
Supermassive black holes and multiphase gas in early-type galaxies

Maximilian Eisenreich



München 2018

Supermassive black holes and multiphase gas in early-type galaxies

Maximilian Eisenreich

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

vorgelegt von
Maximilian Eisenreich
aus Rostock

München, den 29.06.2018

Erstgutachter: Prof. Dr. Simon D. M. White

Zweitgutachter: PD Dr. Klaus Dolag

Tag der mündlichen Prüfung: 14.09.2018

Contents

Zusammenfassung	xi
Abstract	xiii
1 Introduction	1
1.1 Early-type galaxies	2
1.1.1 Galaxies in general	2
1.1.2 Galaxy classification and properties of early-type galaxies	3
1.1.3 Formation and evolution of early-type galaxies	5
1.2 Supermassive black holes	9
1.2.1 Coevolution of galaxies and their supermassive black holes	11
1.3 Circumnuclear gas discs	12
1.4 Galaxy evolution simulations	14
1.4.1 Numerical methods	17
1.5 Purpose and structure of this thesis	21
2 AGN feedback, quiescence and CGM metal enrichment in ETGs	23
2.1 Introduction	23
2.2 Simulations	29
2.2.1 Numerical code and sub-resolution models	29
2.2.2 Initial conditions	34
2.3 Star formation history & black hole growth	38
2.4 Black-hole-governed ISM evolution in the galactic centre	43
2.5 Large-scale metal enrichment through feedback-driven winds	46
2.6 X-ray properties of the gas	53
2.7 Limitations of the AGN feedback model	59
2.8 Summary & conclusions	61
3 CNDs in the centres of ETGs	65
3.1 Introduction	66
3.2 Simulations	71
3.2.1 Numerical code and sub-resolution models	71
3.2.2 Initial conditions	76

3.2.3	Simulation runs	82
3.3	Star formation suppression	83
3.3.1	Comparing initial conditions and star formation efficiencies	83
3.3.2	Resolution comparison	87
3.3.3	Alternative star formation suppression mechanisms	90
3.4	Morphology, stability and velocity dispersion	97
3.5	Molecular gas fractions	108
3.6	Black hole accretion and feedback	113
3.7	Discussion	116
3.7.1	Constraints & caveats	122
3.8	Summary & conclusions	127
4	Pitfalls of AGN sub-resolution models in hydrodynamical simulations	131
4.1	Introduction	132
4.2	Simulations	136
4.2.1	Numerical code and sub-resolution models	136
4.2.2	Hydrodynamical solvers	138
4.2.3	Time-step-limiting methods	140
4.2.4	Initial conditions	142
4.3	AGN wind efficiency and CGM metal enrichment: a parameter study	143
4.4	Hydrodynamical variation tests	147
4.5	Summary & conclusions	160
	Summary	165
	Acknowledgements	181

List of Figures

1.1	Hubble classification scheme	3
1.2	ETG component overview	6
1.3	Baryon conversion efficiency	8
1.4	Molecular and atomic gas detections in ETGs	13
1.5	Image of the Illustris simulation	15
2.1	Initial density and cooling time profiles	37
2.2	Specific star formation rate	39
2.3	Stellar and black-hole mass evolution	41
2.4	SFR – stellar mass relation	42
2.5	Black-hole duty cycles	44
2.6	Central gas maps and phase diagrams	45
2.7	CGM temperature and metallicity maps	47
2.8	Metallicity profiles	49
2.9	O/Fe profiles	50
2.10	Gas flow evolution	52
2.11	Metal-to-gas flow ratios	54
2.12	X-ray scaling relations	56
2.13	X-ray luminosity evolution	57
2.14	X-ray surface brightness map	58
2.15	X-ray cavity	60
3.1	Observed CO maps of NGC 4429	68
3.2	Observed Toomre Q profile of NGC 4429	70
3.3	CND initial condition surface density profile	78
3.4	CND initial condition surface density maps	79
3.5	CND initial condition rotation curves	81
3.6	SF efficiency comparison: Kennicutt-Schmidt relation	84
3.7	SF efficiency comparison: star formation rate	86
3.8	SF efficiency comparison: phase diagrams	88
3.9	Resolution comparison: star formation	89
3.10	Resolution comparison: Kennicutt-Schmidt relation	91
3.11	Resolution comparison: phase diagrams	92

3.12	UV field and cosmic ray comparison: Star formation	93
3.13	UV field and cosmic ray comparison: Kennicutt-Schmidt relation	94
3.14	UV field and cosmic ray comparison: phase diagrams	95
3.15	SF efficiency comparison: surface density & temperature maps after 150 Myr	97
3.16	SF efficiency comparison: surface density & temperature maps after 15 Myr	99
3.17	SF efficiency comparison: Toomre Q & velocity dispersion maps after 150 Myr	100
3.18	Resolution comparison: surface density & temperature maps after 150 Myr	102
3.19	Resolution comparison: Toomre Q & velocity dispersion maps after 150 Myr	103
3.20	CR and SN comparison: surface density & temperature maps after 15 Myr	104
3.21	CR and SN comparison: Toomre Q & velocity dispersion maps after 15 Myr	106
3.22	SF efficiency comparison: Hydrogen species mass fractions	107
3.23	Resolution comparison: Hydrogen species mass fractions	108
3.24	SF efficiency comparison: Hydrogen species maps after 150 Myr	110
3.25	SF efficiency comparison: Hydrogen species maps after 150 Myr	111
3.26	CR and SN comparison: Hydrogen species maps after 15 Myr	112
3.27	Black-hole duty cycles	114
3.28	Star formation rate with BH accretion & feedback	115
3.29	Black-hole mass – stellar mass evolution	117
3.30	BH influence: surface density & temperature maps after 150 Myr	118
3.31	Kennicutt-Schmidt relation of Martig et al. (2013)	121
3.32	Initial driving: star formation rate	125
3.33	Initial driving: surface density & temperature maps after 15 Myr	126
4.1	Varying wind efficiency: Metallicity profiles	144
4.2	Varying wind efficiency: CGM temperature and metallicity maps	145
4.3	Varying wind efficiency: stellar and black-hole mass evolution	146
4.4	Metallicity profiles	148
4.5	Gas flow evolution	150
4.6	CGM temperature and metallicity maps after 1.5 Gyr	151
4.7	CGM temperature and metallicity maps after 4.5 Gyr	153
4.8	Central gas surface density and temperature maps	155
4.9	Specific star formation rate	157
4.10	Stellar mass evolution	158
4.11	Black-hole accretion rate evolution	159
4.12	Accreted gas temperature distribution	161

List of Tables

2.1	Initial condition parameters	36
2.2	Simulation run summary	36
2.3	Quiescent time fractions	38
3.1	CND initial condition parameters	82
3.2	CND simulation runs, part 1	83
3.3	CND simulation runs, part 2	85
4.1	Varying wind efficiency: simulation run summary	143
4.2	Different hydrodynamical implementations: simulation run summary	143

Zusammenfassung

In dieser Doktorarbeit untersuchen wir die Entwicklung des Gases elliptischer Galaxien (ETGs) und wie diese durch deren zentrale supermassive schwarze Löcher (SMBHs) beeinflusst wird. Zur Modellierung der physikalischen Prozesse, die auf das Gas einwirken, führen wir hydrodynamische Simulation mit einem modernen Smoothed-Particle-Hydrodynamics-Code (SPH-Code) durch. Wir präsentieren Simulationen einer isolierten, massiven ETG mit Modellen für Gaskühlung, Sternentstehung, stellares Feedback (in Form von Typ-Ia und Typ-II Supernovae (SNe) und Winden von AGB-Sternen), Anreicherung des Gases mit Metallen durch dieses Feedback, sowie Akkretion auf das und (kinetisches und Strahlungs-) Feedback vom zentralen SMBH. Beide Formen des SMBH-Feedbacks sind notwendig um die Sternentstehungsrate (SFR) und das SMBH-Wachstum mit Beobachtungen vereinbar zu halten. Das kinetische SMBH-Feedback verursacht außerdem Winde metallreichen Gases, die das zirkumgalaktische Medium der ETG bis zu einem Radius von ~ 30 kpc anreichern.

Dann präsentieren wir hochaufgelöste Simulationen einer dichten, molekularen zirkumnuklearen Scheibe (CND) im Zentrum einer ETG, die wir mit den beobachteten CNDs von NGC 4429 (Davis et al., 2018) und ähnlichen Systemen vergleichen. Neben den bereits oben aufgelisteten Prozessen enthalten diese Simulationen Non-Equilibrium-Kühlung, Wasserstoffchemie, interstellare UV-Strahlung, Abschirmung des Gases gegen diese, Ionisation durch kosmische Strahlung (CR), stellares Photoionisationsfeedback, ein verbessertes Sternentstehungsmodell und ein neues SMBH-Akkretionsmodell. Zudem entwickeln wir ein neues “mechanisches” SN-Feedback-Modell. Unter einer Vielzahl getesteter Umstände (verschiedene Gravitationspotentiale, UV-Intensitäten, CR-Ionisationsraten, SN-Modelle, SMBH-Akkretion und -Feedback) ist die SFR der simulierten CND immer signifikant höher als in beobachteten Systemen mit äquivalenter Oberflächendichte. Um dies zu verhindern ist ein physikalischer Prozess notwendig (z.B. Magnetfelder), der den Kollaps von Gas zu hohen Dichten verhindert.

Zuletzt vergleichen wir Simulationen von SMBH-Feedback in isolierten ETGs mit verschiedenen Feedback-Effizienzparametern, sowie verschiedenen hydrodynamischen Solvern (zwei Varianten von SPH und meshless-finite-mass), um den Effekt der grundlegenden hydrodynamischen Modelle auf die Resultate zu untersuchen. Sowohl Änderungen der Feedback-Effizienz als auch der Solver verändern den Effekt des SMBH-Feedbacks auf die Struktur, ausfließenden Winde und SFR des Gases. Während die Abhängigkeit der Resultate von der Effizienz nachvollziehbar ist, zeigt die von den Solvern eine fundamentale

Schwäche der numerischen Modellierung auf. Wir schließen daraus, dass die Resultate hydrodynamischer Simulationen mit unaufgelösten hochenergetischen Prozessen (wie SMBH-Feedback) mit Vorsicht interpretiert werden müssen, da man ihre starke Abhängigkeit von den numerischen Grundlagen in Betracht ziehen muss.

Abstract

In this thesis, we study the evolution of the multiphase gas in and around massive, quiescent early-type galaxies (ETGs), and how it is affected by their central supermassive black holes (SMBH). To model the physical processes acting on the gas, we perform hydrodynamical simulations with a modern smoothed particle hydrodynamics (SPH) code. We present simulations of an isolated, massive ETG that include models for gas cooling, star formation, stellar feedback (in the form of type Ia and type II supernovae (SNe) and winds from asymptotic giant branch stars), metal enrichment of the gas by the stellar feedback, and accretion onto and feedback (in kinetic and radiative form) from the central SMBH. We find that both forms of SMBH feedback together are necessary to keep the star formation rate (SFR) and the black hole growth within observational limits. We also find that the kinetic feedback of the SMBH is able to drive outflows of metal-rich gas into the circumgalactic medium of the ETG, enriching it out to a radius of ~ 30 kpc.

We then present high-resolution simulations of a dense, molecular circumnuclear disc (CND) in the centre of an ETG, which we compare to the observed CNDs of NGC 4429 (Davis et al., 2018) and similar systems. Besides the processes listed above, these simulations also include non-equilibrium cooling, hydrogen chemistry, interstellar UV radiation, shielding of the gas from it, cosmic ray (CR) ionisation, stellar photo-ionisation feedback, an improved star-formation model, and a new SMBH accretion model. We also implement a new “mechanical” SN feedback model. We find that, under a large range of conditions (different gravitational potentials, UV field strengths, CR ionisation rates, SN models, SMBH accretion and feedback), the simulated CND is more star-forming than the observed systems at equivalent gas surface densities. To prevent this, a physical mechanism (such as magnetic fields) is needed to stop the collapse of gas to high densities.

Finally, we compare simulations of SMBH feedback in isolated ETGs with different feedback efficiencies, as well as different hydrodynamical solvers (two flavours of SPH, and meshless-finite-mass) to study the effect of the underlying hydrodynamical models on the results. We find that changing either the SMBH feedback efficiency or the hydrodynamical solver significantly alters the effect of the SMBH feedback on the structure, outflows, and SFR of the gas. While the dependence of the results on the efficiency is straightforward, that on the hydrodynamic solver shows a fundamental weakness in the numerical modelling. We conclude that results of hydrodynamical simulations with unresolved highly energetic processes (such as SMBH feedback) need to be interpreted carefully, taking into account their strong dependence on the simulation’s fundamentals.

Chapter 1

Introduction

Since the discovery in the 1920s that many of the diffuse nebulae that can be observed in the night sky are far away systems of stars similar to our own Milky Way, i.e. galaxies (e.g. Opik, 1922; Hubble, 1926), astronomers and astrophysicists have made great strides in understanding what they are, how they differ from each other, how they came to be, and what influences their evolution. Ever better telescopes (both on the Earth and in orbit around it) have allowed us to measure the light of ever more galaxies in ever more wavelengths and in ever better resolution (both spatial and spectral). This made it possible to map the distribution of galaxies in space and time (and thereby to infer their history, at least in a statistical sense), to measure their internal dynamics in enough detail to discover the supermassive black holes at their centres, and to gain insight into the properties of their stars and their gas.

In parallel to this, computer technology was developed and rapidly improved. The resulting exponential increase in computing power not only helped to process the increasingly large amounts of observational data on galaxies, it also enabled enormous improvements in our ability to theoretically model the physical processes that govern the evolution of galaxies. In particular, numerical simulations progressed from only modelling gravity to including not just the hydrodynamical interactions of the gas, but even important complex processes such as gas cooling, star formation and the injection of thermal energy and momentum into the interstellar gas by supernovae. At the same time, the possible resolution increased drastically, resulting in modern simulations (e.g. IllustrisTNG, Nelson et al., 2018) that are modelling the evolution of thousands of galaxies while resolving the internal structure of each of them at significant detail to compare them to modern observations.

These improvements in the numerical modelling of the evolution of galaxies have resulted in a state where many of their observed statistical characteristics at several epochs of the evolution of the Universe can be reproduced within a single simulation. Nevertheless, many unsolved problems and puzzles remain in the field (see the review by Naab & Ostriker, 2017). Some of these are about the details of galactic evolution; for example, how and to what degree the energy released by growing supermassive black holes affects their host galaxies at various stages in their evolution. Others are more fundamental, e.g. how reliable the models we use for these simulations actually are; how much small changes to

their underpinnings might affect the outcomes, and if they affect them in a controllable, predictable way.

In this thesis, we do our part to answer some of these questions in the context of the late evolution of early-type galaxies, the gas they contain, the effects that their central supermassive black holes might have on that gas, and how all of this can depend on how one decides to numerically solve the hydrodynamic equations governing the gas. To provide the necessary background for these studies, the remainder of this chapter contains short overviews of the known properties of (early-type) galaxies and their evolution (section 1.1); of supermassive black holes, their accretion and feedback (section 1.2); of circumnuclear gas discs in the centres of (early-type) galaxies (section 1.3); and of the current state of galaxy simulations and the numerical methods used in them (and in our simulations specifically; section 1.4). We conclude the introduction with a summary of the structure of this thesis in section 1.5.

1.1 Early-type galaxies

1.1.1 Galaxies in general

Galaxies are usually defined as collections of a large number of stars (somewhere in the large range from 10^7 to 10^{12}) that are gravitationally bound and embedded in a halo of dark matter (DM). In this definition, the requirement for a dark matter component (or more generally that the dynamics of the system cannot be explained by the observed baryons—stars and gas—if Newtonian gravity is assumed, compare Willman & Strader, 2012) discriminates between low mass dwarf galaxies (which contain dark matter) and globular clusters (which do not). Besides the two defining components of stars and dark matter (where the former dominates the mass in the central regions, while the latter makes up the overwhelming majority of the total mass of the system), most galaxies also contain some amount of gas and dust (as well as magnetic fields and cosmic rays), which is collectively called the interstellar medium (ISM). The space between galaxies is filled with the sparse, hot gas of the intergalactic medium (IGM). The gas in the outskirts of galaxies, where the ISM transitions to the IGM, is commonly called the circumgalactic medium (CGM, see e.g. Tumlinson et al., 2017, for a review of its properties).

Finally, supermassive black holes (SMBHs) with masses between about $10^6 M_\odot$ and $10^9 M_\odot$ have been observed in the centres of many galaxies (e.g. via measurements of the stellar or gas dynamics, see Kormendy & Richstone, 1995; Kormendy & Ho, 2013). These central SMBHs are believed to be a common feature of all massive galaxies (with smaller dwarf galaxies being suspected of containing less massive intermediate mass black holes, IMBHs), and to grow chiefly through the accretion of interstellar gas. During this process of accretion they are expected to release enormous amounts of energy as radiation, turning the centre of their host galaxy into an active galactic nucleus (AGN), which can be as bright or even brighter than all of the galaxy’s stars combined (see e.g. Ho & Kormendy, 2000; Padovani et al., 2017, for reviews of AGN properties). SMBHs and their effects on




Figure 1.1: Illustration of the Hubble galaxy classification scheme with exemplary images for the main types (source: NASA/ESA).

galaxy evolution are a main focus of this thesis; hence their properties are described in more detail in section 1.2.

The sizes, masses and luminosities of galaxies vary by several orders of magnitude, from tiny dwarf galaxies up to enormous ellipticals in the centres of galaxy clusters. The overall distribution $\phi(L)$ of galaxy luminosities L , which—in contrast to their masses—can be measured directly, is well described by the Schechter function (Schechter, 1976):

$$\phi(L) = \frac{\phi_*}{L_*} \left(\frac{L}{L_*} \right)^\alpha \exp(-L/L_*), \quad (1.1)$$

with $\alpha = -1.1$, $\phi_* \approx 5 \times 10^{-3} \text{ Mpc}^{-3}$, and $L_* \approx 3 \times 10^{10} L_\odot$ for R-band luminosities (Binney & Tremaine, 2008). The vast majority of galaxies have luminosities below L_* (which is similar to the Milky Way's luminosity); hence dwarf galaxies, tiny and dim as they may be, are also the most common galaxies in the Universe.

1.1.2 Galaxy classification and properties of early-type galaxies

A common and useful way of organising galaxies is to classify them by their morphology, a scheme first introduced by Edwin Hubble (1926) and often called the Hubble tuning fork

(see Fig. 1.1 for an illustration). The scheme distinguishes between four different classes of galaxies based on their optical morphology: ellipticals (E0 to E7, where the number indicates their ellipticity, E0s being spherical), lenticulars (S0), spirals, subdivided into those without central bar structures (Sa, Sb, etc.), and those with bars (SBa, SBb, etc.), and finally irregulars (Irr, which is a catch-all category for every galaxy not fitting into any of the other categories).

Spiral galaxies, as their name suggests, are defined by their prominent spiral structure. They generally contain a central stellar bulge, which is in some cases elongated into a bar-shape (making them barred spirals). Most spirals contain significant amounts of cold, dense gas, show signs of ongoing star formation, and (as a consequence of the young stellar population), appear blue.

In contrast, elliptical galaxies are smooth and ellipsoidal with little discernible substructure. Their overall shape ranges from almost perfectly spherical (E0) to very flattened and elongated (E7, denoting an ellipticity of about 0.7). Most of them contain only small amounts of cold gas, show very few signs of ongoing star formation, and have overall very old (> 10 Gyr) stellar populations, making them appear red. While spiral galaxies are common in the lower density regions of the Universe, dense clusters of galaxies are predominantly populated by elliptical and lenticular galaxies (Dressler, 1980).

The latter (S0 galaxies) have morphologies in between those of ellipticals and spirals: They have both a disc and a bulge, but their disc is smooth and featureless instead of divided into spiral arms. Like ellipticals, they are generally old systems with few signs of recent star formation and only small amounts of cold gas. Because of these similarities, they are often grouped together with ellipticals as early-type galaxies (ETGs), while the spirals are referred to as late-type (LTGs). In recent times, more categories have been added to the classification, most notably in the addition of separate classes for dwarf galaxies (dE, dIrr, dS for dwarf ellipticals, irregulars, and spirals, respectively). This underlines the rather distinct properties of dwarf galaxies—defined as all galaxies fainter than $M_B > -18$ in absolute B-band magnitude—in relation to their more massive counterparts of similar morphology.

ETGs, which are in the focus of this thesis, share a number of properties and scaling relations. One of the most significant of these relations is the so-called fundamental plane of elliptical galaxies (Bender et al., 1992), a very tight relation between the effective radius R_e (i.e. the radius of the circle within which half of the stellar luminosity of the galaxy originates), the average surface brightness within this radius I_e , and the line-of-sight velocity dispersion within the same radius σ_{los} :

$$\log R_e = 1.24 \log \sigma_{\text{los}} - 0.82 \log I_e + \text{constant}, \quad (1.2)$$

with the pre-factors depending on the observed wavelength band (Binney & Tremaine, 2008). Projections from this plane produce the Faber-Jackson relation (Faber & Jackson, 1976) between luminosity L and velocity dispersion ($L \propto \sigma_{\text{los}}^4$), as well as the Kormendy relation between luminosity and effective radius ($L \propto R_e^{1.2}$). Most ETGs contain only little or no cold gas (about 40% may contain up to $\sim 10^9 M_\odot$ of molecular and/or atomic gas,

see Young et al., 2014, and section 1.3), but essentially all of them are embedded in hot (i.e. $T \sim 10^6 - 10^8$ K) gaseous haloes observable through its emission of X-ray photons (first discovered by Forman et al., 1979). How these large reservoirs of hot gas are kept from cooling and forming new stars, thereby changing the observed old and red ETGs into younger, star-forming galaxies, is one of the major questions about the evolution of early-type galaxies. We address this question in detail in chapter 2, where we argue that AGN feedback plays a major role in preventing these cooling flows.

Summarising its most important properties, a massive ETG can be described as follows (compare our sketch in Fig. 1.2). It is a red ellipsoid of old stars embedded in a massive halo of dark matter and surrounded by a halo of hot, ionised, X-ray emitting gas. As examples, in Fig. 1.2, we show an optical image of the stellar distribution of a typical giant elliptical galaxy (ESO 325-G004, top right), and an X-ray image of a hot gaseous halo (bottom left). In its centre, the ETG harbours a supermassive black hole, which accretes gas that cools and flows into the centre, thereby turning into an AGN, which can affect the galactic gas through highly energetic radiation, winds, and jets (see section 1.2 and chapter 2). In the top left of Fig. 1.2 we show a radio-image of the jet of the giant elliptical M87, while one of the effects of AGN on the gas can be seen in the form of cavities in the X-ray image in the bottom left image of the figure. Sometimes, the ETG contains a small, but significant amount of cold, neutral gas in its central region, often in the form of a circumnuclear disc (see section 1.3). The properties and evolution of this cold gas may be influenced by the massive stellar ellipsoid and the central SMBH (see chapter 3). We show an image of one such disc of cold gas in the centre of the ETG NGC 4429 in the bottom right of Fig. 1.2.

1.1.3 Formation and evolution of early-type galaxies

All galaxies form in the context of the cosmological evolution of the Universe. To our current best understanding, this evolution starts with the Big Bang, about 13.8 billion years ago, but for the purpose of understanding galaxy formation we can neglect the earliest phases of the Universe and take as our initial conditions the state of the Universe as it is encoded in the cosmic microwave background (CMB), about 400,000 years after the Big Bang, corresponding to a redshift of $z \sim 1100$. The CMB data (e.g. from the Planck Collaboration et al., 2016) show that the Universe was, at this point, flat, homogeneous, and isotropic, with only small density perturbations. With these measurements, and assuming the standard Λ cold dark matter (Λ CDM) model of cosmology, the large-scale evolution of the Universe is then characterised by its accelerated expansion during which the initially small density perturbations are amplified under the influence of gravity, finally collapsing into the haloes that are the birthplaces of galaxies. Larger initial over-densities collapse at earlier times. As these correspond to smaller spatial scales, the first haloes formed are small, and later merge with other haloes into increasingly large structures, in a process called hierarchical structure formation (e.g. Mo et al., 2010). The gravitational collapse of any given halo is eventually stopped by the counteracting dynamical pressure from its internal velocity dispersion in a violent relaxation process called virialisation.

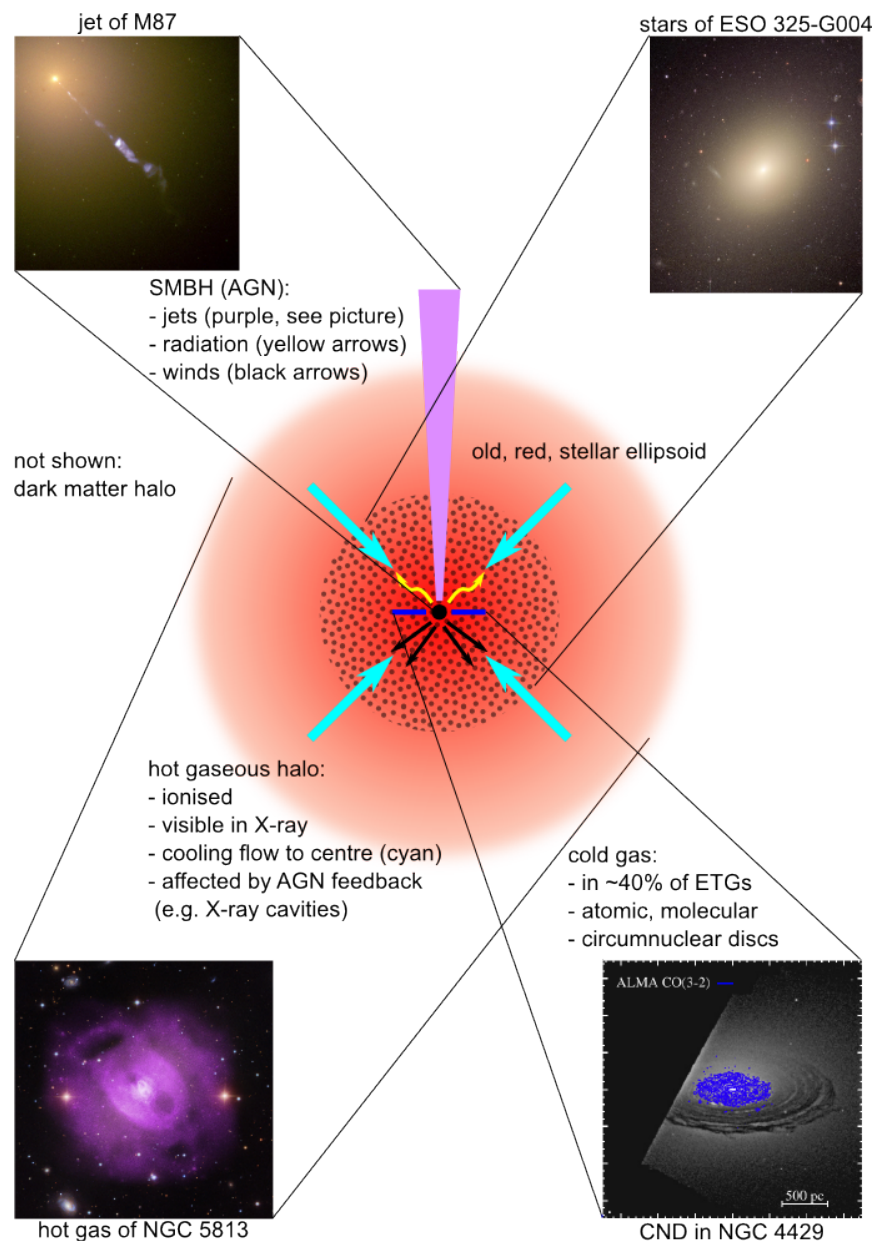


Figure 1.2: Illustration of the main components of a typical ETG (except the dark matter halo). In the corners, we show example observations. Top left: Optical image of the giant elliptical M87 overlaid with a radio image of its jet (source: NASA, Hubble Heritage Team (STScI/AURA)). Top right: Optical image showing the stars of the giant elliptical ESO 325-G004 (source: NASA, ESA, Hubble Heritage Team (STScI/AURA); J. Blakeslee). Bottom left: X-ray image of the hot gas in NGC 5813 overlaid on an optical image of the stars (source: X-ray: NASA/CXC/SAO/S.Randall et al., Optical: SDSS). AGN-created cavities can be seen in the X-ray image. Bottom right: Optical image of the central region of NGC 4429 overlaid with $^{12}\text{CO}(3-2)$ intensity contours from sub-mm observations, showing the central disc of molecular gas and dust (source: Davis et al., 2018).

The overall mass of the Universe is dominated by dark matter, which does not interact via the electromagnetic force, and hence is entirely described by gravity on astrophysical scales. Due to this, the above Λ CDM model works well to explain the early formation of galactic haloes (e.g. Binney & Tremaine, 2008). Of course, the galaxies themselves are predominantly made up of gas and stars that were once formed out of gas. Early on (and on large scales), the hydrodynamical forces acting on the gas are negligible, and it follows the flow of the dark matter. However, once the dark matter collapses into a virialised halo, the gas evolution begins to differ. As the gas is condensed in the converging flow of the collapsing halo, it can be shock-heated and stalled, forming a hot halo, which then slowly cools and further condenses through various processes depending on its density, temperature and composition: At high temperatures ($T > 10^7$ K), cooling happens predominantly via Bremsstrahlung, i.e. deceleration of free electrons in the gas. At lower temperatures ($T \sim 10^4 - 10^6$ K), collisional ionisation/excitation followed by radiative recombination/de-excitation of gas atoms dominate the cooling. Once the gas is so cold that it is mostly neutral, it needs to be enriched with metals (i.e. elements heavier than Lithium) for fine-structure excitation, or have a significant fraction of molecular gas for molecular vibrations to allow for efficient further cooling (see Binney, 1977; Rees & Ostriker, 1977; Silk, 1977, for the original works on the role of cooling gas in galaxy formation).

As an alternative channel to the cooling of shock heated gas, it is possible for cold gas to flow directly into the centre of the halo if the cooling is effective at distances farther away from the halo centre than those at which the gas would be shocked, e.g. in filaments leading into the halo (so-called cold-flow accretion, see Kereš et al., 2005; Dekel et al., 2009). In any case, eventually the gas in the halo centre cools and collapses into ever smaller, denser and colder clumps in a process that ends with the formation of stars once a clump of gas becomes dense enough for nuclear fusion to occur in its centre. Together, these stars then form the galaxy.

The gas collapsing into a halo will have some amount of net angular momentum. This leads to the natural formation of a disc when the gas condenses in the centre, which in turn leads to the formation of spiral structures through various instabilities. Through mergers of galaxies during the process of hierarchical structure formation, these spiral and disc structures can then be destroyed and the stellar orbits randomised, leading to the formation of ETGs (see the review by Naab & Ostriker, 2017). The ETGs can then grow by accreting stars formed in other galaxies by merging with them. Most of the accreted stars will settle at large distances from the galactic centre, so that ETGs grow significantly in both size and mass during their evolution, while most of their constituent stars are formed very early in their formation—just mostly not within one galaxy (e.g. Oser et al., 2010; Rodriguez-Gomez et al., 2016; Qu et al., 2017).

The above description of galaxy formation explains how gravitationally bound systems of stars can form from small density perturbations in a primordial soup of dark matter and gas, but it has a major flaw: Already at high redshifts, many small dense haloes exist in which gas cooling is very efficient. This should lead to a rapid conversion of almost all gas into stars early on in the evolution of the Universe, leaving it filled with supermassive galaxies devoid of gas, which contradicts our observations. For example, Fig. 1.3 (taken




Figure 1.3: Baryon conversion efficiency over dark matter halo mass for different redshifts (source: Moster et al., 2013).

from Moster et al., 2013) shows the baryon conversion efficiency, i.e. the fraction of baryonic mass (gas and stars) that is in the form of stars for different halo masses and redshifts. Even today ($z = 0$), its maximum is only at about 20%, i.e. the vast majority of the baryons in any halo are not in stars.

This overcooling problem (and many related problems in the theory of galaxy evolution) can be solved by including processes that counterbalance the cooling of the gas. Such processes exist in the form of feedback from stars (ionising radiation, stellar winds, and supernovae), as well as from AGN. In many numerical simulations, stellar feedback has been shown to effectively quench star formation in low-mass galaxies (to the left of the peak in Fig. 1.3), while AGN feedback is likely the mechanism responsible for preventing the overcooling problem in more massive systems (e.g. Puchwein & Springel, 2013; Vogelsberger et al., 2014; Schaye et al., 2015; Khandai et al., 2015; Davé et al., 2016, also compare chapter 2).

1.2 Supermassive black holes

A black hole (BH) is formed when a massive object becomes so compact that its gravitational field creates a region of spacetime in which the escape velocity v_{esc} surpasses the speed of light. No matter, including light, can move from within this region—the black hole—beyond its outer boundary, the event horizon, at which $v_{\text{esc}} = c$. Observational evidence for the existence of these extreme objects is abundant: Stellar-mass black holes, which are predicted to be the final stage in the evolution of many massive stars (Chandrasekhar, 1935; Oppenheimer & Volkoff, 1939; Janka, 2012, the latter for a more recent review on the fates of massive stars), have been detected in numerous X-ray binaries via emission caused by accretion from their binary partner (e.g. Remillard & McClintock, 2006; Casares & Jonker, 2014). Recently, the detection of gravitational waves from a merger of two black holes (Abbott et al., 2016) added to the evidence for stellar-mass BHs.

Supermassive black holes ($M_{\text{BH}} \gtrsim 10^5 M_{\odot}$), meanwhile, have been inferred to lie at the centres of many massive galaxies by dynamical modelling of the stellar bulges of their host galaxies, as well as other methods (e.g. Kormendy & Ho, 2013). In the Milky Way, the existence and mass of the central SMBH (Sagittarius A*) have been determined via the measurement of the proper motions and accelerations of individual stars of the nuclear star cluster surrounding the SMBH (Genzel et al., 2000; Ghez et al., 2000). The formation of SMBHs is less well constrained than that of their stellar-mass relatives. In particular, there is disagreement over the origin and properties of their initial seeds in the early universe. These might be stellar-mass black holes in dense star clusters, or intermediate mass black holes formed by the direct collapse of massive clouds of primordial gas (see e.g. Smith et al., 2017; Valiante et al., 2017, for recent reviews of the various scenarios).

Independent of their precise origin, SMBHs grow through the accretion of interstellar gas in their surroundings; a process during which a large amount of energy is converted into radiation in a broad range of frequencies (Soltan, 1982). This radiation turns the centre of the SMBH’s host galaxy into an immensely bright object, sometimes outshining

the rest of the galaxy entirely, an active galactic nucleus. One can estimate the (maximum) luminosity of an AGN by comparing the outward force of radiation pressure F_{rad} from the luminous point source that is the accreting SMBH with the inward gravitational force F_{grav} that is acting on the gas in the SMBH's neighbourhood. Assuming spherical symmetry and that the radiation pressure is due to Thompson scattering, these two forces are

$$F_{\text{rad}} = \frac{\sigma_{\text{T}} L n_{\text{e}}(r)}{4\pi r^2 c}, \quad (1.3)$$

and

$$F_{\text{grav}} = \frac{GM_{\text{BH}}\rho(r)}{r^2}, \quad (1.4)$$

where r is the distance from the source, L its luminosity, σ_{T} the Thompson cross-section, n_{e} the electron number density of the gas, ρ its density, and M_{BH} the mass of the SMBH. Demanding $F_{\text{rad}} = F_{\text{grav}}$ then leads to the maximum luminosity for the accreting black hole, the Eddington luminosity (Eddington, 1916):

$$L_{\text{Edd}} = \frac{4\pi G c m_{\text{p}}}{\sigma_{\text{T}}} M_{\text{BH}}, \quad (1.5)$$

with m_{p} the proton mass. For higher luminosities, the gas is stopped from accreting (assuming spherical symmetry), reducing the luminosity again.

How the accretion of gas onto a black hole works in detail is still a topic of ongoing research, complicated by the small scales and extreme conditions under which it takes place. However, it is reasonably certain that infalling gas first forms an accretion disc around the the black hole (Shakura & Sunyaev, 1973). Within this disc, turbulent and magnetic forces then transport angular momentum outwards, and the gas spirals towards the centre. As part of this process, a complex (and not fully understood) interplay of various processes—such as synchrotron radiation, thermal and line emissions, and inverse Compton scattering—leads to the emission, absorption, re-emission, reflection, and scattering of photons, and produces the broad spectral energy distribution of the AGN. According to the unified AGN scheme (e.g. Antonucci, 1993; Urry & Padovani, 1995), different states of accretion (e.g. high or low accretion rates), as well as varying viewing angles relative to the accretion disc then result in a host of observed types of AGN (e.g. Quasars, Blazars, Seyferts, LINERS, etc; compare e.g. Beckmann & Shrader, 2012; Heckman & Best, 2014), even when the overall accretion model is the same.

In numerical simulations of galaxies, the accretion disc around an SMBH is generally much too small to be resolved. Instead, most galaxy simulations that include black hole growth use a simplified spherical accretion model originally developed for accretion of gas onto the Sun by Hoyle & Lyttleton (1939); Bondi & Hoyle (1944); Bondi (1952). In this model, the mass growth rate of the black hole is simply

$$\dot{M}_{\text{BH}} = \frac{4\pi G^2 M_{\text{BH}}^2 \rho_{\text{gas}}}{(c_{\text{s}}^2 + v_{\text{rel}}^2)^{3/2}}, \quad (1.6)$$

where M_{BH} is the mass of the SMBH, ρ_{gas} is the density of the surrounding gas, c_s is its sound speed, and v_{rel} its relative velocity towards the SMBH. Recently, arguments have been made that the accretion onto the SMBH should not be modelled by this Bondi-Hoyle-Lyttleton formula, as it is neither spherical nor mostly determined by the (thermal and dynamical) dispersion of the gas, but instead by its angular momentum and gravitational torques (e.g. Hopkins & Quataert, 2011). Simulations using a torque-driven accretion model appear to have an advantage over those with Bondi accretion in that they manage to reproduce scaling relations between the SMBH mass and stellar bulge properties (see section 1.2.1) without the need for AGN feedback (Anglés-Alcázar et al., 2013, 2017).

1.2.1 Coevolution of galaxies and their supermassive black holes

The energy released by an AGN is enormous ($L_{\text{Edd}} \sim 10^{46} \text{ erg s}^{-1}$ for $M_{\text{BH}} = 10^8 M_{\odot}$), sometimes comparable to the binding energy of its entire host galaxy (e.g. Silk & Rees, 1998), hence it is plausible that it can significantly influence the evolution of its host galaxy, even if only a small fraction of the radiated energy couples with the galactic gas. The AGN can affect its surrounding gas through multiple channels: Direct interaction between the released photons (especially UV and X-rays) and the gas in the form of photo-ionisation, Compton scattering and radiation pressure can strongly heat the gas and accelerate it away from the galactic centre in massive blast waves. Particularly at high accretion rates (and therefore high AGN luminosities), this radiative heating can potentially unbind the entire gas content of a galaxy, creating massive outflows and quenching the star formation in the affected galaxy, as well as preventing further SMBH growth.

Besides radiation, the AGN also drives out gas from the central accretion disc in mechanical form, as winds and jets. While the latter are strongly collimated streams of relativistic plasma that can pierce through the galactic gas and reach hundreds of kiloparsecs out into the IGM before they dissipate their energy (as is, for example, the case for the radio jet of M87, see Fig. 1.2), winds from the broad line emission region of the accretion disc (e.g. Crenshaw et al., 2003; Moe et al., 2009; Dunn et al., 2010) are less rapid ($\sim 10,000 \text{ km s}^{-1}$) and wider, injecting momentum directly into the ISM surrounding the SMBH. The impact of AGN feedback on the ISM can be observed directly in the form of X-ray cavities expanding outwards from galactic centres (e.g. Churazov et al., 2000; Bîrzan et al., 2004; Forman et al., 2007, and see Fig. 1.2).

Aside from these direct signatures of AGN feedback, the interaction of SMBHs and their host galaxies is evidenced by several scaling relations that show how the stellar spheroid of a galaxy and its central black hole coevolve. The mass of the SMBH M_{BH} is tightly correlated to both the mass of the stellar bulge M_{bulge} (which corresponds to the total stellar mass in ETGs) and its velocity dispersion σ_{bulge} (e.g. Kormendy & Ho, 2013; Saglia

et al., 2016, the former of which the values in the equations below are taken from):

$$\frac{M_{\text{BH}}}{10^9 M_{\odot}} = 0.310 \left(\frac{\sigma_{\text{bulge}}}{200 \text{ km s}^{-1}} \right)^{4.38}, \quad (1.7)$$

$$\frac{M_{\text{BH}}}{10^9 M_{\odot}} = 0.49 \left(\frac{M_{\text{bulge}}}{10^{11} M_{\odot}} \right)^{1.17}. \quad (1.8)$$

1.3 Circumnuclear gas discs

Circumnuclear gas discs (CNDs) consist of multiphase gas and dust, are commonly observed in the centres of many galaxies, and especially associated with particularly active systems, such as (ultra)-luminous infrared galaxies (ULIRGs), AGN, and galaxy mergers (e.g. Downes & Solomon, 1998; Hicks et al., 2013; Medling et al., 2014). In these active galaxies, in which they are significantly more common than in quiescent galaxies (Hicks et al., 2013), CNDs have masses between approximately $10^7 M_{\odot}$ and $10^{10} M_{\odot}$, and sizes of a few tens to several hundreds of parsec (Hicks et al., 2009; Medling et al., 2014). They are connected to starbursts and cospatial young stellar discs (e.g. Davies et al., 2007; Medling et al., 2014), as well as SMBH accretion and AGN activity (e.g. Prieto et al., 2005; Hicks et al., 2009).

These observations lead to a scenario in which CNDs form from inflowing gas from larger scales (through secular evolution or due to mergers), grow in mass until they become gravitationally unstable, and then rapidly turn into stars whose feedback then starts to expel most of the gas from the centre. Meanwhile, the gas disc might form mini-spirals (Prieto et al., 2005), which fuel the accretion disc of the central SMBH on sub-parsec scales, resulting in AGN activity. However, the validity and details (e.g. the causal and temporal relation between circumnuclear starburst and AGN activity) of this scenario are still uncertain and the subject of ongoing research, not only observationally, but also with numerical simulations (e.g. Wada, 2001; Hopkins & Quataert, 2010; Wada et al., 2016; Schartmann et al., 2018).

While CNDs are much more common in active, starbursting galaxies (Hicks et al., 2013), multiphase (including cold, molecular) gas discs have also been observed in a significant fraction of quiescent early-type galaxies. ETGs were originally believed to be essentially devoid of any ISM, but this view changed considerably over time with detections of interstellar ionized gas via optical spectroscopy (e.g. Minkowski & Osterbrock, 1959; Osterbrock, 1960), distributed neutral gas via radio observations (e.g. Balkowski et al., 1972; Knapp et al., 1985), hot gaseous haloes via X-ray emission (e.g. Forman et al., 1979), dust via both infrared emission (e.g. Knapp et al., 1985) and optical absorption (e.g. Ebneter et al., 1988), and finally molecular gas via CO (carbon monoxide) emission (e.g. Sage & Wrobel, 1989; Wiklind et al., 1995).

