
Simulating structure formation with high precision: numerical techniques, dynamics and the evolution of substructures

Francesca Iannuzzi



München 2012

Simulating structure formation with high precision: numerical techniques, dynamics and the evolution of substructures

Francesca Iannuzzi

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

vorgelegt von
Francesca Iannuzzi
aus Mailand, Italien

München, den 9. März 2012

Erstgutachter: Simon White

Zweitgutachter: Andreas Burkert

Tag der mündlichen Prüfung: 11. Mai 2012

*Well, I know not
What counts harsh Fortune casts upon my face,
But in my bosom shall she never come
To make my heart her vassal.*

Shakespeare

Abstract

Numerical simulations are of paramount importance in the study of structure formation in a cosmological context: only they can provide accurate predictions to be tested against observations. Within the computation, the cosmic fluid is represented by a set of point particles. These are far less numerous than the actual building blocks of the system and represent a Monte Carlo sampling of the underlying phase-space density. In simulations involving only the dark-matter component of the cosmic fluid, these particles are made to interact solely through gravity. The presence of a limited number of particles and, consequently, the significant mass that each one of them carries, is at the origin of unwanted *discreteness* effects; these have an impact on the performance and the accuracy of the modelling and need to be tackled with care. Softening of the gravitational force at small scales is introduced to moderate the effect of spurious collisions on the dynamics of the system; as a consequence, the interaction loses its Newtonian form below a scale given by the *softening length* and this sets a resolution limit below which the results of the simulation cannot be trusted. This scale is normally fixed in space and time throughout the simulation; this is not an ideal choice in a cosmological scenario, though, where the high level of inhomogeneity in the matter distribution requires, in principle, a different treatment in each region. In the first part of this thesis, we introduce the problem of adapting the gravitational softening according to the density of the environment and present the implementation of the algorithm in the cosmological simulation code GADGET-3. From the analysis of its impact on simulations of structure formation, we conclude that adaptive gravitational softening enhances the clustering of particles on small scales and anticipates the results obtained from simulations employing a considerably higher number of particles. As will be shown, the extension of the method to scenarios where the baryonic component of the cosmic fluid is also taken into account is very promising, but not trivial.

Even though ordinary, baryonic matter is believed to constitute a small fraction of the matter in the Universe, its modelling is of fundamental importance to the purpose of comparing the outcome of simulations to observational results. Yet, the accurate treatment of hydrodynamics and the handling of various astrophysical processes at work is a problem of great complexity on which little consensus has been achieved. Semi-analytic models of galaxy formation and evolution provide a powerful approach to the problem. Instead of following the evolution of dark matter and baryons self-consistently, they post-process existing, dark-matter-only simulations and model the behaviour of the baryonic component by means of a set of physically-motivated, analytical prescriptions. In the second part of this thesis, we make use of a state-of-the-art semi-analytic model to investigate the relationship between the orbital properties of satellite galaxies and their star-formation activity. Galaxy clusters are the largest virialised structures observed in our Universe; when galaxies join such a dense environment and become *satellites* orbiting within it, several processes act in changing their internal properties. We study the impact that the initial orbits of these galaxies have on the subsequent evolution of star-formation properties or related quantities. We conclude that this impact is strong and its imprint can still be observed at late times in the evolution of the cluster members.

Zusammenfassung

Für die Untersuchung der Strukturbildung in der Kosmologie sind numerische Simulationen von herausragender Bedeutung, da nur sie präzise Voraussagen erlauben, welche mit Beobachtungsdaten verglichen werden können. In der Simulation wird das kosmische Fluid durch punktförmige Teilchen beschrieben. Diese Teilchen sind weit weniger zahlreich als die eigentlichen Bausteine des Systems und stellen eine Monte-Carlo-Stichprobe der Phasenraum-dichte dar, welche diesem System zugrunde liegt. In dem Fall, dass lediglich diejenige Komponente des kosmischen Fluids simuliert wird, welche die Dunkle Materie beschreibt, kann die Wechselwirkung dieser Teilchen ausschließlich durch Gravitationskräfte beschrieben werden. Aufgrund der endlichen Teilchenzahl und der daraus folgenden signifikanten Masse eines jeden dieser Teilchen, ergeben sich für unerwünschte diskrete Effekte. Diese Effekte haben Einfluß auf die Effizienz und Genauigkeit der Simulation und müssen daher sorgfältig behandelt werden. Um den Einfluß von störenden Teilchen-Kollisionen auf die Dynamik des Systems zu mildern, wurde ein sogenanntes Softening der Gravitationskraft auf kleinen Skalen eingeführt. Dadurch verliert die Wechselwirkung unterhalb einer gewissen Skala, welche durch die Softening-Länge definiert ist, ihre Newton'sche Natur. Gleichzeitig resultiert daraus eine Auflösungsgrenze, unterhalb welcher den Ergebnissen der Simulationen nicht mehr vertraut werden kann. Üblicherweise bleibt die erwähnte Skala innerhalb der gesamten Simulation in Raum und Zeit konstant. Diese Vorgehensweise ist jedoch nicht ideal für kosmologische Szenarien, in denen es die hohe Inhomogenität in der Verteilung der Materie im Prinzip nötig macht, jeden Region unterschiedlich zu behandeln. Im ersten Teil dieser Arbeit widme ich mich dem Problem, wie das Gravitations-Softening der Teilchendichte der Umgebung angepasst werden kann. Ferner präsentiere ich die Implementierung des daraus resultierenden Algorithmus in den kosmologischen Simulationscode GADGET-3. Der Einfluß des Algorithmus auf Simulationen der Strukturentstehung wird untersucht und führt

zu der Schlussfolgerung, dass adaptives Gravitations-Softening das Klumpen von Teilchen auf kleinen Skalen verstärkt und es außerdem erlaubt, die Ergebnisse von Simulationen mit erheblich höheren Teilchenzahlen vorwegzunehmen. Wie gezeigt wird, ist der Erweiterung dieser Methode auf Szenarien, in denen zusätzlich die baryonische Komponente des kosmologischen Fluids berücksichtigt wird, viel versprechend wenn auch nicht-trivial.

Obwohl man annimmt, dass gewöhnliche baryonische Materie nur einen kleinen Teil der Gesamtmaterie im Universum ausmacht, ist deren Modellierung fundamental wichtig um Simulationsergebnisse mit Beobachtungen vergleichen zu können. Allerdings ist die exakte Behandlung der Hydrodynamik und der verschiedenen astrophysikalischen Prozesse eine sehr komplexe Aufgabe, bei der noch wenig Einigkeit herrscht. Semi-analytische Galaxienentstehungs- und Evolutionsmodelle bieten einen sehr leistungsfähigen Lösungsansatz für dieses Problem. Anstatt der Evolution von Dunkler Materie und Baryonen selbstkonsistent zu folgen, wird das Verhalten der baryonischen Komponente modelliert, ausgehend von bereits existierenden Simulationen mit rein Dunkler Materie sowie physikalisch motivierten, analytischen Beschreibungen. Im zweiten Teil dieser Doktorarbeit verwenden wir ein semi-analytisches Modell auf dem aktuellen Stand der Forschung, um die Beziehung zwischen den Orbiteigenschaften von Satellitengalaxien und deren Sternentstehungsaktivität zu untersuchen. Galaxienhaufen sind die größten beobachteten virialisierten Strukturen in unserem Universum; wenn Galaxien in solch dichte Umgebungen eintreten und darin zu umlaufenden Satelliten werden, führen verschiedene Prozesse, zu eine Änderung ihrer internen Eigenschaften. Wir untersuchen, welchen Einfluß der Anfangsorbit dieser Galaxien auf die spätere Entwicklung der Sternentstehungseigenschaften oder anderer verwandter Größen hat. Zusammenfassend lässt sich sagen, dass dieser Einfluß stark ist und Auswirkungen auch zu späteren Zeiten in der Entwicklung von Haufenmitgliedern beobachtet werden kann.

Contents

Abstract	vii
Zusammenfassung	ix
1 Elements of Cosmology	1
1.1 The geometry of the Universe	1
1.2 The dynamics of the Universe	3
1.3 Observational constraints	6
1.4 Evolution of cosmic structures	8
1.4.1 Jeans Instability and the Linear Theory	9
1.4.2 Application to the cosmological case	11
1.4.3 Hot vs. cold dark matter	13
1.4.4 Beyond linear theory	14
1.5 The Standard Model - Summary	18
2 Numerical techniques	21
2.1 Introduction	21
2.2 Initial conditions	25
2.3 Simulations	27
2.3.1 Force computation: gravity	27
2.3.2 Force computation: hydrodynamics	30
2.3.3 Time integration	34
2.3.4 Codes	36
2.4 Structure finders	37
2.4.1 The FOF method and the substructure finder SUBFIND	39
2.5 Semi-analytic models of galaxy formation and evolution	42
2.6 Summary and outline of the thesis	44

3 Adaptive gravitational softening I: the single-species case	47
3.1 Introduction	48
3.2 The formalism	50
3.3 Implementation in Gadget	52
3.4 Tests	53
3.4.1 Evolution of a Plummer sphere	54
3.4.2 Evolution of a Hernquist model	56
3.4.3 Equilibrium structure of a polytrope	58
3.5 Performance in a cosmological environment	60
3.5.1 Global behaviour I - Clustering	63
3.5.2 Global behaviour II - Mass function	64
3.5.3 Internal halo properties	67
3.5.4 Comments	69
3.6 Simulating a mini version of the Millennium-II	72
3.7 Conclusions	74
4 Adaptive gravitational softening II: the multiple-species case	77
4.1 Introduction	77
4.2 A spherical galaxy with dark matter and stars	79
4.3 Cosmological simulations I: two dark matter species	82
4.4 Cosmological simulations II: dark matter and gas	89
4.5 A simulation of dynamical friction	92
4.6 Conclusions	99
5 On the orbital and internal evolution of cluster galaxies	101
5.1 Introduction	102
5.2 The semi-analytic models and the selected sample	104
5.3 Results	105
5.3.1 Anisotropy at redshift zero	106
5.3.2 Anisotropy at high redshift	107
5.3.3 Anisotropy at infall	109
5.3.4 Evolution of anisotropy in time	111
5.4 Summary	116
5.5 Discussion	118
6 Conclusions	123
A Relevant quantities for cubic spline softening	127
Bibliography	129
Acknowledgements	140
Publications	145

Elements of Cosmology

In this chapter we will briefly review the fundamentals of Cosmology, the study of the Universe in its entirety. We will start by analysing the consequences of the application of Einstein's theory of gravitation and subsequently unveil the issue of structure formation, towards the definition of the Standard Cosmological Model, believed to provide the best description of the Universe developed to date. The concepts introduced here will serve as a theoretical background for the discussion of the following chapters, where we will progressively restrict the analysis to some specific aspects of today's cosmological studies.

1.1 The geometry of the Universe

Since the early years of the 20th century, when the first steps were taken towards the construction of a scientific theory of the Universe, it became clear that some sort of simplifying assumption was needed in order to overcome the difficulties due to our poor understanding and knowledge about the distribution of matter in the Universe. The starting point is the observational evidence that the Universe, as observed from our point of view, appears isotropic (see, e.g., Fig. 1.1). When combining this evidence with the so-called *Copernican Principle*, stating that we are not privileged observers, one concludes that the Universe must appear isotropic from every point of observation. This last statement is known as the *Cosmological Principle*.

In General Relativity, the properties of symmetry in the distribution of matter directly translate into geometrical properties of space-time. The simplification introduced by the Cosmological Principle is to restrict the possible space-time metrics for the Universe to the so-called *Robertson-Walker metric* (Robertson 1935; 1936a; 1936b; Walker 1937):

$$ds^2 = (cdt)^2 - dl^2 = (cdt)^2 - a(t)^2 \left[\frac{dr^2}{1 - Kr^2} + r^2(d\theta^2 + \sin^2\theta d\phi^2) \right], \quad (1.1)$$

where r , θ and ϕ are the comoving, spherical coordinates and t is the cosmic time. The metric is completely determined by a "radius" $a(t)$, with the dimension of a




Figure 1.1: Map of the galaxy distribution in the nearby universe, obtained from the Two-degree-Field Galaxy Redshift Survey (2dFGRS; Colless et al. 2001). The volume covered by the survey is approximately $10^8 h^{-1} \text{ Mpc}^3$ (for a definition of h , see Footnote 1; $1 \text{ pc} = 3.08568025 \times 10^{18} \text{ cm}$). The most faraway object is found at a distance of $\approx 600 h^{-1} \text{ Mpc}$. Statistical tests applied to these maps confirm that isotropy applies to different viewpoints within them.

length, and a dimensionless constant K ; these quantities are referred to as the *cosmic scale factor* and the *curvature parameter*, respectively. It is possible to scale K so that it takes the values $1, 0, -1$ only. For $K = 0$ the spatial line element dl^2 reduces to:

$$dl^2 = a(t)^2 [dr^2 + r^2(d\theta^2 + \sin^2\theta d\phi^2)], \quad (1.2)$$

which is the line element of an Euclidean space. For $K = 1$, substituting $r = \sin\chi$, we obtain:

$$dl^2 = a(t)^2 [d\chi^2 + \sin^2\chi(d\theta^2 + \sin^2\theta d\phi^2)], \quad (1.3)$$

which is the line element of a hypersphere of radius $a(t)$. Finally, for $K = -1$ we substitute $r = \sinh\chi$ and obtain:

$$dl^2 = a(t)^2 [d\chi^2 + \sinh^2\chi(d\theta^2 + \sin^2\theta d\phi^2)], \quad (1.4)$$

which is the line element of a pseudosphere. Under the assumption of the Cosmological Principle, the three possible geometries for the Universe are therefore the *flat*, *spherical* and *hyperbolic* ones.

The cosmic scale factor accounts for the time evolution of the spatial line element: in case it is not constant, the space-time is either expanding or contracting. The value and behaviour of $a(t)$ and K are intimately related to the matter content of the Universe and we will discuss this further in the following section. For the moment, we just mention that observational evidence supports the scenario of an expanding universe. This was discovered at the beginning of the last century from

observations relating the distance of galaxies to the features of the emitted radiation (Lemaître, 1927; Hubble, 1929; Lemaître, 1931). The light coming from more distant objects was found to be shifted towards lower frequencies (“red-shifted”), indicating a recessive motion with increasing velocity at larger distances. The relation between distance and recession velocity is linear and is termed *Hubble Law*; the exact proportionality between the two quantities is regulated by the *Hubble constant*, $H(t) \equiv \dot{a}/a$ ¹. The cosmological *redshift*, relating the wavelength of the radiation emitted at time t (λ_t) to that observed at the present time (λ_0), is defined as:

$$z_t \equiv \frac{\lambda_0 - \lambda_t}{\lambda_t} = \frac{a_0 - a_t}{a_t}. \quad (1.5)$$

Redshifts have the advantage to be directly observable and are more extensively used in the description of the evolution of the Universe than scale factors are.

1.2 The dynamics of the Universe

We will now discuss the application of General Relativity to the cosmological case and the link between matter content, geometry and evolution of the Universe. In Einstein’s theory of gravitation the field equations relating the matter-energy content of the Universe to the curvature of spacetime are:

$$G_{\mu\nu} = \frac{8\pi G}{c^2} T_{\mu\nu}, \quad (1.6)$$

where $G_{\mu\nu}$ is the Einstein tensor, containing information on the components of the space-time metric, and $T_{\mu\nu}$ is the energy-momentum tensor describing the source term. The relevant components of $G_{\mu\nu}$ for the cosmological application are derived from the Robertson-Walker metric. Regarding the energy-momentum tensor $T_{\mu\nu}$, in cosmology it is usually assumed that the matter-energy content of the Universe can be described by a *perfect fluid*; this choice has its justification in that the mean free path of the fluid elements under consideration is much smaller than the scales of physical interest, a limit where the fluid can indeed be regarded as perfect. Such a fluid is entirely described in terms of its rest-frame energy density ρ and isotropic pressure P ; the corresponding energy-momentum tensor takes the form:

$$T_{\mu\nu} = \left(\rho + \frac{P}{c^2} \right) u_\mu u_\nu - \frac{P}{c^2} g_{\mu\nu}, \quad (1.7)$$

where u_α is the fluid four-velocity and $g_{\mu\nu}$ is the metric tensor.

By combining these assumptions, i.e. the perfect fluid approximation and the Cosmological Principle, into Eq. 1.6, we obtain the following field equations:

$$\ddot{a} = -\frac{4\pi G}{3} \left(\rho + \frac{3P}{c^2} \right) a, \quad (1.8)$$

¹The value of the Hubble constant at the present time is expressed as $H_0 = 100 h \text{ km s}^{-1} \text{ Mpc}^{-1}$, where the dimensionless parameter h carries along the observational uncertainties.

$$\dot{a}^2 + Kc^2 = \frac{8\pi G}{3}\rho a^2, \quad (1.9)$$

better known as the *Friedmann equations*. Along with the expression for energy conservation²

$$d(\rho a^3) = -3\frac{P}{c^2}a^2 da, \quad (1.10)$$

and the equation of state for the perfect fluid, these two equations describe the time evolution of $a(t)$, $\rho(t)$ and $P(t)$.

We will now investigate some general properties of the standard cosmological models described by Eqs. 1.8 and 1.9. These come under the definition of *Friedmann models*, after Alexander Friedmann, who first proposed them in 1922 (Friedmann, 1922). We will start by specifying the equation of state for the cosmic fluid; this is widely assumed to take the form

$$P = w\rho c^2, \quad (1.11)$$

which means that the description of the cosmic fluid is reduced to the value of a single parameter - the state parameter w . By combining the expression for the equation of state and that for the conservation of energy, we obtain the evolution of the energy density ρ as a function of a and w :

$$\rho_w a^{3(1+w)} = \text{const} = \rho_{0w} a_0^{3(1+w)}, \quad (1.12)$$

where the subscript 0 indicates present time. We will only list the three most relevant cases for cosmology:

- $w = 0$: pressureless material (*dust*). This is a good approximation to any form of non-relativistic fluid, where the effect of pressure can be neglected. The matter component of the Universe is regarded as being pressureless and its density evolution is described by the relation

$$\rho_m a^3 = \rho_{0m} a_0^3 \implies \rho_m = \rho_{0m} (1+z)^3; \quad (1.13)$$

- $w = \frac{1}{3}$: radiative fluid. Such a state parameter describes a fluid of relativistic particles in thermal equilibrium. For this radiative component of the Universe, the evolution of ρ is then given by:

$$\rho_r a^4 = \rho_{0r} a_0^4 \implies \rho_r = \rho_{0r} (1+z)^4; \quad (1.14)$$

- $w = -1$: negative pressure. The equation of state implied by this choice of w is generally associated to the perfect fluid equivalent of a cosmological constant (see next section). It is immediate to note that the energy density of such a component undergoes no evolution with a .

²This is obtained from the expression $T_{;\nu}^{\mu\nu} = 0$, where the symbol $;$ stand for the covariant four-divergence.




Figure 1.2: Schematic representation for the time evolution of the Friedmann models described in Sec. 1.2. All models start with a singularity at $t = 0$, corresponding to a null scale factor (grey circles). As explained in the text, the late-time evolution of a depends on the curvature parameter K . All models have been normalised to the present time t_0 .

There is no need to specify w in order to derive the main features of the Friedmann models and we can proceed without assuming a composition for the cosmic fluid. The two major results are (i) that the *past* of models with $-1/3 < w < 1$ is characterised by a singularity, a point in time where a vanishes and the density diverges and (ii) that their *future* depends critically on their curvature. It is quite straightforward to reach these conclusions by just having a look at the Friedmann equations. We know that at present $a > 0$ (by definition) and $\dot{a} > 0$ (by observation); the first of the equations then tells us that, provided w is within the range $(-1/3, 1)$, the acceleration \ddot{a} will be negative: this means that the curve $a(t)$ vs. t is concave downwards and so must have reached $a(t) = 0$ at some finite time in the past. This instant is called *Big Bang singularity* and is labelled with $t = 0$. Because of the concave shape of $a(t)$, the time elapsed between the singularity and the epoch t must always be less than the *Hubble time*, H^{-1} . Regarding the final destiny of these models, we recall, from Eq. 1.10, that the density ρ must decrease with increasing a at least as fast as a^{-3} , provided the pressure does not become negative. This implies that the right hand side of Eq. 1.9 will vanish at least as fast as a^{-1} ; the behaviour of $a(t)$ for $t \gg$ will therefore be completely determined by the value of the curvature constant: for $K = -1$ the expansion proceeds indefinitely; for $K = 0$ the velocity of the expansion tends to zero asymptotically; for $K = 1$ the expansion eventually ceases and is followed by a contraction back to a singular state (see Fig. 1.2).

The final fate of the Universe in the Friedmann models is therefore inseparable from its spatial geometry. We will now see how these two properties relate to the amount of matter that fills the Universe.

If we set $K = 0$ in Eq. 1.9 and solve for the density we obtain:

$$\rho = \rho_{cr} \equiv \frac{3H^2}{8\pi G}, \quad (1.15)$$

where ρ_{cr} is the *critical density*, or the density of a flat model at every value of t . The density parameter Ω is defined as

$$\Omega \equiv \frac{\rho}{\rho_{cr}}. \quad (1.16)$$

If we now rewrite Eq. 1.9 in terms of Ω , we obtain

$$H^2(\Omega - 1) = \frac{Kc^2}{a^2}. \quad (1.17)$$

We can therefore conclude that the Universe will have a negative, null or positive curvature and will evolve accordingly if, at any value of the cosmic time, Ω results to be smaller, equal, or larger than 1, respectively.

The geometrical properties of the Universe and the details of its evolution are therefore determined by how the density of its content relates to the critical density at any time.

Another feature of the Friedmann models is in that they possess a *particle horizon*, $R_H(t)$. This is defined for every single observer and encompasses the points that have been causally connected to it up to some time t . The particle horizon has a finite size in these models, meaning that each observer is in causal contact with a finite portion of the Universe.

1.3 Observational constraints

So far we have followed a purely theoretical derivation of the properties of the Universe, starting from few fundamental assumptions like the Cosmological Principle, the field equations of General Relativity with the energy-momentum tensor of a perfect fluid and a form for the equation of state. We have derived models describing the evolution of the cosmological expansion for different values of the curvature parameter K and related this to the density parameter Ω describing the amount of the matter-energy content of the Universe. In this section we will rapidly outline the fundamental observational evidence that has allowed to discriminate among these models and select those most likely to constitute the correct description of our Universe.

The existence of the *Cosmic Microwave Background* radiation (CMB) was discovered in 1965 (Penzias and Wilson, 1965) and since then has been considered one of the major supporting evidence of the Big Bang model. This radiation fills the observable universe and has a thermal spectrum with uniform temperature $T \sim 2.7$ K. The origin of this background radiation is traced back to the event of *recombination*, when the Universe was around 300 000 years old ($z \approx 1100$) and was characterised by a low enough temperature (≈ 4000 K) for Thomson scattering to weaken and for the

formation of neutral atoms to set in³. The striking uniformity of the temperature distribution represents, by itself, the most convincing validation of the Cosmological Principle. Besides this, CMB studies also state that the Universe is flat (de Bernardis et al., 2000; Stompor et al., 2001; Spergel et al., 2007). This result implies $\Omega = 1$, in contradiction with solid observational evidences stating that baryonic, ordinary matter accounts for up to 4% of the critical density (Geiss and Reeves, 1972; Rogerson and York, 1973; Yang et al., 1984). The existence of non-baryonic, *dark* matter, is quite well established nowadays and is a fundamental element in the theory of structure formation we are going to revise in the following section. However, the combination of baryonic and dark matter does not change the situation in terms of matching the condition $\Omega = 1$, since $\Omega_{matter} \equiv \Omega_b + \Omega_{DM} \sim 0.3$ (White et al., 1993; Verde et al., 2002). The agreement with observations is recovered once some sort of *dark energy* is allowed to join the matter component, supplying the 70% of missing density. This seems to be an *ad hoc* solution to the problem, but we need to stress that there is, indeed, a deeper motivation for the introduction of the dark energy term: the observed acceleration of the cosmological expansion. In the last decade a big effort was made in the observation of type Ia supernovae, standard candles of known luminosity and - therefore - known distance. The result of these studies was that far away objects look fainter than they should be, given their redshift, and this is being interpreted as a sign of an accelerating expansion (Riess et al., 1998; Perlmutter et al., 1999). We recall that, in the simple Friedmann models discussed above, this cannot occur, as the expansion is always decelerated, at least for $-1/3 < w < 1$. The easiest way to account for an accelerated expansion is the introduction of a so-called *cosmological constant* (or Λ) with equation of state $P = -\rho c^2$; this term has to be dominant over all the other components of the cosmic fluid at the present time, but may have not been so in the past: its density in fact shows no evolution with the scale factor a , unlike that of matter and radiation (see Sec. 1.2).

To summarise, observations favour a flat universe whose content is strongly dominated by a dark energy component and whose behaviour is well approximated by a Friedmann model as long as matter or radiation dominate the energy budget. Fig. 1.3 gives a pictorial view of the constraints introduced by the aforementioned observation.

We will now proceed further in describing the theoretical framework which is most likely to account for the observational properties of the Universe; we will finally investigate a central issue we have overlooked so far, namely the undeniable evidence that the distribution of matter in the nearby Universe is far from being homogeneous and is instead organised into structures.

³The complete *decoupling* between radiation and matter occurs when the timescale for Thomson scattering exceeds the Hubble time; this occurs at later times, when the Universe is few million years old ($z \approx 300$).

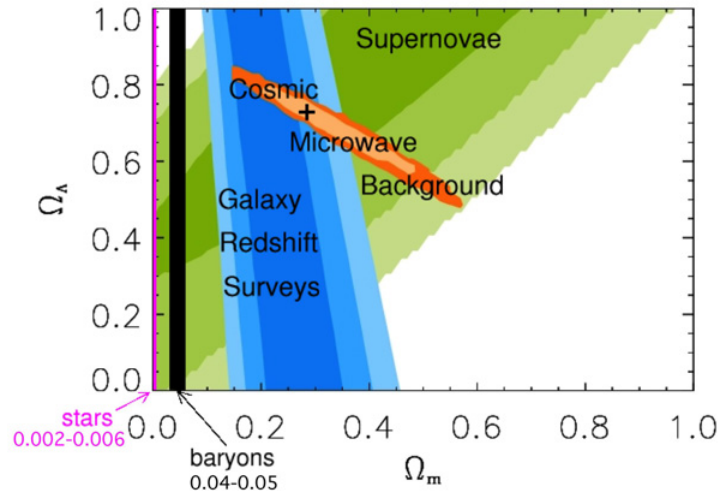


Figure 1.3: Observational constraints on the cosmological parameters Ω_{matter} , Ω_{Λ} and Ω_b updated to 2004 (Figure from Gondolo (2004)). Three types of observations - supernova measurements of the recent expansion history of the Universe, CMB measurements of the degree of spatial flatness, and measurements of the amount of matter in galaxy structures - provide predictions that overlap in a region around $\Omega_{matter} \approx 0.3$, $\Omega_{\Lambda} \approx 0.7$. Constraints on the baryon density in the Universe come from the analysis of the CMB spectrum and from studies on primordial nucleosynthesis; these agree on $\Omega_b \lesssim 0.05$ (black, vertical band). The thin, pink line marks the contribution to the total energy density provided by luminous matter.

1.4 Evolution of cosmic structures

Galaxies are regarded as the fundamental building blocks of the Universe, the constituent particles of the cosmic fluid, and it is with respect to these that we evaluate the characteristics of the distribution of matter on large scales. It is found that on scales smaller than ~ 100 Mpc galaxies are not distributed randomly in space, but rather show strong clustering in the form of filaments and sheets with large voids in between (visible also in Fig. 1.1); on larger scales we recover the expected uniform distribution, while the results on CMB temperature fluctuations confirm that the Universe was extremely close to being homogeneous and isotropic $\approx 300\,000$ yrs after the Big Bang. The question arises as to how the Universe evolved from these early stages, where fluctuations in its radiation and baryonic matter components were of the order of one part in 10^5 , up to the present day, where the distribution of matter is highly clustered. This is the issue we will address in this section, starting from the theoretical foundations of linear theory to the development of sophisticated tools to account for the late stages of structure formation.

We will assume that primordial density fluctuations existed in the cosmic fluid. The problem of their origin is beyond the purpose of this review and constitutes one of the motivations that lead to the development and establishment of inflationary theories (starting from Guth, 1981). Here we just report the predicted form of the power

spectrum describing these perturbations:

$$P(k) = Ak^n, \quad (1.18)$$

where both the normalisation A and the spectral index n are (or are related to) fundamental cosmological parameters (see Table 1.1).

1.4.1 Jeans Instability and the Linear Theory

The physical mechanism which is believed to drive the formation of cosmic structures is gravitational instability. The theoretical framework describing this process is that of Jeans instability, developed at the beginning of the last century as an attempt to explain the formation of stars and planets from an initially smoother cloud of gas (Jeans, 1902). According to this criterion, whether or not a fluctuation will undergo collapse depends on its length scale and on the *Jeans length*, λ_J , of the fluid. A small, positive spherical overdensity of radius λ , sitting in a background fluid of mean density ρ , will collapse if the following condition is met:

$$\lambda > \lambda_J \simeq \frac{v_s}{\sqrt{G\rho}}, \quad (1.19)$$

where v_s is the sound speed of the fluid particles. An equivalent statement is that the system is unstable on a scale λ if the free-fall time, the timescale of unimpeded collapse, is smaller than the hydrodynamical time, the time needed for a sound wave to traverse the system (and therefore to balance the inhomogeneities).

We will now introduce the fundamentals of this analysis for a collisional fluid in a static background, which was the case Jeans studied himself, and later comment on the effect of the expansion of the Universe and the presence, in the cosmic fluid, of collisionless matter components.

The set of equations that rule the evolution of a perfect fluid are:

- **Continuity equation**, for the conservation of mass:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0; \quad (1.20)$$

- **Euler equation**, for the conservation of momentum:

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v} \cdot \nabla) \mathbf{v} = -\frac{1}{\rho} \nabla P - \nabla \Phi; \quad (1.21)$$

- **Poisson equation**, the classical field equation:

$$\nabla^2 \Phi = 4\pi G \rho; \quad (1.22)$$

- **Equation of state**, relating pressure, density and entropy:

$$P = P(\rho, S); \quad (1.23)$$

• **Entropy evolution:**

$$\frac{dS}{dt} \tag{1.24}$$

These are five linear differential equations of first and second order in space and time, describing the evolution of density ρ , velocity \mathbf{v} , entropy S , pressure P and gravitational potential Φ of the fluid.

There is observational evidence⁴ that primordial fluctuations were of *adiabatic* type, involving matter and radiation density components such that $\delta S = 0$: in this case the last two equations reduce to $P = P(\rho)$ and $S = \text{const.}$ We want to use these equations to study how the fluid responds to small perturbations around the equilibrium configuration, which in this case represents a static, homogeneous and isotropic universe. The induced perturbations obey:

$$\delta = \delta(\mathbf{r}, t) \equiv \frac{x(\mathbf{r}, t) - x_0}{x_0} = \frac{\delta x(\mathbf{r}, t)}{x_0} \ll 1 \tag{1.25}$$

where x_0 stands for the equilibrium value of any of the quantities ρ , \mathbf{v} , p and Φ . The next step is to substitute this perturbed solution into the fluid equations (Eq. 1.20 to 1.24) and *linearise* the results, i.e. neglect second-order term and subtract the unperturbed solution. We look for solutions in the form of plane waves:

$$\delta f(\mathbf{r}, t) = \delta f_k \exp(i\mathbf{k} \cdot \mathbf{r} + i\omega t), \tag{1.26}$$

where f stands for the generic quantity which is being perturbed. This choice is equivalent to Fourier-transform the problem and characterise the perturbations in terms of their wavenumber $k = 2\pi/\lambda$. The perturbation in configuration space will result in a superposition of such plane waves, evolving independently from one another, at least as long as the linear regime holds. The set of differential equations becomes an algebraic one in Fourier space and reduces to one final expression, summarising the behaviour of the entire set:

$$\omega^2 = k^2 v_s^2 - 4\pi G\rho, \tag{1.27}$$

termed *dispersion relation*.

When the scale of the fluctuation is small ($k \gg$), Eq. 1.27 reduces to $\omega^2 \simeq k^2 v_s^2 \geq 0$ and the perturbation behaves like a wave of constant amplitude, propagating with phase velocity ranging from 0 to v_s . For perturbations on larger scales the value of ω^2 can become negative, inducing a temporal dependence of the solution 1.26 of the type $\exp(\pm\omega t)$, corresponding to a stationary wave with exponentially growing or decaying amplitude. The presence of the growing solution implies that the system is unstable under the effect of its self-gravity. The physical scale above which gravitational instability takes over is again the Jeans scale, λ_J :

$$\lambda_J = \left(\frac{\pi v_s^2}{G\rho} \right)^{1/2}. \tag{1.28}$$

⁴From the study of the angular power spectrum of temperature fluctuations present in the CMB.

1.4.2 Application to the cosmological case

We have just seen how a perturbation on a scale $\lambda > \lambda_J$ grows exponentially above a homogeneous and static collisional fluid. We expect a modification of this behaviour in the case of an expanding background and also need to assure that the analysis can be extended to collisionless fluids, for which a pressure term cannot be defined.

A full treatment of the linear evolution of perturbations in an expanding universe requires General Relativity, at least to account for scales larger than the cosmological horizon R_H and for the existence of relativistic species (e.g. radiation), where the newtonian analysis breaks down. The equivalent of the dispersion relation for the case of a matter-dominated, flat universe is:

$$\ddot{\delta}_k + 2H\dot{\delta}_k + \delta_k[v_s^2 k^2 - 4\pi G\rho] = 0. \quad (1.29)$$

As we can easily figure out from the sign of the second term, the cosmological expansion acts against the collapse, damping the exponential growth and reducing it to a power law. This result also holds for a non-baryonic, dark matter fluid, provided we specify what is opposing gravity in this case; since dark matter is collisionless, it does not exert an opposing pressure but instead its velocity dispersion σ_v can be responsible for the dissipation of the perturbation. A relativistic analysis gives us the equivalent of Eq. 1.29 for a radiation-dominated, flat universe:

$$\ddot{\delta}_k + 2H\dot{\delta}_k + \delta_k \left[v_s^2 k^2 - \frac{32}{3}\pi G\rho \right] = 0. \quad (1.30)$$

An insight into this expression will soon reveal that as long as radiation dominates the energy content of the Universe, the Jeans scale for collapse will be larger than the horizon, resulting into the absence of gravitational instability. In this regime: perturbations in baryonic matter will oscillate along with those in the radiation component, coupled through Thomson scattering; fluctuations in dark matter will suffer *stagnation* (Meszaros, 1974) as their free-fall time is larger than the characteristic time of expansion, or Hubble time.

Without going into the details of the calculation, we will now summarise the overall trend of the perturbations in the three components δ_R (for radiation), δ_B (for baryonic matter) and δ_{DM} (for dark matter) of the cosmic fluid over the relevant eras and scales:

- **Radiation era** ($t < t_{eq}$), where t_{eq} is the time of the *equivalence* between the matter and radiation ($t_{eq} \approx 50000 \text{ yrs}$; $z_{eq} \approx 3600$); during this epoch the main contribution to the energy density of the Universe came from the radiative component ($\rho \approx \rho_R$):

Outside the horizon the fluctuations in any components are gravitationally tightened to the dominant one, which sets the trend:

$$\delta_R \propto \delta_{DM} \propto \delta_B \propto a^2;$$

Inside the horizon we recall that, for different reasons, perturbations undergo no substantial evolution:

$$\begin{aligned}\delta_R \propto \delta_B \text{ oscillates,} \\ \delta_{DM} \propto \text{const;}\end{aligned}$$

- **Matter era, between equivalence and decoupling** ($t_{eq} < t < t_{dec}$), when dark matter started dominating the energy density of the Universe ($\rho \approx \rho_{DM}$) and baryonic matter was still coupled to radiation:

Outside the horizon the behaviour is common and set by the dominant component:

$$\begin{aligned}\delta_{DM} \propto \delta_R \propto \delta_B \propto a \quad (\Omega = 1), \\ < a \quad (\Omega < 1), \\ > a \quad (\Omega > 1),\end{aligned}$$

where the symbols \propto do not have to be taken literally and they are just meant to suggest the different rapidity of the growth.

Inside the horizon dark matter fluctuations grow, while the baryon-radiation fluid keeps oscillating:

$$\begin{aligned}\delta_R \propto \delta_B \text{ oscillates,} \\ \delta_{DM} \propto a \quad (\Omega = 1), \\ < a \quad (\Omega < 1), \\ > a \quad (\Omega > 1);\end{aligned}$$

- **Matter era, after decoupling** ($t > t_{dec}$), when radiation and ordinary matter started evolving separately:

Outside the horizon the behaviour is as in the previous case;

Inside the horizon perturbations in the radiation and dark matter components keep behaving as in the previous era, whereas those in the baryonic components undergo an accelerated evolution due to gravitational attraction of pre-existent dark matter structures:

$$\delta_B = \delta_{DM}(1 - a_{dec}/a).$$

The first structures to form originate from dark matter perturbations and are referred to as dark matter *halos*. Fluctuations in any component are frozen until matter-radiation equality is reached and only then those in dark matter start growing. Once baryonic matter and radiation decouple, protogalaxies start forming from the infall of gas into the deep potential wells of the halos.

1.4.3 Hot vs. cold dark matter

The presence of some form of non-baryonic matter in the Universe was suggested already in 1933, to account for the large velocity dispersion of galaxies within the Coma Cluster (Zwicky, 1933); later on, another strong piece of evidence in favour of its existence progressively stood out, namely the shape of the rotation curve for spiral galaxies (Babcock, 1939; Bosma, 1978; Rubin et al., 1980). From a cosmological point of view, the presence of dark matter is needed to account for the existence at $z = 0$ (today) of structures corresponding to $\delta \gg 1$, fact that would be unexplainable if fluctuations were allowed to grow from an amplitude of order 10^{-5} at $z \sim 1100$, as stated by CMB observations (Smoot et al., 1992). The prior evolution of dark matter perturbations and the following catch-up by the baryonic component reconciles the scenario with this fundamental observational evidence.

The problem of the composition of dark matter was addressed since its introduction and it is still an open question. Depending on the mass of the constituent particles, we discriminate between *hot* and *cold* dark matter (HDM and CDM, respectively): the impact on the physics of the pre-equivalence era is so different in these two cases, that it brings to opposite scenarios for the evolution of structures.

Hot dark matter decouples from the rest of the cosmic fluid when still relativistic, the opposite holding for cold dark matter. “Derelativisation” occurs when the temperature of the Universe equals the rest energy of the particle:

$$kT(a_{\text{NR},X}) \simeq m_X c^2. \quad (1.31)$$

A small value of m_X implies late derelativisation, possibly after the decoupling of dark matter, while heavy particles derelativise much earlier. The definition of hot/cold dark matter therefore reduces to a difference in the mass of their constituent particles ($\sim \text{eV}$ vs. $\sim \text{GeV}$, respectively). This has strong repercussion on the value of the Jeans scale at the equivalence, when fluctuations start growing. The Jeans mass (mass within a spherical perturbation of radius λ_J) reaches the following values, for the two cases:

$$\begin{aligned} M_{\text{J,HDM}}(a_{\text{eq}}) &\simeq 10^{12} - 10^{14} M_{\odot}, \\ M_{\text{J,CDM}}(a_{\text{eq}}) &\simeq 10^5 - 10^6 M_{\odot}; \end{aligned}$$

these represent the minimum mass scale of collapse⁵. If the primordial spectrum of perturbation (Eq. 1.18) has an amplitude decreasing with k , fluctuations on the smallest possible physical scales will be the most likely to collapse and form structure first; in the HDM case these will be fluctuations on scales of ~ 7 orders of magnitude larger than in the CDM case, resulting in a *anti-hierarchical* scenario for the formation of structures. Conversely, cosmological models based on the CDM assumption provide *hierarchical* structure formation, where small structures form first.

We point out that, in reality, fluctuations on scales smaller than the Jeans length at

⁵Hereafter, masses are expressed in solar masses (M_{\odot} ; $1 M_{\odot} = 1.98892 \times 10^{33} \text{g}$)

the equivalence do not propagate as sound waves and are instead completely obliterated by dissipation phenomena (Silk, 1968). A similar process acts on dark matter perturbations and causes the dissipation of inhomogeneities on scales smaller than the *free streaming* length, i.e. the distance dark matter particles can travel until they become non-relativistic. The different values of the Jeans mass for hot and cold dark matter can be interpreted as a consequences of the different efficiency of dissipation in the two cases.

We have seen that perturbations are subject to alterations in their growth once they enter the horizon so that it is natural to expect a deformation of their primordial spectrum up to the time of matter-radiation equivalence. As long as they are outside the horizon, fluctuations grow regardless of their scale; as the size of the horizon increases, perturbations on smaller scales will either undergo stagnation or oscillate and will have their growth suppressed until equivalence. In addition to this, dissipation phenomena will act smoothing out inhomogeneities below the Jeans scale. We will therefore expect the power spectrum to lose its original shape on scales larger than that of the cosmological horizon at the matter-radiation equivalence. This deformation is accounted for by the introduction of the so-called *Transfer function*, $T(k)$, which acts on the primordial spectrum as a filter suppressing the large- k power. In the case of HDM the suppression will be drastic and $T(k)$ will approach zero exponentially for large k (due to a combination of stagnation and free streaming); in the alternative CDM scenario, significant power survives dissipation and the action of $T(k)$ will cause a softer bend in the power-spectrum ($P(k) \propto k^{-3}$, due essentially to stagnation). Examples of these effects are shown in Fig. 1.4. Several power spectra are plotted, each corresponding to a different cosmological scenario. The curves show the dependence $\propto k$ for small values of k , down to the scale of the horizon at the equivalence (corresponding to the peak in the curves); beyond this scale power is suppressed, particularly in the HDM case (red curve).

There is evidence that clusters of galaxies, the largest virialised structures in the Universe, are also the youngest objects we observe and that therefore the evolution of structures follows a bottom-up, hierarchical, trend. This allows to discard the hypothesis of hot dark matter, in favour of the CDM model (Blumenthal et al., 1984).

We note that in recent years and in order to solve the problems that the CDM scenario has in reproducing observation at small scales, growing interest has developed in the study of *warm* dark matter models ($m_\chi c^2 \sim \text{keV}$).

1.4.4 Beyond linear theory

Linear theory breaks down at $\delta \sim 1$. The structures we observe today are characterised by a mean density which is several hundreds above the mean density of the Universe and to follow the history of their formation we need a formalism valid in more advanced stages of the gravitational collapse. The problem is a complex one and we will see that analytical solutions are difficult to achieve without strong simplifying assumptions. Even then, the results may fail to provide detailed predictions to be tested against observations; this is something that will only be




Figure 1.4: The matter power spectrum for a range of models of structure formation. In blue is the result for a CDM model with $\Omega_{matter} = 1$, $n = 1$ and $h = 0.5$. The power spectrum of a HDM model with $m_{\chi}c^2 = 22 \text{ eV}$ is displayed in red. Finally, the results for a Λ CDM model characterised by $\Omega_{\Lambda} = 0.8$, $\Omega_{matter} = 0.2$ and $h = 1$ are plotted in green. All the curves have been normalised to 1 at the scale $k = 0.2 \text{ h Mpc}^{-1}$. Figure adapted from White et al. (1994).

achievable by means of numerical simulations.

The Spherical Collapse

Gunn and Gott (1972) developed a simple model to describe non-linear evolution. The assumption is that of a spherical perturbation of constant density, expanding with the background universe in such a way that the initial peculiar velocity at the edge is zero. These simplifications allow a special treatment of the problem, which turns out to be very powerful: the perturbation is regarded as a spherical universe embedded in a flat background universe. The sphere evolves like a Friedmann model whose initial density parameter $\Omega_p(t_i)$ is given by:

$$\Omega_p(t_i) = \Omega(t_i) (1 + \delta_i), \quad (1.32)$$

where Ω is the density parameter of the background universe and δ_i ($\ll 1$) is the density contrast of the perturbation. A structure will form if, at some time t_m , the spherical region ceases to expand and begins to collapse; this will always be the case, at some point, in a matter universe with a flat geometry (the background model assumed in the original derivation). For the embedded spherical universe,

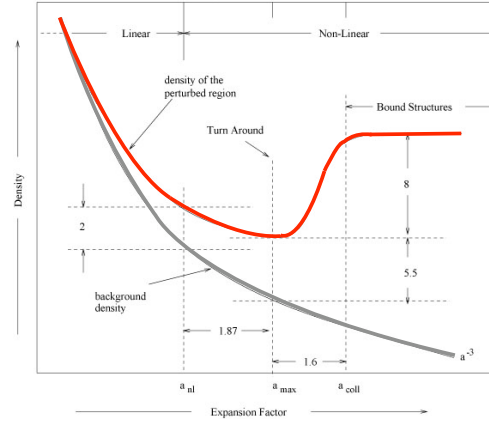


Figure 1.5: Growth of density in an over-dense region, according to the Spherical Collapse Model. The gray curve shows the evolution of the background density in a matter-dominated universe. The red curve represents the density of the spherical inhomogeneity. In the linear region, the growth is $\propto a$; in the non-linear regime, the spherical region collapses faster, virialises and forms a bound structure. Adapted from Fig. 8.1 of Padmanabhan (1993).

one can easily derive the value of the density at the moment of maximum expansion and compare it with that of the background universe, obtaining:

$$\chi = \frac{\rho_p(t_m)}{\rho(t_m)} \simeq 5.6 \Rightarrow \delta(t_m) \simeq 4.6, \quad (1.33)$$

which is a factor of four above the prediction of linear theory. If the contribution of pressure keeps being neglected we should conclude that at a time $t = 2t_m$ the spherical overdensity collapses into a point of infinite density. What happens in reality is that the collapse will proceed and form an extended structure of radius $R \simeq R_m/2$, which will reach the virial equilibrium at $t_{vir} \simeq 3t_m$ (see Fig. 1.5). At these points, the value of the density contrast will have grown to ~ 180 and ~ 400 , respectively, while an extrapolation from linear theory would have estimated it around ~ 1.68 and ~ 2.20 .

