaviour. There is particularly violent mixing in the halo G5, where a late time passage of a massive satellite galaxy takes place (see Fig. 2 of Marinacci et al. 2014a, for
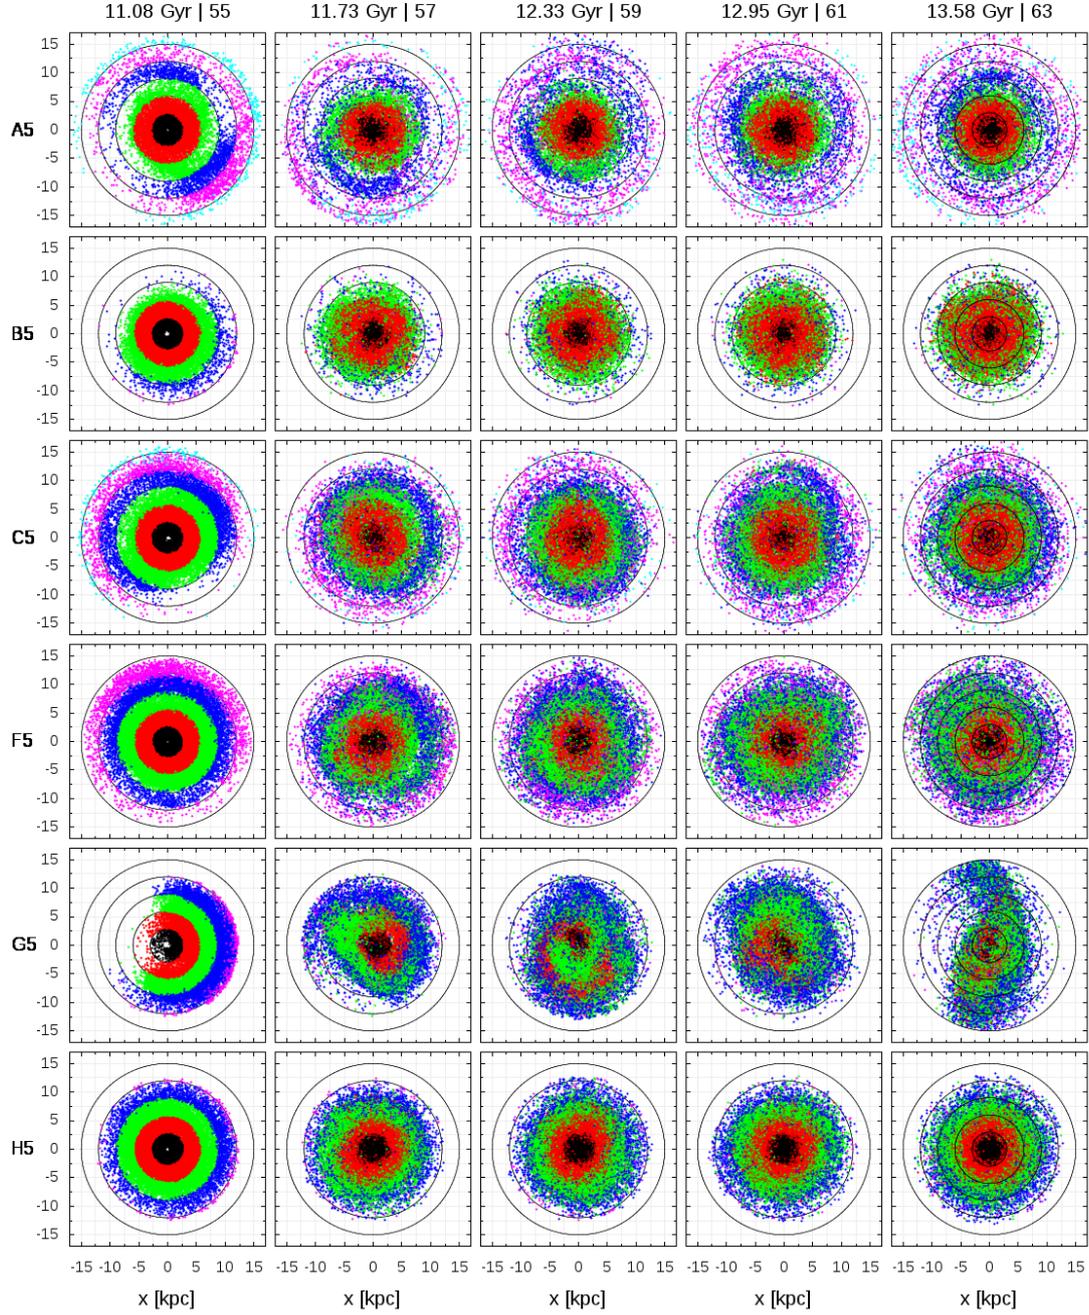
Figure 4.1: This plot gives a visual impression of the evolution of the radial distribution of disk stars in a face-on projection. Each row shows a different disk galaxy. The stars are drawn with different colour according to 7 radial bins defined at $t = 11.08\,\mathrm{Gyr}$ (left column), and only stars with circularity parameter larger than 0.8 are included. Subsequent columns then show how the stars in different initial radial bins become mixed with time.
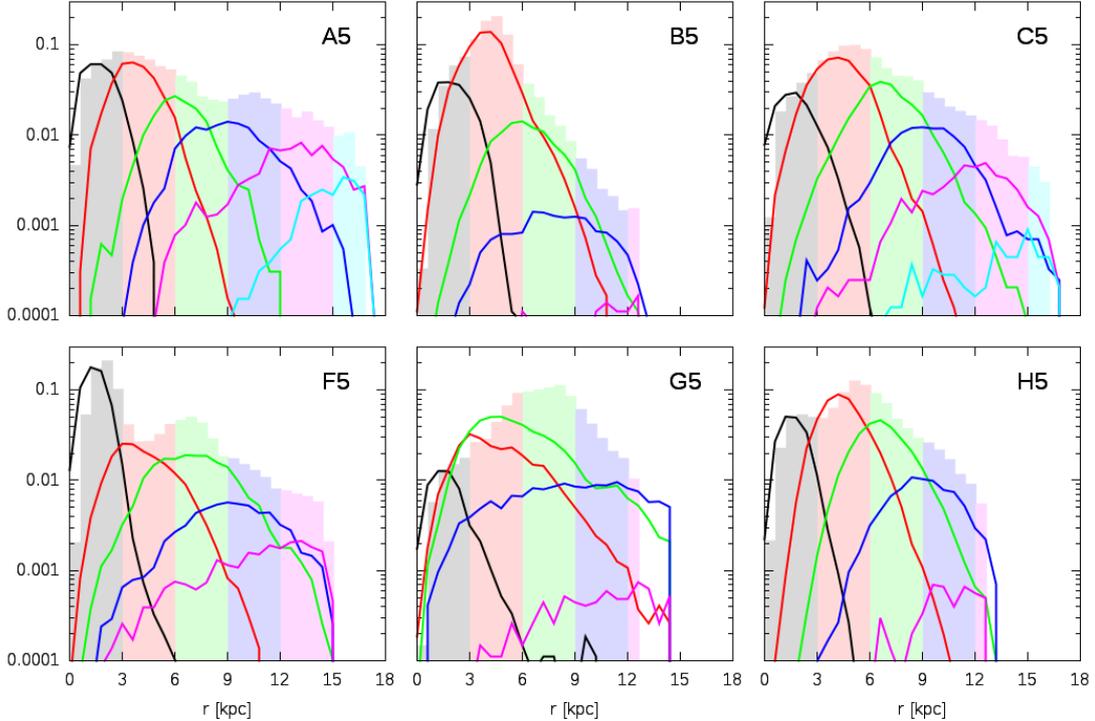
135

Figure 4.2: Change of the radial redistribution of disk stars between times $t = 11.08\,\mathrm{Gyr}$ and $t = 13.58\,\mathrm{Gyr}$. The shaded histograms indicate the initial radial ranges that are examined, and have an area proportional to the amount of stellar mass in the corresponding bin. The solid lines give the radial distributions of the same stars at the present epoch, normalized again to the corresponding stellar mass.

a time sequence of this event). In the other cases, the stars progressively diffuse to the neighbouring radial bins in a more or less gradual fashion. Notice also that in some cases the strength of radial mixing varies with azimuthal angle, implying that an azimuthally averaged analysis of radial migration glosses over some of the finer details of this process.