The dense, multiphase gas of ETGs has been studied in detail in two recent surveys: SAURON (de Zeeuw et al., 2002) and ATLAS^{3D} (Cappellari et al., 2011). In the SAURON survey, Combes et al. (2007) detected CO emission (a tracer of molecular gas) in 28% of




Figure 1.4: Measured molecular hydrogen (left panel, source: Young et al., 2011) and atomic hydrogen (right panel, source: Serra et al., 2012) masses in ETGs of the ATLAS^{3D} sample over their respective K-band magnitudes. The upper limits of non-detections are shown as small crosses in the left panel, and as downward-pointing empty triangles in the right panel; all other symbols represent detections.

the sampled ETGs. Within the ATLAS^{3D} survey, Young et al. (2011) found CO emission in about 22% of their sample galaxies, indicating molecular gas masses mostly in between $10^7 M_\odot$ and $10^9 M_\odot$ (compare Fig. 1.4, left panel). Furthermore, comparable masses of atomic hydrogen (HI) have been detected in about 40% of the sample’s ETGs that are outside of the Virgo galaxy cluster (Serra et al., 2012, compare Fig.

reffig:young2011, right panel). Overall, at least 40% of local ETGs appear to contain atomic and/or molecular gas (Young et al., 2014). This gas (especially the molecular part) is mostly located within the central kiloparsec of its host galaxy (e.g. Boizelle et al., 2017). While it can have various different morphologies (from discs and rings to irregular disrupted shapes), in most cases both the molecular and the atomic gas are in the form of regularly rotating discs that are aligned with the central stellar population (Serra et al., 2012; Alatalo et al., 2013). The molecular gas discs are further generally cospatial with dust discs (Alatalo et al., 2013). In short, a significant fraction of quiescent ETGs contains CNDs, which are similar to those observed in active galaxies in many properties (such as masses and sizes). We show an image of one of these in Fig. 1.2.

However, there are some important differences: While the CNDs in active galaxies are associated with starbursts, those in ETGs are, of course, in quiescent systems and show lower than average star formation efficiencies (Davis et al., 2014). Hence, while the CNDs in active galaxies appear to be unstable to gravitational collapse and star formation, those in ETGs must be stabilised against collapse, even though they can reach similar masses and densities. A common theory is that they are stabilised by the deeper gravitational potential

well in the centres of their massive host galaxies, as it results in steep velocity curves and correspondingly high shear in the rotating gas, preventing collapse (“morphological quenching”, see Martig et al., 2009, 2013). However, this stabilisation mechanism requires the generation of significant amounts of turbulence in the gas, but the CNs in ETGs appear to be dynamically cold (i.e. they have low levels of turbulence, Boizelle et al., 2017), with some of them showing extremely low velocity dispersions (Davis et al., 2017, 2018). We study this puzzle in more detail in chapter 3 of this thesis.

1.4 Galaxy evolution simulations

The physical processes that determine the evolution of galaxies and their ISM are—with the exception of a few simple cases—too non-linear, complex, and interconnected with each other to describe them analytically. Instead, numerical simulations are employed to study these astronomical phenomena from a theoretical perspective. By self-consistently evolving the distribution and properties of matter in discrete timesteps using numerical descriptions of the relevant fundamental forces (i.e.

gravity and hydrodynamics), these tools enabled the successful theoretical investigation of the evolution of galaxies on many scales.

On the largest scales, cosmological simulations evolve a large, representative fraction of the volume of the observable Universe from initial conditions at very high redshifts ($z \gtrsim 100$), at which the density fluctuations still follow linearly from those encoded in the CMB, to more recent periods, most often to the current age of the Universe ($z = 0$). Originally only able to track the dark matter evolution (e.g. in the Millenium simulations, Springel et al., 2005b), more modern cosmological simulations now also include baryonic matter, through hydrodynamics for the gas, as well as approximate models for the more involved baryonic physics such as star formation and feedback (see Dubois et al., 2014; Vogelsberger et al., 2014; Schaye et al., 2015; Bocquet et al., 2016; Nelson et al., 2018, for some of the most recent of these simulations). The large volumes covered by these simulations allow for the study of the statistical properties of galaxies (e.g. their mass function and their distribution in the larger structure of the Universe) while their long timescales give insight into the long-term formation and evolution of these galaxies (e.g. their star formation and merger histories). As an example, a visualisation of the dark matter and gas distribution in the Illustris simulation (Vogelsberger et al., 2014) at $z = 0$ is shown in Fig. 1.5.

However, these large-scale simulations come with a significant drawback: the enormous spatial and temporal scales limit their computationally feasible resolution. Structures smaller than $\sim 1 - 10$ kpc or less massive than $\sim 10^6 - 10^8 M_\odot$ are unresolved. Due to this, the physical processes that happen on the smaller scales of the ISM but affect the global properties of the galaxies, such as star formation and stellar feedback, have to be modelled with sub-resolution schemes that estimate their larger-scale impact depending on the larger-scale properties of the system. These sub-resolution models generally come with some freedom in how to implement them, particularly in the form of free parameters that




Figure 1.5: Dark matter (left half) and gas (right half) density projection of the Illustris simulation volume at $z = 0$ (source: Illustris Collaboration, 2014)

determine the efficiencies of the unresolved processes. These parameters are usually tuned such that the simulation results fit certain important statistical galaxy properties (e.g. the stellar mass – halo mass relation), which weakens the simulations’ predictive power as it becomes questionable to what degree the results depend on the specific set of parameters chosen.

On somewhat smaller scales are simulations of individual galaxies or galaxy mergers, which come in two forms: cosmological zoom-in simulations and idealised, isolated galaxy simulations. In the former, a galactic halo is selected from the end-results of a low resolution, pure dark matter, cosmological simulation. The selected halo is then traced back through time to the initial condition, taking note of all matter that ever comes into close contact with it (e.g. all smaller haloes that merge into it during the simulation). This information is used to define a region of interest in the initial condition that corresponds to the selected final halo. In this region, the resolution is then increased, and the dark matter is partially replaced by gas (according to the observed cosmological baryon fraction) to allow the inclusion of baryonic physics. By rerunning the simulation, the evolution of the galaxy that forms in the selected halo (and its progenitors) can now be simulated at a much higher resolution than what is possible in a full cosmological hydrodynamical simulation, though at the cost of losing the statistical power of simulating thousands of galaxies over a large range of masses. A recent example of zoom-in cosmological simulations are the FIRE (Hopkins et al., 2014) and FIRE-2 (Hopkins et al., 2017) simulations, which evolve a sample of galaxies between dwarf and Milky-Way masses at mass resolutions reaching $m_{\text{gas}} \sim 200 M_{\odot}$ (where m_{gas} is the mass of individual gas fluid elements) for the dwarfs, and about $m_{\text{gas}} \sim 5000 M_{\odot}$ for Milky-Way-sized galaxies.

Another way of simulating individual galaxies is to construct idealised initial conditions of an isolated galaxy with the observed properties of the type of galaxy one is interested in (e.g. massive ETGs) and then evolve these with gravity, hydrodynamics, and other relevant baryonic processes. The resolution that can be reached in these simulations depends on the mass (and size) of the galaxy in question, with dwarf galaxy simulations reaching the best resolutions (compare e.g. the $m_{\text{gas}} \sim 4 M_{\odot}$ in Hu et al., 2016), but is generally comparable to that of zoom-in simulations of similar galaxies. The main distinction between idealised and zoom-in simulations is that the former allows for more control over the specifics of the simulated galaxy properties and enables testing the influence of these properties in a constrained manner, while the latter has the advantage of naturally including environmental effects (e.g. mergers, or the galaxy falling into a massive cluster) on galactic evolution. The simulations we use in chapter 2 are of idealised, isolated ETG initial conditions. Those in chapter 3, while only of a CND in a galactic centre instead of a whole galaxy (and therefore on much smaller scales and with significantly better resolution), are conceptually similar to these isolated galaxy simulations, in that we construct an idealised initial condition and ignore effects from larger scales.

Simulations of individual galaxies can reach much higher resolution than large-scale cosmological simulations, which allows the simulators to resolve the internal structure of galaxies in greater detail and to model important processes such as star formation with a reduced number of free parameters tuned to large-scale properties (see e.g. Hopkins

et al., 2014). However, star formation and other processes (such as feedback or black hole accretion) still happen below the resolution limit of these simulations, and as such have to be implemented as effective, sub-resolution schemes, still leaving some freedom of choosing how to formulate their implementations. The fundamental challenge of all these simulations is the large dynamical range of galaxy evolution physics. Galaxies themselves extend over tens of kiloparsecs, but their evolution is strongly influenced both by processes on the scales of Astronomical Units (e.g. star formation and black hole growth, as well as the respective feedback processes), as well as their wider environment of neighbouring galaxies and clusters on Megaparsec scales. As it is impossible to capture this whole dynamic range with the currently available computing power, some effects always have to be modelled sub-resolution, and these have to be constrained either by information from other, smaller-scale models or observations, or by their effects on larger, resolved scales. The optimal approach to any given simulation is not obvious (resulting in lively debates), but should depend on the specific questions one is trying to answer with it.

1.4.1 Numerical methods

In this section, we give a short overview of the numerical methods most commonly used for simulations of galaxies. We describe the specific methods we used in the different chapters of this thesis in more detail in the “simulations” sections of the corresponding chapters. The dark matter, stellar, and black hole components of a simulated system are generally assumed to only interact gravitationally. In the commonly used N-body method, they are described by a finite number N of discrete particles, which are assumed to be collisionless, i.e. two-body interactions between particles are negligible on the time-scale of interest. Hence, the dynamical evolution of each particle i is governed entirely by the gravitational potential Φ_i that results from all other particles j :

$$\Phi_i = -G \sum_{j=0 (j \neq i)}^N \frac{m_j}{\sqrt{|\vec{x}_i - \vec{x}_j|^2 + \epsilon^2}}, \quad (1.9)$$

where m_j is the mass of particle j , \vec{x}_i and \vec{x}_j are the comoving coordinates of particles i and j , respectively, and ϵ is the gravitational softening length, which is introduced to avoid artificial two-body interactions and the resulting extreme accelerations when two N-body particles come close to each other. The N-body particles are usually much more massive and much fewer than the objects they represent (e.g. star particles often have the masses of whole stellar populations instead of individual stars), which necessitates the softening.

Directly solving the gravitational accelerations using the N-body formula in equation 1.9 scales quadratically with the number of particles, which drastically limits the computationally possible dynamical range. To alleviate this problem, a tree-based approximation method (Barnes & Hut, 1986) is commonly used, which relies on a multi-pole expansion of equation 1.9 after the inter-particle distance $|\vec{x}_i - \vec{x}_j|$ of which only the linear term is kept for particles further away than a chosen distance (a free parameter controlling the accuracy of the model). This method improves the scaling to $\mathcal{O}(N \log N)$.

If gas is included in the simulations, the gravity calculations have to be supplemented by hydrodynamical interactions. The two “classic” approaches used in galaxy simulations for solving these interactions are Eulerian grid-based methods on the one hand, and Lagrangian particle-based methods on the other. In the grid-based schemes, the simulation volume is discretized into a grid of stationary cells across which the fluid equations (i.e. mass, momentum and energy conservation) are then solved in their Eulerian form. To enable a larger dynamical range, most grid codes used for astrophysical simulations employ “adaptive mesh refinement” (AMR), a technique that allows the refinement and de-refinement of cells over the course of the simulation according to pre-set conditions (e.g. demanding a maximum mass per cell, resulting in better resolution in denser regions).

In particle-based methods, the matter is instead described through massive fluid elements, which are moving with the fluid according to their Lagrangian equations of motion. Each of these fluid elements is represented by a particle i with a mass m_i , a position \vec{r}_i , and a velocity \vec{v}_i , as well as a measure of its thermodynamic state (usually either its specific internal energy u_i or its entropic function A_i). Most particle-based codes used in modern galaxy simulations (including those performed in this thesis) then employ the method of smoothed-particle hydrodynamics (SPH) to solve the equations of motions of the particles. A general derivation of the SPH equations of motion from the Lagrangian is given for example in Hopkins (2013). Most importantly, a scalar property of a given particle i (usually either the density ρ_i or the pressure P_i) is estimated by averaging over a corresponding property of its neighbouring particles within a distance of h_i (the smoothing length), weighing each contribution with a smoothing kernel function $W(|\vec{r}_i - \vec{r}_j|, h_i)$. For example, choosing ρ_i as the average scalar property results in:

$$\rho_i = \sum_{j=1}^N m_j W(|\vec{r}_i - \vec{r}_j|, h_i), \quad (1.10)$$

where N is the number of neighbours of particle i . The kernel function is generally chosen to be spherically symmetric, as well as such that it gives higher weights to closer neighbouring particles and zero weight to neighbours j for which $|\vec{r}_i - \vec{r}_j| > h_i$. The smoothing length h_i is usually constrained by demanding the (effective) number of particles within the kernel to be constant, which makes it dependent on the particle position \vec{r}_i .

Equation 1.10 can then be used to solve the discretised Lagrangian equations of motion. Using ρ_i as the kernel-averaged scalar and the entropy A_i as the primary thermodynamic property, these take the form

$$\frac{d\vec{v}_i}{dt} = - \sum_{j=1}^N m_j \left[\frac{f_i P_i}{\rho_i^2} \nabla_i W_{ij}(h_i) + \frac{f_j P_j}{\rho_j^2} \nabla_i W_{ij}(h_j) \right], \quad (1.11)$$

where $W_{ij}(h_i) = W(|\vec{r}_i - \vec{r}_j|, h_i)$, and the pressure $P_i = A_i \rho_i^\gamma$ (with the polytropic index γ) defined via the entropy. The $f_{i/j}$ are correction terms necessary to ensure the conservation properties of the equation for variable smoothing lengths $h_{i/j}$. Defining the pressure via the entropy (as opposed to via the internal energy, i.e. $P_i = (\gamma - 1)u_i \rho_i$) and the density

(and not, e.g. the pressure) through the kernel average are in principle arbitrary choices, as they result in different, but equivalent, expressions of the same equations of motion. Different choices here correspond to different “flavours” of SPH, generally named after their principal variables, i.e. density-entropy SPH (the example in equations 1.10 and 1.11), pressure-energy SPH, etc.

While all of these are formally equivalent, they lead to different outcomes in practice, with some formulations performing clearly worse than others under certain circumstances. For example, the kernel average of the density in the density-entropy formulation leads to densities being averaged out at contact discontinuities in the fluid, while the entropy stays discontinuous, resulting in a steep artificial pressure gradient that stabilises the discontinuity. If instead the pressure is chosen as the kernel-averaged quantity (e.g. in the pressure-entropy formulation), the pressure is now smooth by definition at the contact discontinuity, preventing the artificial surface tension. In shocks fronts, where both the density and the pressure should be discontinuous, both formulations fail, necessitating the introduction of a numerical, artificial viscosity to SPH codes (see e.g. Springel, 2005; Hu et al., 2014, for details). We describe the improvements to SPH that are used in this thesis in section 2.2.1. In chapter 4, we compare two different SPH “flavours” (pressure-entropy and pressure-energy) in the context of energetic AGN feedback, and explain their differences in section 4.2.2.

Recently, the classic schemes of AMR and SPH have been getting competition by methods that are conceptually somewhere in between the two. The first of these new schemes is the moving, unstructured mesh method, which was popularised in astrophysics by the AREPO code (Springel, 2010). In this method, the fluid is represented by an unstructured grid of moving cells whose shapes are defined by a Voronoi tessellation of discrete points. The hydrodynamic equations are then solved over the volumes of the cells as in an AMR code, but the cell-generating points can be moved. This allows the method to combine many of the advantages of both Eulerian and Lagrangian schemes, e.g. they can have continuously changing resolution like SPH codes, while also treating shocks and fluid instabilities more accurately than SPH (like AMR codes).

Another newcomer with advantages over both SPH and AMR are the meshless finite volume (MFV) and meshless finite mass (MFM) methods, both introduced into numerical astrophysics by Gaburov & Nitadori (2011), and later popularised in the field by the GIZMO code (Hopkins, 2015). Like SPH, these methods are Lagrangian, as the fluid is traced by moving particles. Unlike SPH though, the hydrodynamic forces affecting a particle are not calculated by a weighted average over the contribution of its neighbouring particles as if they were all point masses, but instead by the construction of effective faces between the particle and its nearest neighbours, followed by solving the resulting Riemann problems across these faces. As the only difference between these methods and SPH is how the accelerations of particles are calculated, they can be implemented into an SPH code with only a few changes, which allows for a very close comparison between the different methods. We include the MFM method into our comparison study in chapter 4, and describe how it works in more detail in section 4.2.2.

All of the described methods of numerically approximating and solving the hydrody-

dynamic equations have varying strengths and weaknesses, as well as circumstances in which they perform better and those in which they perform worse. For example, SPH couples well to the N-body gravity calculations, allows for a continuously adaptive resolution, and is overall great at conserving mass, energy and momentum, but it also has problems properly capturing fluid mixing and shocks (e.g. Agertz et al., 2007; Springel, 2010), converges slowly, and produces significant amounts of numerical noise due to a constant (zeroth-order) error in the discretisation (e.g. Read et al., 2010; Price, 2012). Many of the problems of SPH have been strongly mitigated in “modern” implementations (e.g. Read et al., 2010; Hopkins, 2013; Hu et al., 2014, , see also section 2.2), but they are not gone. AMR codes are consistent to a higher order than SPH and better at resolving shocks, but they do not conserve angular momentum, introduce errors in their coupling to gravity, and require significantly more (computational) effort to refine their resolution (see e.g. Hopkins, 2015). The newer methods used in AREPO and GIZMO appear to outperform both SPH and AMR in most areas (e.g. they show better fluid mixing than SPH, and allow for automatic spatial resolution refinement unlike AMR, see Springel, 2010; Hopkins, 2015), but they are still numerical approximations, which have also been tested much less extensively in astrophysical contexts than the two old schemes.

While many of the problems of the various methods only come into play in certain circumstances (e.g. shocks or contact discontinuities), simulations of the evolution of galaxies are generally so complex that most or all of these circumstances eventually apply somewhere in the simulated volume. This, combined with the necessity for sub-resolution models of unresolved processes (some of which, like supernovae and AGN, are extreme energy sources), which results from the enormous range of relevant spatial and temporal scales, makes it very important to understand how different numerical schemes affect the results of simulations of the same phenomena. We address this problem in chapter 4.

The simulations we performed for this thesis required a range of different sub-resolution models, which we shortly mention here, referring to the method sections of the individual chapters in which they are introduced in detail. We model the formation of stars with a Schmidt-type model (Schmidt, 1959), in which the star formation rate is proportional to the reciprocal of the free-fall time of the gas, once it has crossed a certain threshold of density and temperature. This model is explained in more detail in section 2.2.1, while an alternative star-formation threshold is introduced in 3.2.1. The cooling of the gas is based on precomputed equilibrium cooling rate tables for various elemental abundances of the gas, as described in section 2.2.1. For the colder gas ($T < 3 \times 10^4 \text{K}$) in the high-resolution simulations of chapter 3, this model is replaced by a non-equilibrium cooling model taking into account processes in a chemical network of different hydrogen species. We describe this model in section 3.2.1.

In all of our simulations, we take into account feedback from the stars in the form of supernovae, which inject mass (in a distribution of various tracked chemical elements), energy and momentum into the surrounding gas. We use three different models for the supernova feedback: In chapters 2 and 4, we use a model of mixed momentum and thermal energy injection by Núñez et al. (2017), which we describe in section 2.2.1. In chapter 3, we either use a slight alteration of the above model, in which only thermal energy is injected,

or our own implementation of “mechanical” supernova feedback in which we inject purely momentum. This model is introduced in section 3.2.1. In the same section, we also describe the additional model for photo-ionisation by stars that we include in the simulations of chapter 3.

Finally, we include sub-resolution models for the accretion of gas and feedback by supermassive black holes. The accretion model we use in chapters 2 and 4 (a modified Bondi-Hoyle-Lyttleton model developed by Choi et al., 2012) is described in section 2.2.1. In chapter 3, we introduce our own alternative black hole accretion model, which we explain in section 3.2.1. We use the combined wind and radiative AGN feedback model by Choi et al. (2012) in all of our simulations that include feedback from the SMBH. It is described in section 2.2.1.

The output of our hydrodynamical simulations are the properties (position, velocity, mass, etc.) of all individual particles at predefined times during the simulation (every ~ 3 Myr in the simulations of chapters 2 and 4, and every ~ 150 kyr in those of chapter 3). For our analysis of these simulation data (e.g. to generate maps of the gas surface density) we make use of the GADGET-analysing framework Pygad (Röttgers & Arth, 2018).

1.5 Purpose and structure of this thesis

The overall goal of this thesis is to advance our understanding of the evolution of the gas in massive early-type galaxies, and in particular how this evolution is affected by the presence of supermassive black holes in the centres of these galaxies. For this purpose, we perform numerical simulations with a modern N-body hydrodynamics code, which includes state-of-the-art models for the cooling of the gas, its metal content, its molecular chemistry, the formation of stars out of it, the feedback effects of the formed stars on it, its accretion onto the central SMBH, and the effects on it caused by the active galactic nucleus the accreting black hole turns into.

With these simulations we first investigate the long-term influence of the AGN feedback on the gas, specifically if and how the feedback prevents the gas from cooling and forming stars, as well as how it affects outflows of gas into the CGM (chapter 2). In chapter 3, we then study the dynamical and chemical evolution of a circumnuclear gas disc in the centre of a massive ETG, where the SMBH dominates the gravitational potential in a large part of the disc. Finally, in chapter 4, we explore how the results of chapter 2 are affected by changes to the underlying method used for solving the hydrodynamic equations, and what this means for astrophysical simulations involving high-energy processes on unresolved scales in general.

Chapter 2

AGN feedback, quiescence and CGM metal enrichment in early-type galaxies

We present three-dimensional hydrodynamical simulations showing the effect of kinetic and radiative AGN feedback on a model galaxy representing a massive quiescent low-redshift early-type galaxy of $M_* = 8.41 \times 10^{10} M_\odot$, harbouring a $M_{\text{BH}} = 4 \times 10^8 M_\odot$ black hole surrounded by a cooling gaseous halo. We show that, for a total baryon fraction of $\sim 20\%$ of the cosmological value, feedback from the AGN can keep the galaxy quiescent for about 4.35 Gyr and with properties consistent with black hole mass and X-ray luminosity scaling relations. However, this can only be achieved if the AGN feedback model includes both kinetic and radiative feedback modes. The simulation with only kinetic feedback fails to keep the model galaxy fully quiescent, while one with only radiative feedback leads to excessive black-hole growth. For higher baryon fractions (e.g. 50% of the cosmological value), the X-ray luminosities exceed observed values by at least one order of magnitude, and rapid cooling results in a star-forming galaxy. The AGN plays a major role in keeping the circumgalactic gas at observed metallicities of $Z/Z_\odot \gtrsim 0.3$ within the central ~ 30 kpc by venting nuclear gas enriched with metals from residual star formation activity. As indicated by previous cosmological simulations, our results are consistent with a model for which the black hole mass and the total baryon fraction are set at higher redshifts $z > 1$ and the AGN alone can keep the model galaxy on observed scaling relations. Models without AGN feedback violate both the quiescence criterion as well as CGM metallicity constraints. This chapter is based on Eisenreich et al. (2017).

2.1 Introduction

Massive early-type galaxies in the local universe are in their vast majority old, quiescent stellar systems, which formed almost all of their stars ~ 10 Gyr ago, and show little to no signs of ongoing star formation (e.g. Kauffmann et al., 2003a). This quiescence is a

puzzle in two ways: First, in how it came to be, i.e. what mechanism quenched the star formation in these systems at high redshifts ($z \sim 2$), leaving them quiescent since then. Theoretical arguments and numerical simulations both point towards powerful feedback from the galaxies' supermassive black holes ejecting and unbinding much of their gas (the fuel for their star formation) from their haloes as the cause for quenching (e.g. Silk & Rees, 1998; Di Matteo et al., 2005; Croton et al., 2006; Bower et al., 2006; Hopkins et al., 2008; McCarthy et al., 2010; Dubois et al., 2010, 2013, 2016; Martizzi et al., 2012; Le Brun et al., 2014; Vogelsberger et al., 2014; Sijacki et al., 2015; Schaye et al., 2015), see also the relevant sections in recent reviews on galaxy formation (Somerville & Davé, 2015; Naab & Ostriker, 2017). Observations clearly indicate that most active galactic nuclei in the local universe live in ETGs (Kauffmann et al., 2003b). The direct connection between AGN activity and quenching is less clear: While some measurements in individual galaxies show clear correlations between fast, AGN-driven outflows and star formation suppression (e.g. Cano-Díaz et al., 2012; Brusa et al., 2015; Carniani et al., 2016), some show the opposite (enhanced star formation in connection with AGN winds, e.g. Cresci et al., 2015a,b). A similar division is seen in statistical studies: some find evidence for star formation suppression through AGN winds (e.g. Page et al., 2012), others do not (e.g. Delvecchio et al., 2015). Furthermore, some theoretical and numerical works also show enhanced star formation associated with AGN feedback (e.g. Gaibler et al., 2012; Silk, 2013; Zubovas et al., 2013). Nevertheless, it is observationally clear that feedback from AGN has significant impact on their host galaxies and the gas within them (see e.g. Fabian, 2012, for a review of the corresponding evidence). Alternative quenching mechanisms usually involve interactions of galaxies with their environment (e.g. Peng et al., 2010; Smith et al., 2012), e.g. through the effects of merging with another galaxy, and it is very possible that various processes aid each other in the quenching (e.g. galaxy mergers causing favourable conditions for efficient AGN feedback, see e.g. Hopkins et al. (2005, 2008)).

While AGN feedback is currently the most fashionable – and indeed probably the dominant – process for quenching galaxies, the physical explanations for quenching in massive systems have varied over time, and a thoughtful analysis would indicate that several other processes have contributed to comparable degrees:

- Firstly, there is an intrinsic physical effect, pointed out by three papers in 1977 (Rees & Ostriker, 1977; Silk, 1977; Binney, 1977) that tends to make any heating process more effective for these systems than for lower mass galaxies: Cooling rates for a hydrogen-helium plasma are reduced significantly for gas at the higher virial temperatures of massive ETGs, and these papers pointed out that there is a critical mass above which the cooling time is longer than the dynamical time. This mass is important in setting the upper mass scale for galaxies even though more modern calculations, which include the cooling effects of metal lines, have significantly altered the simple, original cooling criterion.
- Secondly, we have stripping by ambient gas in clusters and groups. The typical ETG lives in a dense cluster of galaxies, moving through the hot gas envelope seen

in thermal X-ray emission and indicative of the hot, relatively dense environment within which these systems live.

- Thirdly, type I supernovae are effective in blowing out the processed gas from the outer parts of normal ETGs. Papers by Renzini et al. (1993) and others have indicated that this is a major effect – but of course it could not prevent *central* cooling flows and central starbursts.
- Finally, in the cosmological context, gravitational heating from infall can be important, and Johansson et al. (2009) showed that this can add nearly 10^{60} erg over cosmic epochs. This heating can balance a significant fraction of the normally expected thermal gas cooling and greatly delay star formation.

While each of these processes would significantly reduce the star formation which might otherwise occur, recent simulations (see above) have conclusively shown that a central SFR of several solar masses per year would still occur absent effective suppression of central cooling flows and star bursts. It is the properly implemented AGN feedback that was found to be the dominant quenching mechanism in this domain.

In any case, galactic evolution leads to a population of elliptical galaxies that, at low redshifts, spend their time mostly quiescent, forming very few stars, while still containing significant amounts of hot gas (e.g. Forman et al., 1979; Humphrey et al., 2006; Diehl & Statler, 2007; Mulchaey & Jeltama, 2010; Danielson et al., 2012; Kim & Fabbiano, 2013, for some X-ray observations of the hot gas around ETGs). Here, an important question is exactly how much hot gas these massive, local ETGs actually contain, as the gas mass (and its density, which is related), can have a significant impact on the efficiency of different feedback processes. The hot gas mass, through its Bremsstrahlung and metal line emission, can be measured using X-ray telescopes, which by now quite a few groups have done, both for galaxies and larger-scale systems. They correlate either the X-ray luminosity of the gas directly, or related quantities like the system’s baryon fraction (its mass of stars and gas divided by its total mass including dark matter), to a variety of other quantities of the system, e.g. its stellar luminosity or mass, stellar velocity dispersion, total mass, or gas temperature (e.g. Diehl & Statler, 2007; Mulchaey & Jeltama, 2010; Dai et al., 2010; Boroson et al., 2011; Laganá et al., 2013; Kim & Fabbiano, 2013; Sarzi et al., 2013; Kim & Fabbiano, 2015; Anderson et al., 2015). These correlations show that more massive systems, i.e. systems in deeper potential wells, contain more gas, with the most massive—galaxy clusters—having a baryon fraction about equal to that measured from the cosmic microwave background (the so-called “cosmological” baryon fraction).

On the smaller scales of individual galaxies, the most commonly measured relation was between the X-ray luminosity of the hot gas and tracers of the stellar luminosity (L_B or L_K). These relations indicate that quiescent ETGs contain less gas compared to their total masses than groups and clusters, but they also show an enormous scatter of up to several orders of magnitude (e.g. Boroson et al., 2011; Sarzi et al., 2013). Only recently, large galaxy surveys like the Sloan Digital Sky Survey (SDSS) made it possible to put tighter constraints on the gas mass in the less massive haloes on the scales of individual, massive

galaxies. Anderson et al. (2015) used stacked SDSS data to measure the X-ray luminosity–mass relation down to stellar masses of $M_* = 10^{10.8} M_\odot$, while the Planck Collaboration et al. (2013) used the same data to constrain the hot gas mass via the Sunyaev-Zeldovich (SZ) effect. Together, these observations indicate that massive galaxies contain enough hot gas within their virial radii to match the cosmological baryon fraction, but that their hot haloes have to be much less concentrated than in more massive systems. Anderson et al. (2015) also show that the X-ray luminosity–mass relation follows an unbroken power law down to the scales of massive galaxies, indicating that any influence of the AGN should be through a gentle, self-regulated feedback mode, and not through much more disruptive episodes of powerful “thermal blast” feedback events.

The existence of massive hot gaseous haloes around ETGs leads to the second part of the puzzle: how do low-redshift ETGs *stay* quiescent? Without some energetic process stopping it, the hot gas in an ETG would flow towards the centre triggered by efficient gas cooling. The gas would become dense enough to form new stars continuously, breaking the galaxy’s quiescence, which is not observed (e.g. Mathews & Brighenti, 2003). Essentially the same issue exists in larger-scale structures, i.e. galaxy groups and clusters, where it is commonly known as the “cooling-flow problem”. Numerical simulations on all scales from clusters down to individual ETGs show exactly the described behaviour when no feedback processes, or only feedback from stars (i.e. supernovae and stellar winds) are included, but they also show that feedback from the central SMBH can prevent a cooling catastrophe (e.g. Ciotti & Ostriker, 1997, 2001, 2007; Brighenti & Mathews, 2006; Nagai et al., 2007; Ciotti et al., 2010; McCarthy et al., 2010; Le Brun et al., 2014).

In these simulations, the AGN feedback prevents a continuous cooling flow and keeps the system quiescent by heating the central gas and/or generating winds that decrease the central gas density. Although, at these late times, these winds are not necessarily capable of ejecting the gas from the halo completely (the systems are generally too massive and their potential wells too deep), they can still prevent any further cooling and star formation. If, in simulations, AGN-driven outflows keep ETGs quiescent, can we detect these outflows in the real universe? Especially with the rise of integral-field spectroscopy over the recent years, there have indeed been various detections of likely AGN-driven outflows of gas in different phases (ionized, atomic, molecular) from the centres of ETGs (e.g. Kehrig et al., 2012; Cano-Díaz et al., 2012; Alatalo et al., 2011, 2015; Dasyra et al., 2015; Brusa et al., 2015; Cresci et al., 2015a,b; Gomes et al., 2016; Carniani et al., 2016; Wylezalek & Zakamska, 2016), which are often explained with interactions of the AGN-driven radio jet with the surrounding medium, but also with nuclear winds and radiation pressure (Cicone et al., 2014; Morganti et al., 2015; Cheung et al., 2016).

Even with modern supercomputers, it is impossible to simultaneously resolve the scales on which black hole accretion and feedback actually take place (i.e. the \sim AU scales of accretion discs), and include a whole galaxy (or even larger scales) into a simulation. Therefore, numerical implementations of AGN feedback are always sub-resolution models designed to catch the overall effect of the unresolved feedback processes on larger scales. In the past, various groups tried several different such implementations: One of the easiest ways to model the AGN feedback is probably in the form of pure thermal feedback as

first implemented by Springel et al. (2005a). There, a fixed fraction ϵ_r of the rest-mass energy accreted by the SMBH is assumed to be converted into radiation, which then couples with the surrounding gas with a constant efficiency of ϵ_f , leading to isotropic heating of the interstellar medium in the neighbourhood of the black hole. The thermal energy input of this heating is strictly proportional to the black hole’s mass accretion rate: $\dot{E}_{\text{FB}} = \epsilon_f \epsilon_r \dot{M}_{\text{BH}} c^2$.

Already this rather simple AGN feedback model leads to a regulation of the SMBH growth and the star formation that produces galaxies, groups, and clusters much more in line with observations than simulations without any AGN feedback, but it comes with some limitations. For example, this simple thermal feedback model has a resolution-dependant, numerical “overcooling” problem: The surrounding gas can be heated to a temperature at which it cools rapidly, immediately radiating away the feedback energy, and making the feedback inefficient. To avoid this problem, Booth & Schaye (2009) modified the model to let the feedback energy accumulate until it is enough to heat a certain amount of gas by a set temperature such that losses through radiative cooling are reduced. This model was also used in the large-scale cosmological simulations cosmo-OWLS (Le Brun et al., 2014) and EAGLE (e.g. Schaye et al., 2015).

Furthermore, some authors argue that—following both theoretical and observational indications—the AGN feedback should be composed of two different modes: a radiative, “quasar” mode at high accretion rates, which corresponds to the radiative-thermal feedback described above, and a “radio” mode relating to the creation of large X-ray cavities of hot, under-dense gas by radio jets (see e.g. Churazov et al., 2005). Sijacki et al. (2007) implemented such a two-mode feedback model by combining the Springel et al. (2005a) thermal feedback at high accretion rates above a set threshold with the creation of large, hot “bubbles” (with properties dependent on the accretion rate) in the gas for accretion rates below that threshold. However, this model has led to too powerful AGN-driven ejection of gas in massive haloes in the large-scale cosmological Illustris simulation (Genel et al., 2014), leaving the haloes almost devoid of gas, in contradiction to observations.

Steinborn et al. (2015) instead chose to mimic the transition between quasar and radio mode by using the sum of two feedback efficiencies (one for each mode) with different dependencies on the accretion rate to calculate the total feedback energy, which was then distributed thermally to the surrounding gas. In a different approach, Dubois et al. (2012) used a combination of pure thermal feedback in the “quasar” mode at high accretion rates, and a jet-like feedback model in the “radio” mode at lower accretion rates. The jet feedback is modelled by distributing mass, momentum and energy in a small cylinder with the momentum aligned to the cylinder axis. This model was used in the large-scale cosmological simulation suite Horizon-AGN (Dubois et al., 2016). More recently, Weinberger et al. (2017) combined the usual purely thermal energy injection for the “quasar mode” with a purely kinetic energy and momentum injection in random directions for the “radio” mode. Yet another approach (e.g. Churazov et al., 2000; McNamara et al., 2000) assumes that the feedback energy is injected in the form of relativistic particles filling radio-bright bubbles. In this model, the amount of injected momentum is small, and most of the energy is first stored as the enthalpy of the buoyantly rising bubbles. As they cross several

pressure scale-heights, the bubbles gradually release their energy into the surrounding gas, thereby heating it (Churazov et al., 2001, 2002).

Another problem common to purely thermal models of AGN feedback is that, when they are strong enough to drive outflows and affect the gas on large scales, they often heat the interstellar, circumgalactic, and intergalactic medium to overly high temperatures compared to what is observed (e.g. Choi et al., 2014, 2015). If, instead, the feedback is implemented in a more physically motivated way, by directly modelling the interaction of the observed radiation and winds of the AGN with its surrounding gas, and specifically taking into account the momentum transfer that takes place in these interactions, this problem can be solved. Following this approach, extensive work has been done in a series of papers: Starting with first simple (Ciotti & Ostriker, 1997, 2001), then much more sophisticated (Sazonov et al., 2005; Ciotti & Ostriker, 2007) models of radiative feedback (Compton and photoionization heating plus the corresponding radiation pressure), later also including kinetic momentum feedback modelling the broad-line region winds observed in AGN (Ciotti et al., 2009, 2010; Shin et al., 2010), these authors investigated the influence of AGN feedback on the evolution of elliptical galaxies in one-dimensional (1D) simulations. Comparing both types of feedback, they found that both together are necessary to match the observed properties of ETGs, as pure radiative feedback leads to much too high SMBH masses, while pure momentum feedback is incapable of keeping the galaxy quiescent without expelling too much gas. More recently, they expanded their studies to two-dimensional simulations (Ostriker et al., 2010b; Novak et al., 2011; Gan et al., 2014), finding much more stochastic and less efficient AGN feedback than in 1D simulations, and also (as for their 1D work) testing the effects of different (constant and accretion-dependent) feedback efficiencies.

Similar models have also been developed and explored by other authors: Debuhr et al. (2011) implemented AGN feedback in the form of pure momentum input through radiation pressure, and Debuhr et al. (2012) added a kinetic feedback model describing broad-line region winds. Kim et al. (2011) used a combination of radiative feedback (using a radiative transfer technique) and mechanical feedback in the form of bipolar jet-like winds. More recently, Bieri et al. (2017) investigated the impact of winds driven by radiative AGN feedback, using radiative-hydrodynamical simulations to determine how the AGN radiation couples to the ISM. Finally, Hopkins et al. (2016) studied the effects of both broad-line region wind feedback and radiative heating (a model similar the one used in this paper) in small-scale simulations of the central region of a massive galaxy.

Choi et al. (2012) implemented the momentum and radiative feedback models from Ciotti, Ostriker et. al (see above) into the fully three-dimensional hydrodynamical code GADGET-3 (Springel, 2005), and compared it to the “standard” thermal feedback model by Springel et al. (2005a), first in simulations of isolated spiral galaxies and mergers of such (Choi et al., 2014), and later in cosmological zoom-in simulations of ETGs (Choi et al., 2015). They find that their momentum feedback implementation is much more successful at driving outflows and preventing recent star formation in local ETGs than the thermal feedback model, while also producing galaxies with X-ray (i.e. hot gas) luminosities similar to those observed, while the thermal feedback results in much higher X-ray luminosities.

In this work, we use an improved black-hole feedback model of Choi et al. (2012) to investigate the influence of AGN feedback on the late evolution of massive, quiescent elliptical galaxies, as they are observed in the local universe. To this purpose, we run controlled hydrodynamical simulations of an isolated spherical galaxy set up to closely resemble a typical observed local, massive ETG in its major properties. We examine the effects of the different parts of the AGN feedback (i.e. momentum and radiative feedback) on the star formation history of the galaxy (especially with regard to maintaining its quiescence), and on its ISM and CGM properties (especially the generation of large-scale outflows, their effect on the gaseous metal distribution, and the X-ray characteristics of the galactic hot gas; see e.g. Sijacki et al. (2007); McCarthy et al. (2010) on these feedback effects on the scales of galaxy clusters and groups), and compare the results to observational constraints.

This chapter is organized as follows: In section 2.2, we briefly describe the numerical code and the physical sub-resolution models used for our simulations, present the initial conditions of our simulated galaxy, and give an overview of the simulation runs. Afterwards, we present and discuss our simulation results by comparing the star formation history and black-hole growth (section 2.3), the ISM evolution (section 2.4), the metal enrichment of the CGM through gas flows from and to the centre (section 2.5), and the X-ray properties of the galaxy (section 2.6) between the different runs and with observational constraints. Finally, in section 2.8, we summarize our work and the conclusions we draw from its results.

2.2 Simulations

2.2.1 Numerical code and sub-resolution models

For our simulations, we use the N-body smoothed particle hydrodynamics code SPHGal (Hu et al., 2014), an improved version of GADGET-3 (see Springel, 2005, for GADGET-2, the last public version of this code). The improvements introduced in SPHGal are described below. The use of a Wendland C^4 kernel instead of the cubic spline kernel commonly used in older SPH codes including GADGET-3 allows for a higher number of neighbours in the kernel (here $N_{\text{ngb}} = 100$). This reduces the zeroth-order error in the SPH calculations and thereby improves the convergence rate of the code.

As described in section 1.4.1 of the introduction, the density-entropy formulation of SPH used in GADGET-3 and many other, older SPH codes produces artificial pressure gradients at contact discontinuities in the simulated fluid. In SPHGal, the pressure-entropy formulation of SPH is used instead, which prevents this spurious surface tension and thereby improves the accuracy of fluid mixing at contact discontinuities in the simulation. In this

formulation, the hydrodynamic equations of motion take the form of:

$$\frac{d\vec{v}_i}{dt} = - \sum_{j=1}^N m_j (A_i A_j)^{1/\gamma} \times \left[\frac{f_{ij} P_i}{P_i^{2/\gamma}} \nabla_i W_{ij}(h_i) + \frac{f_{ji} P_j}{P_j^{2/\gamma}} \nabla_i W_{ij}(h_j) \right], \quad (2.1)$$

$$f_{ij} = 1 - \left(\frac{h_i}{3A_j^{1/\gamma} m_j n_i} \frac{\partial P_i^{1/\gamma}}{\partial h_i} \right) \left[1 + \frac{h_i}{3n_i} \frac{\partial n_i}{\partial h_i} \right]^{-1}, \quad (2.2)$$

where \vec{v}_i is the velocity of particle i , N the number of its neighbours, m_j the neighbour particle's mass, A_i the entropic function, γ the polytropic index, $W_{ij} = W(\vec{x}_i - \vec{x}_j)$ the smoothing kernel, h_i the smoothing length, and $n_i = \sum_{j=1}^N W_{ij}(h_i)$ is the kernel-based particle number density. Finally, P_i is the kernel-averaged pressure, defined as:

$$P_i = \left[\sum_{j=1}^N m_j A_j^{1/\gamma} W_{ij}(h_i) \right]^\gamma \quad (2.3)$$

To improve the accuracy of the fluid mixing at strong shocks by reducing the noise of the pressure estimate (equation 2.3) there, SPHGal includes an artificial thermal energy conduction if such a shock is detected. The shock treatment of the code is further improved by changes to the implementation of the artificial viscosity. Finally, a time-step-limiting scheme by Durier & Dalla Vecchia (2012) was added to the code to ensure that neighbouring gas particles have time steps of similar lengths, and that those particles affected by a feedback process have their dynamical state updated as soon as possible. This also leads to a more accurate treatment of shocks. See section 4.2.3 (and generally chapter 4) for a more detailed discussion of the time-step limiter and its importance. Together, these changes significantly reduce the numerical artefacts present in the fluid mixing of the original GADGET version and also improve the convergence rate noticeably.

To include physical processes beyond gravity and hydrodynamics in our simulations, the code is supplemented by sub-resolution models for metallicity-dependent gas cooling, star formation, energy and momentum feedback from stars and the central SMBH, as well as metal production and diffusion.

The model for gas cooling, star formation, stellar feedback and metal enrichment was originally implemented by Scannapieco et al. (2005, 2006), and further improved and extended by Aumer et al. (2013). Each gas particle that falls below a temperature threshold set to 12,000 K and above a density threshold of $1.94 \times 10^{-23} \text{ g cm}^{-3}$ has a probability of $1 - e^{-p_{\text{SF}}}$ to turn into a star particle in the current time-step of size Δt , where p_{SF} is defined as

$$p_{\text{SF}} = \epsilon_{\text{SFR}} \sqrt{4\pi G \rho} \Delta t \quad (2.4)$$

which is larger for higher gas densities ρ . The star-formation efficiency ϵ_{SFR} is a free parameter that is necessary because the actual scales of star formation are not resolved

in the simulation. It is set to $\epsilon_{\text{SFR}} = 0.02$ in our simulations. Compare also Springel & Hernquist (2003) for the ultimate origin and motivation of this star-formation model.