The starting assumption of this approach is of course extremely restraining: the idea of a spherical, homogeneous perturbation is itself unrealistic and even if that were the case at the beginning, the subsequent evolution of the collapse would disrupt such a symmetrical configuration and lead to anything but the formation of a spherical structure. The only limit where this analysis could be valid is that of small fluctuations (i.e. just above the Jeans scale), where the effect of pressure cannot be neglected and the overall balance of gravitational and pressure forces can effectively lead to more or less spherical proto-objects.

The Zel'dovich Approximation

If one wants to study the evolution of large fluctuations, the spherical collapse model is of no relevance and a different tool has to be used. The Zel'dovich approach (Zel'Dovich, 1970) gives an approximate solution for the growth of fluctuations when the effect of pressure can be neglected and matter can be regarded as dust, which is the case when the scale of the perturbation is large.

The answer regarding the evolution of perturbations is given in terms of displacement from the initial, *lagrangian* coordinate \mathbf{q} of a set of uniformly distributed particles. The actual, *Eulerian* position \mathbf{r} is a function of both \mathbf{q} and time and is assumed to be given by:

$$\mathbf{r} = a(t) \mathbf{q} + F(\mathbf{q}, t) = a(t) \mathbf{q} - a(t) b(t) \nabla_{\mathbf{q}} \Phi(\mathbf{q}), \quad (1.34)$$

where the first term describes the cosmological expansion, the second accounts for the perturbation and the underlying assumption is that the temporal and spatial dependences can be separated. The term $b(t)$ is known as *linear growth factor* and measures the relative growth of perturbations at different times as predicted by linear theory. The link to the initial perturbation is through the potential Φ , which is in turn related to δ by Poisson equation. The gradient of the perturbed potential identifies the direction along which particles start moving; since the potential is not updated, particles will keep moving of uniform motion until they will eventually cross each other generating a region of infinite density known as *caustic*. Let us see this in more detail. Eq. 1.34 defines a map between the coordinate \mathbf{r} and \mathbf{q} which is unique as long as trajectories do not cross; conservation of mass requires:

$$\rho(\mathbf{r}, t) d^3r = \langle \rho(t_i) \rangle d^3q, \quad (1.35)$$

where t_i sets the beginning of the approximation, when particles are still uniformly distributed. The term d^3r/d^3q is the determinant of the Jacobian of the mapping between \mathbf{q} and \mathbf{r} ($J(\mathbf{r}, t) = \partial \mathbf{r} / \partial \mathbf{q}$) and can easily be derived from Eq. 1.34. The matrix J is symmetric and can be locally diagonalised. As a result, the above equation can be rewritten as:

$$\rho(\mathbf{r}, t) = \frac{\langle \rho(t_i) \rangle}{a(t)^3} \prod_{i=1}^3 [1 + b(t) \alpha_i(\mathbf{q})]^{-1}, \quad (1.36)$$

where α_i are the eigenvalues of the matrix $\frac{\partial \Phi}{\partial q_i \partial q_j}$. This equation indicates that at some time t_{sc} one or more of the terms in the product can vanish and the density becomes infinite; this event is referred to as *shell crossing* and occurs when $b(t_{sc}) = -1/\alpha_j$. This condition corresponds to the situation where points with different lagrangian coordinates end up having the same Eulerian coordinate: the mapping between \mathbf{r} and \mathbf{q} is not unique anymore as trajectories cross each other. For collapse to occur, at least one of the three values of α_j must be negative: if more than one is so, collapse will first occur along the axis characterised by the most negative one. In general, one expects collapse to occur along one axis and form a 2D sheet-like structure (pancake); only in rare cases will the process involve

WMAP-7+BAO+ H_0 for Λ CDM		
H_0	70.2	Hubble constant ($\text{km s}^{-1} \text{Mpc}^{-1}$)
Ω_b	0.0455	Baryonic matter density
Ω_{DM}	0.227	Dark matter density
Ω_Λ	0.728	Dark energy density
n	0.961	Primordial spectral index
σ_8	0.807	Variance on a scale of $8 h^{-1} \text{Mpc}$

Table 1.1: Fiducial cosmological parameters from the last release of the WMAP mission (Komatsu et al., 2011), combined with BAO (Baryonic Acoustic Oscillation) constraints and local measurements of the Hubble constant.

more than one axis giving rise to filamentary or spherical structures.

The Zel'dovich approximation accounts very well for the evolution of density perturbations and provides the same results as full N-body calculations, at least until the moment of shell crossing, where it suddenly breaks down: particles are blind to each other and the structures dissolve as they pass through one another. Another limit of the approximation is that, being purely kinematic, it does not account for close-range forces and therefore is inadequate to describe anything beyond a mildly non-linear regime.

The only way to follow the formation of structures at the late stages of their evolution is to rely on numerical simulations. These only can provide detailed predictions regarding the distribution of matter in different cosmological models and at different times, to be tested against observations. The importance of the numerical approach in cosmological studies and the continuous development of sophisticated techniques to simulate the evolution of the cosmic fluid will be the topic of the next chapter.

1.5 The Standard Model - Summary

We have seen how under the assumptions of the Cosmological Principle and the field equations of General Relativity one can derive a powerful theoretical framework for the study of the properties of the Universe. Observational evidence stating the flatness of the spatial geometry and the acceleration of the cosmic expansion further constrain the model, and suggest the existence of a dark energy and dark matter component in the cosmic fluid (see Fig. 1.3). The observation of the characteristics of clustering allows to discriminate between the two main classes of candidates for the dark matter component in favour of the cold dark matter one. These elements are at the basis of the so-called Λ CDM model, which has established itself as the standard framework for our understanding of the Universe. The current cosmological parameters describing this model are listed in Table 1.1.

As a final note, we stress that this is not the end of the story and that ques-

tions have been left unanswered, most noticeably regarding the nature of dark energy and dark matter, as well as the origin of primordial fluctuations which seeded structure formation. In addition to this, the model also suffers from intrinsic weaknesses, partially alleviated by the introduction of inflationary theories (for an overview, see Liddle and Lyth, 2000; Liddle, 2002).

Numerical techniques

The formation and evolution of structures in a cosmological context is a problem of high complexity that cannot be entirely described by analytical techniques. Numerical simulations are the primary tool employed in these studies and considerable effort has been invested in developing ever more refined techniques since the precursory works of the 1970s. In this chapter we revise the basics of this field and outline the state-of-the-art algorithms used in the different stages of a simulation, from the set-up of the initial conditions to the post-processing of the results. More emphasis will be put on aspects which are relevant for the work presented in the rest of this thesis, whose outline we introduce at the end of the chapter.

2.1 Introduction

As pointed out in Sec. 1.4.4, the evolution of cosmic structures beyond the linear regime is addressed by analytical techniques only under a prohibitive amount of simplifying assumptions. The results are of little use if one aims at constraining cosmological models by comparing their prediction for the distribution of matter in the Universe to existing observations. This task is accomplished by numerical simulations; these follow the evolution of the matter component from the very early stages, where small density fluctuations perturb an otherwise homogeneous distribution, to the current time, characterised by the presence of highly non-linear structures (at least on small scales; see discussion at the beginning of Sec. 1.4). The features introduced by the specific cosmological model under consideration are imprinted on the initial conditions of the simulation, i.e. on the primordial perturbation field. The system is then evolved forward in time according to the forces acting on the simulated cosmic fluid and to the evolution of the background universe. At the end of the simulation one is left with a prediction for the large-scale distribution of matter within the assumed cosmological model; this could in principle be tested against existing observations in order to assess the validity of the model or to constrain its parameters.

In the standard cosmological picture, roughly 80% of matter in the Universe consists of weakly-interacting, “dark” matter particles and only the remaining 20% of

standard, baryonic matter. The former interact only gravitationally, whereas the latter is subject also to small-scale, hydrodynamic forces and to numerous astrophysical processes. Given that what we observe in the Universe is the baryonic, luminous constituent of the cosmic structures, modelling this secondary component is of paramount importance to the purpose of comparing the simulation results to actual data. However, this represents a rather difficult task and no consensus has been achieved yet on the best approach to the problem. The general trend so far has been to perform simulations of the dark-matter component only; the link between the predictions on the properties of these dark structures and the consequences on the distribution of luminous matter is not at all obvious and is itself an object of study.

Modelling the behaviour of dark matter means following the evolution of a collisionless fluid subject to the action of gravity only. The term collisionless refers to the negligible importance of two-body encounters in determining the evolution of the fluid. The governing equation of such a system is the *collisionless Boltzmann equation*:

$$\frac{df}{dt} = \frac{\partial f}{\partial t} + \sum_{i=1}^3 \left(v_i \frac{\partial f}{\partial x_i} - \frac{\partial \Phi}{\partial x_i} \frac{\partial f}{\partial v_i} \right) = 0, \quad (2.1)$$

where Φ stands for the gravitational potential and $f = f(\mathbf{x}, \mathbf{v}, t)$ is the *distribution function* of the system. This quantity represents the phase-space density of fluid elements having position in the small volume $d\mathbf{x}^3$, centred on \mathbf{x} , and velocity in the small range $d\mathbf{v}^3$, centred on \mathbf{v} ; the distribution function provides a full description of the state of a collisionless system at any time t (see, e.g., Binney and Tremaine, 1987; Dehnen and Read, 2011). In the general case, Eq. 2.1 is impossible to solve, being a non-linear, partial differential equation (PDE) in seven dimensions. The problem is overcome by utilising N-body techniques. The distribution function is sampled in phase-space with a finite number N of tracer particles $(\mathbf{x}_j, \mathbf{v}_j)$, with $j = 1 \dots N$. The condition $df/dt = 0$ in the Boltzmann equation implies that the phase-space density around the phase point of a given particle remains constant in time, i.e. along the trajectory of the particle. Finding the trajectories of the particles means, in fact, solving Eq. 2.1 by the method of characteristics; the PDE reduces to a set of ordinary differential equations (ODEs) ruling the evolution of the phase-space coordinates \mathbf{x} and \mathbf{v} for each of the point particles:

$$\frac{d\mathbf{x}}{dt} = \mathbf{v}; \quad \frac{d\mathbf{v}}{dt} = -\nabla\Phi. \quad (2.2)$$

In integrating these ODEs, one has to evaluate the gravitational potential Φ ; this is done out of the sampling particles by using one of the many available techniques (see Sec. 2.3.1).

The problem of modelling the collisionless, dark-matter component is therefore solved by evolving a number N of sampling particles according to Newton's law of gravitation and by solving the field equation in order to determine the underlying gravitational potential. From the overview presented in the previous chapter, one




Figure 2.1: Example of a cosmological simulation box. Shown is the evolution in the distribution of dark matter from $z = 30$, when the age of the Universe was less than 1% of the current age, to the present epoch (leftmost and rightmost boxes, respectively). Figure adapted from http://www.lsw.uni-heidelberg.de/users/mcamenzi/Week_7.html.

would expect the evolution of the cosmic fluid to be better described by General Relativity, rather than by Newtonian mechanics. It turns out, however, that the Newtonian approximation is acceptable, at least on scales smaller than the Hubble radius (c/H). The simulation of the cosmological evolution of dark matter, from the small initial perturbations to the highly clustered state observed at late times, is generally carried out within a cubic box of size L ; this scale varies according to the purpose of the simulation, but is generally large enough to encompass a representative volume of the Universe (i.e. over which the homogeneity of the matter distribution is appreciable; e.g. $L \gtrsim 100 h^{-1}$ Mpc). Periodic boundary conditions (Ewald, 1921; Hernquist et al., 1991) are applied to compensate for the finiteness of the box and for its otherwise artificial edges. To account for the large-scale dynamics of the Universe, the integration is performed in an expanding background. A visual representation of a typical cosmological simulation box is given in Fig. 2.1: as the simulation proceeds, the matter distribution evolves from a state of homogeneity to a more and more pronounced filamentary structure.

The number of particles used in a cosmological simulation is a parameter of crucial importance. Besides setting the mass resolution - namely the mass scale below which the simulation is “blind” - it also influences the “goodness” of the dynamical representation of the system; in fact, the higher the number of particles, the better they will reproduce the dynamics of a collisionless system. In general, the sampling particles used in a dark matter simulation will have masses exceeding those of the alleged dark matter particles by tens of orders of magnitudes; this results in a distorted representation of the small-scale dynamics, namely in a fictitious granularity of the gravitational potential and an induced collisionality in the particles’ behaviour. These so-called *discreteness* problems in N-body simulations need to be tackled with care; more will be said on this topic in the following sections and in Chapter 3. In general, cosmological simulations have always strived to increase

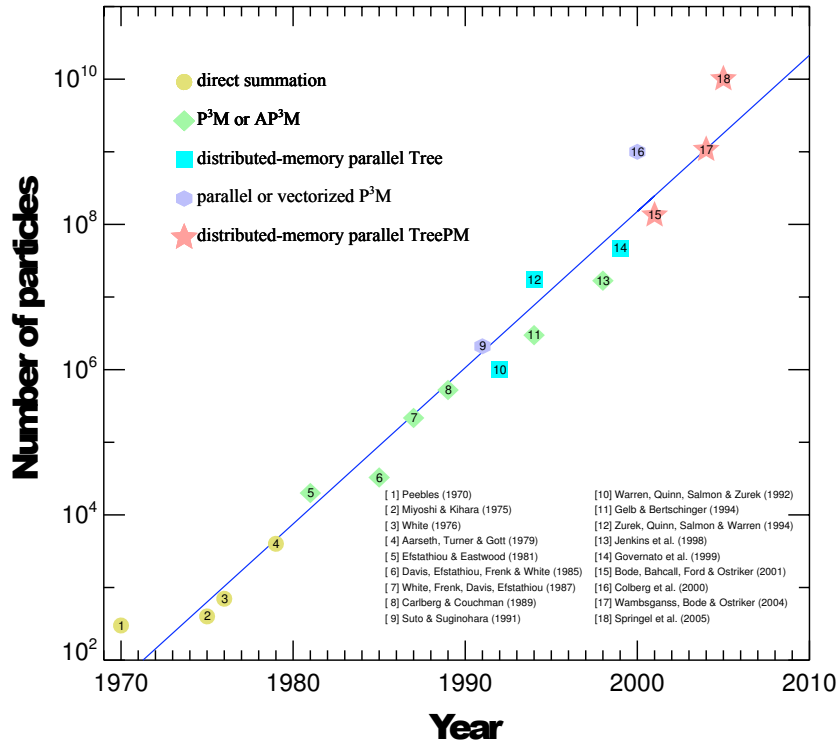


Figure 2.2: Number of particles used in high-resolution simulations of structure formation, as a function of publication date. Different symbols are used for different classes of computational algorithms (see Sec. 2.3.1). Figure adapted from Springel et al. (2005) (supplementary information).

their size and their resolution by adopting ever more efficient algorithms and sophisticated hardware. Fig. 2.2 shows how, since the early 1970s, the number of particles in a simulation has increased exponentially. The largest cosmological, N-body simulation ever performed is the Millennium XXL¹ (not shown in Fig. 2.2). This simulation follows 6720^3 dark matter particles in a cosmological box of side $L = 4.1$ Gpc; each of the sampling particles carries a mass of $\approx 5 \cdot 10^9 M_{\odot}$ (equivalent to that of a small galaxy). An example of a simulation where high resolution was prioritised with respect to the box-size is the Millennium II² (Boylan-Kolchin et al., 2009); this follows 2160^3 particles in a box of side $L = 100 h^{-1}$ Mpc, resulting in a mass resolution of $\approx 7 \cdot 10^6 h^{-1} M_{\odot}$. Already in the mid-1980s, simulations employing as little as 32^3 and 60^3 particles (points 6 and 7 in Fig. 2.2) have been crucial in the establishment of the cold-dark-matter model (CDM; see Sec. 1.4.3) as the best explanation of the available observational results on the distribution of matter in the Universe (Davis et al., 1985a; White et al., 1987; Frenk et al., 1988).

In the following sections we will outline the most important concepts at the basis

¹<http://galformod.mpa-garching.mpg.de/dev/mrobs/mxxl.html>

²<http://www.mpa-garching.mpg.de/galform/millennium-II/>

of the numerical modelling of cosmological systems: the generation of the initial conditions (Sec. 2.2), the force computation and time integration (Sec. 2.3), and the isolation of bound structures in the simulation outputs (Sec. 2.4). In parallel to this, we will also discuss issues related to the modelling of ordinary, luminous matter. In Sec. 2.5 we will report on an alternative and recently developed technique to account for the evolution of baryonic matter. Finally, in Sec. 2.6 we will briefly summarise and introduce the following chapters of the thesis.

2.2 Initial conditions

The purpose of initial conditions is to provide a discrete representation of the primordial density field predicted by the cosmological model under consideration. They serve as the starting point of the actual cosmological simulation and specific codes have been developed to generate them. The problem can be split into two main steps, namely: the generation of an unperturbed distribution of particles (also known as pre-initial conditions) and the computation of a displacement field on an appropriate grid. The initially unperturbed configuration will eventually be displaced according to this field by means of some sort of interpolation technique. To set down a finite number of particles in a suitably uniform distribution is not an easy task. The problem is to avoid introducing spurious features in the subsequent evolution of the perturbations. In this sense, the initial distribution must be as neutral as possible. Randomly placing the particles within the computational box is not a viable method, due to the unwanted introduction of “white-noise”, defined in terms of a spectrum $P(k) \propto k^n$, with $n = 0$; this will eventually lead to the formation of spurious structures even if no cosmological fluctuations are imposed on the system. The other extreme is to set down a regular distribution of equally-spaced particles; this is a somewhat better method, but still the configuration leaves an imprint on the subsequent computation in terms of strongly preferred directions at all scales and due to the artificial introduction of a fixed, minimum spacing. The best solution to the problem is to adopt a glass-like configuration (White, 1994), where the force on each particle is approximately zero. This is obtained by evolving a set of randomly distributed particles with the sign of acceleration reversed, so that their mutual gravitational interaction becomes repulsive; as the resulting configuration shows no preferred direction and no artificial perturbations, this technique has soon become the standard approach to the problem. Fig. 2.3 gives a visual representation of the three techniques used to generate pre-initial conditions.

Once the unperturbed particle distribution is set, one needs to generate the desired linear fluctuation field that will eventually be superimposed on it. This is done by means of the Zel’dovich approximation (Sec. 1.4.4). As seen in Sec. 1.4, the predictions of linear theory for the matter distribution at early times are usually expressed in terms of the power spectrum of density fluctuations, $P(k)$; this can be expressed as the product of a primordial contribution and a transfer function, the latter accounting for the modifications introduced by dissipation processes in the pre-equivalence era. A realisation of this spectrum is generated at each point of a




Figure 2.3: Two-dimensional representation of the pre-initial particle distribution generated by the three methods outlined in the text. Figure from the Ph.D. thesis of Craig Booth (<http://www.craigbooth.com/publications>).

suitable computational grid³. The results are then scaled at the desired redshift by means of the linear growth factor $b(t)$; this initial redshift is chosen such that the fluctuation field is well within the linear regime (to give an idea, a typical starting redshift may lie in the range [60, 120]). The next step is to derive the displacement field; this is defined at every grid point and will eventually be applied to the glass, pre-initial particle distribution in order to superimpose the primordial density fluctuations. Let us write Eq. 1.34 in comoving coordinates (i.e. divide by the scale factor) and derive the displacement $\mathbf{s}(\mathbf{q}, t)$ as well as the final peculiar velocity \mathbf{v} :

$$\mathbf{r} = \mathbf{q} + \mathbf{s}(\mathbf{q}, t) = \mathbf{q} - b(t) \nabla_{\mathbf{q}} \Phi(\mathbf{q}), \quad (2.3)$$

$$\mathbf{v} = -\dot{b}(t) \nabla_{\mathbf{q}} \Phi(\mathbf{q}) = \frac{\dot{b}(t)}{b(t)} \mathbf{s}(\mathbf{q}, t). \quad (2.4)$$

The relation between δ and \mathbf{s} in linear theory is

$$\mathbf{s}(\mathbf{q}, t) = i \frac{\delta}{k^2} \mathbf{k}, \quad (2.5)$$

and therefore the displacement and peculiar velocity fields are straightforwardly obtained from the generated density fluctuations. At this point, there is one more thing to be done, namely to move the glass particles according to this field. This is

³If, as usually assumed, the statistics of the primordial field is Gaussian, one could just perform a random-phase realisation of the spectrum, regardless of the specific statistics assumed for the amplitudes of the modes: the Central Limit Theorem would guarantee the resulting superposition in real space to be normally distributed. However, this is strictly true only if the simulated box approaches infinite size and, as a consequence, the Fourier spectrum is continuous. For finite models, in addition to the random-phase requirement, the amplitudes need to be drawn from a Rayleigh distribution with variance given by the expected power at each scale; only under this additional requirement will the density field be Gaussian at each point in real space.

done in real space by means of some interpolation procedure (e.g. Cloud-in-Cell).

At the end of this procedure, one is left with the phase-space information (i.e. position and velocity) of N computational particles. These carry information about the initial power spectrum and the cosmological parameters of the assumed model. In the next section we will see how this configuration is evolved from the initial redshift up to the present day.

2.3 Simulations

Given initial positions and velocities, the aim of the simulation is to evolve the particle system according to the relevant interactions. Gravity is the main driver of structure formation in the Universe and both dark and baryonic matter are subject to its action. In Sec. 2.3.1 we will deal with the most widespread algorithms to evaluate this fundamental interaction. The techniques used in the simulation of the baryonic component are outlined in Sec. 2.3.2. The discretisation of time and the methods used to integrate positions and velocities are discussed in Sec. 2.3.3. Finally, in Sec. 2.3.4 we will give some examples of relevant cosmological simulation codes.

2.3.1 Force computation: gravity

The most accurate way to evaluate the gravitational force acting on particle k is *direct summation*:

$$\mathbf{F}_k = \sum_{j \neq k} G m_j \frac{\mathbf{r}_j - \mathbf{r}_k}{|\mathbf{r}_j - \mathbf{r}_k|^3}, \quad (2.6)$$

where the sum is on all the other particles in the system. However, this is not a viable way due to its prohibitive computational cost ($\approx N^2$ operations per timestep). Other techniques have been developed that allow a much faster evaluation of the gravitational interaction while maintaining an acceptable accuracy.

Hierarchical tree

This method reduces to direct summation for the contribution of nearby particles, but involves a great simplification in the evaluation of long-range interactions. This result in a reduction in the number of operations per timestep from $\approx N^2$ to $\approx N \log N$.

In one of the most popular versions of the algorithm (Barnes and Hut, 1986), the computational domain is recursively partitioned in a sequence of cubes. Starting from the simulation box (the root node), each cube is split into eight child nodes of half the side length each; the procedure is repeated on each child node until one ends up with a cube containing one single particle (the leaf node). Once this *hierarchical tree* is built, the force computation on particle k proceeds starting from the

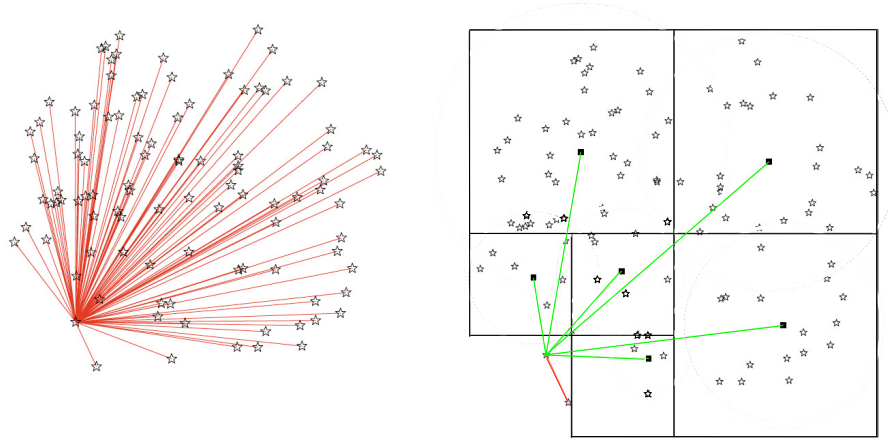


Figure 2.4: Computation of the gravitational force acting on one of 100 particles (marked as stars) in two dimensions. The left plot shows the case for direct summation; each red line represents one evaluation of the force. The right plot shows the hierarchical-tree method; each green line represents a particle-node interaction, whereas the red line a direct particle-particle interaction. Figure adapted from Dehnen and Read (2011).

root node, down to the branches and - in some cases - to the leaves. The following criterion decides how deep into the hierarchy it is necessary to go:

$$r > \frac{l}{\theta}, \quad (2.7)$$

where r is the particle-node distance, l is the size of the node and θ is the *opening angle*, a parameter controlling the accuracy of the force computation. If the inequality is satisfied, the node is considered distant enough for the following simplification to apply: the contribution of all the member particles to the gravitational force acting on particle k is reduced to the computation of one single force provided by a particle of mass equal to the sum of these particles and positioned in their centre of mass. Clearly, this method provides only an approximation to the true force, better or worse depending on the choice of the opening angle. In the limit $\theta = 0$, the method reduces to direct summation.

Fig. 2.4 gives a visual idea of the different effort required by direct and tree methods to compute the gravitational interaction on one target particle.

Particle-Mesh (PM) techniques

Another way to obtain the evolution of the system under the effect of gravity is to solve the Poisson equation:

$$\nabla^2\Phi = 4\pi G\rho. \quad (2.8)$$

This approach is at the basis of PM methods (Hockney and Eastwood, 1981). The first step is to compute the density ρ out of the position of the particles. This is done at a set of regularly spaced points - the nodes of a mesh. The process of estimating the density at each of the grid points is called *mass assignment*: the mass of each particle is assigned to one or more nearby nodes, via a number of possible techniques (nearest-grid-point - NGP, cloud-in-cell - CIC, triangular-shaped-cloud - TSC, etc.). Once the density information is available, Eq. 2.8 is solved to obtain the value of the gravitational potential at each grid point. This is done in Fourier space using the Fast Fourier Transform technique (FFT; Cooley and Tukey 1965); the transformed density field is multiplied by Green's function for the potential ($-4\pi G/k^2$) and the result is then inverse-transformed in real space. The resulting gravitational potential is differentiated to obtain the corresponding force at each grid point. These are finally interpolated to the particle position using the same technique as for the mass assignment. Note that the Fourier approach implicitly assumes periodic boundary conditions for the computational domain, exactly as desired for cosmological simulations of structure formation.

The PM method comes with a computational cost comparable to that of the tree ($\approx N\log N$). It provides an exact evaluation of the gravitational force, modulo inaccuracies introduced by the interpolation techniques. Its main shortcoming lies in the resolution limit set by the grid size; in some cases this cannot be reduced down to the desired scale without incurring memory problems.

The hybrid, TreePM method

A very popular technique used in state-of-the-art cosmological simulation codes is the so-called "TreePM" (Xu, 1995; Bode et al., 2000; Bagla, 2002); as the name suggests, it consists of a hybrid between the tree and PM methods. The scheme is implemented such that small-range interactions are evaluated by means of a hierarchical tree, leaving the long-range part to the PM algorithm. This allows:

- a fast and accurate evaluation of the gravitational interaction between distant particles and the automatic inclusion of periodic boundaries;
- a greater flexibility in the desired accuracy for the evaluated force on small scales.

A final comment on the resolution of the force computation in N-body codes. By definition, the gravitational potential of a collisionless system is smooth - i.e. the dynamics is not influenced by particle-particle encounters. By representing the

system with a number of particles infinitely smaller than its actual building blocks, we are boosting its collisionality. There is nothing that can be done to limit this problem, other than increasing the number of particles. The enhanced collisionality on small scales can have additional, tedious effects; a close approach between two of the sampling particles can lead to divergences in the evaluation of Eq. 2.6. The traditional way to solve this problem is to *soften* the gravitational interaction. This means evaluating, instead of Eq. 2.6, an expression of the kind:

$$\mathbf{F}_k = \sum_{j \neq k} \frac{Gm_j}{(|\mathbf{r}_j - \mathbf{r}_k|^2 + \epsilon^2)^{3/2}} \frac{\mathbf{r}_j - \mathbf{r}_k}{|\mathbf{r}_j - \mathbf{r}_k|}, \quad (2.9)$$

where ϵ is called the *softening length* and sets the scale below which the gravitational force deviates from the classical form of Eq. 2.6. The effect of this softening procedure is to smoothen the potential on scales smaller than ϵ ; the spurious effects due to particle-particle collisions are attenuated and less computational time is needed to integrate the encounters (see Sec. 2.3.3). Gravitational force softening is the subject of the next two chapters and more will be said on the topic. Note that an implicit form of softening is present in the PM method too, even though Eq. 2.9 is not computed directly. This is given by the presence of a finite grid spacing, below which the evaluation of the force is not reliable.

2.3.2 Force computation: hydrodynamics

The cosmological evolution of dark matter is rather well understood; the use of a specific gravity-solver or the other does not alter an overall consistent picture. The same is not true when it comes to modelling the baryonic component of the cosmic fluid. The advantages and disadvantages of the specific method adopted will have an impact on the final results. The situation becomes even worse when considering that a number of astrophysical processes (e.g. gas cooling, star formation and related feedback) need to be accounted for by means of “sub-grid” prescriptions, as this physics occurs on scales too small to be followed by the simulation and has to be introduced “by hand” with ad-hoc recipes. As most of the time the astrophysical process itself is poorly understood, there exist several competing approaches to its modelling, each leading to different results. In the following, we will first discuss the two most important philosophies at the basis of the hydrodynamic modelling of baryonic matter in cosmological simulations and then briefly outline the treatment of the additional astrophysical processes.

The equations ruling the behaviour of a baryonic, perfect fluid are:

- **Continuity equation**

$$\frac{d\rho}{dt} + \rho \nabla \cdot \mathbf{v} = \frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0; \quad (2.10)$$

- **Euler equation**

$$\frac{d\mathbf{v}}{dt} = \frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v} \cdot \nabla) \mathbf{v} = -\frac{1}{\rho} \nabla P - \nabla \Phi; \quad (2.11)$$

- **Energy equation**

$$\frac{du}{dt} = \frac{\partial u}{\partial t} + \mathbf{v} \cdot \nabla u = -\frac{P}{\rho} \nabla \cdot \mathbf{v}. \quad (2.12)$$

This last equation is just an expression for the first law of thermodynamics in the adiabatic limit. The system is closed by the field equation for Φ , accounting for the gravitational term in the equation of motion 2.11, and by the equation of state

$$P = A\rho^\gamma, \quad (2.13)$$

where A is a constant related to the specific entropy of gas and γ is the polytropic index.

These equations can be solved either at the position of the sampling particles or on the nodes of a grid. In the first case we speak of *Lagrangian* codes and in the second of *Eulerian* codes.

Smoothed Particle Hydrodynamics (SPH)

The scene of Lagrangian codes is dominated by the use of SPH algorithms (Lucy 1977; Gingold and Monaghan 1977; Monaghan and Lattanzio 1985. For a review see Monaghan 1992; Rosswog 2009). At the heart of this method is an interpolation technique which allows the density ρ to be defined for each of the sampling particles.

The integral interpolant of any function $A(\mathbf{r})$ is defined by

$$A_{II}(\mathbf{r}) = \int A(\mathbf{r}') W(|\mathbf{r} - \mathbf{r}'|, h) d\mathbf{r}'^3, \quad (2.14)$$

where the integral is performed over the entire space and the contribution at each point \mathbf{r}' is calibrated by the value of the interpolating kernel at $r = |\mathbf{r} - \mathbf{r}'|$. Although its form can vary, the main properties of the kernel W are set by:

$$\int W(|\mathbf{r} - \mathbf{r}'|, h) d\mathbf{r}'^3 = 1 \quad (2.15)$$

and

$$W(|\mathbf{r} - \mathbf{r}'|, h) \xrightarrow{h \rightarrow 0} \delta(\mathbf{r} - \mathbf{r}'), \quad (2.16)$$

where h regulates the spatial extent of the function.

More interesting, from a numerical point of view, is the discretised version of Eq. 2.14. Once the integral interpolant is expressed as

$$A_{II}(\mathbf{r}) = \int \frac{A(\mathbf{r}')}{\rho(\mathbf{r}')} W(|\mathbf{r} - \mathbf{r}'|, h) \rho(\mathbf{r}') d\mathbf{r}'^3, \quad (2.17)$$

it is easy to replace the integral by a sum over a set of interpolation points (the particles), whose masses originate from the $\rho(\mathbf{r}') d\mathbf{r}'^3$ term. The result is:

$$A_{SI}(\mathbf{r}) = \sum_b m_b \frac{A_b}{\rho_b} W(|\mathbf{r} - \mathbf{r}_b|, h), \quad (2.18)$$

where SI stands for summation interpolant, the subscript b denotes the value of any quantity associated to particle b and the summation is performed over all the particles. The density ρ_a associated to particle a is therefore given by ⁴:

$$\rho_a = \sum_b m_b W(|\mathbf{r}_a - \mathbf{r}_b|, h). \quad (2.19)$$

The key point is that, once the kernel W is chosen to be differentiable, it is possible to construct a differentiable interpolant of a function from its values at a set of points. For example, the expression for ∇A_{SI} would simply reduce to:

$$\nabla A_{SI}(\mathbf{r}) = \sum_b m_b \frac{A_b}{\rho_b} \nabla W(|\mathbf{r} - \mathbf{r}_b|, h), \quad (2.20)$$

where ∇W is known analytically. The discretised version of the hydrodynamic equations can be differentiated easily without the need of a grid; the only derivatives present, once the quantities are expressed in terms of kernel interpolation, will involve W . As an example, here is the typical discretised form of the Euler equation, as used in SPH:

$$\frac{d\mathbf{v}_a}{dt} = - \sum_b m_b \left(\frac{P_a}{\rho_a^2} + \frac{P_b}{\rho_b^2} \right) \nabla W(|\mathbf{r}_a - \mathbf{r}_b|, h), \quad (2.21)$$

where the gravitational part has been omitted⁵.

The main advantages of SPH are its conservation properties and spatial adaptivity. Physical quantities like mass, energy and momentum are conserved, by construction, along the flow. The resolution naturally increases in overdense regions, something that makes the method particularly targeted for modelling the highly inhomogeneous matter fields in cosmological simulations. The weaknesses of SPH lie in the inaccurate treatment of shocks and in the suppression of fluid instabilities at contact discontinuities.

Eulerian methods

Instead of defining the gas properties at the location of the sampling particles, Eulerian codes use the vertices of a grid. The evolution of the system is obtained by solving the *Riemann problem* at the boundaries between each grid cell, by means of schemes such as the one proposed by Godunov (1959).

In order to overcome the problem of limited spatial resolution, state-of-the-art Eulerian codes make use of Adaptive Mesh Refinement algorithms (AMR; Berger and Colella 1989); the grid is recursively refined in regions where higher resolution is needed and left coarser in the rest of the domain. This allows grid codes to compete

⁴In principle, the density can also be calculated via integration of a discretised version of Eq. 2.10. However, the summation form given by Eq. 2.19 is numerically more robust. If particle masses are kept fixed, mass conservation is guaranteed and there is no need to solve the continuity equation.

⁵The expression in Eq. 2.21 is a symmetrised version of the discretised Euler equation. A straightforward discretisation of Eq. 2.11 does not conserve linear and angular momentum.

in resolution with Lagrangian codes, but results in poor conservation of fundamental physical quantities. In general, the Eulerian approach to hydrodynamics results in an accurate treatment of shocks and discontinuities. Its application to the cosmological problem is hampered by the lack of Galilean invariance (making the results sensitive to the presence of bulk motions) and by the difficulties in the treatment of gravitational instabilities.

In order to combine the advantages of the Eulerian and Lagrangian methods, a hybrid approach has been recently proposed. We will briefly discuss this in Sec. 2.3.4.

Sub-resolution processes

An accurate description of all the astrophysical processes shaping the observable properties of the Universe is impossible to obtain at once; it would mean a simultaneous modelling of scales ranging from the interior of a star to the largest cosmic structures - differing by tens of orders of magnitudes. Yet, neglecting processes like gas cooling, star formation and supernova feedback (among others) would mean a rather inadequate modelling of the baryonic component. The solution so far has been to provide the codes with specific modules accounting for additional physics by means of pre-defined recipes; this obviously constitutes a deviation from the otherwise self-consistent treatment of the evolution of the fluid. Modelled astrophysical processes generally include:

- **Radiative cooling** - assuming a specific composition for the baryonic matter and a range of possible reactions (collisional excitation and ionisation, recombination, bremsstrahlung, etc.; see, e.g., Katz et al. 1996) cooling rates are computed that bring the gas temperature to a low enough level for it to condense at the centre of dark matter halos, where it will eventually become Jeans-unstable and form stars;
- **Star formation** - rapidly cooling, Jeans-unstable gas is progressively converted into star particles (one star particle containing, in fact, a large ensemble of stars), which are then evolved under the effect of gravity as an additional collisionless component (see, e.g., Cen and Ostriker 1992);
- **Stellar feedback** - the impact that the formation and evolution of stars has on the environment is modelled via prescriptions for galactic winds, metal enrichment and supernova explosions and its main effects are to heat and change the compositions of the surrounding gas (see, e.g., Springel and Hernquist 2003);
- **SMBH feedback** - the presence of super massive black holes (SMBHs) at the centre of galaxies is modelled via prescription for their growth and the feedback associated with their accretion of gas (see, e.g., Di Matteo et al. 2005).




Figure 2.5: Schematic representation of the leapfrog-integration method. Figure taken from the web (<http://www.drexel.edu/physics/>)

Comparison projects aimed at comparing the performances of different codes have been carried out. When concentrating on purely hydrodynamical simulations (i.e. no sub-grid physics; Frenk et al. 1999), these confirmed a very good agreement on the predictions for the dark matter component in collapsed objects, while pointing out some disagreement in properties like gas temperature, gas mass fraction (within about 10%) and X-ray luminosity (within a factor of 2); also, even when a good match was obtained in overdense regions, results for the temperature, pressure and entropy of the gas were found to significantly differ in underdense environments (O’Shea et al., 2005). More recently, Scannapieco et al. (2011) undertook a comparison project aimed at testing the performance of nine different codes on a simulation of galaxy formation; besides varying in the philosophy adopted to tackle hydrodynamics, the codes also differed in the implemented sub-grid physics. Major differences were found in the results; the authors concluded that although the Eulerian or Lagrangian nature of each code did contribute the final discrepancies, the fluctuations introduced by different formulations of the star-formation and feedback schemes are likely to be the fundamental source of the differences seen.

2.3.3 Time integration

Another crucial aspect of a cosmological simulation is the choice of timesteps and the time-integration method. A standard solution is the use of the *leapfrog* scheme, a second-order accurate method described by the following set of equations:

$$\mathbf{v}_{i+\frac{1}{2}} = \mathbf{v}_{i-\frac{1}{2}} + \mathbf{a}_i \Delta t, \quad (2.22)$$

$$\mathbf{r}_{i+1} = \mathbf{r}_i + \mathbf{v}_{i+\frac{1}{2}} \Delta t. \quad (2.23)$$

Positions and velocities are evaluated at different points in time (see illustration in Fig. 2.5), but only one force evaluation per timestep is carried out. The advantage of this scheme lies in its symplectic nature: energy is conserved during the integration. This is evident when integrating the Kepler problem (Fig.2.6): even though a small precession of the orbit may be registered, no changes in the energy occur (the semi-major axis remains unchanged).

There exist several criteria to set Δt . Due to the large dynamic range encoun-

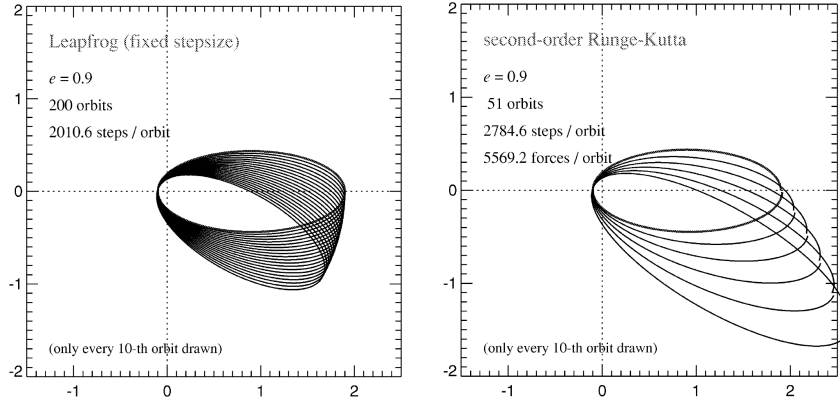


Figure 2.6: Integration of a Keplerian orbit with eccentricity $e = 0.9$. The plot on the left shows the performance of the leapfrog scheme; on the right, as a reference, the results obtained with the Runge-Kutta method at the same order. Figure adapted from Springel (2005).

tered in a cosmological simulation, adopting a timestep fixed in space and time is not feasible. Following the dynamics in an overdense region needs a more accurate integration than in the rest of the box; a timestep small enough to model these regions, besides being not necessarily known a priori, would result in an unsustainable slow-down of the whole computation. For these reasons, the common choice is to adopt variable, individual timestep; this means Δt can vary in space and time according to the evolution of the simulation⁶. A commonly used criterion sets the timestep for each particle according to its force softening ϵ and the modulus of the gravitational acceleration a :

$$\Delta t = \sqrt{\frac{2\eta\epsilon}{a}}, \quad (2.24)$$

where η is an accuracy parameter. The timestep becomes shorter in regions of higher accelerations; the dependence on ϵ underlines the reduced accuracy needed in the integration of close encounters when softening of the gravitational force is introduced in the computation. In the case of baryonic matter, the chosen timestep has to obey the so-called Courant-Friedrichs-Lewy (CFL) condition:

$$\Delta t < \Delta t_{CFL} = \frac{\ell C_{CFL}}{c_s}, \quad (2.25)$$

⁶When the timestep is allowed to vary in space and time, the leapfrog integration will no longer be symplectic. However, for a collisionless system the evolution can still be expected to reach comparable accuracy to a symplectic integration (see discussion in Springel 2005).

where C_{CFL} is a parameter of order unity, ℓ is the size of the resolution element in the simulation (softening or grid cell) and c_s is the local sound speed. The idea behind the CFL condition is that the time resolution should not exceed the time taken by the local sound speed to traverse a spatial resolution element.

2.3.4 Codes

Among the most popular Lagrangian codes, GADGET⁷ (Galaxies with Dark matter and Gas intERACT; Springel et al. 2001b; Springel 2005) is certainly worth mentioning. The most challenging, dark-matter simulations of the last decade have been performed with this code. Among these are the Millennium XXL and Millennium II, mentioned in the Introduction. Besides this, GADGET is also the relevant code for all the work presented in this thesis. In its default operative mode, GADGET employs a TreePM method for the evaluation of the gravitational part of the force and SPH for solving the hydrodynamical contribution. The time integration is performed with a leapfrog scheme, while allowing for individual timesteps.

An example of Eulerian code is provided by ENZO⁸ (Bryan and Norman, 1997). The gravitational part is solved by means of a PM algorithm aided by the TSC interpolation technique. Hydrodynamics is based on the piecewise parabolic method (PPM) of Woodward and Colella (1984) - a higher-order extension of Godunov's method. The underlying mesh is adaptive according to a number of possible refinement criteria. Both GADGET and ENZO are fully parallelised using the MPI message-passing library and their basic versions are publicly available.

Finally, very recently, a code has been developed that belongs to a somewhat intermediate category and could represent the future in the field of cosmological simulations. Based on a moving, unstructured mesh, AREPO (Springel, 2010) aims at combining the advantages of the Eulerian and Lagrangian approaches. At the core of the method is the construction of a moving mesh, defined by the Voronoi tessellation of a set of discrete points moving with the flow. The mesh is used to solve the hydrodynamic equations in a way similar to standard Eulerian codes. The gravitational part, instead, is dealt with by a TreePM algorithm no different from the one implemented in GADGET. Unlike ordinary Eulerian codes, AREPO is fully Galilean-invariant - something which is highly desirable in an astrophysical context, where the presence of bulk motions is common. In addition, the new scheme can adjust its spatial resolution automatically and continuously, and hence inherits the principal advantage of SPH for simulations of cosmological structure growth. At the same time, the high accuracy of Eulerian methods in the treatment of shocks is also retained, while the treatment of contact discontinuities is improved.

⁷<http://www.mpa-garching.mpg.de/gadget>

⁸<http://lca.ucsd.edu/software/enzo/>

2.4 Structure finders

Based also on my contribution to

A. Knebe et al.

MNRAS, 2011, 415, 2293

At the end of a simulation, one is left with the final phase-space information for the N sampling particles. How should one make use of this data? The information contained in a cosmological simulation can be compared to observations of our own Universe only in a statistical sense. The initial distribution of particles represents just one realisation of the Gaussian, primordial perturbation field and adopting a different seed for the random-number generator would result in a completely different appearance of the structures within the final computational box. What does not change from one realisation to the other are the statistical properties of the final particle distribution.