A more quantitative view of the radial redistribution process of stars is provided by Figure 4.2, where we show the distribution of the final radii of stars for different selections of initial radii. In this figure, we subdivide the radial range $0 \leq R < 18\,\mathrm{kpc}$ into 6 different regions of width $3\,\mathrm{kpc}$. Disk
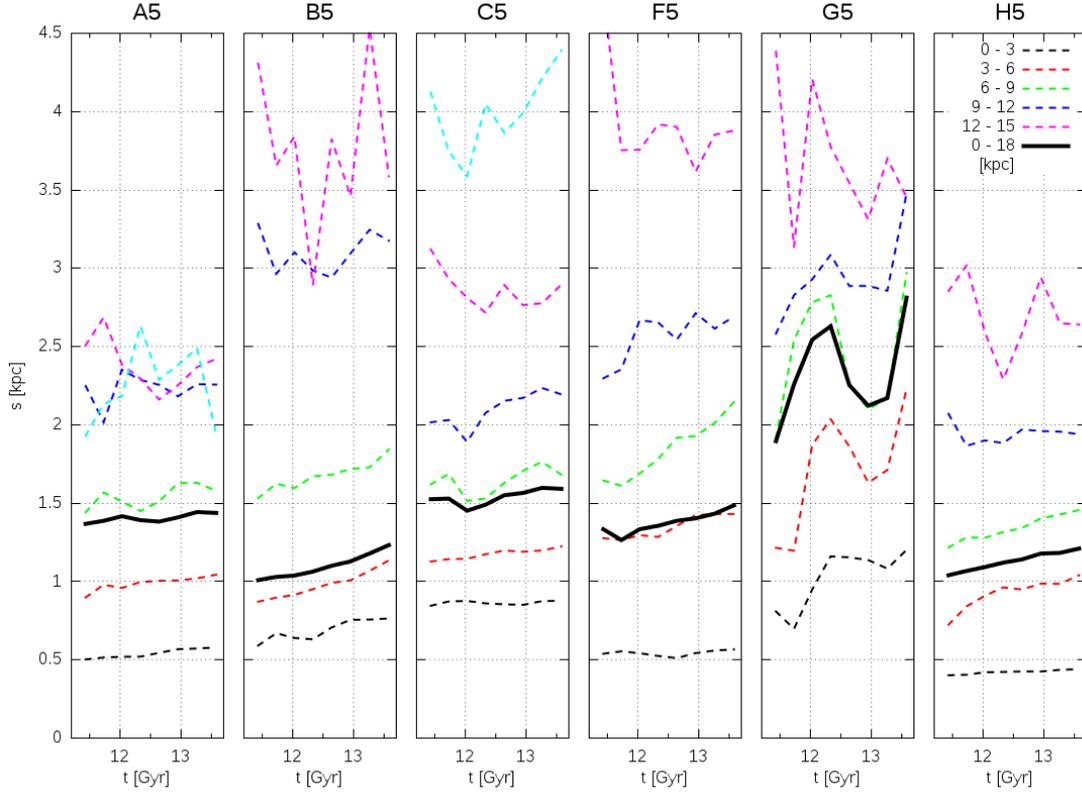
Figure 4.3: Effective diffusion rate due to radial migration. The individual plots show the root-mean-square difference between current and initial radial positions of selected stars whose mass-weighted circularities (defined as in Marinacci et al. 2014a, Eq. 1) are greater than 0.7. The dashed lines refer to different radial ranges, with a colour key given in the top right. The solid thick lines show the mean result for all stars within 18 kpc of the center. The different panels are for the 6 examined galaxies, as labelled.

stars that are on highly circular orbits (with circularity parameter larger than 0.8) in each of the bins are identified at $z = 0.22$ (i.e. at a lookback time of 2.5 Gyr ago). Their relative number as a function of radius is indicated through the shaded histograms. Finally, we show the distribution of the same stars at the present epoch with solid lines. Note that the vertical scale is logarithmic. Whereas clearly tails of these distributions develop into neighbouring radial bins at late time, the distributions remain in the majority of cases strongly peaked at their original annuli, suggesting that the degree of radial migration is quite limited globally. An exception is halo G5, which shows a relatively broad redistribution of stars, consistent with the results of the visual inspection of the stars in Fig. 4.1.

If we assume that stars migrate radially through random resonant scattering effects that are uncorrelated, we expect a random walk process in which the mean squared radial displacement growth proportional to time. This suggest that we can define an effective radial diffusion constant

$$D \equiv \frac{\langle (R_{\mathrm{final}} - R_{\mathrm{initial}})^2 \rangle}{t_{\mathrm{final}} - t_{\mathrm{initial}}} \tag{4.1}$$

and use it to measure a net radial diffusion speed, as one quantitative measure of the strength of radial migration.

In order to test whether this definition can provide a reasonable description of the radial migration, we show in Figure 4.3 measurements of $s \equiv \langle (R_{\mathrm{final}} - R_{\mathrm{initial}})^2 \rangle^{1/2}$ as a function of the final time t_{final}. We see that there is indeed an approximately linear relation between s^2 and the elapsed time, at least at late times. This is especially true for the inner regions of the galaxies, whereas the measurements for the outer regions are affected by substantial counting statistics noise. When the mean for all disk stars is considered, it is evident again that G5 is an outlier, showing a much larger radial redistribution rate than the other galaxies. The averaged values we obtain are quite similar for all systems, except for G5. If the latter is excluded the average value we measure for our models is $D \simeq 0.8 \, \mathrm{kpc}^2 \, \mathrm{Gyr}^{-1}$.

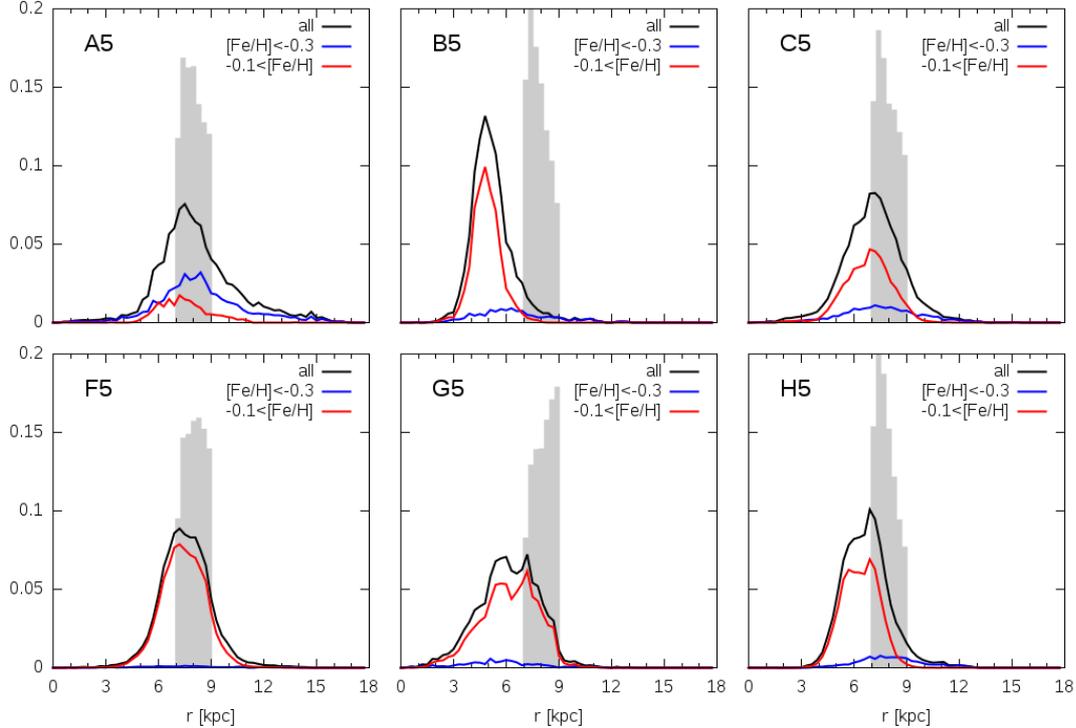In Figure 4.4, we focus instead on a radial region corresponding to the

Figure 4.4: Birth radii distributions of the disk stars in the "Solar neighbourhood" of our disk galaxies. We here select all disk stars in the radial range $7 \, \mathrm{kpc} < R < 9 \, \mathrm{kpc}$ and show the distribution of their birth radii. The latter is measured for the first output time where these stars are found. The solid line is for all stars, the dashed and dotted lines are for low and high metallicity stars, respectively.

"Solar circle" today, i.e. at the stars on nearly circular orbits in the radial region $7 \, \mathrm{kpc} < R < 9 \, \mathrm{kpc}$. Here we ask the question at which radii the corresponding stars were born, and show the corresponding radial distributions. We also show the distributions for two subsamples of different metallicity in order to examine whether the trends depend on metallicity. Interestingly, whereas A5, C5, and F5 draw the stars in this "Solar neighbourhood" mostly from a narrow region more or less centred on the fiducial solar circle, system B5 has apparently seen a significant expansion of its disk, such that most stars now in the range $7 \, \mathrm{kpc} < R < 9 \, \mathrm{kpc}$ used to be

located closer to the center by $2 - 3\,\text{kpc}$. A milder effect of this kind is also seen in H5 and in G5. In the latter case the distributions are however also broader, consistent with our earlier findings that G5 is a special, highly perturbed case.

Comparing our measurements in Figure 4.4 with results reported by Roškar et al. (2008) for their isolated disk simulation, we find a qualitatively similar behaviour, with typically more of the stars ending up in $7\,\text{kpc} < R < 9\,\text{kpc}$ originating at smaller radii than at larger radii. However, our distributions are clearly narrower than those of Roškar et al. (2008), with considerably fewer stars originating at very small radii. In particular, in our simulations there are virtually no stars that are born at $R < 3\,\text{kpc}$ and make it into the fiducial solar neighbourhood, whereas this a sizeable number of such stars in the simulation of Roškar et al. (2008).

4.4 Metallicity structure

4.4.1 Age-metallicity relation

It is often argued that a high diversity in the stellar populations at different locations in disk galaxies can be viewed as an important signature of stellar migration (Roškar et al. 2008; Schönrich & Binney 2009; Roškar & Debattista 2014). In Figure 4.5, we show the age-metallicity relation for all disk stars in our simulations (black dots), together with the mean (thick lines) and 1σ dispersion around the mean (dashed lines).