The gas particles cool with a rate dependent on their current temperature, density and metal abundances. This cooling rate is precomputed for a large number of values of the above-mentioned quantities and then stored in tables, which are read out during the simulation, taking the tabulated rate for the temperature, density and abundances closest to those of the gas particle in question. The chemistry is taken into account by tracking the abundances of Hydrogen, Helium, and the nine metals most important for the cooling rate (C, N, O, Ne, Mg, Si, S, Ca and Fe) in each gas (and star) particle. Each chemical element contributes separately to the cooling rate of the gas (see Aumer et al. (2013) for details). The abundances in the gas change over time due to enrichment from stellar feedback, as well as diffusion between gas particles.

The stellar feedback implemented in the code has three major effects: it enriches the gas with metals, accelerates it (kinetic feedback), and heats it up (thermal feedback). If a star particle is tagged as giving feedback (either through supernova type Ia or II explosions, or through winds from asymptotic giant branch (AGB) stars) it distributes mass (in a mix of all tracked elements depending on the particle’s metallicity and the feedback type) to its 10 closest neighbouring gas particles. The fraction of mass that each gas particle gains is weighted with the SPH smoothing kernel of the feedback-giving star particle, and therefore depends on the distance between the two.

Following the simplifying assumption that all stars in a single population exploding in supernovae type II (SNII) do so at the same age τ_{SNII} , each newly created star particle undergoes a SNII feedback event exactly once at a time τ_{SNII} after its creation, where τ_{SNII} is only dependent on the particle’s metallicity. While this is a valid approximation for SNII feedback because the SNII progenitor stars all have very short lifetimes, the supernova type Ia (SNIa) progenitors are old white dwarfs, which can have vastly different ages at the time of their explosion. Therefore, the SNIa feedback is modelled quasi-continuously: each star particle older than 50 Myr undergoes SNIa feedback events repeatedly every 50 Myr until a maximum age of 10 Gyr, after which the feedback stops. The feedback effects of these SNIa events (i.e. the released mass, energy and momentum) decline with the age τ of the particle as τ^{-1} , thereby following the delay time distribution of the SNIa rate presented in Maoz & Mannucci (2012). The AGB feedback follows the same procedure as the SNIa feedback, but with smaller yields and a different elemental distribution in the released mass.

Through the stellar feedback, energy and momentum are injected into the surrounding gas in the following way, using a model implemented by Núñez et al. (2017). A supernova event (of each type) is assumed to eject mass in an outflow with a velocity $v_{\text{out,SN}} = 4000$ km/s, corresponding to an energy of

$$E_{\text{SN}} = \frac{1}{2} m_{\text{ejected}} v_{\text{out,SN}}^2 \quad (2.5)$$

where m_{ejected} is the total mass ejected by the star particle undergoing a supernova. It depends on the age (for SNIa) the metal composition and the total mass of the star particle,

following tabulated mass yields from Woosley & Weaver (1995) for SNII and Iwamoto et al. (1999) for SNIa. Depending on the distance between the supernova-undergoing star particle (SN particle) and the affected gas particle, this outflow is then assumed to be in one of three characteristic phases of interaction with the ambient gas. This supernova-ejecta phase then determines which fraction of E_{SN} is injected into the gas particle as kinetic and thermal energy, respectively. These phases are (in order of increasing distance to the SN particle) the momentum-conserving free-expansion phase, the Sedov-Taylor phase where 30% of E_{SN} are imparted as kinetic and 70% as thermal energy (Sedov, 1959), and the snow-plough phase where radiative cooling becomes important and reduces the total injected energy (see Núñez et al., 2017, for details).

For the AGB feedback, the feedback energy and momentum are always imparted to the neighbouring gas particles like they are in the free-expansion phase of the supernova feedback. Here, the feedback energy is calculated according to equation 2.5, but with $v_{\text{out,SN}} = 4000$ km/s replaced by $v_{\text{out,AGB}} = 25$ km/s, and is therefore much lower.

For the growth of the central SMBH of the simulated galaxy by accretion of gas and the resulting feedback we use a (slightly modified) model by Choi et al. (2012), whose accretion model is in turn based on the model by Springel et al. (2005a). The black hole is represented by a single collisionless “sink” particle that grows by absorbing nearby gas particles. The accretion rate follows the Bondi-Hoyle-Lyttleton formalism (Hoyle & Lyttleton, 1939; Bondi & Hoyle, 1944; Bondi, 1952) with several adjustments to its SPH implementation (the soft-Bondi criterion, the free-fall modification and an alternative averaging method) added by Choi et al. (2012). The inflow rate of gas onto the black hole is then given by

$$\dot{M}_{\text{BHL}} = \left\langle \frac{4\pi G^2 M_{\text{BH}}^2 \rho}{(c_s^2 + v^2)^{3/2}} \right\rangle, \quad (2.6)$$

where M_{BH} is the mass of the black hole, ρ the surrounding gas density, c_s the sound speed in the gas, v the relative speed between the black hole and the gas, and the angle brackets stand for the SPH kernel averaging. The probability of a gas particle to be part of the inflow onto the black hole in a given time step is then dependent on its kernel weight, the length of the time step, the inflow rate given above, and modifying factors corresponding to the soft-Bondi criterion and the free-fall time.

These modifiers are the following: For the soft-Bondi criterion, the accretion probability is multiplied by a factor rising linearly from 0 at $r = r_{\text{B}} + h$ to 1 (if $r_{\text{B}} \geq h$) or to $(r_{\text{B}}/h)^3$ (if $r_{\text{B}} < h$) at $r = r_{\text{B}} - h$, where r is the distance between the black hole and the gas particle, r_{B} is the Bondi radius of the black hole, and h is the SPH smoothing length of the gas particle. This takes into account the limited resolution of the SPH simulation by only allowing gas particles to be accreted that are statistically within the Bondi radius. The free-fall modification takes into account the relative free-fall times of the gas particles neighbouring the black hole to make it more likely for particles closer to the black hole (i.e. with shorter free-fall times) to be accreted than particles further away.

The inflow rate \dot{M}_{inf} is limited by the Eddington rate (Eddington, 1916)

$$\dot{M}_{\text{edd}} \equiv \frac{4\pi G M_{\text{BH}} m_{\text{p}}}{\epsilon_{\text{r}} \sigma_{\text{T}} c}, \quad (2.7)$$

which depends on the proton mass m_p , the Thompson cross-section σ_T and the radiative efficiency of the AGN feedback $\epsilon_r = L_r/\dot{M}_{\text{acc}}c^2$, which is assumed to have a fixed value of 0.1, as is commonly done (e.g. Shakura & Sunyaev, 1973). L_r is the radiated luminosity of the accreting black hole, with \dot{M}_{acc} being the rate of gas mass accreted by the black hole (i.e. incorporated into the BH mass). The inflow rate is then $\dot{M}_{\text{inf}} = \min(\dot{M}_{\text{BHL}}, \dot{M}_{\text{edd}})$.

The AGN feedback is a combination of kinetic-thermal wind feedback, and radiative feedback modelling the interaction of the SMBH's X-ray flux with the surrounding gas. It works as follows: We assume that only some of the inflowing gas \dot{M}_{inf} is ultimately accreted onto the black hole while a large part of it is blown out again due to the feedback in the form of broad-line winds. The rate with which this blown out gas is flowing out from the accretion region of the black hole is then simply $\dot{M}_{\text{outf}} = \dot{M}_{\text{inf}} - \dot{M}_{\text{acc}}$, and the energy and momentum fluxes of the AGN wind follow from conservation laws (Ostriker et al., 2010b):

$$\dot{E}_w \equiv \epsilon_w \dot{M}_{\text{acc}} c^2 = \frac{1}{2} \dot{M}_{\text{outf}} v_w^2, \quad (2.8)$$

$$\dot{p} = \dot{M}_{\text{outf}} v_w, \quad (2.9)$$

where ϵ_w is the feedback efficiency (essentially a free parameter) and v_w is the speed of the wind immediately after ejection by the AGN. From equation 2.8 follows:

$$\frac{\dot{M}_{\text{outf}}}{\dot{M}_{\text{acc}}} = \frac{2\epsilon_w c^2}{v_w^2}, \quad (2.10)$$

i.e. the ratio between outflowing and accreted mass is only dependent on the wind speed v_w and the feedback efficiency ϵ_w . Following Choi et al. (2012), we fix the wind speed at $v_w = 10,000 \text{ km s}^{-1}$, a typical velocity for observed broad line winds (e.g. Crenshaw et al., 2003; Moe et al., 2009; Dunn et al., 2010). The outflow-to-inflow ratio is now fully determined by choosing a feedback efficiency. We choose $\epsilon_w = 0.005$, which leads to 90% of the inflowing gas mass being ejected in the wind, while only 10% ultimately contribute to the black-hole growth. Numerically, this wind is implemented in such a way that gas particles that are part of the black-hole accretion region have a probability corresponding to the mass ratio $\dot{M}_{\text{outf}}/\dot{M}_{\text{inf}}$ to be ejected each time step. Similarly, they are swallowed with a probability of $\dot{M}_{\text{acc}}/\dot{M}_{\text{inf}}$. The broad line winds represented by this feedback collide with the gas in the immediate surroundings of the SMBH and create a larger, momentum-driven outflow. As this happens on scales not resolved by the simulation, we mimic it by sharing the momentum of the ejected particle equally with two of its neighbouring gas particles. The excess energy of the feedback is then distributed among the three particles as thermal energy. The direction of the momentum is set to be either parallel or anti-parallel to the angular momentum of the corresponding particle relative to the black hole (with 50% chance for each) before the ejection. We choose this direction for the momentum because the broad line outflows we mimic are stronger perpendicular to the accretion disc than parallel to it according to simulations by Proga & Kallman (2004). While the accretion disc is not resolved in our simulations, its angular momentum has to be that of the accreted gas particle.

In addition to the mechanical feedback described above, radiative feedback is used to both heat and accelerate the gas. This part of the feedback has its source in the X-ray luminosity $L_r = \epsilon_r \dot{M}_{\text{acc}} c^2$ of the AGN, which produces a flux $F_r = L_r / 4\pi r^2$ at the position of each gas particle with the distance r to the central black hole. From this flux, the heating rate $\dot{E}_{\text{X-ray}}$ of the gas is calculated using formulae by Sazonov et al. (2005) to describe the Compton and photo-ionization heating of the gas (see Choi et al., 2012, for details). Other than the accretion luminosity and the distance to the AGN, the heating rate also depends on the temperature, the proton number density and the metallicity of the influenced gas particles. The metallicity dependence was not in the original Sazonov et al. (2005) formulae, but added in Choi et al. (2017) to account for metal line absorption. Note that, while L_r is the bolometric luminosity of the AGN, only the effects of hard X-rays are taken into account. Beside the heating, the X-ray flux from the black hole also creates a radiation pressure, which is modelled as a momentum change of the gas particles

$$\dot{p}_{\text{X-ray}} = \frac{\dot{E}_{\text{X-ray}}}{c} \quad (2.11)$$

radially away from the the black hole. Last but not least, the actual Eddington force is also included in this feedback model. It is based on Thompson scattering of the AGN's radiation with the surrounding electrons in the gas, and is implemented as a momentum change of the gas particles radially away from the black hole defined as

$$\dot{p}_{\text{Edd}} = \frac{F_r N_e \sigma_T}{c} \quad (2.12)$$

where N_e is the number of electrons in the considered gas particle.

2.2.2 Initial conditions

We simulate the evolution of a spherically symmetric, isolated, large, and local early-type galaxy consisting of an old stellar population embedded in a hot gaseous halo, a dark matter halo, and containing a central supermassive black hole. These four components of our model galaxy are represented by four different particle types, all of which but the gas particles are collisionless and only interact gravitationally. The gas particles additionally interact via hydrodynamic forces.

The stellar component of the ETG has a mass of $M_* = 8.41 \times 10^{10} M_\odot$, and is modelled with $N_* = 841,000$ particles with individual masses of $m_* = 10^5 M_\odot$ and gravitational softening lengths of $\epsilon_* = 20$ pc, which trace a Hernquist density profile (Hernquist, 1990) with a scale length of $a_* = 2.21$ kpc, corresponding to an effective (i.e. projected half-mass) radius of $R_e = 4.01$ kpc. The mass and size of the stellar spheroid are scaled according to a relation by Williams et al. (2010). Each star particle has an initial age τ_{ini} and metallicity (tracked by multiple elemental abundances as described above). Ages are randomly assigned to the stars in a log-normal distribution with a mean age of $\mu = 6$ Gyr and a logarithmic spread of $\sigma = 0.1$, corresponding to the predominantly very old stellar population found in most ETGs (e.g. Thomas et al., 2005; McDermid et al., 2015, esp. Fig.

10 and Fig. 14, respectively). Note that our simulations cover a time of about 4.5 Gyr, hence an initial stellar age of 6 Gyr corresponds to 10.5 Gyr at the end of the simulation, i.e. redshift 0. The stellar metallicities follow the observations of Greene et al. (2013): they are solar at the galactic centre and decline exponentially with increasing radius such that $[\text{Fe}/\text{H}] = -0.3$ at $r = 2R_e \approx 8$ kpc, giving a linear slope for $[\text{Fe}/\text{H}]$ of 0.0375 dex/kpc.

The mass of the central SMBH is set to $M_{\text{BH}} = 4 \times 10^8 M_\odot$, following the observed $M_{\text{BH}} - M_*$ relation by Kormendy & Ho (2013). Its softening length is identical to the stellar one: $\epsilon_{\text{BH}} = \epsilon_*$. The virial dark matter mass (i.e. the dark matter mass within the virial radius) follows from the stellar mass via an abundance matching relation (Moster et al., 2013), yielding $M_{\text{DM, vir}} = 6.92 \times 10^{12} M_\odot$. This dark matter mass is distributed spherically symmetrically with a Hernquist (1990) density profile where the scale length $a_{\text{DM}} = 74.7$ kpc is determined by associating it with a NFW profile (Navarro et al., 1996) of the same $M_{\text{DM, vir}}$ and a concentration factor of $c_{\text{con}} = 9$ (see Springel et al., 2005a). This leads to a virial radius $R_{\text{vir}} = 402$ kpc (where $R_{\text{vir}} = R_{200, \text{crit}}$ i.e. the radius at which the density is 200 times the critical density). With these choices for the stellar and dark matter distribution, the dark matter mass fraction within the half-mass radius of the galaxy is about 50%, which is on the high end of, but still fully consistent with, the spread of observational values (e.g. Barnabè et al., 2011). The dark matter density profile is traced by $N_{\text{DM}} = 10^6$ dark matter particles of masses $m_{\text{DM}} = 9.71 \times 10^6 M_\odot$ and softening lengths $\epsilon_{\text{DM}} = 200$ pc. The angular momentum of the DM halo is given by the spin parameter $\lambda_{\text{DM}} = 0.033$.

The last component of the model ETG is the hot gas halo whose mass M_{gas} is determined from the stellar and DM mass by setting the total baryon fraction of the galaxy within its virial radius $f_b = (M_{\text{gas}} + M_*)/M_{\text{DM}}$. Due to the difficulty of detecting low-density gas at large distances from the galactic centre, the total baryon fraction of elliptical galaxies is not particularly well constrained (e.g. Sarzi et al., 2013). Still, there are several X-ray observations of hot gaseous haloes that show the existence of a scaling relation between the hot gas X-ray luminosity of a given ETG with some of its other properties (e.g. total mass, stellar luminosity), albeit mostly with a very large scatter (see e.g. Boroson et al., 2011; Kim & Fabbiano, 2013; Sarzi et al., 2013; Anderson et al., 2015). As $L_{\text{X, gas}}$ depends on the total hot gas mass, we use one of these relations to estimate the baryon fraction to be roughly 20% of the cosmological value (see Fig. 2.12), $f_{b, \text{cosm}} = 0.1864$ (Planck Collaboration et al., 2014). This corresponds to a hot gas mass of about $M_{\text{gas, vir}} = 1.74 \times 10^{11} M_\odot$. The hot halo is sampled with particles of the same mass and softening length as are used for the stellar component (i.e. $m_{\text{gas}} = 10^5 M_\odot$, $\epsilon_{\text{gas}} = 20$ pc), resulting in a gas particle number of $N_{\text{gas}} = 2.12 \times 10^6$. These particles trace the radial dependence of the gas density in form of a β -profile (Cavaliere & Fusco-Femiano, 1976; Jones & Forman, 1984; Eke et al., 1998) with a slope parameter $\beta = 2/3$ (following Jones & Forman, 1984) and a core radius $r_c = 0.22R_S = 9.8$ kpc (following Makino et al., 1998), where $R_S = R_{200}/c_{\text{con}} \approx 44.7$ kpc is the scale radius of the dark matter profile. The gas density profile is cut at a radius of $r_{\text{cut}} = 50r_c \approx 492$ kpc, i.e. there are no gas particles beyond this radius. From the density distribution, the temperature profile follows with the assumptions that the gas is in hydrostatic equilibrium and distributed isotropically (see Moster et al., 2011). Every

Table 2.1: Overview of important parameters of the initial conditions.

Parameter	Stars	Gas	Dark matter	Black hole
total mass M_{tot} [M_{\odot}]	8.41×10^{10}	2.12×10^{11}	9.71×10^{12}	4×10^8
virial mass M_{vir} [M_{\odot}]	8.33×10^{10}	1.74×10^{11}	6.92×10^{12}	
particle mass m [M_{\odot}]	10^5	10^5	9.71×10^6	4×10^8
number of particles N	8.41×10^5	2.12×10^6	10^6	1
softening length ϵ [kpc]	20	20	200	20

Table 2.2: Summary of simulations showing the run name and the active black-hole model components.

Name of run	no-BH	BH-W	BH-WR
BH accretion		✓	✓
BH Eddington force		✓	✓
BH wind feedback		✓	✓
BH radiative feedback			✓

gas particle is also given initial metal abundances, such that the radial metallicity profile follows that of the stars, but with a slightly lower metallicity peak of 93% solar at the centre (following Konami et al., 2014). The hot gas halo has the same spin parameter as the dark matter halo ($\lambda_{\text{gas}} = 0.033$).

With a gas particle mass of $m_{\text{gas}} = 10^5 M_{\odot}$, we are able to resolve the Jeans mass of all gas except the densest and coldest (compare Fig. 2.6, bottom row), almost all of which is star-forming and hence turned into star particles after some time (if it is not affected by accretion or feedback processes). Therefore, our resolution should be sufficient for investigating the impact of AGN feedback on the global properties of the galaxy, and the larger-scale behaviour of the hot gas. An overview of the most important initial condition parameters is given in Table 2.1. The initial density profiles of the components of the model galaxy (solid lines), as well as the cooling time of its gas (dashed line, calculated with the same cooling rate that is used in the simulation, and neglecting heating by feedback processes) are shown in Fig. 2.1, on the left and right ordinate, respectively. Within the central 5 kpc (\sim the stellar half-mass radius), the cooling time varies between ~ 100 and ~ 200 Myr.

We compare three different simulation runs in this work, one without black-hole accretion and feedback (no-BH), one with BH accretion and the wind-feedback model described above, but without the radiative feedback (BH-W), and one with the full BH feedback model (radiative heating and radiation pressure in addition to the wind feedback; BH-WR). We compare the no-BH model to the feedback models to investigate if BH feedback is necessary and/or sufficient to keep the ETG quiescent, and then compare the feedback models to probe the importance and effects of different feedback implementations on the galactic evolution. Table 2.2 summarizes these simulation runs.

To check how the mass of the hot halo affects the AGN feedback’s ability to influence the

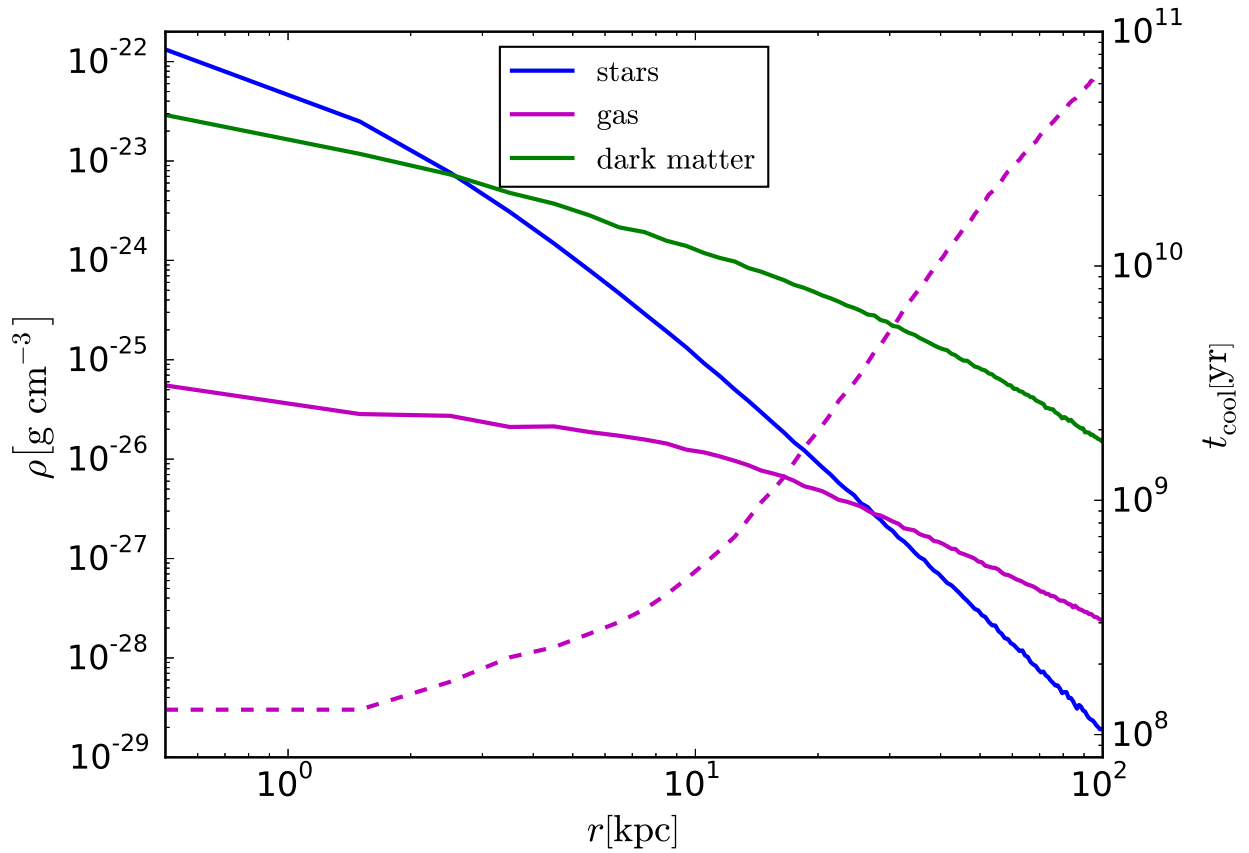


Figure 2.1: Initial radial density profiles of the stars (blue), gas (magenta), and dark matter (green) in the simulated galaxy (solid lines, left y-axis), as well as the initial cooling time profile of the gas (dashed magenta line, right y-axis). The cooling time is calculated using the same cooling rate that is used in the simulation itself, and is not taking into account gas heating due to feedback processes.

Table 2.3: Quiescent time fractions: Total time above & below the quiescence limit, as well as fraction of total time below the quiescence limit (quiescent fraction) for the different runs

Name of run	no-BH	BH-W	BH-WR
time active (Gyr)	4.12	1.63	0.60
time quiescent (Gyr)	0.23	2.80	3.83
quiescent fraction	0.05	0.63	0.87

galactic evolution, we include an additional run (BH-WR₅₀) using the BH-WR feedback model, but containing 50% of the cosmological baryon fraction in the initial condition (raising the virial gas mass to $M_{\text{gas,vir}} = 5.61 \times 10^{11} M_{\odot}$, and the number of gas particles to $N_{\text{gas}} = 6.86 \times 10^6$), which is otherwise identical to the other runs. This run lies an order of magnitude above the observed X-ray scaling relation already in the initial condition (see Fig. 2.12, top panel). It also completely fails to be quiescent (see Fig. 2.2), which shows that the amount of hot gas in ETGs has to be reduced to a low baryon fraction (or the gas density significantly reduced, compare Anderson et al. 2015) at higher redshifts for AGN feedback to be capable of regulating the late-time evolution of the galaxy. In the further discussion in this paper, we will concentrate on the runs with 20% of the cosmological baryon fraction (see Tab. 2.2).

2.3 Star formation history & black hole growth

In Fig. 2.2, we show the evolution of the total specific star formation rate (sSFR) over time for all of our runs to investigate if the black-hole feedback can efficiently keep the galaxy quiescent. As the sSFR oscillates quite strongly between individual snapshots, and is therefore noisy, we plot its running mean, averaged over 10 snapshots (about 30 Myr) each, to get a clearer picture of its time evolution. The black, horizontal line in this figure is the quiescence limit according to Franx et al. (2008), i.e. $\text{sSFR} = 0.3\tau_{\text{Hub}}^{-1}$ where $\tau_{\text{Hub}} \equiv H_0^{-1}$ is the Hubble time.

Using this definition, the galaxy is actively star-forming for the whole simulation time except the first ~ 200 Myr (5% of the total time) if we use the no-BH model, where the short quiescent period at the beginning of the simulation is just the time the initial hot gas needs to cool down in the centre of the galaxy (compare Fig. 2.1). However, if we include black-hole feedback, i.e. in the BH-W and BH-WR models, the overall star-formation rate of the galaxy is reduced significantly; the ETG is now quiescent for about 63% and 87% of the total simulation time in the BH-W and BH-WR models, respectively (Tab. 2.3). This shows that black-hole feedback is necessary to keep an isolated early-type galaxy quiescent at low redshifts.

The star-formation history of the galaxy differs distinctly in the two feedback models: The additional radiative feedback in the BH-WR model not only reduces the overall star-formation of the galaxy compared to the BH-W model, and keeps it quiescent for a




Figure 2.2: Running mean with a step size of about 30 Myr (10 snapshots) of the specific star formation rate over time for the different models (see legend). The black horizontal line is the quiescence limit (Franx et al., 2008). Without black-hole feedback, the galaxy is actively star-forming at almost all times; including the feedback reduces the overall star formation significantly. With only wind feedback, there are alternating long periods of complete quiescence and star formation, while the full feedback model including radiation leads to fast oscillation between star-forming and quiescent states of the galaxy. For 50% of the cosmological baryon fraction, the AGN feedback fails to prevent star formation.

significantly longer fraction of the simulation time (compare Tab. 2.3), it also increases the time-variability of the sSFR. While in the BH-W model, the sSFR rises and falls moderately over long times, in the BH-WR model, it oscillates much faster and over a much higher range of values (with more than two orders of magnitude between the lowest and highest sSFR values in the 30 Myr running mean. Looking at individual snapshots without the running mean, the picture is largely the same: the sSFR repeatedly drops to zero in the BH-WR model, while its scatter is always less than one order of magnitude (usually below a factor of 2) in the BH-W model and only in the 10% range in the no-BH model).

This difference is caused by the additional isotropic heating provided by the radiative feedback: All the potentially star-forming dense gas in the galactic centre is heated to high temperatures by Compton scattering whenever a feedback event happens, strongly reducing the star-formation rate until the central gas can cool again. In contrast, in the BH-W model, the wind feedback only accelerates and heats the gas directly surrounding the SMBH, which then flows out of the central region only heating some of the cold, star-forming gas through shocks. Hence, in the BH-W model, the sSFR changes less rapidly and not as strongly as in the BH-WR model.

Figure 2.3 shows the time-evolution of both the total stellar mass (solid lines, left ordinate) and the mass of the central black hole (dashed lines, right ordinate) for all models. Without black-hole accretion and feedback (no-BH model), the stellar mass grows by about 10%. Both black-hole feedback models reduce the overall growth of the stellar mass significantly, compared to the no-BH model (to about 4% in the BH-W, and about 2% in the BH-WR model). While the BH-WR model reduces the stellar mass growth more than the BH-W model, it leads to a slightly stronger growth of the black-hole mass (about 7.5%, compared to about 5% in the BH-W model), i.e. more gas accretion. Overall, both feedback models keep the black-hole and the stellar mass growth small, and the galaxy quiescent for the majority of the time (though the BH-WR feedback is much more successful on that front).

In Fig. 2.4, we compare our simulations to the observed present day population of galaxies in a SFR – stellar mass diagram. For our simulations, we use the median star-formation rate and stellar mass over the whole simulation time (with the 1st and 3rd quartile of the spread in SFR given as error bars). For the observations, we show data from Chang et al. (2015) (see also Renzini & Peng, 2015) as a grey-scale histogram of the abundance of observed galaxies in certain M_* – SFR bins. We use the 1σ upper limit of the observed data without V_{\max} -correction to highlight the bimodal distribution between star-forming and quiescent galaxies. To guide the eye, we also show the a fit for the star-forming sequence by Renzini & Peng (2015) and its 1σ scatter with cyan solid and dashed lines, respectively.

In the no-BH model, our simulated galaxy falls into the lower end of the star-forming sequence at all times during the simulation. In the BH-W model, the galaxy has star-formation rates roughly corresponding to the “green valley” between the blue and the red sequence of star-forming and quiescent galaxies. Finally, in the BH-WR model, the galaxy’s star-formation rate is firmly in the range of the quiescent galaxy population. All in all, this confirms the conclusions made above: in our simulations, black-hole feedback




Figure 2.3: Total stellar mass (solid lines, left y-axis) and black hole mass (dashed lines, right y-axis) of the galaxy over time for the same models. Without black-hole feedback, the stellar mass grows by about 10% over the whole simulation time; the feedback models reduce this growth to 4% (BH-W) and 2% (BH-WR) respectively. In the BH-WR model, the black-hole mass grows more (about 7.5%) than in the BH-W model (about 5%), but in both feedback models the growth of the stellar and the black-hole mass is very small.




Figure 2.4: SFR vs. stellar mass relation: Of the three models of this paper, the median values, as well as the scatter between the 25th and 75th percentile over the simulation time are plotted. For comparison, observed data from Chang et al. (2015)—using the SDSS and WISE surveys—is shown as a grey-scale histogram depicting the number of galaxies in a given M_* – SFR bin, with a fit for the star-forming sequence by Renzini & Peng (2015) (solid) and its 1σ scatter (dashed) marked by the cyan lines. To emphasise the bimodal split of the galaxy population into star-forming and quiescent galaxies, we plot the 1σ upper limits for the observed star-formation rates and neglect V_{\max} corrections.

is necessary to keep the galaxy quiescent, and including both wind and radiative feedback leads to significantly less overall star-formation than including only wind feedback.

Lastly, we take a look at the accretion behaviour of the central SMBH in Fig. 2.5, which shows the Eddington ratio $f_{\text{Edd}} = \dot{M}_{\text{BH}}/\dot{M}_{\text{Edd}}$ distribution of our two models including black-hole feedback compared with observational data from various sources (Ho, 2009; Greene & Ho, 2007; Kauffmann & Heckman, 2009; Heckman et al., 2004), which were compiled by Novak et al. (2011). We plot the cumulative fraction of simulation time during which the black hole accreted above a given f_{Edd} over the Eddington ratio. Both models produce time fractions at high Eddington ratios that are too high compared to the observations, or only in agreement with the highest observational estimates, i.e. the SMBH in the simulations has a rather high duty cycle. There is only little difference between the two models, though the BH-WR model has a slightly higher duty cycle than the onlyWind model, which matches the correspondingly higher SMBH mass growth (compare Fig. 2.3).

2.4 Black-hole-governed ISM evolution in the galactic centre

In the previous section, we have shown that feedback from the central supermassive black hole is both necessary to, and capable of keeping an isolated early-type galaxy quiescent. The more physically complete feedback model, including both wind and radiative feedback, is more efficient at this than the pure wind-feedback model. In this section and the next, we will now examine how the feedback achieves the quiescence through its influence on the galactic gas, and how it creates metal-rich, large-scale outflows that enrich the circumgalactic medium.

We begin with Fig. 2.6, which shows colour maps of the surface density (top row) and the density-weighted average temperature (central row) in a 10 kpc thick, 10x10 kpc wide region around the galactic centre (and therefore around the central SMBH), as well as the temperature-density phase diagrams of all the galactic gas (bottom row). In the maps, the line of sight is chosen to be parallel to the initial angular momentum of the hot halo. The solid vertical and horizontal lines in the phase diagrams mark the star formation thresholds in density and temperature, respectively, i.e. gas in the bottom right corner of the diagrams is star-forming. The individual temperature-density bins are colour-coded for gas particle number density, which is rising from blue to red colours. The dashed line marks the resolution of the simulation limit following the Jeans mass criterion: Above the dashed line, the Jeans mass of the gas is resolved, i.e. it is larger than the kernel mass $M_{\text{ker}} = N_{\text{ngb}}m_{\text{gas}} = 10^7 M_{\odot}$, while it is unresolved below the line, meaning that gas in this region could collapse under self-gravity due to numerical noise. Almost all of the gas particles with unresolved Jeans mass are in the star-forming region of the phase space, and will be turned into star particles if they stay in this region. Hence, while this resolution does not allow us to make predictions about the substructure of the dense, cool gas, it is high enough to investigate the global and larger-scale effects of AGN feedback that interest




Figure 2.5: Time distribution of different accretion rates in units of the Eddington limit (so-called “black-hole duty cycles”) for the different models (see legend) compared to observations (square: Ho (2009), circles: Greene & Ho (2007), upward-pointing triangles: Kauffmann & Heckman (2009), downward-pointing triangles: Heckman et al. (2004), all taken from Novak et al. (2011)). Both simulated models fall above all of the observed values at Eddington ratios of 1% and 10%.

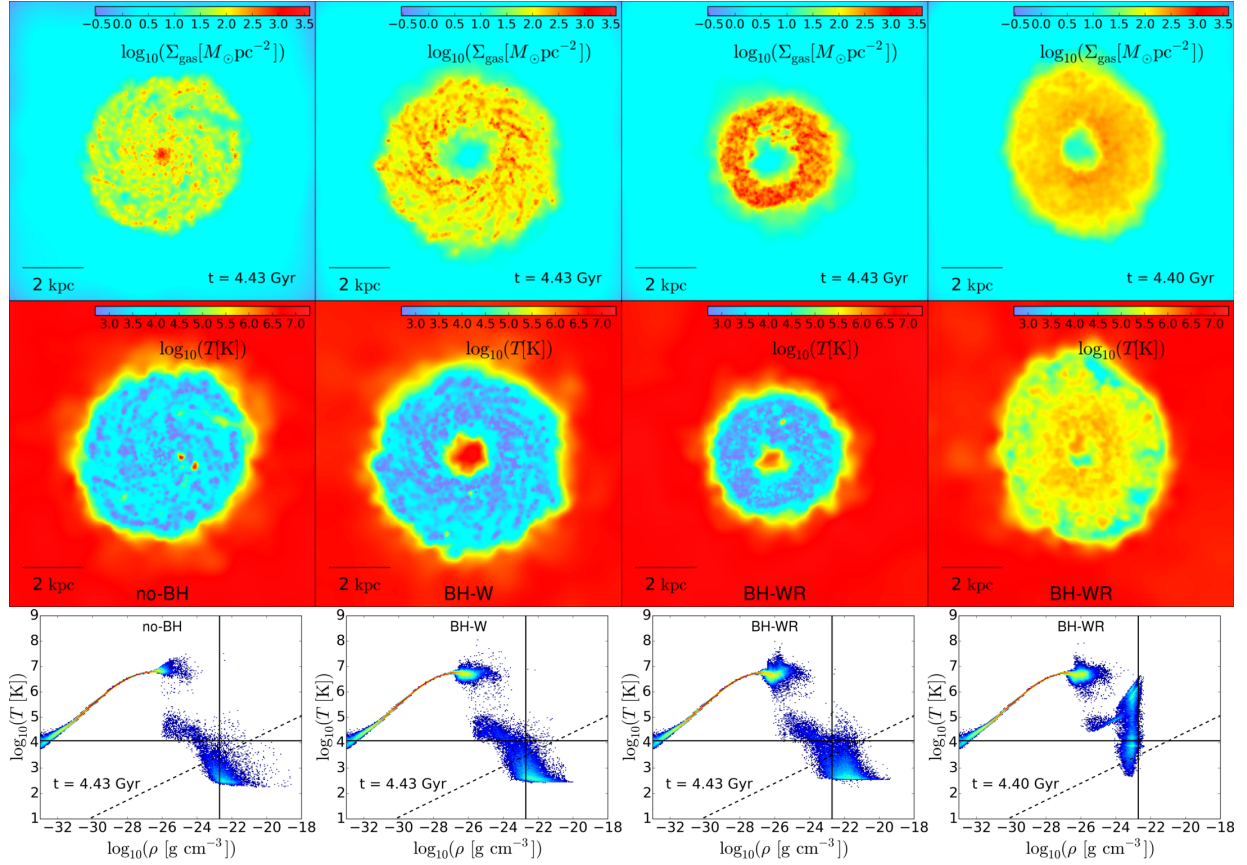


Figure 2.6: Maps of the gas surface density (top row) and the density-weighted average temperature (central row) in a 10 kpc thick, 10x10 kpc wide region around the galactic centre, as well as temperature-density phase diagrams of the gas (bottom row). The columns show, from left to right, the states at the end of the simulation time for the no-BH, BH-W and BH-WR models, as well as the state of the BH-WR model 30 Myr before the end of the simulation. In the phase diagrams, the solid vertical and horizontal lines mark the density and temperature thresholds for star formation, respectively. The dashed line marks the resolution limit of the simulation, i.e. the Jeans mass is resolved for all gas that lies to the top left of the dashed line in the phase space, and the colour shows the number density of gas particles in the temperature-density bins (rising number density from blue to red). In all simulations, a central disc of cold, dense gas is formed. While in the no-BH model, this disc includes an especially dense central core, in the other two models, the very centre is instead evacuated of gas by the wind feedback. Additionally, in the BH-WR model, the whole disc is periodically heated and enlarged by the radiative feedback, before condensing and cooling down again.

us in this paper. The three left columns show the state of the three models (from left to right: no-BH, BH-W, BH-WR) at the end of the simulation time, while the last column to the right shows the state of the BH-WR model about 30 Myr before the end of the simulation.

In the no-BH model, we see a typical cooling-flow: the gas from the hot halo of the galaxy flows to the centre, cools down, and forms a disc of cold, dense gas in the central ~ 6 kpc. An especially dense core of up to a few times 10^{-19}g/cm^3 forms at the very centre and leads to significant star-formation. As the stellar feedback is too weak to affect the gas much, this central core and disc are stable over the whole simulation time, while some of the gas is constantly being turned into stars, and is in turn replenished by further cooling from the hot halo.

In both the BH-W and the BH-WR model, the gas cooling down into the centre forms a dense, ring-like disc with a hole of under-dense gas in the very centre of about 1 kpc size. This central hole is produced by the wind feedback, which both heats the central gas and accelerates it parallel to the angular momentum of the gas disc, creating a fast, biconal outflow, which evacuates the centre and thereby significantly reduces the star-formation rate. In the BH-WR model, the star-formation rate is further reduced by the radiative heating: While, in the BH-W model, the cold, dense disc around the centre is mostly undisturbed and forms stars constantly (though at a lower rate than the core in the no-BH model), in the BH-WR model, the whole disc is repeatedly heated and expanded (and thereby made less dense) by the radiative feedback (compare the two rightmost columns in Fig. 2.6). While the gas disc re-compresses and cools again quickly after being heated, the radiative feedback nevertheless reduces the galactic star-formation rate greatly, leading to the overall very quiescent state of the BH-WR model.

2.5 Large-scale metal enrichment through feedback-driven winds

In Fig. 2.7, we show the density-weighted temperature (top row) and metallicity (bottom row) in a 100×100 kpc wide, 3 kpc thick region around the galactic centre. Unlike in Fig. 2.6, the line of sight in this figure is perpendicular to initial angular momentum of the hot gas, allowing for the depiction of the biconal outflows produced by the black-hole feedback. The left most column shows the initial condition; the other columns show, from left to right, the state of the no-BH, BH-W, and BH-WR models at the end of the simulation time, and the black circle marks $r = 0.01 R_{\text{vir}}$.

In the no-BH model, we again see that a thin, cold disc is formed in the galactic centre. This disc is surrounded above and below the plane by shock-heated gas that is being accreted onto it. The cold disc and its surrounding accretion-shock region are the only part of the galaxy which remain metal-enriched by the end of the simulation: All the metal-rich gas that is initially distributed throughout the galaxy cools down into the centre, where it is further enriched with metals produced by newly formed stars. The lack

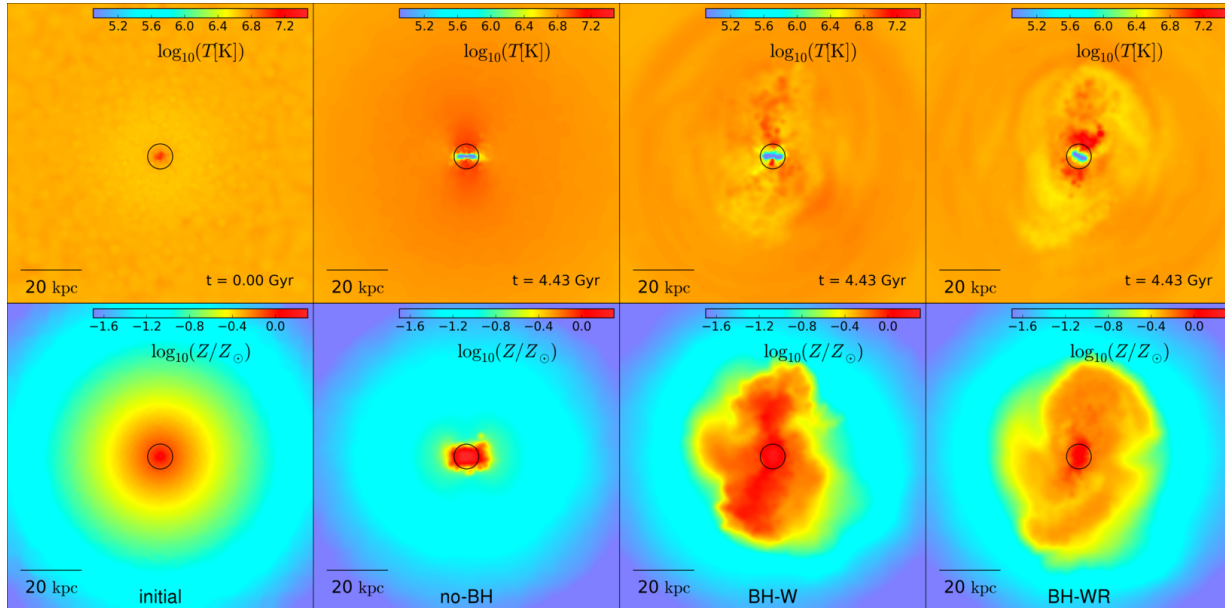


Figure 2.7: Maps of the density-weighted average temperature (top row) and metallicity (bottom row) of the gas in a 3 kpc thick, 100x100 kpc wide region around the galactic centre. The columns show, from left to right, the initial condition and the states at the end of the simulation time for the no-BH, BH-W and BH-WR models. The black circle marks 1% of the virial radius. In the no-BH model, a very metal-rich and cold central disc is formed, surrounded by a hot accretion shock of infalling gas. All gas outside this central region is depleted of metals. In the other two models, the wind feedback drives large-scale, hot, metal-rich outflows, which keep the CGM of the galaxy metal-enriched out to ca. 30 kpc.

of strong feedback (the stellar feedback being too weak to affect the central gas much beside the metal-enrichment) then leads to a highly metal-enriched central region of the galaxy and a total lack of any metals in the outskirts and the CGM.

The situation looks very different in the two models that include black-hole feedback, BH-W and BH-WR: While we still see cold, central discs, in these models, the wind feedback produces biconal, hot, and metal-rich outflows perpendicular to the disc plane. These outflows not only evacuate the galactic centre from gas that would otherwise form stars, they also transport metals produced in the centre up to ca. 30 kpc away from it, into the CGM. Most of the outflowing, metal-rich gas falls back towards the centre within a few hundreds of Megayears, but the constantly renewed, black-hole feedback driven outflows from the centre still keep the CGM (up to ca. 30 kpc) enriched with metals.