Useful statistics routinely extracted from simulations are the power spectrum of the matter distribution, along with its equivalent in real space - the *two-point correlation function*. These quantities convey information on the clustering properties of matter and trace the increase in small-scale power due to the hierarchical assembly of structures (see the example in Fig. 2.7).

In this section we will focus on a crucial step of the standard post-processing of cosmological simulations: the identification of collapsed structures. These range from isolated, dark-matter halos to smaller sub-structures embedded within them. Studying how matter organises itself into individual objects, both in terms of global properties of these objects (e.g. mass) and their internal features, can provide a powerful way to compare the outcome of the simulation to observational results. Structure-finding is performed by numerical codes specifically developed to this purpose. The first codes were assembled in the early 1970s, when the field of numerical simulations was taking its first steps. Although visually recognising structures and patterns, in general, may be a relatively easy task, developing targeted algorithms is more problematic. The definition of what a structure is, along with the determination of its centre and extent, are all subject to some level of arbitrariness, on top of being affected by the graininess of the particle representation of the system. The importance of obtaining ever more accurate predictions is reflected by the outstanding increase in the number of different codes that have been developed in the last two decades (see Fig. 2.8) - 29 known to date. Even though they differ in the adopted algorithms, most of these codes are based on either the spherical-overdensity method (SO; Press and Schechter 1974) or on the friends-of-friends technique (FOF; Davis et al. 1985b). The family of codes based on the former philosophy can be regarded as “density-peak locators”, while those based on the latter as “particle collectors”. In the first case, the idea is to identify peaks in the matter density field and to grow spheres of increasing radius around them. When the enclosed density falls below a pre-defined threshold (generally set by the spherical collapse model - e.g. 180 times the critical density), a spherical structure




Figure 2.7: The dimensionless power spectrum ($\Delta^2(k) = P(k)d^3k$) of the dark matter distribution in a $500 h^{-1}$ Mpc box (Millennium Simulation; (Springel, 2005)). In gray are the predictions of linear theory for the five redshifts under consideration, while in blue are the results of the simulation. Note the deviation from linear theory, advancing to larger and larger scales (small k) as the simulation proceeds. Error bars are shown for the $z = 0$ case only and at large scales (they become negligible for larger values of k). Figure from Springel (2005) (supplementary information).

is identified. Most of the codes adopting this method differ just in the way they locate density peaks. In the second case, particles close to each other in either configuration space (3D) or phase-space (6D) are connected and the resulting assembly is regarded as an individual object.

In both cases, this initial phase is followed by a procedure aiming at discarding gravitationally-unbound particles from the identified structures.

As shown by Knebe et al. (2011), most codes provide comparable results when applied to a given cosmological simulation. The number of objects as a function of mass, their spatial position and bulk velocities, as recovered by different halo finders, show an impressive agreement. More complicated is the case of substructures; identifying overdensities embedded within a larger halo has proven to be a more difficult task, whose outcome shows some level of dependency on the adopted technique. Discrepancies on the recovered mass can rise up to 50%, with some codes showing also a strong dependence of their results on the radial position of the object within the density field of the host.

In what follows we describe in greater detail the FOF-based code SUBFIND (Springel et al. 2001a). Being the companion code of GADGET, SUBFIND has been used to process the results of the largest and most influential simulations ever

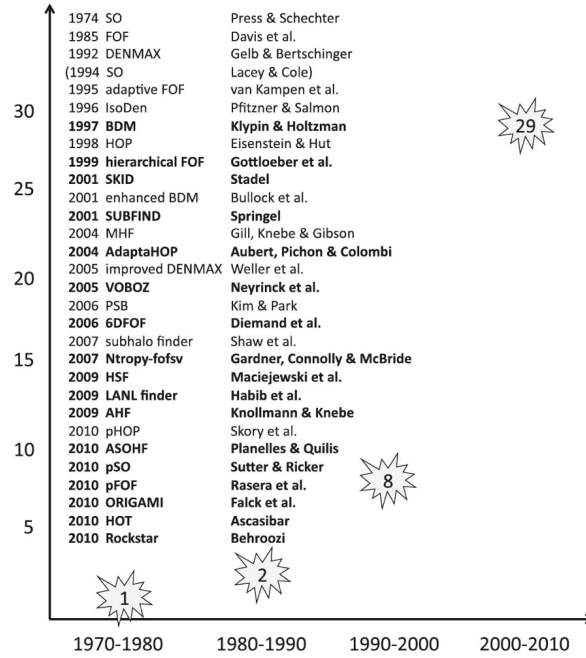


Figure 2.8: Cumulative number of (known) halo finders as a function of time. Figure from Knebe et al. (2011).

performed. Besides this, FOF and SUBFIND are the halo and subhalo finder used throughout the rest of this thesis.

2.4.1 The FOF method and the substructure finder SUBFIND

Given an arbitrary distribution of particles, the FOF algorithm acts linking particles separated by a spatial distance $r < \ell d$, where d stands for the mean inter-particle separation in the box and ℓ is an adjustable parameter - the *linking length*. The usually adopted value $\ell = 0.2$ corresponds to a lower limit of 125 for the density contrast of the resulting structure. In the 3D version of the algorithm, only the spatial geometry of the particle distribution is at the basis of the grouping procedure; there exist 6D extensions, though, that include an additional proximity condition in phase-space (Diemand et al., 2006). A lower limit is generally set for the number of particles belonging to an FOF structure; a usual choice is to set $N_{min(Halo)} = 20$. The centre of the structure is generally defined as the position of the member particle characterised by the lowest potential energy. Once this is known, spherical-overdensity quantities like the virial mass and radius, defined for a pre-defined enclosed overdensity (generally ≈ 200) with respect to the background or ρ_{crit} , can also be easily computed.

At the end of the procedure, one is left with a catalogue of all the FOF halos identified in the computational box. A typical object will look like the one in the top-left panel of Fig. 2.9; an irregularly-shaped assembly of particles characterised by strong internal inhomogeneities. The isolation of these *substructures* is the main task of subhalo finders.




Figure 2.9: Example of a SUBFIND-detected halo. The top-left panel shows the parent FOF halo. The top-right panel presents the smooth, “host” component (accounting for around 90% of the mass), while the isolated substructures are shown in the bottom-left. Plotted in the last panel are the unbound particles discarded from the group. Figure from Springel et al. (2001a).

SUBFIND identifies gravitationally bound, locally overdense regions within an input parent halo, traditionally provided by the FOF group finder. The densities are estimated based on the initial set of particles via adaptive kernel interpolation, based on a number N_{dens} of smoothing neighbours (see Eq. 2.19). Local overdensities are identified through a topological approach that searches for saddle points in the isodensity contours within the global field of the halo. As shown in Fig. 2.10, this is done in a top-down fashion, starting from the particle with the highest associated density and adding particles characterised by progressively lower densities. If a particle only has denser neighbours already belonging to a certain structure, it gets added to the same group. If, instead, the particle represents a local maximum in density, a new structure is grown around it. Finally, if the particle has denser neighbours from two different structures, an isodensity contour that traverses a saddle point is identified. In this case, the two involved structures are joined and registered as candidate subhaloes (if they contain at least $N_{min(Sub)}$ particles). These candidates, selected according to the spatial distribution of particles only, are later processed for gravitational self-boundedness. Particles with positive total energy are iteratively dismissed until only bound particles remain. If the remaining bound number of particles is at least $N_{min(Sub)}$, the candidate is ultimately recorded as a subhalo. The set of initial substructure candidates forms a nested hierarchy that is processed from inside out, allowing the detection of substructures within substructures. Particles not bound to any genuine substructure are assigned to the

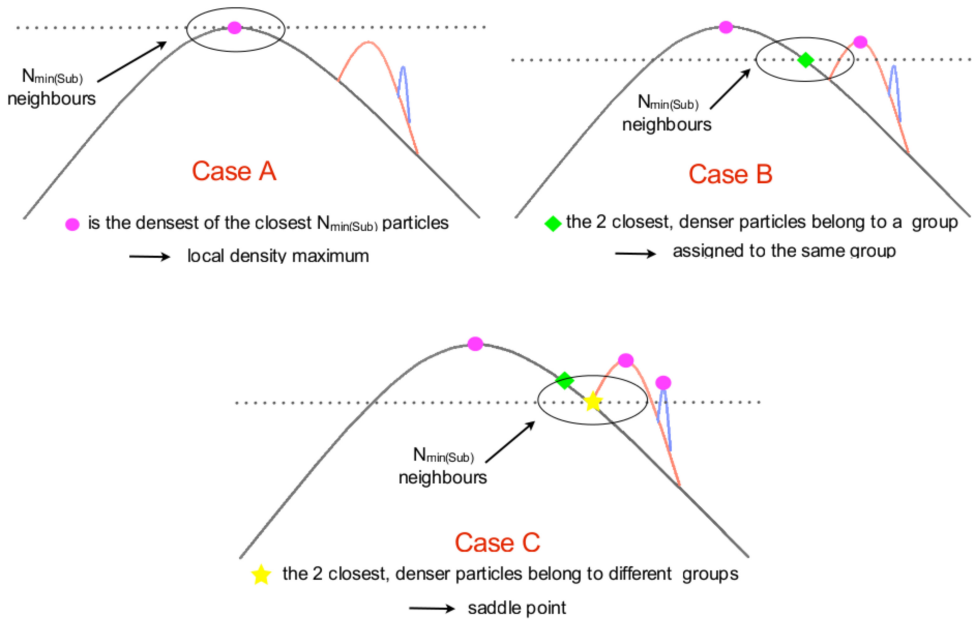


Figure 2.10: A pictorial representation of the algorithm employed by SUBFIND to identify embedded substructures (in red and blue) within a host halo (in black). The objects are represented as density peaks, whose maxima are marked by a pink circle. The search starts from the absolute peak in density (top-left figure) and proceeds to particles with a lower associated density (as for the one marked by the green diamond in the top-right figure); when a saddle is found (yellow star in the bottom figure), a new candidate substructure is identified (in this case the one marked in red).

“background halo”. This component is also checked for self-boundedness and in the end some particles may remain that are not bound to any of the structures. The top-right and bottom-left panels of Fig. 2.9 show, respectively, examples of the background halo and embedded substructures; the bottom-right panel shows, instead, discarded particles.

The results of a simulation are generally saved at a number of intermediate times during its evolution. Once the information on the halos and subhalos present at each of these so-called “snapshots” is made available, one can reconstruct the merging history leading to the formation of the final structures. These are referred to as *merger trees* (see a visual example in Fig. 2.11) and follow the evolution of collapsed objects throughout the simulation, from the time they are isolated structures, to the stage where they merge and become substructures in larger halos. As will be explained in the next section, this information is fundamental to construct models for the formation and evolution of the galaxy population associated to these dark structures.




Figure 2.11: A schematic view of the merger tree for a dark matter halo. The horizontal lines set specific points in time (from the earliest t_1 down to t_5). The circles represent dark-matter halos merging and growing to eventually give rise to the single object at t_5 . Figure from Baugh (2006).

2.5 Semi-analytic models of galaxy formation and evolution

Several times in the course of this chapter it was pointed out how difficult it is to model the baryonic component of the cosmic fluid. Yet, it is against observations of luminous structures that we have to compare the predictions of a numerical simulation. A brief account of how hydrodynamical simulations deal with the astrophysical processes related to the formation and evolution of galaxies was already given in Sec. 2.3.2. Even though the coeval evolution of dark matter and gas is followed self-consistently at early times and in large regions of the computational domain, the condensation of gas at the centre of halos, the subsequent formation of a stellar component and the impact this has on the environment are all modelled by means of pre-defined recipes.

In the last two decades, a powerful tool for the study of galaxy formation and evolution in a cosmological context has grown as a noticeable alternative to the purely numerical approach. These so-called “semi-analytic” models treat the behaviour of baryonic matter solely by means of physically-motivated, theoretical prescriptions, although the outcome of these recipes depends on the specific features of the underlying dark-matter halo (mass, internal structure, formation history).

The idea that galaxies formed from the condensation of gas within the potential well of collapsed, dark-matter halos was first proposed by White and Rees (1978). The first semi-analytic models based on this picture have been proposed by White and

Frenk (1991) and Cole (1991); these already contained prescriptions for the modelling of feedback and stellar populations, but their treatment for the formation history of the underlying dark-matter halos was still rather unrealistic. Soon after, Kauffmann et al. (1993) and Cole et al. (1994) noticeably improved this aspect by following the hierarchical assembly of halos along their merger histories; these were not yet extracted from numerical simulations and represented a Monte Carlo implementation of the Press-Schechter formalism⁹ (Press and Schechter, 1974). More recently, it has become a common practice to make use of state-of-the-art, cosmological simulations in order to obtain a detailed prediction of the formation history of the underlying dark-matter halos (Kauffmann et al., 1999; Springel et al., 2001a).

There exist a wealth of models currently used to derive predictions on the properties of the observed galaxy population. In what follows, we give an overview of their basic functioning by outlining the main features of one of the most sophisticated among them, presented in Guo et al. (2011). In this case, the backbone is provided by the merger trees obtained out of the 64 and 67 outputs from the Millennium¹⁰ (Springel, 2005) and Millennium II simulations, respectively. The former follows the evolution of dark matter in a cosmological box of side length $500 h^{-1}$ Mpc and as predicted by a standard Λ CDM model. A gas distribution mirroring that of dark matter is assigned to each of the identified objects at early times; the fraction of gas to dark-matter mass is varied according to the mass and formation time of the host halo, to account for the photoheating induced by the presence of a UV background. The evolution of this baryonic component is regulated by recipes for gas cooling, star formation, supernova feedback, black-hole growth and related effects. The efficiency of these processes is modulated by the properties of the host halo and its specific history within the simulation. This is particularly true for another set of modelled processes, namely those that act on *satellite* objects. When, during the evolution, a halo joins a larger structure (and therefore becomes embedded within it) its associated galaxy starts being referred to as a satellite. Additional processes are then switched on, that act on the gaseous and stellar component of the object. Tidal stripping is accounted for, by mirroring the effects that accretion into the new environment has on the dark matter halo (whose behaviour is known from the simulation). Ram-pressure stripping (Gunn and Gott, 1972), due to the satellite's motion through the gas distribution within the host, is also accounted for. The dynamics of mergers is followed self-consistently until the point where the dark-matter mass of the object falls below that of the associated baryonic component; this may happen due to aggressive stripping affecting the subhalo¹¹. After

⁹Under the assumption of (i) the validity of the spherical collapse model (Sec. 1.4.4) and (ii) a Gaussian statistics for the primordial perturbations, the Press-Schechter model provides predictions for the number density of objects of different masses and at different redshifts, given a cosmological model.

¹⁰<http://www.mpa-garching.mpg.de/galform/virgo/millennium/>

¹¹On top of this, the capability of halo finders to identify substructures decreases when the objects are found more and more towards the central regions of the host - as in the late stages of a merger event; generally, the recovered mass will be an under-estimate of the real mass. This is particularly true for SUBFIND (see Knebe et al. 2011).

this point, the dynamical evolution provided by the simulation is not considered adequate anymore; the subsequent positions and velocities of the merging object are traced by the most bound particle at this time, modified by a orbit-shrinking factor modelling the orbital decay induced by dynamical friction (see Sec. 4.5). The merger event is “imposed” after a certain time, again according to dynamical-friction arguments.

The model of Guo et al. (2011), as well as all the others, contains a substantial number of free parameters; these have a physical grounding and are intimately related to the details of the different implemented processes. These parameters are generally tuned to reproduce observations of the galaxy population in the Local Universe (i.e. $z \approx 0$), such as those provided by, e.g., the Sloan Digital Sky Survey (; Abazajian et al. 2003). This does not guarantee agreement at higher redshifts and indeed this is generally not the case - one of the mostly quoted weaknesses of the semi-analytic approach.

Applied to the Millennium and Millennium II simulations, the semi-analytic model of Guo et al. (2011) provides an impressive match to the abundance and large-scale clustering of redshift-zero galaxies, as a function of stellar mass and luminosity. Non-negligible discrepancies, especially in the prediction of low-mass, passive galaxies are also registered; in general, it is believed that an improved treatment of star formation should relax the tension with observational results.

2.6 Summary and outline of the thesis

The development of ever more refined numerical techniques has played a fundamental role in our current understanding of structure formation in a cosmological context. The gravitational collapse of matter cannot be followed by analytical means beyond the linear regime and this limits the possibility of obtaining strictly theoretical predictions for the distribution of matter in the Universe. This information is of primary importance in order to assess and discard different cosmological models on the basis of a thorough comparison with observational results. Detailed predictions of the properties of cosmological structures can only be provided by numerical simulations.

The first simulations were carried out in the 1970s and already by the mid-1980s they had provided the strongest evidence yet in favour of the cold-dark-matter paradigm. In more recent years, increasing effort is being invested in the modelling of the baryonic component of the cosmological, matter fluid. If the behaviour of dark matter is relatively well understood, the same is not true for ordinary matter; significant discrepancies are still registered in the results obtained by different techniques, especially when the treatment of astrophysical processes like gas cooling, star formation and feedback on the environment is also included. Semi-analytic models of galaxy formation and evolution provide a hybrid approach to the problem; they combine dynamical histories of the build-up of dark-matter structures (extracted from numerical simulations) to analytical models for the evolution of the baryonic component within them. Even though they are capable of reproduc-

ing the main properties of the observed galaxy population in the nearby Universe (i.e. at the present time), significant discrepancies are registered at earlier times, not to mention the necessity of tuning a considerable amount of free parameters to regulate the importance of the different processes in shaping galaxy properties.

In this thesis we present two works involving different aspects of numerical techniques in cosmological studies. In the first part (Chapters 3 and 4) we are concerned with the evaluation of the gravitational interaction in N-body calculations. Specifically, we are dealing with the techniques that need to be adopted in order to moderate the impact of discreteness effects associated to the Monte Carlo nature of the system. What is the best way to model a collisionless system, given a limited number of particles? How can we suppress particle noise whilst maintaining the desired resolution in the evaluation of the gravitational interaction? Adaptive gravitational softening is introduced and its implementation in the cosmological simulation code GADGET is described. The effects of this technique on simulations involving dark matter only are presented in Chapter 3; the extension to hydrodynamical simulations featuring also the presence of ordinary, baryonic matter is discussed in Chapter 4.

In the last part of the thesis (Chapter 5) we make use of state-of-the-art semi-analytic models of galaxy formation and evolution in order to investigate the link between the orbital and internal properties of galaxies (e.g. colour, star formation activity, age of the stellar population). The former is traced by the underlying cosmological simulation, while the latter is modelled by means of analytical prescriptions. When, during their evolution, galaxies stop being isolated objects and become satellites orbiting within the matter distribution of a larger system, additional processes start playing a role in the evolution of their internal properties. How does the efficiency of these processes depend on the orbit of the galaxy? Does the initial orbit of a satellite determine its late-time evolution? How do the orbital properties of satellites change with time? We tackle these problems and also account for the existing observational evidence on the subject.

Finally, our main conclusions are summarised in Chapter 6.

Adaptive gravitational softening I: the single-species case

Based on

F. Iannuzzi, K. Dolag
MNRAS, 2011, 417, 2846

Cosmological simulations of structure formation follow the collisionless evolution of dark matter starting from a nearly homogeneous field at early times down to the highly clustered configuration at redshift zero. The density field is sampled by a number of particles in number infinitely smaller than those believed to be its actual components and this limits the mass and spatial scales over which we can trust the results of a simulation. Softening of the gravitational force is introduced in collisionless simulations to limit the importance of close encounters between these particles. The scale of softening is generally fixed and chosen as a compromise between the need for high spatial resolution and the need to limit particle noise. In the scenario of cosmological simulations, where the density field evolves to a highly inhomogeneous state, this compromise results in an appropriate choice only for a certain class of objects, the others being subject to either a biased or a noisy dynamical description. We have implemented adaptive gravitational softening lengths in the cosmological simulation code GADGET; the formalism allows the softening scale to vary in space and time according to the density of the environment, at the price of modifying the equation of motion for the particles in order to be consistent with the new dependencies introduced in the system's Lagrangian. We have applied the technique to a number of test cases and to a set of cosmological simulations of structure formation. We conclude that the use of adaptive softening enhances the clustering of particles at small scales, a result visible in the amplitude of the correlation function and in the inner profile of massive objects, thereby anticipating the results expected from much higher resolution simulations.

3.1 Introduction

In collisionless N-body simulations the matter field is represented by a number of point-particles in number considerably reduced with respect to the actual building blocks of the system. These particles have no physical counterpart and should be simply regarded as Monte-Carlo samplings of the probability density distribution in position and velocity. When during the evolution two of these particles are found at small spatial separations, the gravitational attraction between them can become arbitrarily large and following the encounter properly can turn out to be prohibitively expensive from the computational point of view; in addition, this collisional behaviour is an artifact of the granularity of the mass distribution in the simulation and would not manifest in an intrinsically-smooth, collisionless system. A way around this problem, namely the discreteness of the particle representation, is obtained by smoothing the density distribution from a collection of spikes to a continuous field; to this purpose, every particle is assigned a finite volume over which its mass is spread. This translates into the gravitational interaction between the computational particles being “softened” at small separations, i.e. being attenuated and prevented to diverge. In this way close encounters are impeded and the simulation can proceed at a regular pace. The gravitational smoothing does not considerably affect the artificial two-body relaxation though, as this is driven by both close and distant encounters and the former cannot be avoided by softening techniques (Chandrasekhar, 1942; Spitzer and Hart, 1971; Hernquist and Barnes, 1990; Theis, 1998; Dehnen, 2001; Diemand et al., 2004). Suppressing relaxation is primarily a condition on the number of sampling particles and depends only weakly on the value of the softening (Power et al., 2003).

The details of this smoothing procedure can be reduced to the choice of a softening kernel $W(r, h)$, which determines the functional form of the modified density profile of the particle, and of a softening length h , which controls the spatial extent within which the modification applies. In the classic case of “Plummer” softening, the gravitational potential of each particle is that of a Plummer sphere whose scale length is given by the value of the softening h ; this results in the following form for the gravitational force $\mathbf{F}(\mathbf{r})$ between two particles of masses m_i and m_j separated by a distance r :

$$\mathbf{F}(\mathbf{r}) = -G \frac{m_i m_j}{(r^2 + h^2)^{3/2}} \mathbf{r}. \quad (3.1)$$

An evident shortcoming of this choice is that the gravitational force is modified at all separations; several works (e.g. Dyer and Ip 1993; Dehnen 2001) have shown that a better choice is to use kernels with a compact support, meaning that the interaction recovers its Newtonian, original form at separations greater than the softening length. The most commonly used one is the cubic spline of Monaghan and Lattanzio (1985):

$$W(r, h) = \frac{8}{(\pi h^3)} \begin{cases} 1 - 6q^2 + 6q^3, & 0 \leq q < 0.5 \\ 2(1 - q)^3, & 0.5 \leq q < 1 \\ 0, & 1 \leq q \end{cases} \quad (3.2)$$

where $q = r/h$, r being the distance from the centre of the kernel, i.e. from the particle's position. The particles assume the density profile given by Eq. 3.2 and the gravitational potential and force fields are modified accordingly.

It is evident that the introduction of softening leads to an unphysical modification of the gravitational interaction at interparticle separations below h ; the inverse square law is replaced by a gentler interaction, whose strength approaches zero in the limit $r \rightarrow 0$. This results in a systematic misrepresentation of the force at small scales, an effect commonly referred to as *bias* (Merritt, 1996; Dehnen, 2001). Ideally one would like to limit this effect to the smallest possible scales by reducing h accordingly; at the same time the softening lengths cannot be made arbitrarily small without running into the discreteness problem previously addressed. It is then clear how the choice of h assumes a critical importance for the reliability of the force computation and indeed a number of studies (Merritt, 1996; Romeo, 1998; Athanassoula et al., 2000; Dehnen, 2001) have been devoted to the subject.

In general there is an overall agreement that no such a thing as an “optimal” softening length exists when one has to deal with highly inhomogeneous systems, as too big a value would degrade the description of over-dense regions whilst too small a value would enhance collisionality in under-dense environments. Generally, the gravitational softening scale is set to a fixed value, indeed decided as a compromise between the desired spatial resolution and the need to moderate the noise arising from particle-particle interactions. As a matter of fact, in most cases this means fine-tuning the softening parameter to follow properly the dynamics of a certain class of objects - the densest - at a certain time; in a cosmological simulations, where the initially uniform matter density field evolves to form highly differentiated structures, this implies following poorly not only the moderately-dense regions, but also the progenitors of the “target”, densest objects.

Ideally it would be desirable to vary the softening scale in space and time according to the density evolution of the different environments: this would remove the aforementioned trade-off between spatial resolution and particle noise, thus allowing to increase the former and reduce the latter depending on the local features of the density field. The problem with a varying softening length though, as will be shown in Section 3.2, is that it introduces an additional dependence on the spatial position r in the definition of the gravitational potential; if this is not accounted for in the derivation of the equation of motion, the energetics of the system can manifest unwanted behaviours. In this context, Price and Monaghan (2007) (hereafter PM07) have proposed a formalism to adapt the gravitational softening lengths in a simulation while retaining conservation of both momentum and energy, which, as just mentioned, would be lost when adopting an adaptive scheme. In their method the softening lengths are computed in the same fashion as for the smoothing length in smoothed particle hydrodynamics (SPH - Gingold and Monaghan 1977; Lucy 1977), where h varies with the local density ρ according to $h \propto \rho^{-1/3}$. The adoption of this scheme involves a number of modifications in the overall structure of a typical cosmological simulation code, most noticeably the introduction of an additional term in the equations of motion. The technique has already been employed in simulation codes like the TreePM code of Bagla and Khandai (2009) (hereafter BK09) and the

TreeSPH EvoL code of Merlin et al. (2010). BK09 apply the method to collisionless simulations of structure formation and find an increase in clustering in over-dense regions accompanied by a somewhat loss in resolution in the under-dense ones; Merlin et al. (2010) performed standard hydrodynamical tests in order to test the overall performance of their code and found good energy-conservation properties along with an improved behaviour in tests like the isothermal collapse of Bate and Burkert (1997). Although we are here concerned about particle-based codes, where the gravitational interaction is computed at the particle's location, we note that grid-based codes employing adaptive mesh refinement (Bryan and Norman, 1997; Truelove et al., 1998; Abel et al., 2000; Knebe et al., 2001; Kravtsov et al., 2002; Teyssier, 2002; O'Shea et al., 2004; Quilis, 2004) are intrinsically spatially adaptive; these codes perform a recursive refinement of an initial regular grid, thereby increasing their resolution in regions of interest. For similar reasons as those outlined above, conservation of energy is not strictly guaranteed in these codes, as it would not be when using adaptive gravitational softening without changing the equation of motion accordingly.

In this chapter we discuss the implementation of the adaptive softening length formalism by PM07 in the cosmological simulation code GADGET (Springel et al. 2001b; Springel 2005) and the main effects that it has on dark matter simulations of the large-scale structure of the Universe. Our results qualitatively confirm those of BK09 in over-dense regions, but show no drawbacks in under-dense ones; in particular, we do not register a significant loss of low-mass structures and we considerably improve the results in term of particle clustering when comparing to standard simulations employing fixed softening.

In Sections 3.2 and 3.3 we briefly describe the technique and its implementation in GADGET; in Section 3.4 we show two of the tests performed to check the implementation and its result in controlled scenarios for which we know the solution, whereas in Sections 3.5 and 3.6 we apply the technique to cosmological simulations of structure formation. We summarise the results and conclude in Section 3.7.

3.2 The formalism

In this section we will briefly report the gist of the algorithm. We refer the reader to the pivotal work PM07 for a thorough treatment of the problem and for a detailed derivation of the new equation of motion.

The Lagrangian for a system of N particles interacting only through gravity reads:

$$L = \sum_{i=0}^N m_i \left(\frac{1}{2} \mathbf{v}_i^2 - \Phi_i \right), \quad (3.3)$$

where

$$\Phi_i(\mathbf{r}_i) = -G \sum_{j=0}^N m_j \phi(|\mathbf{r}_i - \mathbf{r}_j|, h). \quad (3.4)$$

The expression for the potential Φ departs from its Newtonian definition due to the introduction of softening and hence of the kernel ϕ and the related length scale

h . When this quantity has no further dependences and is kept fixed the resulting equation of motion will resemble the Newtonian original, but for the inheritance of this additional scale; its expression is derived by applying the Euler-Lagrange equations to (3.3) and for particle i it will read:

$$m_i \frac{d\mathbf{v}_i}{dt} = \sum_{j=0}^N \mathbf{F}_{ij}(h), \quad (3.5)$$

where

$$\mathbf{F}_{ij}(h) = -\nabla\Phi_{ij} = -Gm_j\phi'(|\mathbf{r}_i - \mathbf{r}_j|, h) \frac{\mathbf{r}_i - \mathbf{r}_j}{|\mathbf{r}_i - \mathbf{r}_j|}. \quad (3.6)$$

Here $\phi' = \partial\phi/\partial|\mathbf{r}_i - \mathbf{r}_j|$ is the force kernel. The expressions for both the potential and the force kernel in the cubic-spline case are given in Appendix A (Equation A.1 and A.2). The force $\mathbf{F}_{ij}(h)$ reduces to the Newtonian gravitational force for separations $|\mathbf{r}_i - \mathbf{r}_j| > h$ and to the kernel-dependent “softened” version otherwise.

When h is allowed to vary from particle to particle according to the density of the environment a few complications arise. In order to maintain its translational invariance and hence ensure that the resulting dynamics would conserve linear momentum, the Lagrangian needs to be symmetrised. One way to do this is by averaging over the potentials evaluated with the softenings of the interacting pair of particles; the new Lagrangian would read:

$$L = \sum_{i=0}^N \frac{1}{2} m_i \mathbf{v}_i^2 - \frac{G}{2} \sum_{i=0}^N \sum_{j=0}^N m_i m_j \left(\frac{\phi_{ij}(h_i) + \phi_{ij}(h_j)}{2} \right), \quad (3.7)$$

where $\phi_{ij}(h_i) \equiv \phi(|\mathbf{r}_i - \mathbf{r}_j|, h)$. When deriving the equation of motion from Eq. 3.7, the additional dependence on the position r through h results in an extra term besides the classical inverse square law. The evolution of a particle i subject only to the action of gravity and supplied with a variable softening length will in fact obey:

$$\begin{aligned} \frac{d\mathbf{v}_i}{dt} &= -G \sum_{j=1}^N m_j \left(\frac{\phi'_{ij}(h_j) + \phi'_{ij}(h_i)}{2} \right) \frac{\mathbf{r}_i - \mathbf{r}_j}{|\mathbf{r}_i - \mathbf{r}_j|} \\ &\quad - \frac{G}{2} \sum_{j=1}^N \left(\frac{\zeta_i}{\Omega_i} \frac{\partial W_{ij}(h_i)}{\partial \mathbf{r}_i} + \frac{m_j \zeta_j}{m_i \Omega_j} \frac{\partial W_{ij}(h_j)}{\partial \mathbf{r}_i} \right), \end{aligned} \quad (3.8)$$

where

$$\zeta_i \equiv \frac{\partial h_i}{\partial \rho_i} \sum_{k=0}^N m_k \frac{\partial \phi_{ik}(h_i)}{\partial h_i}, \quad (3.9)$$

$$\Omega_i \equiv 1 - \frac{\partial h_i}{\partial \rho_i} \sum_{k=0}^N \frac{\partial W_{ik}(h_i)}{\partial h_i}. \quad (3.10)$$

The additional contribution, which we will refer to as the “correction term”, is attractive in nature and will therefore act towards increasing the resulting gravitational force. If this term were not accounted for when using adaptive softening,

the particles would be evolved according to a law inconsistent with the system's Lagrangian, i.e. energy conservation would be lost.

We specify that Equations 3.8 and 3.10 are valid in this form only if softening lengths are determined by fixing the number of *neighbours* within the softening sphere (see Sec. 3.3), as will always be the case in the examples shown in this and the next chapter. In the cases where softening lengths are determined by fixing the total *mass* within the softening sphere, the two equations become, respectively:

$$\begin{aligned} \frac{d\mathbf{v}_i}{dt} &= -G \sum_{j=1}^N m_j \left(\frac{\phi'_{ij}(h_j) + \phi'_{ij}(h_i)}{2} \right) \frac{\mathbf{r}_i - \mathbf{r}_j}{|\mathbf{r}_i - \mathbf{r}_j|} \\ &\quad - \frac{G}{2} \sum_{j=1}^N m_j \left(\frac{\zeta_i}{\Omega_i} \frac{\partial W_{ij}(h_i)}{\partial \mathbf{r}_i} + \frac{\zeta_j}{\Omega_j} \frac{\partial W_{ij}(h_j)}{\partial \mathbf{r}_i} \right) \end{aligned} \quad (3.11)$$

and

$$\Omega_i \equiv 1 - \frac{\partial h_i}{\partial \rho_i} \sum_{k=0}^N m_k \frac{\partial W_{ik}(h_i)}{\partial h_i}. \quad (3.12)$$

In simulations as those presented in this chapter, where particles have all the same mass, choosing either of these two approaches results into the same final softening. However, when multiple particle species coexist, as will be the case for the simulations in Chap. 4, this will not generally hold anymore.

For the full expression of the quantities $\partial\phi/\partial h$, $\partial W/\partial h$ and $\partial W/\partial \mathbf{r}$, see Appendix A.

3.3 Implementation in Gadget

We have implemented the adaptive softening formalism on the latest version of the GADGET code, namely GADGET-3. This computes gravitational forces via the TreePM method (Xu 1995; Bode et al. 2000; Bagla 2002) and hydrodynamical interactions by means of SPH; it differs from the previous versions of the code in that it features a more flexible domain decomposition, something that made it suitable for simulations involving extreme levels of clustering, as the Millennium-II (Boylan-Kolchin et al., 2009).

Implementing adaptive softening required modifications of the Tree-algorithm and of the timestep criterion along with the introduction of the machinery to compute the softening lengths. Besides the obvious change in the expression for the gravitational acceleration, the opening criterion for the nodes has been modified in order to avoid the presence of smoothed particle-node interaction; in other words, the correction term can be non-zero only within a particle-particle interaction (more details in Sec. 3.2.4 of BK09). The timestep criterion employed by GADGET for collisionless particles reads:

$$\Delta t = \left(\frac{2\eta\epsilon}{a} \right)^{1/2} \quad (3.13)$$

where η is an accuracy parameter, ϵ is the Plummer equivalent softening ($h \simeq 2.8\epsilon$, where h is the support of the cubic spline kernel) and a the acceleration of the particle. We kept the criterion itself unchanged and substituted ϵ with the individual value of the adaptive softening.

As to the computation of the softening lengths the method is identical to the one GADGET uses for setting the smoothing lengths in SPH calculations. Qualitatively, they represent the radius of the sphere centred on the particle and containing a specific number of companion particles, generally referred to as “neighbours”; formally, this translates into the following relation:

$$\frac{4\pi}{3} h_i^3 \sum_{j=0}^N W_{ij}(h_i) = N_{ngbs} \quad (3.14)$$

where N_{ngbs} stands for the number of neighbours and N is the number of particles within distance h_j from the target particle j . The above equation is solved iteratively with a Newton-Raphson method until the difference between the two sides falls below a certain tolerance threshold.

Adaptive softening lengths can be activated both when the code works in TreePM mode and when it uses the Tree-only algorithm. In the latter case the softening lengths are left varying without boundaries according to the local features of the particle distribution; in the former case we do instead allow for the presence of a minimum and maximum value for the softening lengths. As explained in Sec. 3.2.1 of BK09, the existence of a minimum value is not crucial and only prevents the simulation from becoming overly expensive in terms of computational time; conversely, the upper bound is introduced to ensure that the long-range force (the particle-mesh contribution) is negligible on the scales where softening is important, so that errors arising from the non-modification of the long range force are under control. These bounds are expressed in terms of the splitting scale r_s , the scale (generally of order the grid spacing) where the splitting of the potential in a long-range and short-range component is performed; choosing $h_{max} \simeq \frac{r_s}{2}$ results in the long-range contribution being below 1% of the total force at scales where softening is important. Although we always impose a lower limit to the softening length when using the code in its TreePM mode, we did not find the presence of an upper limit to have dramatic consequences on the results, especially when using the adaptive formalism in its full, conservative version.

3.4 Tests

In this section we present some of the tests performed in order to check the correctness of the implementation and explore the general effects of adaptive softening when simulating different physical scenarios. We will initially show the behaviour of the code in simulating simple systems, of well-known properties; the force profile of a Plummer and Hernquist models are investigated, and their temporal evolution in and out of equilibrium; the density profile of a polytrope and the behaviour of its total energy in time is also shown. Most of these examples are present already in

PM07 and were specifically chosen to test our implementation.

In all the numerical simulations presented in this section we adopt units of mass $[M] = 1$, length $[R] = 1$ and $G = 1$. As a result, the energy per unit mass is measured in units of GM/R and time in $(GM/R^3)^{-1/2}$.

3.4.1 Evolution of a Plummer sphere

The system considered here consists of a set of N particles distributed according to a ‘‘Plummer’’ profile:

$$\rho(r) = \frac{3GMr_s^2}{4\pi(r_s^2 + r^2)^{5/2}} \quad (3.15)$$

Here the total mass M and the scale radius r_s are set equal to unity. The idea is to evaluate the resulting gravitational force profile and investigate its dependence on the choice of softening. Once this is accomplished we concentrate on the behaviour of the total energy as the system is let evolve in time.

This test is identical to that presented by PM07 and we refer to section 4.3 of their paper or alternatively to Aarseth et al. (1974) for details on the setup of the initial conditions.

The test was run using different number of particles and both Tree-only and TreePM algorithms for the evaluation of the gravitational force; the results behave as expected in the different cases and here we show only those obtained using the pure Tree method on $N = 1000$ particles.

Fig. 3.1 shows the averaged square errors (ASE) of the simulated force field corresponding to different choices of both fixed and adaptive gravitational softening. This quantity measures the deviation of the force experienced by particles at different radii from the analytical value, given by

$$f(r) = \frac{GMr}{(r_s^2 + r^2)^{3/2}}, \quad (3.16)$$

and it is defined as

$$\text{ASE} = \frac{1}{f_{max}^2 N} \sum_{i=1}^N |f_i - f_{exact}(\mathbf{r}_i)|^2, \quad (3.17)$$

where f_i is the force on particle i and f_{max} is the maximum value of the exact solution. For a discussion on the use of the ASE or related quantities to assess the error on a force profile we refer to PM07, Merritt (1996) and Dehnen (2001).

The results show how less sensitive the representation of the force field is to the choice of N_{ngbs} than to the choice of ϵ ; choosing from 30 to 500 neighbours does not shift the global ASE much away from the minimum value corresponding to the ‘‘optimal’’ choice of fixed softening ($\epsilon_{opt} \approx 0.2$, of the order the mean interparticle separation within the scale radius). As already commented by PM07, the slightly worse behaviour obtained when the new force definition is used and $N_{ngbs} < 100$

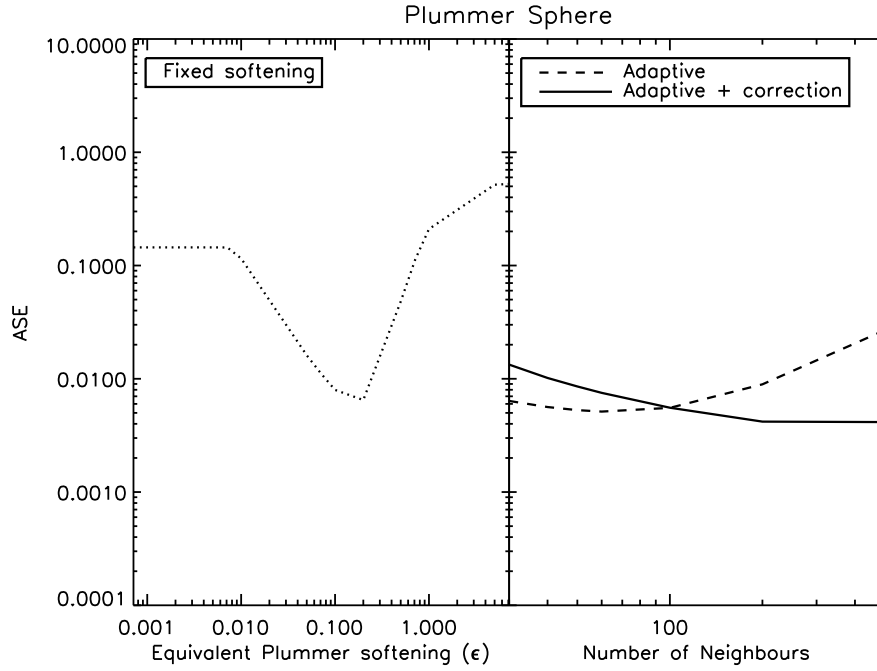


Figure 3.1: Average squared errors (ASE) of the force field generated by distributing $N = 1000$ particles according to a Plummer profile. The left panel shows the behaviour of the ASE as a function of the choice of (fixed) softening; the right panel shows the results of adaptive softening varying the number of neighbours.

can be ascribed to noise-induced gradients in the softening lengths altering the evaluation of the correction term; notwithstanding this, the ASEs for the “adaptive+correction” case are remarkably close to the minimum corresponding to the optimal choice of fixed softening throughout the full range of N_{ngbs} . These results compare very well with those of PM07 (see their Fig. 2 and the second set of lines from the top). Choosing the optimal value for ϵ and $N_{ngbs} = 60$ we have evolved the system in time and checked the conservation of the total (kinetic plus potential) energy. Fig. 3.2 and 3.3 show the results for two different initial velocity setups, corresponding to a dynamical equilibrium and a perturbed state respectively. In the first case we expect no major evolution of the system (at least for several relaxation times) whereas in the second an initial overall expansion occurs before a state of equilibrium is reached. We refer again to PM07 and Aarseth et al. (1974) for details on the setup of the velocity profiles. The energy fluctuates or increases, reflecting the global changes in the softening lengths, when the adaptive formalism is used without changing the equation of motion accordingly; we register fluctuations of order per cent in the equilibrium set-up and an increase of $\approx 6\%$ in the perturbed set-up. Conservation is instead re-established down to time-stepping accuracy as soon as the correct equation of motion is used. Note that the initial total energies are different in the adaptive and fixed softening cases; this is due to the definition and evaluation of the potential energy being different in the two scenarios (see Sec. 3.2).

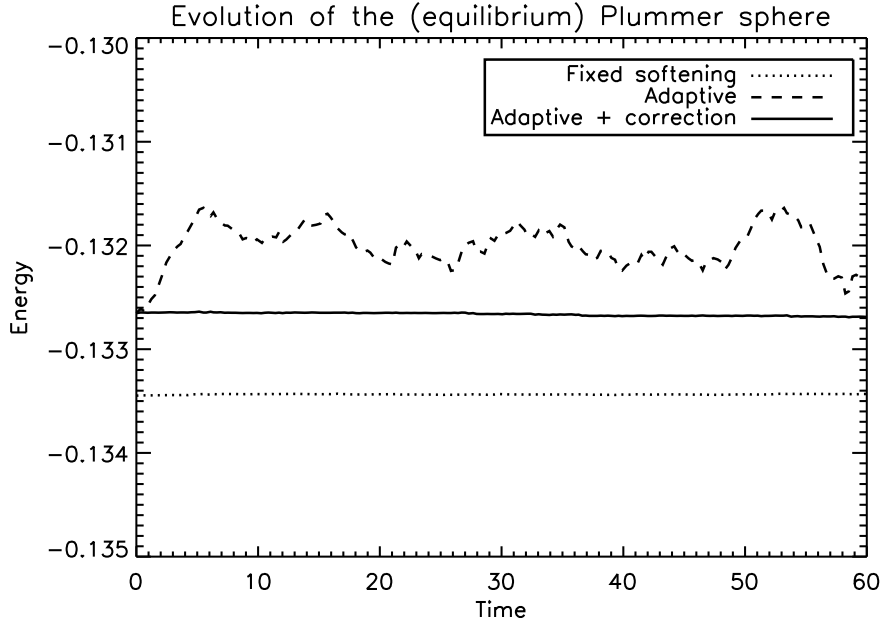


Figure 3.2: Behaviour of the total energy (kinetic plus potential) of the Plummer sphere as a function of time. The velocities were generated according to the equilibrium distribution function. For a description of the units, see the introduction to Sec. 3.4.