There is clearly a strong correlation of the mean metallicity with age at early times, where the galaxies quickly self-enrich and climb up to average metallicities around solar. This average metallicity stays then nearly flat, for periods of around $8\,\text{Gyr}$, with the exception of the G5 system again, where the rise to high metallicity is more gradual. In the latter case, the distorted structure of the galaxy may also mean that our selection of disk stars ends up being not very representative for all the stars of G5, which may be the reason of the different shape of the age metallicity relation in
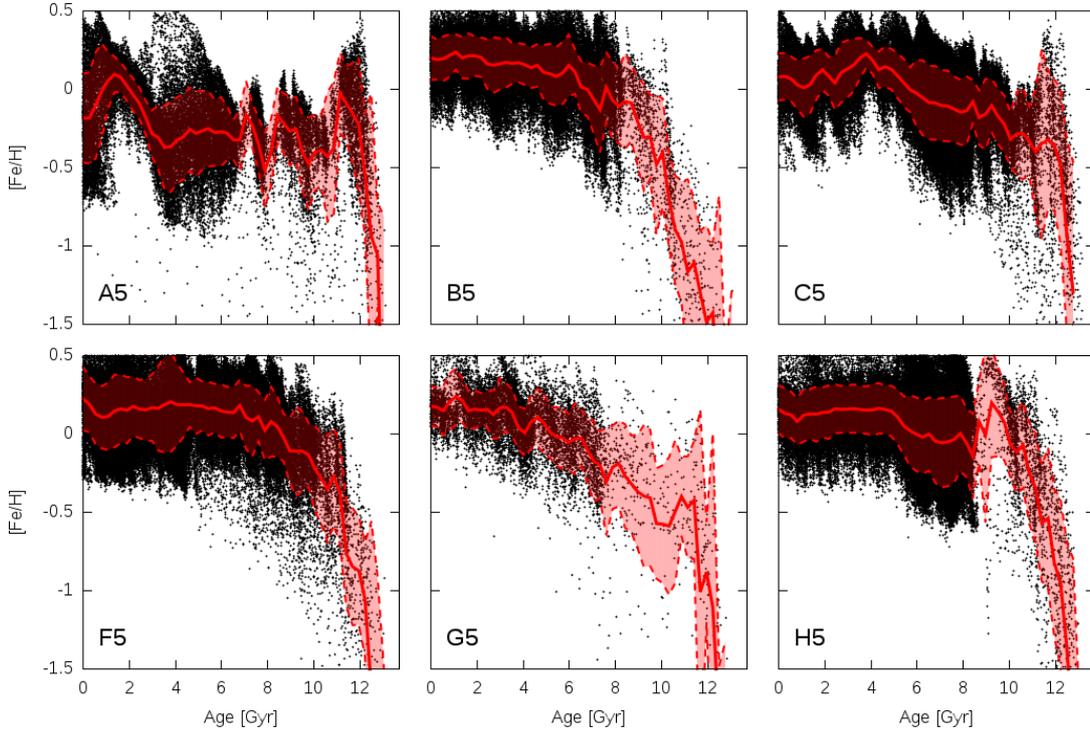
Figure 4.5: Age-metallicity relation for disk stars in our simulation models. The dots show the individual disk stars in the metallicity-age plane, while the thick lines show the mean in the different bins. The dotted lines indicate the 1σ dispersion around the mean.

this system.

Arguably the most striking result of Figure 4.5 is however that there is a large metallicity spread at any given age, of about $\sim 0.3 - 0.4\,\mathrm{dex}$. Given that radial migration is not overly strong in our models, this casts significant doubts about the notion that such a dispersion indicates the presence of substantial radial migration. In fact, our results rather suggest that radial migration is not a significant driver in creating this scatter at a given radius, unlike proposed by Roškar et al. (2008). Instead, it seems much more likely that the complicated fountain-like gas flows and galactic winds that permanently exchange material between the interstellar and circum-galactic media (e.g. Marinacci et al. 2014b) play a major role in establishing the large metallicity spread. Note also that the vast majority of disk stars form later
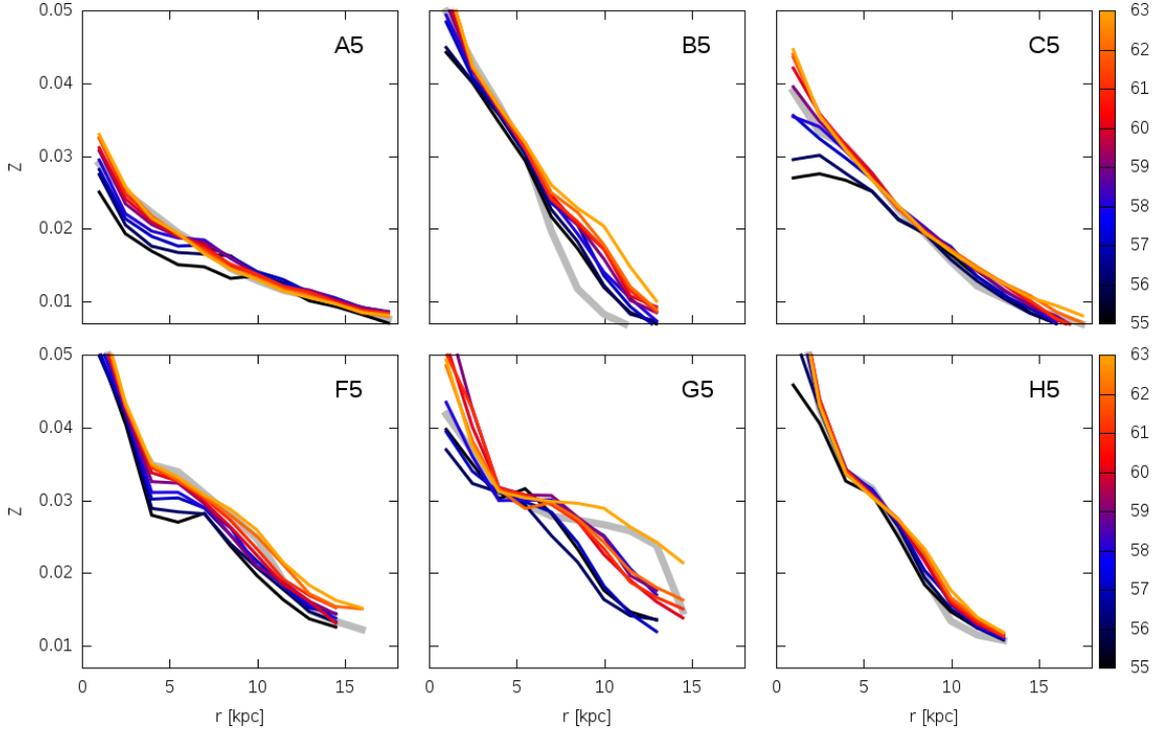
Figure 4.6: Evolution of metallicity profiles in the six examined disk galaxies. The different coloured lines show the metallicity profiles of the set of "clean" disk stars (i.e. excluding the bulge stars) in the galaxies at different output times over the recent 2.5 Gyrs. The thick grey line shows the profile if the stars are instead placed at their birth radii, i.e. at the radius where they first can be identified. This in some sense should 'undo' the effects of radial migration.

than 8 Gyr ago, at which point the age-metallicity relation is essentially flat, demonstrating that metallicity is not a particularly useful age indicator for these stars at all.

4.4.2 Metallicity profiles

Further evidence for a subdominant influence if radial migration is obtained by an examination of the metallicity profiles in the disk galaxies. In Figure 4.6, we show radial metallicity profiles of all of our simulations at a set of different times. Again, we restrict the measurements to a "clean" sample of disk stars on highly circular orbits in order to highlight effects of radial migration. The metallicity gradient at $z = 0$ (snapshot dump 63) is shown

with a thin yellow line, the profiles at earlier times with different colours, as labeled. The black line gives the profile at $z = 0.22$ (snapshot dump 55). The thick grey line shows the profile if the $z = 0$ stars are instead placed at their birth radii.

Overall, there appears to be little systematic change in the steepness of the metallicity profiles over this timescale. In particular, a significant trend towards a flattening of the profiles with time is not evident, although this may still operate at a weak level in a subset of the systems. We also see that the grey thick profile is generally very close to the yellow profile, apart from some larger distortions in the outer parts in B5 and G5. Hence, radial migration has apparently not affected the shape of the radial metallicity profile of the $z = 0$ stars in a significant way, demonstrating also that it must have been a weak effect overall for the structural evolution of the galaxies.

4.5 Vertical structure and radial migration

4.5.1 Different causes for population thickening

If stars migrate radially, their mean vertical displacement above the disk plane is expected to be changed, an effect presumably governed by action conservation (e.g. Loebman et al. 2011). Indeed, Roškar et al. (2013) showed for simulations of isolated disk formation that the vertical thickening is both a function of stellar age and the amount of radial displacement. The vertical thickening is expected to increase with the age of stars, because stellar populations born close to the mid-plane of disks are expected to be heated through various channels, for example by the bombardment with dark matter substructures or by minor mergers. However, radial migration can also affect the thickness of a stellar population as it changes its radius.