Figure 2.8 demonstrates the same point by showing the 3D radial profiles of the gas metallicity at the end of the simulation for all three models, compared to the initial metallicity profile. In the no-BH model, the star-formation and subsequent stellar feedback in the cold, dense disc and core raises the metallicity in the centre up to two to three times the solar value, but the constant inflow of gas from the outer parts of the galaxy, combined with the lack of any large-scale outflows, leads to a total depletion of metals in the CGM gas (with less than 10% solar metallicity outside of a radius of 20 kpc). On the other hand, with black-hole feedback (in the BH-W and BH-WR models), the central metallicity stays lower, close to the initial value, while the outer parts of the galactic gas become significantly enriched with metals out to a radius of about 30 kpc. There is little difference between the two models that include black-hole feedback, as the wind feedback is the relevant part to create large-scale outflows capable of enriching the CGM. Outside of a radius of ca. 30 kpc, the gas becomes depleted in metals even in the feedback-including models, as the AGN-driven outflows lose their momentum and stop progressing before they can reach further out. The two models with black-hole feedback also agree much better with the observations of O’Sullivan et al. (2007), at least out to ~ 30 kpc, but considering the observational uncertainties, as well as the slope of the initial metallicity profile in the simulations (which is set to follow the stellar profile, and therefore not necessarily realistic), one should be cautious before reading too much into this agreement.

In Fig. 2.9 we show the radial profiles of the oxygen-to-iron ratio in solar units (O/Fe) at the simulation’s end for all three models, compared with the initial condition (which is set to the solar value, i.e. 1). As oxygen is mainly produced in type-II supernovae of young, massive stars, while iron mostly originates from old white dwarfs undergoing type-Ia supernovae, and the initial stellar population in our simulations is made up exclusively of old stars, the radial O/Fe profile shows the combined effect of star-formation and feedback-driven metal distribution.

In the no-BH model, significant star-formation in the centre of the galaxy leads to an enhancement of O/Fe in the central regions up nearly twice the initial value. None of the enriched gas from the star-forming centre is transported out beyond a few kiloparsecs, so no new, star-formation-dependent oxygen reaches the outskirts of the galaxy, while there are still old stars producing iron through SNIa in these regions. Thereby, O/Fe decreases significantly beyond a radius of ca. 6 kpc, though not as drastically as the overall metallicity




Figure 2.8: Radial gas metallicity profiles at the end of the simulation, for the different models, as well as in the initial condition (see legend). Also shown are observed values from O’Sullivan et al. (2007) for NGC 7796 (circles) and NGC 57 (squares). In the no-BH model, only the centre of the galaxy stays metal enriched, while the volume further out than ca. 6 kpc becomes strongly depleted in metals. In contrast, both in the BH-W and the BH-WR models, the gas stays enriched at more than 10% solar metallicity out to about 30 kpc.




Figure 2.9: Radial gas oxygen-to-iron ratio (O/Fe) profiles in solar units at the end of the simulation for the different models, as well as in the initial condition (see legend). The no-BH model shows enhanced O/Fe in the star-forming galactic centre, and decreased O/Fe further out. The BH-W model still has enough star-formation for a slightly enhanced O/Fe in the centre, and also shows higher O/Fe in the outskirts than the no-BH model due to its feedback-driven outflows. The BH-WR model leads to decreased O/Fe at all radii due to the strong lack of star formation, but still shows higher O/Fe at large radii than the no-BH model due to large-scale outflows.

(compare Fig. 2.8).

In the BH-W model, the central star formation (and therefore oxygen production) is significantly lower than in the no-BH model, but still high enough to raise O/Fe slightly above the initial value in the galactic centre. The wind-feedback driven outflow also transports some of the centrally-produced oxygen out into the CGM (up to about 30 kpc, the reach of the outflows), leading to O/Fe in these regions about twice as high as in the no-BH model, but still lower than in the initial condition.

In the BH-WR model, the star-formation rate is so low that the production of iron by old stars outweighs that of oxygen by newly formed ones everywhere in the galaxy: O/Fe is lower than initially at all radii. Still, what little oxygen is produced in the slightly star-forming centre is distributed by the feedback-driven outflows into the CGM, leading to O/Fe in the region between about 6 kpc and 30 kpc lower than in the BH-W model, but still slightly higher than in the no-BH model.

Observed oxygen-to-iron ratios in massive elliptical galaxies mostly range from about 60% to ca. 100% of the solar value (e.g. Pinto et al., 2014; Grange et al., 2011, and references therein), and have roughly flat radial profiles (Grange et al., 2011). The results of our BH-W and BH-WR runs are consistent with these observations, while those of the no-BH run fit less well to the data, especially in the centre, which is very oxygen-rich in the simulation.

Using Fig. 2.10, we examine the gas flow rates in the different models directly. It shows, for all three models, the time evolutions of the gas mass inflow rate (IFR) through a shell at 20 kpc distance from the centre (top panel), and the ratio of outflow to inflow rate (OFR/IFR, bottom panel). The flow rates are estimated taking the current positions and velocities of all gas particles on one side of the shell, and then determining which of them will be on the other side of the shell 3 Myr (about one snapshot-interval) later, assuming constant velocity during that time. The total mass of all these gas particles is then the inflow or outflow rate, respectively. Same as for the specific star-formation rate, we use a running mean over about 30 Myr for the flow rates to reduce the noise of the plots.

In the models that include black-hole feedback, BH-W and BH-WR, the flow rates are very similar. Both inflow and outflow rate oscillate around about $8M_{\odot}/\text{yr}$, leading to an OFR/IFR ratio that fluctuates around 1. Periods of net-outflow alternate with those of net-inflow: gas is blown out by the black-hole-wind feedback, and falls back towards centre, leading to a quasi steady state (at least within the simulation time). While the black-hole wind produces outflows, these are compensated for by gas flowing back into the galaxy, so the galactic centre is not completely depleted. Instead, a so-called “galactic fountain” is created.

In the no-BH model, there is overall much less movement of the gas (at this distance from the centre) than in the other models: The inflow rate is lower by a factor of about 3-4, while the outflow rate is lower by more than an order of magnitude, leading to an OFR/IFR ratio that is constantly below 1, usually oscillating around ~ 0.2 . Without the black-hole feedback that drives gas out from the central regions, which then slows down and starts flowing into the centre again, the only movement of the gas is an inflow to the centre where it is converted into stars. The CGM gas is also not replenished by




Figure 2.10: Inflow (IFR, top panel), and the ratio of outflow to inflow rate (OFR/IFR, bottom panel) through a shell at 20 kpc distance from the centre over time for the different models (see legend). For all flow rates, the running mean with a ca. 30 Myr (10 snapshot) step size is plotted. The horizontal black line in the bottom panel shows a 1:1 ratio of OFR to IFR, i.e. zero net-flow of gas. The BH-W and BH-WR models show very similar flow rates, with OFR/IFR oscillating around 1, while the no-BH model produces generally much lower flow-rates and constant OFR/IFR ratios far below unity.

feedback-driven outflows, hence the total inflow rate is also significantly lower than in the feedback-including models. As the hot gas halo and its metallicity are slowly depleted, the inflow rate shrinks further, though the effect is minor (from about $3M_{\odot}/\text{yr}$ to about $2M_{\odot}/\text{yr}$ in the last 4 Gyr).

In the BH-W and BH-WR models, the total mass flow of the gas is in equilibrium. The same is not true of the total flow of metals in the galaxy, as the AGN winds drive out gas that is on average more metal-rich than its inflowing counterpart. To demonstrate this, we plot the distribution of ratios of metal mass outflow (mOFR) to inflow (mIFR) rate, i.e. the total mass of all elements except hydrogen and helium in the corresponding flows, against the distribution of the same ratio for the total gas mass (i.e. OFR/IFR). This is shown in Fig. 2.11 for flow rates through a shell at 20 kpc distance from the centre (top panel), and at 50 kpc distance (bottom panel). The distribution is visualized with contours enclosing (from the inside out) 33%, 67%, and 95% of the total number of values in all snapshots (i.e. values at different, evenly spaced simulation times).

In the no-BH model, the mOFR/mIFR ratio is about equal to the OFR/IFR at all times, both at 20 kpc and at 50 kpc. This confirms that there is no outflow of centrally produced metals beyond 20 kpc in this model. In contrast, in both the BH-W and the BH-WR model, there is a significantly higher ratio of outflow to inflow rate for metals than for the total gas mass at a distance of 20 kpc from the galactic centre, demonstrating that, although the overall gas flow is balanced, the black-hole feedback still drives out a significant amount of metals into the CGM beyond 20 kpc. This is different further away from the centre: At a distance of 50 kpc, the two models that include black-hole feedback have mOFR/mIFR ratios close to the total OFR/IFR at most times. The AGN-driven outflows have lost all of their outward momentum before reaching this far out from the galactic centre, and are therefore not able to enrich this region with metals.

2.6 X-ray properties of the gas

Finally, we take a look at the effects of the black-hole feedback on the X-ray luminosity of the galactic hot gas content. The X-ray emission of the hot gas, produced by bremsstrahlung and metal-line cooling, is the main observable by which galactic hot haloes are detected. It is also dependent on the temperature, on the metallicity, and (heavily) on the density of the gas, all of which are influenced by feedback from the central SMBH, potentially making it a good tool to investigate the effect of the feedback.

To estimate the X-ray luminosity of the gas in our simulations, we calculate $L_{X,0}(T, Z = 0.4Z_{\odot})$ (the normalized X-ray luminosity at 0.4 solar metallicity and temperature T), and $(dL_X/dZ)(T)$ (the first-order metallicity dependence of the X-ray luminosity at temperature T) for a range of gas temperatures between 0.1 and 60 keV with XSPEC¹ (Arnaud, 1996). Taking the tabulated temperature T_i that is closest to the particles true temperature, we then use these pre-calculated tables to compute the X-ray luminosities of all gas

¹see <http://heasarc.gsfc.nasa.gov/docs/xanadu/xspec/> for more information




Figure 2.11: Ratio of outflow over inflow rates of gaseous metals ($m\text{OFR}/m\text{IFR}$) vs. the same for all gas, metal or not (OFR/IFR), at 20 kpc (top) and 50 kpc (bottom) from the centre; the contours enclose 33%, 67%, and 95% (from the inside out) of the distribution of values for all snapshots of all different models (see legend). The solid black line marks a one-to-one relation between the two quantities, and the dashed black lines mark ratios of 1 (equal in- and outflow).

particles ($L_{X,i}$) via:

$$L_{X,i} = \mathcal{N}(\rho_i, n_{e,i})[L_{X,0}(T_i, Z = 0.4Z_\odot) + (dL_X/dZ)(T_i) \cdot (Z_i - 0.4Z_\odot)], \quad (2.13)$$

where $\mathcal{N}(\rho_i, n_{e,i})$ is a normalization factor dependent on the mass density and electron density of the gas particle. With the luminosity of each particle known, we then compute the total X-ray luminosity of the galactic hot gas by summing over all particles, excluding those whose densities surpass the star-formation limit $\rho_i > 1.94 \times 10^{-23} \text{ g cm}^{-3}$. This very dense gas is assumed to reside in star-forming regions, where it is very likely to be obscured because of the high dust content.

Using this estimate, we then compare the median total X-ray luminosity and its scatter in our three models with observations in two different scaling relations (Fig. 2.12). The upper panel shows the $L_X - M_{\text{tot}}(R < 5R_e)$ relation with observations by Kim & Fabbiano (2013), while the lower panel shows the $L_X - M_*$ relation of Anderson et al. (2015). The black-hole feedback does not influence the median X-ray luminosity of the hot gas much, compared to the already rather large scatter between individual observed galaxies. The strongest effect can be seen in the BH-WR model. Here, the radiative heating of the central gas both raises the median luminosity by a factor of a few (compared to the initial condition), and also causes strong fluctuations in L_X over time (visible through the large scatter for the model, and in Fig. 2.13), as the heated gas quickly cools down again (compare Fig. 2.6). Nevertheless, the median values of all three models are broadly consistent with the observations by Kim & Fabbiano (2013) in the $L_X - M_{\text{tot}}(R < 5R_e)$ relation.

On the other hand, they all fall significantly too high in the $L_X - M_*$ relation, compared to the stacked observations of Anderson et al. (2015), as does the initial condition. The reason for this contradiction might be that the stellar mass we assume in our initial condition is too low compared to the dark-matter (or total) mass. We scaled our stellar mass to the dark-matter mass using the abundance matching relation of Moster et al. (2013); had we used the relation of Kravtsov et al. (2018) instead, our stellar mass would be larger by a factor of ~ 2.5 , enough to explain the difference between observed and simulated values, at least for the initial condition. In the $L_X - M_{\text{tot}}(R < 5R_e)$ relation, the no-BH model shows a significantly lower $M_{\text{tot}}(R < 5R_e)$ than the other models. This is caused by a shrinking of the effective radius due to the comparatively strong central star-formation in this model.

In Fig. 2.13, we show the time evolution of L_X in the 0.5-2 keV band for all of our models. The median luminosities only vary by a factor of a few between the different models (which is comparable to the scatter in L_X between observed galaxies), but in the BH-WR model, the X-ray luminosity has a much larger scatter than in the other two cases. This is due to the radiative AGN feedback rapidly changing the density and temperature of the gas in the galactic centre, which is the densest part of the ISM and therefore contributes the most to the total L_X .

In our simulations, the median X-ray luminosity of the gas fails as an observational diagnostic for the influence of the black-hole feedback on the gas, as the long-time feedback-effects on L_X are small compared to the effects of the total amount of gas available in

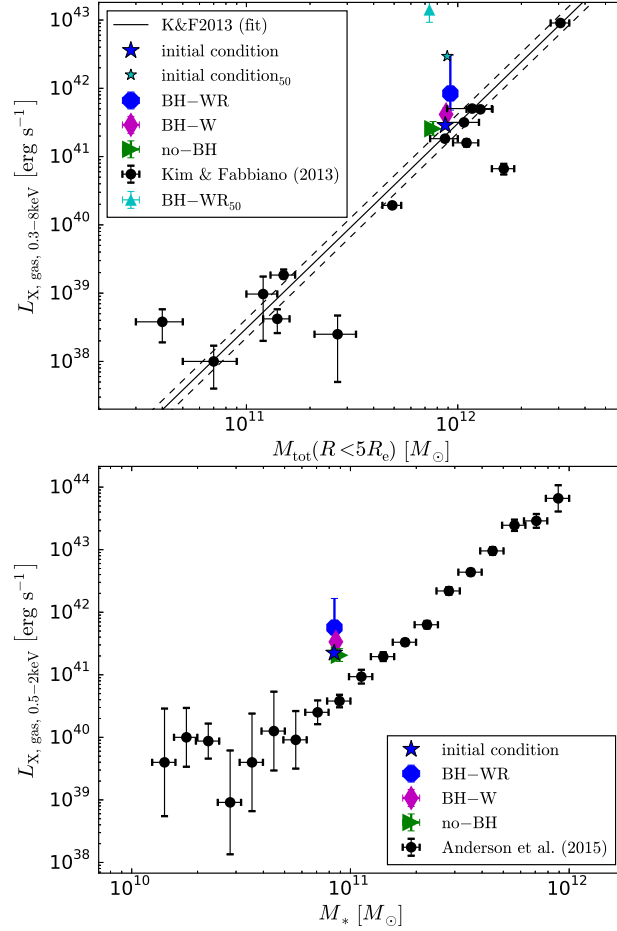


Figure 2.12: Top: The X-ray luminosity of the hot gas in the 0.3-8 keV band over the total mass within five times the effective radius for the initial condition used in this work, the median value over the whole simulation time for all three models, and observations by Kim & Fabbiano (2013). Additionally, we show the initial and median values of the simulation with 50% of the cosmological baryon fraction (BH-WR₅₀). Also shown are the best fit (solid line) and its logarithmic scatter (dashed lines) for observational data. The error bars on the simulation symbols represent the scatter between the 25th and 75th percentile over the whole simulation time. See the text for how L_X was calculated. Bottom: The X-ray luminosity of the hot gas in the 0.5-2 keV band (different band than in the top panel!) over the total stellar mass, with observations by Anderson et al. (2015). There are only comparatively small differences between the median X-ray luminosities of the three different models with 20% of $f_{b, \text{cosm}}$, compared to the observational scatter. The simulated values fit well to the observations in the $L_X - M_{\text{tot}}(R < 5R_e)$ relation (top), while they are all too high in the $L_X - M_*$ relation (bottom). In the BH-WR₅₀ run, the X-ray luminosity is an order of magnitude too high for the observed $L_X - M_{\text{tot}}(R < 5R_e)$ relation already in the initial condition, which increases to about 2 orders of magnitude for the median.




Figure 2.13: X-ray luminosity in the 0.5-2 keV band over time for the different models (see legend). In the BH-WR model, the heating of the central dense gas by the radiative AGN feedback causes a strong oscillation in the X-ray luminosity by more than order of magnitude, while L_X changes much less over time in the BH-W and no-BH models.




Figure 2.14: Map of the X-ray surface brightness (0.5-2 keV band) of the gas in a 3 kpc thick, 100x100 kpc region around the galactic centre. The columns show, from left to right, the initial condition and the states at the end of the simulation time for the no-BH, BH-W and BH-WR models. The black circle marks 1% of the virial radius. While the no-BH model produces a very bright core corresponding to the accretion region around the central star-forming disc, in the BH-W and BH-WR models, the X-ray bright gas is more spread out throughout the ISM and CGM, though the bulk of the luminosity still comes from the central core region.

the galaxy, and the total galactic mass that determines its overall density profile. Ciotti et al. (2017) find a similar lack of long-term influence of the AGN feedback on the X-ray luminosity in their two-dimensional simulations of ETGs. AGN feedback could have a larger impact on the total X-ray luminosity by permanently reducing the hot gas density, e.g. by expelling large quantities of gas to beyond the virial radius, into the intergalactic medium. This is a possible mechanism to quench forming ETGs in the first place, but would need to happen at higher redshifts (compare e.g. McCarthy et al., 2011). For example, Choi et al. (2015), who investigate the influence of different AGN feedback mechanisms on the formation of ETGs in cosmological simulations, find that the X-ray luminosity of their ETGs is $\sim 2 - 3$ orders of magnitude lower in simulations with the same mechanical-radiative AGN feedback model that was used in our BH-WR run than in simulations without AGN feedback. This is the case because the AGN wind feedback drives out large amounts of gas from the galaxies at high redshifts, strongly reducing the density and thereby the X-ray luminosity of the remaining hot halo gas. Therefore, the total X-ray luminosity of a galaxy can be an observational constraint on the effect of AGN feedback on the galactic gas at high redshifts.

However, under the conditions investigated in our simulations, i.e. massive, local elliptical galaxies at low redshift, the feedback seems only capable to drive outflows up to ~ 30 kpc, beyond which they start cooling down into the centre again. This, together with central heating, is enough to keep the galaxy quiescent, but it is not enough to permanently expel significant amounts of gas from the galaxy, and lower the overall X-ray luminosity, which is mostly determined by the initial condition, i.e. the gas density distribution and gravitational potential of the galaxy about 4.5 Gyr before now.

While the total X-ray luminosity fails to distinguish the models with and without

feedback from one another, we can still see the effects of the feedback in the X-ray surface brightness distribution, which we show in Fig. 2.14 for the initial condition and the final states of our three runs. In all three runs, the formation of the central dense disc leads to a very bright core region of the X-ray emitting gas. However, in the models including wind feedback (BH-W and BH-WR), the gas in the surrounding CGM (out to ~ 30 kpc) is kept about an order of magnitude brighter, as well as much more structured, than in the no-BH run.

The hot outflows caused by the wind feedback in both the BH-W and BH-WR models result in several kpc large, outward moving cavities in the X-ray surface brightness. Such X-ray cavities have been observed and associated with AGN activity in many galaxies, and their sizes are roughly comparable to those of the cavities found in our simulations (e.g. Churazov et al., 2000; Bîrzan et al., 2004; Forman et al., 2007). They are somewhat visible in the relevant panels of Fig. 2.14, but we also show one much clearer example of a forming X-ray cavity in the BH-WR model in Fig. 2.15.

2.7 Limitations of the AGN feedback model

The AGN feedback model we use in this paper contains both a radiative and kinetic wind feedback mode, which together are efficient at limiting the SMBH growth, inhibiting star formation, and driving metal-enriched outflows up to about 30 kpc from the centre. However, the model does not include every mode of AGN feedback that can be observationally motivated and might have an impact on the galactic gas content. In particular, we neglect the potential effects of collimated radio jets on the evolution of the ISM and CGM.

It is difficult to directly model the highly relativistic synchrotron jets that are observed in many massive galaxies (e.g. M87) in galaxy-scale simulations as the ultra-high velocities in the jet lead to very small time steps. Nonetheless, several groups have implemented jet-like AGN feedback in the form of non-relativistic (jet velocity $\sim 10,000$ km/s), strongly collimated winds mimicking outflows launched from the accretion disc around the SMBH into galaxy evolution simulations (e.g. Omma et al., 2004; Dubois et al., 2010), especially in the context of the radio mode in combined quasar/radio mode feedback models (e.g. Dubois et al., 2012, 2013).

The simulations presented in this paper do not include a dedicated jet feedback mechanism, even though the SMBH spends most of its time at low accretion rates ($\sim 90\%$ of the total time at $f_{\text{Edd}} < 1\%$, compare Fig. 2.5), i.e. in the radio-mode regime where such a feedback mechanism might be relevant. The wind feedback included in our simulations is similar to the non-relativistic jet feedback models described and used in the papers cited above, so its effect would likely be qualitatively the same, i.e. driving out gas perpendicular to the central gas disc plane.

However, a jet feedback mode would be even more strongly collimated (with all gas in the jet being accelerated along the same axis) than our wind feedback (which would only be perfectly collimated if all gas flowing onto the SMBH had a uniform direction of angular momentum). This might enable it to push the central gas farther away from the centre




Figure 2.15: n

Map of the X-ray surface brightness (0.5-2 keV band) of the gas in a 3 kpc thick, 30x30 kpc region around the galactic centre for the BH-WR model. The black circle marks 1% of the virial radius. One can see the formation of an X-ray cavity in an AGN-driven outflow.

than the wind feedback, enriching the CGM beyond ~ 30 kpc with metals. On the other hand, an additional jet feedback mode is unlikely to change the results of our simulations as far as the central region of the galaxy is concerned. The wind feedback is already able to evacuate the very centre, inhibiting the star formation there, and a jet would be just as incapable of affecting the surrounding dense gas disc (which is kept mostly quiescent by the radiative feedback) as the wind feedback due to its collimation.

2.8 Summary & conclusions

We investigated the influence of different forms of AGN feedback on the evolution of isolated early-type galaxies at low redshift. For this purpose, we used the SPH code SPHGal, an improved version of GADGET-3, to run simulations of an idealized, isolated, massive, local elliptical galaxy with three different models: one without any accretion onto or feedback from the central supermassive black hole (model no-BH), one where the black-hole feedback is implemented as a kinetic wind (model BH-W), and finally one where radiative X-ray heating is added to the wind feedback from the black hole (model BH-WR). We compared the outcomes of these three simulations between each other, as well as to observations, with respect to their star-formation and black-hole growth history, the central density and temperature evolution of the ISM, the large-scale flow patterns and metal distribution in the gas, and the X-ray properties of the hot galactic gas.

Our results can be summarized as follows:

- Without black-hole feedback, a classical cooling-flow problem occurs in our simulation: gas constantly flows into the central region and cools onto a cold, dense disc with a very dense core in the centre, resulting in permanent significant star formation. Once gas had time to cool down from the initial condition, the galaxy is not quiescent anymore for the rest of the simulation time.

Including black-hole feedback reduces the star formation rate significantly: Pure wind feedback (the BH-W model) only reduces the SFR moderately, resulting in a galaxy roughly in the “green valley”, changing between slightly quiescent and slightly active states. This reduction of the SFR is achieved by bipolar outflow from the central, cold, dense disc, that is caused by the wind feedback and evacuates the very centre of the galaxy, which would otherwise (as in the no-BH model) be the region with the highest SFR.

Adding radiative heating to the feedback as well (the BH-WR model) results in a very quiescent, “red and dead” galaxy that shows only short bursts of significant star formation between longer periods of very little star formation. Here, the dense, potentially star-forming gas around the wind-feedback-evacuated galactic centre is periodically heated and expanded by the radiative feedback, interrupting ongoing star formation, before it cools down again. This feedback-driven alternation of heating & expansion and cooling & compression leads to a strongly oscillating, overall much lower SFR in the galaxy.

- The growth of the black hole is quite limited in both feedback models, its mass increasing by only about 5% and 7.5% over the whole simulation time of ~ 4.35 Gyr in the BH-W and BH-WR models, respectively. The duty cycle of the black hole is tendentially too high in both models.
- Without black-hole feedback, the outer parts of the ISM and CGM are deprived of metals over time, as all metal-rich gas slowly flows into the galactic centre, leading to very low metallicity in the outskirts and very high, supersolar metallicity in the central region. In contrast, black-hole-wind feedback leads to large-scale outflows, flattening the metallicity profile of the galactic gas out to about 30 kpc, while reducing it in the centre. Additional radiative feedback does not affect this outcome.

The limited range (~ 30 kpc distance from the centre) in which the AGN feedback has a significant effect on the CGM in our simulations is strongly influenced by the chosen initial condition, specifically the deep potential well of this massive, low-redshift ETG. In cosmological simulations (e.g. Choi et al., 2015), where the conditions are different (especially at higher redshifts), the AGN feedback can drive winds to much higher distances from the galaxy, or even eject gas from the galactic potential well completely. Furthermore, a collimated jet mode of the AGN feedback, which is not included in our model, might also be able to move gas farther away from the centre.

- The median total X-ray luminosity of the gas is only slightly affected by the different black-hole feedback models, compared to the observational scatter. The radiative feedback has the largest effect, leading to strong oscillation of the X-ray luminosity as the central dense gas is constantly heated and cooling. Overall though, all models fit reasonably well to the observed $L_X - M_{\text{tot}}(R < 5R_e)$ relation, making the total gaseous X-ray luminosity a bad diagnostic to investigate black-hole feedback in low-redshift, massive ETGs. At higher redshifts, where the AGN feedback can more efficiently reduce the overall hot gas density of the galaxy, its impact on the total X-ray luminosity can be much more significant (compare the cosmological simulations of Choi et al., 2015), making L_X a more useful diagnostic at these earlier times. Both the BH-W, and the BH-WR feedback models result in the formation of X-ray cavities in the feedback-induced hot outflows from the galactic centre, confirming the observational link between these cavities and AGN activity.

We conclude that feedback from the central supermassive black hole is necessary to keep massive, isolated elliptical galaxies quiescent during the late stages of their evolution, as the feedback from their old stellar population (the only plausible alternative mechanism for maintaining the quiescence in isolation) is too weak to have any significant effect on the thermal state of the ISM at these galactic mass ranges. How the AGN feedback is implemented has a major influence on its effect on the galactic gas: While radiative feedback is very efficient in maintaining a low star formation rate through heating the central ISM, it is incapable of driving significant outflows by itself and does not affect the growth of the SMBH much, while, on the other hand, the kinetic momentum feedback

resulting from broad-line region winds drives large-scale outflows that can enrich the nearby CGM of the galaxy, but is less effective at keeping the galaxy quiescent. Both forms of feedback need to act together to maintain quiescence, prevent excessive SMBH growth, and distribute metals through the galactic gas via an AGN-driven galactic fountain, and thereby produce galactic properties in accordance with the observational constraints.

It is worth noting that, while the AGN feedback in our simulations is perfectly capable of keeping the ETG quiescent when we assume 20% of the cosmological baryon fraction for the initial mass of the hot halo, if we instead take an initial condition with 50% of the cosmological baryon fraction, corresponding to an about three times more massive gaseous halo, the central gas density becomes too high for the AGN feedback to have much of an effect, resulting in a constantly star-forming galaxy even when the full feedback model is enabled. This implies (for our assumed gas density profile) that local quiescent ETGs must have lost much of their gas at higher redshifts, likely during the time when they were quenched in the first place, possibly because it was ejected by still more effective (due to shallower potential wells and much higher quasar activity) AGN-driven winds. Alternatively, the hot gas could have been not fully ejected, but instead distributed much more evenly throughout the CGM of the galaxy, leading to a much shallower, and hence less X-ray bright and less efficiently cooling gas density profile.

While our isolated initial conditions allow us to study the effects of AGN feedback in a controlled way, free from degeneracies with effects from intergalactic interactions, this of course also means that these interactions—and more generally, the effects of the galactic environment on the ETG’s evolution and the AGN feedback’s efficiency—are not included. These isolated simulations also do not provide any statistics: we look at only one possible ETG out of a range of possible ones that are consistent with observational constraints on their masses, etc. To address these issues, cosmological (zoom-in) simulations are necessary. For example, Choi et al. (2015) did such a study for massive ETGs, and their results in regard to the necessity of AGN feedback are consistent with ours.

Another weakness of our approach in this work is the limited resolution: While it is good enough to investigate the effects of the AGN feedback on the galactic gas on large, galaxy-wide scales, it limits our ability to resolve the central dense gas. This makes it impossible to attempt more detailed comparisons of AGN-driven outflows of different phases in the simulation with observed ones, as, for example, warm-ionized outflows of $\sim 10^5 M_\odot$ (e.g. those observed by Cheung et al., 2016) can simply not be well resolved with gas particles of $10^5 M_\odot$. To investigate the effect of AGN feedback on the multi-phase ISM in the centre of ETGs in detail, it will therefore be necessary (at least with current computing technology) to focus simulations on smaller scales than a whole galaxy.

Chapter 3

Circumnuclear gas discs in the centres of early-type galaxies

We present three-dimensional hydrodynamical simulations showing the evolution of circumnuclear molecular gas disc in the central 400 pc of a massive early-type galaxy containing a supermassive black hole in its centre for up to 150 Myr. The initial properties of the disc are based on the observed circumnuclear disc of NGC 4429 (Davis et al., 2018). Our simulations include models for star-formation, non-equilibrium gas cooling, a chemical network for the gas (including interstellar UV radiation and cosmic ray ionisation), stellar feedback with several different implementations of supernovae, SMBH accretion, and AGN feedback. We find that the simulated CND has a significantly higher ratio of SFR surface density to gas surface density than NGC 4429 and similar observed CNDs in ETGs. This is caused by the fragmentation of the disc into very dense ($\rho > 10^{-20} \text{ g cm}^{-3}$), highly star forming gas clumps, which is not prevented by the additional stability against fragmentation provided by the deep gravitational potential of the massive ETG. We show that—within the tested parameter range—neither UV radiation nor ionisation by cosmic rays is able to prevent the formation of dense gas clumps, and the subsequent high star formation rates. We introduce a new model for supernova feedback in which the SN energy is distributed directly in the form of radial momentum instead of thermal energy. This new supernova model leads to a better agreement of the morphology of the simulated CND with the observed one, but does not prevent very dense gas clumps and high SFRs. We finally include accretion onto, and feedback from the SMBH, which leads to the evacuation of the central ~ 100 pc, but does not affect the evolution of the CND outside of this radius. A higher SFR leads to a higher black hole accretion rate, as more gas is driven into the centre by supernovae. Physical mechanisms not included in the simulations, e.g. magnetic fields, are likely responsible for stabilising the CND against collapse to high densities and thereby for suppressing their star formation.

3.1 Introduction

A significant fraction of local ETGs (at least 40% according to the ATLAS^{3D} survey, see Young et al., 2014) contain cold, dense, molecular and atomic gas within the central region around their SMBH, often in the shape of regularly rotating discs (e.g. Combes et al., 2007; Young et al., 2011; Serra et al., 2012, see also section 1.3). Most of these detected circumnuclear gas discs have masses between $10^7 M_\odot$ and $10^9 M_\odot$ (Young et al., 2011), and radii of a several hundred parsecs (e.g. Davis et al., 2014; Boizelle et al., 2017). Davis et al. (2014) find the star formation in these systems to be suppressed when compared to the central regions of blue spiral galaxies, i.e. while they have comparable star formation rate surface densities, the CNDS in ETGs have much higher gas surface densities than spiral centres, placing them on average a factor of ~ 2.5 below the Kennicutt-Schmidt (KS) relation (Kennicutt, 1998) for normal spiral galaxies.

They propose that this star formation suppression could be due to the dynamical stability of these CNDS, which may be caused by the deep potential wells and rising rotation curves exhibited by the centres of ETGs (in contrast to spirals). This explanation is closely related to the concept of “morphological quenching” introduced by Martig et al. (2009). The basic argument is this: A rotating gas disc is stable against axis-symmetric instabilities (i.e. against breaking into rings and then clumps) if $Q > 1$ where Q is the Toomre parameter (Toomre, 1964) defined as:

$$Q = \frac{\kappa\sigma}{\pi G \Sigma_{\text{gas}}}. \quad (3.1)$$

Here, Σ_{gas} is the surface density of the gas disc, σ its velocity dispersion, and κ its epicyclic frequency given by:

$$\kappa = \sqrt{4\Omega^2 + R \frac{d\Omega^2}{dR}}, \quad (3.2)$$

where Ω is the angular velocity of the gas disc at a given cylindrical radius R from its centre. It follows from this that a disc of a given surface density Σ_{gas} is more stable if it has more turbulent support (i.e. higher σ), a faster rotation (i.e. higher Ω), and/or more shear (proportional to $d\Omega/dR$). The disc could also be supported thermally, in which case the sound speed of the gas c_s would be the relevant quantity instead of σ . In contrast to spiral galaxies, ETGs are dominated by a massive stellar spheroid, which results both in a deeper potential well (causing a higher Ω) as well as a more steeply rising rotation curve (and therefore a higher $d\Omega/dR$). Hence, if the morphology of a galaxy changes from a spiral to a bulge-dominated ETG in which the stellar mass is more centrally concentrated, the star formation in an otherwise similar gas disc can be suppressed, and the galaxy undergoes “morphological quenching”.

In Martig et al. (2009), the authors test their proposed mechanism for star formation suppression with cosmological zoom-in simulations of a Milky-Way-like galaxy that undergoes a major merger late in its evolution, causing the formation of a massive central bulge. They find that, indeed, the formation of the bulge leads to a more stable, less starforming gas disc than in a comparison run, which excluded the major merger (and therefore the

bulge growth). However, these simulations are of poor resolution (gas particle mass of $2.1 \times 10^4 M_\odot$), and do not include any models for either stellar or AGN feedback.

Within the context of the ATLAS^{3D} project, Martig et al. (2013) then examined morphological quenching in high-resolution AMR simulations (5.3 pc minimum cell size with $\lesssim 3.7 \times 10^3 M_\odot$ gas mass per cell) of isolated elliptical galaxies compared to equivalent spirals. As in their previous work (Martig et al., 2009), the authors find that the gas disc is much more stable in the elliptical galaxy than in the spiral one, preventing fragmentation into small dense clumps, and thereby reducing the star formation surface density Σ_{SFR} for equal gas surface densities. However, this is only the case in their simulations with a low gas mass and low initial gas surface density ($\Sigma_{\text{gas}} \lesssim 200 M_\odot \text{pc}^{-2}$). In their pair of simulations with a higher initial surface density profile (falling from $\Sigma_{\text{gas}} \sim 700 M_\odot \text{pc}^{-2}$ in the centre to $\Sigma_{\text{gas}} \sim 100 M_\odot \text{pc}^{-2}$ at the outer rim), the gas disc fractures into dense clumps in both the spiral and the elliptical galaxy, resulting in much more similar star formation rates. Finally, they compare the observed $\Sigma_{\text{SFR}} - \Sigma_{\text{gas}}$ relation of the lenticular galaxy NGC 524 (Crocker et al., 2011) to the results of a simulation in which the initial conditions match the observed properties of the gas disc in NGC 524 (which has $\Sigma_{\text{gas}} \lesssim 100 M_\odot \text{pc}^{-2}$). Their simulation matches the observations reasonably well, though the slope of the simulated $\Sigma_{\text{SFR}} - \Sigma_{\text{gas}}$ relation is significantly steeper than the observed one.

Recently, a few CNDs in the centres ETGs have been measured and dynamically modelled in detail, using high resolution CO observations from the Atacama Large Millimeter/submillimeter Array (ALMA). Boizelle et al. (2017) measured the masses, sizes, molecular gas surface densities, and kinematics of five CO-bright CNDs in ETGs. They are generally around $10^8 M_\odot$ in mass, extend over about one kiloparsec, have surface densities of a few hundred $M_\odot \text{pc}^{-2}$, rotation velocities of several hundred km s^{-2} , and low velocity dispersions of the order of 10 km s^{-2} . The authors use the kinematic data to analyse the stability of the observed CNDs by estimating their Toomre Q parameters within several radial bins. All five of the observed galaxies appear to be stable ($Q > 1$) with the minimum Q ranging from about 3 to about 15 for the various galaxies. Boizelle et al. (2017) argue that—as the rotation curves are fairly flat in all of their galaxies—this stability is mostly driven by the high angular velocities Ω in the CNDs. Furthermore, the discs become more stable both at their outer borders (due to Σ_{gas} falling more quickly than κ) and close to the centre where the central SMBH dominates the gravitational field. Despite being stable according to the Toomre criterion, their CNDs show low levels of star formation, indicating that some gas does collapse into dense, starforming clumps.

In an effort to determine SMBH masses from the dynamics of their surrounding molecular gas—the mm-Wave Interferometric Survey of Dark Object Masses (WISDOM) project (Onishi et al., 2017; Davis et al., 2017, 2018)—CO observations of nearby CNDs in ETGs were used to model morphology and kinematics of these molecular gas discs. Onishi et al. (2017) measured the CO brightness distribution of the fast-rotating ETG NGC 3665 with the Combined Array for Research in Millimeter Astronomy (CARMA). The authors find that the central molecular CND of this galaxy (which has a molecular gas mass of $\sim 8.1 \times 10^8 M_\odot$ Young et al., 2011) is regularly rotating with a rotation velocity of about 300 km s^{-1} . It is centrally concentrated with an exponentially declining surface density,




Figure 3.1: Left panel: Observed $^{12}\text{CO}(3-2)$ intensity map of the molecular gas disc in the centre of NGC 4429. The black ellipse in the lower left corner shows the beam size (i.e. the resolution) of the ALMA observation. Right panels: Velocity (upper panel) and velocity dispersion (lower panel) maps of $^{12}\text{CO}(3-2)$ emitting gas. The maps show a regularly rotating gas disc with a central hole. This figure is taken from Davis et al. (2018).

extends to a radius of about one kiloparsec, and appears to have a hole in the central ~ 65 pc, possibly created by the low-luminosity AGN ($f_{\text{Edd}} = L_{\text{X}}/L_{\text{Edd}} \sim 2 \times 10^{-7}$) that is observed in the galaxy (via its central X-ray luminosity L_{X} as well as a detected radio jet). The disc shows very little non-circular motion with a velocity dispersion of about 12.5 km s^{-1} , and only tiny amounts of star formation.

Davis et al. (2017) found a small molecular disc with a mass of $1.6 \times 10^7 M_{\odot}$ in the fast-rotating ETG NGC 4697 using CO observations of ALMA. The disc is regularly rotating with around 200 km s^{-1} , has an exponential surface density profile, and a radius of about 200 pc. Its velocity dispersion appears to be extremely low ($\sigma < 3 \text{ km s}^{-1}$), while its star formation surface density is also very low ($\Sigma_{\text{SFR}} \sim 0.001 - 0.017 M_{\odot} \text{ yr}^{-1} \text{ kpc}^{-2}$), i.e. the disc seems to be dynamically very cold and also quite stable against collapse into dense clouds and subsequent starformation. The authors speculate that morphological quenching might be responsible for stabilising the disc while keeping it dynamically cold. They also mention the possibility that hard X-ray radiation from the surroundings of the disc keeps the CO formation limited to the densest (and dynamically coldest) regions of the gas, reducing their observed (via CO) velocity dispersion.

Finally, Davis et al. (2018) analysed CO observations from ALMA of the fast-rotating ETG NGC 4429. As in the two previous cases, the galaxy contains a central, regularly rotating molecular CND with an exponentially declining surface density profile. The CND orbits a SMBH, which is believed to be a low-luminosity AGN due to observed unresolved

radio continuum emission (Nyland et al., 2016), as well as $H\alpha$ from a broad-line emission region (Constantin et al., 2015). Fig. 3.1 shows the observed intensity (a tracer of the surface density), velocity, and velocity dispersion maps of the CND. Like the one in NGC 3665, the CND in NGC 4429 has a central void of detected molecular, $^{12}\text{CO}(3-2)$ emitting gas, which has radius of about 50 parsec. The $^{12}\text{CO}(3-2)$ emitting disc then extends out to a radius of about 400 parsec, outside of which it abruptly ends, save for a few isolated low-surface brightness clouds. It has a total gas mass of circa $10^8 M_{\odot}$, no detected atomic hydrogen (Serra et al., 2012), and a surface density of a few hundred solar masses per square parsec. The disc’s rotation velocity varies between about 150 km s^{-1} close to the centre and about 250 km s^{-1} at the outer rim. Again, the velocity dispersion of the gas is very low with an estimated 2.2 km s^{-1} , and the star formation surface density, while not quite as low as in the case of NGC 4697, is still only about $0.1 M_{\odot} \text{ yr}^{-1} \text{ kpc}^{-2}$ (Davis et al., 2014), putting NGC 4429 well below the Kennicutt-Schmidt relation together with most of the other molecular-gas-containing ETGs measured in the ATLAS^{3D} survey.

Using their kinematically derived estimate of the central SMBH mass ($M_{\text{BH}} \sim 1.5 \times 10^8 M_{\odot}$), as well as of the stellar mass distribution (with a total stellar mass of $M_* \sim 1.5 \times 10^{11} M_{\odot}$), the authors calculate the Toomre parameter of the disc at different radii to analyse its stability, shown in Fig. 3.2. They find that the disc is marginally stable (with $Q \sim 2.5$) over most of its inner extent, with rising stability both towards the SMBH-dominated centre and towards the outer parts of the disc.

Overall, these observations show that there are CNDs in the centres of ETGs that are mostly (or even almost entirely) molecular and form stars at a much lower rate than their cousins in spiral galaxies at equivalent surface densities. Some of them also appear to have very low velocity dispersions and show a surprising level of dynamical stability given their high surface densities and low levels of turbulent support. In this chapter, we try to reproduce and explain these astonishing qualities with hydrodynamical simulations, focussing primarily on the CND in NGC 4429 (which we described in detail above). We set up our initial conditions to mimic the gas distribution of this disc and then analyse if we can reproduce its various properties, such as its star formation surface density, its velocity dispersion, its stability, and its molecular gas fraction with different assumed physical models, allowing us to (for example) test the effectiveness of morphological quenching in this scenario.

This chapter is structured as follows: We describe the numerical code and models used in our simulations, present our chosen initial conditions, and give an overview of the various simulation runs and their properties in section 3.2. Afterwards, we analyse first the evolution of the star formation and its suppression in our simulations in section 3.3, then the morphology and stability of the CND in section 3.4, followed by the evolution of the abundance of the various hydrogen species in section 3.5, and finally the impact of SMBH black hole accretion and feedback on the CND in section 3.6. We discuss our results in comparison to the observations and other simulations, and address our simulations’ caveats in section 3.7, before concluding with a summary of this chapter in section 3.8.




Figure 3.2: The estimated Toomre Q parameter for the CO-emitting molecular CND of NGC 4429. In the $^{12}\text{CO}(3-2)$ emitting region (between the dotted and the dashed vertical lines), the disc is marginally stable with $Q \sim 2.5$, while its stability increases both towards the very centre and towards the outer regions. This figure is taken from Davis et al. (2018).