3.4.2 Evolution of a Hernquist model

We now perform a similar analysis on a different density distribution; we generate a Monte Carlo realisation of a Hernquist model (Hernquist, 1990), whose density profile reads:

$$\rho(r) = \frac{GMr_s}{2\pi r(r_s + r)^3}. \quad (3.18)$$

The model is cusped near the origin and provides a more realistic representation of the matter distribution in collapsed objects of cosmological interest, but the steep rise in the density profile makes the distribution more difficult to resolve than the Plummer case addressed in the previous section. We investigate the ASE and the energy evolution for a Hernquist model in equilibrium, sampled with $N = 1000$ particles; as in the previous section, the total mass M and the scale radius r_s are set equal to unity. Fig. 3.4 shows the averaged square errors (ASE) of the simulated force field corresponding to different choices of both fixed and adaptive gravitational softening. Again we see how adapting softening provides relatively stable results when varying the number of neighbours from 30 to 500; conversely, one needs to choose the fixed softening accurately to avoid a strong misrepresentation of the force profile. We note that the optimal value for ϵ is smaller here than for the Plummer sphere ($\epsilon_{opt} \approx 0.05$ vs. $\epsilon_{opt} \approx 0.2$), a consequence of the different density distribution in the two models: the Hernquist being cusped and the Plummer being

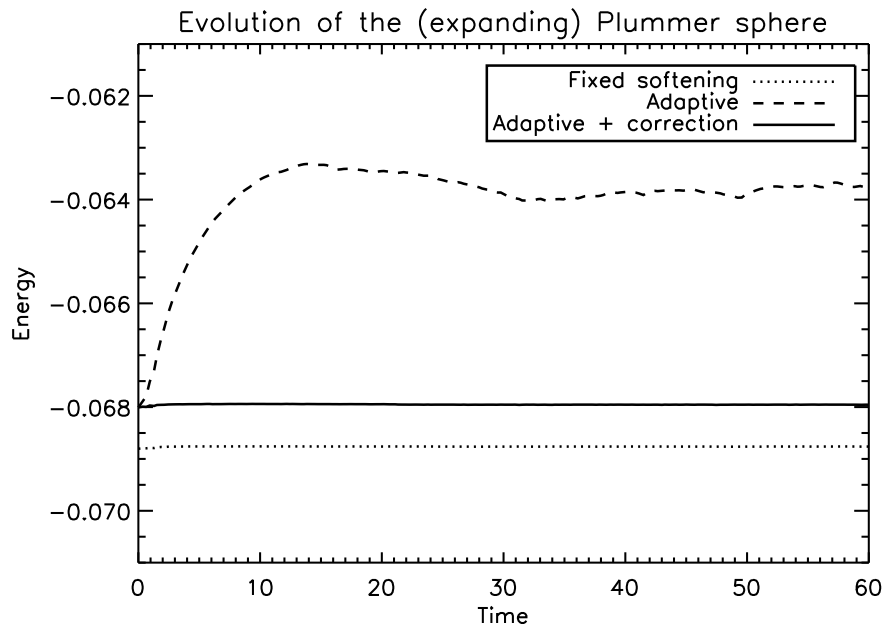


Figure 3.3: Behaviour of the total energy (kinetic plus potential) of the Plummer sphere as a function of time. The equilibrium velocity distribution has been multiplied by a factor 1.2, leading to an initial global expansion of the system. For a description of the units, see the introduction to Sec. 3.4.

flat at the centre. A number of neighbours $N_{ngbs} \approx 60$ is enough to overcome the noise in the evaluation of the correction term and to provide better results than in the case where plain adaptive softening is used. For any number of neighbours, however, the force profile of the Hernquist model is reproduced better by adapting the gravitational softening. The results are in qualitative agreement with those of PM07 (see their Fig. 3 and the second set of lines from the top). As before, we now choose the optimal value for ϵ and $N_{ngbs} = 60$ to evolve the system in time, checking the behaviour of the total energy. The results are shown in Fig. 3.5 for the equilibrium model; again, the energy fluctuates ($\approx 1\%$) when the adaptive formalism is used without changing the equation of motion accordingly, whereas conservation is re-established down to time-stepping accuracy as soon as the correct equation is used.

A more demanding test for energy conservation is to follow the collision of two such particle distributions and check if the adaptive formalism maintains its conservative properties even in this more dynamically active scenario. To this purpose we have simulated the head-on collision of two different realisations of the Hernquist model, where the two systems were placed at a distance r of 150 scale radii and made approach each other at a speed $v = \sqrt{(2GM/r)}$. The evolution of the total energy is shown in Fig. 3.6; the collision between the two spheres at $t \simeq 900$ leaves a clear imprint in the total energy, when the softening is allowed to vary and the standard equation of motion is used, while as soon as the correction term is taken into account conservation is recovered.

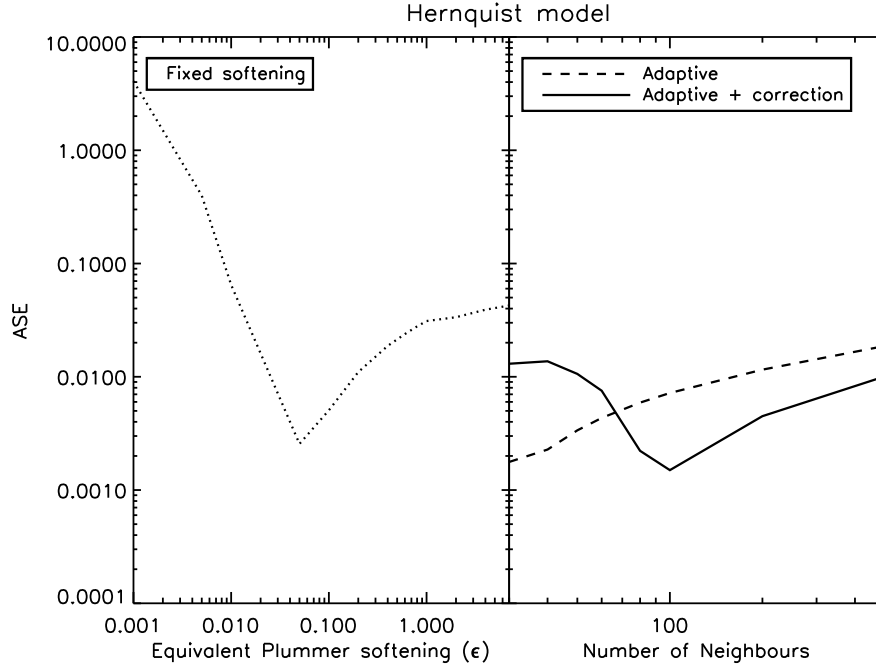


Figure 3.4: Average squared errors (ASE) of the force field generated by distributing $N = 1000$ particles according to a Hernquist profile. The left panel shows the behaviour of the ASE as a function of the choice of (fixed) softening; the right panel shows the results of adaptive softening varying the number of neighbours.

We now want to investigate whether or not the use of adaptive softening helps maintaining the original slope of the density profile in time. To this purpose we let the equilibrium system evolve for several dynamical times and compare the matter distribution among the different models. Fig. 3.7 shows the cumulative density profile at four different times, marked in the upper left corner. In all the panels the gray line represent the initial profile; the solid lines correspond to the optimal choice for fixed softening ($\epsilon_{opt} \approx 0.05$) and to $N_{ngbs} \approx 60$ for adaptive softening; the dotted lines correspond to $\epsilon = 0.001$ and $N_{ngbs} = 30$, whereas the dashed ones to $\epsilon = 1.0$ and $N_{ngbs} = 500$. Again we see how sensitive the results are to the choice of ϵ and, correspondingly, how little they depend on N_{ngbs} , especially when the correct equation of motion is used. The density profile in the runs with adaptive softening is compatible with the original one at all times (and for all N_{ngbs} , when the correction term is added), with a tendency to outperform the results obtained when using fixed softening.

3.4.3 Equilibrium structure of a polytrope

Here we consider a system evolving under the action of both gravity and hydrodynamical forces. A homogeneous gas sphere of initial radius $r_0 = 1$ with equation of state $P = K\rho^\gamma$, $\gamma = 5/3$ and initial null internal energy is released to the action of self-gravity and pressure force until hydrostatic equilibrium is reached. The process is speeded up by means of the standard SPH artificial viscosity together

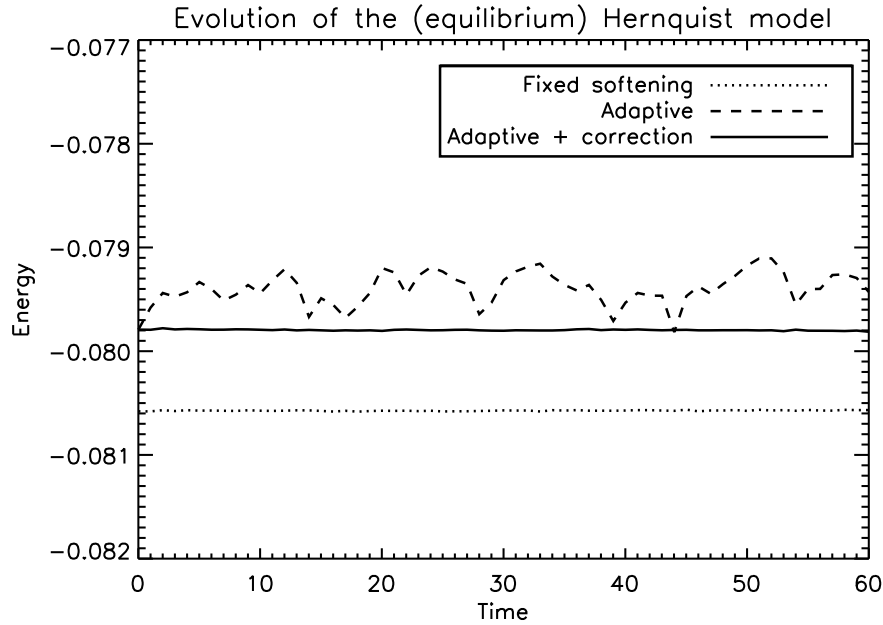


Figure 3.5: Behaviour of the total energy (kinetic plus potential) of the Hernquist model as a function of time. The velocities were generated according to the equilibrium distribution function. For a description of the units, see the introduction to Sec. 3.4.

with a damping term in the force equation; their effect is to deliberately remove kinetic energy helping the system settling faster to equilibrium. An analytical solution for the internal structure of the system does not exist when $\gamma = 5/3$; the exact radial density profile was then obtained by numerically integrating the corresponding Lane-Emden equation:

$$\frac{\gamma K}{4\pi G(\gamma - 1)} \frac{d^2}{dr^2}(r\rho^{\gamma-1}) + r\rho = 0. \quad (3.19)$$

The reference for the setup and realisation of the test is Sec. 4.4 of PM07 and there we refer the reader for additional details on the generation of initial conditions and features of the run.

Fig. 3.8 shows the density profile of the simulated system at $t = 40$, plotted against the equilibrium solution given by Eq. 3.19. The units are as those introduced in the previous section; the time was chosen by following the oscillations in the density at $r = 0.2$: when their amplitude reduced of a factor of 20 with respect to the initial, we assumed equilibrium was reached.

Three cases are shown, namely fixed softening ($\epsilon \approx 1/40$ of the mean interparticle separation, as chosen by PM07) and adaptive softening ($N_{ngbs} = 60$) with and without the use of the correction term. In all cases we have used 60 neighbours for the SPH computations. As found by PM07, the use of adaptive softening provides a better representation of the density profile, especially the inner regions; the higher central densities can be related to a smaller gravitational softening than in the fixed case (at least when the correction term is employed), while the slightly better agreement at large radii comes from much larger softening reducing the noise of the

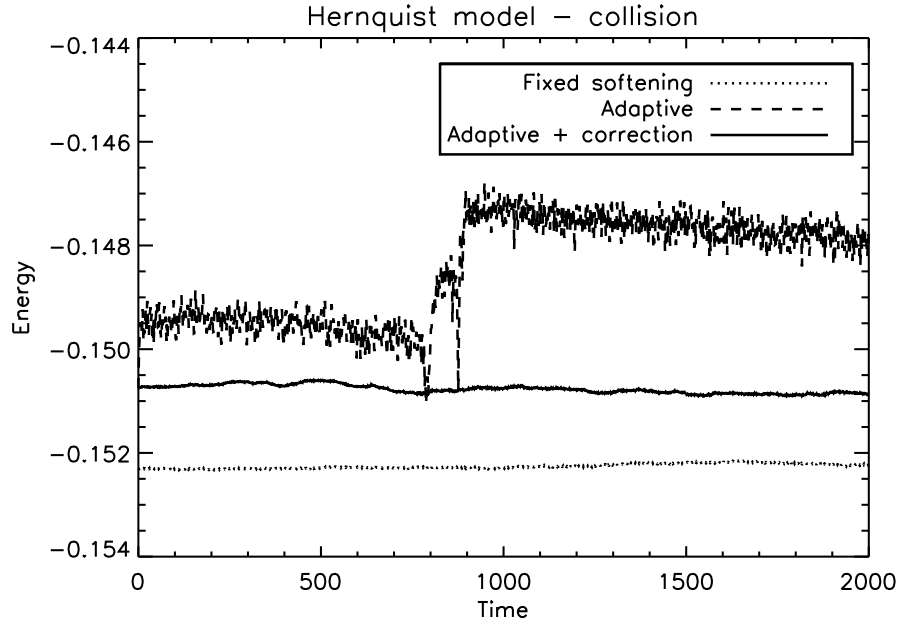


Figure 3.6: Behaviour of the total energy of two colliding Hernquist spheres as a function of time. The collision is head-on, with the two systems at an initial separation $r = 150 r_s$ and approaching each other at a speed $v = \sqrt{(2GM/r)}$. For a description of the units, see the introduction to Sec. 3.4.

particle distribution.

The equilibrium solution just found has then been perturbed by inducing radial oscillations on the sphere. The system was let evolve under gravity and pressure force only, the behaviour of its total energy being now of interest. The results are displayed in Fig. 3.9. The system oscillates with a period $\tau \approx 4$, close to the expected 3.82 for the fundamental mode of oscillation of such a polytrope. This is seen clearly in the behaviour of the energy when adaptive softening is used and the correction to the equation of motion is neglected; the evaluated potential becomes cyclically deeper and shallower following, respectively, the overall decrease and increase of the softening lengths. Energy conservation is re-established at the level of the runs with fixed softening once the system is evolved according to the appropriate equations.

3.5 Performance in a cosmological environment

The problem of the formation and evolution of the large-scale structure of the Universe has no analytical solution and this complicates the analysis when different simulations of the same system are compared; to assess the results and determine which technique is dealing best with the problem is not obvious anymore. Increasing the resolution of a cosmological simulation generally goes along with extending the representation of the initial perturbation field to smaller scales, resulting in new fluctuations entering the initial conditions and therefore, strictly speaking, in

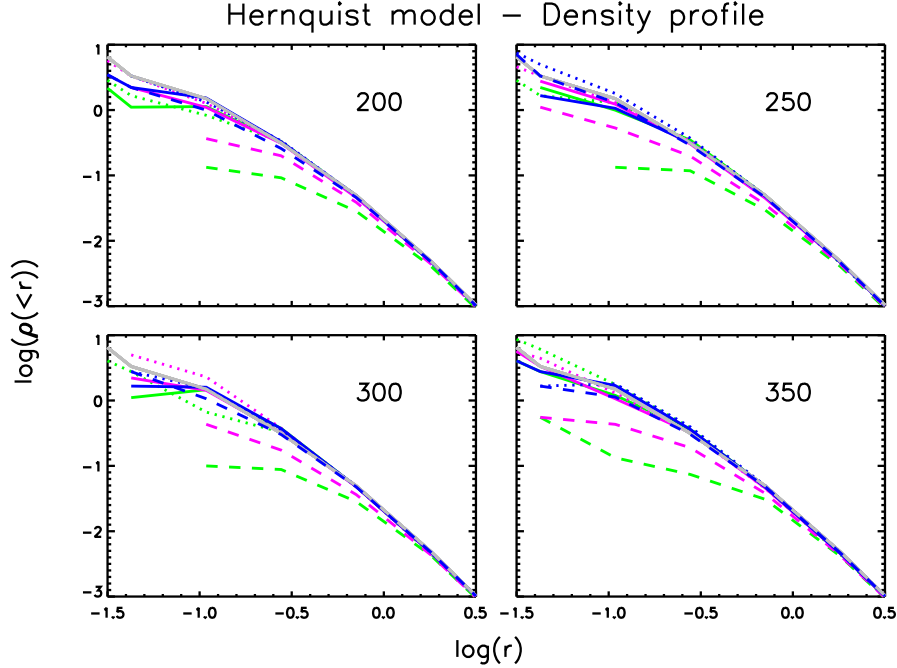


Figure 3.7: Cumulative density profile for the 1000-particles Hernquist model at four different times ($t = 200, 250, 300, 350$). The gray lines represent the initial profile at $t = 0$; solid lines correspond to the fiducial choice for $\epsilon(0.05)$ and $N_{ngbs}(60)$; dotted lines correspond to $\epsilon = 0.001$ and $N_{ngbs} = 30$, while dashed lines to $\epsilon = 1$ and $N_{ngbs} = 500$.

a new system. In order to address the problem of structure formation and of how this is dealt with by the different softening formulations we have decided to perform a series of simulations sharing the same initial power spectrum of fluctuations, but differing in the total number of particles (Binney 2004): the same structures form at all the resolution levels, but are more accurately described as more and more particles populate them. If the use of adaptive softening allowed to anticipate the behaviour shown by the standard runs at higher resolution, this would assess the superiority of the method over the usual choice of fixed softening.

The simulated system here consists of a $100 h^{-1}$ Mpc side periodic box in a Λ CDM cosmology defined by the following choice of parameters:

$$\begin{aligned} \Omega_{\text{tot}} = 1, \quad \Omega_m = 0.3, \quad \Omega_b = 0.04, \quad \Omega_\Lambda = 0.7, \\ h = 0.7, \quad \sigma_8 = 0.9, \quad n_s = 1, \end{aligned} \quad (3.20)$$

where h and σ_8 are the values, at redshift zero, of the dimensionless Hubble parameter and rms of the mass fluctuations smoothed on a scale of $8 h^{-1}$ Mpc, whereas n_s is the index of the primordial spectrum of fluctuations. Two sets of simulations have been performed: three runs with 64^3 , 128^3 and 256^3 particles using fixed softening and two runs with 64^3 and 128^3 particles using adaptive softening (with and without the correction term). The initial conditions were generated at $z = 60$ and differ among them only in the number of particles resolving the same initial fluctuation field. The pre-initial uniform distribution, represented by an equally-spaced

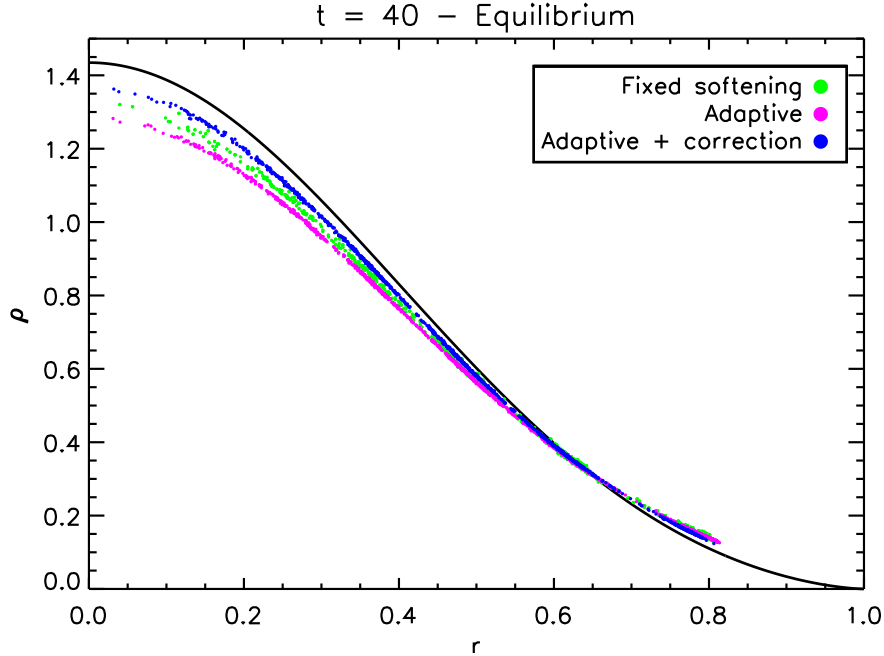


Figure 3.8: Radial density profile of a polytropic sphere in hydrostatic equilibrium. The solid black line represents the numerical solution of the Lane-Emden equation (Eq. 3.19); the points are the results obtained by evolving a $N = 1445$ -particles system with different choice of softenings.

grid of particles, is applied a displacement field generated on 64^3 , 128^3 and 256^3 grids, depending on the resolution level. The power spectra are anyway identical in all the three cases and they are truncated at the smallest frequency resolved on a 64^3 grid. The simulations follow only the evolution of the dark matter component of the density field and the resulting mass associated to a single particle varies from ≈ 32 to 4 and $0.5 \cdot 10^{10} h^{-1} M_{\odot}$ depending on the resolution level of the run. In all the cases we have used the TreePM algorithm varying the grid from 64^3 to 128^3 and 256^3 , according to the resolution level. The features of the simulations are summarised in Table 3.1.

The value of fixed softening was chosen to be $1/40$ of the mean interparticle separation in the system, i.e. the smallest advisable value according to the generally accepted criterion that limits this scale to avoid too negative binding energies in close pairs. The choice of the number of neighbours for the adaptive runs corresponds to the minimum number ensuring a robust evaluation of the correction term. To this purpose we have monitored the changes in the two crucial quantities entering the evaluation of the correction term, namely ζ and Ω (Eq. 3.9 and 3.10, respectively), when varying N_{ngbs} . The code was run on the $z = 0$ snapshot of the 128^3 simulation employing fixed softening with N_{ngbs} varying from 16 to 80 with a step $\Delta N_{ngbs} = 8$. Fig. 3.10 and 3.11 show the distribution of the two quantities when varying N_{ngbs} ; we can conclude that a number $N_{ngbs} = 60$ is enough to ensure a converged evaluation of ζ and Ω and that the resulting correction term would be robustly estimated. We have also run the simulations using adaptive softening and

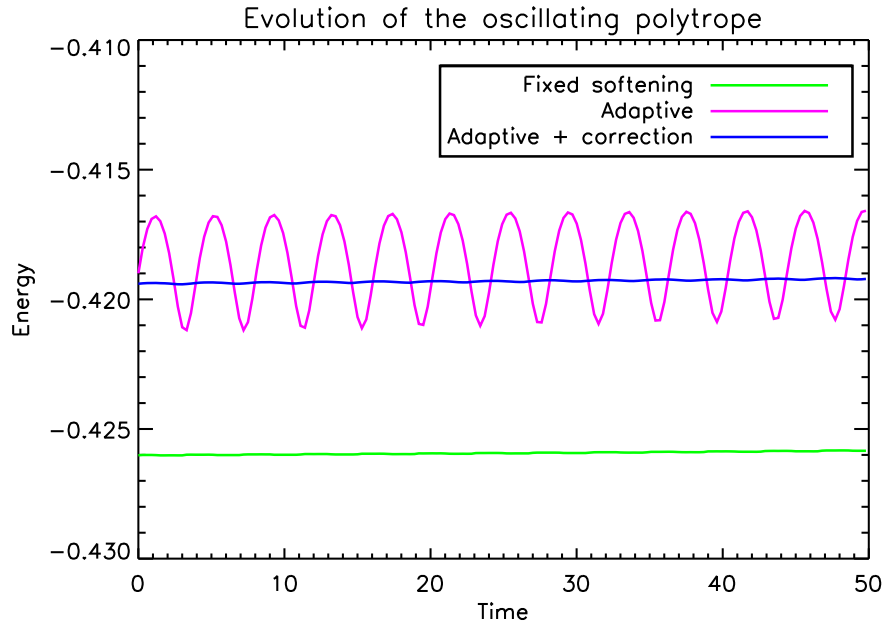


Figure 3.9: Behaviour of the total energy (kinetic, potential and internal) of the polytropic sphere as a function of time. The system has been perturbed from the equilibrium state at $t = 40$ by the induction of radial oscillations. For a description of the units, see the introduction to Sec. 3.4.

no correction to the equations of motion, using the same number of neighbours as the fully adaptive ones; although the results of the previous section suggest that the introduction of the correction term is crucial to the proper functioning of the adaptive formalism, we will anyway show the behaviour of these simulations for completeness.

In the following subsections we summarise the main results.

3.5.1 Global behaviour I - Clustering

The level of clustering in the simulations has been first assessed qualitatively by investigating the densities achieved throughout the box. Fig. 3.12 shows the particles with an associated density greater than 10^5 , $5 \cdot 10^5$ and 10^6 times the average density; the results are from the redshift zero snapshots of the 128^3 and 256^3 simulations. The densities have been computed in the SPH fashion, using 60 neighbours for all the runs. As one can see, the runs with adaptive softening reach higher particles densities. Noticeably, the regions where this enhancement of clustering is observed correspond to the high density regions of the 256^3 simulation; indeed, moving from left to right (i.e. from the fixed softening case to the fully adaptive one) the same areas can be seen to be more and more populated. The numbers in the upper-right corner corresponds to the total number of particles surviving the density threshold; interestingly, the number of these high-density particles in the 256^3 simulation correspond to a mass of ≈ 23000 , 10000 , 300 particles at the 128^3 resolution level.

Table 3.1: Basic features of the simulations presented in Sec. 3.5.

Name	ϵ [h^{-1} kpc]	N_{ngbs}	m_p [$h^{-1} M_{\odot}$]	k_{min} [h kpc $^{-1}$]	k_{max} [h kpc $^{-1}$]
Fixed - 64^3	39	-	$31.76 \cdot 10^{10}$	$6.28 \cdot 10^{-5}$	$2 \cdot 10^{-3}$
Fixed - 128^3	19.5	-	$3.97 \cdot 10^{10}$	$6.28 \cdot 10^{-5}$	$2 \cdot 10^{-3}$
Fixed - 256^3	10	-	$0.5 \cdot 10^{10}$	$6.28 \cdot 10^{-5}$	$2 \cdot 10^{-3}$
Adapt - 64^3	-	60 ± 0.1	$31.76 \cdot 10^{10}$	$6.28 \cdot 10^{-5}$	$2 \cdot 10^{-3}$
Adapt+corr - 64^3	-	60 ± 0.1	$31.76 \cdot 10^{10}$	$6.28 \cdot 10^{-5}$	$2 \cdot 10^{-3}$
Adapt - 128^3	-	60 ± 0.1	$3.97 \cdot 10^{10}$	$6.28 \cdot 10^{-5}$	$2 \cdot 10^{-3}$
Adapt+corr - 128^3	-	60 ± 0.1	$3.97 \cdot 10^{10}$	$6.28 \cdot 10^{-5}$	$2 \cdot 10^{-3}$

On a more quantitative level, we have also analysed the clustering by means of the two-point correlation function $\xi(r)$. This statistics represents the excess probability, compared to a uniform random distribution, of finding pairs of particles at a given spatial separation. Fig. 3.13 shows the ratio of the correlation functions obtained in the different runs to the result of the 256^3 simulation; the curves are shown at $z = 0$ and starting from separations of $10 h^{-1}$ kpc, which correspond roughly to the resolution scale at the 256^3 level. Other than for some discrepancies at separation of several megaparsecs¹, the correlation functions of the three runs with fixed softening overlap almost perfectly down to $100 h^{-1}$ kpc. Below that scale the differences in resolution result in a progressively larger amplitude. Since the behaviour of $\xi(r)$ at the small- r end is mainly determined by the distribution of particles within individual halos, we could conclude that the internal distribution of particles in a structure becomes denser as the resolution of the simulation is increased. The runs with adaptive softening produce a correlation function which is in full agreement with that of the “fixed” run at the same resolution down to roughly $100 h^{-1}$ kpc; at smaller separation the amplitude grows instead larger, approaching the results obtained at higher resolution when using fixed softening. This is particularly evident in the 64^3 -run using adaptive softening and the correction of the equation of motion (“Adapt+corr - 64^3 ”, blue curve): not only is the amplitude at small separation ($10 - 50 h^{-1}$ kpc) higher than the other two 64^3 runs, but the behaviour at scales between 50 and $200 h^{-1}$ kpc is indistinguishable from the reference 128^3 -run (black, dashed curve). Both the “adaptive” runs at the 128^3 resolution level behave well in reproducing the correlation function of the 256^3 simulation, the one with correction (purple curve) even slightly exceeding the results for the 256^3 case on scales below $100 h^{-1}$ kpc.

¹These discrepancies have no physical meaning and are due to a somewhat less accurate determination of the correlation function at large separations. The method we use is Monte-Carlo based and tuned to provide very solid results at small separations, at the expense of accuracy on large scales. We checked that the differences we register are within the typical uncertainties due to the method when the correlation function is computed at these separations.




Figure 3.10: Number of particle with ζ (Eq. 3.9) in a certain range as a function of ζ . The code was run on the same clustered distribution of particles using different values for the number of neighbours.

3.5.2 Global behaviour II - Mass function

The search for structures in the simulations has been carried out with the algorithms FOF (Davis et al. 1985b) and SUBFIND (Springel et al. 2001a). Structures are first identified as collections of $N > N_{min}$ particles separated by mutual distances smaller than some fraction b of the mean interparticle separation. These so-called “FOF halos” are later examined by SUBFIND for the identification of self-bound substructures and the removal of spurious background particles. In the simulations presented here the halos have been searched using a linking length $b = 0.16$ and a minimum threshold of 32 particles. The internal structure of these candidate objects has then been probed in order to identify local, gravitationally-bound overdensities; those containing at least 20 particles were addressed to as “subhalos” leaving the others as part of the smooth halo component. Finally, particles not gravitationally bound to any substructure of the parent FOF halo were dismissed. Note that SUBFIND was modified in the unbinding part so that the evaluation of the gravitational potential takes into account the individual softening length of the particles. Fig. 3.14 shows the mass function of the FOF halos at $z = 0$; the number of objects per logarithmic mass interval is plotted for the different runs down to the mass limit corresponding to 32 particles. When the correct equation of motion is used, the agreement between the simulations with adaptive and fixed softening at the same resolution level is striking throughout the full mass range; when the correction term is instead ignored, the mass functions drop visibly at low




Figure 3.11: Number of particle with Ω (Eq. 3.10) in a certain range as a function of Ω . The code was run on the same clustered distribution of particles using different values for the number of neighbours.

masses, the number of objects containing a number of particles less than of order N_{ngbs} being significantly underestimated. Overall, the runs with adaptive softening and correction underestimate the total number of FOF halos by $\simeq 3\%$, as opposed to a $\simeq 25\%$ for those without the correction term; if we considered only halos containing more than 100 particles, these percentages would fall to $\simeq 1\%$ and $\simeq 14\%$, respectively. Similar results are obtained when considering the “spherical overdensity” masses M_{Δ}^2 ; this hints to the fact that the global shape of the structures in the simulations is not changed significantly by the adoption of the adaptive softening formalism.

The upturn in the curve for the 256^3 simulation at the low-mass end resembles the effect of spurious halos seen in warm dark matter simulations (see, e.g. Wang and White 2007 and their Fig. 9); these objects form from the artificial fragmentation of filaments and have masses much smaller than the free-streaming mass of the model being adopted, their origin being entirely due to the specifics of the pre-initial conditions. Having our simulations a truncated power-spectrum, we might in principle be hitting the same problem and have a contribution to our mass functions coming from such spurious structures. This could indeed be the case for the 256^3 simulation at the very low-mass end, a regime we are not interested in for our com-

²These are defined as the masses of the spherical regions centred on the potential minimum of the smooth halo and corresponding to an overdensity of Δ , typically ≈ 200 , with respect to either the critical density ρ_{crit} or to the mean background density $\rho_m = \Omega_m \rho_{crit}$. Hereafter, when referring to either the virial radius or the virial mass of an object, we assume the overdensity to be defined with respect to the mean background density.

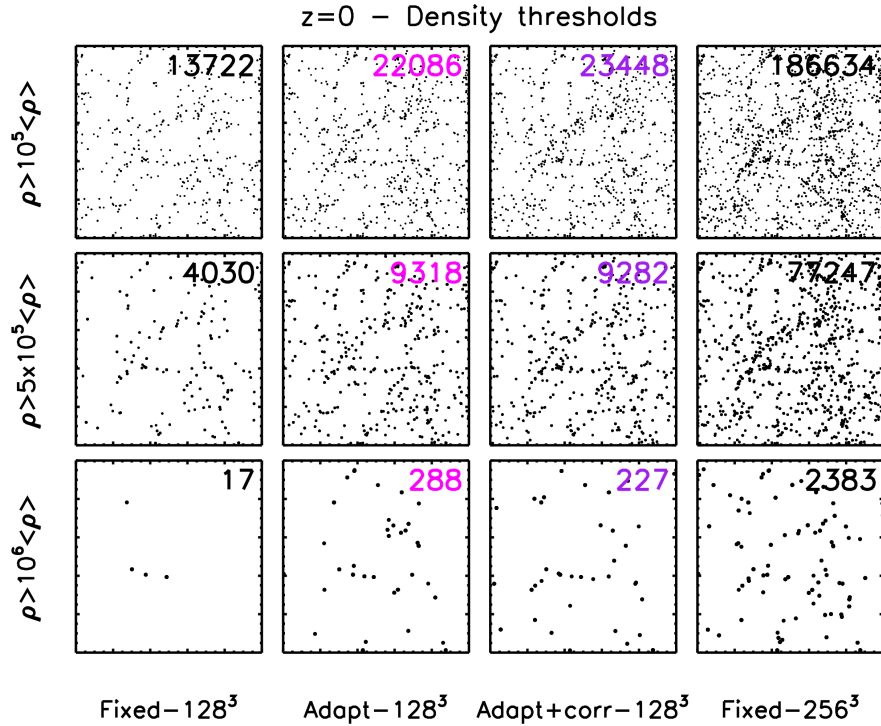


Figure 3.12: Particles with an associated density greater than a threshold. Their total number is shown in the upper-right corner. The full $100 h^{-1}$ Mpc side box is shown and all the particles are plotted; some are not distinguishable from their neighbours due to their proximity, which in some cases can reach $0.5 h^{-1}$ kpc. The densities were computed as the position-weighted sum of the nearest neighbours, in the usual SPH fashion; the same number of neighbours (60) has been used for all the simulations. The results are shown at $z = 0$.

parison. As for the 128^3 simulations, the problem would affect objects with masses corresponding to a handful of particles that do not survive the 32-particles limit and therefore do not enter the evaluation of the mass functions. We also specify that this problem does not invalidate the results we displayed in the previous section in terms of correlation functions; the number of particles in the alleged spurious structures in the 256^3 simulation is insufficient to affect the evaluation of the correlation function at the scales we are interested.

3.5.3 Internal halo properties

The results in terms of correlation function and mass function suggest that the main differences between the runs are likely to manifest in the internal structures of the collapsed objects. We will investigate here what effects the different definitions of softening have in the representation of the most massive halo formed in the simulations. Fig. 3.15 shows the cumulative mass distribution in the halo out to the virial radius. The values differ by fractions of per cent in the various runs and assess around $2.4 h^{-1}$ Mpc. Again it is evident how the use of adap-

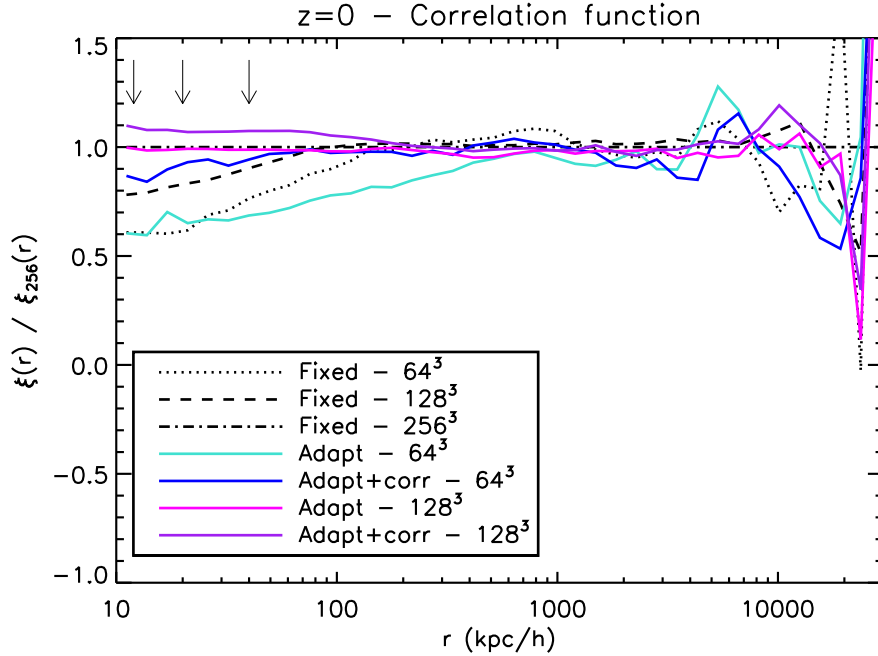


Figure 3.13: Two-point correlation function at $z = 0$ for a series of simulations of a $100 h^{-1}$ Mpc side box in a Λ CDM universe; shown is the ratio to the results of the 256^3 simulation. The runs differ only in the choice of softening and in the total number of particles used to sample the initial perturbation field. The arrows indicate the equivalent-Plummer values of the gravitational softening in the standard runs.

tive softening enhances the clustering of particles anticipating the behaviour of higher resolution simulations. The results of the two runs with adaptive softening and correction (“Adapt+corr - 64^3 ”, blue curve; “Adapt+corr - 128^3 ”, purple curve) are particularly worth of notice: the first reproduces the higher-resolution results down to $20 h^{-1}$ kpc, whereas the second is perfectly compatible with the highest resolution run down to less than $5 h^{-1}$ kpc. Another way to interpret this result is in terms of mean inner density as a function of the enclosed number of particles. As noted by Moore et al. (1998) and Power et al. (2003), obtaining robust results in regions closer to the centre, where the density is higher, demands increasingly large particle numbers. As in Fig. 14 of Power et al. (2003) we have compared the mean inner density measured at different fractions of the virial radius as a function of the enclosed particle number for all the runs. The results are displayed in Fig. 3.16. The 64^3 simulation using fixed softening provides converged results at most down to $\simeq 6\%$ of the virial radius. At smaller radii the results start to diverge from those obtained at higher resolution and cannot be trusted anymore. The two runs with adaptive softening behave instead very well down to 3% of the virial radius. Similarly, the 128^3 simulation is reliable down to $\simeq 2\%$ of the virial radius when using fixed softening and down to 1% of the virial radius when adaptive softening is employed. Power et al. (2003) relate the converged side of the enclosed- ρ vs. enclosed- N plot to the regions where the average collisional relaxation time t_{rel} exceeds some fraction of the age of Universe t_0 (between 0.6 and 1). In our case,

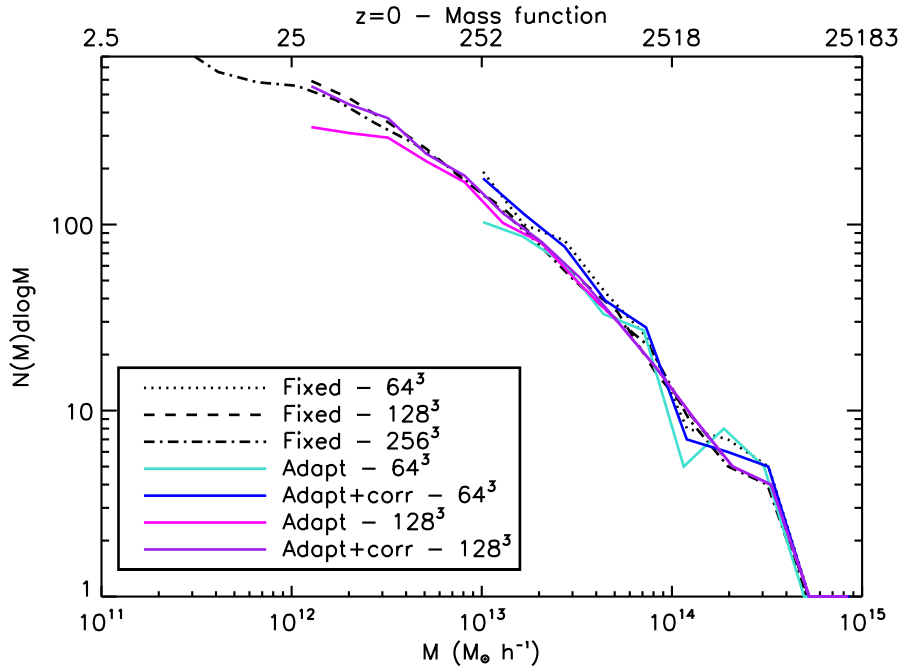


Figure 3.14: FOF mass function at $z = 0$ for a series of simulations of a $100 h^{-1}$ Mpc side box in a Λ CDM universe. The upper x-axis displays the number of particles in the 128^3 run corresponding to the masses in the lower axis. The mass functions are truncated at a mass corresponding to 32 particle masses in every resolution level (i.e. at $\approx 10^{13}$, $1.3 \cdot 10^{12}$, $1.7 \cdot 10^{11} h^{-1} M_{\odot}$ for the 64^3 , 128^3 and 256^3 run, respectively). The simulations differ only in the choice of softening and in the total number of particles used to sample the initial perturbation field.

depending on whether we consider the results at 6% of the virial radius converged or not, we could extend the regions down to the radii where $t_{rel} \approx 0.4t_0$. In any case, it is not clear whether the 256^3 run can be considered converged according to this criterion. It is anyway worth of notice that the 128^3 adaptive run with correction reproduces perfectly the enclosed density of the 256^3 at 1% of the virial radius. Similar results hold when investigating the properties of substructures. Fig. 3.17 shows the number of subhalos with mass greater than a certain value; the curves represent the average over the five biggest halos in the simulations and the masses have all been normalised to the virial mass of the host. Due to the poverty of substructures we are not showing the results for the lowest resolution simulations and we are in general limited to qualitative considerations. As can immediately be seen though, the runs with adaptive softening lie closer to the higher resolution simulation than does the run with fixed softening.

3.5.4 Comments

All the results shown so far hold at redshift zero. As shown in Fig. 3.19, even at this redshift only $\approx 1\%$ of the particles have softening smaller than $1/40$ of the mean interparticle separation and at higher redshifts this behaviour becomes even

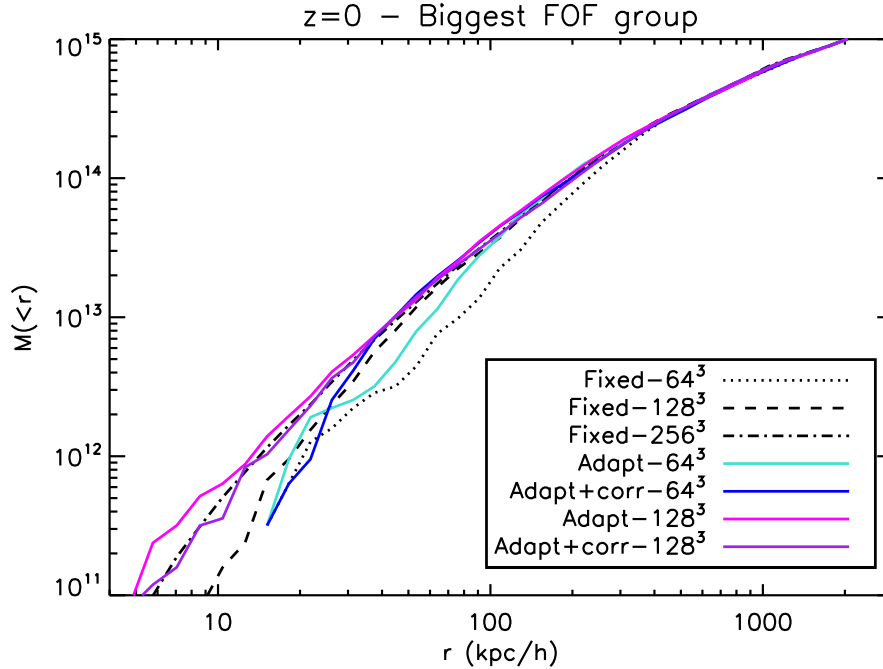


Figure 3.15: Cumulative mass distribution for the most massive FOF group at $z = 0$.

more extreme. This has the effect of reducing the amplitude of the correlation function more rapidly than in the standard runs when going back in time. A more quantitative representation of this effect will be given in the next section (Fig. 3.20). The mass functions are not substantially affected though, at least not for objects more massive than the 32-particles limit (Fig. 3.21).

As the particle timesteps depends on the value of the gravitational softening (see Eq. 3.13), it is natural to expect that they would now span a wider range of values. Fig. 3.18 shows that this is the case; plotted are the distributions of particles in time bins³ at redshift zero for the three 128^3 simulations: the “adaptive” runs tend to have particles in lower bins than in the “fixed” simulation and, at the same time, more particles in the highest bin. This is a consequence of what was just mentioned, namely that only a very small fraction of the particles obtains an adaptive softening smaller than the fiducial choice for the standard simulations at the same resolution level. As finer time bins get populated, the overall number of timesteps increases when adaptive softening is used (along with the number of operations performed within them); at the same time, since the highest bins become more crowded, the average number of active particles - i.e. those that are advanced in a timestep - decreases instead, almost compensating for this overhead. Increasing N_{ngbs} has the effect that one intuitively imagines, namely that of increasing the softening associated to the particles; the differences again stand out more evident

³In GADGET the simulated timespan is mapped onto the integer interval $[0, 2^{N_{timebins}}]$. This interval is split recursively into timebins, the smallest of which has length 2. Each time, the code computes individual timesteps for the particles and distributes them in the correspondent bin; the smallest populated bin sets the next timestep for the simulation.