Interestingly, different claims have been made as to whether a stellar population will thicken or become narrower when it migrates outwards. Minchev et al. (2012) has argued, based on vertical action conservation, that
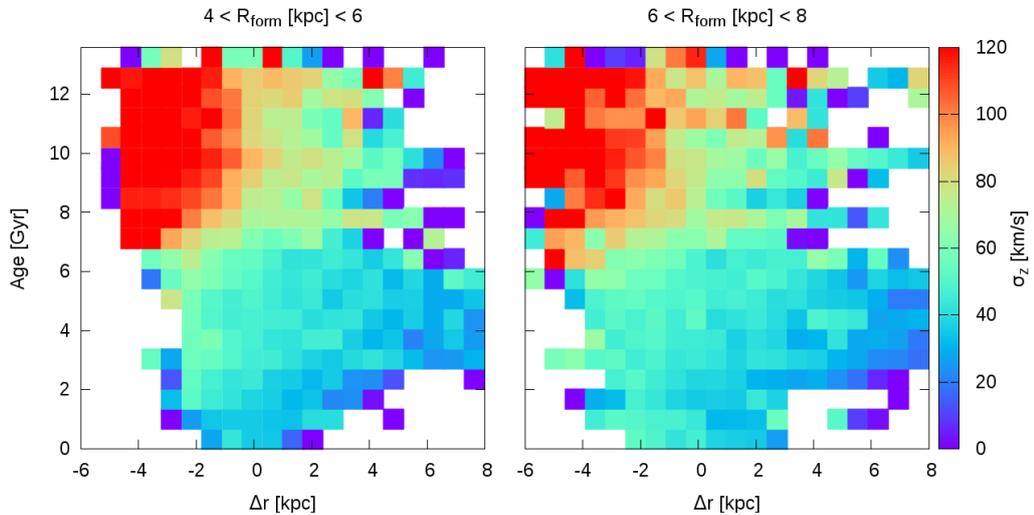
Figure 4.7: Mean velocity dispersion σ_z (colour-coded) of disk stars as a function of their age and the amount of radial migration they have experienced since birth. Only stars forming in a narrow radial are selected, as labelled (i.e. $4\,\mathrm{kpc} < r < 6\,\mathrm{kpc}$ in the left panel, and $6\,\mathrm{kpc} < r < 8\,\mathrm{kpc}$ in the right panel).

stars would not thicken as the migrate outwards, whereas just the opposite has been claimed by (Schönrich & Binney 2009), who instead assumed energy conservation in the vertical distribution of stars in their analytic model. Loebman et al. (2011) present N-body simulations that seem to support the analysis of Schönrich & Binney (2009), but Minchev et al. (2012) have supported their dissenting conclusions with a set of different N-body simulations (Chilingarian et al. 2010; Martig et al. 2012). Also, Solway et al. (2012) found that the vertical action rather than the vertical energy is conserved on average during radial migration.

In light of these disagreeing results it is particularly interesting to analyse our simulations with respect to this question, taking advantage of the fact that they are the first fully cosmological simulations of disk formation available for studying this topic. In Figure 4.7, we show for stars born in a small radial interval their mean vertical velocity dispersion, both as function of age and as a function of the amount of radial displacement they have experienced. In other words, we examine for stars originally born at

a fixed radius how their velocity dispersion depends on their age and the shift away from their birth location. The left and right panels of the figure refer to different radial ranges for the birthplace, $4\,\mathrm{kpc} < R_{\mathrm{birth}} < 6\,\mathrm{kpc}$ and $6\,\mathrm{kpc} < R_{\mathrm{birth}} < 8\,\mathrm{kpc}$, respectively. The two-dimensional distributions of the panels show the mean vertical velocity dispersion of disk stars on nearly circular orbits, with a colour-scale as given in the legend, both as a function of the radial migration shift ΔR in the horizontal direction and the age in the vertical direction. We have here stacked our different simulation models to reduce statistical noise.

At a given ΔR, there is a clear trend for stars that are older to show a higher vertical velocity dispersion. This is expected and consistent with a gradual vertical heating of ageing stellar populations. More interesting is the trend seen for fixed age. Here we see that stars moving to smaller radii actually increase their velocity dispersion, while stars moving outward become colder. This is qualitatively very similar to recent results by Roškar et al. (2013). It agrees with the analytic analysis of Minchev et al. (2012), but not with that of Schönrich & Binney (2009).

4.5.2 Correlation between birth place and vertical scale height

Because the vertical velocity dispersion is directly related to the vertical thickness of a stellar population, we in principal expect that corresponding trends about the impact of radial migration can also be seen directly in the thickness of stellar populations.

In Figure 4.8, we test this idea by showing for a fiducial population of 'Solar neighbourhood' stars (those today in the radial range $7\,\mathrm{kpc} < R < 9\,\mathrm{kpc}$) their root mean square vertical height above the disk plane, $z_h = \sqrt{\langle z^2 \rangle}$, as a function of the radial migration ΔR these stars have experienced. To improve the statistics, all 6 galaxy models have been stacked. We show two different measurements; the blue line uses simply the current z-coordinate of the stars above the disk plane, whereas the red line uses
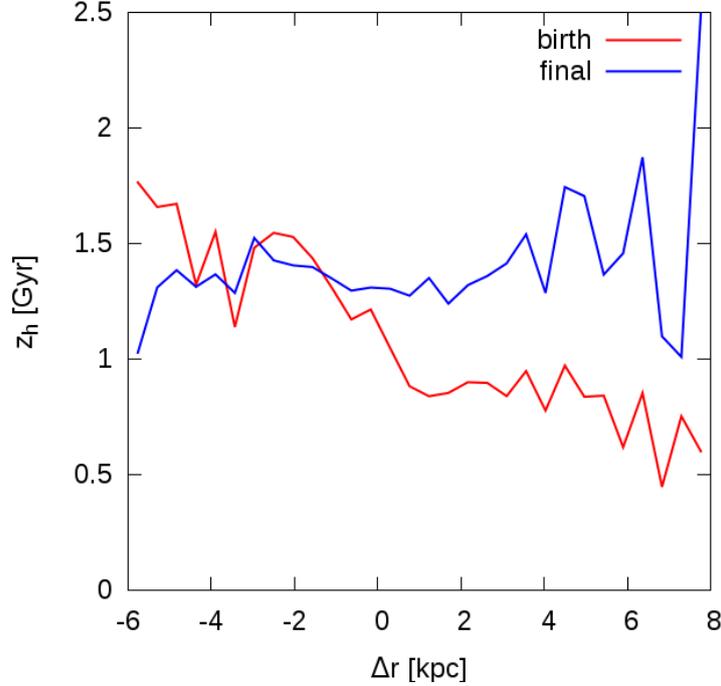
Figure 4.8: Correlation between mean height $z_h = \sqrt{\langle z^2 \rangle}$ of stars found in the solar neighbourhood today ($7 \, \text{kpc} < r < 9 \, \text{kpc}$), and the amount of radial migration they have experienced. The blue line is for the current (i.e. final) height, whereas the red line uses the height always at birth time of the stars. We see that stars that have migrated in to the Solar position had initially greater average height, while stars that moved out reduced their vertical scale height. To improve statistics, this plot stacked the different galaxy models.

the z_{birth} coordinate, i.e. the values of the z-coordinate of the stars the first time we find them in a snapshot file after they have been born.

Interestingly, the red line shows a pronounced trend of decreasing z_h with growing ΔR. Stars that had to migrated outward to reach the Solar circle have originally formed in an on average narrower component. This can be understood by the flaring of an approximately isothermal gas disk with radius when the surface density drops. The scale-height of stars forming in the outer parts of disks should then have a tendency to be higher than those forming in the inner regions.

As the blue line shows, this signal is however largely erased when the

instantaneous scale heights of the stars are computed as a function of the amount of radial migration they experienced. Here we might in principle still expect a negative correlation between average scale height and migration distance, because if stars reduce their vertical velocity dispersion as they migrate outwards, one may also expect that they reduce their vertical scale length as well. However, this assumes that the strength of the vertical gravitational field stays the same. If the surface density declines with radius (which it does), this is however not the case, and the narrowing of the stellar distributions can be prevented. Apparently, this effect is indeed operating in our simulations. In fact, if the stars begin with a relatively high velocity dispersion in the center, and if the surface density drops sufficiently with radius, then outward stellar migration could still lead to a thickened population, as argued by Roškar et al. (2013). The idea that stellar migration contributes to the formation of a thick disk component is therefore not fully excluded, but our results do not really favour such a scenario.

4.6   Resolution dependence

The radial migration phenomenon may in principle depend sensitively on resolution, because the effect is expected to be strongest at the corotation resonance (Sellwood & Binney 2002), and resolving resonances well is notoriously difficult (Weinberg & Katz 2002). Unfortunately, the exact numerical requirements on particle number and gravitational softening length in order to reliably account for radial migration are unknown. However, simulations of the same initial conditions carried out with widely different resolutions allow us to at least empirically check the consistency of our primary results with respect to numerical resolution.