3.2 Simulations

3.2.1 Numerical code and sub-resolution models

We have performed simulations of a CND in a massive ETG using the smoothed particle hydrodynamics code SPHGal (Hu et al., 2014), a modern version of GADGET-3 (Springel, 2005). This is the same basic code as in chapter 2, but using the pressure-energy formulation of SPH instead of the pressure-entropy one, which was used in the previous chapter. This switch is motivated by the better energy conservation properties in pressure-energy compared to pressure-entropy (see Hu et al., 2016, and also the discussion in chapter 4). In pressure-energy SPH, the equation of motions becomes:

$$\frac{d\vec{v}_i}{dt} = - \sum_{j=1}^N (\gamma - 1)^2 m_j u_i u_j \times \left[\frac{f_{ij}}{P_i} \nabla_i W_{ij}(h_i) + \frac{f_{ji}}{P_j} \nabla_i W_{ij}(h_j) \right], \quad (3.3)$$

$$f_{ij} = 1 - \left(\frac{h_i}{3(\gamma - 1)m_j u_j n_i} \frac{\partial P_i}{\partial h_i} \right) \left[1 + \frac{h_i}{3n_i} \frac{\partial n_i}{\partial h_i} \right]^{-1}, \quad (3.4)$$

with the kernel-averaged pressure P_i now defined via the specific internal energy u_i :

$$P_i = \sum_{j=1}^N (\gamma - 1) m_j u_j W_{ij}(h_i) \quad (3.5)$$

As before, \vec{v}_i is the velocity of particle i , N the number of its neighbours, m_j the neighbour particle's mass, γ the polytropic index, $W_{ij} = W(\vec{x}_i - \vec{x}_j)$ the smoothing kernel, h_i the smoothing length, and $n_i = \sum_{j=1}^N W_{ij}(h_i)$ the kernel-based particle number density.

The sub-resolution models for gas cooling, star formation, and feedback have also been adjusted from our previous work: For gas with a temperature above 30,000 K, we still use the metal-dependent equilibrium cooling tables implemented by Aumer et al. (2013), for which we track the abundances of 11 elemental species in the gas, which can change due to stellar feedback as well as metal diffusion between neighbouring gas particles. However, once the gas cools below the threshold of $T = 30,000$ K, we switch to a non-equilibrium cooling model using a simplified chemical network. This non-equilibrium chemistry and cooling model was implemented into our code by Hu et al. (2016) and is based on extensive earlier work by Nelson & Langer (1997), Glover & Mac Low (2007), and Glover & Clark (2012). Essentially the same model was used in the small-scale ISM simulation suite of the SILCC project (e.g. Walch et al., 2015; Girichidis et al., 2016).

The model tracks the formation and destruction of several chemical forms of hydrogen (molecular H_2 , atomic H^0 , ionized H^+), carbon and oxygen (CO , C^+ , O^0), as well as the abundance of free electrons e^- . However, the carbon and oxygen evolution is turned off in our simulations, since the densities at which their molecule formation becomes relevant are

not resolved. Of the three hydrogen species, only two (H_2 and H^+) are tracked directly, while the abundance of H^0 follows from hydrogen atom number conservation.

Several formation and destruction processes for the hydrogen species are modelled: Hydrogen can be ionized through collisions with free electrons and with cosmic rays (CRs), and recombines both in the gas phase and on dust grains. Similarly, molecular hydrogen is formed on the surfaces of dust grains and destroyed by photo-dissociation via photons from the interstellar radiation field, cosmic rays, or through collisions with other hydrogen molecules, atoms, or free electrons. These processes are summarised in Table 1 of Micic et al. (2012), which also contains references to the individual reaction rates for the listed interactions. Besides the density and temperature of the gas, the reaction rates also depend on the specific chemical abundances of the included species. The calculation of individual, species-abundance dependent reaction rates is what makes the model ‘non-equilibrium’ in contrast to the equilibrium cooling rates that are used for the hot gas (as well as for all gas in chapters 2 and 4).

All of the mentioned processes result in heating (photo-dissociation, H_2 formation, cosmic ray ionisation) or cooling (collisional dissociation and ionisation, H^+ recombination) of the gas. Together with cooling through various line-emissions (e.g. rotational and vibrational H_2 lines, Lyman- α), heating from the photo-electric effect on dust grains, and excitation of H_2 by UV photons, these processes make up the non-equilibrium cooling and chemistry model.


The evolution of H_2 (and the corresponding heating and cooling) is heavily affected by shielding from interstellar radiation, both due to dust, and due to the molecular hydrogen itself (self-shielding). These two effects depend on the column densities of dust and molecular gas, respectively. The column densities around each particle are calculated using the TREECOL algorithm (Clark et al., 2012), which utilises the HEALPIX algorithm (Górski & Hivon, 2011) to define $N_{\text{pix}} = 12$ pixels of equal area on a spherical surface around the particle and then calculates the column density for each of these pixels out to a shielding length set to $L_{\text{sh}} = 50$ pc. The shielding arises from the attenuation of the interstellar radiation flux as photons get absorbed or lose energy when they interact with dust and H_2 molecules while they travel deeper into a cloud of molecular gas. The attenuated flux then results in a decreased reaction (i.e. dissociation) rate of the interstellar radiation and the molecular gas.

Many of the mentioned processes (e.g. the photo-dissociation and shielding of molecular gas described above) require estimates of the interstellar radiation field and of the dust mass. The latter is defined via the dust-to-gas mass ratio, which we assume to be solar in our simulations. For the former, we make (in most of our simulations) the simplified assumption of a constant UV radiation field with an intensity of $G_0 = 6.8$ (given in Habing units, i.e. in units of the local solar neighbourhood value of $1.6 \times 10^{-3} \text{ ergs}^{-1} \text{ cm}^{-2}$, as measured by Habing, 1968). We arrive at this value from linearly scaling the UV intensity with the star-formation-rate surface density (following both Hu et al., 2016; Ostriker et al., 2010a), comparing the latter in the galactic centre of our quiescent ETG simulations in chapter 2 to the SFR surface density in the solar neighbourhood. This results in a very low estimate for G_0 , as the actual SFR surface density in our CNDS simulations of this chapter

is much higher than in the previous simulations (resulting from a much higher gas surface density).

To test the impact of the radiation field on the simulation results, we also perform runs with a much higher field strength of either $G_0 = 1000$ or $G_0 = 10^6$. $G_0 = 1000$ is in the range of observed UV field strengths in (the nuclei of) starburst galaxies (e.g. Ivison et al., 2010; Loenen et al., 2010), while it can be even higher ($G_0 \sim 10^4$) in AGN host galaxies (Tamura et al., 2007). The UV luminosities of low-luminosity AGN are observed to be in the range of $\sim 10^{40}$ - 10^{45} erg s $^{-1}$ (e.g. Telfer et al., 2002; Maoz, 2007; Malkan et al., 2017), corresponding to a UV intensity of $G_0 \sim 2.6 \times 10^2$ - 2.6×10^7 at distance of $r = 50$ pc from the SMBH, or $G_0 \sim 4.1$ - 4.1×10^5 at distance of $r = 400$ pc (which is about the extent of the CNB in NGC 4429), as the intensity declines as r^{-2} with increasing distance from the source. As mentioned above, the radiation field is assumed to be constant in our non-equilibrium cooling model; hence we do not take into account the decrease of AGN UV intensity with distance from the centre. Given this, together with the fact that the above estimate neglects any absorption and energy loss of UV photons along the way through the disc, the value of $G_0 = 10^6$ is most likely an overestimate of the actual UV radiation field intensity in NGC 4429's CNB, and therefore provides an upper limit for G_0 in our simulations.

Besides UV radiation, the chemistry and heating of the gas can be significantly influenced by cosmic rays, i.e. high energy protons and atomic nuclei. Within our chemical network, this influence is due to three reactions in particular (see Micic et al., 2012):



i.e. ionisation of atomic hydrogen, and dissociation of molecular hydrogen (without and with additional ionisation of one of the atoms) by cosmic rays. The rate of the first reaction is given by the CR ionisation rate ζ_{CR} , a free parameter in the model. The rates for reactions 3.7 and 3.8 are $2.22\zeta_{\text{CR}}$ and $0.037\zeta_{\text{CR}}$, respectively (Micic et al., 2012).

In most of our simulations, we ignore the role of cosmic rays by setting the CR ionisation rate to $\zeta_{\text{CR}} = 0$. However, we do perform a few runs in which we test the effect of cosmic rays on the evolution of our simulated CNB, assuming ionisation rates of $\zeta_{\text{CR}} \sim 10^{-14}$ s $^{-1}$, 10^{-13} s $^{-1}$ and 10^{-10} s $^{-1}$. The CR ionisation rates in AGN and nuclear starburst systems might be as high as a few hundred times that of the Galactic disc ($\zeta_{\text{CR}, \text{MW}} \sim 3 \times 10^{-16}$ s $^{-1}$, see González-Alfonso et al., 2013; van der Tak et al., 2016), hence $\zeta_{\text{CR}} = 10^{-14}$ s $^{-1}$, $\zeta_{\text{CR}} = 10^{-13}$ s $^{-1}$ are reasonable estimates for ζ_{CR} in our simulations. A rate of $\zeta_{\text{CR}} = 10^{-10}$ s $^{-1}$ is much higher, but we include it as a limiting case.

For the star formation, we use the same stochastic approach as in chapter 2 in which each gas particle that crosses a certain threshold in density and temperature has a probability to turn into a star particle such that the average star formation rate follows the

Schmidt relation (Schmidt, 1959), i.e:

$$\text{SFR} = \frac{\epsilon_{\text{SFR}} M_{\text{gas}}}{t_{\text{ff}}}, \quad t_{\text{ff}} = \frac{1}{\sqrt{4\pi G \rho_{\text{gas}}}}, \quad (3.9)$$

where SFR is the star formation rate, M_{gas} is the mass of star-forming gas, ϵ_{SFR} is the star formation efficiency (a free parameter), and t_{ff} is the free-fall time of the gas dependent on its density ρ_{gas} . As part of our studies in this paper, we vary the efficiency between $\epsilon_{\text{SFR}} = 2\%$ and $\epsilon_{\text{SFR}} = 0.2\%$.

Following Hu et al. (2017), and in contrast to our previous work, we exchange the separate temperature and density thresholds used in chapter 2 for a star formation threshold based on the definition of the Jeans mass:

$$M_{\text{J}} = \frac{\pi^{5/2} c_{\text{s}}^3}{6G^{3/2} \rho^{1/2}}, \quad (3.10)$$

where $c_{\text{s}} \propto \sqrt{u}$ is the speed of sound (with u the specific internal energy), and ρ the density of a given gas particle. The Jeans mass M_{J} is the maximum mass a gas cloud of a given size can have without collapsing under its own self-gravity. As long as M_{J} is numerically resolved in the simulation, such a collapse is followed directly by the SPH code. However, once a gas particle becomes dense and/or cold enough, its Jeans mass becomes smaller than the mass enclosed within its smoothing kernel (its kernel mass $M_{\text{ker}} = N_{\text{ngb}} m_{\text{gas}}$, where $N_{\text{ngb}} = 100$ is the number of neighbours, and m_{gas} is the gas particle mass), and is not resolved anymore. Once this condition, $M_{\text{J}} < M_{\text{ker}}$, is fulfilled for a gas particle, and it is situated within a converging flow (i.e. the velocity divergence within its neighbours is negative), we assume that it will collapse under its self-gravity, and thus enable the star formation model for it.

Once star particles form in the simulation, they can affect their surrounding gas via several feedback effects. As in chapter 2, we include models for feedback from supernovae type II and type Ia, as well as from AGB stars. Here we use the same implementation as in our previous work with one exception: Instead of injecting both momentum and thermal energy into the gas, depending on an estimate of the SN phase (see Núñez et al., 2017), we inject only thermal energy. As before, these feedback processes also inject metals into the neighbouring gas, where the composition of the ejecta depends on the type of the feedback.

The modelling of supernova feedback via the injection of thermal energy has a significant drawback: As shown by Hu et al. (2016), if the mass resolution of the gas is worse than about $m_{\text{gas}} \sim 1 M_{\odot}$, the blast wave generated by a supernova explosion is not properly resolved. With worse mass resolution, the conversion of thermal to kinetic energy in the Sedov-Taylor phase of the supernova blast wave progresses too slowly. Due to this, much of the thermal energy is radiated away by the cooling gas before it has time to accelerate the gas outwards, leading to a supernova feedback lacking in kinetic energy and ejecta velocity, which is less effective at driving away gas.

Hence, we experiment with an alternative supernova model in one of our simulation runs, in which we inject radial momentum (and velocity) directly into the gas, foregoing the

thermal Sedov-Taylor blast wave phase, which is likely not properly resolved. In this new model, which we call “mechanical” supernova feedback in this work, we assume that the conversion of thermal into kinetic energy during the Sedov-Taylor phase already happened on unresolved scales when we inject the feedback energy. We then take total momentum generated at the end of this phase p_{sf} from the analytical estimates of Kim & Ostriker (2015):

$$p_{\text{sf}} = 2.17 \times 10^5 M_{\odot} \text{ km s}^{-1} E_{51}^{0.93} n_0^{-0.13}, \quad (3.11)$$

where E_{51} is the energy of the supernova explosion in units of 10^{51} erg (i.e. the same energy we inject in thermal form in our previous SN feedback model), and n_0 is the hydrogen number density of the gas surrounding the supernova, i.e. the kernel-weighted average density of the SN-particle’s 10 closest neighbours into which the feedback energy is injected. We distribute this momentum kernel-weighted among the 10 closest gas particles to the stellar particle that is undergoing the supernova event, directly increasing their velocities radially away from the SN-particle accordingly. In this model, we do not inject any thermal energy into the gas particles affected by the supernova, though of course some of the kinetic energy will be transformed into heat via interactions with the surrounding gas.

In addition to the supernova feedback, we include a model for photo-ionisation heating of the gas by massive stars, implemented by Hu et al. (2017), into all of our simulations in this chapter. This model follows a Stroemgren approach, but takes into account possible inhomogeneities in the density of gas surrounding a star particle, as well as potential overlapping ionisation bubbles from several star particles. Assuming an ionising photon rate of 10^{48} photons per second per $100M_{\odot}$ of stellar mass, an ionisation radius is estimated iteratively by matching the ionising photon rate with the recombination rate in the surrounding gas, taking into account gas particles that are already ionised by a different star particle (see Hu et al., 2017, for details). Gas particles within the final estimate of the ionisation radius are flagged as ionised and heated to 10^4 K should their temperature be lower. Furthermore, their ionised hydrogen fraction is set to 1. Since this model ignores angular information, unrealistic outcomes may occur for particularly anisotropic density distributions. For example, a distant dense gas cloud could be ionised using most of the flux of a star particle that is otherwise surrounded by low density gas, even though it should only receive a small fraction of the radiation due to its small angular cross section. To mitigate this problem, the ionisation radius is artificially limited to a maximum of 50 pc.

In some of the simulations we performed in this paper, we include models for black hole accretion and feedback. The black hole feedback model is the same as in chapter 2, a combination of radiative and mechanical wind feedback implemented by Choi et al. (2012), using a wind efficiency of $\epsilon_w = 0.005$. However, we have removed the modified Bondi-Hoyle-Lyttleton accretion model we used before in favour a newly implemented simple “sink particle” approach to the accretion: We assume that every gas particle whose distance to the central black hole r decreases below an accretion radius of $r_{\text{acc}} = 10 \text{ pc}$, and whose relative velocity to the black hole is smaller than its circular velocity $v_{\text{circ}} = \sqrt{GM_{\text{BH}}/r}$ (with M_{BH} the black hole mass), is accreted immediately. The accretion rate of the black

hole is then defined as the sum of the masses of the gas particles accreted in a given time-step divided by the length of that time-step.

This model relies on the fact that the spatial resolution of the gas in the simulations of this chapter is high enough to directly track the inflow of gas towards the black hole until the innermost 10 parsecs. If a gas particle comes this close to a black hole, the accretion condition is simple: We assume that it is accreted immediately if it is slow enough to be bound to the SMBH when crossing the accretion radius. Due to the requirement that a gas particle has to come within a fixed distance to the black hole to be accreted—instead of having a variable, gas-density dependent accretion radius as in the modified Bondi-Hoyle-Lyttleton model used in chapter 2—this accretion model prevents the accretion of gas from the inner rim of the initial CND (which has a cavity of 50 pc radius in its centre).

Besides this advantage, the new accretion model also has errors that arise from its simplifying assumptions. Gas particles are accreted with a probability of 1 or 0 based on their relative distance to the black hole. This ignores that they represent an extended gaseous medium spread within the smoothing kernel of the particle. Due to this, our model underestimates the overall accretion rate of the black hole, particularly of low-density gas, which is represented by particles with large smoothing lengths. However, the error is less significant for dense gas, which dominates the accretion onto the black hole in these simulations of a dense CND. Another error arises from the accretion of discrete particles. Instead of computing a continuous accretion rate, which is used to determine AGN feedback, like in the model used in the other chapters, in this approach the accretion rate is discrete, given only by the total mass of accreted gas particles in a given time-step. However, due to the small gas particle masses (particularly compared to the SMBH mass) and overall low accretion rates in our simulations, this is not a significant problem either.

To test how strongly the results of our simulations are affected by the smoothness and lack of turbulence in our initial condition (see section 3.2.2), we perform one run in which we generate turbulence in the initial CND before switching on the star formation and stellar feedback models. This initial turbulence generation is achieved by the same method as used in Hu et al. (2017), an artificial supernova driving: while star formation and stellar feedback are disabled, we use the same Jeans-threshold to determine the dense gas particles that would otherwise be starforming. We then inject 10^{51} erg thermal energy into these particles and their neighbours within their smoothing kernel (as is done for the standard thermal SN feedback we use) with a probability proportional to $\epsilon_{\text{SFR}}/t_{\text{ff}}$ ($\epsilon_{\text{SFR}} = 0.2\%$ in the relevant run), i.e. the same probability with which the gas particle would be turned into a star if star formation were enabled. This method allows us to mimic the generation of turbulence in the gas by exploding supernovae without reducing the gas content in the CND due to star formation.

3.2.2 Initial conditions

We simulate the short term evolution of a circumnuclear gas disc embedded in the very centre of a massive early-type galaxy. The initial condition of our CND is based on the observed properties of the ETG NGC 4429 (Davis et al., 2018), which are described in

section 3.1. It is a dense, regularly rotating disc of molecular gas with a total mass of $M_{\text{gas}} = 10^8 M_{\odot}$, extending out to 400 pc from the centre and surrounding a central hole of 50 pc radius. The disc has an exponential surface density profile of the form:

$$\Sigma_{\text{gas}}(R) \propto \begin{cases} 0, & R < 50 \text{ pc} \\ \exp(-R/R_{\text{scale}}), & 50 \text{ pc} \leq R \leq 400 \text{ pc} \\ 0, & R > 400 \text{ pc} \end{cases}, \quad (3.12)$$

where R is the cylindrical radius, i.e. the distance from the centre in the plane of the disc, and $R_{\text{scale}} = 370 \text{ pc}$ is the scale radius (compare equation 1 in Davis et al., 2018). Together with the total gas mass, this gives a surface density of about $350 M_{\odot} \text{ pc}^{-2}$ at the inner border of the disc, and of about $150 M_{\odot} \text{ pc}^{-2}$ at its outer border as is shown in Fig. 3.3 (also compare the left panel of Fig. 3.4, which shows surface density maps of the initial condition from a face-on and from an edge-on perspective). For the vertical distribution of the gas we assume a one-dimensional density profile of the form:

$$n_{\text{gas}}(z) \propto \frac{2}{z_0} \frac{e^{2z/z_0}}{(1 + e^{2z/z_0})^2}, \quad (3.13)$$

where z is the vertical distance from the plane of the disc, and the scale height is set to $z_0 = 10 \text{ pc}$. This results in a geometrically thin disc around the centre (compare Fig. 3.4, right panel). The specifics of the initial vertical distribution are not important for the evolution of the disc, as (due to the lack of significant turbulent support) it will always contract into a thin disc in the beginning, before star formation and stellar feedback set in, which then dominate the vertical gas distribution for the rest of the simulation time.

The measured velocity dispersion of the circumnuclear gas disc in NGC 4429 is very low ($\sigma_{\text{gas}} \sim 2.2 \text{ km s}^{-1}$ Davis et al., 2018), hence we set up our initial condition with only tangential velocities in the gas, i.e. the gas rotates around the centre with a velocity equal to its circular velocity, and has zero velocity perpendicular to the disc plane. We assume the gas to initially have solar metallicity (a reasonable basic estimate for an ETGs central gas phase metallicity), a temperature of $T = 10^4 \text{ K}$, and a neutral hydrogen fraction of 1 (i.e. no ionized or molecular hydrogen). The latter two are arbitrary choices with very little influence on the evolution of the CND, as the gas cools down to a few tens of Kelvin within a few time-steps and forms corresponding amounts of H_2 simultaneously, irrespective of the precise initial values chosen.

The CND is randomly sampled with N_{gas} SPH particles of mass m_{gas} . Each particle is given a tangential velocity, which is only dependent the total enclosed mass $M_{\text{enc}}(r)$ at its radius r , defined to be the circular velocity at that radius:

$$v_{\text{circ}}(r) = \sqrt{\frac{GM_{\text{enc}}(r)}{r}}. \quad (3.14)$$

Besides the enclosed gas mass (which is given at any radius by the total gas mass and equations 3.12 and 3.13), the total enclosed mass in equation 3.14 also includes the mass




Figure 3.3: Initial face-on radial surface density profile of the CND. The red, dashed line shows the analytical profile given in equation 3.12, while the black, solid line is calculated directly from the gas particles in the initial condition.

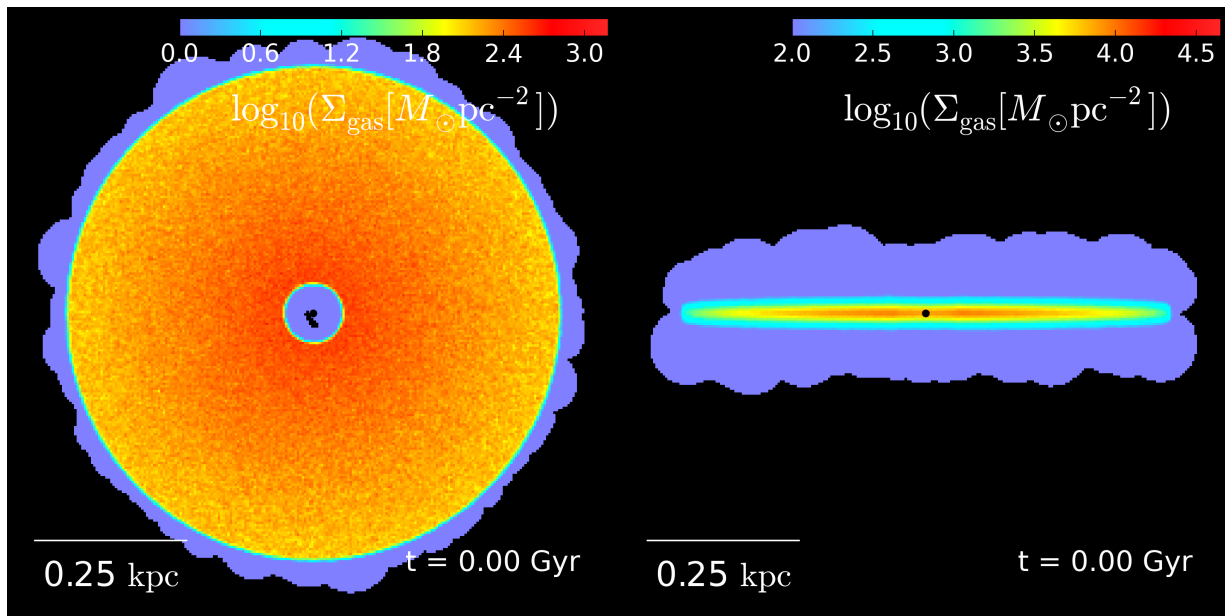


Figure 3.4: Maps of the gas surface density of the initial condition, from a face-on (left panel) and from an edge-on (right panel) perspective. The black circle (only visible edge-on) shows the position of the central SMBH. The gas disc is vertically thin, and has an exponentially declining radial surface density profile within the disc plane with a 100 pc wide hole in the centre.

of the central SMBH, as well as the enclosed masses of the stellar bulge and the dark matter halo. The black hole mass is set to $M_{\text{BH}} = 4 \times 10^8 M_{\odot}$ in most of our simulation runs, following our choice in chapter 2. We call the initial condition with this SMBH mass **IC_4e8BH**. In one set of runs, we instead choose $M_{\text{BH}} = 1.5 \times 10^8 M_{\odot}$, which is the measured mass for the SMBH in the centre of NGC 4429 (Davis et al., 2018). The initial condition with this lower black hole mass is called **IC_1.5e8BH**. To simplify the simulations, we fix the position of the SMBH in the centre of the galactic potential in of our simulations run, except for those in which we allow it to accrete gas. This choice does not significantly affect the overall evolution of the simulated CND outside of the central few tens of parsecs.

We model both the dark matter component of the galaxy, as well as its stellar bulge, as static background potentials mimicking spherical Hernquist density profiles (Hernquist, 1990) for both. Using static potentials instead of N-body particles allows us to capture the main influence of the DM and bulge components on the central gas evolution (i.e. their gravity), while reducing the computation time compared to what is necessary to simulate a large number of N-body particles. As our simulations only cover a relatively short time scale of about 150 Myr within which neither the stellar nor the DM distribution in the centre of an ETG varies much, and a small spatial scale of about half a kiloparsec around the centre, where the stellar component of ETGs is dominated by an old, spherical bulge, the use of static potentials is a reasonable approximation.

For the dark matter density profile, we assume the same parameters as in chapter 2: a virial mass of $M_{\text{DM, vir}} = 6.92 \times 10^{12} M_{\odot}$, and a scale length of $a_{\text{DM}} = 74.7$ kpc. For the stellar bulge density profile, we use two different parameter sets for the two different initial conditions. In IC_4e8BH, we again follow chapter 2, assuming a mass of $M_{\text{bulge}} = 8.41 \times 10^{10} M_{\odot}$, and a scale length of $a_{\text{bulge}} = 2.21$ kpc, thereby connecting the SMBH mass and the bulge mass via the relation of Kormendy & Ho (2013). In IC_1.5e8BH (where we use the SMBH mass of NGC 4429), we also take NGC 4429’s observed stellar mass of $1.5 \times 10^{11} M_{\odot}$ (Cappellari et al., 2011), and assume a scale length of $a_{\text{bulge}} = 1.5$ kpc, which leads to a very large concentration of stellar mass in the centre of the galaxy, resulting in a particularly high and steep circular velocity profile as observed in Davis et al. (2018).

In Fig. 3.5, we show the rotation curves (and the contributions of the various galactic components to it) for the two initial conditions (IC_4e8BH on the left and IC_1.5e8BH on the right). The black hole and the bulge dominate the gravitational potential (for which the circular velocity is a proxy) in the central 400 pc of the galaxy in which the CND resides, while the gas disc itself and the dark matter halo only contribute little to the potential. For IC_4e8BH, the gas disc rotates with a similar velocity of $\sim 170 \text{ km s}^{-1}$ at all radii, with only a moderate gradient. In contrast, for IC_1.5e8BH, the very concentrated, massive stellar bulge results in a steep rise of the rotation velocity towards larger radii, while also causing larger circular velocities overall (up to 275 km s^{-1} at 400 pc distance from the centre).

We compare simulations with two different gas particle masses ($m_{\text{gas}} = 100 M_{\odot}, 20 M_{\odot}$), corresponding to $N_{\text{gas}} = 10^6, 5 \times 10^6$ gas particles, respectively. We name the initial conditions and summarise their properties in Table 3.1. The naming scheme follows the




Figure 3.5: Radial profiles of the circular velocity (see equation 3.14) for the initial conditions. The left panel shows the rotation curves for IC_4e8BH assuming SMBH and bulge properties as in chapter 2, the right panel those for IC_1.5e8BH, which assumes NGC 4429-like SMBH and bulge mass distributions. The total circular velocity profile (green line), and the individual contributions by the various components (i.e. what v_{circ} would be if only the mass of the respective component existed) are shown. The solid black line shows the SMBH contribution, the red line that of the stellar bulge, cyan that of the gas disc, and the dashed black line that of the dark matter halo. The initial velocities of the gas particles are set such that they exactly match the total rotation curve. For IC_4e8BH, the SMBH dominates the potential in the inner ~ 150 pc. Beyond this distance, the bulge dominates, leading to a relatively flat rotation curve around 170 km s^{-1} . In contrast, for IC_1.5e8BH, the bulge is dominant everywhere except in the innermost 50 pc, and the resulting rotation curve is much steeper, rising from about 170 km s^{-1} at the inner edge of the gas disc to about 275 km s^{-1} at the outer edge.

Table 3.1: Overview of the initial conditions and their varying properties.

Name	$m_{\text{gas}} [M_{\odot}]$	N_{gas}	$M_{\text{BH}} [10^8 M_{\odot}]$	$M_{\text{bulge}} [10^{10} M_{\odot}]$	$a_{\text{bulge}} [\text{kpc}]$
IC_4e8BH_r100	100	10^6	4	8.41	2.21
IC_1.5e8BH_r100	100	10^6	1.5	15	1.5
IC_4e8BH_r20	20	5×10^6	4	8.41	2.21

simple rule of adding ‘_r’ and the mass resolution in solar masses to the name of the IC type given above (for the different bulge and SMBH properties).

3.2.3 Simulation runs

We simulate the evolution of our CND as described above with various alterations to our numerical model and its parameters, which determine the name of each run. From left to right, the names of the simulation runs are a string of handles for these parameters and models, each separated by a $_$, as explained below. We compare runs with two different mass resolutions, $m_{\text{gas}} = 100 M_{\odot}$ and $m_{\text{gas}} = 20 M_{\odot}$. The names of these runs start with **r100** and **r20**, respectively. The gravitational softening length is varied between the runs from $\epsilon_{\text{gas}} = 10 \text{ pc}$, over $\epsilon_{\text{gas}} = 4 \text{ pc}$, to $\epsilon_{\text{gas}} = 1 \text{ pc}$, corresponding to **s10**, **s4**, and **s1** in the names of the runs, respectively. The runs with $\epsilon_{\text{gas}} = 10 \text{ pc}$ and $\epsilon_{\text{gas}} = 4 \text{ pc}$ are run for $\sim 150 \text{ Myr}$, while those with $\epsilon_{\text{gas}} = 1 \text{ pc}$ are shorter test runs that only cover $\sim 15 \text{ Myr}$. We perform runs with two different star-formation efficiencies, $\epsilon_{\text{SFR}} = 2\%$ and $\epsilon_{\text{SFR}} = 0.2\%$, marked in the run names by **sf2** and **sf0.2**, respectively. The stellar potential and SMBH mass are varied between the two options described section 3.2.2, and referred to in the names of the runs by **BH4** for the initial condition based on chapter 2, and **BH1.5** for the one based on the observational data from NGC 4429.

In two test runs, we increase the constant UV field strength of the interstellar radiation field that is used in the non-equilibrium cooling model from the fiducial $G_0 = 6.8$ used in the other simulations to $G_0 = 10^3$ and $G_0 = 10^6$, marked by **UV1e3** and **UV1e6**, respectively, in the names of the runs. In three more test runs, we include ionisation by cosmic rays with ionisation rates of $\zeta_{\text{CR}} = 10^{-14} \text{ s}^{-1}$, $\zeta_{\text{CR}} = 10^{-13} \text{ s}^{-1}$, and $\zeta_{\text{CR}} = 10^{-10} \text{ s}^{-1}$, which in the run names correspond to **CR-14**, **CR-13**, and **CR-10**, respectively.

We test our new mechanical supernova feedback implementation in one run, which is marked in its name with **mFB**. In four runs, we allow the SMBH to move away from the centre. In two of those, we activate the black hole accretion model described in section 3.2.1, but not the AGN feedback. These runs’ names are marked with **acc**. In the other two, both the accretion and the AGN feedback model are active. They have **AGN** added to their names. Finally, in one run (whose name ends in **TD**), we let the simulation run for $\sim 15 \text{ Myr}$ without star formation, but instead with an artificial supernova driving to inject turbulence into the initial condition, before switching on the proper star formation and stellar feedback model instead. We give an overview of all the runs and their differing properties in the tables 3.2 (containing the used initial condition, softening length, star formation efficiency, UV field strength, and CR ionisation rate of the runs) and 3.3 (speci-

Table 3.2: Overview of the simulation runs and their varying properties, part 1.

Name	initial condition	ϵ_{gas}	ϵ_{SFR}	G_0	ζ_{CR}
r100_s10_sf2_BH4	IC_4e8BH_r100	10 pc	2%	6.8	0
r100_s10_sf0.2_BH4	IC_4e8BH_r100	10 pc	0.2%	6.8	0
r100_s10_sf2_BH1.5	IC_1.5e8BH_r100	10 pc	2%	6.8	0
r100_s10_sf0.2_BH1.5	IC_1.5e8BH_r100	10 pc	0.2%	6.8	0
r20_s4_sf0.2_BH4	IC_4e8BH_r20	4 pc	0.2%	6.8	0
r100_s1_sf0.2_BH4	IC_4e8BH_r100	1 pc	0.2%	6.8	0
r100_s1_sf0.2_BH4_UV1e3	IC_4e8BH_r100	1 pc	0.2%	10^3	0
r100_s1_sf0.2_BH4_UV1e6	IC_4e8BH_r100	1 pc	0.2%	10^6	0
r100_s1_sf0.2_BH4_CR – 14	IC_4e8BH_r100	1 pc	0.2%	6.8	10^{-14} s^{-1}
r100_s1_sf0.2_BH4_CR – 13	IC_4e8BH_r100	1 pc	0.2%	6.8	10^{-13} s^{-1}
r100_s1_sf0.2_BH4_CR – 10	IC_4e8BH_r100	1 pc	0.2%	6.8	10^{-10} s^{-1}
r100_s1_sf0.2_BH4_mFB	IC_4e8BH_r100	1 pc	0.2%	6.8	0
r100_s10_sf2_BH4_acc	IC_4e8BH_r100	10 pc	2%	6.8	0
r100_s10_sf0.2_BH4_acc	IC_4e8BH_r100	10 pc	0.2%	6.8	0
r100_s10_sf2_BH4_AGN	IC_4e8BH_r100	10 pc	2%	6.8	0
r100_s10_sf0.2_BH4_AGN	IC_4e8BH_r100	10 pc	0.2%	6.8	0
r100_s1_sf0.2_BH4_TD	IC_4e8BH_r100	1 pc	0.2%	6.8	0

fying which supernova feedback model the runs use, as well as if black hole accretion, AGN feedback, or preceding turbulent SN driving are active).

3.3 Star formation suppression

3.3.1 Comparing initial conditions and star formation efficiencies

One of the major peculiarities of observed CNDs in ETGs is their low SFR surface densities Σ_{SFR} when compared to spiral galaxies of comparable gas mass surface densities Σ_{gas} . We compare the position of our simulated CND in the $\Sigma_{\text{SFR}} - \Sigma_{\text{H}+\text{H}_2}$ plane (where $\Sigma_{\text{H}+\text{H}_2}$ is the mass surface density of the neutral (atomic and molecular) hydrogen) to those of observed neutral gas discs in ETGs (from Davis et al., 2014). To start with, we compare the four runs r100_s10_sf2_BH4, r100_s10_sf0.2_BH4, r100_s10_sf2_BH1.5, and r100_s10_sf0.2_BH1.5 at the end of the simulation time (i.e. after ~ 150 Myr) in Fig. 3.6. These runs differ in their star formation efficiency parameter ϵ_{SFR} , as well as their initial condition (i.e. the mass of the central SMBH and the shape and depth of the stellar gravitational potential well). Beside the simulation results, Fig. 3.6 also shows the observations of Davis et al. (2014)—highlighting the position of NGC 4429, which our initial conditions are based on—and the KS relation for star-forming galaxies for comparison.

In all of the compared simulations, $\Sigma_{\text{H}+\text{H}_2}$ is lower and Σ_{SFR} higher than the value for NGC 4429. The runs with $\epsilon_{\text{SFR}} = 2\%$ (squares) have much lower $\Sigma_{\text{H}+\text{H}_2}$, placing them above

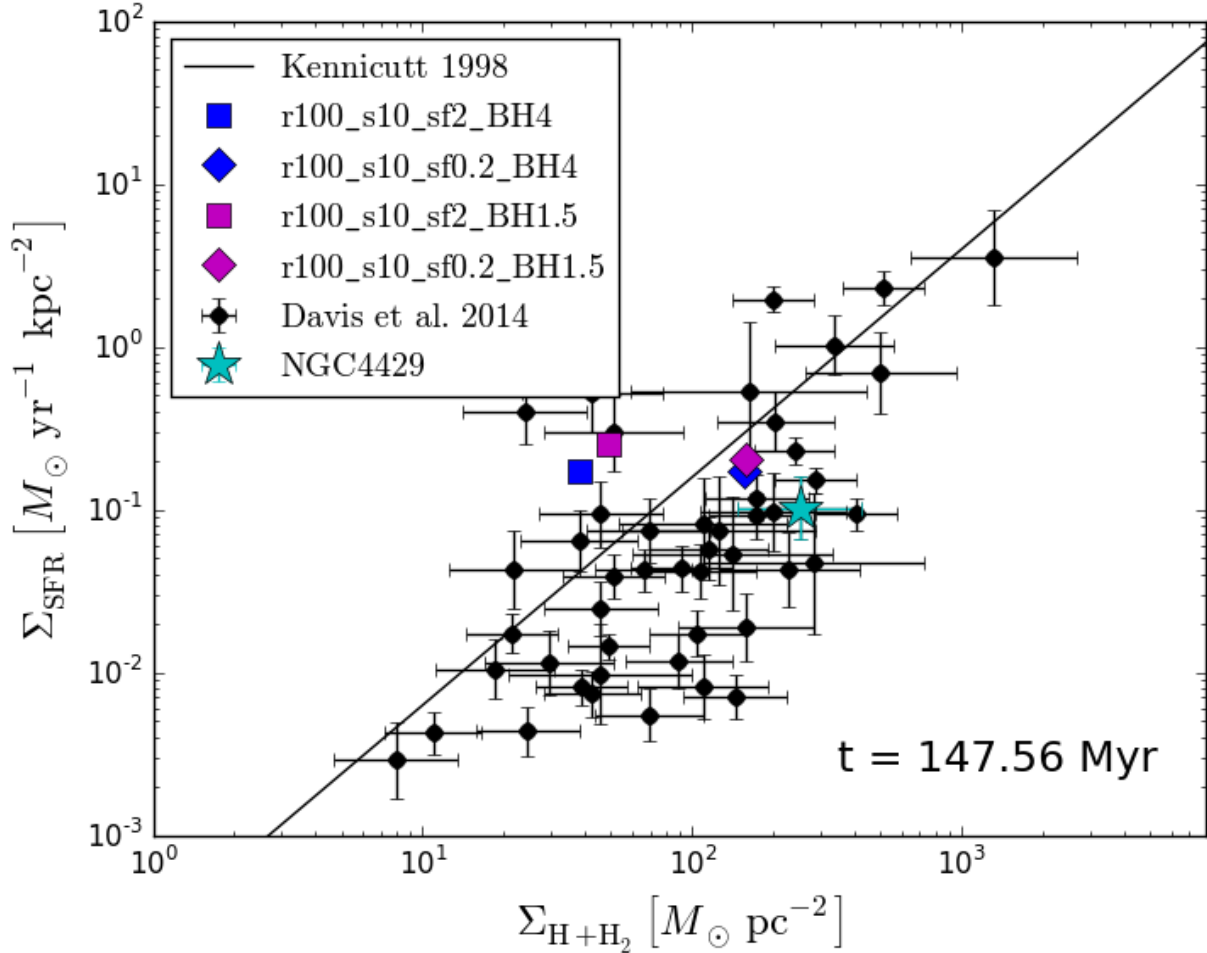


Figure 3.6: Star formation rate surface density over neutral hydrogen mass surface density of the CND at the end of the simulation time for the runs with different star formation efficiencies and initial conditions (see legend). The black line shows the Kennicutt (1998) relation, the black dots with error bars the observed values of cold gas in ETGs from Davis et al. (2014), with the cyan star marking the observations for NGC 4429 from the same work. For $\epsilon_{\text{SFR}} = 2\%$, the simulated CND falls above the Kennicutt relation, far away from the majority of the observed systems. It is closer to those in the $\epsilon_{\text{SFR}} = 0.2\%$ case, but still at a lower gas surface density and higher SFR surface density than the observed NGC 4429. The chosen initial condition has little impact.

Table 3.3: Overview of the simulation runs and their varying properties, part 2.

Name	SN feedback	BH accretion	AGN feedback	SN driving
r100_s10_sf2_BH4	thermal	no	no	no
r100_s10_sf0.2_BH4	thermal	no	no	no
r100_s10_sf2_BH1.5	thermal	no	no	no
r100_s10_sf0.2_BH1.5	thermal	no	no	no
r20_s4_sf0.2_BH4	thermal	no	no	no
r100_s1_sf0.2_BH4	thermal	no	no	no
r100_s1_sf0.2_BH4_UV1e3	thermal	no	no	no
r100_s1_sf0.2_BH4_UV1e6	thermal	no	no	no
r100_s1_sf0.2_BH4_CR – 14	thermal	no	no	no
r100_s1_sf0.2_BH4_CR – 13	thermal	no	no	no
r100_s1_sf0.2_BH4_CR – 10	thermal	no	no	no
r100_s1_sf0.2_BH4_mFB	mechanical	no	no	no
r100_s10_sf2_BH4_acc	thermal	yes	no	no
r100_s10_sf0.2_BH4_acc	thermal	yes	no	no
r100_s10_sf2_BH4_AGN	thermal	yes	yes	no
r100_s10_sf0.2_BH4_AGN	thermal	yes	yes	no
r100_s1_sf0.2_BH4_TD	thermal	no	no	yes

the KS relation and more than order of magnitude away from the bulk of the observations. If the star-formation efficiency is set an order of magnitude lower ($\epsilon_{\text{SFR}} = 0.2\%$, diamonds), Σ_{SFR} is about the same at the end of the simulation, but $\Sigma_{\text{H}+\text{H}_2}$ is about a factor of 4-5 higher, placing these runs much closer to observations, but still towards the high end of their distribution. In both cases, the chosen initial condition (blue or purple colour) has only a small impact, and does not change the overall result.

In Fig. 3.7 we show the time evolution of the SFR for the same four runs. From this we can discern the cause of the differences between the runs with $\epsilon_{\text{SFR}} = 2\%$ and $\epsilon_{\text{SFR}} = 0.2\%$. In the runs with low star-formation efficiency, the SFR is quite constant over the whole simulation time, starting with a short peak at $\text{SFR} \sim 0.2 M_{\odot} \text{yr}^{-1}$ and then slowly declining to about half that value by the end of the simulation, as gas is consumed to form stars, and the surface density of the CNB is slowly decreasing. At the end of the simulation, about 15% of the original gas mass was transformed into stars.

On the other hand, in the runs with high star-formation efficiency, the SFR starts of an order of magnitude higher, at about $\text{SFR} \sim 1 M_{\odot} \text{yr}^{-1}$, consuming gas at a much higher rate and leading to more powerful SN feedback, which in turn drives some gas out of the CNB entirely and some gas into denser substructures, increasing the SFR further. All of this leads to a decrease in the gas mass by 40-50% and a corresponding drop in surface density. Eventually, the density decrease leads to a decline in the SFR to levels similar to those of the $\epsilon_{\text{SFR}} = 0.2\%$ runs, but by this point $\Sigma_{\text{H}+\text{H}_2}$ declined so much that the system now places above the KS relation. Again, we see that the chosen IC has little influence on




Figure 3.7: Star formation rate over time comparing runs with different star formation efficiencies and initial conditions (see legend). The star formation rate is almost constant (slowly declining) at about $0.1 M_{\odot} \text{ yr}^{-1}$ for a star-formation efficiency of $\epsilon_{\text{SFR}} = 0.2\%$, while it is an order of magnitude higher for $\epsilon_{\text{SFR}} = 2\%$ until much of the gas is consumed and ejected from the CND and the SFR falls towards the $\epsilon_{\text{SFR}} = 0.2\%$ run levels. The black hole mass, and the stellar mass and size differences between the two initial conditions have no significant impact on the star formation.

the evolution of the simulation.