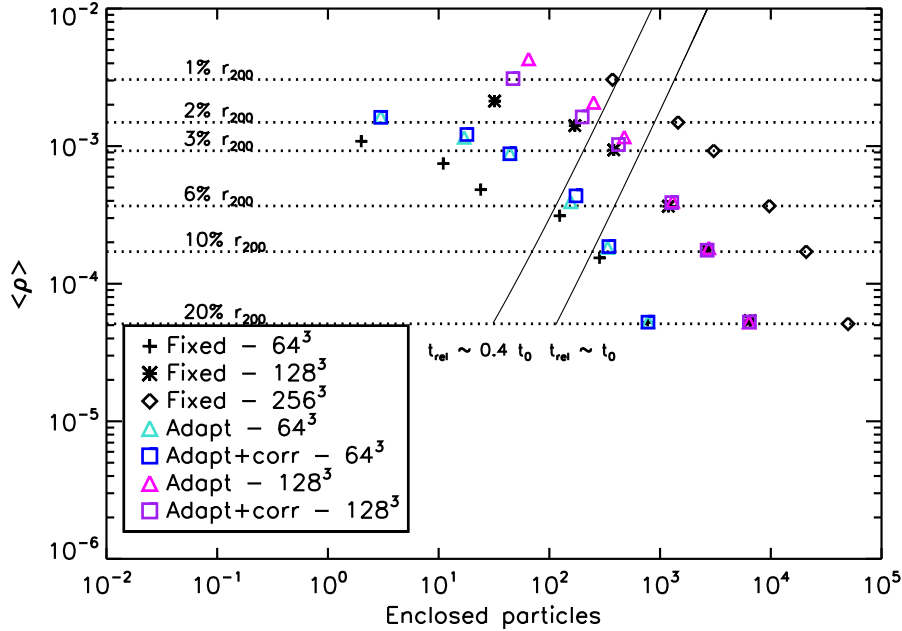


Figure 3.16: Mean inner density as a function of the enclosed number of particle at different fractions of the virial radius. The halo under consideration is the same as in Fig. 3.15. The solid lines separate the regions, at their right, where the average collisional relaxation time (t_{rel}) exceeds some fraction of the age of Universe (t_0).

when looking at the correlation functions, whose amplitude at small separations reduces, whereas the mass functions are confirmed to be considerably less sensitive to variations in N_{ngbs} .

All the simulations discussed in this section the gravitational softening was prevented to fall below 0.1% of the TreePM splitting scale r_s (corresponding to 558 pc for the 64^3 case and to 279 pc for the 128^3 one; in both cases, this translates to $\approx 1.5\%$ of the softening scale in the “fixed” runs); this implies that the maximum density achievable in the simulation corresponds to roughly $10^9 \bar{\rho}$. As pointed out in Sec. 3.3, the presence of a lower bound is not crucial and it is introduced mainly to prevent few particles in over-dense regions from slowing down the simulation. No upper limit to the gravitational softening was set in these simulations; we do not find substantial differences between the runs with and without such a constraint, especially if the correct equation of motion is used. We have also checked both sets against Tree-only runs and found perfectly compatible results.

Although we both observe an increased level of clustering in massive objects, our global results in terms of particle correlation function and halo mass function somewhat differ from those of BK09; when comparing the results from the adaptive runs to the standard case with fixed softening, BK09 register a deficit in the number of small-mass objects and a slight lowering in the amplitude of the two-point correlation function at small separation, whereas we notice no change in the mass function and instead an overall enhanced level of particle clustering. The simulated background cosmology differ, but this should have no effect in the analysis

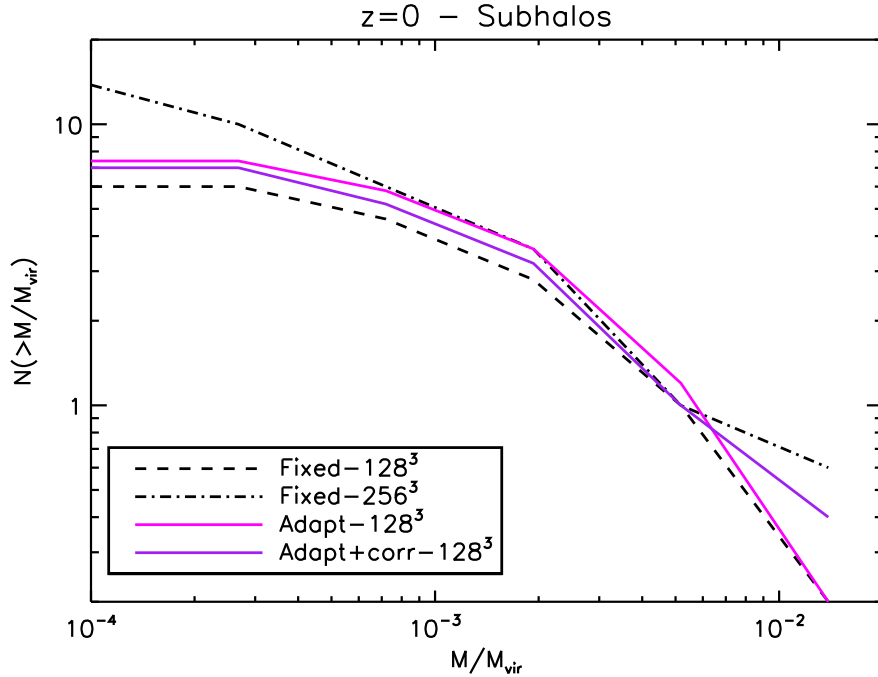


Figure 3.17: Average subhalo mass function at $z = 0$ for the five most massive halos in the simulations. The subhalo masses have been normalised to the virial mass of the host and vary from $8 \cdot 10^{11}$ to $4 \cdot 10^{13} h^{-1} M_{\odot}$, corresponding to ≈ 20 and 1000 particles, respectively.

we are interested here, which is focused on relative differences between the runs more than on absolute results of cosmological interest. The larger linking length we have employed is not at the origin of the discrepancy between the FOF mass functions, nor we think is the different evaluation of the correlation function responsible for the antithetical outcomes.

The results differ also on the more technical timing level; in our simulations we do not register substantial saving of computing time when adaptive softening is used, but we do notice increasingly better performances when the number of neighbours is increased. BK09 halve the wallclock time of the reference run when using adaptive softening with 32 neighbours and leave it unchanged when using 48. If we set an upper limit to the softening of order the splitting scale r_s (as in BK09), we progressively cut the timing when going from 32 up to 60 neighbours and leave it essentially unaltered by increasing the number even further; the minimum wallclock time is reached for $N_{ngbs} \simeq 60$ and it equals that of the simulation with fixed softening. If no upper limit is set for the softening, the wallclock times tend to increase; a minimum is reached for a number of neighbours $\simeq 40$ and it is around $\sim 30\%$ higher than in the standard run.

We ascribe the responsibility for the different behaviour of the implementation in the various aspects to the respective mother codes.

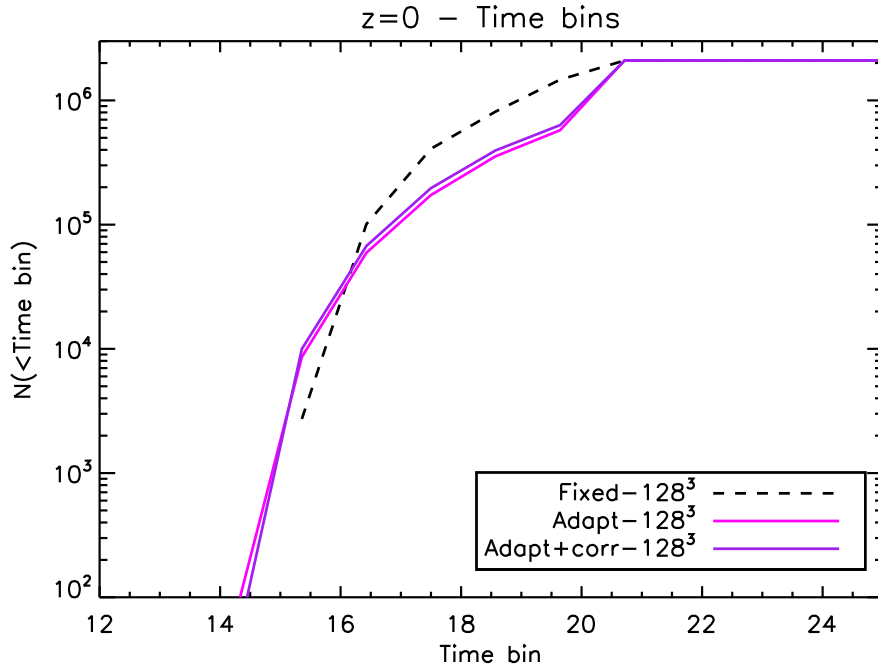


Figure 3.18: Cumulative distribution of particles in time bins at $z = 0$. The simulations were run using 29 time bins. Smaller numbers corresponds to finer intervals.

3.6 Simulating a mini version of the Millennium-II

As a more demanding application, we have investigated the effect of adaptive softening on a cosmological simulation whose small-scale clustering is known in detail by an extremely high resolution run. We have used the modified GADGET-3 on the initial conditions of the “mini”-Millennium-II simulation (hereafter mmII), a low-resolution version of the better known Millennium-II (hereafter mII, Boylan-Kolchin et al. 2009). The mmII follows the evolution of 432^3 particles (as opposed to 2160^3 particles for the mII simulation) within a box of side $100 h^{-1}$ Mpc. The underlying cosmology is a Λ CDM with parameters

$$\begin{aligned} \Omega_{\text{tot}} &= 1, \quad \Omega_m = 0.25, \quad \Omega_b = 0.045, \quad \Omega_\Lambda = 0.75, \\ h &= 0.73, \quad \sigma_8 = 0.9, \quad n_s = 1. \end{aligned} \quad (3.21)$$

This results in the particles having masses $m = 8.61 \cdot 10^{10} h^{-1} M_\odot$, a factor 125 larger than in the mII simulation, corresponding to the same mass resolution of the original Millennium simulation (Springel et al. 2005).

The initial conditions for the mmII are identical to those for the main (mII) run. This means that although they share the power spectrum of the fluctuation field, this is sampled with a factor of 125 less particles in the case of the mmII; it is reasonable to expect that the smallest perturbations are then not represented. This reverses the situation with respect to the series of simulations analysed so far in this section. Fig. 3.20 shows the ratio of the correlation functions from the mmII and “adaptive”-mmII to the original mII at four different redshifts. At high red-

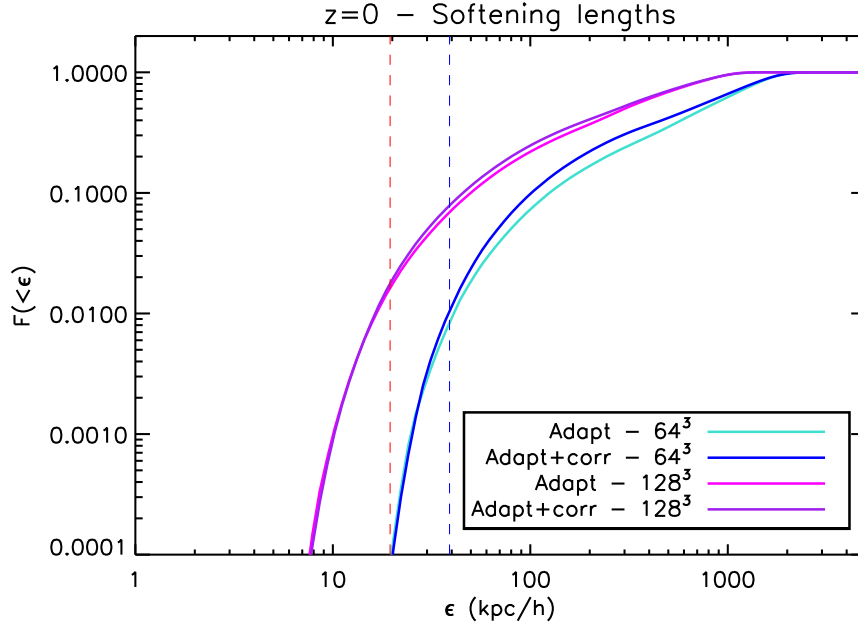


Figure 3.19: Cumulative distribution of softening lengths at $z = 0$. The softening lengths are expressed in terms of equivalent-Plummer values (ϵ). The red (blue) dashed curve represents the value of the gravitational softening in the 128^3 (64^3) “fixed” simulations.

shifts, down to approximately $z = 2$, the amplitude of $\xi(r)$ at scales smaller than $\simeq 20 h^{-1} \text{kpc}$ is lower in the “adaptive”-mmII than in the mmII; the result is not surprising though, as at those redshifts we expect almost all the particles to have softening lengths considerably larger than $\epsilon = 5 h^{-1} \text{kpc}$, the value adopted in the original mmII. At low redshifts the situation reverses, leading to a correlation function which approaches more and more tightly the results of the mII simulation; at $z = 0$ and on scales below $\sim 10 h^{-1} \text{kpc}$, the “adaptive” correlation function is a factor ~ 1.3 above that of the mmII. Fig. 3.21 shows the mass functions for the three simulations at the same redshifts as in Fig. 3.20. The halos are identified by the FOF algorithm, using $b = 0.2$ and a minimum of 20 particles. Again, the mass functions are in agreement at all redshifts.

As for the tests presented before, this demonstrates that using adaptive gravitational softening while incorporating the correction term into the equation of motion allows to resolve the smallest scales better than simulations with fixed softening at the same resolution. The results converge to what is expected from much higher resolution simulations and no substantial degrading of the results is observed at high redshifts, where the formal resolution of the simulations using adaptive gravitational softening is naturally smaller than in the simulations with fixed softening.

3.7 Conclusions

We have implemented adaptive gravitational softening in the cosmological TreePM code GADGET-3. The formalism was introduced first by Price and Monaghan (2007)




Figure 3.20: Comparison of the two-point correlation functions from the mII (black curve), mmII (green curve) and “adaptive”-mmII (purple curve) simulations at different redshifts; shown is the ratio to the results from the mII. See the text for a description of the three cases.

and features the same technique used by SPH to determine individual softening lengths for each of the simulation particles. The spatial variation of this scale requires a modification of the equation of motion, governing the evolution of the particle trajectories, in order to be consistent with the new dependencies introduced in the system’s Lagrangian.

We have applied this technique to several test cases in order to check the behaviour of the total energy when the new equation of motion is used; we then moved to the cosmological scenario and, specifically, to simulations of the large-scale structure of the Universe, where the evaluation of the effects of adaptive softening is complicated by the lack of an analytical solution to the problem.

Our main conclusions are:

- The inclusion of the correction term in the equation of motion is essential to ensure the conservation of total energy.
- In cosmological simulations of the large-scale structure a number of neighbours ≈ 60 is needed to obtain a converged estimation of the correction term.
- With such a choice we show that – on the contrary to previous claims in the literature – the adoption of adaptive softening does not lead to an under-representation of halos at low masses, if the correct equation of motion is used.
- Using the adaptive scheme effectively increases the dynamical range of cosmo-




Figure 3.21: FOF mass functions from the mII (black curve, diamonds), mmII (green curve) and “adaptive”-mmII (purple curve) simulations at different redshifts. See the text for a description of the three cases.

logical simulations, while the computational costs only mildly increase. Especially, the amplitude of the two-point correlation function at small scales and the subhalo mass function improve compared to simulations with the same number of particles and fixed gravitational softening, anticipating the results obtained in higher resolution simulations.

- The convergence of the inner density profile for the most massive object found in the simulations improves significantly when adaptive softening and the correct equation of motion are used.
- When re-simulating a low-resolution version of the Millennium-II simulation and comparing the results obtained with fixed and adaptive softening, we notice again perfect agreement in the mass functions at all times and an evolution of the “adaptive” correlation function towards the higher amplitude of the Millennium-II’s at late times.

In this chapter we have applied the method to scenarios where only one particle species was being followed. Either gas-only or dark-matter-only simulations were performed, where matter was discretised by means of equal-mass particles. The case of “hybrid” simulations, where multiple matter fields interact and co-evolve, will be investigated in the next chapter.

Adaptive gravitational softening II: the multiple-species case

Adaptive gravitational softening provides an improved representation of the clustering of particles on small scales. This has been thoroughly tested on a number of scenarios, including cosmological simulations of the evolution of dark matter throughout time (see Chapter 3). Simulations including both dark matter and gas should also considerably benefit from the adoption of the scheme: the adaptive behaviour of the resolution scale would allow to follow the collapse of gas down to scales currently unachievable in standard simulations at comparable resolution. This is of particular importance considering that cold gas and star particles are likely to clump at lengthscales below those typical chosen for the gravitational softening. With the ultimate purpose of running hydrodynamical simulations of structure formation, we have studied the behaviour of the new softening formalism on a number of test cases to calibrate its effect on computations involving multiple particle species. We discuss the strengths and weaknesses of the method, whose behaviour, this time, requires some caution.

4.1 Introduction

In the previous chapter we have investigated the effects of adaptive softening in various kinds of simulations, ranging from isolated structures to the cosmological scenario. Some of these featured the action of gravitation only (dark matter simulations), while others involved the presence of pressure forces in addition to gravity (gas-only simulations)¹. In all cases the discretisation of the matter field was done by means of particles with identical masses. As we have seen in Chapter 2, increasingly large dark-matter-only simulations have been performed in the last decades and they have been crucial in advancing our knowledge on the clustering properties of collisionless matter within an expanding background; today, we regard the subject as mostly understood and major efforts and interest are being shifted towards

¹In addition to the test presented in Sec. 3.4.3, also the classical “Evrard” test (Evrard, 1988) and “Bertschinger” test (Bertschinger, 1985) were performed. The results were not shown as they did not differ significantly from those obtained using standard, fixed softening.

modeling the coeval behaviour of the collisional, baryonic component and its evolution on both galactic and cosmological scales. These *hydrodynamical* simulations can either simply follow the evolution of dark matter and gas under the effect of gravitational and pressure forces (in this case they are referred to as *non-radiative*) or they may include simplified recipes to account for the formation of stars and related feedback mechanisms. Clearly the latter are the simulations of interest and ever-increasing effort is put in developing an appropriate treatment of the various baryonic processes known to take place, especially on galactic scales. In these simulations, therefore, one has to deal with two collisionless species (dark matter and stars) and a collisional one (gas), each characterised by a different spatial distribution; moreover, these will most likely be sampled by particles of different masses. It is therefore not surprising that, as a follow-up of the work presented in Chapter 3, we decided to investigate the behaviour of the adaptive-softening formalism in simulations with multiple particles species; it would be interesting to test whether the method could bring improvements in the treatment of the gravitational interaction also in this type of simulations. We start our investigation with a somewhat intermediate case, where two different collisionless components, representing dark matter and stars, co-exist with different spatial distributions; these simulations are the subject of Sec. 4.2. Even though stars and dark matter are represented by particles of identical mass, the presence of two components with markedly different spatial distributions justifies the presence of these simulations in this chapter, rather than in the previous. In Sec. 4.3 and 4.4 we move to the cosmological case. We consider two matter fields distributed similarly at the initial conditions and being discretised by particles of different masses; these represent two collisionless components in the former case (dark matter “1” and “2”), while in the latter they stand for one collisionless and one collisional species (dark matter and gas). It is well known that the discreteness introduced by the particle representation of the simulated system will cause spurious energy transfers when species with different masses are present, the flow taking place from the heaviest to the lightest components (Steinmetz and White, 1997; Binney and Knebe, 2002). This effect is to be expected if two-body relaxation is ruling the evolution of the system, something which is clearly unwanted in a collisionless simulation. The overall effect is that of an artificial “evaporation” of light particles from collapsed objects; considering that the component sampled by the least massive particles is generally the gaseous one, this translates into a systematic and artificial underrepresentation of the baryonic component within dark matter halos. The effect is particularly strong in objects of small-intermediate masses, whose potential well is not deep enough to keep the gas particles bound notwithstanding this artificial heating. Obviously, the spurious loss of baryons from dark matter halos has direct consequences on the predictions regarding the amount of gas able to cool and form stars in these objects. Increasing the gravitational softening length in these simulations helps alleviating the problem, at the expense of a loss in resolution. Considering that the adaptive method induces an overall increase of the individual softening lengths whilst enhancing the resolution, we expect it to outperform the standard, fixed softening approach even more in these simulations. Besides the attenuation of the spurious, two-body

heating of gravitational origin, we expected a further improvement from the use of adaptive softening; this concerns the behaviour of gas, regardless of the mass of its particles relative to the dark matter ones. In calculations including self-gravity of the gaseous component, such as cosmological simulations of structure formation, the relative size of gravitational softening and hydrodynamic smoothing length can have tangible consequences on the results of the simulation at those spatial scales. As shown by Bate and Burkert (1997) and Sommer-Larsen et al. (1998), the gravitational collapse of gas is inhibited if the resolution in the gravitational force is poorer than in pressure forces (as is the case when the softening length is larger than the smoothing length) and can instead be artificially induced in the opposite regime (when the softening length is smaller than the smoothing length). Bate and Burkert (1997) conclude that the resolution of the simulation is given by the larger of the two scales and that any results below is of no physical meaning. They also remark that having the scales of gravitational and hydrodynamical resolutions set equal would be the ideal approach to the problem. Considering that cold gas is expected to clump on scales below those of dark matter and definitely below those typically chosen for the gravitational softening, having this quantity fixed results in a waste of resolution, heavily affecting the behaviour of the gaseous component on small scales. A gravitational softening which shrinks during collapse, in the same way as the hydrodynamical smoothing length, would allow the behaviour of gas to be followed to a greater extent. The expectation is therefore that, thank to the adaptivity of the softening scale, not only will the segregation of the gaseous component be considerably attenuated, but also its collapse will be followed down to scales which are currently unachievable in standard simulations at comparable mass resolution; this would ultimately provide a more reliable representation of the properties of the baryonic component within galaxy-like substructures. We will show how these expectations are, in fact, only partially fulfilled; in at least one of the two adaptive approaches, the extension of the method to hydrodynamical simulations comes with subtle problems that need to be thoroughly understood and kept under control. Finally, in Sec. 4.5 we show an intriguing result obtained in a simple simulation of dynamical friction, where a massive body orbiting within a sea of light particles progressively decays towards the density peak of their distribution.

4.2 A spherical galaxy with dark matter and stars

In this section we investigate the evolution of a galaxy-like object, made of a concentrated bulge of stars embedded within a more extended dark matter halo. The distribution of the two components is spherically symmetric and follows a Hernquist profile (see Sec. 3.4.2); the scale radii differ, being 1 kpc for the bulge and 5 kpc for the halo. The object extend for around 400 kpc and consists of 1.1×10^6 equal-mass particles, 10^5 of them contributed by the stellar component. We let the object evolve passively, adopting different choices for the gravitational softening, and monitored the changes in its internal structure. The integration spanned five gigayears, i.e. a few dynamical times and a small fraction of the relaxation time of

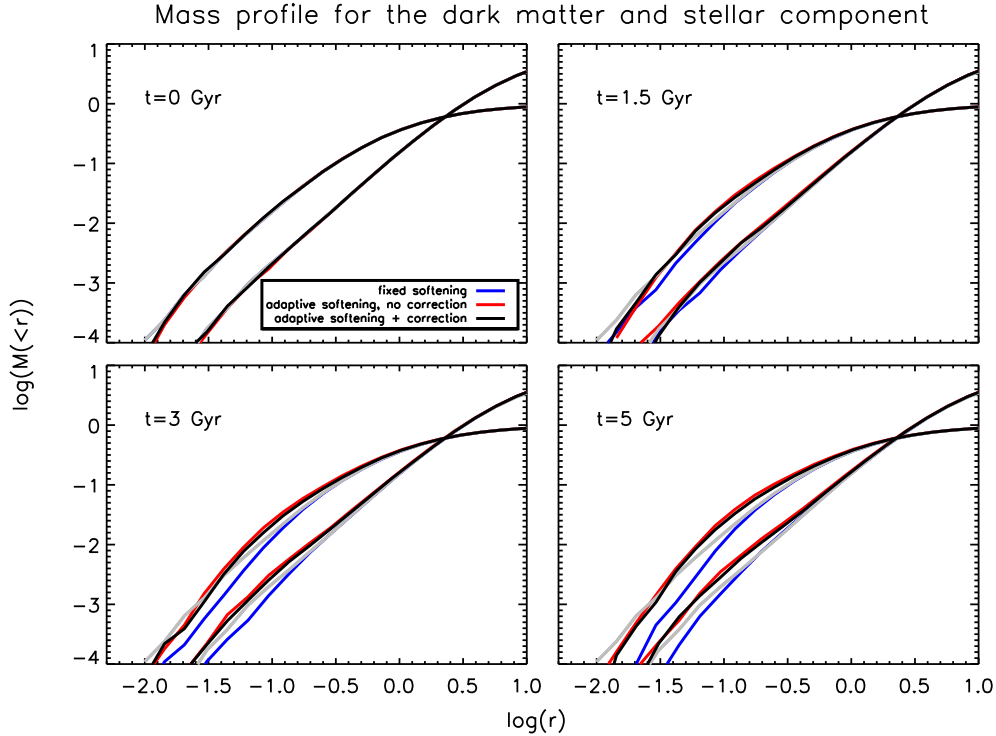


Figure 4.1: Cumulative mass profile for the two components of a spherical system in equilibrium and evolving passively in isolation. The most concentrated component represents a stellar bulge embedded within a dark matter halo. The former is sampled by 10^5 particles, while the latter by ten times as many. Both species follow a Hernquist profile at the initial conditions. The blue curve represents the results obtained with fixed softening, set to the same value for both components; the black (red) curve shows the behaviour of the simulation using adaptive softening with (without) correction of the equation of motion (in both cases $N_{ngbs} = 60$). In each panel, the gray lines represent the distribution of the two components at the initial condition and serve as a reference for the expected result.

the object; in other words, we do not expect changes in its structure from the initial condition setup. Fig. 4.1 shows the cumulative mass profile of the stellar and dark matter components at four different times during the integration. The gray curves represent the distribution at the initial condition and serve as a reference for the expected results at each time. The blue lines show the profiles obtained using fixed softening; this was set to 20 parsecs, around $1/20$ of the mean interparticle separation within the half-mass radius. This choice is representative of the range of values used in these simulations (Michael Hilz, private communication). Both curves, for the bulge and halo component, show a progressive detachment from the original profiles: the inner cusp is not maintained in time and particles are lost from the central region. Conversely, both the adaptive softening approaches, with and without the correction to the equation of motion (black and red curves, respectively), provide a remarkably good behaviour at all times during the simulation; the bulge and halo profiles are maintained close to the original down to the innermost regions. The origin of such a different behaviour lies, obviously, in the distribution of

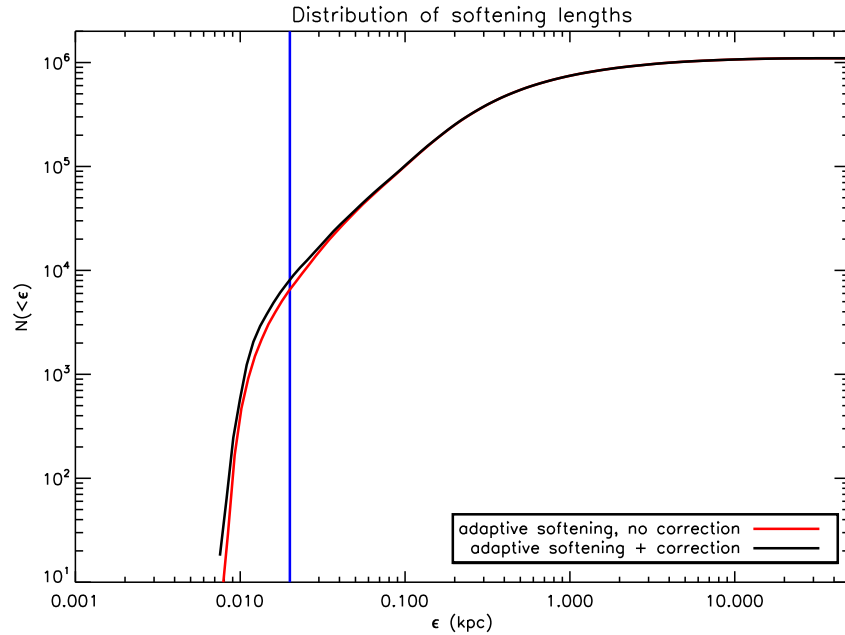


Figure 4.2: Cumulative distribution of softening lengths for the simulation of the spherical system with stellar bulge and dark halo. The results are from the last snapshot of the simulation, after five gigayears of passive evolution. The vertical, blue line marks the softening length chosen in the standard simulation. The black (red) curve shows the behaviour of the simulations using adaptive softening with (without) correction of the equation of motion (in both cases $N_{ngbs} = 60$). In these cases, for consistency, the Plummer-equivalent softening is plotted.

individual softening lengths. As shown in Fig. 4.2, in the last snapshot of the simulation particles have softenings ranging from 0.07 to several tens of kiloparsecs. Around ten thousand particles have this individual scale set smaller than the value adopted in the simulations using fixed softening (marked by the blue, vertical line). Even if this is just 1% of the total number of particles, it is a big enough value to make a difference in the inner distribution of the object; indeed, the regions where the mass profiles in Fig. 4.1 deviate from the expectation correspond to an enclosed mass of few per cent the total mass. We therefore remark once more that the use of adaptive softening increases the force resolution in regions where, normally, the effect of the gravitational interaction would be washed out by the use of a too large, fixed softening length.

We also performed simulations of mergers between two of these objects. Varying the softening approach, we found no significant difference in the mass distribution of the objects before, during and just after the merger event took place. We expect these to manifest in the late evolution of the remnant, for the same reasons outlined above.

4.3 Cosmological simulations I: two dark matter species

In this section we investigate the effects of adaptive softening on cosmological simulations featuring two collisionless fields sampled by particles of different masses. Hereafter we will refer to this set of simulations as “DM1+DM2”. The initial conditions were generated starting from those used to perform the simulations at intermediate resolution described in Sec. 3.5; this means that we are considering the same periodic box of $100 h^{-1}$ Mpc side-length in a Λ CDM cosmology, although we now discretise its mass by means of twice as many particles, i.e. 2×128^3 . In order to do this, we split each of the original particles into two distinct components and displace them in order to maintain the centre-of-mass position and velocity unchanged. By doing so, the global fluctuation field arising from the distribution of particles is not affected and remains a faithful representation of the assumed cosmology. The original mass is split between the pair with a ratio given roughly by $\Omega_b/\Omega_{dm} \approx 0.15$; this means that, even though we are dealing with collisionless species, the total mass contributed by each of them is of the order the total expected dark matter and baryonic mass within the cosmological volume under consideration. Starting from these new initial conditions, we ran several simulations with the different softening approaches. On each of the fifteen snapshots saved between redshift five and zero we ran the halo and subhalo finder SUBFIND to identify gravitationally bound structures in the simulations. The expectation is that the mass of these halos, regardless of its value, should be contributed by both particle species according to the ratio between their masses; alternatively, we can say that we expect a roughly equal number of particles from both species within any object identified in the simulation. We do expect significant scatter, especially for small halos made of a handful of particles; the mean behaviour should anyway follow the expectation, if the simulation is collisionless and insensitive to the underlying discretisation. If, conversely, two-body relaxation were at work we would register a systematic shortage of light particles as a consequence of the induced mass segregation; this would affect mainly small objects, incapable of counteracting this spurious effect with their gravitational attraction. Fig. 4.3 shows the results of the various simulations, when confronted with this test. The different lines represent the median fraction of light particles in units of the expected value, plotted as a function of the mean number of particles within the halos in the corresponding bin. The blue line shows the case for fixed softening; in this simulation the value is set to $1/20$ of the mean interparticle separation, double the size used for the runs of Sec. 3.5 and corresponding to $39 h^{-1}$ kpc. Clear signs of segregation are evident at the low-mass end, where the curve progressively deviates from the expectation, marked in orange. Even such a large softening is not enough to alleviate collisionality in the simulation. The set of red curves shows the results of the simulations performed with adaptive softening and no correction of the equation of motion; each line correspond to a different choice in the number of neighbours used to set the softening scale. The lines are clearly indistinguishable and all show a remarkable agreement with the expectation. This result is very promising, but unfortunately it is not fully

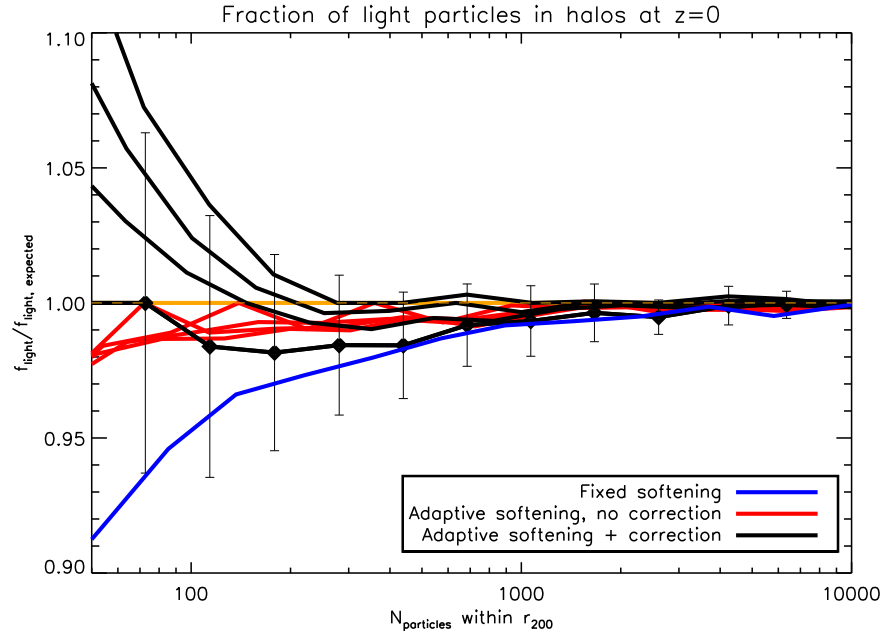


Figure 4.3: Fraction of light particles in collapsed objects for the suite of cosmological simulations with two collisionless species (DM1+DM2 set). Plotted is the median fraction of light particles in units of the expected value, given by the global ratio in the simulated box. The results are given as a function of the mean number of particles within the virial radius of the objects in the bin. The orange, dashed line represents the expectation, namely a fraction equal to the global value for the box regardless of the size of the halo. In blue are the results from the simulation using fixed softening. The black curves show the case where adaptive softening is used, along with the correction term; the results obtained without correction term are instead displayed in red. Four lines are overplotted for each of the adaptive approaches, corresponding to different choices in the number of neighbours (90,120,150,180). The red curves are almost indistinguishable; the black ones are not, though, and the higher the number of neighbours, the higher the low-mass-end fraction. Error bars, corresponding to one standard deviation, are shown only for the “adaptive+correction” case with 90 neighbours, to avoid overcrowding. The results are shown at redshift zero.

maintained when adaptive softening is used along with the new equation of motion. The results for those simulations are shown in black. Even though all curves lie closer to the expectation with respect to the blue one, at least for halos containing more than ≈ 200 particles (corresponding to $\approx 4 \cdot 10^{12} h^{-1} M_{\odot}$, in this case), they diverge at the low-mass end; not only some dependence of the results on the number of neighbours is evident, but also the increase in the fraction of light particles above the expectation is puzzling. The spread in the results (shown in one case only, to avoid overcrowding) is large, especially at the low-mass end, but still the finding requires some further investigation. We start by wondering how the trend, observed at redshift zero, sets in with time. Figures 4.4, 4.5 and 4.6 show the redshift evolution in the fraction of light particles in halos for the three softening approaches. The curves span the evolution from redshift around five down to redshift zero and are colour coded according to the rainbow spectrum (violet for redshift five, red for

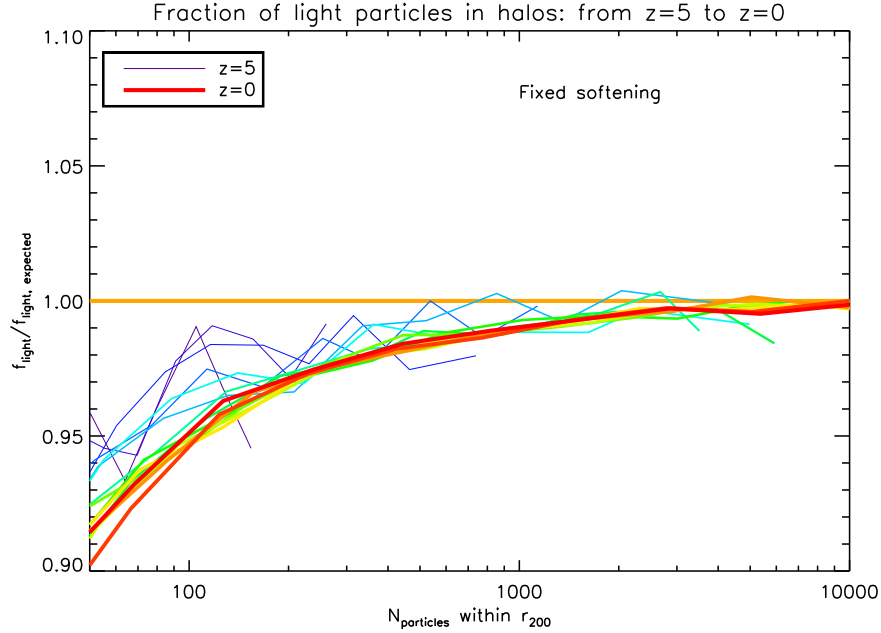


Figure 4.4: Fraction of light particles in collapsed objects for the suite of cosmological simulations with two collisionless species (DM1+DM2 set). Plotted is the median fraction of light particles in units of the expected value, given by the global ratio in the simulated box. The results are given as a function of the mean number of particles within the virial radius of the objects in the bin. Shown is the redshift evolution of this fraction for the simulation employing fixed softening. Starting from redshift around five (violet, thin line), fifteen curves show the results at different times, down to redshift zero (red, thick curve). The orange, horizontal line represents the expectation, namely a fraction equal to the global value for the box regardless of the size of the halo.

redshift zero). In the simulation using fixed softening (Fig. 4.4) we see that the depletion of light particles occurs progressively in time, even though already at high redshifts this component results underrepresented in halos of small-intermediate mass. When adaptive softening is used without correcting the equation of motion (Fig. 4.5), no redshift evolution is registered: all curves are in remarkable agreement with the expectation (although the scatter, especially at high redshift, is very large). When analysing the other adaptive approach (Fig. 4.6), the time dependence shows up again in a similar fashion as for the fixed-softening case, namely: at a given, small-intermediate halo mass, the fraction of light particles is systematically higher at earlier times. The difference is that, in this case, the values start out higher than expected. The results shown for the adaptive-softening case are those obtained with $N_{ngbs} = 90$. We already saw (Fig. 4.3) that varying this parameter has no effects on the results obtained when the old equation of motion is maintained and this remains true also when analysing the redshift evolution. Increasing N_{ngbs} in the other case, instead, does influence the presence of light particles in small halos: the higher N_{ngbs} , the higher their fraction. In terms of redshift evolution we register a mild reduction in the spread among the lines, when increasing N_{ngbs} . Basically, we are left with the following facts: (i) the use of adaptive softening with

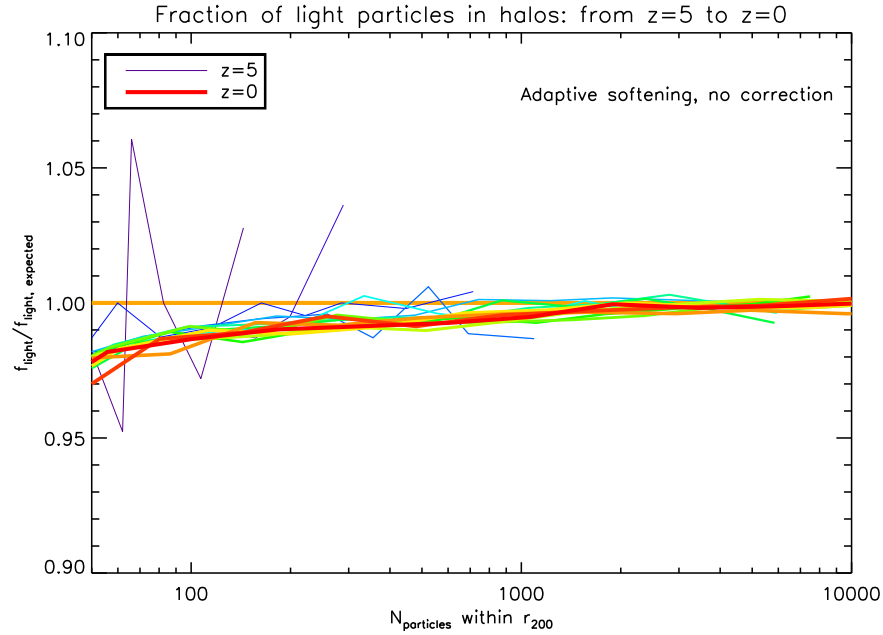


Figure 4.5: Fraction of light particles in collapsed objects for the suite of cosmological simulations with two collisionless species (DM1+DM2 set). Plotted is the median fraction of light particles in units of the expected value, given by the global ratio in the simulated box. The results are given as a function of the mean number of particles within the virial radius of the objects in the bin. Shown is the redshift evolution of this fraction for the simulation employing adaptive softening without correction term and with $N_{ngbs} = 90$. Starting from redshift around five (violet, thin line), fifteen curves show the results at different times, down to redshift zero (red, thick curve). The orange, horizontal line represents the expectation, namely a fraction equal to the global value for the box regardless of the size of the halo.

the new equation of motion induces an excess of light particles within low-mass, collapsed objects; *(ii)* the importance of this effects shows a clear dependence on the number of neighbours adopted to evaluate individual softenings; *(iii)* the excess of light particles sets in at high redshifts and tends to attenuate with time. The progressive “evaporation” of light particles from halos as time goes by could well be related to two-body effects: even though these seem to disappear when adaptive softening is used without correcting the equation of motion, their imprint may still be seen in the fully-consistent, adaptive approach. The fact that this effect reduces in magnitude when the number of neighbours is raised (and therefore the softening scale is larger) is in agreement with this interpretation. The most puzzling finding is definitely the “inverted segregation” causing excess of light particles within low-mass halos. In interpreting this unexpected behaviour, we are aided by a striking piece of evidence: the upturn of the black curves in Fig. 4.3 occurs at $N_{particles} \lesssim N_{ngbs}$. This means that the halos affected by a surplus of light particles are those made of a total number of particles less than the number of neighbours used to set the individual softenings.

In order to confirm this and further investigate the finding, we performed some

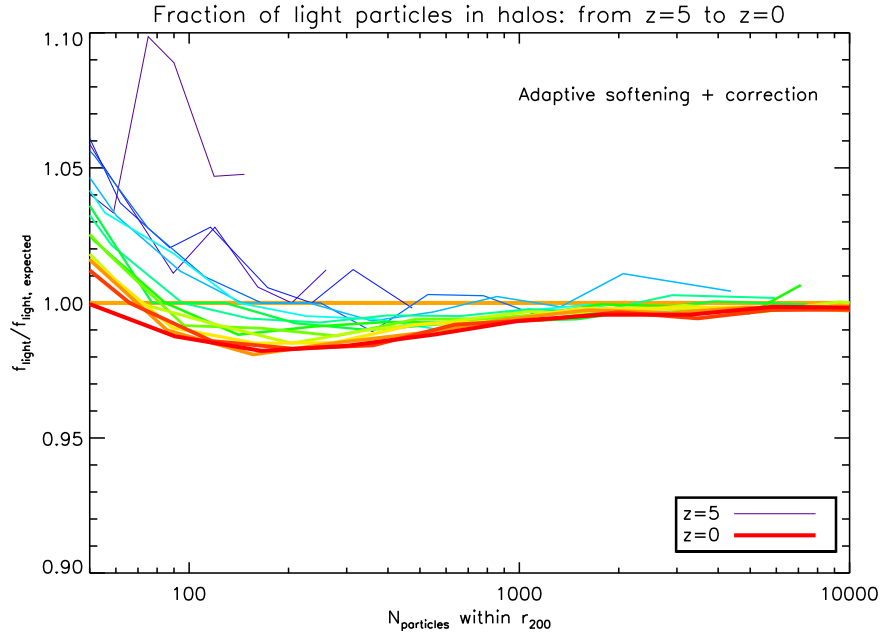


Figure 4.6: Fraction of light particles in collapsed objects for the suite of cosmological simulations with two collisionless species (DM1+DM2 set). Plotted is the median fraction of light particles in units of the expected value, given by the global ratio in the simulated box. The results are given as a function of the mean number of particles within the virial radius of the objects in the bin. Shown is the redshift evolution of this fraction for the simulation employing adaptive softening with correction term and $N_{ngbs} = 90$. Starting from redshift around five (violet, thin line), fifteen curves show the results at different times, down to redshift zero (red, thick curve). The orange, horizontal line represents the expectation, namely a fraction equal to the global value for the box regardless of the size of the halo.

targeted simulations of isolated halos. We set up initial conditions consisting of two superimposed Plummer spheres (see Sec. 3.4.1), each made of the same number of particles, but differing in total mass; the mass ratio between the spheres and, hence, between the particles sampling them, is chosen to match that of the DM1+DM2 simulations. The object is let evolve in equilibrium and the behaviour of the two different species is monitored. We performed several simulations, varying the number of particles in the spheres, the number of neighbours, the softening approach. Here we show just one representative example. We consider two superimposed Plummer distributions in equilibrium, each sampled by 500 particles; we evolve the system in time and check the behaviour of the cumulative mass profile for the two species. In simulations with fixed softening (not shown) we register the evaporation of the light component as a progressive increase of its half-mass radius. This effect, as expected, is strongly dependent on the value of ϵ and vanishes for $\epsilon \approx 4$ (for the definition of the units, see Sec. 3.4.1); this correspond to the radius containing $\approx 90\%$ of the initial mass. In simulations with adaptive softening and corrected equation of motion, we recover the same puzzling behaviour already shown by the black curves of Fig. 4.3. Fig. 4.7 shows the cumulative mass profile for the two species at different times. The blue curves represent the initial con-

dition, where the two components share the same spatial distribution. The green and red sets show the situation after 10 and 20 dynamical times (corresponding to ≈ 1 and ≈ 2 relaxation times for the initial setup), respectively. The system has been evolved using as many as 5000 neighbours to set the individual softening². This number has intentionally been chosen large to emphasise the final effect. The whole system undergoes an initial expansion, due to an underestimation of the real potential; this is caused by the excessively large softenings resulting from the choice of 5000 neighbours. Both populations are affected similarly and we can ignore this feature in our discussion. What is interesting is the relative behaviour of the light and heavy species, showed by the dashed and solid curves, respectively: the former moves inward and its half-mass radius reduces to two thirds that of the other component. By running several other simulations of the same system, varying N_{ngbs} , we found that this behaviour sets in at $N_{ngbs} > 1000$. This confirms that the anomalous baryon fractions registered at the low-mass end of the mass function in the cosmological, DM1+DM2 simulations may indeed concern objects where $N_{particles} \lesssim N_{ngbs}$. Although the reasons at the origin of this behaviour are still not clear, we can now at least isolate and confine the problem to certain scales, below which results cannot be trusted. As a comparison, Fig. 4.8 shows the results obtained with adaptive softening and old equation of motion; the initial expansion, due to the large softening lengths induced by the choice $N_{ngbs} = 5000$, affects both component; these evolve similarly, though, and show no sign of segregation.