To this end, we consider in Figure 4.9 the resolution dependence of results we obtained for the radial migration of disk stars obtained for model Aquarius-C, as simulated by Marinacci et al. (2014a). The three different resolution levels we consider differ in mass resolution by factors of eight,
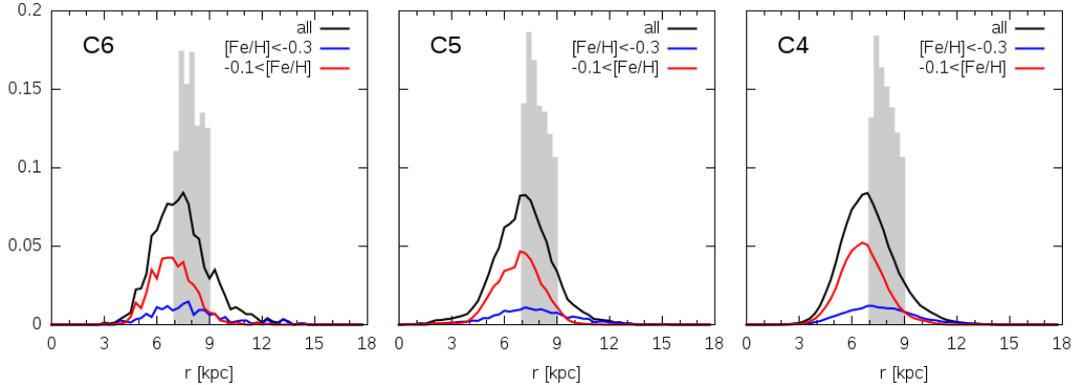
Figure 4.9: Resolution dependence of the birth radii distributions in the models Aq-C-6, Aq-C-5, and Aq-C-4. This plot repeats the measurements shown in the different panels of Fig. 4.4, except that here three different resolution runs of the C model are shown. It is reassuring that these different resolution agree rather well.

so that the total dynamic range covers in particle number (or equivalently mass resolution) is 64, that in spatial resolution (i.e. gravitational softening length) is a factor of 4. Reassuringly, the quantitative results obtained in Fig. 4.9 agree very well. This shows that not only the stellar mass and radial scale length of these forming disk galaxies agree quite well, the convergence also extends to more subtle dynamical properties of the stellar distributions. The employed numerical procedures appear hence to be quite robust and are not a significant source of uncertainty. Note however that it is unclear whether the underlying model for star formation and its regulation through feedback processes is reliable and reasonably close to reality. This therefore remains as a substantial systematic uncertainty.

4.7 Discussion and conclusions

The analysis we have presented in this chapter is the first study of radial migration in disk galaxies forming in self-consistent hydrodynamical cosmological simulations. Such simulations have only become available very

recently, and all previous work on radial migration had to rely on isolated toy simulations with ad-hoc initial conditions.

We have quantified the level of radial migration occurring for stars on nearly radial orbits in our disk galaxies. We have found that the radial redistribution of stars is consistent with a slow diffusion process that can be characterized by a diffusion constant of order $D \simeq 0.8\,\mathrm{kpc}^2\,\mathrm{Gyr}^{-1}$. At this rate, the impact of radial migration on radial metallicity gradients appears to be rather small, making it unlikely that radial migration significantly contributes to the scatter in the age-metallicity relation. We have also examined the change of the vertical structure of stellar populations as they migrate radially. Our results confirm that stellar populations reduce their vertical velocity dispersion as they migrate outwards. This cooling is consistent with theoretical models that assume conservation of the vertical action, but disagrees with the assumption that the vertical energy is conserved. Interestingly, at the same time we find that outwardly migrating stellar populations roughly maintain their thickness. While this does not strictly preclude the idea that radial migration contributes to the formation of a thick disk component, it also does not favor this in any way.

While the analysis of self-consistent hydrodynamic simulations eliminates any dependence on arbitrary choices for initial conditions, there are nevertheless important limitations of our approach. For example, the relatively large gravitational softening length used may artificially puff up the vertical disk structure, rendering the migration mechanisms less effective, and hiding some of the induced changes in the vertical dispersion. A similar concern applies to the effective equation of state subgrid model for star formation (Springel & Hernquist 2003) employed in the simulations, which may lead to artificially thick, flared stellar disks in the outer parts, and possibly to a distorted velocity distribution of the disk stars. Furthermore, the feedback processes assumed in the hydrodynamical simulations carry significant uncertainties, making the star formation histories predicted by the simulations uncertain, although the fact that the simulation models agree

quite well with global properties of late-time galaxies such as the Milky Way suggests that the model predictions are plausible. We note that our resolution study suggests that our results for radial migration should be robust with respect to numerical resolution.

In the coming years, better observational data from missions such as GAIA should become available that make a better theoretical understanding of radial migration quite pressing. We note that this is also required for the proper interpretation of chemical tagging techniques (e.g. Freeman & Bland-Hawthorn 2002). Numerical simulations are a highly useful tool for advancing our understanding of this topic and of galactic dynamics in general, and fully cosmological simulations of disk galaxy formation like the ones studied here are in principle ideal for this. It will be very interesting to push these calculations to much higher resolution such that also very subtle effects of galactic dynamics become accessible with them.

5

Concluding remarks

The exponential growth of computer performance in the recent past has dramatically increased the capabilities of numerical methods in Astronomy. Advanced cosmological simulations based on the ΛCDM paradigm supply us with crucial tests of the cosmological model and extremely detailed predictions of galaxy formation. Increasingly, the results of numerical simulations also inform and motivate observational searches for theoretically predicted effects. For the first time in history, we are now able to experiment with astrophysical processes through synthetic analogs created as computer simulations.

In this thesis we have carried out a set of controlled experiments with the halos of the Aquarius project of Milky Way-sized galaxies. We placed fully formed "hand made" disk galaxies into the centers of assembling dark matter halos and let them co-evolve. We were interested in testing whether thin disk galaxies can survive in realistic dark matter halos, despite their triaxiality and large content of substructures. If this could be answered positively, it would justify hopes that the failure to produce the right amount of such galaxies in all simulations thus far is due to an inaccurate treatment of the relevant physics rather than a more fundamental issue in the realm of cosmology. We found that, while thin disk galaxies that are carefully inserted in the centers of dark matter halos may be significantly tilted and strong instabilities like bars and spiral arms may be generated, their overall morphology as thin disks remains largely unchanged. We also confirmed that the presence of a central stellar bulge and/or a light stellar disk may stabilize against bar formation to some extent.

In the process of this work, we were confronted with the task to provide

steady-state galaxy models with prescribed properties that are in the best possible equilibrium with a given dark matter halo. To address this problem, we invented a new iterative method to derive stationary solutions of the collisionless Boltzmann equation. We developed a new code based on it, that we called GALIC. Our code creates a particle realization of the prescribed density distribution and then iteratively adjusts the velocities of individual particles, so that the time-averaged superposition of their orbits reproduces the prescribed density distribution. Solving Jeans equations for the second velocity moments, we developed a scheme that allows good guesses for the velocity adjustments through a stochastic process that prevents the method from getting stuck in local minima. As an additional benefit it allow us to converge to a solution with desired properties, for example with a certain velocity anisotropy, in a sometimes large space of degenerate solutions. In our numerical experiments with the set of created galaxy models, we found that all of them showed good stability going beyond what can be achieved with conventional moment-based methods that are widely used in the field. We made our code publicly available so that other researchers in the community can benefit from it and ideally improve it further. We have also done some exploratory work in this direction ourselves (not reported in this thesis), where we have extended the method from axisymmetry to full three dimensions, allowing us to also create triaxial equilibrium models.

Finally, we used recent hydrodynamical simulation of galaxy formation that successfully formed disk galaxies with realistic properties to study the origin of the spread of metallicities in the age-metallicity relation. We found that radial migration of stars, generally suspected to potentially have a strong influence on the metallicity distribution, has only a modest impact on metallicity gradients and therefore is unlikely to be responsible for the scatter in the age-metallicity relation. We caution that our findings are based on a specific simulation model that uses phenomenological parameterizations of star formation and associated feedback processes, which introduces systematic uncertainties. But it is very encouraging that our results

153

for radial migration in galaxies studied with this model at widely different numerical resolution are in excellent agreement.

Bibliography

Abel, T., Bryan, G. L., Norman, M. L. (2002), *The Formation of the First Star in the Universe*, Science, 295, 93

Agertz, O., Teyssier, R., Moore, B. (2011), *The formation of disc galaxies in a ΛCDM universe*, MNRAS, 410, 1391

Allgood, B., Flores, R. A., Primack, J. R., Kravtsov, A. V., Wechsler, R. H., Faltenbacher, A., Bullock, J. S. (2006), *The shape of dark matter haloes: dependence on mass, redshift, radius and formation*, MNRAS, 367, 1781

Alpher, R. A., Bethe, H., Gamow, G. (1948), *The Origin of Chemical Elements*, Physical Review, 73, 803

Alpher, R. A., Herman, R. C. (1948), *On the Relative Abundance of the Elements*, Physical Review, 74, 1737

Aragón-Calvo, M. A., Jones, B. J. T., van de Weygaert, R., van der Hulst, J. M. (2007), *The multiscale morphology filter: identifying and extracting spatial patterns in the galaxy distribution*, A&A, 474, 315

Athanassoula, E. (2002), *Bar-Halo Interaction and Bar Growth*, ApJL, 569, L83

Athanassoula, E., Misiriotis, A. (2002), *Morphology, photometry and kinematics of N-body bars - I. Three models with different halo central concentrations*, MNRAS, 330, 35

Aumer, M., White, S. D. M., Naab, T. (2014), *The diverse formation histories of simulated disc galaxies*, MNRAS, 441, 3679

Aumer, M., White, S. D. M., Naab, T., Scannapieco, C. (2013), *Towards a more realistic population of bright spiral galaxies in cosmological simulations*, MNRAS, 434, 3142

Barnes, J., Hut, P. (1986), *A hierarchical O(N log N) force-calculation algorithm*, Nature, 324, 446