We can further see the differences caused by changing ϵ_{SFR} by comparing the density-temperature phase diagrams of the gas in the four runs at the end of the simulation time (Fig. 3.8, where the colour marks the number density of gas particles at a given temperature and density, from cyan at the low end to purple at the high end). The black dashed line in this figure marks the star-formation threshold for these four simulations, i.e. gas that is to the lower right of this line is allowed to form stars. The gas is split into three phases: cold, dense, mostly molecular, star-forming gas; warm, mostly ionised gas at overlapping, but slightly lower densities; and hot, fully ionised gas at the lowest densities.

The hot phase is created by SN explosions and then cools down to the warm phase, which then further cools, replenishing the cold, star-forming phase. In all of the four runs, the cold phase reaches similar maximum densities. As the SFR is only dependent on the density and on ϵ_{SFR} (once the gas passes the star-formation threshold), this results in higher SFR for the runs with $\epsilon_{\text{SFR}} = 2\%$ compared to those with $\epsilon_{\text{SFR}} = 0.2\%$. The subsequently stronger SN feedback then leads to the creation of more populated hot and warm phases in these high star-formation efficiency runs. Which initial condition is chosen barely matters at all for the phase distribution of the gas.

From these results, particularly those shown in Fig. 3.6, we can draw some conclusions. Our simulations are not able to match the low SFR surface densities of most observed CNDs in ETGs, and in particular not that of NGC 4429, even though the gas disc resides in a deep potential well and has large rotation velocities. In other words, morphological quenching—the suppression of star formation by the deep, concentrated stellar potential—is not working (at least not to the extent necessary to match the observations). This is true for both initial conditions, which vary significantly in both the slope and the overall height of the rotation curve (i.e. the important factors for stabilising the disc through the stellar morphology). It is possible that even more extreme rotation speeds (e.g. some of the observed CNDs of Boizelle et al. (2017) have rotation velocities in excess of 500 km s^{-1}) could have an impact, but the stellar and SMBH potential in IC.1.5e8BH_r100 is based directly on the observed NGC 4429, so more extreme values should not be necessary to match its star-formation properties.

3.3.2 Resolution comparison

While it could be argued that the r100_s10_sf0.2_BH4 run fits the position of NGC 4429 in the $\Sigma_{\text{SFR}} - \Sigma_{\text{H}+\text{H}_2}$ plane at least within the latter’s observational error range, the disagreement between observation and simulation is actually significantly worse, as the SFR in the r100_s10_sf0.2_BH4 run is artificially suppressed due to the limited resolution. To understand this, we compare this run (with $m_{\text{gas}} = 100 M_{\odot}$, $\epsilon_{\text{gas}} = 10 \text{ pc}$) to one with a better mass resolution of $m_{\text{gas}} = 20 M_{\odot}$ and therefore a shorter gravitational softening length for the gas particles of $\epsilon_{\text{gas}} = 4 \text{ pc}$ (run r20_s4_sf0.2_BH4), as well as to a run (r100_s1_sf0.2_BH4) with the same mass resolution, but a ten-times shorter softening length of $\epsilon_{\text{gas}} = 1 \text{ pc}$. In Fig. 3.9 we compare the time evolution of the SFR in these three runs. Shortening the gravitational softening length results in a higher SFR (until gas consump-




Figure 3.8: Density-temperature phase diagrams of the gas at the end of the simulation time for the runs with 2% (left column) or 0.2% (right column) star formation efficiency, and using the IC_4e8BH_r100 (top row) or IC_1.5e8BH_r100 (bottom row) initial condition. The dashed line shows the temperature-density pairs at which the Jeans mass is equal to the kernel mass. The gas is split into three distinct phases: Hot, fully ionized gas at $T \sim 10^6 - 10^7$ K, warm ionized gas at $T \sim 10^4$ K, and cold neutral (mostly molecular) starforming gas at $T \sim 10 - 100$ K. The colder gas phases reach higher densities. At $\epsilon_{\text{SFR}} = 2\%$, more gas resides in the warm and (especially) the hot phase than in the $\epsilon_{\text{SFR}} = 0.2\%$ runs. The differences between the runs using different initial conditions are minor.




Figure 3.9: Star formation rate over time comparing runs with different mass resolutions and gravitational softening lengths at $\epsilon_{\text{SFR}} = 0.2\%$ (see legend). At better mass resolution and with shorter softening lengths, the SFR is higher (until it declines again due to the diminishing gas supply).

tion by star formation and ejection by SN feedback reduce it towards the end).

Unsurprisingly, this leads to worse match of these higher-resolution simulations to the observations in the $\Sigma_{\text{SFR}} - \Sigma_{\text{H}+\text{H}_2}$ relation plot, as is shown in Fig. 3.10 for the end of the simulation time. The runs with shorter ϵ_{gas} lie right on the KS relation for star-forming galaxies, significantly above the majority of the observed systems, and have a neutral gas surface density more than a factor of three below that of NGC 4429. The reasons for the difference are the same as in the runs with a higher star-formation efficiency: higher SFR leading to more consumed and expelled gas, reducing $\Sigma_{\text{H}+\text{H}_2}$.

We can understand why the change in softening length leads to increased star formation by looking at the phase diagrams of the three runs, which are shown in Fig. 3.11. Comparing the distribution of the gas in the cold, star-forming phase, we can see that it reaches higher densities in the runs with shorter softening lengths: In r100_s10_sf0.2_BH4, the gas does not extend beyond $\rho \sim 10^{-20} \text{ g cm}^{-3}$, while in r20_s4_sf0.2_BH4, the maximum density is $\rho \sim 10^{-19} \text{ g cm}^{-3}$, and in r100_s1_sf0.2_BH4, it is as high as $\rho \sim 8 \times 10^{-18} \text{ g cm}^{-3}$.

The denser gas is the cause of the higher SFRs. The gravitational softening length determines the smallest scale on which the gravitational attraction between neighbouring gas particles can increase if their distance decreases. This essentially limits to how small (and therefore to how dense) a structure the gas in the simulation can collapse. Hence, a larger gravitational softening length decreases the maximum density of the matter in the simulation, thereby affecting the SFR in a Schmidt-type, density-dependent star-formation model.

What is a reasonable softening length for a simulation depends on its resolution, i.e. the mass of its particles in an SPH code, with better resolved simulations (i.e. those with lower particle masses) needing smaller softening lengths. While there is no consensus about the “correct” softening length for a given mass resolution, the simulations are clearly not resolving the star formation properly. If the star formation were properly resolved, the runs with varying star-formation efficiency should result in about the same SFR, because, in the run with the lower efficiency, the gas—being consumed less quickly by star formation—would reach higher densities, raising the SFR until it is counteracted by stellar feedback. However, in our r100_s10_sf0.2_BH4 run (with the large softening length), the star formation is not balanced by the feedback, but by the artificial pressure from the gravitational softening that prevents further collapse of the gas. This is confirmed by the better resolved r20_s4_sf0.2_BH4 run showing a higher SFR and denser gas.

3.3.3 Alternative star formation suppression mechanisms

Are there other effects that might suppress the star formation in the CND, which we are not including in our simulations? One possibility is that the interstellar UV radiation field is stronger than the field strength of $G_0 = 6.8$ that we have assumed so far. The CND we simulate is dense, and situated in an extreme environment: a galactic nucleus with a massive, concentrated stellar bulge and a SMBH at the centre (that is possible accreting at a low rate); so a stronger UV field is plausible. Therefore, we perform two test runs with a very high ($G_0 = 1000$, run r100_s1_sf0.2_BH4_UV1e3), and an extremely high ($G_0 = 10^6$,




Figure 3.10: Star formation rate surface density over neutral hydrogen mass surface density of the CND at the end of the simulation time for the runs with different mass resolutions and gravitational softening lengths at $\epsilon_{\text{SFR}} = 0.2\%$ (see legend). The black line shows the Kennicutt (1998) relation, the black dots with error bars the observed values of cold gas in ETGs from Davis et al. (2014), with the cyan star marking the observations for NGC 4429 from the same work. In the runs with better mass resolution or a shorter softening length, the higher SFR leads to a decline in the gas surface density by the end of the simulation, taking the simulated CND further from NGC 4429, and most of the other observed values.




Figure 3.11: Density-temperature phase diagrams of the gas at the end of the simulation time for the runs with different mass resolutions and gravitational softening lengths at $\epsilon_{\text{SFR}} = 0.2\%$. From the left to the right are the runs with $m_{\text{gas}} = 100 M_{\odot}$, $\epsilon_{\text{gas}} = 10$ pc; with $m_{\text{gas}} = 10 M_{\odot}$, $\epsilon_{\text{gas}} = 4$ pc; and with $m_{\text{gas}} = 100 M_{\odot}$, $\epsilon_{\text{gas}} = 1$ pc. The dashed line shows the temperature-density pairs at which the Jeans mass is equal to the kernel mass. Shorter gravitational softening lengths result in denser gas, which increases both the star formation rate and (due to the increased stellar feedback) the amount of hot gas.

run r100_s1_sf0.2_BH4_UV1e6) UV field strength to test the possible impact of strong UV radiation on the star formation (see section 3.2.1 for a more detailed motivation of these values for G_0).

Another process that could potentially suppress the star formation are cosmic rays that might (fractionally) ionise the gas, and in the process stabilise it to a degree. The CR ionisation rates in AGN and nuclear starburst systems might be as high as a few hundred times that of the Galactic disc ($\zeta_{\text{CR}} \sim 3 \times 10^{-16} \text{ s}^{-1}$, see González-Alfonso et al., 2013; van der Tak et al., 2016, compare also section 3.2.1). Hence, we include two test runs in which we include the CR ionisation module of the chemical network with ionisation rates of $\zeta_{\text{CR}} = 10^{-14} \text{ s}^{-1}$ (run r100_s1_sf0.2_BH4_CR – 14) and $\zeta_{\text{CR}} = 10^{-13} \text{ s}^{-1}$ (run r100_s1_sf0.2_BH4_CR – 13). For all of these test runs, we use a gas mass resolution of $m_{\text{gas}} = 100 M_{\odot}$ and a softening length of $\epsilon_{\text{gas}} = 1$ pc. The lower mass resolution keeps the simulations relatively quick, while the small softening length lets us test the effect of the tested physical processes in simulations where the SFR is less artificially suppressed. We also only run these test simulations for a tenth of the time of the previous simulations, i.e. ~ 15 Myr, as we are interested in the immediate effect of the changes on the early high SFR, which is, in any case, determining the latter evolution of the CND through gas depletion and SN feedback in any case.

In Fig. 3.12, we compare the time evolution of the SFR in the four runs just described to the corresponding “standard” run, r100_s1_sf0.2_BH4. Aside from the first ~ 2 Myr, in which a CR ionisation rate of $\zeta_{\text{CR}} = 10^{-13} \text{ s}^{-1}$ leads to a slight decrease of SFR, neither a stronger UV radiation field, nor CR ionisation has any significant effect on the star formation. As a result, the position of the simulated CND after these ~ 15 Myr in the $\Sigma_{\text{SFR}} - \Sigma_{\text{H}+\text{H}_2}$ relation (shown in Fig. 3.13) barely differs between the five runs, with all of them showing SFR surface densities above the KS relation, and far above NGC 4429.




Figure 3.12: Star formation rate over time comparing runs with varying UV field strengths and CR ionisation rates, as well as with mechanical—instead of thermal—SN feedback (see legend). Within the range of tested values shown, changing either the UV field strength or the CR ionisation rate has little effect on the star formation rate. Using mechanical instead of thermal SN feedback changes the precise time evolution of the SFR, but does not change its overall high values.




Figure 3.13: Star formation rate surface density over neutral hydrogen mass surface density of the CND after ~ 15 Myr for the runs with varying UV field strengths and CR ionisation rates, as well as with mechanical—instead of thermal—SN feedback (see legend). The black line shows the Kennicutt (1998) relation, the black dots with error bars the observed values of cold gas in ETGs from Davis et al. (2014), with the cyan star marking the observations for NGC 4429 from the same work. Within the range of tested values shown, the changes to the UV field and CR ionisation rate do not significantly affect the gas or the SFR surface density of the simulated CND, which is far above the observed NGC 4429. The same is true if the thermal SN feedback model is replaced by the mechanical model.




Figure 3.14: Density-temperature phase diagrams of the gas after ~ 15 Myr for the runs with varying UV field strengths and CR ionisation rates. In the top row are the runs without CR ionisation and varying UV field: From the left to right, G_0 varies from 6.8, over 1000, to 10^6 . In the bottom row $G_0 = 6.8$, but the cosmic ray ionisation rate varies (from left to right) between $\zeta_{\text{CR}} = 10^{-14} \text{ s}^{-1}$, $\zeta_{\text{CR}} = 10^{-13} \text{ s}^{-1}$, and $\zeta_{\text{CR}} = 10^{-10} \text{ s}^{-1}$. The dashed line shows the temperature-density pairs at which the Jeans mass is equal to the kernel mass. The UV field strength has only little influence on the phase distribution of the gas: the hot phase starts cooling at slightly higher densities for higher G_0 and the very dense gas is slightly warmer at $G_0 = 10^6$. Higher CR ionisation rates lead to warmer dense gas whose temperature distribution is tightened. At extremely high ionisation rates ($\zeta_{\text{CR}} = 10^{-10} \text{ s}^{-1}$), all the gas is constantly in the warm ionised phase, with neither a cold, not a hot phase existing. Note that the straight lines at 10^4 K as well as the curved lines below that temperature are the result of the photo-ionisation heating and subsequent cooling of dense gas.

We can see the effect of the UV field and the CR ionisation on the gas directly in the temperature-density phase diagrams shown in Fig. 3.14. A higher UV field strength (top centre and top right panels) leads to a slightly hotter and more populated warm phase, while reducing the amount of gas on the low-density end of the cold phase somewhat. Its effect on the star-forming gas (below the dashed black line) is not very significant though. Especially the very dense gas ($\rho > 10^{-20} \text{ g cm}^{-3}$), which dominates the SFR, is barely effected by raising the UV field strength to $G_0 = 1000$, and only slightly heated (but still beyond the star-formation threshold) for $G_0 = 10^6$. Including CR ionisation (lower panels) has a larger effect on the gas phase distribution: Higher ionisation rates lead to significantly warmer dense gas with a less spread-out temperature distribution. Nonetheless, for ionisation rates of both $\zeta_{\text{CR}} = 10^{-14} \text{ s}^{-1}$ and $\zeta_{\text{CR}} = 10^{-13} \text{ s}^{-1}$, the very dense gas is unaffected, and hence the influence of the cosmic rays on the SFR is negligible. This only changes when we increase the ionisation rate even further, to $\zeta_{\text{CR}} = 10^{-10} \text{ s}^{-1}$ (run r100_s1_sf0.2_BH4_CR – 10, shown in the bottom right panel in Fig. 3.14). For this extremely high ionisation rate, all of the gas stays ionised at all times, shutting of star formation in the CND completely. Of course, this also means that, in this run, the CND is fully ionised instead of molecular, which completely contradicts the observations.

Neither UV background radiation nor CR ionisation (at least within realistic limits and while keeping the disc molecular) are capable of suppressing the star formation in the CND. Another reason for the disparity between the simulations and the observations could be the implementation of the SN feedback. In all of the previous runs, the SN feedback is implemented as a kernel-weighted dump of thermal energy in the gas particles neighbouring the star particle that undergoes a supernova. The problem with this implementation is that the Sedov-Taylor phase of the supernova blast wave, in which most of the thermal energy of the ejecta is converted to kinetic energy, is not properly resolved at our mass resolution for SNe exploding in the dense gas (which is most of them). Because of this, the blast wave cools down before enough energy can be converted. Compare Hu et al. (2016), where the authors argue that a mass resolution of $\sim 1 M_{\odot}$ is necessary to resolve the blast wave in gas with a density of $\rho \sim 10^{-22} \text{ g cm}^{-3}$. This leads to an underestimate of the kinetic energy and velocity of the supernova ejecta. Therefore, the effect of the SN feedback might be underestimated in our simulations. Hence, we test how the SFR changes if we exchange the thermal SN implementation with our new mechanical model (run r100_s1_sf0.2_BH4_mFB, see section 3.2.1). We show the SFR evolution of this run in Fig. 3.12, and its position in the Kennicutt-Schmidt diagram in Fig. 3.13. Using the mechanical SN feedback model changes the evolution of the SFR in detail, but it does not prevent it from rising to $\sim 1 M_{\odot} \text{ yr}^{-1}$ and correspondingly high SFR surface densities. While the mechanical SN feedback does lead to a different gas morphology than the thermal implementation (see section 3.4), it is not able to prevent the formation of very dense clumps of gas. However, it is this very dense tail of the gas density distribution that dominates the SFR, and hence the mechanical SN feedback does not help in suppressing the SFR.

Overall, neither morphological quenching, nor any of the other mechanisms we tried, is able to reduce the SFR in the CND enough for the simulation to agree with the observed value, if the gas is not artificially stabilised against collapse due to a lack of resolution.




Figure 3.15: Maps of the gas surface density (top row) and the density-weighted average temperature (bottom row) in a 1 kpc thick, 1x1 kpc wide region around the galactic centre. The columns show, from left to right, the states at the end of the simulation time for the runs using IC_4e8BH_r10 with $\epsilon_{\text{SFR}} = 2\%$ and $\epsilon_{\text{SFR}} = 0.2\%$, as well as those using IC_1.5e8BH_r100 with $\epsilon_{\text{SFR}} = 2\%$ and $\epsilon_{\text{SFR}} = 0.2\%$. The black dot in the centre marks the position of the supermassive black hole. In all cases, the initial void in the centre of the CND is filled in by the end of the simulation time. At low star-formation efficiency, the CND turns into a dense, cold, very tightly wound, starforming spiral, still with an exponential radial surface density profile. At high star-formation efficiency, the surface density of the gas has drastically decreased by the end of the simulation, as a high SFR (and the following SN feedback) causes the destruction of the initial CND, and the left-over gas later settles into a lower-density, looser spiral. The initial condition has no significant influence on the final state of the CND.

An extremely high cosmic ray ionisation rate is able to prevent star formation, but it does so entirely (while the observed CNDs still have low levels of ongoing star formation) and by keeping the entire disc warm and ionised, which does not agree with the observed discs being predominantly molecular. While less extreme cosmic ray ionisation can heat some of the dense gas without ionising it completely, it is not able to prevent the formation of significant amounts of very dense star forming gas, which dominates the total SFR.

3.4 Morphology, stability and velocity dispersion

After analysing the star formation and phase distribution of the gas in our simulations, we now investigate how its morphology develops and influences the stability of the CND in

the different runs. First, as in section 3.3, we again compare the simulations with different star-formation efficiencies and initial conditions. In Fig. 3.15, we show the gas surface density (top row) and temperature (bottom row) at the end of the simulation time for the four runs. In the runs with $\epsilon_{\text{SFR}} = 2\%$ (first and third column), the initial dense disc spanning the inner 400 pc (compare Fig. 3.4) with a void in the inner 50 pc is mostly gone by the end of the simulation. What is left is a small, very dense disc in the very centre, surrounded by loose spiral arms, and with an overall much lower surface density. The very centre, as well as the spiral arms are a few tens of Kelvin cold, while the less dense regions between the arms are warmer, reaching a few hundred Kelvin. Hot spots of up to a few tens of parsec in radius are created by SN explosions.

In contrast, in the $\epsilon_{\text{SFR}} = 0.2\%$ runs (second and last column), even though the disc does not stay smooth, but fractures into a tightly wound spiral, its surface density decreases a lot less and keeps its exponential radial profile. Almost all of the disc is now cold, $T \lesssim 30$ K. For all four runs, the central void is filled in by cold gas from the disc over the course of the simulation. As is to be expected from the very similar SFR evolutions and phase distributions, the choice of initial condition has no significant effect on the gas morphology, with the exception of the very centre: a more massive SMBH leads to a denser central gas accumulation.

The differences in morphology between runs with different star-formation efficiencies are driven by star formation and SN feedback, as can be seen in Fig. 3.16, which is the same as Fig. 3.15, but for the state of the simulations after ~ 15 Myr (a tenth of the total simulation time). With low star-formation efficiency, the CND looks essentially the same at this early time as it does by the end of the simulation, with the exception of a less filled in central void. On the other hand, with the high star-formation efficiency, we see the effects of the initial high SFR on the gas morphology. Large cavities are being created by many SN explosions in close succession, evacuating parts of the disc and driving gas into dense rings where the SN bubbles meet. In this process, the initial surface density profile of the CND is destroyed.

In Fig. 3.17, we show how the stability of the gas (estimated by the Toomre Q , top row) and the line-of-sight velocity dispersion of the molecular hydrogen (bottom row) at the end of the simulation time differ among the four discussed simulations with varying star-formation efficiency and initial condition. As the gas in these simulations is very cold, it is stabilised by its velocity dispersion σ , instead of its thermal pressure (represented by the sound speed c_s). Hence, we use the Toomre Q as defined in equation 3.1, i.e.

$$Q = \frac{\kappa\sigma}{\pi G \Sigma_{\text{gas}}}. \quad (3.15)$$

We estimate the epicyclic frequency κ in 100 evenly spaced cylindrical radial bins within the inner 400 pc (with the rotation axis assumed to be along the z-axis, identical to the setup of the initial condition), assign each gas particle a κ according to the bin it falls into, and then map the particle data onto a 2D grid taking into account the SPH kernels to generate the Toomre Q map. Due to the binning of κ , there are no properly defined values for Q outside of a radius of 400 pc, which is why the maps show only white beyond this




Figure 3.16: Maps of the gas surface density (top row) and the density-weighted average temperature (bottom row) in a 1 kpc thick, 1x1 kpc wide region around the galactic centre. The columns show, from left to right, the states after ~ 15 Myr for the runs using IC_4e8BH_r10 with $\epsilon_{\text{SFR}} = 2\%$ and $\epsilon_{\text{SFR}} = 0.2\%$, as well as those using IC_1.5e8BH_r100 with $\epsilon_{\text{SFR}} = 2\%$ and $\epsilon_{\text{SFR}} = 0.2\%$. The black dot in the centre marks the position of the supermassive black hole. At this early stage in their evolution, the CNDs are already in a state very similar to that at the end of the simulation time, if the star-formation efficiency is low (with the exception of the central void, which is not fully filled in yet). For the high star-formation efficiency runs, the early evolution is characterized by the CND starting to be torn apart by hot supernova bubbles.




Figure 3.17: Maps of the Toomre Q parameter of the gas (top row) and the line-of-sight velocity dispersion (bottom row) of the molecular gas in a 1 kpc thick, 1x1 kpc wide region around the galactic centre. The columns show, from left to right, the states at the end of the simulation time for the runs using IC_4e8BH_r100 with $\epsilon_{\text{SFR}} = 2\%$ and $\epsilon_{\text{SFR}} = 0.2\%$, as well as those using IC_1.5e8BH_r100 with $\epsilon_{\text{SFR}} = 2\%$ and $\epsilon_{\text{SFR}} = 0.2\%$. The black dot in the centre marks the position of the supermassive black hole. In the Toomre Q maps, white marks $Q = 1$, while blue is $Q < 1$ (unstable) and red $Q > 1$ (stable). Q is only defined in the maps within 400 pc from the centre (within the radius of the initial CND). The white region outside of this radius is undefined in the map and not $Q = 1$. For $\epsilon_{\text{SFR}} = 0.2\%$, the disc shows a spiral structure on the threshold of Toomre stability with a mostly very low velocity dispersion $\sigma_{\text{los}} < 6 \text{ km s}^{-1}$, while in the $\epsilon_{\text{SFR}} = 2\%$ runs, the final spiral structure is completely Toomre-stable (due to the lower surface density) with a higher velocity dispersion around $\sigma_{\text{los}} \sim 15 \text{ km s}^{-1}$. Supernova bubbles cause localised regions of very high velocity dispersion.

radius. The border to the white region is not a sharp cut at 400 pc, because we take into account the smoothing kernels of the gas particles when creating the maps. Due to this, some gas particles just inside the border provide some value beyond it. In reality, as Σ_{gas} decreases quickly beyond 400 pc, the gas would be stable at these radii.

In the $\epsilon_{\text{SFR}} = 2\%$ runs, the surface density of the gas disc has dropped so far that the entire structure is completely Toomre-stable ($Q \sim 10$) at the end of the simulation. The velocity dispersion of the gas has risen to $\sigma \sim 10\text{-}20 \text{ km s}^{-1}$ in much of the gas (much higher in areas directly affected by SN feedback). If the star-formation efficiency is reduced to 0.2%, the disc is dynamically colder, with velocity dispersions mostly around $\sigma \sim 3 \text{ km s}^{-1}$, and less stable, with $Q \sim 1\text{-}6$ in most of the spiral structure. Notably, the stability increases in the centre towards the SMBH, as is the case for the observed CNDs in ETGs (e.g. Boizelle et al., 2017; Davis et al., 2018). Both Q and σ are broadly consistent with the observed values for NGC 4429 in the $\epsilon_{\text{SFR}} = 0.2\%$ runs, indicating that preventing the collapse of gas to very high densities (in these simulations achieved artificially due to the resolution limits) results not only in a reduced SFR, but also a low velocity dispersion combined with marginal stability of the disc.

If we increase the resolution, as is the case with the star formation, the simulation results move away from the observations. In Fig. 3.18, we show the surface density and temperature maps at the end of the simulation time for the runs r100_s10_sf0.2_BH4, r20_s4_sf0.2_BH4, and r100_s1_sf0.2_BH4 to compare how changes in the resolution and gravitational softening length impact the final morphology of the gas. With a smaller softening length, the overall surface density of the gas decreases, it becomes more centrally concentrated and falls off more quickly towards larger radii. Outside of the very centre, the developing spiral pattern is much less tightly wound than in the r100_s10_sf0.2_BH4 run with $\epsilon_{\text{gas}} = 10 \text{ pc}$. Due to this, larger areas with warm ($T \sim 10^4 \text{ K}$) gas form in the runs with a smaller softening length.

In Fig. 3.19, we show the Toomre Q and velocity dispersion maps for the three runs with different softening lengths. As is the case for the runs with higher star-formation efficiency, the higher SFR in the runs with lower softening lengths lead to a more stable disc (even if not quite to the same degree). For both the r100_s1_sf0.2_BH4 and the r20_s4_sf0.2_BH4 run, a larger fraction of the gas has (very) high velocity dispersions ($\sigma > 30 \text{ km s}^{-1}$), raising the average σ of the CND from $\sigma \sim 4 \text{ km s}^{-1}$ (in the r100_s10_sf0.2_BH4 run) to $\sigma \sim 8 \text{ km s}^{-1}$ (in r100_s1_sf0.2_BH4) and $\sigma \sim 17 \text{ km s}^{-1}$ (in r20_s4_sf0.2_BH4), even though most of the area still has $\sigma \lesssim 6 \text{ km s}^{-1}$.

Of the other variations of our simulation that we compared in section 3.3, most change the morphology of the gas just as little as they do the evolution of the SFR. The exceptions are the runs with CR ionisation and those using the mechanical SN feedback implementation. We compare the runs r100_s1_sf0.2_BH4_CR – 14 with a reasonable CR ionisation rate, r100_s1_sf0.2_BH4_CR – 10 with the extremely high CR ionisation rate of $\zeta_{\text{CR}} = 10^{-10} \text{ s}^{-1}$, and r100_s1_sf0.2_BH4_mFB with mechanical SN feedback to the corresponding run with thermal SN feedback and no CR ionisation (r100_s1_sf0.2_BH4). Fig. 3.20 shows the density and temperature maps of these four runs after $\sim 15 \text{ Myr}$. In the run with $\zeta_{\text{CR}} = 10^{-10} \text{ s}^{-1}$ (third column), the CND stays completely smooth and warm

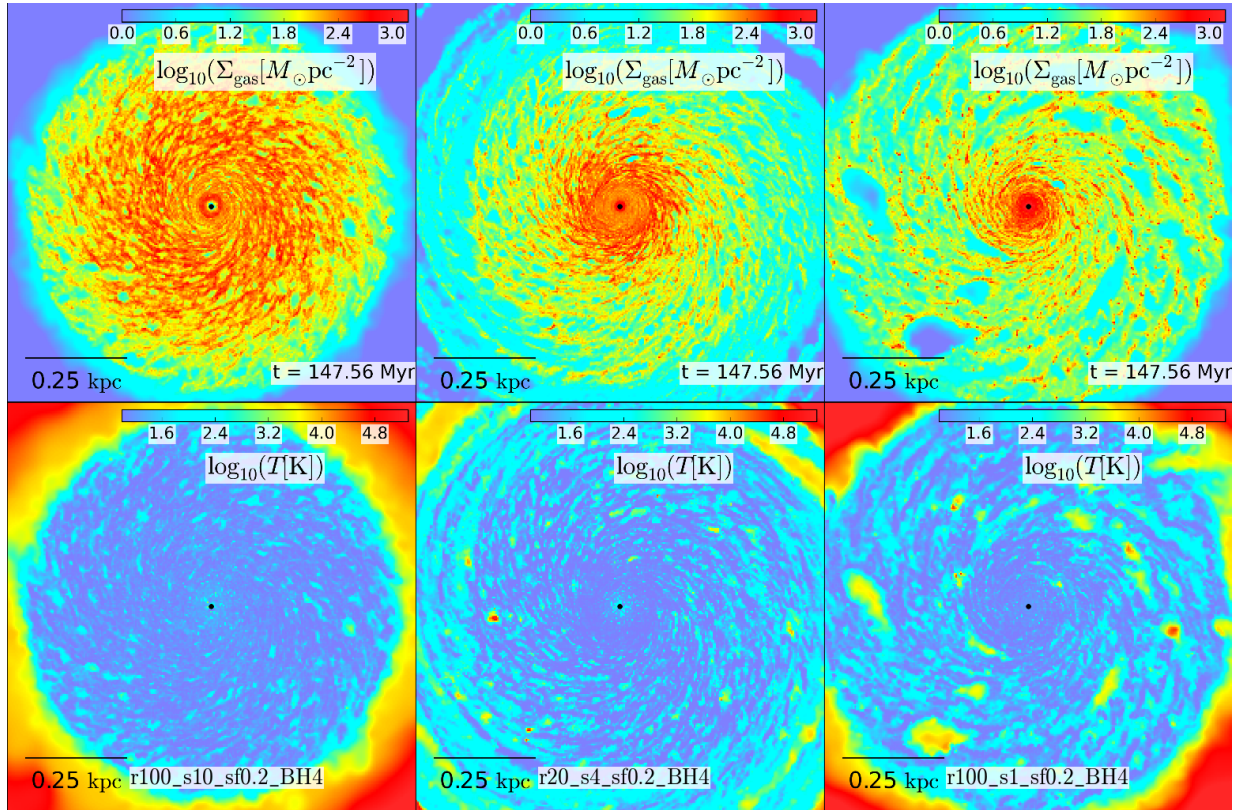


Figure 3.18: Maps of the gas surface density (top row) and the density-weighted average temperature (bottom row) in a 1 kpc thick, 1x1 kpc wide region around the galactic centre. The columns show, from left to right, the states at the end of the simulation time for the runs with $m_{\text{gas}} = 100 M_{\odot}$, $\epsilon_{\text{gas}} = 10 \text{ pc}$; with $m_{\text{gas}} = 10 M_{\odot}$, $\epsilon_{\text{gas}} = 4 \text{ pc}$; and with $m_{\text{gas}} = 100 M_{\odot}$, $\epsilon_{\text{gas}} = 1 \text{ pc}$. The black dot in the centre marks the position of the supermassive black hole. In the runs with shorter softening lengths, the final surface density of the CND declines much more steeply from the centre, as the higher SFR and SN feedback compress gas into the centre, while dispersing the disc further out, creating a looser spiral structure than in the $m_{\text{gas}} = 100 M_{\odot}$, $\epsilon_{\text{gas}} = 10 \text{ pc}$ run.




Figure 3.19: Maps of the Toomre Q parameter of the gas (top row) and the line-of-sight velocity dispersion (bottom row) of the molecular gas in a 1 kpc thick, 1x1 kpc wide region around the galactic centre. The columns show, from left to right, the states at the end of the simulation time for the runs with $m_{\text{gas}} = 100 M_{\odot}$, $\epsilon_{\text{gas}} = 10$ pc; with $m_{\text{gas}} = 10 M_{\odot}$, $\epsilon_{\text{gas}} = 4$ pc; and with $m_{\text{gas}} = 100 M_{\odot}$, $\epsilon_{\text{gas}} = 1$ pc. The black dot in the centre marks the position of the supermassive black hole. In the Toomre Q maps, white marks $Q = 1$, while blue is $Q < 1$ (unstable) and red $Q > 1$ (stable). Q is only defined in the maps within 400 pc from the centre (within the radius of the initial CND). The white region outside of this radius is undefined in the map and not $Q = 1$. The runs with smaller softening lengths have an overall higher Toomre Q . Similarly, the velocity dispersion is rising if ϵ_{gas} is decreased. Both of these are effects of the higher SFR in the runs with smaller softening lengths leading to a decrease in surface density and more SN-induced turbulence.




Figure 3.20: Maps of the gas surface density (top row) and the density-weighted average temperature (bottom row) in a 1 kpc thick, 1x1 kpc wide region around the galactic centre. The compared simulations all have $\epsilon_{\text{SFR}} = 0.2\%$, $\epsilon_{\text{gas}} = 1 \text{ pc}$, and the initial condition IC_4e8BH_r100. The columns show, from left to right, the states after $\sim 15 \text{ Myr}$ for the runs with thermal SN feedback and no CR ionisation; with thermal SN feedback and $\zeta_{\text{CR}} = 10^{-14} \text{ s}^{-1}$; with thermal SN feedback and $\zeta_{\text{CR}} = 10^{-10} \text{ s}^{-1}$; and with mechanical SN feedback and no CR ionisation. The black dot marks the position of the supermassive black hole. In the run with $\zeta_{\text{CR}} = 10^{-10} \text{ s}^{-1}$, the extremely high ionisation rate causes all of the gas to be warm and ionised, resulting in a completely smooth, stable disc. In all other runs, the disc fractures into cold, starforming clumps and/or spiral arms. Mechanical SN feedback makes the dense substructures of the CND smaller on average, resulting in more clumps and less large spiral arms. The cosmic ray ionisation (with a moderate $\zeta_{\text{CR}} = 10^{-14} \text{ s}^{-1}$) causes the cold, dense parts of the disc to be slightly warmer than in the runs without it.

($T \sim 10^4$ K) with essentially no substructure developing.

In the thermal SN feedback run with no (left column) or less extreme (second column) CR ionisation, the disc fractures into a spiral with SN-driven, large, hot, under-dense cavities between the dense spiral arms in the outer parts. The CR ionisation raises the temperature of the dense, cold gas (compare also the phase diagrams in Fig. 3.14, bottom row). Under the influence of the SN feedback, this later develops into the loose spiral around a central core that is shown in Fig. 3.18 (right column). In contrast, in the run with mechanical SN feedback (fourth column), the disc is not fracturing into spiral arms (outside of the central region), but instead into more homogeneously distributed dense clumps. This difference is caused by the change in the SN feedback implementation. It appears that the mechanical feedback stabilises the gas more against the formation of spiral arms, but does little to prevent gas on small spatial scales from collapsing to very high densities, resulting in a morphology of many small, dense clumps.

Maps of the Toomre Q and the line-of-sight velocity dispersion for the four just compared runs after ~ 15 Myr are shown in Fig. 3.21. In the run with an extremely high CR ionisation rate (third column), the velocity dispersion is very low at $\sigma \sim 2 \text{ km s}^{-1}$ (except for the very centre), while the temperature is much higher than in the other runs. Due to this, in contrast to all other runs, in this one the temperature of the CND is so high that it is stabilised thermally. Hence, in this run, we take into account the sound speed c_s instead of the velocity dispersion when calculating the Toomre Q , which results in $Q \sim 10$, comfortably in the stable regime. A more moderate CR ionisation rate (second column) has essentially no influence on either the stability or the velocity dispersion of the gas.

Exchanging the thermal SN feedback (run shown in the left column) for the mechanical one (right column), the unstable regions in the CND change from spiral arms to small, dense clumps, which follows from the altered surface density morphology. Furthermore, the mechanical SN feedback leads to the creation of many more regions with very high velocity dispersions of $\sigma > 35 \text{ km s}^{-1}$.

The very low observed velocity dispersion of the CND in NGC 4429 ($\sigma \sim 2.2 \text{ km s}^{-1}$) can only be matched by those of our simulations in which the star formation is sufficiently suppressed, which we only achieve in runs with artificial prevention of gas collapse through low resolution, or in the run with an extremely high CR ionisation rate, which does not leave the disc molecular. If the collapse of gas to very high densities is not prevented, the SFR becomes too high and the resulting SN feedback leads to too high average velocity dispersions. Furthermore, in the long term, the high SFR leaves the CND at a too low surface density compared to the observations. While the change from thermal to mechanical SN feedback has a significant impact on the morphology of the disc (indicating that the SNe are indeed unresolved), it does not prevent the formation of dense gas, and hence does not improve the agreement between the simulation and the observations.




Figure 3.21: Maps of the Toomre Q parameter of the gas (top row) and the line-of-sight velocity dispersion (bottom row) of the molecular gas in a 1 kpc thick, 1x1 kpc wide region around the galactic centre. The compared simulations all have $\epsilon_{\text{SFR}} = 0.2\%$, $\epsilon_{\text{gas}} = 1 \text{ pc}$, and the initial condition IC_4e8BH_r100. The columns show, from left to right, the states after $\sim 15 \text{ Myr}$ for the runs with thermal SN feedback and no CR ionisation; with thermal SN feedback and $\zeta_{\text{CR}} = 10^{-14} \text{ s}^{-1}$; with thermal SN feedback and $\zeta_{\text{CR}} = 10^{-10} \text{ s}^{-1}$; and with mechanical SN feedback and no CR ionisation. The black dot marks the position of the supermassive black hole. In the Toomre Q maps, white marks $Q = 1$, while blue is $Q < 1$ (unstable) and red $Q > 1$ (stable). Q is only defined in the maps within 400 pc from the centre (within the radius of the initial CND). The white region outside of this radius is undefined in the map and not $Q = 1$. With an extremely high ionisation rate of $\zeta_{\text{CR}} = 10^{-10} \text{ s}^{-1}$, the entire CND is stable while having a very low velocity dispersion throughout, while a more moderate rate of $\zeta_{\text{CR}} = 10^{-14} \text{ s}^{-1}$ has little influence on Q or $\sigma_{\text{H}_2, \text{los}}$. Note that for $\zeta_{\text{CR}} = 10^{-10} \text{ s}^{-1}$, we defined Q using c_s instead of σ , as the gas is stabilised by its high thermal energy in this run. Mechanical SN feedback causes instabilities in the disc to be more distributed into small clumps than into large spiral arms. It also causes more small regions of high velocity dispersion than thermal SN feedback.




Figure 3.22: Total gas mass fractions of molecular hydrogen (H_2 , left panel), atomic hydrogen (H , middle panel), and ionised hydrogen (H^+ , right panel) over time comparing runs with different star formation efficiencies and initial conditions (see legend). For $\epsilon_{\text{SFR}} = 0.2\%$, about 80% of the gas mass is in H_2 throughout all of the simulation time (except the very beginning, where it has not formed yet), with about 20% being atomic, and less than 1% ionised. For $\epsilon_{\text{SFR}} = 2\%$, however, the molecular fraction decreases significantly over time to about 10-15% by the end of the simulation. The atomic fraction rises to about 40%, and between 40% and 50% of the gas mass are ionised. Note that, while the total mass fraction of the ionised hydrogen is high, very little of this gas remains inside the central disc. If IC_1.5e8BH_r100 is used instead of IC_4e8BH_r100, less gas is ionised, and (for the high star-formation efficiency) more gas stays molecular, but the general trends are the same with both initial conditions.




Figure 3.23: Total gas mass fractions of molecular hydrogen (H_2 , left panel), atomic hydrogen (H , middle panel), and ionised hydrogen (H^+ , right panel) over time comparing runs with different mass resolutions and gravitational softening lengths at $\epsilon_{\text{SFR}} = 0.2\%$ (see legend). Shortening the softening length and/or reducing the gas particle mass leads to a reduction in the molecular mass fraction from about 80% to about 50%, with a corresponding increase in the ionised fraction from less than 1% to about 30%.

3.5 Molecular gas fractions

The final property in which we compare our simulations to the observed CNDS in ETGs (and particularly in NGC 4429), is their relative abundance of molecular hydrogen. The CND in NGC 4429 has a molecular gas mass of about $M_{\text{H}_2} \sim 10^8 M_{\odot}$, while its atomic hydrogen mass is measured to be below $M_{\text{H}} < 1.3 \times 10^7 M_{\odot}$, i.e. at most about 13% of the molecular mass. Most other observed CNDS in ETGs appear to have similar mass ratios between the hydrogen species, with at most 10-20% of the total gas mass being atomic within the region where H_2 is detected (Boizelle et al., 2017).

In Fig. 3.22 we show the time evolution of the total mass fractions of molecular (left panel), atomic (middle panel), and ionised (right panel) hydrogen for the four runs comparing the different initial conditions and star-formation efficiencies. As usual, the differences between runs with different initial conditions are small compared to the effect of the star-formation efficiency. In the runs with $\epsilon_{\text{SFR}} = 0.2\%$, the molecular mass fraction is roughly constant at about 80% for the whole simulation time, with the other 20% being almost entirely atomic, and less than 1% of the gas mass being ionised. In the runs with $\epsilon_{\text{SFR}} = 2\%$, on the other hand, the molecular mass fraction drops to only 10-15% by the end of the simulation, while the atomic fraction rises to about 40%, and a whole 40-50% of the gas are ionised. The high ionised gas fraction is somewhat misleading, as almost none of the ionised gas stays in the disc (compare Fig. 3.24). The central gas disc itself is dominated by the atomic hydrogen in these runs.

In Fig. 3.23, we compare the time evolution of the hydrogen mass fractions for the three runs with varying gravitational softening length and mass resolution. While the atomic hydrogen mass fraction at the end of the simulation is similar ($\sim 20\text{-}30\%$) between all three simulations, in the runs with smaller softening lengths, the molecular gas fraction is

reduced from 80% to about 50%, while the ionised fraction correspondingly rises from less than 1% to about 30%. To keep the CNM sufficiently dominated by molecular hydrogen, the SFR needs to be low, as the stellar feedback otherwise heats, ionises, and expels a significant part of the molecular gas.

To get a better grasp on the distribution of the various hydrogen species, we show maps of the H_2 (left column), H (central column), and H^+ (right column) surface densities at the end of the simulation time for the runs r100_s10_sf2_BH4 (top row), and r100_s10_sf0.2_BH4 (bottom row) in Fig. 3.24. Note that the colour scale for the H^+ map shows a smaller range, as ionised hydrogen is much rarer inside the disc than the other two species. While molecular hydrogen dominates in all parts of the CNM for the run with lower star-formation efficiency, in the run with higher star-formation efficiency, everywhere except the very centre is dominated by atomic hydrogen instead. With a higher star-formation efficiency, the amount of ionised gas inside the disc is also significantly higher than for the lower efficiency, though still more than an order of magnitude lower than that of atomic hydrogen in most parts of the CNM. As the total mass fraction of the ionised gas is about 50% in this run (see Fig. 3.22, right panel), this means that most of the ionised gas is pushed out of the inner region.