From the results we have discussed so far in this section we would have to conclude that the use of adaptive softening does indeed bring a remarkable improvement in the treatment of multiple particle species in simulations, but that the correction introduced in the equation of motion leads to unwanted effects which, albeit small, make the use of the old equation of motion preferable. However, we need to remind the reader that this use of adaptive softening (keeping the old law for the gravitational acceleration) has its own, unwanted side-effects. As we have already discussed in the previous chapter, this approach leads to a decrease in the number of halos identified within the box (see Sec. 3.5.2). This affects essentially the low-mass end of the mass function and is due to the softening within these objects being large enough to wash out their internal structure, to the point that they are not detected by the halo-finding algorithm anymore. This effect remains visible in the DM1+DM2 simulations, as shown by Fig. 4.9. Plotted is the number density of bound halos with masses greater than the corresponding value given by the x axis. It is evident that the shape of the red curves (corresponding to the simulations with adaptive softening and old equation of motion) at the low-mass end: (i) depends on the adopted number of neighbours and (ii) is flatter than in all other simulations. The decrease in the identified objects can be as low as 48% for the simulation with 180 neighbours and a milder 30% in the run with 90 neighbours. The black curves (corresponding to adaptive softening and new equation of motion) show, instead, remarkable agreement among themselves and with the results obtained using fixed softening. It is tempting to relate the excellent behaviour in reproducing the ex-

²This is possible even in the presence of 1000 particles (see Eq. 3.14).

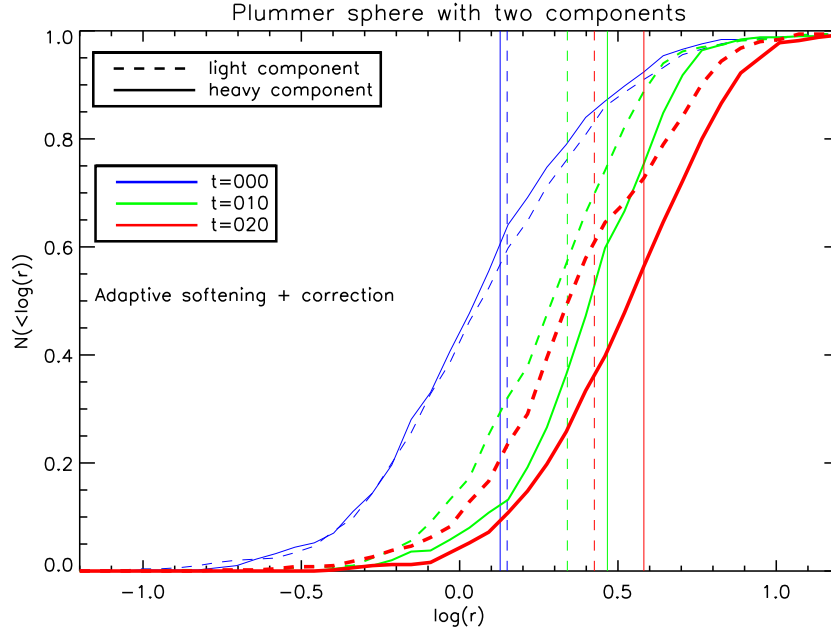


Figure 4.7: Cumulative mass profiles for the heavy and light component of the Plummer sphere. The object consists of the superposition of two 500-particle Plummer halos, sampled by particles of different masses (mass ratio ≈ 0.15). The curves represent the time evolution of the mass profiles when the object is let evolve passively. Solid (dashed) curves show the results for the heavy (light) component. The different colours correspond to different times (units are the same as in Sec. 3.4.1). The vertical lines mark the mean radial distance from the centre for each component at each time. The simulation was performed using adaptive softening with the correct equation of motion. Individual softenings were set using 5000 neighbours.

pected fraction of light particles in small halos, provided by the simulations with adaptive softening and old equation of motion, to the lack of objects that these simulations produce at the low-mass end of the mass function. One could wonder whether the objects we are missing represented a biased population, namely one characterised by a low presence of light particles; in this case, their inclusion would change the results of Fig. 4.3 in worse. We have good reason to think this is not the case: the improved behaviour shown by the red curves begins at halo masses much higher than those affected by and underrepresentation in the mass function. To give an idea, 100 particles in Fig. 4.3 correspond to $\approx 10^{12.3} M_{\odot}$ in Fig. 4.9; in this mass regime there is virtually no difference in the mass functions and, yet, the much improved behaviour of the simulations with adaptive softening is already in place. The fact that the two effects are unrelated is supported by another piece of evidence: in simulations with 60 neighbours (not shown), the fraction of light particles in halos remains excellent and indistinguishable from the results obtained with a higher N_{ngbs} , while the shortage of collapsed objects reduces to 16%.

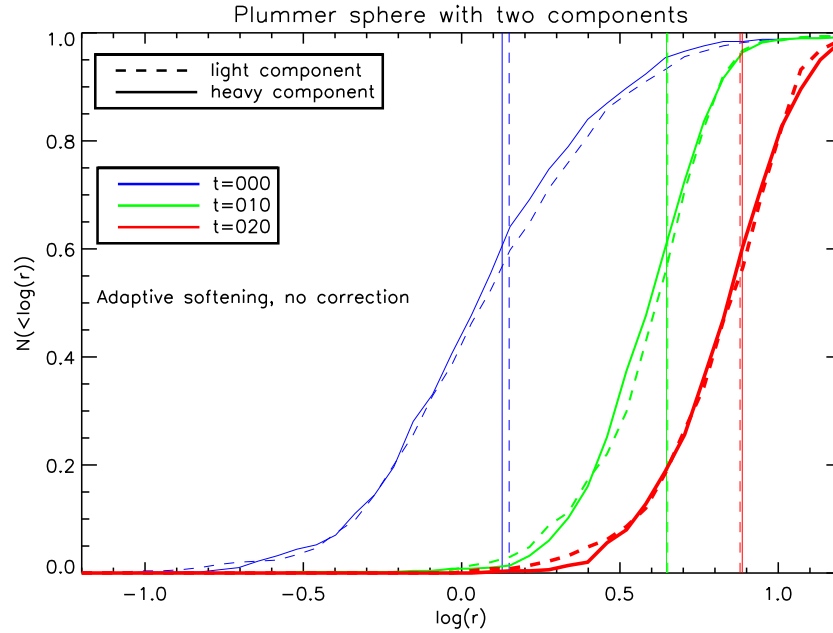


Figure 4.8: Cumulative mass profiles for the heavy and light component of the Plummer sphere. The object consists of the superposition of two 500-particle Plummer halos, sampled by particles of different masses (mass ratio ≈ 0.15). The curves represent the time evolution of the mass profiles when the object is let evolve passively. Solid (dashed) curves show the results for the heavy (light) component. The different colours correspond to different times (units are the same as in Sec. 3.4.1). The vertical lines mark the mean radial distance from the centre for each component at each time. The simulation was performed using adaptive softening and old equation of motion. Individual softenings were set using 5000 neighbours.

4.4 Cosmological simulations II: dark matter and gas

In this section we deal with non-radiative, hydrodynamical simulations following the coeval evolution of dark matter and gas particles. We will refer to this set of simulations as “DM+GAS”. The initial conditions are identical to those used for the DM1+DM2 set, but for the fact that the second species, sampled by the lightest particles, now represents a gaseous component subject to both gravity and pressure forces. Having thoroughly analysed the results of the collisionless equivalent of these simulations in the previous section, we can already guess what the results of this new study will be like. We expect that pressure forces will aid the spurious evaporation of structures, leading to even lower baryon fractions in halos than what is expected from the cosmic value (i.e. the ratio taken over the whole box). We expect similar relative behaviours from the three softening approaches as found in Sec. 4.3. This is indeed true for the resulting mass functions, which we decided to omit. We discuss, instead, the baryon fractions as a function of halo size. These are shown in Fig. 4.10 at four different redshifts and colour-coded as in Figures 4.3 and 4.9. Again, the behaviour of adaptive softening without correction term (red curves) is remarkably good at all redshifts. We show two curves for this specific




Figure 4.9: Mass function of collapsed objects for the suite of cosmological simulations with two collisionless species (DM1+DM2 set). Plotted is the number density of objects with virial masses greater than the corresponding value on the x-axis. In blue are the results from the simulation using fixed softening. The black curves show the case where adaptive softening is used, along with the correction term; the results obtained without correction term are instead displayed in red. Four lines are overplotted for each of the adaptive approaches, corresponding to different choices in the number of neighbours (90,120,150,180). The black curves are almost indistinguishable; the red ones are not, though, and the higher the number of neighbours, the lower their low-mass-end amplitude. The results are shown at redshift zero.

case, corresponding to the choice of 60 and 120 neighbours: the agreement is striking, confirming the insensitivity of this approach to the parameter N_{ngbs} . The simulation using adaptive softening with correction term (black curve) provides baryon fractions comparable to those obtained in the simulation with fixed softening (blue curve) until $z \approx 2$; at later times the behaviour worsens and by the end of the run the presence of gas in halos is severely underestimated, with respect to the other simulations. Varying the number of neighbours does introduce variations, most noticeably in the tail of the curve (note that the upturn at $N_{particles} \lesssim N_{ngbs}$ is still present in this simulation). However, since the softening distribution for the two adaptive approaches is similar, at a given number of neighbours, we attribute the differences in the baryon fractions to some dynamical effect introduced by the correction term. In general, the fraction of light, gaseous particles bound in halos confirms to be lower than what found in the previous section for the collisionless equivalent of these simulations. Indeed, taking the runs with fixed softening as a reference, the presence of light particles in the smallest identified halos drops from 95% to 65%.

Although we acknowledge the behaviour of the simulations with adaptive soften-

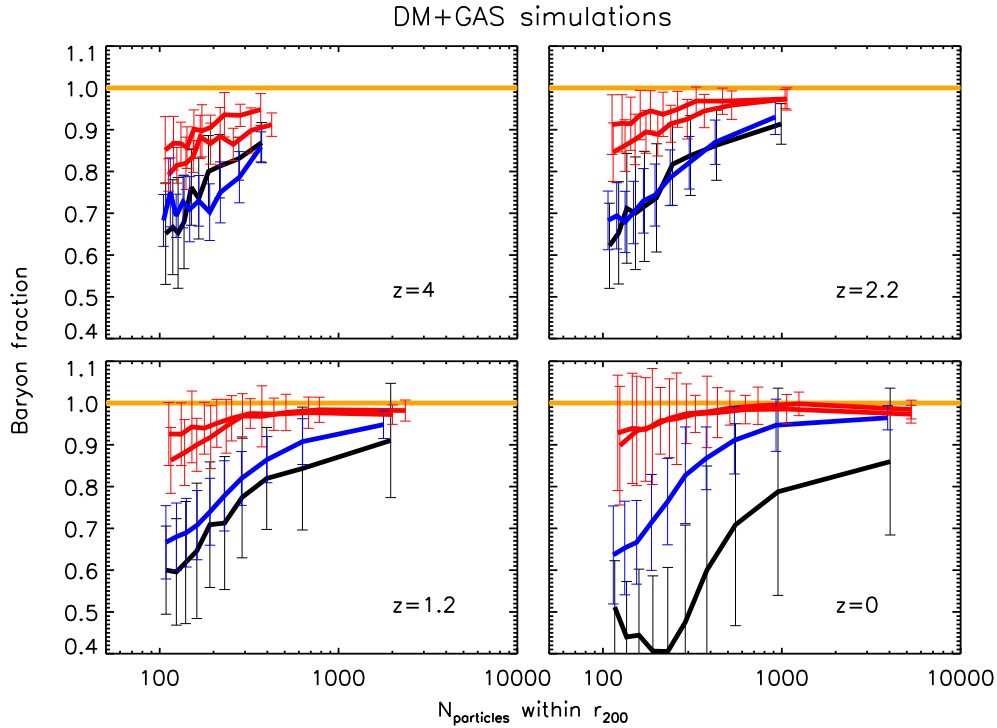


Figure 4.10: Baryon fractions in collapsed objects for the suite of cosmological simulations with dark matter and gas (DM+GAS set). Plotted is the median fraction of gas particles in units of the expected value, given by the global ratio in the simulated box. The results are given as a function of the mean number of particles within the virial radius of the objects in the bin. The orange line represents the expectation, namely a fraction equal to the global value for the box regardless of the size of the halo. In blue are the results from the simulation using fixed softening. In black is the case where adaptive softening is used, along with the correction term (one curve, corresponding to $N_{ngbs} = 120$); the results obtained without correction term are instead displayed in red (two curves, corresponding to $N_{ngbs} = 60, 120$). The error bars correspond to one standard deviation.

ing and correction of the equation of motion to be somehow problematic, we have to admit that the comparison with the standard simulation, using fixed softening, has not been entirely fair. Nowadays, most of these hydrodynamical simulations are, in fact, performed with a “trick”, adopted to compromise between the need of large softenings at high redshifts, to avoid the evaporation of gas, and the need of resolution at low redshifts, when structures collapse to form dense halos. This solution consists in preventing the softening scale to exceed a certain value, given in physical units. From the beginning of the simulation the softening length is kept constant, in comoving units, until its value, in physical units, reaches a certain threshold; from this point onwards the softening scale is kept fixed to this physical size and its comoving equivalent shrinks. The change is applied to all particles at the same time; the exact moment depends on the chosen threshold.

Before closing this topic, we would like to compare the behaviour of the adaptive-softening approach to the results obtained with these more realistic settings. In

doing this, we change setup and apply the adaptive method on a different simulation. This is given by the Box3 of the Magneticum suite³, a simulation of a $128 h^{-1} \text{ Mpc}$ -side cube with WMAP7 cosmology (Komatsu et al., 2011). The version we are using discretises dark matter and gas components by means of 216^3 particles each; the mass ratio is given by $\Omega_b/\Omega_{dm} \approx 0.2$ and is around 30% higher than in the DM+GAS simulations. The improved resolution allows the identification of smaller halos, with masses down to $\approx 10^{10} h^{-1} M_\odot$, leading to a total number of bound objects four times bigger than those found in the DM+GAS runs. The simulation with fixed softening was performed using $\epsilon_{com} = 30 h^{-1} \text{ kpc}$ and $\epsilon_{max,phys} = 10 h^{-1} \text{ kpc}$. This means that the softening is initially set to the value given by ϵ_{com} and that when $\epsilon_{phys} = a \cdot \epsilon_{com} = \epsilon_{max,phys}$ it gets changed into $\epsilon_{max,phys}$ until the end of the simulation. Translated into fractions of the mean interparticle separation, these numbers mean that the softening length is set to $\approx 1/20$ of this scale (as in the DM+GAS simulations) until around $z = 2$ (the time where the switch happens) and that it progressively reduces afterwards, reaching $1/50$ of the mean interparticle spacing at redshift zero. As we have had several occasions to assess the robust behaviour of the simulations with adaptive softening and old equation of motion, here we just consider the other adaptive approach. Fig. 4.11 shows the baryon fractions recovered for these simulations at four different redshifts. We chose redshifts lower than in Fig. 4.10 to focus on the time range where the behaviour of the adaptive, DM+GAS simulation started to worsen. At all times, the adaptive run (black curve) performs better than the standard one (blue curve). The baryon fractions in the simulation with fixed softening are only slightly lower than those of the DM+GAS set, given the same number of particles in the halo (as a reference: in Fig. 4.11 a mass of $\approx 8 \times 10^{11} h^{-1} M_\odot$ corresponds to around one hundred particles). The effects of a more aggressive softening prescription are partially balanced by the increase in the mass ratio between gas and dark matter particles. The difference between the two cases is then due to an improvement in the behaviour of the simulations with adaptive softening, rather than a deterioration of the results of the fixed-softening run. We regard the more moderate difference between the particle masses to be responsible for this behaviour; the mass ratio enters the evaluation of the correction term (see Eq. 3.8) and can be responsible of numerical effects associated to its evaluation. This example was shown to remark that which one of the two approaches performs best, whether it is fixed softening or adaptive softening with correction, might depend on the system under consideration and, more specifically, on details as the mass ratio between particles and on the recipes adopted to set the fixed softening scale.

4.5 A simulation of dynamical friction

In this last section we show an interesting result obtained in a simulation following the evolution of a massive body orbiting within a sea of light particles of mass m . The expectation in this case is that, in the limit $M_{body} \gg m$, the object experi-

³<http://www.mpa-garching.mpg.de/~kdolag/Simulations/>

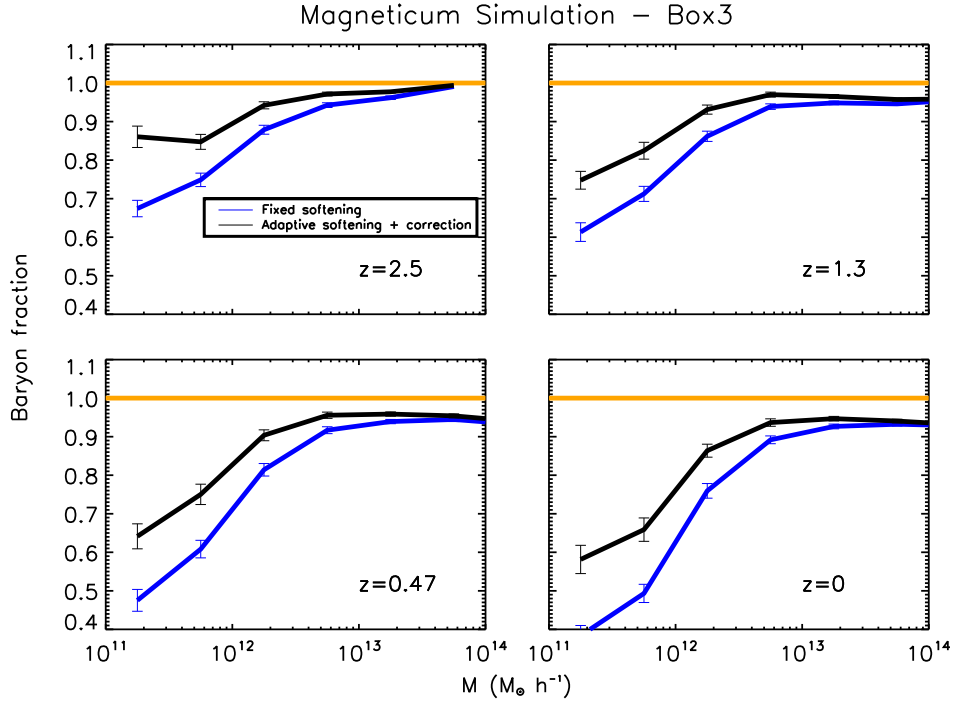


Figure 4.11: Baryon fractions in collapsed objects extracted from simulations of the Box3 from the Magneticum suite. Plotted is the median fraction of gas particles in units of the expected value, given by the global ratio in the simulated box. The results are given as a function of the mean virial mass of the objects in the bin. The orange line represents the expectation, namely a fraction equal to the global value for the box regardless of the size of the halo. In blue are the results from the simulation using fixed softening. In black is the case where adaptive softening is used, along with the correction term (one curve, corresponding to $N_{ngbs} = 128$). The error bars correspond to one standard deviation.

ences *dynamical friction*, a force leading to the transfer of energy and momentum to the field particles (see, e.g. Binney and Tremaine, 1987; Mo et al., 2010). This arises from the cumulative effect of two-body encounters between the massive object and the light particles surrounding it. The system is driven towards a state of equipartition, where each population is characterised by the same mean kinetic energy per particle ($m_x \langle v_x^2 \rangle$); in this process, the more massive components lose energy and momentum to lighter particles, until the balance is reached. In the case under consideration, the effect translates into a progressive decay of the orbit followed by the massive body. The details of the dynamical evolution of the object depend on the specifics of the system and are approximated by the Chandrasekhar formula (Chandrasekhar, 1943):

$$\mathbf{F}_{df} = -4\pi \left(\frac{GM_{body}}{v_{body}} \right)^2 \ln \Lambda \rho(< v_{body}) \frac{\mathbf{v}_{body}}{v_{body}}, \quad (4.1)$$

where \mathbf{F}_{df} is the dynamical friction force acting on the body, $\rho(< v_{body})$ is the density of field particles with speed less than v_{body} and $\ln \Lambda$ is the Coulomb logarithm. This

last quantity is not well constrained and represents a problematic aspect of Eq. 4.1. In estimates of orbital decay rates Λ is taken to be the ratio of the host halo mass to that of the orbiting body, i.e.

$$\ln \Lambda \approx \ln \left(\frac{M_{halo}}{M_{body}} \right). \quad (4.2)$$

From a first look at Eq. 4.1 we learn that the drag force has a strong dependence on M_{body} , with more massive objects being more heavily affected by dynamical friction. Also, from Eq. 4.2 we can say that, given M_{body} , the exerted force is larger in more massive hosts.

Chandrasekhar's formula was derived under a stringent set of assumptions, none of them applicable to a real physical system; caution is therefore needed in interpreting its prediction. The first assumption lies in regarding all objects as point particles; this is clearly a problem when considering an actual physical system, but not when comparing to the results of a simulation. The second approximation consists in neglecting any evolution in the density distribution of the underlying system, due to the presence of the massive body; this is clearly not the case neither in real systems nor in simulations. Finally, it is assumed that the distribution of the field particle is infinite, homogeneous and isotropic. Intimately related to this last assumption is the need to introduce the rather obscure Coulomb logarithm term in the derivation of the final formula; indeed, this allows to avoid a divergence whose origin lies in the assumption of a homogeneous and infinite medium. For all practical purposes, we expect two-body encounters to occur at a maximum distance given by the size of the host halo and for this reason it is fair to assume Eq. 4.2. However, this is not consistent with other aspects of the derivation of the formula, which instead rely on the assumption of an infinite medium. Despite these facts, Eq. 4.1 provides an accurate description of the drag force experienced by massive bodies orbiting within a system of light particles. This is especially true in the case we are about to discuss, where the objects are point particles and the entire orbit lies inside the host (limiting the impact of the second assumption).

The system we decided to simulate represents a classical setup to test the effects of dynamical friction (see, e.g., Sec. 12.3.1 of Mo et al. 2010). We first generate a distribution of particles with density profile of a singular isothermal sphere:

$$\rho(r) = \frac{V_c^2}{4\pi G r^2}, \quad (4.3)$$

where V_c^2 is the circular velocity, independent of the radial distance from the centre. The velocities follow a locally Maxwellian distribution, with dispersion $\sigma = V_c/\sqrt{2}$. We let the system evolve for several dynamical times, until its lagrangian radii⁴ stabilise and the distribution can be considered to be in equilibrium. We then place a massive body at a certain distance from the centre and set it on a circular orbit (as a consequence, $v_{body} = V_c$). Given this global setup, the prediction for the

⁴These are the radii of the spheres, centred on the object, containing a certain percentage of its total mass.

radial decay of the orbit, as derived from Chandrasekhar's formula, is obtained integrating the following expression:

$$r \frac{dr}{dt} = -0.428 \ln \Lambda \frac{GM_{body}}{V_c}. \quad (4.4)$$

Along with the fundamental assumption $M_{body} \gg m$, the other properties the radial decays depends on are the mass of the body and how this compares to the mass of the host. We have simulated this system varying the number of particles in the isothermal halo, the relative values of m , M_{body} , M_{halo} , the position of the object; for each case, we performed three runs with the usual, different softening approaches, varying also the adopted softening length and the number of neighbours. The results change, but the relative behaviours of the different cases remain qualitatively similar. Here we discuss the simulations performed with 1.5×10^5 particles for the isothermal halo, $M_{body}/m = 100$ and $M_{halo}/M_{body} \approx 400$; for these choices $\ln \Lambda \approx 6$, although we stress again that this should be taken as no more than a guess for what the real value should be. The results of the run with fixed softening are shown in Fig. 4.12. The softening length was chosen to be around the mean interparticle separation in the central region, where the orbital evolution of the massive object takes place. However, we performed simulations using a range of values for this scale and in no cases we found hints of improvements in the results. As shown in the bottom right panel, the first phases of the decay (until around $t = 120$) show a reasonable trend, when compared to the analytical prediction with $\ln \Lambda = 4$. After that point the results of the simulation deviate from the analytical prediction and the decay proceeds slower than it started. The results of the simulation with adaptive softening and no correction of the equation of motion are shown in Fig. 4.13. We have performed several runs varying N_{ngbs} , without major changes or improvements in the outcome; the results shown here were obtained setting $N_{ngbs} = 120$. The oscillations in the radial distance of the massive objects reveal deviation from a purely circular motion; in general, the decay is slower than in the previous case and, at least in the initial stages, it would correspond to $\ln \Lambda < 4$. Finally, in Fig. 4.14 we show the results obtained using adaptive softening along with the new equation of motion. As in the previous case, the simulation was performed with 120 neighbours. The shape of the orbital decay is in remarkable agreement with the expectation for $\ln \Lambda = 6$. Not only does the simulation reproduce the dynamical behaviour of the massive body as predicted assuming a single value of $\ln \Lambda$ throughout time, but this value of the Coulomb logarithm happens to be the favoured one for the system under consideration. Among all the simulations that we have performed, with different softening prescriptions and properties of the system, those with adaptive softening and correction of the equation of motion have always provided the results closer to the expectation. Even though the exact solution to the problem is not known, due to the uncertainties on the value of the Coulomb logarithm, it is still remarkable that these simulations can follow the orbital decay of the object reproducing at least the shape corresponding to a given value for $\ln \Lambda$. It seems that, in this case, the use of adaptive softening, aided by the correction to the equation of motion, provides the better description of the gravitational interaction ruling the evolution of the system.




Figure 4.12: Trajectory and radial decay of a massive body orbiting within an isothermal sphere. The host is sampled by 1.5×10^5 identical particles; the mass ratio between the object and the field particles is 100; the ratio between the mass contained within $r = 1$ and the mass of the body is ≈ 400 . Dynamical friction drives the orbital decay of the object, initially on a circular motion around the host. The first three panels show the projection of the trajectory on the different planes. The red (purple) triangle marks the position of the object at the beginning (end) of the integration; the colour code follows the rainbow spectrum from violet (at $t = 0$) to red (at $t = 280$). The bottom-right panel shows the distance of the body from the centre of the isothermal sphere, as a function of time. The orange curves mark the prediction of the Chandrasekhar’s formula for different values of the Coulomb logarithm; the solid curve corresponds to $\ln \Lambda = 6$ (favoured value for this setup), while the dotted curve to $\ln \Lambda = 10, 4$ (from left to right). The blue line gives the result obtained from the simulation employing fixed softening, set equal to the mean interparticle separation within $r < 1$.

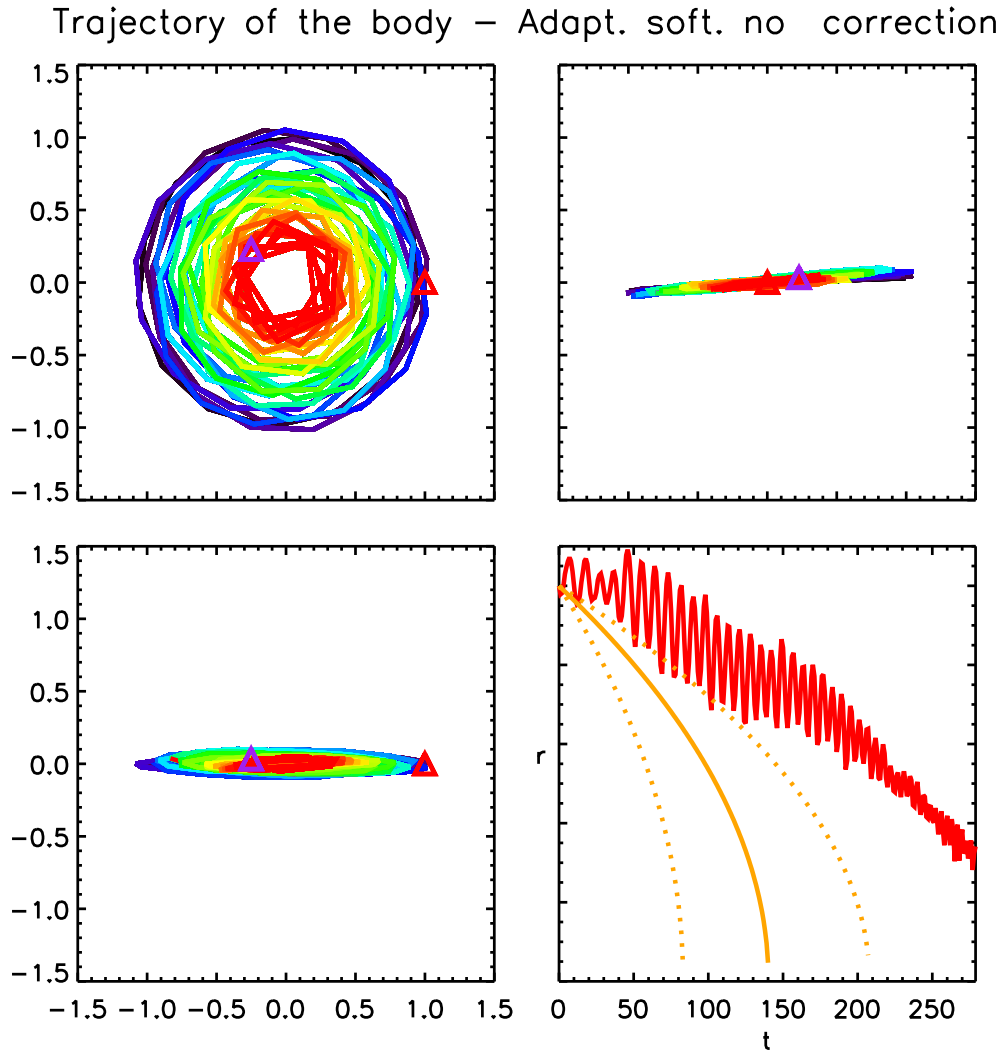


Figure 4.13: Trajectory and radial decay of a massive body orbiting within an isothermal sphere. The host is sampled by 1.5×10^5 identical particles; the mass ratio between the object and the field particles is 100; the ratio between the mass contained within $r = 1$ and the mass of the body is ≈ 400 . Dynamical friction drives the orbital decay of the object, initially on a circular motion around the host. The first three panels show the projection of the trajectory on the different planes. The red (purple) triangle marks the position of the object at the beginning (end) of the integration; the colour code follows the rainbow spectrum from violet (at $t = 0$) to red (at $t = 280$). The bottom-right panel shows the distance of the body from the centre of the isothermal sphere, as a function of time. The orange curves mark the prediction of the Chandrasekhar's formula for different values of the Coulomb logarithm; the solid curve corresponds to $\ln \Lambda = 6$ (favoured value for this setup), while the dotted curve to $\ln \Lambda = 10, 4$ (from left to right). The red line gives the result obtained from the simulation with adaptive softening and no correction of the equation of motion; the individual softenings were computed by means of 120 neighbours.




Figure 4.14: Trajectory and radial decay of a massive body orbiting within an isothermal sphere. The host is sampled by 1.5×10^5 identical particles; the mass ratio between the object and the field particles is 100; the ratio between the mass contained within $r = 1$ and the mass of the body is ≈ 400 . Dynamical friction drives the orbital decay of the object, initially on a circular motion around the host. The first three panels show the projection of the trajectory on the different planes. The red (purple) triangle marks the position of the object at the beginning (end) of the integration; the colour code follows the rainbow spectrum from violet (at $t = 0$) to red (at $t = 140$). The bottom-right panel shows the distance of the body from the centre of the isothermal sphere, as a function of time. The orange curves mark the prediction of the Chandrasekhar's formula for different values of the Coulomb logarithm; the solid curve corresponds to $\ln \Lambda = 6$ (favoured value for this setup), while the dotted curve to $\ln \Lambda = 10, 4$ (from left to right). The red line gives the result obtained from the simulation with adaptive softening and correction of the equation of motion; the individual softenings were computed by means of 120 neighbours.

4.6 Conclusions

In this chapter we have investigated the behaviour of the adaptive-softening formalism on simulations involving different particle species. These included collisionless components, characterised by either different density distributions or particle masses, as well as hydrodynamical simulations following the evolution of dark matter and gas. Our main conclusions are:

- The inner cusp of a density distribution is maintained considerably more stable in time, when either of the two adaptive approaches is adopted; this holds for all components (e.g.: stellar bulge, dark matter halo), regardless of their concentration (see Sec. 4.2).
- In a cosmological simulation with two collisionless species, the spurious evaporation of light particles from collapsed objects is remarkably attenuated when adaptive softening is used without correcting the equation of motion; improvements are registered also when using the fully-conservative, adaptive approach, but in this case the results show some dependence on the parameter N_{ngbs} (see Sec. 4.3, Fig. 4.3).
- Similar results are obtained in a non-radiative, hydrodynamical simulation, namely: the baryon fractions in halos of small-intermediate masses are considerably improved when adaptive softening is used without correcting the equation of motion. The other adaptive approach, instead, performs worse than the standard simulations with fixed softening, albeit just at low redshifts (see Sec. 4.4, Fig. 4.10).
- The differences in the results of the two adaptive approaches is not related to the distribution of softening lengths being discrepant, but rather to some dynamical effect introduced by the correction term in the equation of motion. When the ratio between the mass of the light and heavy particles is raised, the results of the fully-conservative approach improve considerably (see Sec. 4.4, Fig. 4.11).
- In general, when using the adaptive approach with correction of the equation of motion in these simulations, the structure of objects with a number of particles less than N_{ngbs} is not reliable (see Sec. 4.3, Figures 4.3 and 4.7).
- In a simulation of dynamical friction, following the orbital decay of a massive object moving within an isothermal host, the use of adaptive softening along with the correction term provides the result closer to the expected behaviour (see Sec. 4.5).

Even though there are cases where the fully-conservative, adaptive formalism behaves as well as or better than the other adaptive method or the standard, fixed-softening approach, its results are not entirely predictable and are subject to dependencies on numerical parameters (as N_{ngbs} or the mass ratio between the par-

ticles) and on the specifics of the simulated system. We have performed simulations varying the definition of softening from the radius of the sphere containing a certain number of neighbours, to the radius of the sphere containing a certain total mass. As discussed at the end of Sec. 3.2, this constitutes another viable approach to the definition of individual softening lengths and in simulations with multiple species, sampled by particles of different masses, it will generally lead to different softenings. However, its adoption induced no significant changes in the results. It has been argued that, in simulations with multiple components, the best approach would be to define individual softenings out of neighbours of the same species (Daniel Price, David Hubber - private communication). This would have the desirable consequence to set the gravitational softening of gas particles equal, by definition, to the SPH smoothing length (see discussion in the introduction to this chapter). We tried this variant on simulations without correction term and in that case - unsurprisingly - we found no significant changes in the results. The other simulation could not be performed, as the correction term applied to the equation of motion is not consistent with this new definition of softening lengths. A new conservative formalism would have to be derived, starting from another modified Lagrangian.

Considering the results in their entirety, we conclude that the use of adaptive softening provides remarkable and promising results when applied to various kind of collisionless and hydrodynamical simulations involving multiple species. We recommend its use without changes in the equation of motion or, alternatively, to consider the derivation of a new formalism where individual softening are determined out of neighbours of the same species only. The only drawback registered when using adaptive softening without correction to the equation of motion is the loss of halos in the low-mass tail of the mass function (see Sec. 4.3, Fig. 4.9). However, the importance of this effect reduces when decreasing N_{ngbs} . Considering that, in this case, we are not limited by the need of a minimum number of neighbours⁵, we could push N_{ngbs} down to values where the loss of objects becomes tolerable (for $N_{ngbs} = 60$ we register 16% less object; a further decrease to the still reasonable value $N_{ngbs} = 30$ could reduce the deficit further to few percent).

Improved baryon fractions and enhanced resolution in the representation of gravity within collapsed objects will have a strong impact on the star-formation properties of halos in cosmological simulations. It would be interesting to investigate, by instance, the effects on zoomed simulations of galaxy clusters; specifically, one could test the prediction regarding the diffuse stellar component and the low-mass end of the galaxy stellar mass function, quantities which are strongly affected by numerical effects and are still not-converged in state-of-the-art simulations (see, e.g. Saro et al., 2006; Murante et al., 2007; Dolag et al., 2010).

⁵We are, instead, when using the correction to the equation of motion; a minimum number of neighbours is necessary for a robust evaluation of the term (see discussion in Sec. 3.5).

On the orbital and internal evolution of cluster galaxies

Based on
F. Iannuzzi, K. Dolag
2012, to be submitted to MNRAS

Galaxies inhabiting a cluster environment experience significant evolution in their orbital motions throughout time; this is accompanied by changes in the anisotropy parameter, measuring the relative importance of radial and tangential motions for a given class of objects. Along with orbital changes, galaxies in clusters are well known to undergo severe alteration in their hot/cold gas content and star formation properties. Understanding the link between the changes in the internal properties of galaxies and their orbital motion is of crucial importance in the study of galaxy evolution, as it could unveil the primary mechanism responsible for its environmental dependence. Do the changes in the internal properties happen in parallel with those in the orbital motion? Or are the orbital features at the time of infall what determines the fate of the member galaxies? Alternatively: are the properties of galaxies at a given time related to the coeval orbital anisotropy or are they better related to the anisotropy at infall? In order to answer these questions, we studied the orbital evolution of different galaxy populations in the semi-analytic models of Guo et al. (2011) applied on to the Millennium Simulation. For each class of objects, characterised by different internal properties (such as age, star formation rate and colour), we studied the anisotropy profile at redshift zero and its evolution by tracing the progenitors back in time. We conclude that *(i)* the anisotropy of all the galaxy populations increases with time, after falling inside the cluster environment, *(ii)* the orbital properties at infall strongly influence the subsequent evolution of the internal features of galaxies.

5.1 Introduction

In the currently favoured scenario for the formation of cosmological structure in the Universe, dark matter halos form and merge giving rise to a hierarchy of objects assembling in a bottom-up fashion. In this context, galaxy formation is pictured to arise from cooling and condensation of baryons within the potential well associated to these dark-matter structures (White and Rees, 1978). Once the galaxy is formed, its evolution will be driven by *(i)* “nature”, i.e. the object’s intrinsic features (essentially stellar mass) and *(ii)* “nurture”, external processes related to the environment the galaxy inhabits during different stages of its history. It is indeed well known that, although the structural properties of galaxies are mainly determined by their stellar mass (Kauffmann et al., 2003; Tanaka et al., 2004; van den Bosch et al., 2008), the existence of an environmental dependence cannot be disregarded (Hogg et al., 2004; Balogh et al., 2004; Kauffmann et al., 2004; Blanton et al., 2005, among the others). A combination of these effects is responsible for the observed “bimodality” in the galaxy properties (Baldry et al., 2004; Kauffmann et al., 2004), namely the existence of two well-distinguished classes of objects characterised by either red colour/high mass/old stellar population or blue colour/lower masses/young stellar population, the former residing preferentially in over-dense environments (Oemler, 1974; Dressler, 1980; Bower and Balogh, 2004; Balogh et al., 2004; Ball et al., 2008; Bamford et al., 2009; Skibba et al., 2009), while the latter being found mainly in the field. A clear and thorough physical picture of galaxy evolution is yet to come; understanding what mechanisms play the leading role in shaping galaxy properties is a topic of ever increasing interest and research activity, but the results are still controversial. Among the environmental processes, four broad classes of mechanisms are generally considered:

- **galaxy mergers**, negligible in massive clusters, but important in small groups, which drive morphological changes and can affect star formation (Toomre and Toomre, 1972; Farouki and Shapiro, 1981);
- **strangulation** (Larson et al., 1980), causing the removal of the hot-gas halo associated to a galaxy when this is accreted onto a larger structure;
- **ram-pressure stripping** (Gunn and Gott, 1972), important in massive clusters where the density of the intra-cluster medium is highest, it leads to the progressive stripping of the cold-gas component of the satellite galaxies;
- **tidal effects**, arising from the gravitational interaction with other members and with the cluster potential itself, they cause stripping and heating (Richstone, 1976; Moore et al., 1996).

Some of these processes are better understood and formalised (e.g. ram-pressure stripping), while others still lack a strict, physical description and are referred to in more generic terms (e.g. strangulation); on top of this, the regimes where each of the mechanisms is active/unimportant are only broadly assessed. In galaxy clusters, the densest possible environments, the contribution of mergers can be

neglected, due to the large velocity dispersion of the system; strangulation occurs rapidly and devoids the hot-gas reservoir, while ram-pressure stripping and tidal interactions proceed as the satellite plunges into its host. The importance of these last two processes increases with local density (gas density for ram-pressure stripping, total matter density for tidal effects), but while in the first case this results into enhanced stripping, the effect of tidal interactions is believed to be mainly that of an induced gas consumption following the increase of nuclear activity (Boselli and Gavazzi, 2006). Given the dependence upon gas and matter density and how these are, in turn, related to the radial distance from the cluster centre, it is obvious to expect these processes to strongly depend on the orbital history of the satellites, namely how often these happen to transit the innermost regions with respect to the outskirts. This will depend on the initial orbit of the galaxy at the time of accretion, as well as on its time evolution within the dynamically-active cluster region. There exist observational evidences that HI-deficient galaxies in nearby clusters are early-type spirals on radial orbits, while gas-rich galaxies are characterised by more tangential motions (Dressler, 1986; Solanes et al., 2001); other works studied the differences in the velocity distribution of early-type and late-type galaxies and found evidences of the latter moving on slightly more radial orbits, especially at large clustercentric radii (Mahdavi et al., 1999; van der Marel et al., 2000; Katgert et al., 2004; Biviano and Katgert, 2004; Biviano, 2006, 2008; Biviano and Poggianti, 2009; Wojtak and Łokas, 2010). Deriving this information from a sample of observed cluster members is not an easy task and various techniques were developed appositely. By instance, knowing the projected number density profile and line-of-sight velocity dispersion of the selected galaxies, one could perform a Jeans dynamical analysis (Binney and Tremaine, 1987) to derive the cluster mass and velocity anisotropy profiles (quantifying the relative importance of radial and tangential orbits as a function of the clustercentric radius). This technique relies on the assumption of spherical symmetry, collisionless dynamics and dynamical equilibrium, in addition to being hampered by the so-called “mass-anisotropy” degeneracy (even if considerable progress has recently been done to remove it, see Łokas and Mamon 2003; Battaglia et al. 2008; Wojtak et al. 2009). This limitations, along with the difficulties in performing observations in the outskirts of clusters, should induce caution in the interpretation of the results. Yet it would be interesting to know the extent to which the evolution of galaxies is linked to their specific orbital motion. Is the type of orbit determining the efficiency of the environmental processes in shaping the internal properties of the object? In this case we would expect severely affected galaxies to move on more markedly radial orbits, as these cross the cluster right down to its innermost regions. Or do the changes in the galaxy properties happen in parallel to those in the orbital motion? In this case we would expect the orbits of the objects having suffered major environmental influences to differentiate from those of just-accreted satellites.

The importance of the numerical approach in the study of galaxy evolution need not be stressed; cosmological simulations provide a fundamental tool to understanding the assembly of structures throughout time, but the reliability regarding their treatment of interactions other than gravity is subject of debate (see, e.g., Scannapieco

et al. 2011). Semi-analytic models (White and Frenk 1991; Cole 1991; Kauffmann et al. 1993; Cole et al. 1994; Kauffmann et al. 1999; Springel et al. 2001a; see Baugh 2006 for a review) present a powerful, hybrid approach to the problem of galaxy formation and evolution: they make use of accurate, dark-matter-only simulations to account for the growth of structures in the cosmological context and regulate the formation and evolution of galaxies according to a set of analytical prescriptions encompassing the physics of the baryonic component. In this work we use one of the most advanced semi-analytic models developed to date to study the link between orbital and internal properties of galaxies belonging to massive clusters in a Λ CDM cosmology. We will show how the results of this analysis suggest a scenario where the specific orbital features of the satellites have major consequences on their evolution. The chapter is organised as follows: Sec. 5.2 describes the simulation, the semi-analytic model and the choice of the cluster sample used in our analysis; Sec. 5.3 reports the results for the anisotropy parameter of the member galaxies at redshift zero (5.3.1), higher redshifts (5.3.2), at the time of the last infall inside their host (5.3.3) and regarding its time evolution (5.3.4); finally, Sections 5.4 and 5.5 contain a brief summary and discussion of the results.