Barnes, J. E., Hernquist, L. (1992), *Dynamics of interacting galaxies*, ARA&A, 30, 705

Berentzen, I., Shlosman, I. (2006), *Growing Live Disks within Cosmologically Assembling Asymmetric Halos: Washing Out the Halo Prolateness*, ApJ, 648, 807

Bergemann, M., Ruchti, G. R., Serenelli, A., Feltzing, S., Alves-Brito, A., Asplund, M., Bensby, T., Morbidelli, L., Sacco, G., Kordopatis, G., Tautvaišienė, G. (2014), *The Gaia-ESO Survey: radial metallicity gradients and age-metallicity relation of stars in the Milky Way disk*, A&A, 565, A89

Binney, J., Burnett, B., Kordopatis, G., Steinmetz, M., Gilmore, G., Bienayme, O., Bland-Hawthorn, J., Famaey, B., Grebel, E. K., Helmi, A., Navarro, J., Parker, Q., Reid, W. A., Seabroke, G., Siebert, A., Watson, F., Williams, M. E. K., Wyse, R. F. G., Zwitter, T. (2014), *Galactic kinematics and dynamics from Radial Velocity Experiment stars*, MNRAS, 439, 1231

Binney, J., Merrifield, M. (1998), *Galactic Astronomy*, Princeton University Press

Binney, J., Tremaine, S. (2008), *Galactic Dynamics: Second Edition*, Princeton University Press

Bird, S., Vogelsberger, M., Haehnelt, M., Sijacki, D., Genel, S., Torrey, P., Springel, V., Hernquist, L. (2014), *Damped Lyman-alpha absorbers as a probe of stellar feedback*, ArXiv e-prints, 1405.3994

Bovy, J., Rix, H.-W. (2013), *A Direct Dynamical Measurement of the Milky Way's Disk Surface Density Profile, Disk Scale Length, and Dark Matter Profile at $4\,kpc < R < 9\,kpc$*, ApJ, 779, 115

Bibliography

Brandt, A. (1977), *Multi-Level Adaptive Solutions to Boundary-Value Problems.*, Problems. Math. Comp., 31, 333

Bromm, V., Coppi, P. S., Larson, R. B. (2002), *The Formation of the First Stars. I. The Primordial Star-forming Cloud*, ApJ, 564, 23

Büdenbender, A., van de Ven, G., Watkins, L. L. (2014), *The tilt of the velocity ellipsoid in the Milky Way disk*, ArXiv e-prints, 1407.4808

Buyle, P., van Hese, E., de Rijcke, S., Dejonghe, H. (2007), *Radial-orbit instability of a family of anisotropic Hernquist models with and without a supermassive black hole*, MNRAS, 375, 1157

Cen, R. (1992), *A hydrodynamic approach to cosmology - Methodology*, ApJS, 78, 341

Chilingarian, I. V., Di Matteo, P., Combes, F., Melchior, A.-L., Semelin, B. (2010), *The GalMer database: galaxy mergers in the virtual observatory*, A&A, 518, A61

Cretton, N., de Zeeuw, P. T., van der Marel, R. P., Rix, H.-W. (1999), *Axisymmetric Three-Integral Models for Galaxies*, ApJS, 124, 383

Debattista, V. P., Mayer, L., Carollo, C. M., Moore, B., Wadsley, J., Quinn, T. (2006), *The Secular Evolution of Disk Structural Parameters*, ApJ, 645, 209

DeBuhr, J., Ma, C.-P., White, S. D. M. (2012), *Stellar discs in Aquarius dark matter haloes*, MNRAS, 426, 983

Dehnen, W. (2009), *Tailoring triaxial N-body models via a novel made-to-measure method*, MNRAS, 395, 1079

Di Matteo, P., Haywood, M., Combes, F., Semelin, B., Snaith, O. N. (2013), *Signatures of radial migration in barred galaxies: Azimuthal variations in the metallicity distribution of old stars*, A&A, 553, A102

Dicke, R. H., Peebles, P. J. E. (1979), *The big bang cosmology - enigmas and nostrums.*, in *General Relativity: An Einstein centenary survey*, edited by S. W. Hawking, W. Israel, 504–517

Diemand, J., Moore, B., Stadel, J. (2004), *Velocity and spatial biases in cold dark matter subhalo distributions*, MNRAS, 352, 535

D'Onghia, E., Springel, V., Hernquist, L., Keres, D. (2010), *Substructure Depletion in the Milky Way Halo by the Disk*, ApJ, 709, 1138

D'Onghia, E., Vogelsberger, M., Hernquist, L. (2013), *Self-perpetuating Spiral Arms in Disk Galaxies*, ApJ, 766, 34

Dressler, A. (1980), *Galaxy morphology in rich clusters - Implications for the formation and evolution of galaxies*, ApJ, 236, 351

Dubinski, J., da Costa, L. N., Goldwirth, D. S., Lecar, M., Piran, T. (1993), *Void evolution and the large-scale structure*, ApJ, 410, 458

Duffell, P. C., MacFadyen, A. I. (2011), *TESS: A Relativistic Hydrodynamics Code on a Moving Voronoi Mesh*, ApJS, 197, 15

Edvardsson, B., Andersen, J., Gustafsson, B., Lambert, D. L., Nissen, P. E., Tomkin, J. (1993), *The Chemical Evolution of the Galactic Disk - Part One - Analysis and Results*, A&A, 275, 101

Efstathiou, G., Lake, G., Negroponte, J. (1982), *The stability and masses of disc galaxies*, MNRAS, 199, 1069

Fedorenko, R. P. (1962), *A relaxation method for solving elliptic difference equations.*, Computational Math. and Math. Phys., 1, 1092

Freeman, K., Bland-Hawthorn, J. (2002), *The New Galaxy: Signatures of Its Formation*, ARA&A, 40, 487

Frenk, C. S., White, S. D. M. (2012), *Dark matter and cosmic structure*, Annalen der Physik, 524, 507

Bibliography

Gamow, G. (1948), *The Evolution of the Universe*, Nature, 162, 680

Gauthier, J.-R., Dubinski, J., Widrow, L. M. (2006), *Substructure around M31: Evolution and Effects*, ApJ, 653, 1180

Genel, S., Vogelsberger, M., Springel, V., Sijacki, D., Nelson, D., Snyder, G., Rodriguez-Gomez, V., Torrey, P., Hernquist, L. (2014), *Introducing the Illustris project: the evolution of galaxy populations across cosmic time*, MNRAS, 445, 175

Ghigna, S., Moore, B., Governato, F., Lake, G., Quinn, T., Stadel, J. (1998), *Dark matter haloes within clusters*, MNRAS, 300, 146

Gingold, R. A., Monaghan, J. J. (1977), *Smoothed particle hydrodynamics - Theory and application to non-spherical stars*, MNRAS, 181, 375

Governato, F., Brook, C., Mayer, L., Brooks, A., Rhee, G., Wadsley, J., Jonsson, P., Willman, B., Stinson, G., Quinn, T., Madau, P. (2010), *Bulgeless dwarf galaxies and dark matter cores from supernova-driven outflows*, Nature, 463, 203

Guedes, J., Callegari, S., Madau, P., Mayer, L. (2011), *Forming Realistic Late-type Spirals in a ΛCDM Universe: The Eris Simulation*, ApJ, 742, 76

Guo, Q., White, S., Li, C., Boylan-Kolchin, M. (2010), *How do galaxies populate dark matter haloes?*, MNRAS, 404, 1111

Guth, A. H. (1981), *Inflationary universe: A possible solution to the horizon and flatness problems*, Phys. Rev. D, 23, 347

Hansen, S. H., Moore, B. (2006), *A universal density slope Velocity anisotropy relation for relaxed structures*, New Astronomy, 11, 333

Harrison, E. R. (1970), *Fluctuations at the Threshold of Classical Cosmology*, Phys. Rev. D, 1, 2726

Hayashi, E., Navarro, J. F., Springel, V. (2007), *The shape of the gravitational potential in cold dark matter haloes*, MNRAS, 377, 50

Heng, K. (2014), *The Nature of Scientific Proof in the Age of Simulations*, American Scientist, 102, 174

Hernquist, L. (1990), *An analytical model for spherical galaxies and bulges*, ApJ, 356, 359

Hernquist, L. (1993), *N-body realizations of compound galaxies*, ApJS, 86, 389

Hernquist, L., Mihos, J. C. (1995), *Excitation of Activity in Galaxies by Minor Mergers*, ApJ, 448, 41

Hockney, R. W., Eastwood, J. W. (1981), *Computer Simulation Using Particles*

Holley-Bockelmann, K., Mihos, J. C., Sigurdsson, S., Hernquist, L. (2001), *Models of Cuspy Triaxial Galaxies*, ApJ, 549, 862

Hubble, E. (1929), *A Relation between Distance and Radial Velocity among Extra-Galactic Nebulae*, Proceedings of the National Academy of Science, 15, 168

Jaffe, W. (1983), *A simple model for the distribution of light in spherical galaxies*, MNRAS, 202, 995

Jalali, M. A., Tremaine, S. (2011), *Generalized Schwarzschild's method*, MNRAS, 410, 2003

Katz, N., Hernquist, L., Weinberg, D. H. (1992), *Galaxies and gas in a cold dark matter universe*, ApJL, 399, L109