In Fig. 3.25, we show the hydrogen species mass surface density maps for the three runs with different gravitational softening lengths. A smaller softening length leads to less molecular gas (with the exception of the dense inner ~ 100 pc, and more ionised gas. While the difference is not as extreme as for the run with higher star-formation efficiency, it is caused by the same mechanism: a higher SFR leading to more ionising stellar feedback, reducing the dominance of molecular gas within the CNM.

Besides a changing SFR, two effects can cause the hydrogen species fractions to differ significantly in our simulations: the inclusion of cosmic ray ionisation, and the use of mechanical (instead of thermal) SN feedback. In Fig. 3.26, we compare the hydrogen species surface density maps after ~ 15 Myr simulation time of a run with a moderate CR ionisation rate (r100_s1_sf0.2_BH4_CR – 14, central row), and one with mechanical SN feedback (r100_s1_sf0.2_BH4_mFB, bottom row), to those of the corresponding run without CR ionisation and with thermal SN feedback (r100_s1_sf0.2_BH4, top row). Neither CR ionisation nor mechanical SN feedback change the amount of molecular or atomic hydrogen much within the first ~ 15 Myr (though the mechanical feedback alters the distribution of those species due to the changes to the general morphology that it effects), but both increase the ionised gas surface density in the disc significantly, in particular within the SN bubbles for the mechanical SN feedback.

In regards to matching the observations with our simulations, we again find that suppressing the star formation is crucial to keep the CNM dominated by molecular gas. As a complicating factor, some mechanisms that might help with that goal—such as cosmic ray ionisation or more effective SN feedback—also increase the fraction of ionised gas. To match the observations in star formation, surface density, velocity dispersion, and molecular gas fraction, there needs to be a mechanism that prevents the collapse of gas into dense, heavily star-forming clumps without heating it so much that the molecules in the CNM are destroyed.




Figure 3.24: Maps of the molecular (left column), atomic (middle column), and ionised (right column) hydrogen mass surface density in a 1 kpc thick, 1x1 kpc wide region around the galactic centre at the end of the simulation time for the runs with $\epsilon_{\text{gas}} = 10 \text{ pc}$ using IC_4e8BH_r100, and $\epsilon_{\text{SFR}} = 2\%$ (top row), as well as $\epsilon_{\text{SFR}} = 0.2\%$ (bottom row). The black dot in the centre marks the position of the supermassive black hole. Note the smaller range on the colour scale for the H^+ map. At low star-formation efficiency, the molecular gas dominates the entire CND, and there is very little ionised gas within the disc. At high star-formation efficiency, the gas is only predominantly molecular in the central $\sim 50 \text{ pc}$, while the rest of the disc is mostly atomic. This run also produces more ionised gas than the low star-formation efficiency one, but H^+ is still the least common form of hydrogen in the CND (while it is co-dominant with H in the total (not just the disc) mass, compare Fig. 3.22).




Figure 3.25: Maps of the molecular (left column), atomic (middle column), and ionised (right column) hydrogen mass surface density in a 1 kpc thick, 1x1 kpc wide region around the galactic centre at the end of the simulation time for the runs with $m_{\text{gas}} = 100 M_{\odot}$, $\epsilon_{\text{gas}} = 10$ pc (top row); $m_{\text{gas}} = 20 M_{\odot}$, $\epsilon_{\text{gas}} = 4$ pc (central row); and $m_{\text{gas}} = 100 M_{\odot}$, $\epsilon_{\text{gas}} = 1$ pc (bottom row). The black dot in the centre marks the position of the supermassive black hole. Note the smaller range on the colour scale for the H^+ map. In the runs with smaller gravitational softening lengths, the CNB contains more ionised, and less molecular hydrogen (except for the central region).

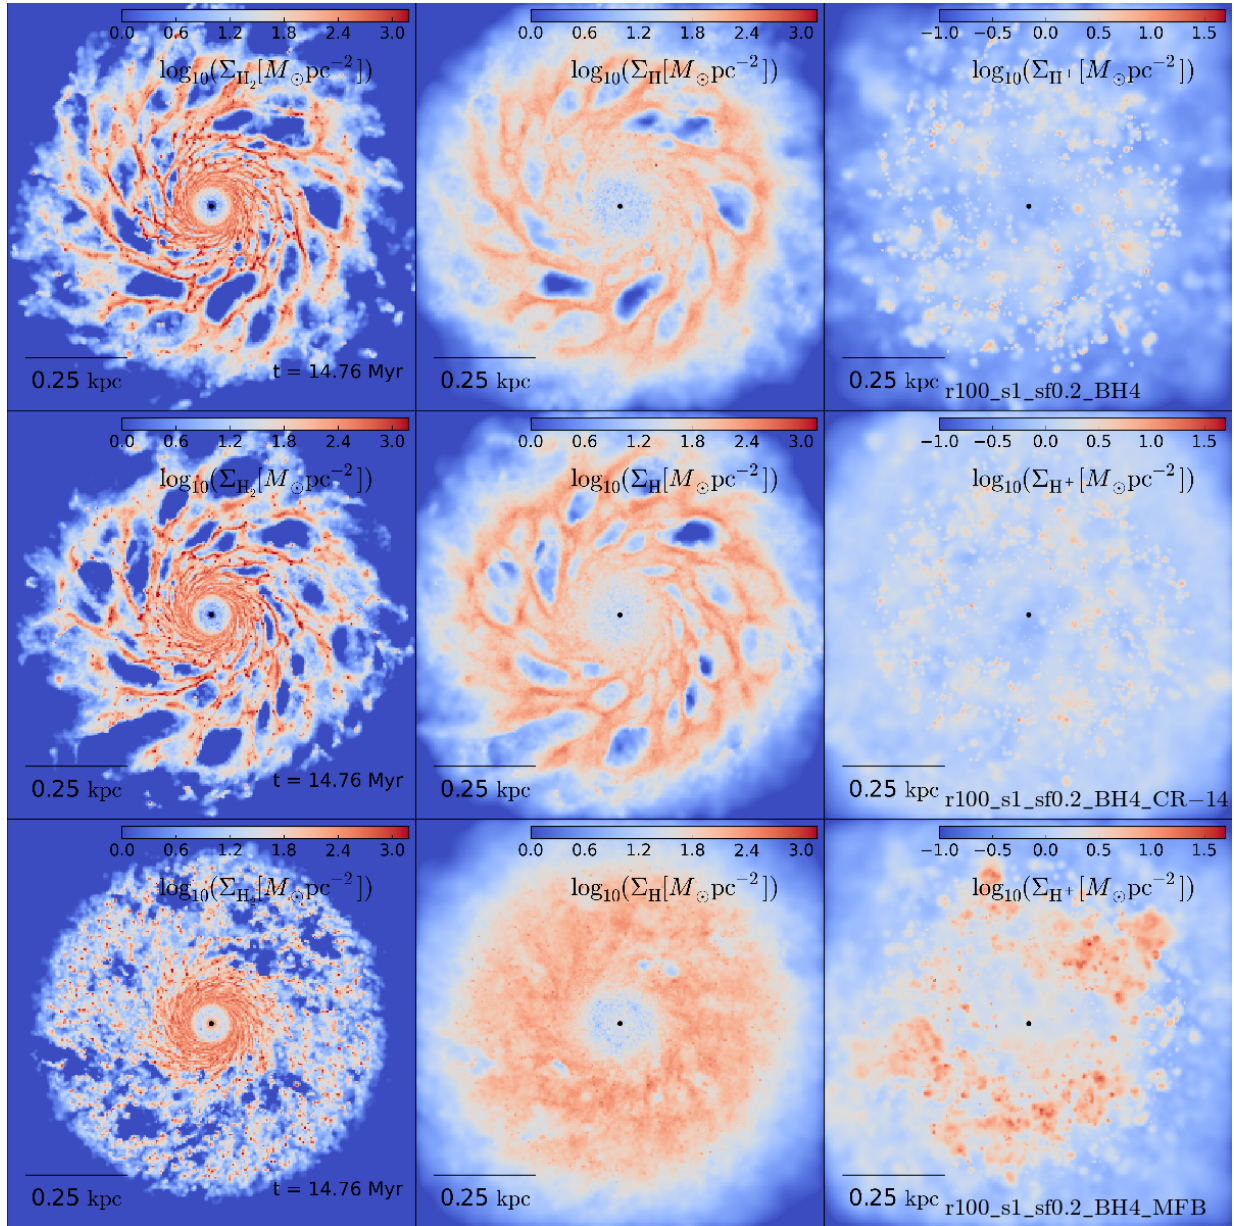


Figure 3.26: Maps of the molecular (left column), atomic (middle column), and ionised (right column) hydrogen mass surface density in a 1 kpc thick, 1x1 kpc wide region around the galactic centre after ~ 15 Myr. The compared simulations all have $\epsilon_{\text{SFR}} = 0.2\%$, $\epsilon_{\text{gas}} = 1$ pc, and the initial condition IC_4e8BH_r100. From top to bottom, the rows show runs with thermal SN feedback and $\zeta_{\text{CR}} = 0$; thermal SN feedback and $\zeta_{\text{CR}} = 10^{-14} \text{ s}^{-1}$; and mechanical SN feedback and $\zeta_{\text{CR}} = 0$. The black dot marks the position of the supermassive black hole. Note the smaller range on the colour scale for the H^+ map. Cosmic ray ionisation leads to more spread-out ionised gas. With mechanical instead of thermal SN feedback, more gas is ionised, especially in SN regions.

3.6 Black hole accretion and feedback

In all the simulations we analysed so far, we have treated the central SMBH merely as a passive concentration of mass, which only affects the gas by its gravitational potential. In reality, this black hole would accrete inflowing gas, turn into an AGN (though possibly of very low luminosity, depending on the accretion rate), and affect the surrounding gas via feedback. In fact, NGC 4429 is believed to host a low-luminosity AGN, as there is radio emission detected in its centre. Hence, in this section, we analyse simulations in which we allowed the central SMBH to accrete gas using the sink accretion model we introduced in section 3.2.1, as well as simulations in which we also turn on our AGN feedback model. We add the black hole accretion and feedback to the runs with the highest and the lowest SFR (excepting the run with an extremely high CR ionisation rate), i.e. the runs `r100_s10_sf2_BH4` and `r100_s10_sf0.2_BH4` (our “low resolution” runs with high and low star-formation efficiency). We call the runs with black hole accretion but without AGN feedback `r100_s10_sf2_BH4_acc` and `r100_s10_sf0.2_BH4_acc`, and those with both accretion and feedback from the SMBH `r100_s10_sf2_BH4_AGN` and `r100_s10_sf0.2_BH4_AGN`. In all four of these runs the SMBH is now also allowed to move if it is accelerated by the neighbouring gas particles, instead of having its position fixed to centre of the potential well.

In Fig. 3.27 we show the Eddington ratio distribution for all of our simulations that include black hole accretion at the end of the simulation time (after ~ 150 Myr), together with some observed values (sources given in the figure caption). With accretion but no AGN feedback, the SMBH accretes gas faster with higher star-formation efficiency than with lower efficiency, because the higher SFR in the former run leads to gas being driven into the centre by the more powerful supernova feedback. If AGN feedback is included, the accretion rate is strongly reduced and independent of the star-formation efficiency. The overall accretion rate is small in all simulations: It is very rarely (less than 0.03% of the time) higher than half, and for at most 10% of the time higher than 10% of the Eddington rate.

In most cases, the accretion rate of the SMBH is so low that its effect on the evolution of the CND as a whole is negligible, as we can see if we compare the time evolution of the SFR for the accretion runs to that of the non-accreting corresponding runs in Fig. 3.28. If the star-formation efficiency is low, neither SMBH accretion nor AGN feedback have any noticeable impact on the SFR, except for a slight increase, particularly in the first half of the simulation time. This increase is caused by the initial movement of the SMBH during which it drags gas behind it, creating connected ring structures around the centre. As the gas that is swept up in the wake of the black hole is colliding into these ring structures, its density is increased, resulting in the slightly enhanced SFR. Later on, the SMBH settles down in the centre of the potential well, moving very little from then on out. If the star-formation efficiency is high, there is almost no effect on the SFR by the black hole accretion or feedback during the early phase of high star formation rate. Only when the SFR begins to decline after about 40 Myr, the effect of the accretion and feedback becomes significant. As it evacuates the centre-most part of the disc, which at




Figure 3.27: Time distribution of different accretion rates in units of the Eddington limit (so-called “black-hole duty cycles”) for the different models with black hole accretion (see legend) compared to observations (square: Ho (2009), circles: Greene & Ho (2007), upward-pointing triangles: Kauffmann & Heckman (2009), downward-pointing triangles: Heckman et al. (2004), all taken from Novak et al. (2011)). The central SMBH accretes more gas and at higher rates in the runs with a higher star-formation efficiency, as the increased stellar feedback pushes gas into the centre. AGN feedback significantly reduces the accretion rate of the SMBH and negates the effect of the star formation on the accretion.




Figure 3.28: Star formation rate over time comparing the different runs that include black hole accretion to the corresponding ones that do not (see legend). In the runs with low star-formation efficiency, there is only a marginal increase of the SFR in the runs including black hole accretion or feedback, compared to the one with neither. In the runs with high star-formation efficiency, the runs with BH accretion or feedback show a ~ 40 - 60% lower late-time SFR than the one without.

these late times is the place of the densest, most star-forming gas, the SFR decreases when compared to the run without accretion.

In Fig. 3.29 we show the evolution of the SMBH mass relative to its initial mass $\Delta M_{\text{BH}}(t) = M_{\text{BH}}(t) - M_{\text{BH}}(0)$ over the same evolution for the stellar mass $\Delta M_*(t) = M_*(t) - M_*(0)$ for the runs with black hole accretion. In all cases, the SMBH grows by less than one percent over the ~ 150 Myr of the simulation. Its growth is completely negligible in the runs including AGN feedback. Without the feedback, it grows by about $\lesssim 1\%$ of its initial mass if the star-formation efficiency is high, and about four times less if it is low. While the effect of the SMBH on the overall stellar mass growth is small (see Fig. 3.28 above), the star formation (and its subsequent feedback) have a significant impact on the SMBH growth if AGN feedback is neglected: in the runs with more star formation, the SMBH accretes more gas. This is due to gas being driven towards the SMBH in the galactic centre by the stellar feedback, which is more effective if the SFR is higher.

We show the effect of the SMBH on the very centre of the CNDS in the surface density and temperature maps displayed in Fig. 3.30 for the pure SMBH accretion and the accretion plus AGN feedback run with high star-formation efficiency, as well as the corresponding run without black hole accretion. The black hole accretion strongly decreases the gas surface density within a radius of approximately 25 pc around it, causing the drop in the SFR as the central 25 pc contain the densest, most star-forming gas in the corresponding run without accretion. AGN feedback leads to the expulsion of the central gas and the creation of a cavity with a radius of about 130 pc. The corresponding maps for the low star-formation efficiency runs are not shown, but they produce similar black-hole-induced regions of reduced density in the inner ~ 130 pc. Even though the black hole accretion and feedback create central cavities, most of the CNDS stays unaffected by the black hole. In the low star-formation efficiency runs, where the dense, star-forming gas is quite evenly spread in the disc, the accretion and AGN feedback have therefore no significant effect on the SFR.

3.7 Discussion

In the preceding sections, we have analysed the evolution of our simulated CNDS in regards to its star formation, surface density, stability, velocity dispersion, phase structure, and chemical hydrogen composition, and compared its properties to those of the observed CNDS in ETGs, and specifically to that of NGC 4429, which we based our initial condition on. Altering the stellar bulge potential and SMBH mass in the initial condition, the star formation efficiency, the resolution and gravitational softening length, the amount of background UV radiation and cosmic ray ionisation, the supernova feedback implementation, and the presence or absence of black hole accretion and feedback, we ran a large set of simulations and found that none of them can reproduce all of the observed properties of NGC 4429's CNDS that we compare. Specifically, in almost all runs the mass fraction of molecular hydrogen in the disc becomes too low compared to the observations, as does the gas surface density, while the star formation rate surface density is higher than in the




Figure 3.29: Evolution of the SMBH mass relative to its initial mass over the corresponding evolution of the stellar mass for the different models with black hole accretion (see legend). The lines show the mass growth since the beginning until a given time for both the SMBH and the stars, starting (by definition) at zero. The squares and diamonds mark the total mass growth over the whole simulation time, i.e. the endpoints of the mass evolution. If AGN feedback is included, the black hole growth is insignificantly small. Without the feedback, the SMBH grows by $\sim 0.25\%$ ($\sim 10^6 M_\odot$) when the star-formation efficiency is low, and by $\lesssim 1\%$ ($\sim 3.9 M_\odot$) when it is high. More star formation leads to more SMBH growth.

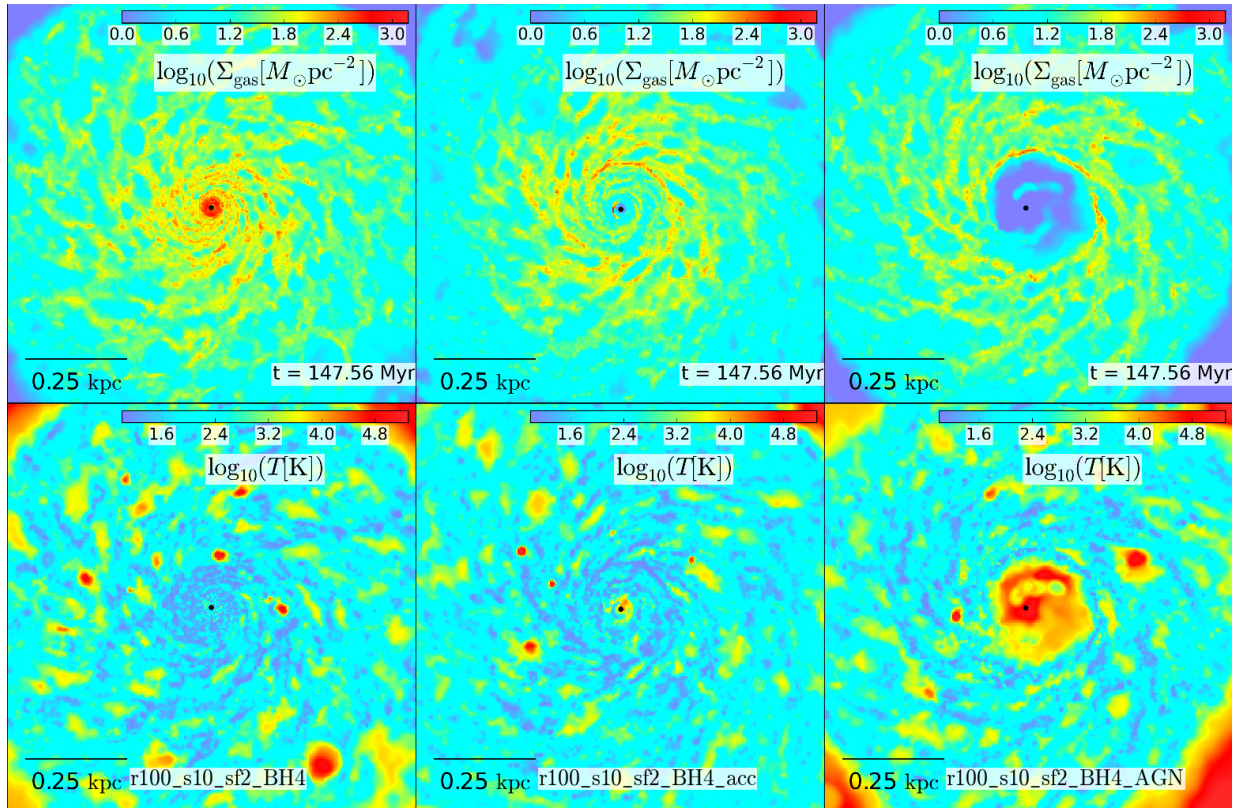


Figure 3.30: Maps of the gas surface density (top row) and the density-weighted average temperature (bottom row) in a 1 kpc thick, 1x1 kpc wide region around the galactic centre for otherwise equivalent simulations with and without BH accretion and AGN feedback. The columns show, from left to right, the states at the end of the simulation for the run without BH accretion or AGN feedback, the run with only BH accretion, and the run with both accretion and AGN feedback. The black dot marks the position of the supermassive black hole. The black hole accretion significantly reduces the surface density of the gas in the central ~ 100 pc around it, significantly loosening the inner spiral structure and removing the dense core in the central ~ 25 pc. Additionally including AGN feedback almost entirely evacuates the inner ~ 130 pc of dense and cold gas.

observed CNDs for equivalent gas surface densities.

Some of the alterations we tested—the initial SMBH mass and stellar potential, and the field strength of the UV radiation—have barely any influence on the properties of our simulated CND at all. The change in rotation speed and shear that is caused by the different initial conditions is not enough to significantly alter the stability of the CND, and in both cases the disc fractures with some gas collapsing to high densities. The UV radiation has very little impact on the evolution of the CND, because this evolution is mostly determined by the star formation and subsequent stellar feedback in the very dense gas. The photo-dissociation (heating) rate of the molecular gas by the UV radiation decreases with increasing H_2 column density N_{H_2} (it is approximately proportional to $N_{H_2}^{-0.75}$, see Draine & Bertoldi, 1996). Due to this, the UV radiation has no significant effect on the dense gas, where N_{H_2} is correspondingly high, and hence impacts the evolution of the CND very little.

Other changes have more of an impact: Cosmic ray ionisation can stabilise the entire CND and prevent any star formation in it if the assumed ionisation rate is high enough ($\zeta_{CR} = 10^{-10} \text{ s}^{-1}$, far above what is reasonable even in low-luminosity AGN), but in the process ionises the whole disc, which makes it incompatible with the observations. At more reasonable ionisation rates ($\zeta_{CR} = 10^{-14} \text{ s}^{-1}$), the CRs still heat the dense gas slightly, and ionise more of the CND (though most of the gas stays neutral), but they are incapable of preventing dense, star forming gas from forming and have very little impact on the SFR because of this.

Exchanging the thermal supernova implementation for our new mechanical one is similarly ineffective in suppressing the formation of dense clumps and stars. It does, however, lead to a more homogeneous CND morphology with small, clumpy substructures instead of the large spiral arms that develop with the thermal supernova feedback (compare Fig. 3.20). As such, the mechanical SN feedback produces a more similar morphology to that of the observed CND in NGC 4429 (see the $^{12}\text{CO}(3-2)$ intensity map in the left panel of Fig. 3.1) than the thermal SN feedback. However, it also increases the average velocity dispersion of the molecular gas to $\sigma \sim 9 \text{ km s}^{-1}$, removing it further from NGC 4429's observed $\sigma \sim 2.2 \text{ km s}^{-1}$.

Allowing the central SMBH to accrete gas and affect it via AGN feedback has a large impact on the central hundred or so parsec of the CND, as the black hole either swallows or ejects gas particles in this region, maintaining a cavity in the centre of the disc. SMBH accretion and feedback might therefore be the mechanism by which CNDs like that of NGC 4429 and similar ETGs develop and maintain their central cavities. Furthermore, the growth of the SMBH depends on the amount of star formation in the CND, as a higher SFR leads to more supernovae driving more gas towards the central black hole, increasing its available fuel for accretion. However, this effect vanishes if AGN feedback is included, as the AGN prevents the accumulation of gas in the very centre even if SN feedback is driving it towards it. Neither the accretion nor the AGN feedback have any significant impact on the evolution of the CND outside of this central region, as the black hole is unable to accrete the far away gas, and neither the wind feedback (which expels gas perpendicular to the disc plane) nor the radiative feedback are able to penetrate deeply into the disc.

We see by far the largest changes in the evolution of our simulated CND if we change the star formation efficiency or the resolution (or more importantly, the gravitational softening length). At low resolution ($m_{\text{gas}} = 100 M_{\odot}$) with a correspondingly large softening length of $\epsilon_{\text{gas}} = 10 \text{ pc}$, varying the star formation efficiency by a factor of 10 also changes the star formation rate by the same factor (at least in the beginning before stellar feedback and gas depletion become noticeable), as the density distribution of the star-forming, cold, dense gas is about the same in both cases. This is the case because the limited resolution of the gravitational force (set by the softening length) limits the maximum density the gas can collapse to. Due to this artificial stabilisation of the gas, reducing the star formation efficiency from $\epsilon_{\text{SFR}} = 2\%$ to $\epsilon_{\text{SFR}} = 0.2\%$ reduces the SFR by an order of magnitude and prevents the CND from large-scale fracturing. The surface density stays relatively high (though the low SFR still leads to a slow decrease in gas mass and, therefore, the surface density), and keeps its shallow exponential shape. The morphology of the disc is one of very densely wound spiral arms, and looks consistent with that of the CND in NGC 4429, as is the velocity dispersion of the molecular gas, which stays low ($\sigma \sim 3 \text{ km s}^{-1}$), as turbulence-inducing supernovae happen only at a low rate. Finally, the molecular hydrogen mass fraction of about 80% is also reasonably close to the $\gtrsim 90\%$ of the CND in NGC 4429. While the SFR surface density is still a bit high compared to the observations even with these conditions, it does agree with the observed value within the margin of error, and could likely be further reduced by assuming an even lower star-formation efficiency.

Both of the simulation runs with a low star-formation efficiency and a large softening length (r100_s10_sf0.2_BH4 and r100_s10_sf0.2_BH1.5) manage to reproduce the observed CND of NGC 4429 reasonably well. Unfortunately, the stability that prevents the formation of very dense, strongly star-forming gas is an artefact of the low resolution, as we see in our comparisons with a better resolved run (r20_s4_sf0.2_BH4), or—reducing the problem to its core—a run with the same mass resolution, but a lowered softening length (r100_s1_sf0.2_BH4).

Instead of this artificial resolution constraint, there needs to be a physical mechanism to keep the star formation in the CND suppressed to the levels observed in systems like NGC 4429 by preventing most of the gas from collapsing into very dense clumps, and/or by keeping very dense clumps from collapsing further into stars. This mechanism needs to be specific to the centres of massive ETGs, as CNDS in spiral galaxies form stars at a much higher rate. It also needs to suppress the star formation while simultaneously allowing for the formation and continued existence of large amounts of molecular gas. From our simulations, it appears that this mechanism can not be interstellar UV radiation, cosmic ray ionisation, supernova feedback, or accretion and feedback (in the form of collimated winds and X-ray heating) from a low-luminosity AGN (though all of these conclusions come with caveats, see section 3.7.1).

Morphological quenching—i.e. the stabilisation of the gas disc against fragmentation into dense clumps by the high and steep rotation curve created by the massive bulge of the ETG—is exactly the kind of non-thermal star-formation suppression mechanism that is thought to cause the low ratio of $\Sigma_{\text{SFR}} - \Sigma_{\text{gas}}$ in the observed CNDS in ETGs (Martig et al., 2009, 2013; Davis et al., 2014). The high-resolution AMR simulations of




Figure 3.31: Star formation rate surface density over gas mass surface density for the isolated simulations of Martig et al. (2013). Coloured symbols show the values in radial bins for the different simulations (circles/pentagons: high gas mass and overall surface density; triangles: low gas mass and surface density; blue: simulations in LTGs, red: in ETGs). The solid coloured lines are fits to the simulation values, the dashed red line is the fit for the spiral galaxies divided by 5. Finally, the shaded grey region is the 1σ area around the Kennicutt (1998) relation. For the low gas mass runs, the SFR is suppressed by about a factor of 5 in the ETGs compared to the LTGs, while for the high gas mass runs, there is only a very small reduction in Σ_{SFR} . This figure is taken from Martig et al. (2013).

Martig et al. (2013), which compare isolated simulations of gas discs in ETGs to those in corresponding LTGs, find that the star formation is less efficient by a factor of ~ 5 in their ETG simulations than in their LTG ones. This is caused by a significant reduction in the amount of very dense gas ($\rho_{\text{gas}} \gtrsim 10^{-20} \text{ g cm}^{-3}$) as the disc is more stable in the ETG simulations than in the LTG ones.

However, this is only true for those of their simulations with a relatively low gas surface density of $\Sigma_{\text{gas}} \lesssim 100 M_{\odot} \text{ pc}^{-2}$. For their simulations with a more massive gas disc of $\Sigma_{\text{gas}} \sim 100\text{--}1000 M_{\odot} \text{ pc}^{-2}$, the difference in the $\Sigma_{\text{SFR}}\text{--}\Sigma_{\text{gas}}$ relation between their ETG and LTG simulations decreases to a few 10% (see Fig. 3.31). This indicates that morphological quenching is effective as long as the gas surface densities in the disc are not too high, but becomes greatly less effective once they reach a few hundred $M_{\odot} \text{ pc}^{-2}$.

Our own simulations, in which we do not find significant morphological suppression of dense gas and star formation, start out with gas surface densities of $\sim 150\text{--}350 M_{\odot} \text{ pc}^{-2}$, closer to the denser of the simulation pairs of Martig et al. (2013). Possibly, this is the reason for the insufficient star formation suppression in our simulations via morphological quenching. Other processes that are not included in our simulations might be necessary to help stabilise the CND against the formation of dense, highly star-forming clumps.

One possible factor that we are neglecting in our simulations is the influence of magnetic fields on the gas evolution. Magnetic fields are a source of non-thermal pressure that can stabilise molecular clouds against further collapse to higher densities, high magnetic field strengths B are commonly associated with the centres of massive galaxies, and the formation efficiency of massive stars (which need particularly high density clouds to form) has recently been observed to be anticorrelated with B in the central gas ring of a galaxy that is in the process of being quenched (Tabatabaei et al., 2018). Magnetohydrodynamical simulations of CNDS in ETGs, as well as magnetic field measurements in these systems would be useful for assessing the role that magnetic fields may play in their own star formation suppression.

Another possibility is that we are missing pressure from turbulence on scales that are unresolved in our simulation. All of the star forming gas, and especially the very dense, highly star forming gas, has Jeans lengths far below our hydrodynamical resolution limit (which is given by the particle mass) in all of our simulations, and hence the density evolution at these scales can not necessarily be trusted. This, as well as other weaknesses in our modelling of the system could be the cause of the excess star formation in our simulations. We discuss these weaknesses and caveats in section 3.7.1.

3.7.1 Constraints & caveats

Our conclusions are drawn from a limited set of hydrodynamical simulations with a number of approximations and potential weaknesses in their modelling. Hence, as with any such simulation, there are caveats to take into account that put constraints on the extent to which our results can be trusted.

The most significant of these constraints is most definitely the limited resolution of our simulations. With a gas particle mass between 100 and 20 solar masses, our resolution is

high when compared to most other hydrodynamical simulations on galactic scales involving similar total amounts of gas, but nevertheless, most of the cold molecular gas that condenses in the CND reaches densities that are not properly resolved. Due to this, the results of the simulations are not converged, i.e. they change significantly when we increase the resolution.

One way in which the resolution heavily affects the simulated gas evolution is in its density distribution, as we discussed extensively over the course of this chapter. The gravitational softening stabilises the gas against further collapse at a resolution-dependent density. As no physical process stabilises the gas before this density is reached, higher resolution simulations have gas at higher densities than lower resolution ones. We see this effect in our comparisons between the runs with 100 and with 20 solar mass gas particles, as well as between runs in which we change the gravitational softening length by a factor of ten. Increasing the resolution and correspondingly decreasing the softening length even further may very well lead to even denser gas.

These high gas densities pose several problems for the simulation. Firstly, the denser the gas becomes, the less well it is resolved, and, due to this, the less reliable its structure and evolution become. Increasing the mass resolution should mitigate this effect, but if the corresponding lower softening length leads to higher maximum densities for the gas, the problem persists.

Secondly, the SFR of the gas depends directly on its density. If the density distribution is not reliable—be it due to artificial stabilisation through gravitational softening or due to the first problem just mentioned—the SFR is as well. In larger-scale galaxy evolution simulations, the star formation efficiency is often tuned such that the population of simulated galaxies fits the Kennicutt-Schmidt relation for whatever resolution is chosen for the simulation. We could certainly do the same in our simulations (that is, tune ϵ_{SFR} to the observed CNDs in ETGs), as we can see by the effect the changing of the efficiency has when we test it. However, for the purpose of understanding *why* the star formation in the CNDs is suppressed in ETGs, such a tuning would defeat the purpose.

Thirdly, and relating to the point above, supernovae exploding in dense gas need a high resolution (about $1 M_{\odot}$ according to Hu et al., 2016) for their blast wave to be properly resolved and not cool away their energy before enough of it is converted from thermal into kinetic form. As the supernova blast waves are not properly resolved, they might be less effective at inhibiting further star formation once a certain SFR (and thereby SN rate) is reached. We include test runs in which we use a mechanical supernova implementation instead, which assumes that the conversion of the initial thermal energy into kinetic energy with its associated radial momentum already happened below the resolution scale. This change prevents the formation of large spiral structures, but it does not suppress the formation of dense gas clumps and the subsequent high star formation rate.

Another set of limitations of our simulations lies in the initial conditions. For one, we only test a very limited range of parameters: two different SMBH masses and stellar bulge profiles (and hence only two different rotation curves), and just a single CND mass, extent and surface density. Changing these parameters over a larger range could reveal how dependent the results are on the specifics of the CND system, particularly in respect

to morphological quenching. We also limit our simulation to the gas within the central, dense, molecular disc (the region of NGC 4429 in which $^{12}\text{CO}(3-2)$ emission is detected by Davis et al. 2018), thereby neglecting the possible inflow of gas from larger radii.

Furthermore, the stellar (and dark matter, though on the scales of the CND, this is a minor contribution) potential is assumed to be completely spherical and static. While this approximation captures the most important effect of the old stellar population on the central gas disc (determining its rotation curve), and some other effects like SNIa feedback should be negligible on the simulation time scale of ~ 150 Myr, it is not impossible that the movement of the old stars and subsequent small deviations from spherical symmetry in their potential might affect the evolution of the gas disc. Effects like these are bound to be minor, however, and as the inclusion of proper star particles would greatly reduce the resolution we could computationally afford (due to the large stellar mass in the simulation region), neglecting them in the interest of increased gas resolution is almost certainly a good trade-off.

Finally, while the gas in our initial conditions is set to be in rotational equilibrium given by the gravitational potential, it has no initial turbulence ($\sigma = 0$). Hence, the CND initially has only thermal pressure, which quickly dwindles as the gas cools to support it against an initial collapse. Is the collapse to high densities and the subsequently high SFR therefore an artefact of the initial conditions? To test this, we performed a simulation in which—before we activated the star formation and stellar feedback models—we let the CND evolve with an artificial supernova driving to generate turbulence for about 15 Myr, after which the disc has an average velocity dispersion of $\sigma \sim 5 \text{ km s}^{-1}$. In Fig. 3.32, we compare the time evolution of the star formation rate of this run with that of the corresponding run without preceding turbulence driving. While the details of the evolution differ, the overall trend of an initially high SFR that later declines as the amount of gas in the CND drops is the same in both runs.

We also compare the morphology (i.e. the surface density and temperature maps) of the two runs in Fig. 3.33 after ~ 15 Myr of simulation with enabled star formation. Again, there are no significant differences between the two runs even as short as ~ 15 Myr after the start of the simulation. Both show the same fracturing into cold, dense spiral arms with supernova-driven cavities in between them. Hence, it appears that the lack of turbulence in the initial condition has no lasting effect on the simulation results.

There are also several potentially relevant physical processes that we are not including in our simulations. A major one is magnetic fields, which (as described above) might be another non-thermal source of pressure to stabilise the CND against collapse to high densities. Cosmic rays are only included as an ionisation and heating term in our sub-resolution cooling and chemistry model. Their proper, direct treatment would also require the use of full magnetohydrodynamics. Finally, while we do include radiative feedback from the SMBH in the form of Compton scattering of hard X-rays in some of the runs, this feedback model is an approximation that assumes that the gas is optically thin, a condition that is not necessarily met in the very dense molecular gas of our simulated CND. Including a full radiative transfer model to track the proper interaction of radiation with the gas would eliminate this modelling weakness, but is currently too computationally




Figure 3.32: Star formation rate over time comparing two otherwise identical simulations with and without preceding artificial supernova driving (see legend). While the exact evolution of the SFR naturally differs between the two runs, it is overall very similar in the maxima it reaches, as well as its long-term declining trend.




Figure 3.33: Maps of the gas surface density (top row) and the density-weighted average temperature (bottom row) in a 1 kpc thick, 1x1 kpc wide region around the galactic centre after ~ 15 Myr of simulation time with active star formation for two otherwise identical simulations with (left column) and without (right column) preceding artificial supernova driving. The black dot in the centre marks the position of the supermassive black hole. Even after a short simulation time, the morphology of the CND is very similar in the two runs.

expensive to include at these scales.

3.8 Summary & conclusions

We studied the evolution of a molecular, circumnuclear gas disc in the centre of a massive early-type galaxy using high-resolution numerical hydrodynamical simulations. We performed these simulations with SPHGal, an improved version of the SPH code GADGET-3, including a number of sub-resolution models to capture relevant physical processes in the gas disc. These include a Jeans-limited Schmidt-type star formation model, a stellar feedback model including photo-ionisation and supernovae (including metal enrichment), metal-dependent equilibrium cooling tables for the hot gas ($T > 10^4$ K), and a sophisticated non-equilibrium chemical network for the cold gas ($T < 10^4$ K) that tracks the abundances of molecular, atomic, and ionised hydrogen, as well as free electrons over a set of formation and destruction equations, as well as corresponding heating and cooling processes. With several alterations of this code, we performed simulations of an isolated, regularly rotating, dense, molecular gas disc in the central 400 parsec of an ETG modelled with a static gravitational potential of a massive stellar spheroid and dark matter halo. A supermassive black hole was placed into the very centre of the galaxy (and of the CND). The properties of our simulated CND were based on the observed properties of the CND in NGC 4429 (Davis et al., 2018).

We investigated the dependence of the star formation, morphology, stability, and chemical hydrogen abundances in the CND on several physical parameters and implementation methods in our simulations. In various combinations, we compared simulation runs

- with two different combinations of the stellar gravitational potential and the SMBH mass used in the initial condition, one based directly on the measured properties of NGC 4429, one on the massive ETG used for the simulations in chapter 2,
- with a star-formation efficiency of $\epsilon_{\text{SFR}} = 2\%$ or $\epsilon_{\text{SFR}} = 0.2\%$,
- with two different mass resolutions and gravitational softening lengths,
- with various field strengths of the interstellar UV radiation field (as part of the chemical network model),
- with several different cosmic ray ionisation rates (also as part of the chemical network),
- with a thermal and a mechanical implementation of the supernova feedback,
- and with and without a model for black hole accretion and AGN feedback,

to the properties of the CND in NGC 4429 and other observed CNDs in ETGs.

Our results can be summarised as follows:

- In all runs, the simulated CND lies significantly above and to the left of the observed $\Sigma_{\text{SFR}} - \Sigma_{\text{gas}}$ relation for ETGs containing cold neutral gas, and in particular the observed value for NGC 4429. The SFR surface density in most of the simulations is almost an order of magnitude above the mean of the observed systems at the same gas surface density. The excess star formation (compared to the observations) is a result of gas collapsing to very high densities ($\rho_{\text{gas}} \gtrsim 10^{-20} \text{ g cm}^{-3}$) at which the SFR is correspondingly high. The chosen initial condition, the UV field strength, and the inclusion or exclusion of SMBH accretion and feedback have no significant effect on the density distribution and SFR of the simulation.
- In our simulations, only two mechanisms are able to prevent the collapse of gas to high densities. The first is artificial stabilisation due to the gravitational resolution limit, which is given by the softening length. This leads to a lower maximum gas density and correspondingly lower SFR in simulations with worse resolution. At a given resolution where the gas density distribution is determined by the softening, the SFR is then proportional to the star-formation efficiency. The second is an extremely high cosmic ray ionisation rate of $\zeta_{\text{CR}} = 10^{-10} \text{ s}^{-1}$ which ionises the entire CND and heats it to $T \sim 10^4 \text{ K}$, thereby stabilising it against any collapse.
- The morphology of the simulated CND is primarily dependent on the initial star formation rate. If the initial SFR is low—which is only the case in the lower resolution simulations with a larger softening length (extreme CR ionisation leads to zero star formation)—the CND forms a very densely wound spiral, keeping its initial radial exponential surface density profile. If the SFR is high in the beginning of the simulation and thermal supernova feedback is used, the disc instead fractures into a loose spiral with arms of dense, star-forming gas, separated by hot cavities generated by successive supernovae. A dense core forms in the centre from gas pushed into it by the surrounding SN feedback. With the mechanical SN feedback implementation, the CND does not fracture into a loose, large-scale spiral, but instead into a roughly homogeneous collection of small, cold, dense clumps. This latter morphology is closer to that observed in the central molecular disc of NGC 4429.
- The simulated CND is initially (after cooling) right at the threshold of instability with a Toomre $Q \sim 1$. If the SFR is low, and the disc forms into a densely wound spiral, it stays around the stability threshold, with the spiral arms being unstable and the gaps in between them stable. The overall velocity dispersion of the gas disc stays low ($\sigma \sim 3 \text{ km s}^{-1}$), in accordance with the observed very low $\sigma \sim 2.2 \text{ km s}^{-1}$ of NGC 4429's CND. If the SFR is high, the depletion of gas by forming stars and ejection through supernovae over time reduces the surface density of the gas to the point where it is stable everywhere but in the densest regions of the remaining spiral arms. In this case, the average velocity dispersion is higher ($\sigma \sim 10 \text{ km s}^{-1}$), as the more common supernovae induce more turbulence into the gas.
- The mass fraction of molecular gas in the CND stays at about 80% (relatively close to

the observed $\sim 90\%$) over the course of the simulation if the SFR is low. Otherwise the disc stops being predominantly molecular in favour of a higher $\sim 40\text{-}50\%$ atomic gas fraction.

- Accretion onto and feedback from the SMBH cause the central ~ 100 pc of the CND to be evacuated of gas, but do not affect its structure or evolution outside of this region of influence. The accretion rate onto the black hole is generally low in the simulations, even without AGN feedback. A higher SFR leads to more black hole accretion as gas is driven towards the centre by supernova feedback.

We conclude from these results that the star formation suppression and low velocity dispersion that is observed in CNDs in the centres of ETGs such as NGC 4429 cannot be explained by morphological quenching, supernova feedback, an enhanced interstellar UV background, cosmic ray ionisation or accretion and feedback from the central SMBH. To suppress the star formation enough to match the observations, there needs to be a mechanism that prevents the cold molecular gas of the CND from collapsing to very high densities. This mechanism also has to leave the disc dominated by molecular gas and with a very low velocity dispersion. The most likely candidate for this that we included in our simulations is morphological quenching, which can increase the disc stability against collapse without increasing its velocity dispersion or ionisation rate significantly. This mechanism was found to successfully suppress the star formation in simulations by (Martig et al., 2013), as long as the gas surface density was below $\sim 200 M_{\odot} \text{pc}^{-2}$. We suggest that it fails in our simulations because the surface density is significantly higher than this threshold. As an alternative star formation suppression mechanism that is not included in our simulations, we propose non-thermal effects such as magnetic fields providing pressure support against the collapse of dense molecular clouds, which was observed to anti-correlate with the formation of massive stars in the centre of NGC 1097 by Tabatabaei et al. (2018).

Our simulations come with a number of weaknesses, which limit the robustness of their results. First and foremost among these is the limited resolution, as the majority of the cold molecular gas reaches densities for which its structure and hydrodynamic evolution is no longer properly resolved, which causes the simulation results to deviate significantly from each other when the resolution is changed. This problem is compounded by the unresolved nature of the supernova blast waves that explode within the dense gas. For the simulations with the thermal supernova feedback implementation, this limits their effect on the surrounding gas. Another potential weakness is the lack of turbulent support in the initial condition (where the velocity dispersion is zero), which might influence the initial evolution of the CND artificially (i.e. as an artefact of the initial condition). This we test by including a simulation run in which the star-forming evolution was preceded by the generation of turbulence through an artificial supernova driving phase. At the point at which the star formation is enabled, this run has a velocity dispersion of $\sigma \sim 5 \text{ km s}^{-1}$, but its further evolution is essentially identical to the corresponding run without preceding turbulence driving. This leads us to conclude that the collapse to high densities and subsequent high SFRs is not an artefact of the initial condition.