5.2 The semi-analytic models and the selected sample

The samples analysed in this work were extracted from the galaxy catalogue obtained by Guo et al. (2011) (hereafter GUO11) after running their galaxy formation models on the Millennium Simulation (hereafter MS; Springel et al., 2005). The MS shows how dark matter structures form and evolve in a Λ CDM cosmological scenario characterised by the parameters $\Omega_{\text{tot}} = 1$, $\Omega_m = 0.25$, $\Omega_b = 0.045$, $\Omega_\Lambda = 0.75$, $h = 0.73$, $\sigma_8 = 0.9$, $n_s = 1$. The evolution of cosmological structures is traced by 2160^3 particles moving in a periodic box of side $500 h^{-1}$ Mpc under the mutual gravitational influence. The results of the simulation were stored at 64 different times starting from redshift 127 down to redshift zero. At each of these times a catalogue of bound structures and substructures was generated, by applying a friend-of-friend technique (Davis et al., 1985b) and the SUBFIND algorithm (Springel et al., 2001a) on the particle distribution; these catalogues constitute the basis for recovering the merger history of structures throughout time, also referred to as the merger tree of the simulation. As seen in Sec. 2.5, semi-analytic models of galaxy formation, as those of GUO11, provide a description of the cosmological evolution of baryonic matter; a gas distribution is associated to each bound structure identified in the simulation and its evolution in time is regulated by a set of recipes implemented on to the dark-matter merger tree identified in the simulation: according to the specific histories of each of the substructures, the evolution of the associated baryonic components will follow its own, peculiar path. The physical processes implemented in the recipes of GUO11 include cooling, star formation, supernova and AGN feedback, hot-gas stripping, metal enrichment and alone they provide a remarkable match to the abundance and large-scale clustering properties of the observed galaxy population at low redshift. We will comment on the limits of

the model in the discussion of our results; we refer the reader to GUO11 for a thorough description of the implemented physics and of the strengths and weaknesses of their approach.

We have selected the 1000 most massive clusters identified in the MS at redshift zero; these have all virial masses¹ greater than $2 \times 10^{14} M_{\odot}$ and contain a total of around one million galaxies. Not all of these galaxies have an associated dark matter halo and follow the dynamics dictated by the underlying simulation; the fate of these *orphan* galaxies (generally referred to as “type 2s” in the galaxy formation model), which have lost their dark-matter component, is to eventually merge with the galaxy at the centre of the cluster they inhabit, within the timescales set out by dynamical friction (Sec. 4.5). During this time, the orbit of the orphan galaxy is traced by the most bound particle present in its dark matter halo before this vanished, modified by a shrinking factor introduced to mimic the orbital decay caused by dynamical friction. Since the orbits of these galaxies are altered, we decided not to include them in our analysis. This leaves us with $\approx 2 \times 10^5$ galaxies in the mass range $10^3 < M_*/M_{\odot} < 10^{12}$. We have stacked the resulting galaxy sample by subtracting bulk motions and by normalising positions and velocities to the virial values. To summarise, we are studying the mean behaviour of galaxies residing in the most massive clusters at redshift zero.

From the galaxy catalogue we have access to a wealth of information regarding the internal properties of the selected objects; those we are primarily interested in are stellar/dark matter masses and colour, the latter defined as the the difference between the rest-frame total absolute magnitude in the SDSS u and i bands. We will eventually discuss the impact on our results of using other quantities, such as mean stellar age and specific star formation rate, to split our galaxy sample into distinct populations.

5.3 Results

In this section we will present the results on the velocity anisotropy for the galaxies in the selected sample. This quantity is defined as:

$$\beta = 1 - \frac{\sigma_t^2}{\sigma_r^2}, \quad (5.1)$$

where σ_t and σ_r are the velocity dispersion in the tangential and radial direction², respectively. The velocity anisotropy can take values ranging from $-\infty$ to 1; the former case corresponds to purely circular motions (no dispersion in the radial direction), while in the latter the orbits are completely stretched along the radial

¹These are defined as the masses of the spherical regions centred on the potential minimum of the smooth halo and corresponding to an overdensity of Δ , typically ≈ 200 , with respect to either the critical density ρ_{crit} or to the mean background density $\rho_m = \Omega_m \rho_{crit}$. Hereafter, when referring to either the virial radius or the virial mass of an object, we assume the overdensity to be defined with respect to the critical density.

²By tangential direction we mean the one described by variations in any of the angles θ or ϕ , i.e. $\sigma_t^2 = (\sigma_{\theta}^2 + \sigma_{\phi}^2)/2$.

direction (no tangential dispersion). The case $\beta \approx 0$ is referred to as *isotropy* and it occurs when the velocity dispersion is of comparable magnitude in both radial and tangential directions. We will mostly discuss global values of β , describing a full population of objects without discriminating the spatial distribution of the members; in one case we will also show the radial profile, to underline both the general properties of the anisotropy parameter and its dependence on the distance from the centre of the cluster. In what follows, we split the galaxy population into a red and blue sample according to the value of the $u - i$ colour indicator, namely whether it is greater or smaller than a certain threshold; this is set to the value at which the colour distribution of the galaxy sample splits into two well-defined, different components. At redshift zero, this happens around $u - i = 2.5$; considering all the galaxies within three virial radii, the resulting “red” sample consists of ≈ 72000 objects, against the ≈ 130000 of the “blue” counterpart³.

5.3.1 Anisotropy at redshift zero

We first analyse the results at redshift zero and as a function of the radial distance of the galaxies from the centre of the stacked cluster. The curves in Fig. 5.1 show the behaviour of the anisotropy parameter in 20 equally-spaced radial shells, extending from roughly 0.1 to 3 virial radii. The error bars were calculated by means of a bootstrapping algorithm and correspond to two standard deviations (this holds for all the uncertainties reported in this chapter, unless explicitly stated otherwise). The black curve shows the result for the full population: the importance of radial motions increases, moving from the central regions outwards and peaks around 1.5 – 2 virial radii. This behaviour is not surprising and it is compatible with other results in the literature (Rasia et al., 2004; Gill et al., 2004; Mamon and Łokas, 2005; Sales et al., 2007; Wojtak et al., 2009; Biviano and Poggianti, 2009; Host et al., 2009; Lemze et al., 2011; Lapi and Cavaliere, 2011). The blue and red curves represent the results for the subgroups of galaxies characterised by blue and red colours, respectively; the clear message coming from the plot is that the blue population has a systematically and significantly lower anisotropy than the red population, at least between 0.5 and 2.5 virial radii. This is confirmed by the global values of β , marked by the horizontal, dotted lines. These were computed out of all the galaxies found at distances less than three virial radii from the centre, with the following result: $\beta_{all} = 0.253 \pm 0.006$, $\beta_{blue} = 0.189 \pm 0.006$, $\beta_{red} = 0.345 \pm 0.008$. In summary, red galaxies in our sample move on more radially-biased orbits than the blue counterparts. A clear explanation of this effect will progressively arise in the following sections.

³The higher fraction of blue objects can be traced back to (i) the exclusion of type 2s and (ii) the choice of three virial radii. When restricting the analysis to 1.5 virial radii and including type 2s, the number of red objects becomes nearly four times as large as that of the blue ones.

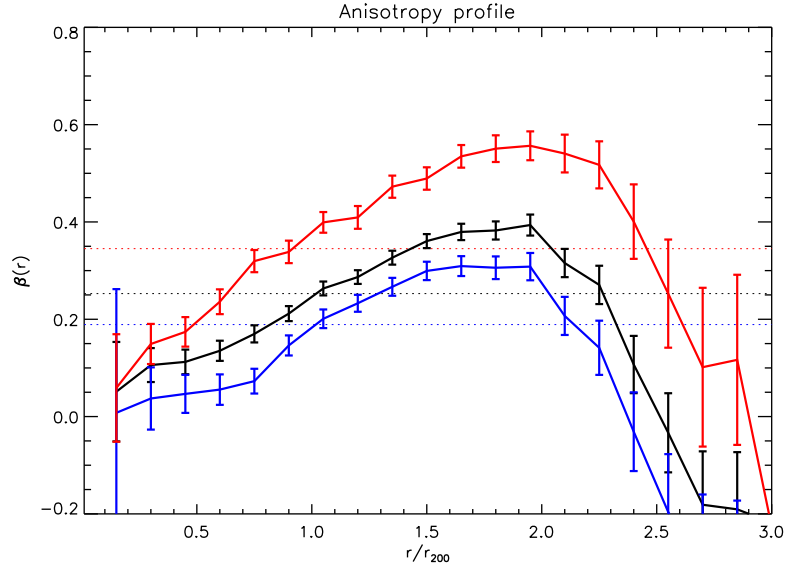


Figure 5.1: Anisotropy profile for the stacked sample at redshift zero. The black, solid curve shows the radial behaviour of the anisotropy parameter for the full population of galaxies found in the selected cluster sample; the red (blue), solid curve represents the result for the subsample of galaxies characterised by an $u - i$ colour greater (smaller) than 2.5. The error bars are evaluated via a bootstrapping algorithm and correspond to two standard deviations. The horizontal, dotted lines mark the global value of the anisotropy parameter for all the galaxies within three virial radii.

5.3.2 Anisotropy at high redshift

In what follows, we focus our attention on the time evolution of the anisotropy parameter. Having not found significant changes in the radial behaviour with respect to the trend shown in Fig. 5.1, we will only refer to global values of β hereafter. We proceeded in two ways: we considered all the member galaxies belonging to the high-redshift progenitors of the selected clusters and, in parallel, we also analysed the progenitors of the redshift-zero galaxies. The two approaches differ in that the first galaxy sample may contain objects that do not survive until redshift zero, whereas in the second case only satellites that have a redshift-zero descendant are considered. We show the results of the first method in Fig. 5.2, where the value of β , computed for the full population of member galaxies regardless of their spatial position, is plotted as a function of redshift. The colour code is the same as for Fig. 5.1; keeping the threshold $u - i = 2.5$ for the splitting between the red and blue population or adjusting it to the variation in the colour distribution of the high-redshift galaxies introduces no substantial changes in the results.

The two main considerations arising from this analysis are (i) that the global value of the anisotropy parameter does not significantly evolve in time and (ii) that the blue population is characterised by a lower degree of anisotropy at all redshifts. A more interesting result is given by the analysis of the second sample of high-redshift galaxies. These are progenitors of redshift-zero objects that are already members of their final, host cluster at the redshift of interest and that will not

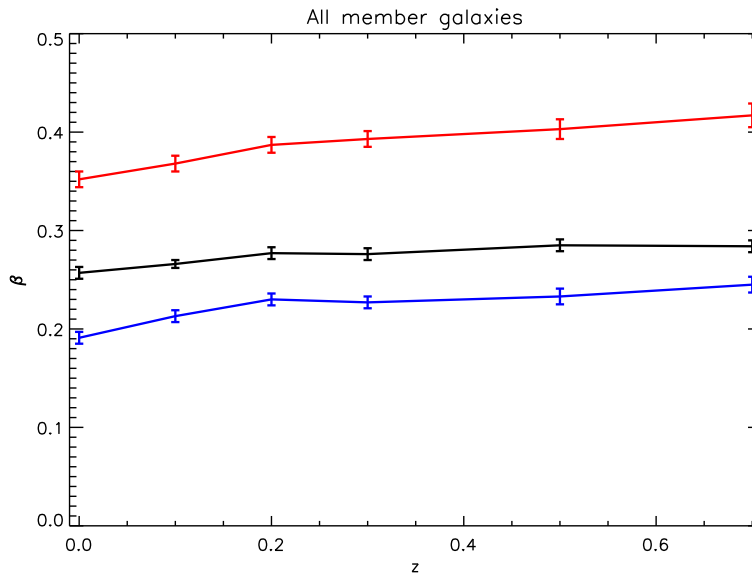


Figure 5.2: Global value of the anisotropy parameter as a function of redshift. Considered are all the member galaxies belonging to the high-redshift progenitors of the redshift-zero cluster sample; these may include objects which do not survive to redshift zero. The black curve shows the result for the full galaxy population at each redshift, whereas the red (blue) curve refer to the subsample of objects with $u-i$ colours greater (smaller) than a threshold; this is set to be the value at which the colour distribution splits into two, well-defined component (e.g. $u-i \approx 2.5$ at redshift zero, $u-i \approx 2.35$ at $z = 0.7$). The error bars are evaluated via a bootstrapping algorithm and correspond to two standard deviations.

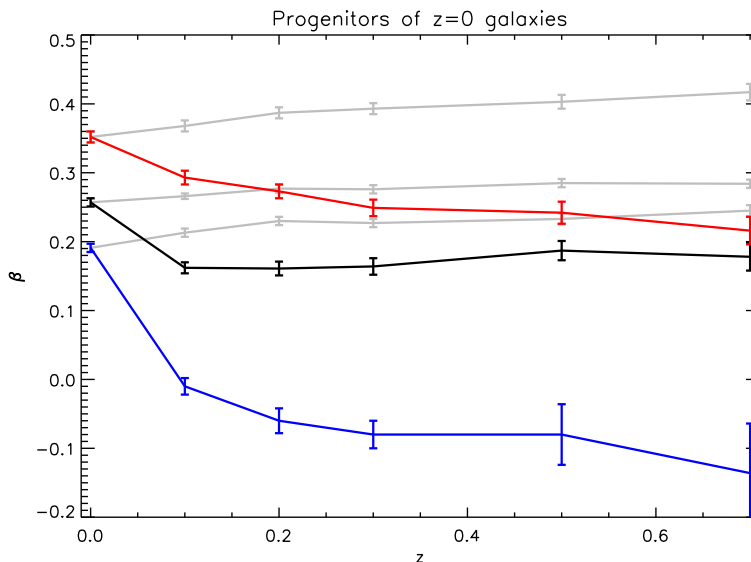


Figure 5.3: Global value of the anisotropy parameter as a function of redshift. Considered are the progenitors of the galaxies belonging to the redshift-zero cluster sample; only progenitors which are already satellites and will remain satellites down to redshift zero are taken into consideration. The black curve shows the result for the full population at each redshift, whereas the red (blue) curve refer to the subsample of objects with redshift-zero descendant characterised by $u-i$ colours greater (smaller) than 2.5. Overplotted in gray are the results displayed in Fig. 5.2. The error bars are evaluated via a bootstrapping algorithm and correspond to two standard deviations.

leave it anytime afterwards. Fig. 5.3 shows the global anisotropy for this class of galaxies; the colour code refers to the redshift-zero population only and not to the progenitors: the red (blue) curve corresponds to progenitors of galaxies that are red (blue) at redshift zero. The plot shows with striking clarity that, at each redshift, the anisotropy of the progenitors of redshift-zero galaxies is considerably lower than the value for the whole, high-redshift population (Fig. 5.2, overplotted in gray); this effect seems to be even stronger for progenitors of galaxies that are blue at redshift zero. These results suggest *(i)* that galaxies on more radial orbits have less chances to survive and are progressively removed from the member population and *(ii)* that a similar selection effect acts on the galaxy colour (namely: galaxies on radial orbits have very low chances to remain blue until redshift zero). The second point will become even clearer when analysing the anisotropy at the time the galaxies fall inside their host cluster.

5.3.3 Anisotropy at infall

We finally study the orbital properties of galaxies at the time they become members of their cluster; by this we mean the moment the dark-matter halo of a galaxy is not an independent structure anymore, but a substructure identified within a larger halo. We want to see how the anisotropy of the infalling population varies with time. Again, we split the analysis into two parts; first, we consider the full

sample of galaxies joining the high-redshift progenitors of our original cluster sample and, second, we also restrict the study to the subsample of galaxies with a descendant at redshift zero. The results of the first approach are displayed in Fig. 5.4, where we plot β against the infall redshift. The red (blue) curve refers again to galaxies that are red (blue) at the redshift under consideration; keeping $u - i = 2.5$, or varying the $u - i$ colour threshold to adapt to the colour evolution of the galaxy population at high-redshift introduces non substantial changes in the results. A clear trend appear, showing that the infall anisotropy increases going towards lower redshifts; this is not surprising, as mass-accretion is expected to occur progressively along small filaments, extending radially outside massive clusters. The second feature which is apparent in the plot is that most of the infalling population consists of blue galaxies, as shown by the vicinity of the black and blue curve, as well as by the size of the error bars. Third we see that, again, the anisotropy of blue galaxies is lower than that of the red galaxies. At a first glance, this may result somewhat unexpected as there is no obvious reason why such a trend should already be in place before entering the cluster environment. We note, though, that the mean mass of the red sample is considerably higher than for the blue counterpart (almost two order of magnitudes for both stellar and dark-matter masses); previous works have already stated that more massive satellites tend to approach the cluster along more eccentric orbits, as they are more likely to reside in dense, radial filaments than the least massive ones (Tormen, 1997). This effect could explain the behaviour observed in Fig. 5.4 and also play a role in the interpretation of Fig. 5.1; an initially higher anisotropy for the red population of infalling galaxies may leave an imprint in the final profile at redshift zero. The impact of this effect is anyway limited by the small number of these galaxies within the infalling population: at $z = 0.7$, out of 19049 infalling galaxies only 419 are red. Clearly, these objects may only represent a contribution to the population of red, member galaxies at low redshifts. Fig. 5.5 shows the results for the subsample of infalling galaxies with a redshift-zero descendant. As usual, in black are the result for the full population whereas the red (blue) curve refers to the subsample of galaxies that are progenitors of objects characterised by a red (blue) colour at redshift zero. Again we see that, at each redshift, the anisotropy of these galaxies is systematically lower than the value for the full, infalling population (Fig. 5.4, overplotted in gray). This suggests that galaxies present in the cluster at redshift zero tend to originate from the subsample of objects entering the cluster environment with the least radially stretched orbit. We then split the full population not according to the colour of the redshift-zero descendant, but on the basis of their colour at the time of infall; the orange (cyan) curve represent the infall anisotropy of galaxies that are red (blue) at the time of infall. We clearly see that the galaxies which are blue at redshift zero are descendant of the subgroup of infalling, blue galaxies with the lowest infall anisotropy. It seems, therefore, that not only does the initial orbit strongly influence how long the galaxy is going to survive in a cluster environment, but that it also plays an important role in determining the evolution of its internal properties.

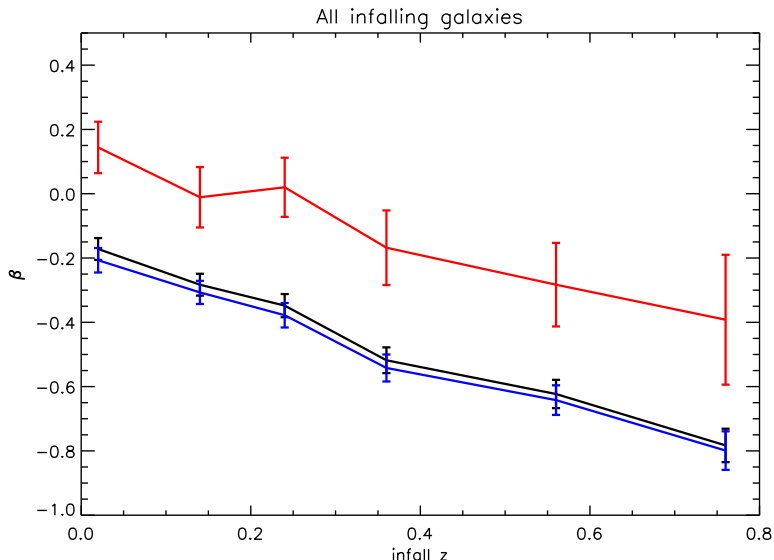


Figure 5.4: Global value of the anisotropy parameter for the infalling population, as a function of infall redshift. Considered are all the galaxies falling inside the progenitors of the redshift-zero clusters at different times; these may include objects which do not survive to redshift zero. The black curve shows the result for the full infalling population at each redshift, whereas the red (blue) curve refer to the subsample of objects with $u - i$ colours greater (smaller) than 2.5. The error bars are evaluated via a bootstrapping algorithm and correspond to two standard deviations.

5.3.4 Evolution of anisotropy in time

When visually comparing the results for the infall anisotropy (Fig. 5.4 and 5.5) to the typical values at redshift zero (Fig. 5.1 and first points of Fig. 5.2) we already guess that β increases, from the time the galaxies enter the cluster to the end of the simulation. This is shown more quantitatively in Fig. 5.6. Each of the curves represents the results at a specific redshift, from 0.7 (in red) to zero (in black); the points on each curve correspond to the anisotropy, at this redshift, of the subgroup of galaxies that fell inside the clusters at the redshift given in abscissa. The first points of the lines show the anisotropy of the galaxies which are just infalling and, taken altogether, they reproduce the black curve of Fig. 5.4. Overall, there is a strong indication of an increase of β with time, after the galaxies become satellites and orbit within the environment of a larger halo; this seems to occur rapidly within the first 2 Gyrs after infall and more gradually afterwards.

In order to understand what this increase was due to, we investigated the evolution of individual orbits. We randomly selected 500 galaxies from the redshift-zero sample and recovered their full history, starting from the last infall inside their host. Knowing the satellite’s mass, its position, velocity and the virial mass of the host at the time of infall, we integrated the orbit forward in time. We used the leapfrog method to solve for the motion of the object, with timesteps of around 350 Myrs. We adopted this time resolution as it corresponds, approximately, to the separation between the different snapshots in the MS. The integration proceeds for a number




Figure 5.5: Global value of the anisotropy parameter for the infalling population, as a function of infall redshift. Considered are the progenitors of the galaxies belonging to the redshift-zero cluster sample; only progenitors which are becoming satellites for the last time in their history are taken into consideration. The black curve shows the result for the full infalling population at each redshift, whereas the red (blue) curve refer to the subsample of objects with redshift-zero descendant characterised by $u - i$ colours greater (smaller) than 2.5. Displayed in orange (cyan) is the anisotropy of the subgroups of infalling satellites which are red (blue) at the time of infall. Overplotted in gray are the results displayed in Fig. 5.4. The error bars are evaluated via a bootstrapping algorithm and correspond to two standard deviations.




Figure 5.6: Global value of the anisotropy parameter at different times and as a function of infall redshift. Each curve represents the results at a specific redshift, as stated in the legend. The points on the curves correspond to the anisotropy of galaxies which have last become satellites at the redshift given on the x axis. The error bars are evaluated via a bootstrapping algorithm and correspond to two standard deviations.

of steps corresponding to the total time elapsing from infall to redshift zero. We adopt two different approaches in the integration: in the first case (hereafter case A) we keep the mass of both the host and the satellite constant at the infall value, whereas in the second case (hereafter case B) we update both masses according to the values provided in each of the snapshots⁴. In both cases we assume the mass of the host to follow a spherical NFW profile (Navarro et al., 1996); the concentration parameters were evaluated from Eq. 5 of Neto et al. (2007), who obtained the concentration-mass relation for the halos identified in the MS at redshift zero; we do also allow the concentration parameter to vary with redshift and assume the dependency found by Duffy et al. (2008) for simulated halos in the redshift range 0 – 2. At each timestep the mass felt by the orbiting satellites varies, according to the distance from the centre of the host (for both case A and B) and to the cosmological accretion (taken into account in case B only). Case A hardly ever reproduces an orbit close to the original, whereas case B provides a better description to the motions occurring in the simulation. Even in this case, however, a perfect match to the original orbit is not guaranteed. Our approximation for the distribution of the host mass, on top of the fact that we are neglecting the effect of interactions with other orbiting galaxies, can sometimes result in evident mismatches. Moreover, our time resolution is considerably rougher than that of the MS. We could certainly improve our integration by interpolating the results for the host and satellite masses

⁴The value for the dark-matter mass of the satellite, recovered from the simulation, is subject to inaccuracies whose importance depends on the spatial position of the object (see Knebe et al. 2011). We have performed the integration updating the mass of the host only and keeping the mass of the satellite fixed at the infall value: the differences with the results obtained in case B are negligible.

between two subsequent snapshots and, thus, by using smaller timesteps, but an exact match to the MS results is not necessarily what we are aiming for: we want to study the evolution of the orbital shapes in a reasonably realistic scenario and see (i) if we can reproduce the increase in anisotropy observed in the MS results and (ii) if it is possible to isolate the responsible mechanism. Fig. 5.7 shows few examples of original orbits from the MS (black curves in the plots on the left column) and the results of our integration for case A and B (overplotted purple and magenta curves, respectively); as mentioned, in some cases our approximation grasps the original dynamics pretty well (e.g. magenta curves in the first three rows), while in others we are clearly missing something (e.g. last row). The right column shows the time evolution of the mass felt by the satellite in both case A and B and explains the differences in the corresponding orbital evolution.

Knowing position, velocity and masses one can straightforwardly compute the orbital energy E , the angular momentum L and, from these, the orbital eccentricity:

$$e = \sqrt{1 + \frac{2EL^2}{(GmM)^2\mu}}, \quad (5.2)$$

where m and M are, respectively, the satellite and host mass and μ is the reduced mass ($\mu = mM/(m + M)$). Closely related to the eccentricity is the orbital circularity

$$\eta = \sqrt{1 - e^2}, \quad (5.3)$$

formally defined as the ratio between the angular momentum L and that of the circular orbit characterised by the same orbital energy E . The circularity parameter has finite values only for bound orbits, i.e. those with $E < 0$ or, alternatively, $e < 1$; the limit $\eta = 0$ corresponds to purely radial orbits, whereas $\eta = 1$ to perfectly circular motions. At each timestep of the integration, we computed the circularity of the orbit and looked at its evolution with time. Fig. 5.8 shows the results of this study, for both integration A and B (left and right columns, respectively). The gray lines in both panels of the first row show the evolution of circularity with time for the subsample of galaxies on bound orbits for at least 20 timesteps (to avoid overcrowding). Overplotted is the median value computed in seven redshift bins and for all 500 galaxies: this remains roughly constant in case A, while it decreases to lower circularities in case B. The blue (green) star shows the median circularity at the first (last) timestep where the orbit is bound, again for all 500 galaxies: the evolution proceeds with opposite trends in the two cases. The last two rows display the distributions behind these median values. In the second row the results for the seven redshift bins are presented; the black curve corresponds to the highest-redshift bin and the ivory to the lowest, following the colour code of the solid lines in the first two panels. While no strong evolution is registered for case A, in case B the distributions clearly evolve towards smaller values of η . Finally, in the last row we see how circularity changes from the beginning to the end of the integration (actually: from the first to the last timestep where the orbit is bound), regardless to the absolute times/redshifts these moments correspond to. The results are very clear and, as mentioned, opposite for the two integrations. In

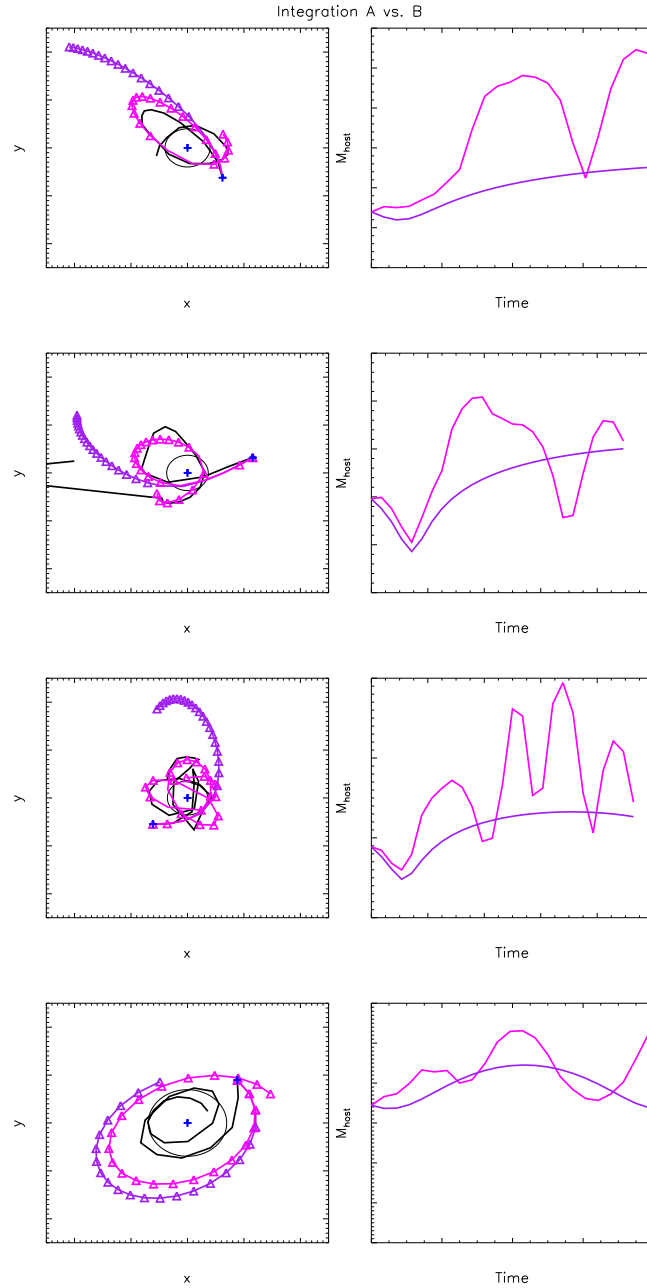


Figure 5.7: Examples of integrated orbits. The left column shows the motion of four different galaxies around their host. The blue crosses represent the centre of the host and the initial position of the galaxy, while the circle gives the size of the virial region at the beginning of the integration. The black curve corresponds to the original orbit from the MS and the purple (magenta) curve shows the result of the integration in case A (B). Whenever the orbit is bound, a triangle is overplotted; this is not done for the original orbit, as the information on its energy is not available. The timestep is $\approx 350\text{Myrs}$. The right column shows the evolution of the mass felt by the galaxy as it orbits around the host in both case A (purple curve) and case B (magenta curve); as spherical symmetry is assumed in the mass distribution of the host, this quantity corresponds to the mass contained within the sphere of radius $\mathbf{r} = \mathbf{r}_{sat} - \mathbf{r}_{host}$.

case A we start from a distribution with median circularity ≈ 0.31 (blue curve) and evolve towards a distribution with median circularity ≈ 0.5 (green curve), meaning that the orbits tend to circularise. In the second case, where the evolution of the host mass is taken into account, the initial orbits have a median circularity ≈ 0.48 (similarly to what found by Tormen 1997), while the final orbits ≈ 0.26 : the motions become more and more radial, leading to an increase of the anisotropy parameter. This small experiment seems to suggest that the origin of the increase in β , which we have registered in the above analysis, is to be connected with the evolution of the orbits in an ever-deeper potential well, such as that of a growing host.

5.4 Summary

We have studied the evolution of the anisotropy parameter for galaxies orbiting within the most massive clusters extracted from the MS; we have related the value of this parameter to the internal properties of the galaxies, as predicted by the semi-analytic models of GUO11. Our findings can be summarised as follows:

- At redshift zero, blue galaxies move on less radial orbits than red galaxies do ($\beta_{blue} \approx 0.19$ vs. $\beta_{red} \approx 0.35$; see Fig. 5.1).
- At higher redshifts, progenitors of redshift-zero objects move on less radial orbits with respect to the full population of member galaxies (see Fig. 5.3); this is particularly true for progenitors of galaxies that are blue at redshift zero.
- Of all the galaxies entering the cluster at a certain time, those that will survive until redshift zero are the subgroup characterised by the most tangential orbits at infall (see Fig. 5.5).
- Of all the blue galaxies having a redshift-zero descendant (either red or blue) and entering the cluster at a certain time, those that remain blue until redshift zero are the subgroup of infalling objects on the most tangential orbits (compare the blue and cyan curves in Fig. 5.5).
- The orbits of infalling galaxies become increasingly radial going towards redshift zero (see Fig. 5.4).
- The orbit of satellite galaxies become increasingly radial in time after the infall inside the cluster environment (see Fig. 5.6).
- This increase can be explained by cosmological accretion, deepening the potential of the host as the galaxies orbit around it (see Sec. 5.3.4).

These findings and their statistical significance strongly suggest that the orbital features at infall have a major influence on the subsequent evolution of the galaxies, i.e. on their survival time within the cluster and on their turning from being star-forming/blue to passive/red objects. Galaxies on more markedly radial orbits




Figure 5.8: Evolution of the circularity parameter for cases A (left column) and B (right column). The integration of the orbits has been performed on a randomly selected subsample of 500 galaxies, given their initial position and velocity. The first row shows the evolution of η with time: the gray, thin curves in background correspond to the subset of galaxies found on bound orbits for more than 20 timesteps; the thick, solid line marks the median behaviours of all galaxies in seven redshift bins; the blue (green) star shows the median circularity at the first (last) time, during the integration, where the orbit is bound. The second row shows the circularity distribution in the seven redshift bins: black to ivory, the curves correspond to the highest and lowest redshift bin, following the color code of the lines in the first two panels. The last row shows the circularity distribution at the first and last time the orbit is bound (blue and green curves, respectively).

will be more easily disrupted and removed from the cluster members as time goes by; even if they survive until redshift zero, these objects are those most likely to undergo changes in their internal properties and turn from being star-forming/blue to passive/red. Galaxies which, conversely, join their host with a significant component of their motion along the tangential direction tend to be less affected by the cluster environment; to this category belong all the blue galaxies found in the redshift-zero sample.

5.5 Discussion

In the previous sections we have been drawing conclusions from a model which, albeit being state-of-the-art, is well known to contain some level of simplification in its treatment of galaxy formation and evolution; caution is therefore needed.

As discussed in Section 4.4 of Guo et al. (2011), only for galaxies in the stellar mass range $9.5 < \log(M_*/M_\odot) < 11$ does the model predict $u - i$ colour distribution in agreement with observations; at lower (higher) masses the galaxies tend to be too red (blue) and the colour distribution strongly deviates from the reference given by the SDSS/DR7 sample. We have performed our analysis on the full galaxy population and after applying a mass-cut leaving objects within the above range only; the results remain qualitatively the same.

As already mentioned, galaxies orphans of their dark-matter halo, or “type 2s”, were not included in the analysis due to the lack of information on their real dynamics. These objects are an important ingredient in the model as their presence allows the predicted galaxy luminosity function and radial number density profile in clusters to match observations; in addition to this, they represent a substantial fraction of the satellite populations (they account for half of all cluster members with $M_* > 10^{10}M_\odot$). Including these objects in our sample has the immediate effect of increasing the median stellar age and $u - i$ colour index, while lowering the median specific star formation rate; briefly, the sample ages. The effect on the anisotropy parameter corresponds to enhanced radial motions: from $\beta \approx 0.25$ for the full, redshift-zero population without type 2s, we move to $\beta \approx 0.35$ when these are included. The relative differences between the blue and red population are maintained.

We have split the original sample of galaxies into two different populations characterised by different values of the $u - i$ colour indicator. As a threshold, we have adopted the value where the colour distribution separates into two distinct component; for the galaxy sample we selected, this occurs at $u - i \approx 2.5$ at redshift zero. We have performed the same analysis changing this threshold and making more extreme cuts, both at redshift zero and at higher redshifts; the results remain coherent with those shown in the previous section. We have also used other properties, besides colour, to identify the two different galaxy populations, namely specific star formation rate and mean stellar age. Not surprisingly, as all these properties are expected to relate to one another, the results have not significantly

changed.

We already mentioned that, according to our results, the type of orbit a galaxy is moving on when it last enters the cluster environment has major consequences on the evolution of the internal properties of the object. The finding must be related to the mechanisms responsible for these environment-induced changes affecting satellite galaxies moving within a cluster potential. The model of GUO11 includes a more sophisticated and realistic treatment of these effects with respect to earlier attempts; indeed, as opposed to previous semi-analytic models (Baldry et al., 2006; Weinmann et al., 2006; Wang et al., 2007; De Lucia and Blaizot, 2007), where the hot gas associated to a galaxy is immediately and entirely removed as the object is accreted on to a larger system, in GUO11 the stripping of gas is performed gradually and modelled to reproduce the combined effect of tidal and ram-pressure forces (extending and integrating the recipes of Font et al. (2008) and Weinmann et al. (2010)). In addition to this, these processes are activated only while the satellite resides within the virial radius of the host; this limits the impact of possible failures in the FOF algorithm to identify physically independent, but spatially close, structures. The improved treatment of environmental effects allows GUO11 to reproduce, with noticeable accuracy, the radial distribution of star-forming, cluster galaxies as found in the SDSS data for a large sample of nearby clusters (see their Fig. 3). Albeit its successes, the model is still incomplete as it lacks any treatment for the stripping of the cold-gas component, an effect known to be important in the inner regions of rich clusters (Gavazzi, 1989; Solanes et al., 2001; Boselli and Gavazzi, 2006). Even though the overall treatment of environmental processes in the model does not entirely encompass the full range of mechanisms at work in real clusters, we think our results are not significantly affected; if anything, the inclusion of cold-gas stripping and therefore a more aggressive implementation of gas removal could only strengthen the differences found in the galaxy properties as a function of their initial orbit.

These speculations are confirmed by the results we found when applying a similar analysis to hydrodynamical simulations, where the co-evolution of dark and baryonic matter is followed self-consistently throughout time. We have considered few among the most massive and relaxed clusters from the Hutt⁵ sample (Dolag et al., 2009), a set of high-resolution, zoomed cluster simulation with implemented cooling, star-formation and feedback processes. We also used the results from the cosmological Magneticum Pathfinder Simulations⁵, isolating the most massive clusters in the $128 h^{-1}$ Mpc-side box (“Box 3” run), as well as in the $896 h^{-1}$ Mpc-side box (“Box 1a” run). In all cases, we found anisotropy profiles in qualitative agreement with those from the semi-analytic models in terms of radial behaviour and global values for β . More importantly, we found that simulated galaxies characterised by either a young stellar population, low $u - i$ colour index or high specific star formation rate ($> 1 \times 10^{-11} M_{\odot} \text{ yr}^{-1}$) present a systematically lower anisotropy parameter in each of the radial bins.

To summarise, we are confident that the intrinsic limits of the semi-analytic model

⁵<http://www.mpa-garching.mpg.de/~kdolag/Simulations/>

by GUO11 do not affect the conclusions drawn in this paper.

A final comment on our choice regarding the cluster sample: we have focussed our analysis on the most massive clusters found in the MS - objects with virial masses greater than $2 \times 10^{14} M_{\odot}$ - but we emphasise that there is no reason to expect the results obtained for this sample to apply at lower host masses. In fact, studying the properties of the dark-matter subhalos identified in the Millennium II Simulation (Boylan-Kolchin et al., 2009), Faltenbacher (2010) shows that their global anisotropy parameter depends both on the host mass and on the environment the host is sitting in, with satellites residing in either massive or isolated clusters moving on more markedly radial orbits. On the other hand, Wetzel (2011) examines high-resolution, N-body simulation and evaluates the orbital parameters of satellites at the time of infall, defined as the first time the object crosses the virial radius of a larger host halo; again, a clear dependence of the initial orbit on the host mass is found, with galaxies characterised by increasingly radial motions at higher host halo masses. We therefore do not expect our results to apply to environments other than those of the most massive structures, at least not on a quantitative level; also the dependence of the internal evolution of the galaxy properties on the initial orbit may not be as strong in smaller objects or groups, where ram-pressure stripping is much less efficient a process than in rich clusters.

Besides the results on the link between orbital and internal properties of galaxies, we also showed our findings on the temporal evolution of the orbital anisotropy itself. We found that this quantity increases in time and that this can be related to the presence of a growing host. The result emerges from several aspects of our analysis; however, in order to further assess the existence of this signal, we performed an even cleaner test. We have taken a set of galaxies at a certain point in time, after their last infall inside the cluster, and followed them through a number of snapshots, monitoring their global anisotropy. In this way we are tracking always the same objects at each time, as opposed to the analysis presented in Fig. 5.6 and 5.8. From the sample of 500 objects used in Sec. 5.3.4, we have taken the subset of galaxies that co-existed through ≈ 20 snapshots, down to redshift zero. This resulted in a sub-sample of 380 galaxies, whose global behaviour is shown in Fig. 5.9. During the 6.8 Gyrs that elapse between snapshot 43 ($z = 0.83$) and 63 ($z = 0$), the global anisotropy of the group increases from ≈ 0 to ≈ 0.3 . The trend of increasing β with time arises quite strongly from all our analysis. However, we do register some tension between this result and existing observations. Analysing two cluster sets at $z \approx 0$ and $z \approx 0.6$, Biviano and Poggianti (2009) depict a scenario where radially-infalling galaxies progressively turn from being star-forming to quiescent, while reducing their anisotropy from positive values down to $\beta \approx 0$. At redshift zero, the resulting anisotropy profile for the star-forming galaxy sample is consistent with more radially biased orbits than those characterising the quiescent counterpart. This is equivalent to saying that the orbits evolve in parallel with the internal properties of the satellites and that the former do not significantly impact the latter. Our analysis does, instead, support the opposite scenario, as thoroughly discussed. Earlier on, Mahdavi et al. (1999) and Biviano and Katgert (2004) reported of observational findings supporting more radial orbits for late-type

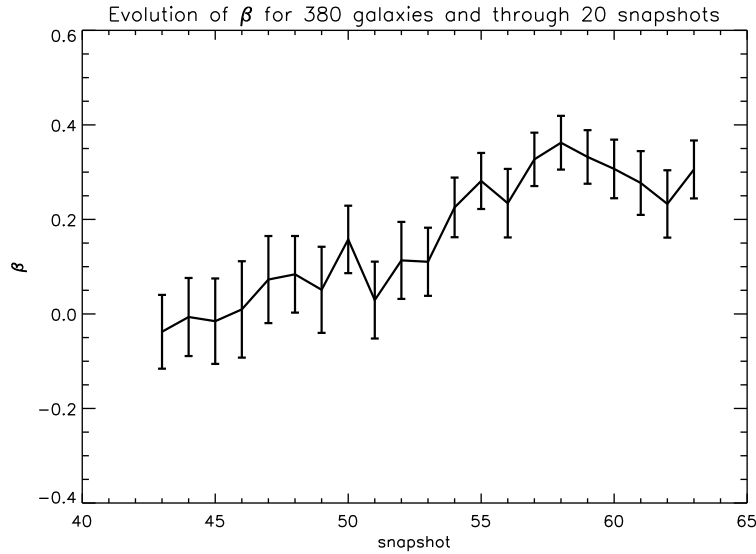


Figure 5.9: Evolution of the anisotropy parameter for a group of 380 galaxies co-existing through 21 snapshots inside the cluster environment. Snapshot 43 corresponds to $z = 0.83$ and snapshot 63 to $z = 0$. The error bars are evaluated via a bootstrapping algorithm and correspond to one standard deviation.

galaxies than for early-type ones; their explanation was, again, in terms of coeval changes in the orbital and internal properties of the objects. Other observations, as mentioned in Sec. 5.1, report of gas-rich galaxies in nearby clusters being found on more tangential orbits (Dressler, 1986; Solanes et al., 2001); this is better reconciled with our findings, at least those regarding the impact of the satellites' orbits on their internal evolution.

Both the observational and the numerical approach come with specific limitations and it would be interesting to pinpoint the origin of the disagreement. One way to do this could be to analyse the simulated sample with the techniques used by the observers to extract dynamical information out of their data. This could point out possible inadequacies in the assumptions and techniques used in the processing of observed data, or, alternatively, rule this option out and restrict the boundary of the problem to the model and its analysis.

Conclusions

In this thesis we have addressed two distinct problems in the field of numerical cosmology. In the first part of the work (Chapter 3 and 4) we have been concerned on how best to simulate the gravitational interaction, which represents the main driver of structure formation in a cosmological context. In the second part (Chapter 5) we made use of a state-of-the-art semi-analytic model to investigate the link between the orbital and internal properties of satellite galaxies - objects moving through the dense medium of a galaxy cluster. In the first introductory chapters we outlined the basics of Cosmology (Chapter 1) and stressed the importance of numerical simulations in our understanding of the observable universe, whilst outlining the most important techniques in use (Chapter 2). In the following, we will concisely summarise the relevant points of each topic and report our main findings.

Softening of the gravitational force is adopted in cosmological simulations in order to moderate the discreteness effects due to a poor representation of the underlying phase-space density. Indeed, the point particles used in a simulation to trace the evolution of the cosmic fluid are considerably less numerous than the real building blocks of the system. This results in a somewhat inaccurate representation of the overall dynamics, with the importance of collisions between particles being severely overestimated. On small scales, these spurious encounters have the aggravating property of being particularly demanding on a computational level. Via softening of the gravitational force, one limits the impact of small-scale collisions on the performance and quality of the simulation. However, by doing so the Newtonian form of the interaction is lost on scales smaller than the softening length and this sets a strong spatial-resolution limit. The softening length is generally set to a fixed value - a fraction of the mean interparticle separation within the computational box. There is no such a thing as an optimal value for the softening length when the density field varies considerably in space and time, as in a cosmological simulation. In fact, keeping the value fixed will certainly result in an inappropriate dynamical description of all regions at several times during the computation. Ideally, the solution would be to adapt the softening scale according to the local density, thus allowing higher resolution in over-dense regions without enhancing collisionality in

under-dense environments. The problem with this approach is the introduction of an additional dependence (on the spatial position, through the softening length) in the Lagrangian of the system; if not accounted for when deriving the equation of motion, this may result in strong fluctuations or drifts in the energetics of the system. In this context, Price & Monaghan (2007) have proposed a formalism to adapt gravitational softening lengths while retaining conservation of both momentum and energy.