Kazantzidis, S., Bullock, J. S., Zentner, A. R., Kravtsov, A. V., Moustakas, L. A. (2008), *Cold Dark Matter Substructure and Galactic Disks.*

Bibliography

I. Morphological Signatures of Hierarchical Satellite Accretion, ApJ, 688, 254

Kazantzidis, S., Zentner, A. R., Kravtsov, A. V., Bullock, J. S., Debattista, V. P. (2009), *Cold Dark Matter Substructure and Galactic Disks. II. Dynamical Effects of Hierarchical Satellite Accretion*, ApJ, 700, 1896

Kuijken, K., Dubinski, J. (1995), *Nearly Self-Consistent Disc / Bulge / Halo Models for Galaxies*, MNRAS, 277, 1341

Larson, R. B. (1974), *Effects of supernovae on the early evolution of galaxies*, MNRAS, 169, 229

Larson, R. B. (2000), *The Formation of the First Stars*, in *Star Formation from the Small to the Large Scale*, edited by F. Favata, A. Kaas, A. Wilson, volume 445 of *ESA Special Publication*, 13

Lemaître, G. (1927), *Un Univers homogène de masse constante et de rayon croissant rendant compte de la vitesse radiale des nébuleuses extragalactiques*, Annales de la Societe Scietifique de Bruxelles, 47, 49

Lemaître, G. (1931), *The Beginning of the World from the Point of View of Quantum Theory.*, Nature, 127, 706

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, 389

Lintott, C. J., Schawinski, K., Slosar, A., Land, K., Bamford, S., Thomas, D., Raddick, M. J., Nichol, R. C., Szalay, A., Andreescu, D., Murray, P., Vandenberg, J. (2008), *Galaxy Zoo: morphologies derived from visual inspection of galaxies from the Sloan Digital Sky Survey*, MNRAS, 389, 1179

Loebman, S. R., Roškar, R., Debattista, V. P., Ivezić, Ž., Quinn, T. R., Wadsley, J. (2011), *The Genesis of the Milky Way's Thick Disk Via Stellar Migration*, ApJ, 737, 8

Lucy, L. B. (1977), *A numerical approach to the testing of the fission hypothesis*, AJ, 82, 1013

Machado, R. E. G., Athanassoula, E. (2010), *Loss of halo triaxiality due to bar formation*, MNRAS, 406, 2386

Marinacci, F., Pakmor, R., Springel, V. (2014a), *The formation of disc galaxies in high-resolution moving-mesh cosmological simulations*, MNRAS, 437, 1750

Marinacci, F., Pakmor, R., Springel, V., Simpson, C. M. (2014b), *Diffuse gas properties and stellar metallicities in cosmological simulations of disc galaxy formation*, MNRAS, 442, 3745

Martig, M., Bournaud, F., Croton, D. J., Dekel, A., Teyssier, R. (2012), *A Diversity of Progenitors and Histories for Isolated Spiral Galaxies*, ApJ, 756, 26

Melvin, T., Masters, K., Lintott, C., Nichol, R. C., Simmons, B., Bamford, S. P., Casteels, K. R. V., Cheung, E., Edmondson, E. M., Fortson, L., Schawinski, K., Skibba, R. A., Smith, A. M., Willett, K. W. (2014), *Galaxy Zoo: an independent look at the evolution of the bar fraction over the last eight billion years from HST-COSMOS*, MNRAS, 438, 2882

Merritt, D. (1985), *Distribution functions for spherical galaxies*, MNRAS, 214, 25P

Minchev, I., Famaey, B. (2010), *A New Mechanism for Radial Migration in Galactic Disks: Spiral-Bar Resonance Overlap*, ApJ, 722, 112

Bibliography

Minchev, I., Famaey, B., Quillen, A. C., Dehnen, W., Martig, M., Siebert, A. (2012), *Radial migration does little for Galactic disc thickening*, A&A, 548, A127

Misner, C. W. (1969), *Mixmaster Universe*, Physical Review Letters, 22, 1071

Mo, H. J., Mao, S., White, S. D. M. (1998), *The formation of galactic discs*, MNRAS, 295, 319

Monaghan, J. J. (1992), *Smoothed particle hydrodynamics*, ARA&A, 30, 543

Moore, B., Ghigna, S., Governato, F., Lake, G., Quinn, T., Stadel, J., Tozzi, P. (1999), *Dark Matter Substructure within Galactic Halos*, ApJL, 524, L19

Moore, G. E. (1965), *Cramming more components onto integrated circuits*, Electronics Magazine, 4

Moster, B. P., Somerville, R. S., Maulbetsch, C., van den Bosch, F. C., Macciò, A. V., Naab, T., Oser, L. (2010), *Constraints on the Relationship between Stellar Mass and Halo Mass at Low and High Redshift*, ApJ, 710, 903

Navarro, J. F., Frenk, C. S., White, S. D. M. (1996), *The Structure of Cold Dark Matter Halos*, ApJ, 462, 563

Navarro, J. F., Frenk, C. S., White, S. D. M. (1997), *A Universal Density Profile from Hierarchical Clustering*, ApJ, 490, 493

Navarro, J. F., Ludlow, A., Springel, V., Wang, J., Vogelsberger, M., White, S. D. M., Jenkins, A., Frenk, C. S., Helmi, A. (2010), *The diversity and similarity of simulated cold dark matter haloes*, MNRAS, 402, 21

Okamoto, T., Isoe, M., Habe, A. (2014), *Cosmic evolution of bars in simulations of galaxy formation*, ArXiv e-prints, 1410.0442

Osipkov, L. P. (1979), *Spherical systems of gravitating bodies with an ellipsoidal velocity distribution*, Soviet Astronomy Letters, 5, 42

Ostriker, J. P., Spitzer, Jr., L., Chevalier, R. A. (1972), *On the Evolution of Globular Clusters*, ApJL, 176, L51

Pakmor, R., Marinacci, F., Springel, V. (2014), *Magnetic Fields in Cosmological Simulations of Disk Galaxies*, ApJL, 783, L20

Penzias, A. A., Wilson, R. W. (1965), *A Measurement of Excess Antenna Temperature at 4080 Mc/s.*, ApJ, 142, 419

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., Project, T. S. C. (1999), *Measurements of Ω and Λ from 42 High-Redshift Supernovae*, ApJ, 517, 565

Planck Collaboration, Ade, P. A. R., Aghanim, N., Alves, M. I. R., Armitage-Caplan, C., Arnaud, M., Ashdown, M., Atrio-Barandela, F., Aumont, J., Aussel, H., et al. (2013), *Planck 2013 results. I. Overview of products and scientific results*, ArXiv e-prints, 1303.5062

Quillen, A. C., Minchev, I., Bland-Hawthorn, J., Haywood, M. (2009), *Radial mixing in the outer Milky Way disc caused by an orbiting satellite*, MNRAS, 397, 1599

Read, J. I., Lake, G., Agertz, O., Debattista, V. P. (2008), *Thin, thick and dark discs in ΛCDM*, MNRAS, 389, 1041

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.,

Bibliography

Leibundgut, B., Phillips, M. M., Reiss, D., Schmidt, B. P., Schommer, R. A., Smith, R. C., Spyromilio, J., Stubbs, C., Suntzeff, N. B., Tonry, J. (1998), *Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant*, AJ, 116, 1009

Rodionov, S. A., Athanassoula, E., Sotnikova, N. Y. (2009), *An iterative method for constructing equilibrium phase models of stellar systems*, MN-RAS, 392, 904

Roškar, R., Debattista, V. P. (2014), *Radial Migration in Spiral Galaxies*, ArXiv e-prints, 1406.4515

Roškar, R., Debattista, V. P., Loebman, S. R. (2013), *The effects of radial migration on the vertical structure of Galactic discs*, MNRAS, 433, 976

Roškar, R., Debattista, V. P., Quinn, T. R., Stinson, G. S., Wadsley, J. (2008), *Riding the Spiral Waves: Implications of Stellar Migration for the Properties of Galactic Disks*, ApJL, 684, L79

Roškar, R., Debattista, V. P., Quinn, T. R., Wadsley, J. (2012), *Radial migration in disc galaxies - I. Transient spiral structure and dynamics*, MNRAS, 426, 2089

Rubin, V. C., Ford, Jr., W. K. (1970), *Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions*, ApJ, 159, 379

Ryan, Jr., R. E., Cohen, S. H., Windhorst, R. A., Silk, J. (2008), *Galaxy Mergers at z 1 in the HUDF: Evidence for a Peak in the Major Merger Rate of Massive Galaxies*, ApJ, 678, 751

Satoh, C. (1980), *Dynamical Models of Axisymmetric Galaxies and Their Applications to the Elliptical Galaxy NGC4697*, PASJ, 32, 41

Scannapieco, C., Wadepuhl, M., Parry, O. H., Navarro, J. F., Jenkins, A., Springel, V., Teyssier, R., Carlson, E., Couchman, H. M. P., Crain, R. A., Dalla Vecchia, C., Frenk, C. S., Kobayashi, C., Monaco, P., Murante, G.,

Okamoto, T., Quinn, T., Schaye, J., Stinson, G. S., Theuns, T., Wadsley, J., White, S. D. M., Woods, R. (2012), *The Aquila comparison project: the effects of feedback and numerical methods on simulations of galaxy formation*, MNRAS, 423, 1726

Scannapieco, C., White, S. D. M., Springel, V., Tissera, P. B. (2009), *The formation and survival of discs in a ΛCDM universe*, MNRAS, 396, 696