In Chapter 3 we discussed the implementation of adaptive gravitational softening in the cosmological simulation code GADGET-3. We have applied the method to several test cases and to a set of cosmological dark-matter simulations of structure formation. Our main finding is that the use of adaptive softening enhances the clustering of particles at small scales, a result visible in the amplitude of the two-point correlation function and in the inner density profile of massive objects, thereby anticipating the results expected from much higher resolution simulations.

Besides the modelling of dark matter, an accurate description of the behaviour of the baryonic component is highly desirable in a cosmological simulation. In the end, we can only directly observe the properties of ordinary, luminous matter. Treating the coeval evolution of dark-matter, gas and stars is not a trivial problem; besides having to account for the physics involved, an additional problem is in that the various matter fields are generally sampled by particles of different masses. Spurious energy transfer occurs from the heaviest to the lightest component, leading to an artificial evaporation of light particles from collapsed objects; eventually, this results in an overall lower baryon fraction in objects of intermediate to small masses, with conceivable consequences on their associated star formation. Increasing the softening length generally alleviates the problem, at the expense of a loss in resolution. Considering that the adaptive method described in Chapter 3 induces an overall increase in the individual softening lengths, whilst enhancing the resolution, we expected it to outperform the standard, fixed softening approach even more in these hydrodynamical simulations. Not only should the segregation of the gaseous component be considerably attenuated, but also its collapse would be followed down to scales which are currently unachievable in standard simulations at comparable mass resolution, thus providing a more reliable representation of the behavior of galaxy-like substructures.

In Chapter 4 we investigated the effects of adaptive softening on simulations featuring species sampled by particles of different masses, whether they represented two dark matter fields or dark matter and gas. Our expectations have only partially been fulfilled. Even though there are cases where the fully-conservative, adaptive formalism provides much improved baryon fractions with respect to the standard, fixed-softening approach, its results are not entirely predictable and are subject to dependencies on numerical parameters and on the specifics of the simulated system. In contrast, from the test cases performed, we register an excellent behaviour of adaptive softening when the equation of motion is left unchanged. The problems identified in the use of this hybrid approach - discussed in both Chapter 3 and 4 - can be easily minimised. Alternatively, a possible way to overcome the issues

in the use of the fully-conservative method may lie in a re-definition of individual softening.

In the last two decades, the importance of modelling the baryonic component of the cosmic fluid led to the progressive development of semi-analytical techniques. These rely on the knowledge of the merger history describing the formation of dark matter halos, a piece of information generally extracted from cosmological simulations. The behaviour of baryons, in the form of gas and stars, within these dark structures is followed by means of analytic prescriptions, regulating the importance of processes like gas cooling, star formation and feedback mechanisms. How these shape the properties of individual objects depends on the properties of the host halo, essentially its mass and assembly history. Semi-analytic models of galaxy formation and evolution provide an impressive match to the properties of the observed galaxy population in the Local Universe, although there is still room for improvement.

In Chapter 5 we made use of a state-of-the-art semi-analytic model to investigate the link between the orbital and internal properties of satellite galaxies. When, during their evolution, galaxies stop being isolated objects and become part of a cluster, a number of environmental processes acts altering their gas content and star formation activity. This happens while galaxies orbit within the dense intra-cluster medium and interact with other members. At least some of these processes depend on the gas or dark-matter density of the environment, and their efficiency is strongest in the innermost regions of a galaxy cluster; the specific orbit followed by each object should then have an impact on the extent to which the latter is affected by the environment. At the same time, encounters with other members and, most importantly, variations in the gravitational potential induced by the cosmological growth of the host cluster, may have an impact on the orbits of the satellite galaxies. We were concerned on how these two aspects related to one another. Do the orbits change as a consequence of the prolonged stay within the cluster environment while the internal properties vary accordingly? In this case one would expect the two to evolve simultaneously and depend essentially on the time elapsed from infall. Does the initial orbit of the satellite galaxy determine its late-time evolution in terms of star-formation activity? In this case, one should be able to see an additional dependence of the galaxy properties on the initial orbital parameters. The quantity we have used to quantify the orbital features of satellite galaxies is the orbital anisotropy; radially anisotropic orbits plunge deeper into the cluster centre, whereas tangentially anisotropic orbits have a more circular shape. Observational evidences suggest that young, star-forming galaxies move on more radially biased orbits than their passive, old counterparts. The explanation given for this is that the former still retain memory of their radial infall, while the latter have had their orbits severely altered during their stay within the cluster. From the model, we recover a radically different picture. We register radially anisotropic orbits for our passive galaxy sample and more tangential motions for the star-forming objects. We relate this to *(i)* the former being accreted on already more radially-biased orbits and *(ii)* a selection effect allowing only the objects with markedly tangential motions

to maintain their star-formation properties throughout time, notwithstanding the action of various environmental processes. We also register a global effect on the orbital properties of satellites, which we prove to stem from the cosmological evolution of the host cluster.

The tension with the observational results is an interesting problem to tackle. It may arise from limits in the analytical model (although this seems unlikely) or from inadequacies in the assumptions and techniques used to extract dynamical information from observations. A step towards understanding the origin of this discrepancy could be the analysis of the galaxy sample, obtained from simulations and semi-analytic models, with observational techniques.

Obtaining ever more accurate predictions on the properties of ordinary, luminous matter is a great challenge in the field of cosmological simulations. As proved by the existence of numerous comparison projects, such as those of Frenk et al. (1999), O'Shea et al. (2005), Knebe et al. (2011), Scannapieco et al. (2011), the need for a consensus on the treatment of hydrodynamics, small-scale astrophysical processes and on the identification of collapsed structures is widespread. At the same time, semi-analytical models continuously refine their techniques in order to provide always a better match to different properties of observed galaxy populations. In this thesis, we have presented a powerful tool expected to provide a much more accurate modelling of star-forming gas, besides being already effective in dark-matter-only simulations. In addition, we have derived predictions from simulations and semi-analytic models that can be tested against observations, to ultimately add another piece of information to our knowledge of galaxy evolution.

Relevant quantities for cubic spline softening

Here we report the full expression of some of the quantities introduced in Sec. 3.2.

If the density associated to the particles is given by $\rho_i = \sum_{j=0}^N m_j W_{ij}$, where W_{ij} is defined in Eq. 3.2, the expressions for the potential and force kernel read, respectively:

$$\phi(r, h) = \begin{cases} \frac{16}{h} \left[\frac{1}{3}q^2 - \frac{3}{5}q^4 + \frac{2}{5}q^5 \right] - \frac{14}{5} \frac{1}{h}, & 0 \leq q < 0.5 \\ \frac{1}{h} \left[\frac{1}{15}q + \frac{32}{3}q^2 - 16q^3 + \frac{48}{5}q^4 - \frac{32}{15}q^5 \right] - \frac{16}{5} \frac{1}{h}, & 0.5 \leq q < 1 \\ -\frac{1}{r}, & 1 \leq q \end{cases} \quad (\text{A.1})$$

and

$$\phi'(r, h) = \begin{cases} -\frac{32}{h^2} \left[\frac{1}{3}q - \frac{6}{5}q^3 + q^4 \right], & 0 \leq q < 0.5 \\ -\frac{1}{h^2} \left[-\frac{1}{15}q^{-2} + \frac{64}{3}q - 48q^2 + \frac{192}{5}q^3 - \frac{32}{3}q^4 \right], & 0.5 \leq q < 1 \\ -\frac{1}{r^2}, & 1 \leq q \end{cases} \quad (\text{A.2})$$

where $q = r/h$.

The other quantities entering the definition of the correction term (Equations 3.8, 3.9 and 3.10) are:

$$\frac{\partial \phi}{\partial h} = \begin{cases} \frac{16}{h^2} \left[-q^2 + 3q^4 - \frac{12}{5}q^5 \right] + \frac{14}{5} \frac{1}{h^2}, & 0 \leq q < 0.5 \\ \frac{1}{h^2} \left[-32q^2 + 64q^3 - 48q^4 - \frac{64}{5}q^5 \right] + \frac{16}{5} \frac{1}{h^2}, & 0.5 \leq q < 1 \\ 0, & 1 \leq q \end{cases} \quad (\text{A.3})$$

$$\frac{\partial W}{\partial h} = \frac{8}{\pi h^4} \begin{cases} 3 \left[-1 + 10q^2 - 12q^3 \right], & 0 \leq q < 0.5 \\ 6 \left[-1 + 4q - 5q^2 + 2q^3 \right], & 0.5 \leq q < 1 \\ 0, & 1 \leq q \end{cases} \quad (\text{A.4})$$

$$\frac{\partial W}{\partial \mathbf{r}} = \frac{48}{\pi h^4} \frac{\mathbf{r}}{r} \begin{cases} -2q + 3q^2, & 0 \leq q < 0.5 \\ -1 + 2q - q^2, & 0.5 \leq q < 1 \\ 0. & 1 \leq q \end{cases} \quad (\text{A.5})$$

Bibliography

Aarseth S. J., Hénon M., Wielen R. (1974). A comparison of numerical methods for the study of star cluster dynamics. *A&A*, 37:183–187.

Abazajian K. et al. (2003). The First Data Release of the Sloan Digital Sky Survey. *AJ*, 126:2081–2086.

Abel T., Bryan G. L., Norman M. L. (2000). The Formation and Fragmentation of Primordial Molecular Clouds. *ApJ*, 540:39–44.

Athanassoula E., Fady E., Lambert J. C., Bosma A. (2000). Optimal softening for force calculations in collisionless N-body simulations. *MNRAS*, 314:475–488.

Babcock H. W. (1939). The rotation of the Andromeda Nebula. *Lick Observatory Bulletin*, 19:41–51.

Bagla J. S. (2002). TreePM: A Code for Cosmological N-Body Simulations. *Journal of Astrophysics and Astronomy*, 23:185–196.

Bagla J. S., Khandai N. (2009). The Adaptive TreePM: an adaptive resolution code for cosmological N-body simulations. *MNRAS*, 396:2211–2227.

Baldry I. K., Balogh M. L., Bower R. G., Glazebrook K., Nichol R. C., Bamford S. P., Budavari T. (2006). Galaxy bimodality versus stellar mass and environment. *MNRAS*, 373:469–483.

Baldry I. K., Glazebrook K., Brinkmann J., Ivezić Ž., Lupton R. H., Nichol R. C., Szalay A. S. (2004). Quantifying the Bimodal Color-Magnitude Distribution of Galaxies. *ApJ*, 600:681–694.

Ball N. M., Loveday J., Brunner R. J. (2008). Galaxy colour, morphology and environment in the Sloan Digital Sky Survey. *MNRAS*, 383:907–922.

- Balogh M. L., Baldry I. K., Nichol R., Miller C., Bower R., Glazebrook K. (2004). The Bimodal Galaxy Color Distribution: Dependence on Luminosity and Environment. *ApJ*, 615:L101–L104.
- Bamford S. P. et al. (2009). Galaxy Zoo: the dependence of morphology and colour on environment. *MNRAS*, 393:1324–1352.
- Barnes J., Hut P. (1986). A hierarchical $O(N \log N)$ force-calculation algorithm. *Nature*, 324:446–449.
- Bate M. R., Burkert A. (1997). Resolution requirements for smoothed particle hydrodynamics calculations with self-gravity. *MNRAS*, 288:1060–1072.
- Battaglia G., Helmi A., Tolstoy E., Irwin M., Hill V., Jablonka P. (2008). The Kinematic Status and Mass Content of the Sculptor Dwarf Spheroidal Galaxy. *ApJ*, 681:L13–L16.
- Baugh C. M. (2006). A primer on hierarchical galaxy formation: the semi-analytical approach. *Reports on Progress in Physics*, 69:3101–3156.
- Berger M. J., Colella P. (1989). Local adaptive mesh refinement for shock hydrodynamics. *Journal of Computational Physics*, 82:64–84.
- Bertschinger E. (1985). Self-similar secondary infall and accretion in an Einstein-de Sitter universe. *ApJS*, 58:39–65.
- Binney J. (2004). Discreteness effects in cosmological N-body simulations. *MNRAS*, 350:939–948.
- Binney J., Knebe A. (2002). Two-body relaxation in cosmological simulations. *MNRAS*, 333:378–382.
- Binney J., Tremaine S. (1987). *Galactic dynamics*.
- Biviano A. (2006). Mass Profiles of Galaxy Clusters from the Projected Phase-space Distribution of Cluster Members. *ArXiv Astrophysics e-prints*.
- Biviano A. (2008). Galaxy systems in the optical and infrared. *ArXiv e-prints*.
- Biviano A., Katgert P. (2004). The ESO Nearby Abell Cluster Survey. XIII. The orbits of the different types of galaxies in rich clusters. *A&A*, 424:779–791.
- Biviano A., Poggianti B. M. (2009). The orbital velocity anisotropy of cluster galaxies: evolution. *A&A*, 501:419–427.
- Blanton M. R., Eisenstein D., Hogg D. W., Schlegel D. J., Brinkmann J. (2005). Relationship between Environment and the Broadband Optical Properties of Galaxies in the Sloan Digital Sky Survey. *ApJ*, 629:143–157.
- Blumenthal G. R., Faber S. M., Primack J. R., Rees M. J. (1984). Formation of galaxies and large-scale structure with cold dark matter. *Nature*, 311:517–525.

- Bode P., Ostriker J. P., Xu G. (2000). The Tree Particle-Mesh N-Body Gravity Solver. *ApJS*, 128:561–569.
- Boselli A., Gavazzi G. (2006). Environmental Effects on Late-Type Galaxies in Nearby Clusters. *PASP*, 118:517–559.
- Bosma A. (1978). *The distribution and kinematics of neutral hydrogen in spiral galaxies of various morphological types*. PhD thesis, PhD Thesis, Groningen Univ., (1978).
- Bower R. G., Balogh M. L. (2004). The Difference Between Clusters and Groups: A Journey from Cluster Cores to Their Outskirts and Beyond. *Clusters of Galaxies: Probes of Cosmological Structure and Galaxy Evolution*, page 325.
- Boylan-Kolchin M., Springel V., White S. D. M., Jenkins A., Lemson G. (2009). Resolving cosmic structure formation with the Millennium-II Simulation. *MNRAS*, 398:1150–1164.
- Bryan G. L., Norman M. L. (1997). Simulating X-Ray Clusters with Adaptive Mesh Refinement. In D. A. Clarke & M. J. West, editor, *Computational Astrophysics; 12th Kingston Meeting on Theoretical Astrophysics*, volume 123 of *Astronomical Society of the Pacific Conference Series*, page 363.
- Cen R., Ostriker J. P. (1992). Galaxy formation and physical bias. *ApJ*, 399:L113–L116.
- Chandrasekhar S. (1942). *Principles of stellar dynamics*.
- Chandrasekhar S. (1943). Dynamical Friction. I. General Considerations: the Coefficient of Dynamical Friction. *ApJ*, 97:255.
- Cole S. (1991). Modeling galaxy formation in evolving dark matter halos. *ApJ*, 367:45–53.
- Cole S., Aragon-Salamanca A., Frenk C. S., Navarro J. F., Zepf S. E. (1994). A Recipe for Galaxy Formation. *MNRAS*, 271:781.
- Colless M. et al. (2001). The 2dF Galaxy Redshift Survey: spectra and redshifts. *MNRAS*, 328:1039–1063.
- Cooley J. W., Tukey J. W. (1965). An algorithm for the machine calculation of complex Fourier series. *Math. Comput.*, 19:297–301.
- Davis M., Efstathiou G., Frenk C. S., White S. D. M. (1985a). The evolution of large-scale structure in a universe dominated by cold dark matter. *ApJ*, 292:371–394.
- Davis M., Efstathiou G., Frenk C. S., White S. D. M. (1985b). The evolution of large-scale structure in a universe dominated by cold dark matter. *ApJ*, 292:371–394.

- de Bernardis P. et al. (2000). A flat Universe from high-resolution maps of the cosmic microwave background radiation. *Nature*, 404:955–959.
- De Lucia G., Blaizot J. (2007). The hierarchical formation of the brightest cluster galaxies. *MNRAS*, 375:2–14.
- Dehnen W. (2001). Towards optimal softening in three-dimensional N-body codes - I. Minimizing the force error. *MNRAS*, 324:273–291.
- Dehnen W., Read J. I. (2011). N-body simulations of gravitational dynamics. *European Physical Journal Plus*, 126:55.
- Di Matteo T., Springel V., Hernquist L. (2005). Energy input from quasars regulates the growth and activity of black holes and their host galaxies. *Nature*, 433:604–607.
- Diemand J., Kuhlen M., Madau P. (2006). Early Supersymmetric Cold Dark Matter Substructure. *ApJ*, 649:1–13.
- Diemand J., Moore B., Stadel J., Kazantzidis S. (2004). Two-body relaxation in cold dark matter simulations. *MNRAS*, 348:977–986.
- Dolag K., Borgani S., Murante G., Springel V. (2009). Substructures in hydrodynamical cluster simulations. *MNRAS*, 399:497–514.
- Dolag K., Murante G., Borgani S. (2010). Dynamical difference between the cD galaxy and the diffuse, stellar component in simulated galaxy clusters. *MNRAS*, 405:1544–1559.
- Dressler A. (1980). Galaxy morphology in rich clusters - Implications for the formation and evolution of galaxies. *ApJ*, 236:351–365.
- Dressler A. (1986). The morphological types and orbits of H I-deficient spirals in clusters of galaxies. *ApJ*, 301:35–43.
- Duffy A. R., Schaye J., Kay S. T., Dalla Vecchia C. (2008). Dark matter halo concentrations in the Wilkinson Microwave Anisotropy Probe year 5 cosmology. *MNRAS*, 390:L64–L68.
- Dyer C. C., Ip P. S. S. (1993). Softening in N-body simulations of collisionless systems. *ApJ*, 409:60–67.
- Evrard A. E. (1988). Beyond N-body - 3D cosmological gas dynamics. *MNRAS*, 235:911–934.
- Ewald P. (1921). *Annalen der Physik*, 369:253.
- Faltenbacher A. (2010). The impact of environment on the dynamical structure of satellite systems. *MNRAS*, 408:1113–1119.

- Farouki R., Shapiro S. L. (1981). Computer simulations of environmental influences on galaxy evolution in dense clusters. II - Rapid tidal encounters. *ApJ*, 243:32–41.
- Font A. S. et al. (2008). The colours of satellite galaxies in groups and clusters. *MNRAS*, 389:1619–1629.
- Frenk C. S. et al. (1999). The Santa Barbara Cluster Comparison Project: A Comparison of Cosmological Hydrodynamics Solutions. *ApJ*, 525:554–582.
- Frenk C. S., White S. D. M., Davis M., Efstathiou G. (1988). The formation of dark halos in a universe dominated by cold dark matter. *ApJ*, 327:507–525.
- Friedmann A. (1922). Über die Krümmung des Raumes. *Z. Phys.*, 10:377–386.
- Gavazzi G. (1989). 21 centimeter study of spiral galaxies in the Coma supercluster. II - Evidence for ongoing gas stripping in five cluster galaxies. *ApJ*, 346:59–67.
- Geiss J., Reeves H. (1972). Cosmic and Solar System Abundances of Deuterium and Helium-3. *A&A*, 18:126.
- Gill S. P. D., Knebe A., Gibson B. K., Dopita M. A. (2004). The evolution of substructure - II. Linking dynamics to environment. *MNRAS*, 351:410–422.
- Gingold R. A., Monaghan J. J. (1977). Smoothed particle hydrodynamics - Theory and application to non-spherical stars. *MNRAS*, 181:375–389.
- Godunov S. K. (1959). A Difference Scheme for Numerical Solution of Discontinuous Solution of Hydrodynamic Equations. *Math. Sbornik*, 47:271–306.
- Gondolo P. (2004). Introduction to Non-Baryonic Dark Matter. *ArXiv Astrophysics e-prints*.
- Gunn J. E., Gott, III J. R. (1972). On the Infall of Matter Into Clusters of Galaxies and Some Effects on Their Evolution. *ApJ*, 176:1.
- Guo Q. et al. (2011). From dwarf spheroidals to cD galaxies: simulating the galaxy population in a Λ CDM cosmology. *MNRAS*, 413:101–131.
- Guth A. H. (1981). Inflationary universe: A possible solution to the horizon and flatness problems. *Phys. Rev. D*, 23:347–356.
- Hernquist L. (1990). An analytical model for spherical galaxies and bulges. *ApJ*, 356:359–364.
- Hernquist L., Barnes J. E. (1990). Are some N-body algorithms intrinsically less collisional than others? *ApJ*, 349:562–569.
- Hernquist L., Bouchet F. R., Suto Y. (1991). Application of the Ewald method to cosmological N-body simulations. *ApJS*, 75:231–240.

- Hockney R. W., Eastwood J. W. (1981). *Computer Simulation Using Particles*.
- Hogg D. W. et al. (2004). The Dependence on Environment of the Color-Magnitude Relation of Galaxies. *ApJ*, 601:L29–L32.
- Host O., Hansen S. H., Piffaretti R., Morandi A., Ettori S., Kay S. T., Valdarnini R. (2009). Measurement of the Dark Matter Velocity Anisotropy in Galaxy Clusters. *ApJ*, 690:358–366.
- Hubble E. (1929). A Relation between Distance and Radial Velocity among Extra-Galactic Nebulae. *Proceedings of the National Academy of Science*, 15:168–173.
- Jeans J. H. (1902). The Stability of a Spherical Nebula. *Royal Society of London Philosophical Transactions Series A*, 199:1–53.
- Katgert P., Biviano A., Mazure A. (2004). The ESO Nearby Abell Cluster Survey. XII. The Mass and Mass-to-Light Ratio Profiles of Rich Clusters. *ApJ*, 600:657–669.
- Katz N., Weinberg D. H., Hernquist L. (1996). Cosmological Simulations with TreeSPH. *ApJS*, 105:19.
- Kauffmann G., Colberg J. M., Diaferio A., White S. D. M. (1999). Clustering of galaxies in a hierarchical universe - I. Methods and results at $z=0$. *MNRAS*, 303:188–206.
- Kauffmann G. et al. (2003). The dependence of star formation history and internal structure on stellar mass for 10^5 low-redshift galaxies. *MNRAS*, 341:54–69.
- Kauffmann G., White S. D. M., Guiderdoni B. (1993). The Formation and Evolution of Galaxies Within Merging Dark Matter Haloes. *MNRAS*, 264:201.
- Kauffmann G., White S. D. M., Heckman T. M., Ménard B., Brinchmann J., Charlot S., Tremonti C., Brinkmann J. (2004). The environmental dependence of the relations between stellar mass, structure, star formation and nuclear activity in galaxies. *MNRAS*, 353:713–731.
- Knebe A., Green A., Binney J. (2001). Multi-level adaptive particle mesh (MLAPM): a c code for cosmological simulations. *MNRAS*, 325:845–864.
- Knebe A. et al. (2011). Haloes gone MAD: The Halo-Finder Comparison Project. *MNRAS*, 415:2293–2318.
- Komatsu E. et al. (2011). Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Interpretation. *ApJS*, 192:18.
- Kravtsov A. V., Klypin A., Hoffman Y. (2002). Constrained Simulations of the Real Universe. II. Observational Signatures of Intergalactic Gas in the Local Supercluster Region. *ApJ*, 571:563–575.

- Lapi A., Cavaliere A. (2011). Self-similar Dynamical Relaxation of Dark Matter Halos in an Expanding Universe. *ApJ*, 743:127.
- Larson R. B., Tinsley B. M., Caldwell C. N. (1980). The evolution of disk galaxies and the origin of S0 galaxies. *ApJ*, 237:692–707.
- Lemaître G. (1927). Un Univers homogène de masse constante et de rayon croissant rendant compte de la vitesse radiale des nébuleuses extra-galactiques. *Annales de la Société Scientifique de Bruxelles*, 47:49–59.
- Lemaître G. (1931). Expansion of the universe, A homogeneous universe of constant mass and increasing radius accounting for the radial velocity of extra-galactic nebulae. *MNRAS*, 91:483–490.
- Lemze D. et al. (2011). Profiles of Dark Matter Velocity Anisotropy in Simulated Clusters. *ArXiv e-prints*.
- Liddle A. R. (2002). Inflationary cosmology: theory and phenomenology. *Classical and Quantum Gravity*, 19:3391–3401.
- Liddle A. R., Lyth D. H. (2000). *Cosmological Inflation and Large-Scale Structure*.
- Łokas E. L., Mamon G. A. (2003). Dark matter distribution in the Coma cluster from galaxy kinematics: breaking the mass-anisotropy degeneracy. *MNRAS*, 343:401–412.
- Lucy L. B. (1977). A numerical approach to the testing of the fission hypothesis. *AJ*, 82:1013–1024.
- Mahdavi A., Geller M. J., Böhringer H., Kurtz M. J., Ramella M. (1999). The Dynamics of Poor Systems of Galaxies. *ApJ*, 518:69–93.
- Mamon G. A., Łokas E. L. (2005). Dark matter in elliptical galaxies - II. Estimating the mass within the virial radius. *MNRAS*, 363:705–722.
- Merlin E., Buonomo U., Grassi T., Piovan L., Chiosi C. (2010). EvoL: the new Padova Tree-SPH parallel code for cosmological simulations. I. Basic code: gravity and hydrodynamics. *MNRAS*, 513:A36+.
- Merritt D. (1996). Optimal Smoothing for N-Body Codes. *AJ*, 111:2462–+.
- Meszaros P. (1974). The behaviour of point masses in an expanding cosmological substratum. *A&A*, 37:225–228.
- Mo H., van den Bosch F. C., White S. (2010). *Galaxy Formation and Evolution*.
- Monaghan J. J. (1992). Smoothed particle hydrodynamics. *ARA&A*, 30:543–574.
- Monaghan J. J., Lattanzio J. C. (1985). A refined particle method for astrophysical problems. *A&A*, 149:135–143.

- Moore B., Governato F., Quinn T., Stadel J., Lake G. (1998). Resolving the Structure of Cold Dark Matter Halos. *ApJ*, 499:L5+.
- Moore B., Katz N., Lake G., Dressler A., Oemler A. (1996). Galaxy harassment and the evolution of clusters of galaxies. *Nature*, 379:613–616.
- Murante G., Giovali M., Gerhard O., Arnaboldi M., Borgani S., Dolag K. (2007). The importance of mergers for the origin of intracluster stars in cosmological simulations of galaxy clusters. *MNRAS*, 377:2–16.
- Navarro J. F., Frenk C. S., White S. D. M. (1996). The Structure of Cold Dark Matter Halos. *ApJ*, 462:563.
- Neto A. F. et al. (2007). The statistics of Λ CDM halo concentrations. *MNRAS*, 381:1450–1462.
- Oemler, Jr. A. (1974). The Systematic Properties of Clusters of Galaxies. Photometry of 15 Clusters. *ApJ*, 194:1–20.
- O’Shea B. W., Bryan G., Bordner J., Norman M. L., Abel T., Harkness R., Kritsuk A. (2004). Introducing Enzo, an AMR Cosmology Application. *ArXiv Astrophysics e-prints*.
- O’Shea B. W., Nagamine K., Springel V., Hernquist L., Norman M. L. (2005). Comparing AMR and SPH Cosmological Simulations. I. Dark Matter and Adiabatic Simulations. *ApJS*, 160:1–27.
- Padmanabhan T. (1993). *Structure Formation in the Universe*.
- Penzias A. A., Wilson R. W. (1965). A Measurement of Excess Antenna Temperature at 4080 Mc/s. *ApJ*, 142:419–421.
- Perlmutter S. et al. (1999). Measurements of Omega and Lambda from 42 High-Redshift Supernovae. *ApJ*, 517:565–586.
- Power C., Navarro J. F., Jenkins A., Frenk C. S., White S. D. M., Springel V., Stadel J., Quinn T. (2003). The inner structure of Λ CDM haloes - I. A numerical convergence study. *MNRAS*, 338:14–34.
- Press W. H., Schechter P. (1974). Formation of Galaxies and Clusters of Galaxies by Self-Similar Gravitational Condensation. *ApJ*, 187:425–438.
- Price D. J., Monaghan J. J. (2007). An energy-conserving formalism for adaptive gravitational force softening in smoothed particle hydrodynamics and N-body codes. *MNRAS*, 374:1347–1358.
- Quilis V. (2004). A new multidimensional adaptive mesh refinement hydro + gravity cosmological code. *MNRAS*, 352:1426–1438.
- Rasia E., Tormen G., Moscardini L. (2004). A dynamical model for the distribution of dark matter and gas in galaxy clusters. *MNRAS*, 351:237–252.

- Richstone D. O. (1976). Collisions of galaxies in dense clusters. II - Dynamical evolution of cluster galaxies. *ApJ*, 204:642–648.
- Riess A. G. et al. (1998). Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant. *AJ*, 116:1009–1038.
- Robertson H. P. (1935). Kinematics and World-Structure. *ApJ*, 82:284.
- Robertson H. P. (1936a). Kinematics and World-Structure II. *ApJ*, 83:187.
- Robertson H. P. (1936b). Kinematics and World-Structure III. *ApJ*, 83:257.
- Rogerson J. B., York D. G. (1973). Interstellar Deuterium Abundance in the Direction of Beta Centauri. *ApJ*, 186:L95.
- Romeo A. B. (1998). Modelling gravity in N-body simulations of disc galaxies. Optimal types of softening for given dynamical requirements. *A&A*, 335:922–928.
- Rosswog S. (2009). Astrophysical smooth particle hydrodynamics. *New A Rev.*, 53:78–104.
- Rubin V. C., Ford W. K. J., Thonnard N. (1980). Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 / $R = 4$ kpc/ to UGC 2885 / $R = 122$ kpc/. *ApJ*, 238:471–487.
- Sales L. V., Navarro J. F., Lambas D. G., White S. D. M., Croton D. J. (2007). Satellite galaxies and fossil groups in the Millennium Simulation. *MNRAS*, 382:1901–1916.
- Saro A., Borgani S., Tornatore L., Dolag K., Murante G., Biviano A., Calura F., Charlot S. (2006). Properties of the galaxy population in hydrodynamical simulations of clusters. *MNRAS*, 373:397–410.
- Scannapieco C. et al. (2011). The Aquila comparison Project: The Effects of Feedback and Numerical Methods on Simulations of Galaxy Formation. *ArXiv e-prints*.
- Silk J. (1968). Cosmic Black-Body Radiation and Galaxy Formation. *ApJ*, 151:459.
- Skibba R. A. et al. (2009). Galaxy Zoo: disentangling the environmental dependence of morphology and colour. *MNRAS*, 399:966–982.
- Smoot G. F. et al. (1992). Structure in the COBE differential microwave radiometer first-year maps. *ApJ*, 396:L1–L5.
- Solanes J. M., Manrique A., García-Gómez C., González-Casado G., Giovanelli R., Haynes M. P. (2001). The H I Content of Spirals. II. Gas Deficiency in Cluster Galaxies. *ApJ*, 548:97–113.

- Sommer-Larsen J., Vedel H., Hellsten U. (1998). The Structure of Isothermal, Self-gravitating, Stationary Gas Spheres for Softened Gravity. *ApJ*, 500:610.
- Spergel D. N. et al. (2007). Three-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Implications for Cosmology. *ApJS*, 170:377–408.
- Spitzer, Jr. L., Hart M. H. (1971). Random Gravitational Encounters and the Evolution of Spherical Systems. I. Method. *ApJ*, 164:399–+.
- Springel V. (2005). The cosmological simulation code GADGET-2. *MNRAS*, 364:1105–1134.
- Springel V. (2010). E pur si muove: Galilean-invariant cosmological hydrodynamical simulations on a moving mesh. *MNRAS*, 401:791–851.
- Springel V., Hernquist L. (2003). Cosmological smoothed particle hydrodynamics simulations: a hybrid multiphase model for star formation. *MNRAS*, 339:289–311.
- Springel V. et al. (2005). Simulations of the formation, evolution and clustering of galaxies and quasars. *Nature*, 435:629–636.
- Springel V., White S. D. M., Tormen G., Kauffmann G. (2001a). Populating a cluster of galaxies - I. Results at $z=0$. *MNRAS*, 328:726–750.
- Springel V., Yoshida N., White S. D. M. (2001b). GADGET: a code for collisionless and gasdynamical cosmological simulations. *NewA*, 6:79–117.
- Steinmetz M., White S. D. M. (1997). Two-body heating in numerical galaxy formation experiments. *MNRAS*, 288:545–550.
- Stompor R. et al. (2001). Cosmological Implications of the MAXIMA-1 High-Resolution Cosmic Microwave Background Anisotropy Measurement. *ApJ*, 561:L7–L10.
- Tanaka M., Goto T., Okamura S., Shimasaku K., Brinkmann J. (2004). The Environmental Dependence of Galaxy Properties in the Local Universe: Dependences on Luminosity, Local Density, and System Richness. *AJ*, 128:2677–2695.
- Teyssier R. (2002). Cosmological hydrodynamics with adaptive mesh refinement. A new high resolution code called RAMSES. *A&A*, 385:337–364.
- Theis C. (1998). Two-body relaxation in softened potentials. *A&A*, 330:1180–1189.
- Toomre A., Toomre J. (1972). Galactic Bridges and Tails. *ApJ*, 178:623–666.
- Tormen G. (1997). The rise and fall of satellites in galaxy clusters. *MNRAS*, 290:411–421.

- Truelove J. K., Klein R. I., McKee C. F., Holliman, II J. H., Howell L. H., Greenough J. A., Woods D. T. (1998). Self-gravitational Hydrodynamics with Three-dimensional Adaptive Mesh Refinement: Methodology and Applications to Molecular Cloud Collapse and Fragmentation. *ApJ*, 495:821–+.
- van den Bosch F. C., Pasquali A., Yang X., Mo H. J., Weinmann S., McIntosh D. H., Aquino D. (2008). Satellite Ecology: The Dearth of Environment Dependence. *ArXiv e-prints*.
- van der Marel R. P., Magorrian J., Carlberg R. G., Yee H. K. C., Ellingson E. (2000). The Velocity and Mass Distribution of Clusters of Galaxies from the CNOC1 Cluster Redshift Survey. *AJ*, 119:2038–2052.
- Verde L. et al. (2002). The 2dF Galaxy Redshift Survey: the bias of galaxies and the density of the Universe. *MNRAS*, 335:432–440.
- Walker A. G. (1937). Kinematics and World-Structure III. *Proceedings of the London Mathematical Society*, 42:90–127.
- Wang J., White S. D. M. (2007). Discreteness effects in simulations of hot/warm dark matter. *MNRAS*, 380:93–103.
- Wang L., Li C., Kauffmann G., De Lucia G. (2007). Modelling and interpreting the dependence of clustering on the spectral energy distributions of galaxies. *MNRAS*, 377:1419–1430.
- Weinmann S. M., Kauffmann G., von der Linden A., De Lucia G. (2010). Cluster galaxies die hard. *MNRAS*, 406:2249–2266.
- Weinmann S. M., van den Bosch F. C., Yang X., Mo H. J., Croton D. J., Moore B. (2006). Properties of galaxy groups in the Sloan Digital Sky Survey - II. Active galactic nucleus feedback and star formation truncation. *MNRAS*, 372:1161–1174.
- Wetzel A. R. (2011). On the orbits of infalling satellite haloes. *MNRAS*, 412:49–58.
- White M., Scott D., Silk J. (1994). Anisotropies in the Cosmic Microwave Background. *ARA&A*, 32:319–370.
- White S. D. M. (1994). Formation and Evolution of Galaxies: Les Houches Lectures. *ArXiv Astrophysics e-prints*.
- White S. D. M., Frenk C. S. (1991). Galaxy formation through hierarchical clustering. *ApJ*, 379:52–79.
- White S. D. M., Frenk C. S., Davis M., Efstathiou G. (1987). Clusters, filaments, and voids in a universe dominated by cold dark matter. *ApJ*, 313:505–516.
- White S. D. M., Navarro J. F., Evrard A. E., Frenk C. S. (1993). The baryon content of galaxy clusters: a challenge to cosmological orthodoxy. *Nature*, 366:429–433.

- White S. D. M., Rees M. J. (1978). Core condensation in heavy halos - A two-stage theory for galaxy formation and clustering. *MNRAS*, 183:341–358.
- Wojtak R., Łokas E. L. (2010). Mass profiles and galaxy orbits in nearby galaxy clusters from the analysis of the projected phase space. *MNRAS*, 408:2442–2456.
- Wojtak R., Łokas E. L., Mamon G. A., Gottlöber S. (2009). The mass and anisotropy profiles of galaxy clusters from the projected phase-space density: testing the method on simulated data. *MNRAS*, 399:812–821.
- Woodward P. R., Colella P. (1984). A piecewise parabolic method for gas dynamical simulations. *J. Comp. Phys*, 54:174.
- Xu G. (1995). A New Parallel N-Body Gravity Solver: TPM. *ApJS*, 98:355–+.
- Yang J., Turner M. S., Schramm D. N., Steigman G., Olive K. A. (1984). Primordial nucleosynthesis - A critical comparison of theory and observation. *ApJ*, 281:493–511.
- Zel'Dovich Y. B. (1970). Gravitational instability: An approximate theory for large density perturbations. *A&A*, 5:84–89.
- Zwicky F. (1933). Die Rotverschiebung von extragalaktischen Nebeln. *Helvetica Physica Acta*, 6:110–127.

Acknowledgements

I thank Klaus Dolag for tirelessly trying to infuse some optimism into my views of things. Either that, or the instinct of survival, had a profound effect on me. In particular, I would like to thank him for his support and vicinity in the last few months. I thank Martin Asplund, who, although not directly involved in my Ph.D, has shown interest in the progress of my work and was capable of empathy in difficult moments. I thank Lauro Moscardini for being supportive throughout these years, even from a distance, and for writing the letters for my postdoc applications. For the same reasons I thank Thorsten Naab, who also patiently read the manuscript of my first paper and helped improving it. For their help on my work on gravitational softening I also thank Daniel Price, Steffen Knollmann, Michael Hilz, Mike Boylan-Kolchin and Jasjeet Bagla. For the work on galaxy orbits I acknowledge suggestions and contributions from Simon White, Chervin Laporte, Gerard Lemson, Alessia Gualandris, Laura Sales and Andrea Biviano. I am grateful to Klaus Dolag, Alessia Gualandris, Silvia Fabello, Laura Porter, Rob Yates, Katarina Markovič and Jacopo Ghiglieri for reading each a part of this thesis and providing comments on different aspects, according to their expertise. Also, I thank the kind Martin Henze, Irina Thaler, Ludwig Oser and Hannelore Hämmerle for their help with the Zusammenfassung. A special thank to Cornelia Rickl, Maria Depner and Gabriele Kratschmann for making everyone's life easier with their kindness and impeccable assistance, often extending beyond their duties.

I thank my family, for their always tactful participation in my life decisions; for their being supportive and close, without being intrusive; for their not really caring about me getting the Ph.D., but rather me being alright. I thank my father, for teaching me rectitude; my mother, for teaching me strength; my brother, for teaching me responsibility (and patience). A special wish to Gabriele for getting through a difficult present with renewed enthusiasm and aspiration for the future.

These years in Munich have been of fundamental importance in my personal

growth. Also, they have been extremely valuable in terms of social life and contact with people from nearly every culture. Leaving this city and all of those who made these years unforgettable is inevitable. However, I will always carry with me this experience and I know most of the relationships I built here are strong enough to survive the absence of everyday life. For making the Munich experience so unique and memorable, I would like to thank (in random order) Rob Yates, Akila Jeesson-Daniel, Irina Zhuravleva, Michael Korbman, Francesco de Gasperin, Nicoletta Krachmalnicoff, Laura Mascetti, Davide Burlon, Ilaria Cantoro, Martin Henze, Victor Silva, Daniela Biton, Renzo Capelli, Stefano Mineo, Marco Baldi, Massimo Dotti, Luca Graziani, Alana Knapman & friends. A thank you also to Angelo and Eva Ciliberti for welcoming me to Munich and keeping an eye on me. For the mutual support of the first weeks, the hilarious year in Großhadern and many nice dinners, I thank Lucia Morganti and Veronica Biffi. Living together with Raquel Asensio and Monica Bortolani, even for just few months or weeks and with the limitations due to very different everyday schedules, has been a very pleasant surprise; the fact that we kept in touch and are still close friends fills me with joy. I thank the LMU English Drama Group for giving me the opportunity to join a colourful assembly of young, eccentric and enthusiastic literature students (an alien world, really!) sharing the passion for theatre. A special thank also to the friends of the theatre group I-talia, especially to Luisa Sartorelli, Mattia Righi, Augusto Giussani and Monica Colloca, for the time spent together rehearsing, panicking, laughing, acting with great success (eventually) and also for their vicinity in the last few months.

Finally, I would like to spend few targeted words on some people whose presence in my life I value particularly. I thank Katarina Marcovič, for what I regard as a very rooted friendship, for many interesting discussions on all possible topics and for our similar sensitivity and approach to life. I thank Laura Porter, for sharing everyday life in the office, for our mutual understanding and for being available every time I needed her help or company. I thank Lucia Morganti, for our ever growing friendship, for listening to me with always the right sensitivity, whether I am cheerful or in tears. I thank Silvia Fabello, for being my mate these years, for sharing doubts and anguish for the present and the future and for allowing me to open up every time I needed it. I thank Alessia Gualandris - “practically perfect in every way” - for being my big sister throughout these years; for investing a considerable amount of her time being my science counsellor or simply listening to me; for being a reference on pretty much every aspect of life. I thank Alessandra and Fabio, for making me feel “at home” every time we speak, for being part of my life even if the circumstances have temporarily brought us in three different cities; the fact that our bond is stronger than ever reinforces my belief that some relationships are just everlasting. I leave the last thought for Jacopo Ghiglieri, thanks to whom my last year has brightened up. How his presence in my life has been crucial in many difficult moments I cannot express in words; his benefic influence on my everyday life is instead for everyone to see.

Most of the work presented in this thesis has been produced under the effect of:

Bach, *Goldberg Variations*

Bach, *Cello Suite #1, #2, #3*

Bach, *Die Kunst der Fuge*

Bach, *Brandenburg Concertos*

Bach, *Orchestral Suites*

Beethoven, *Symphonies*

Beethoven, *Piano Concerto #5*

Beethoven, *Piano Sonatas*

Händel, *Wassermusik*

Händel, *Feuerwerksmusik*

Mendelssohn, *The Hebrides*

Mendelssohn, *Symphony #3, #4*

Mendelssohn, *The Fair Melusina*

Mozart, *Piano Concerto #3, #24, #13, #15, #11, #23, #9*

Mozart, *Piano Concerto #2, #12, #16, #8, #19, #14, #4, #27*

Mozart, *Requiem Mass in D minor*

Mozart, *Piano Sonatas KV 330, KV 331, KV 310*

Verdi, *Messa da Requiem*

Verdi, *Macbeth*

Tchaikovsky, *Symphony #4*

Tchaikovsky, *Romeo and Juliet*

Publications

Journal papers

- **Iannuzzi, F.**; Dolag, K. 2012. *On the orbital and internal evolution of cluster galaxies*. To be submitted to MNRAS
- **Iannuzzi, F.**; Dolag, K. 2011. *Adaptive gravitational softening in GADGET*. MNRAS, 417:2846-2859
- Maio, U.; **Iannuzzi, F.** 2011. *Baryon history and cosmic star formation in non-Gaussian cosmological models: numerical simulations*. MNRAS, 415:3021-3032
- Knebe, A.; Knollmann, S. R.; Muldrew, S. I.; Pearce, F. R.; Aragon-Calvo, M. A.; Ascasibar, Y.; Behroozi, P. S.; Ceverino, D.; Colombi, S.; Diemand, J.; Dolag, K.; Falck, B. L.; Fasel, P.; Gardner, J.; Gottloeber, S.; Hsu, C.-H.; **Iannuzzi, F.**; Klypin, A.; Lukic, Z.; Maciejewski, M.; McBride, C.; Neyrinck, M. C.; Planelles, S.; Potter, D.; Quilis, V.; Rasera, Y.; Read, J. I.; Ricker, P. M.; Roy, F.; Springel, V.; Stadel, J.; Stinson, G.; Sutter, P. M.; Turchaninov, V.; Tweed, D.; Yepes, G.; Zemp, M. 2011. *Halo gone MAD: The Halo-Finder Comparison Project*. MNRAS, 415:2293-2318
- Roncarelli, M.; Moscardini, L.; Branchini, E.; Dolag, K.; Grossi, M.; **Iannuzzi, F.**; Matarrese, S. 2010. *Imprints of primordial non-Gaussianities in X-ray and SZ signals from galaxy clusters*. MNRAS, 402:923-933
- Grossi, M.; Verde, L.; Carbone, C.; Dolag, K.; Branchini, E.; **Iannuzzi, F.**; Matarrese, S.; Moscardini, L. 2009. *Large-scale non-Gaussian mass function and halo bias: tests on N-body simulations*. MNRAS, 398:321-332

Conference proceedings

- **Iannuzzi, F.**; Dolag, K. 2011. *Adaptive gravitational softening in GADGET*. To appear in "Advances in Computational Astrophysics: methods, tools and outcomes", ASP Conference Series, R. Capuzzo-Dolcetta, M. Limongi and A. Tornambè eds
- Carbone, C.; Branchini, E.; Dolag, K.; Grossi, M.; **Iannuzzi, F.**; Matarrese, S.; Moscardini, L.; Verde, L. 2009. *The properties of the dark matter halo distribution in non-Gaussian scenario*. Nuclear Physics B Proceedings Supplements, Volume 194, p. 22-27