Schönrich, R., Binney, J. (2009), *Chemical evolution with radial mixing*, MNRAS, 396, 203

Schwarzschild, M. (1979), *A numerical model for a triaxial stellar system in dynamical equilibrium*, ApJ, 232, 236

Searle, L., Zinn, R. (1978), *Compositions of halo clusters and the formation of the galactic halo*, ApJ, 225, 357

Sellwood, J. A. (2013), *Dynamics of Disks and Warps*, 923

Sellwood, J. A., Binney, J. J. (2002), *Radial mixing in galactic discs*, MNRAS, 336, 785

Sellwood, J. A., Nelson, R. W., Tremaine, S. (1998), *Resonant Thickening of Disks by Small Satellite Galaxies*, ApJ, 506, 590

Shen, S., Mo, H. J., White, S. D. M., Blanton, M. R., Kauffmann, G., Voges, W., Brinkmann, J., Csabai, I. (2003), *The size distribution of galaxies in the Sloan Digital Sky Survey*, MNRAS, 343, 978

Siebert, A., Bienaymé, O., Binney, J., Bland-Hawthorn, J., Campbell, R., Freeman, K. C., Gibson, B. K., Gilmore, G., Grebel, E. K., Helmi, A., Munari, U., Navarro, J. F. (2008a), *Estimation of the tilt of the stellar velocity ellipsoid from RAVE and implications for mass models*, MNRAS, 391, 793

167

Bibliography

Siebert, A., Bienaymé, O., Binney, J., Bland-Hawthorn, J., Campbell, R., Freeman, K. C., Gibson, B. K., Gilmore, G., Grebel, E. K., Helmi, A., Munari, U., Navarro, J. F., Parker, Q. A., Seabroke, G., Siviero, A., Steinmetz, M., Williams, M., Wyse, R. F. G., Zwitter, T. (2008b), *Estimation of the tilt of the stellar velocity ellipsoid from RAVE and implications for mass models*, MNRAS, 391, 793

Silk, J., Rees, M. J. (1998), *Quasars and galaxy formation*, A&A, 331, L1

Smoot, G. F., Bennett, C. L., Kogut, A., Wright, E. L., Aymon, J., Boggess, N. W., Cheng, E. S., de Amici, G., Gulkis, S., Hauser, M. G., Hinshaw, G., Jackson, P. D., Silverberg, R. F., Tenorio, L., Weiss, R., Wilkinson, D. T. (1992), *Structure in the COBE differential microwave radiometer first-year maps*, ApJL, 396, L1

Solway, M., Sellwood, J. A., Schönrich, R. (2012), *Radial migration in galactic thick discs*, MNRAS, 422, 1363

Springel, V. (2005a), *The cosmological simulation code GADGET-2*, MNRAS, 364, 1105

Springel, V. (2005b), *The cosmological simulation code GADGET-2*, MNRAS, 364, 1105

Springel, V. (2010), *E pur si muove: Galilean-invariant cosmological hydrodynamical simulations on a moving mesh*, MNRAS, 401, 791

Springel, V., Di Matteo, T., Hernquist, L. (2005a), *Modelling feedback from stars and black holes in galaxy mergers*, MNRAS, 361, 776

Springel, V., Hernquist, L. (2003), *Cosmological smoothed particle hydrodynamics simulations: a hybrid multiphase model for star formation*, MNRAS, 339, 289

168

Springel, V., Wang, J., Vogelsberger, M., Ludlow, A., Jenkins, A., Helmi, A., Navarro, J. F., Frenk, C. S., White, S. D. M. (2008a), *The Aquarius Project: the subhaloes of galactic haloes*, MNRAS, 391, 1685

Springel, V., White, S. D. M. (1999), *Tidal tails in cold dark matter cosmologies*, MNRAS, 307, 162

Springel, V., White, S. D. M., Frenk, C. S., Navarro, J. F., Jenkins, A., Vogelsberger, M., Wang, J., Ludlow, A., Helmi, A. (2008b), *Prospects for detecting supersymmetric dark matter in the Galactic halo*, Nature, 456, 73

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., Pearce, F. (2005b), *Simulations of the formation, evolution and clustering of galaxies and quasars*, Nature, 435, 629

Springel, V., White, S. D. M., Tormen, G., Kauffmann, G. (2001a), *Populating a cluster of galaxies - I. Results at z=0*, MNRAS, 328, 726

Springel, V., Yoshida, N., White, S. D. M. (2001b), *GADGET: a code for collisionless and gasdynamical cosmological simulations*, New Astronomy, 6, 79

Steinmetz, M., Mueller, E. (1994), *The formation of disk galaxies in a cosmological context: Populations, metallicities and metallicity gradients*, A&A, 281, L97

Stinson, G. S., Bovy, J., Rix, H.-W., Brook, C., Roškar, R., Dalcanton, J. J., Macciò, A. V., Wadsley, J., Couchman, H. M. P., Quinn, T. R. (2013), *MaGICC thick disc - I. Comparing a simulated disc formed with stellar feedback to the Milky Way*, MNRAS, 436, 625

Bibliography

Stone, J. M., Norman, M. L. (1992), *ZEUS-2D: A radiation magnetohydrodynamics code for astrophysical flows in two space dimensions. I - The hydrodynamic algorithms and tests.*, ApJS, 80, 753

Syer, D., Tremaine, S. (1996), *Made-to-measure N-body systems*, MNRAS, 282, 223

Teyssier, R. (2002), *Cosmological hydrodynamics with adaptive mesh refinement. A new high resolution code called RAMSES*, A&A, 385, 337

Toomre, A. (1964), *On the gravitational stability of a disk of stars*, ApJ, 139, 1217

van den Bosch, R. C. E., Gebhardt, K., Gültekin, K., van de Ven, G., van der Wel, A., Walsh, J. L. (2012), *An over-massive black hole in the compact lenticular galaxy NGC1277*, Nature, 491, 729

Vandervoort, P. O. (1984), *On Schwarzschild's method for the construction of model galaxies*, ApJ, 287, 475

Velazquez, H., White, S. D. M. (1999), *Sinking satellites and the heating of galaxy discs*, MNRAS, 304, 254

Vogelsberger, M., Genel, S., Sijacki, D., Torrey, P., Springel, V., Hernquist, L. (2013), *A model for cosmological simulations of galaxy formation physics*, MNRAS, 436, 3031

Vogelsberger, M., Genel, S., Springel, V., Torrey, P., Sijacki, D., Xu, D., Snyder, G., Bird, S., Nelson, D., Hernquist, L. (2014a), *Properties of galaxies reproduced by a hydrodynamic simulation*, Nature, 509, 177

Vogelsberger, M., Genel, S., Springel, V., Torrey, P., Sijacki, D., Xu, D., Snyder, G., Nelson, D., Hernquist, L. (2014b), *Introducing the Illustris Project: simulating the coevolution of dark and visible matter in the Universe*, MNRAS, 444, 1518

Vogelsberger, M., Genel, S., Springel, V., Torrey, P., Sijacki, D., Xu, D., Snyder, G., Nelson, D., Hernquist, L. (2014c), *Introducing the Illustris Project: simulating the coevolution of dark and visible matter in the Universe*, MNRAS, 444, 1518

Walker, I. R., Mihos, J. C., Hernquist, L. (1996), *Quantifying the Fragility of Galactic Disks in Minor Mergers*, ApJ, 460, 121

Weinberg, M. D., Katz, N. (2002), *Bar-driven Dark Halo Evolution: A Resolution of the Cusp-Core Controversy*, ApJ, 580, 627

White, S. D. M., Rees, M. J. (1978), *Core condensation in heavy halos - A two-stage theory for galaxy formation and clustering*, MNRAS, 183, 341

Widrow, L. M., Dubinski, J. (2005), *Equilibrium Disk-Bulge-Halo Models for the Milky Way and Andromeda Galaxies*, ApJ, 631, 838

Wielen, R. (1977), *The diffusion of stellar orbits derived from the observed age-dependence of the velocity dispersion*, A&A, 60, 263

Wielen, R., Fuchs, B., Dettbarn, C. (1996), *On the birth-place of the Sun and the places of formation of other nearby stars*, A&A, 314, 438

Xu, D. D., Mao, S., Wang, J., Springel, V., Gao, L., White, S. D. M., Frenk, C. S., Jenkins, A., Li, G., Navarro, J. F. (2009), *Effects of dark matter substructures on gravitational lensing: results from the Aquarius simulations*, MNRAS, 398, 1235

Yurin, D., Springel, V. (2014), *An iterative method for the construction of N-body galaxy models in collisionless equilibrium*, MNRAS, 444, 62

Zel'dovich, Y. B. (1970), *Gravitational instability: An approximate theory for large density perturbations.*, A&A, 5, 84

Zeldovich, Y. B. (1972), *A hypothesis, unifying the structure and the entropy of the Universe*, MNRAS, 160, 1P

Bibliography

Zhang, Y., Yang, X., Wang, H., Wang, L., Luo, W., Mo, H. J., van den Bosch, F. C. (2014), *Spin alignments of spiral galaxies within the large-scale structure from SDSS DR7*, ArXiv e-prints, 1409.7150

Zwicky, F. (1933), *Die Rotverschiebung von extragalaktischen Nebeln*, Helvetica Physica Acta, 6, 110

Acknowledgements