Higher resolution simulations (possibly including the treatment of magnetic fields and proper radiative transfer) are likely necessary to fully understand the peculiar properties (in particular their low star formation efficiency) of circumnuclear gas discs like that of NGC 4429. Furthermore, a larger range of initial conditions—both in regard to the gravitational potential provided by the SMBH and stellar spheroid, and to the surface density profile of the CND—should be tested to discover how the stability, density distribution, and star formation of the disc depend on these details of the system. The stabilising effect of morphological quenching in particular might significantly depend on these parameters.

Chapter 4

Pitfalls of AGN sub-resolution models in hydrodynamical simulations

We present hydrodynamical simulations of AGN feedback in a massive, quiescent early-type galaxy surrounded by a hot gaseous halo to investigate how changes to the model parameters and to the underlying hydrodynamical methods influence the effect of the feedback on the galactic evolution. Aside from the mentioned changes, the simulations are identical to those performed in chapter 2. First, we perform a parameter study changing the efficiency ϵ_w of the AGN wind feedback. We show that the ability of the AGN wind feedback to launch gaseous outflows from the galactic centre is strongly dependent on the chosen ϵ_w . Reducing the efficiency to $\epsilon_w = 0.25\%$ leads to a complete cessation of AGN-driven outflows, while increasing it to $\epsilon_w = 2\%$ increases the range of the outflows by about 50% compared to the fiducial $\epsilon_w = 0.5\%$. We then compare six otherwise identical simulations in which we test three different hydrodynamical solvers (pressure-entropy SPH, pressure-energy SPH, and MFM), and two different time-step-limiting schemes. The changes to the hydrodynamical model affect the simulation results more than changing ϵ_w . In agreement with earlier works (e.g. Durier & Dalla Vecchia, 2012), we find that the time-step-limiting scheme needs to take into account gas particles affected by feedback immediately to prevent the feedback to be overly effective at driving outflows. Comparing the three solvers, we show that using pressure-entropy SPH leads to very different results than the other two solvers, in that the AGN is much more effective at driving winds and evacuating the central region of gas. This appears to be caused by the energy-conservation issues of the pressure-entropy implementation of SPH (compare Hu et al., 2016). Pressure-energy SPH and MFM produce more similar, but still significantly different results, particularly in regards to gas outflows, black hole accretion and star formation. We conclude that the strong dependence of simulation results on the details of the hydrodynamical implementation needs to be taken into account when interpreting hydrodynamical simulations, especially if they include very energetic processes such as AGN feedback.

4.1 Introduction

The evolution of galaxies through cosmic time is determined by the complex interplay of dark matter, gas, dust, stars, black holes, and the various kinds of radiation produced by (some of) these components. While most of the mass is in the dark matter, and the large-scale distribution of matter is mostly governed by gravity, the hydrodynamic interactions between the baryonic components of the Universe (i.e. everything that is not dark matter or dark energy) have a strong impact on the evolution of structures on the scale of galaxy clusters and below (especially concerning the visible parts of these structures). Gravity and basic hydrodynamics are not enough to explain the observed properties of galaxy evolution, though. From the gas condensing and cooling in the centres of gravity wells (i.e. galaxies), stars and supermassive black holes are formed. Both of these can release large amounts of momentum and energy into their surrounding gas: stars, e.g., when they undergo supernovae; SMBHs through the enormous amounts of radiation that are generated in the process of gas being accreted onto them. This heats and expands the gas, creates winds, and thereby counteracts further star formation and black hole growth, as well as strongly influences the overall properties of the luminous matter.

The results of many numerical galaxy evolution simulations indicate that these feedback mechanisms, both from stars and SMBHs, are necessary to explain many observed properties of galaxies, from the shape of the galactic stellar mass function (e.g. Vogelsberger et al., 2014; Schaye et al., 2015; Somerville & Davé, 2015; Naab & Ostriker, 2017, the latter two being reviews containing many examples), over the distribution of hot gas in massive galaxies (e.g. Choi et al., 2015), to the continued quiescence of massive early-type galaxies (ETGs, see e.g. chapter 2). These numerical simulations, while they are our best way to model galactic evolution and to investigate which processes shape the properties that we observe in real galaxies, come with many difficulties. Feedback, especially from SMBHs (i.e. AGN feedback) can significantly affect the gas on scales up to Megaparsecs (e.g. in the case of relativistic jets), but it is generated on much smaller scales more on the order of tens to hundreds of AU (the scales of a SMBH accretion disc, or an exploding star). It is computationally infeasible to resolve these scales in a numerical simulation that encompasses a whole galaxy or more. Therefore, these highly energetic, influential processes need to be modelled in a sub-resolution fashion: the impact of the feedback on the smallest resolved scales (in the form of momentum and/or energy inputs) has to be estimated and plugged into a simulation that handles the gravitational and hydrodynamical forces on these larger scales.

The reliability of the simulations then depends on the robustness of these sub-resolution models. In the case of AGN feedback, there is no consensus about how it should be modelled, and a variety of different approaches are used, some of which are summarised below. Springel et al. (2005a) introduced a simple implementation in which the accreting black hole isotropically dumps a fraction of the accreted rest mass energy as heat into the surrounding gas. Booth & Schaye (2009) modified this model by storing the feedback energy of the SMBH until it is enough to heat the surrounding gas to a certain temperature, thereby reducing the energy loss from radiative cooling of gas that was heated to a temperature

at which its cooling time is short. Sijacki et al. (2007) implemented an AGN feedback model with two modes separated by accretion rate, meant to represent the observationally motivated “radio” and “quasar” modes of AGN. In this model, at high accretion rates, the “quasar” mode feedback heats the surrounding gas as in the Springel et al. (2005a) implementation, while at low accretion rates, large hot “bubbles” are created in the gas. Other models that clearly separate the two modes of AGN feedback include one by Steinborn et al. (2015), in which both modes work via the input of thermal energy, but with different efficiencies; another one by Dubois et al. (2012), where the low-accretion-rate “radio” mode is represented by the collimated, bipolar injection of momentum into the gas, representing a jet; and a recent one by Weinberger et al. (2017) with thermal energy input for the “quasar” mode, and pure momentum injection in random directions for the “radio” mode. Finally, there is the physically motivated model of combined winds (modelled through momentum input) and radiative heating (modelled through Compton scattering) that was developed by Ciotti & Ostriker (1997); Ciotti et al. (2009) and others (see citations in chapter 2) in one and two dimensions, and implemented in three dimensions by Choi et al. (2012). This is also the model use in chapter 2 and in this chapter.

Newly introduced AGN feedback models, where they are not motivated by observational and/or physical arguments, often refer to problems of previous approaches that the new model overcomes as their motivation (e.g. Booth & Schaye (2009) specifically modified the model of Springel et al. (2005a) to avoid the “overcooling” problem of the older model), but only a few studies have directly compared different AGN feedback implementations to each other regarding their effect on various galactic properties. Choi et al. (2014, 2015) compared their mechanical-radiative feedback model (Choi et al., 2012) to the purely thermal AGN model of Springel et al. (2005a) in various galaxy simulations. They found that their model performs better at matching the observed X-ray luminosity of the galactic gas, as it heats the gas far less than the Springel et al. (2005a) model, while still quenching star formation and generating outflows. In chapter 2, we compared the individual effects of the radiative and the mechanical parts of the model by Choi et al. (2012) on the evolution of an isolated massive ETG. We found that the radiative feedback is most efficient at preventing star formation, while the mechanical wind feedback is necessary to drive outflows. In a more idealised context, simulating a galactic gas disc without gravity, gas cooling, star formation, or stellar feedback, Cielo et al. (2018) compared the effects of a complex radiative AGN feedback mode (using a code with a radiative transfer scheme) with those of a mechanical collimated-jet feedback mode. While the former is more effective at increasing the temperature and decreasing the density of the gas, the latter produces stronger outflows, especially if the jets are parallel to the disc plane.

All of these studies are rather limited in that they compare only a small fraction of the proposed AGN feedback models, and only in a very limited number of astrophysical contexts. Larger comparison projects would be necessary to determine how robust simulation results are to the detailed implementation of AGN feedback. Furthermore, it is not only the sub-resolution models themselves that can vary and thereby change the outcome of a simulation. There are also various different formalisms to solve the basic hydrodynamical equations, which might influence how (even identical) feedback models affect the

galactic gas. The hydrodynamical solvers currently in use in the field can be grouped into the two “classic” formalisms of Eulerian adaptive mesh refinement codes (e.g. RAMSES, Teyssier, 2002) and Lagrangian smoothed particle hydrodynamics codes (e.g. GADGET, Springel, 2005), as well as the more modern approaches of moving-mesh methods (e.g. AREPO, Springel, 2010) and meshless Riemann solvers (meshless finite-mass and meshless finite-volume, e.g. Hopkins, 2015).

Comparison studies between these different solvers are important to determine if results from galaxy evolution simulations are robust, and how sensitive they are to their numerical underpinnings. Many studies like this have been done in different astrophysical contexts, though only a few with simulations including AGN feedback. One of the first comparison studies was the Santa Barbara Cluster Comparison Project (Frenk et al., 1999), which compared the results of simulations without radiative gas cooling or feedback, using 12 different codes (from SPH over fixed grids to AMR) to evolve a galaxy cluster from identical initial conditions. They found overall decent agreement between the simulations, with the largest discrepancies (of around a factor of two) in the X-ray luminosity of the gas. Similar non-radiative, no-feedback comparison studies of cluster simulations using SPH and AMR codes have been performed by Voit et al. (2005) and Mitchell et al. (2009), both of which find that AMR results in significantly higher entropy in the cluster cores than SPH.

Agertz et al. (2007) compared classic SPH and AMR implementations in idealized tests of interacting turbulent fluids. They showed that SPH creates spurious surface tension in regions with large density gradients, leading to the suppression of hydrodynamical instabilities, which are much better resolved in AMR. Since then, various people have improved the SPH formalism to mitigate (among others) the problem of spurious surface tension, creating so called “modern” SPH codes (e.g. SPHGal, Hu et al., 2014).

In the Aquila comparison project (Scannapieco et al., 2012), the formation and evolution of a disc galaxy was simulated from the same initial conditions with 13 different codes (including SPH, AMR, and moving-mesh methods), each using their own set of sub-resolution models for cooling, star formation, and (stellar and/or AGN) feedback. The project found large differences between the different codes for several fundamental properties of the resulting galaxies (like its mass and size), but attributed most of these to the differences in the feedback implementations, as opposed to those in the hydrodynamical solvers. As sub-resolution models and hydrodynamical solvers were altered at the same time, it is difficult to constrain the impact of solver changes on the results in the Aquila project.

The ongoing AGORA project (Kim et al., 2014) aims to eliminate that weakness by comparing simulations with different hydrodynamical solvers using not only identical initial conditions, but also common physical sub-resolution models, including models for stellar feedback, though no AGN feedback. Comparing isolated disc galaxy simulations (Kim et al., 2016), they find that the different codes produce results that agree generally well with each other. While they find some systematic differences between the different solvers, these are small compared to changes caused by varying sub-resolution models (e.g. for stellar feedback).

In his paper introducing the MFM/MFV code GIZMO, Hopkins (2015) compares these

meshless solvers to classic and modern SPH, AMR, and moving-mesh codes; both in numerical test cases and in isolated disc galaxy simulations (the latter only compares SPH and MFM/MFV) using identical sub-resolution models (though no AGN feedback). Interestingly, while MFM and MFV perform better than most of the other solvers in the numerical test cases, in the galaxy simulation including sub-resolution models for stellar feedback etc, MFM, MFV, and the two SPH formulations produce very similar results.

The moving-mesh code AREPO and the SPH code GADGET-3 have been compared to each other extensively in a variety of cosmological and isolated galaxy simulations using the same initial conditions and sub-resolution models, though generally none for AGN feedback (e.g. Vogelsberger et al., 2012; Sijacki et al., 2012; Kereš et al., 2012; Nelson et al., 2013). Hayward et al. (2014) added a comparison between the two codes for isolated disc and merger initial conditions, this time including a model for black-hole accretion and AGN feedback. All of these authors find significant differences in many galactic properties between the two codes, mostly owing to the problems present in the “classic” SPH formulation.

Schaller et al. (2015) compared results from the EAGLE cosmological simulation suite (Schaye et al., 2015) using the modern SPH code ANARCHY with those of reruns of these simulations using the classic SPH code GADGET-2 instead. They find that most galactic properties are not significantly affected by the change of SPH formulation. Only the efficiency of AGN feedback in massive galaxies changes due to the effects of spurious surface tension in the classic SPH code. The feedback efficiency (and the galactic properties it influences) is affected even more by using or not using a mechanism to limit the time-steps of gas particles affected by feedback (see Durier & Dalla Vecchia, 2012). All of these affect the results less than changes in the AGN feedback model itself, though.

Another study about the influence of different hydrodynamical solvers on cosmological simulations including AGN feedback was performed by Richardson et al. (2016). They compared the AMR code RAMSES with the classic SPH code HYDRA (Couchman et al., 1995), and found that the two solvers produce similar results for self-regulated properties such as the black hole growth or the star formation rate, while other properties (e.g. the entropy of the cluster core) show larger differences.

Finally, the nIFTy galaxy cluster simulations compared AMR, (classic and modern) SPH, and moving-mesh codes both in non-radiative simulations (Sembolini et al., 2016a), and in simulations including radiative cooling and (stellar and AGN) feedback models (Sembolini et al., 2016b). In the non-radiative case, they confirm the results of the earlier comparison studies. The mesh codes show cores of high entropy in the cluster centres, while the classic SPH codes produce clusters with low central entropy. For the runs with modern SPH, the central entropy profiles fall between the two extremes. If cooling and feedback processes are included, the clear differences between the different solvers’ results disappear, and are replaced by a large scatter of the properties that is driven by the differences in the (AGN) feedback models. As in the Aquila project, these varied from code to code, making it impossible to clearly constrain the influence of changes in the underlying hydrodynamics.

Overall, the results of these comparison studies are far from conclusive. There seems to be a general trend that changing the hydrodynamical solver of a simulation has significant

effects in non-radiative simulations, but that the inclusion of cooling and feedback results in less dramatic differences of the results between different solvers. Mostly, this appears to be an artefact of different feedback models creating a scatter in the results that makes any possible influence of the hydrodynamical solver undetectable. In studies that took care to keep the sub-resolution models as similar as possible between the different codes, the resulting differences caused by changing the hydrodynamical solver were sometimes small (e.g. Hopkins, 2015; Richardson et al., 2016), and sometimes larger (e.g. Hayward et al., 2014), apparently depending on which particular solvers were compared, and which sub-resolution models were included.

In this chapter, we compare simulations of an isolated ETG with a fixed initial condition and a fixed set of sub-resolution models including black hole accretion and AGN feedback, in which we alter the hydrodynamical solver between two different flavours of modern SPH (pressure-entropy and pressure-energy) and MFM. We also study the effects of changing the time-step limiting scheme on the simulation results. Additionally, as a point of comparison for the effects of changing the hydrodynamical formalism, we show the results of a parameter study altering the AGN wind feedback efficiency.

This chapter is organized as follows: In section 4.2, we give a short overview of the numerical code and the included sub-resolution models used in the simulations, as well as the different formulations of SPH and MFM, and the two time-step limiting methods that we compare. We also briefly present the initial conditions and the set of simulations shown in this chapter. Following this, section 4.3 contains the results of our AGN wind efficiency parameter study, as well as a short discussion of these, while we show and discuss the results of the runs with an altered hydrodynamical solver and/or time-step limiting scheme in section 4.4. Finally, we summarise our results and the conclusions to be drawn from them in section 4.5.

4.2 Simulations

4.2.1 Numerical code and sub-resolution models

As in chapter 2, we use the modern SPH code SPHGal (Hu et al., 2014) for the simulations presented in this chapter. SPHGal is a version of GADGET-3 (Springel, 2005) with improved fluid mixing and convergence properties.

Besides gravity and hydrodynamics, various important physical processes are included in the form of sub-resolution models. These are the same models we used in chapter 2: the gas cooling, star formation, metal enrichment and stellar feedback models by Scannapieco et al. (2005, 2006), Aumer et al. (2013) and Núñez et al. (2017); and the black hole accretion and feedback model by Choi et al. (2012). A more detailed description of these models can be found in section 2.2.1. We briefly summarise them in the following paragraphs.

Gas particles cool according to a pre-tabulated cooling rate that depends on their current temperature, density, and elemental abundances (which are tracked separately for Hydrogen, Helium, and the metals C, N, O, Ne, Mg, Si, S, Ca and Fe). They also change

their temperature adiabatically (i.e. due to changes in their density), and can be heated by shocks, as well as stellar and black-hole feedback.

Star particles (which represent whole stellar populations with a single formation time, as individual stars are not resolved) are formed from gas that has cooled below a threshold temperature of 12,000 K and collapsed above a density threshold of $1.94 \times 10^{-23} \text{ g cm}^{-3}$ at a rate inversely proportional to the dynamical time $\tau_{\text{dyn}} = (4\pi G\rho)^{-1/2}$ (where ρ is the gas density) and proportional to the star-formation efficiency, which is set to $\epsilon_{\text{SFR}} = 0.02$.

Star particles heat, accelerate, and enrich their surrounding gas with metals through stellar feedback, which (in the implemented model) comes in three forms: Supernovae type II are assumed to all happen approximately simultaneously for a given star particle once it reaches a metallicity-dependent age τ_{SNII} , at which point an appropriate amount of momentum, thermal energy, and mass in various elemental species is distributed to the neighbouring gas particles. Supernovae type Ia, which can occur over a much broader range of stellar ages, are implemented in a quasi-continuous way: All star particles undergo a SNIa event every 50 Myr. The distributed energy, momentum, and mass of each event is steadily declining with the star particles' age, following the delay time distribution presented by Maoz & Mannucci (2012). Finally, winds from asymptotic giant branch stars are implemented in the same quasi-continuous way as the SNIa feedback, but with smaller yields and a different elemental composition of the wind. The total momentum and energy injected by a stellar feedback event depends on the star particles mass, age (for SNIa), and metal composition, as given in precomputed tables from Woosley & Weaver (1995) for SNII and Iwamoto et al. (1999) for SNIa, and is scaled with the speed of the outflowing wind, which is set to 4000 km s^{-1} for the supernovae, and to 25 km s^{-1} for the AGB winds.

The accretion of gas by the black hole is calculated according to the Bondi-Hoyle-Lyttleton rate (Hoyle & Lyttleton, 1939; Bondi & Hoyle, 1944; Bondi, 1952):

$$\dot{M}_{\text{BHL}} = \left\langle \frac{4\pi G^2 M_{\text{BH}}^2 \rho}{(c_s^2 + v^2)^{3/2}} \right\rangle, \quad (4.1)$$

where M_{BH} is the mass of the black hole, ρ the surrounding gas density, c_s the sound speed in the gas, v the relative speed between the black hole and the gas, and the angle brackets stand for the SPH kernel averaging. The black-hole accretes individual gas particles stochastically with a probability proportional to the rate given above multiplied by additional factors to take into account the particle's free-fall time onto the black hole and the limited resolution of the gas (see Choi et al., 2012, for details). Additionally, the accretion rate is limited by the Eddington rate.

The feedback from the black hole comes in two parts: first, a mechanical wind feedback, in which a fraction of the gas particles that would be accreted onto the black hole are instead blown away at high velocities in parallel or anti-parallel direction to their axis of rotation around the black hole. Assuming a wind speed of $v_w = 10,000 \text{ km s}^{-1}$ (typical for observed broad-line-region winds, e.g. Crenshaw et al., 2003; Moe et al., 2009; Dunn et al., 2010), and given a feedback efficiency ϵ_w (i.e. how much of the accreted mass-energy is ultimately converted into an outflowing wind), it follows from basic energy conservation how much gas mass should be accreted and ejected, respectively, for a given inflowing gas mass. In

chapters 2 and 3, we chose $\epsilon_w = 0.005$, which corresponds to 90% of the inflowing gas mass being ejected in the wind. In this chapter we vary the wind efficiency between the values of 0.0025, 0.005, 0.01, and 0.02, corresponding to ejecta percentages of 82%, 90%, 95%, and 97%, respectively. Gas particles ejected by the black-hole wind feedback immediately share their momentum with their two nearest neighbours as in an inelastic collision, heating up in the process.

The second part is a radiative feedback, which isotropically heats and accelerates gas proportional to the black hole's accretion rate, and to a radiative efficiency set to $\epsilon_r = 0.1$. This radiative feedback mode represents X-ray heating through Compton scattering. The formulae from Sazonov et al. (2005) are used to calculate the heating rates of the gas particles.

4.2.2 Hydrodynamical solvers

As part of the tests shown in this chapter, we compare results from simulations using identical subresolution physics models (e.g. for star formation, black hole accretion, etc.), but three different methods for solving the basic hydrodynamical equations: two formulations of SPH: pressure-entropy SPH (which is what we used in chapter 2) and pressure-energy SPH (which we used in chapter 3), as well as a meshless-finite-mass solver. We briefly describe the differences between these three approaches in this section.

As stated in section 1.4.1 of the introduction, SPH is a Lagrangian method for numerically solving hydrodynamic equations. The fluid (i.e. the interstellar gas) is tracked by discrete particles, which move according to the pressure forces from the surrounding gas. The forces acting on a fluid particle, as well as its basic hydrodynamic variables (i.e. some combination of density, pressure, internal energy, and entropy) are determined by averaging over the forces and properties of neighbouring particles, each weighted by a kernel function, which takes into account the distance between the particles.

Different formulations of this method arise from choosing the primary variables averaged over the kernel and used in the equations of motion (see Hopkins, 2013, for a derivation and comparison of those formulations). In chapter 2, we used the pressure P and the entropic function A , leading to the equations of motion:

$$\frac{d\vec{v}_i}{dt} = - \sum_{j=1}^N m_j (A_i A_j)^{1/\gamma} \times \left[\frac{f_{ij} P_i}{P_i^{2/\gamma}} \nabla_i W_{ij}(h_i) + \frac{f_{ji} P_j}{P_j^{2/\gamma}} \nabla_i W_{ij}(h_j) \right], \quad (4.2)$$

$$f_{ij} = 1 - \left(\frac{h_i}{3A_j^{1/\gamma} m_j n_i} \frac{\partial P_i^{1/\gamma}}{\partial h_i} \right) \left[1 + \frac{h_i}{3n_i} \frac{\partial n_i}{\partial h_i} \right]^{-1}, \quad (4.3)$$

where \vec{v}_i is the velocity of particle i , N the number of its neighbours, m_j the neighbour particle's mass, A_i the entropic function, γ the polytropic index, $W_{ij} = W(\vec{x}_i - \vec{x}_j)$ the

smoothing kernel, h_i the smoothing length, and $n_i = \sum_{j=1}^N W_{ij}(h_i)$ is the kernel-based particle number density. Finally, P_i is the kernel-averaged pressure, defined as:

$$P_i = \left[\sum_{j=1}^N m_j A_j^{1/\gamma} W_{ij}(h_i) \right]^\gamma \quad (4.4)$$

This pressure-entropy formulation of SPH improves the mixing of the gas compared to the classical density-entropy formulation, because it does not lead to the creation of spurious surface tension along contact discontinuities (e.g. Hu et al., 2014).

The second version of SPH used in this chapter (which we also used in chapter 3) is the pressure-energy formulation, where we replace the entropic function by the specific internal energy u as a primary variable. With this choice, the equations of motion change to:

$$\begin{aligned} \frac{d\vec{v}_i}{dt} &= - \sum_{j=1}^N (\gamma - 1)^2 m_j u_i u_j \\ &\quad \times \left[\frac{f_{ij}}{P_i} \nabla_i W_{ij}(h_i) + \frac{f_{ji}}{P_j} \nabla_i W_{ij}(h_j) \right], \end{aligned} \quad (4.5)$$

$$f_{ij} = 1 - \left(\frac{h_i}{3(\gamma - 1)m_j u_j n_i} \frac{\partial P_i}{\partial h_i} \right) \left[1 + \frac{h_i}{3n_i} \frac{\partial n_i}{\partial h_i} \right]^{-1}, \quad (4.6)$$

with the kernel-averaged pressure P_i now defined via the specific internal energy u_i :

$$P_i = \sum_{j=1}^N (\gamma - 1) m_j u_j W_{ij}(h_i) \quad (4.7)$$

Compared to the pressure-entropy formulation, pressure-energy SPH is more accurate at conserving energy in shocks due to the former formulation having problems with erroneous density estimates when converting entropy into energy (Hu et al., 2016).

Besides the two SPH variants, we use a meshless-finite-mass implementation for our tests. MFM is part of a class of mesh-free Lagrangian methods for solving hydrodynamic equations that were first developed by Larson & Vila (2008a,b), first implemented in an astrophysical context by Gaburov & Nitadori (2011), and finally described and implemented in the form we use by Hopkins (2015).

Instead of using the smoothing kernel to determine the weight of each neighbour's forces onto a gas particle, but otherwise calculating these forces as if the particles were point masses, as is done in SPH, in the MFM method, the kernel is instead used to calculate to which degree each volume element of the gas volume belongs to each of the particles that represent it, i.e.

$$\psi_i(\vec{x}) = \frac{1}{\omega(\vec{x})} W(\vec{x} - \vec{x}_i, h(\vec{x})), \quad (4.8)$$

where $\psi_i(\vec{x})$ is the fraction of the volume element at position \vec{x} associated with particle i , W is the kernel weight, h the smoothing length of the kernel, and ω normalizes such that

the total volume fractions of all particles sum to one. Using these volume fractions, the Euler equations (describing the hydrodynamic evolution of the gas) can be transformed into

$$\frac{d}{dt}(V_i \vec{U}_i) + \sum_{j=1}^N \vec{F}_{ij} \cdot \vec{A}_{ij} = 0, \quad (4.9)$$

where $V_i = \int \psi_i(\vec{x}) d^3 \vec{x}$ is the effective volume of particle i , $\vec{U}_i = (\rho_i, \rho_i \vec{v}_i, \rho_i e_i)$ is the vector of its conserved quantities (mass density ρ_i , momentum density $\rho_i \vec{v}_i$, and energy density $\rho_i e_i$), \vec{A}_{ij} is an effective face area between the particle i and its neighbour j dependent on ψ_i and ψ_j , and \vec{F}_{ij} are the fluxes of the conserved quantities through this effective face. The sum is over all N neighbours of particle i . With this, the hydrodynamical equations can be reduced to a one-dimensional Riemann problem at the quadrature point on the effective face between the two neighbours (see e.g. Hopkins, 2015, for a more detailed derivation). A numerical Riemann solver is then used to determine the evolution of the particles.

According to tests by Hopkins (2015), MFM has several advantages over SPH: it can produce results of comparable accuracy with a much smaller number of neighbours used in the kernel function (and therefore at higher mass resolution at similar computational cost), and is better at resolving shocks and conserving angular momentum. It is, however, much more complex than SPH, and much less extensively tested in the context of galaxy evolution simulations.

4.2.3 Time-step-limiting methods

Besides the hydrodynamics solvers, we also switch the method for limiting the time-steps of neighbouring particles used in our simulation and test the effect of this change on the results. The purpose of time-step-limiting methods is to prevent fast gas particles (e.g. those accelerated by SN or AGN feedback) from moving through much of their surrounding gas particles without interacting with them. This problem arises from the use of individual time-steps for the particles. To save computational time, instead of using global time-steps of a fixed size everywhere, each SPH particle can have a different time-step, which is constrained by its acceleration and its thermal state (through the Courant criterion), i.e. the higher a particle's acceleration and/or sound speed, the shorter its time-step (for details, see Springel, 2005).

As each particle is only interacting with other particles when it is 'active', i.e. when its latest time-step has concluded and its properties are updated, a fast-moving particle on a short time-step can move through particles on longer time-steps before the latter become active again, thereby bypassing any interaction that should take place between the particles. This becomes especially relevant in the case of strong shocks moving through a medium, as the time-step size can vary drastically between the pre-shock medium and the shock material (see Saitoh & Makino, 2009). To combat this problem, two (quite similar) methods were implemented into GADGET-3/SPHGal, the 'wake-up' mechanism (WU), and the 'time-step limiter' (TSL), the latter of which we used in the simulations of chapters 2 and 3. We briefly describe both methods here.

Pakmor et al. (2012) implemented the wake-up scheme, following a proposed method by Saitoh & Makino (2009). The mechanism uses the signal velocities v_{ij}^{sig} between the particles to determine those particles whose time-steps have to be reduced. The signal velocity between two particles i and j is a measure of how quickly changes in the hydrodynamical state of the gas are transmitted between the pair, and is defined as (see Springel, 2005)

$$v_{ij}^{\text{sig}} = c_i + c_j - 3w_{ij}, \quad (4.10)$$

where c_i and c_j are the sound speeds of the respective particles, and w_{ij} is the relative velocity between the particles projected onto their separation vector (which is negative if the particles move towards each other, and set to 0 otherwise). The maximum of the signal velocities of a particle i to all its neighbours j , $v_{ij}^{\text{sig}, \text{max}} = \max_j(v_{ij}^{\text{sig}})$, is used in the Courant criterion to constrain the particle's time-step, and hence stored in the memory. Now, at each time-step, for all active particles, the stored $v_{ij}^{\text{sig}, \text{max}}$ of all neighbours of each active particle are compared to the current signal velocity between the active particle and its neighbour. If the stored $v_{ij}^{\text{sig}, \text{max}}$ of a neighbour is smaller than the current v_{ij}^{sig} between this neighbour and the active particle by a factor of 4.1 or less (an empirically constrained free parameter), the neighbour is 'woken up'.

'Waking up' (or 'activating') a particle means reducing the size of its time-step such that it becomes active at the next 'global' time-step (i.e. it becomes active together with the particles that currently have the shortest time-steps). This procedure ensures that particles react to changes in the hydrodynamical properties of their neighbours as soon as possible, if there is a large enough difference in their signal velocities (i.e. in their thermal states).

The second mechanism is the time-step limiter introduced by Durier & Dalla Vecchia (2012), which was implemented into SPHGal by Hu et al. (2014) and used in our previous simulations. It has two major parts. The first is similar to the WU mechanism described above. The time-step sizes of inactive neighbours of an active particle are reduced to that of the active particle, if their current time-step size is more than 4 times larger than the active particle's and their maximum signal velocity is 4 times smaller than the current signal velocity between them and the active particle. The factor of 4 in these conditions is again an empirically constrained free parameter. There are a few minor differences in this mechanism compared to the WU one (the time-step size is reduced to that of the active particle, not the smallest possible; not only the signal velocities, but also the time-step sizes themselves are compared, and the factor is 4 instead of 4.1), but the effect should be essentially the same: particles whose neighbours have much smaller time-steps than them have their time-steps adjusted such that they can interact with those neighbours before they have passed through.

The second part of the TSL mechanism ensures that gas particles that are affected by feedback (either from the AGN or from stars) immediately become active and update their time-step size according to their new properties (as they have been accelerated and heated). Without this immediate update, a particle with a long time-step that is affected by a feedback event could drift away from the environment in which the feedback took place

before interacting with its surroundings, leading to similar problems as those described above.

4.2.4 Initial conditions

We use the same initial conditions as in chapter 2, i.e. an isolated massive early-type galaxy containing a spheroid of old stars, a hot gaseous halo, a central supermassive black hole, and a dark matter halo. The stellar spheroid has a mass of $M_* = 8.41 \times 10^{10} M_\odot$ and is traced by $N_* = 841,000$ star particles with masses of $m_* = 10^5 M_\odot$ each. The stellar density follows a Hernquist profile (Hernquist, 1990) with a scale length of $a_* = 2.21$ kpc, so that the mass and size of the spheroid follow the relation by Williams et al. (2010). Each star particle is given an initial age (randomly chosen from a log-normal distribution around 6 Gyr, resulting in a very old stellar population typical of massive ETGs), and an initial metallicity (solar in the centre, exponentially dropping with increasing distance from it). This allows us to track metal enrichment from the old stars via AGB winds and SNIa explosions.

Following a relation by Kormendy & Ho (2013) between the stellar and the black hole mass, we set the mass of the central SMBH to $M_{\text{BH}} = 4 \times 10^8 M_\odot$. The virial dark matter mass is determined to be $M_{\text{DM, vir}} = 6.92 \times 10^{12} M_\odot$, using the abundance matching relation by Moster et al. (2013). The dark matter density follows a Hernquist profile with a scale length of $a_{\text{DM}} = 74.7$ kpc, which is traced by $N_{\text{DM}} = 10^6$ dark matter particles of $m_{\text{DM}} = 9.71 \times 10^6 M_\odot$.

Finally, the mass of the hot gaseous halo is set such that the total baryon fraction within the virial radius is 20% of the cosmological value, resulting in a virial gas mass of $M_{\text{gas, vir}} = 1.74 \times 10^{11} M_\odot$. $N_{\text{gas}} = 2.12 \times 10^6$ gas particles with $m_{\text{gas}} = 10^5 M_\odot$ each trace this hot gas component, following a β -profile (Cavaliere & Fusco-Femiano, 1976; Jones & Forman, 1984; Eke et al., 1998) with the slope parameter $\beta = 2/3$ and the core radius $r_c = 9.8$ kpc (following Jones & Forman, 1984; Makino et al., 1998, respectively). Given this density profile, the temperature of the gas then follows from the assumptions of hydrostatic equilibrium and isotropy. As with the stellar component, the gaseous halo is given an initial exponential metallicity profile. Both the dark matter and the gaseous halo are initially rotating with a spin parameter of $\lambda = 0.033$.

We discuss two sets of simulations in this chapter. The first is a parameter study comparing runs with varying AGN wind efficiency ϵ_w that are otherwise identical to each other. These runs are named according to the wind efficiency used. An overview of them is given in Table 4.1.

The second (and main) set of simulations consists of six runs with identical initial conditions and sub-resolution models (including their parameters), but different combinations of the hydrodynamical solvers and time-step-limiting schemes described in sections 4.2.2 and 4.2.3, respectively. They are named by stringing together abbreviations of their hydrodynamical solver (PS for pressure-entropy SPH, PE for pressure-energy SPH, MFM for meshless-finite-mass) and time-step-limiting schemes (TSL for the time-step limiter, WU for the wake-up method), as shown in Table 4.2.

4.3 AGN wind efficiency and CGM metal enrichment: a parameter study 143

Table 4.1: Summary of efficiency-test simulations showing the name of the run, and the AGN wind efficiency ϵ_w .

name of run	ϵ_w
eff0025	0.0025
eff005	0.005
eff01	0.01
eff02	0.02

Table 4.2: Summary of hydro-test simulations showing the name of the run, the hydrodynamical solver used, and the active time-step limiting scheme.

name of run	hydrodynamical solver	time-step-limiting scheme
PS_TSL	pressure-entropy SPH	time-step limiter
PS_WU	pressure-entropy SPH	wake-up
MFM_TSL	meshless-finite-mass	time-step limiter
MFM_WU	meshless-finite-mass	wake-up
PE_TSL	pressure-energy SPH	time-step limiter
PE_WU	pressure-energy SPH	wake-up

4.3 AGN wind efficiency and CGM metal enrichment: a parameter study

We begin with a numerical study that is typical for how sub-resolution models like our AGN feedback model are used to gain insight into the processes that shape galactic evolution. As a follow-up to chapter 2, we vary the efficiency of the AGN wind feedback ϵ_w to investigate how the outflow of gas and the enrichment of the circumgalactic medium with metals depend on this essentially free parameter in our AGN model, and if it could be potentially constrained by observational data of the CGM metal properties.

While it is clear from the start that the latter (constraining the efficiency through the CGM properties) would only be possible with significant caveats (most importantly the lack of an intergalactic medium and neighbouring galaxies in our study, as these might impact the CGM metal properties), the idea appears sound: We have a physically motivated AGN wind feedback model that produces reasonable results for the evolution of galaxies both in our own previous study and in larger, cosmological contexts (Choi et al., 2015). It has a strong influence on the CGM, but only within a limited distance from the galactic centre (~ 30 kpc), and it has an obvious uncertainty in the free parameter of ϵ_w . Hence, varying this parameter should give us valuable insight into the stability of the feedback effects and into the importance of how well the AGN’s radiation can couple to the surrounding gas to create a wind.

In Fig. 4.1, we show the metallicity profiles at the end of the simulation for the different runs with varying AGN wind efficiency, compared to the initial metallicity profile (black line) and some observed values from O’Sullivan et al. (2007). While changing the efficiency




Figure 4.1: Radial gas metallicity profiles at the end of the simulation, for runs with different AGN wind efficiencies, as well as in the initial condition (see legend). Also shown are observed values from O’Sullivan et al. (2007) for NGC 7796 (circles) and NGC 57 (squares). At an efficiency of $\epsilon_w = 0.0025$, the whole galaxy is deprived of metals, especially outside of the central ~ 5 kpc. For $\epsilon_w = 0.005$ and $\epsilon_w = 0.01$, gas out to about 30 kpc is enriched at more than 10% solar metallicity, while an efficiency of $\epsilon_w = 0.02$ is necessary to enrich the CGM out to even larger distances of about 50 kpc.

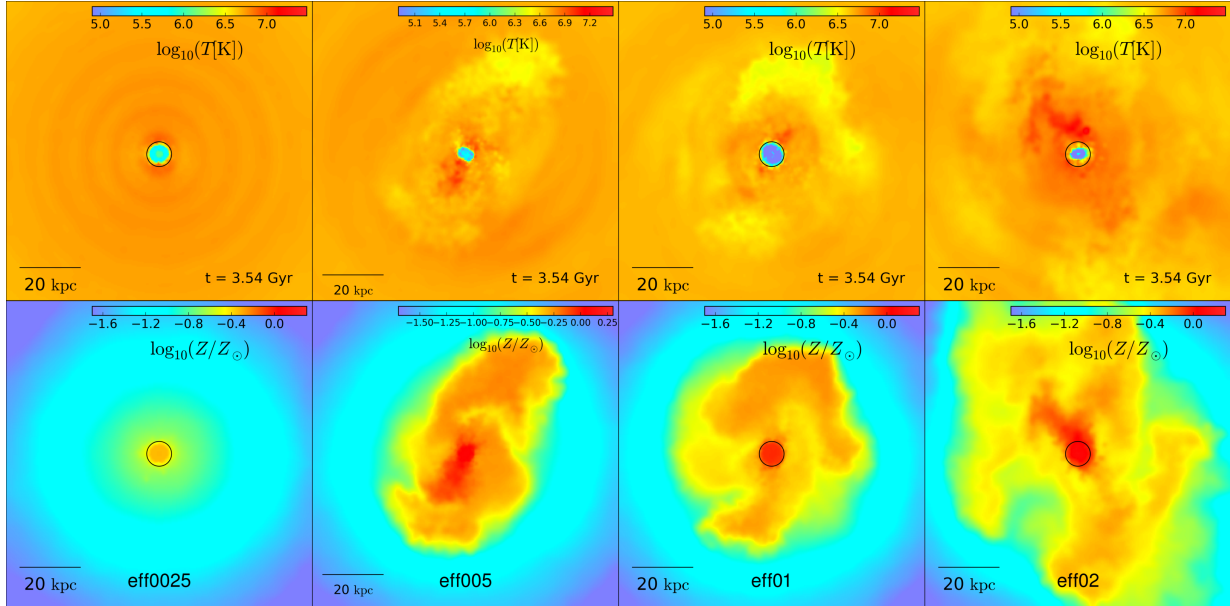


Figure 4.2: Maps of the density-weighted average temperature (top row) and metallicity (bottom row) of the gas in a 3 kpc thick, 100x100 kpc wide region around the galactic centre. The columns show, from left to right, the runs with $\epsilon_w = 0.0025$, $\epsilon_w = 0.005$, $\epsilon_w = 0.01$, and $\epsilon_w = 0.02$ after about 3.5 Gyr of simulation time. The black circle marks 1% of the virial radius. With $\epsilon_w = 0.0025$, the AGN feedback is incapable of affecting the gas much beyond $\sim 1\%$ of the virial radius. At higher efficiencies, hot, metal-rich outflows are produced. These have a similar reach of about 30 kpc for $\epsilon_w = 0.005$ and $\epsilon_w = 0.01$, but go out further to about 50 kpc for $\epsilon_w = 0.02$.

from 0.5% to 1% has little effect on the final metal distribution, halving the fiducial value of $\epsilon_w = 0.5\%$ to 0.25% makes the AGN incapable of driving metals out of the galaxy. On the other hand, if the wind efficiency is strongly increased to 2%, the AGN-driven outflows almost double their range, and enrich the CGM with metals out to at least 50 kpc.

This is further illustrated in Fig. 4.2, which shows the density-weighted temperature (top row) and metallicity (bottom row) in a 100x100 kpc wide, 3 kpc thick region around the galactic centre with the line of sight chosen perpendicular to the initial angular momentum of the gas, i.e. edge-on with respect to the forming cold gas disc. Hot, metal-rich outflows are visible in all runs with more than 0.25% wind efficiency, and they reach out to farther distances in the 2% efficiency run than in the other two.

Clearly, the efficiency of the AGN wind feedback has a strong influence on the outflows produced by the feedback. If it is below some threshold value between 0.25% and 0.5%, the AGN is incapable of creating any outflows, leading to a CGM that is just as deprived of metals as it is in runs without any AGN feedback at all. The effects of changing the efficiency are less straightforward at values above 0.5%: While doubling it to 1% does not change the properties of the outflow and the CGM significantly, increasing it by a factor of four to 2% allows the AGN-driven winds to reach much farther out into the CGM.




Figure 4.3: Total stellar mass (solid lines, left y-axis) and black hole mass (dashed lines, right y-axis) of the galaxy over time for the runs with different wind efficiencies (see legend). In the run with $\epsilon_w = 0.0025$, the black hole grows rapidly, while star-formation is suppressed. In comparison, at higher efficiencies, the black hole growth is strongly reduced, while more stars can form. The black hole growth declines monotonically with higher wind feedback efficiencies, while the stellar mass growth reaches a maximum around $\epsilon_w = 0.005$.

Varying ϵ_w also produces a similar pattern of effects on the stellar and black-hole mass evolution, which are shown in Fig. 4.3. Reducing the efficiency to 0.25% has a very strong effect: As the wind feedback is unable to blow any gas out of the central region, the black hole grows rapidly from the dense gas surrounding it. This leads to strong radiative AGN feedback (as we only changed the *wind* efficiency), keeping star-formation at a very low level, resulting in a net loss of stellar mass due to stellar feedback distributing some of the star particles' mass into the surrounding gas. At higher efficiencies, while the black hole growth is monotonically declining with increasing ϵ_w as more gas is blown out than accreted, the effect on the stellar mass evolution is less clear-cut: increasing the efficiency to 1% reduces the stellar mass growth significantly compared to $\epsilon_w = 0.005$, but increasing it even more to 2% does not lower it any further.

In conclusion, the efficiency of the AGN wind feedback model needs to be above some threshold value for the feedback to be effective at creating outflows. Increasing the wind efficiency beyond that has less drastic effects. The growth of both the black hole and the stellar mass can be reduced, but both are already very low at the fiducial $\epsilon_w = 0.005$. Furthermore, the AGN-powered outflows can enrich the CGM with metals to larger radii, but only if the wind efficiency is increased by a large enough factor. Even at $\epsilon_w = 0.02$, the outflows only reach out to about 50 kpc, while there are observations of metal absorption lines more than a hundred kiloparsec away from the centre of the galaxy (e.g. Werk et al., 2013).

4.4 Hydrodynamical variation tests

How reliable are the results presented in the previous section? To what degree can we trust these simulations to predict the properties of real AGN winds? To investigate these questions, we rerun the simulation in which the AGN feedback has the strongest impact on the ISM and CGM (i.e. the $\epsilon_w = 0.02$ run) with variations of the underlying hydrodynamics code, while keeping the sub-resolution physics models identical (i.e. we do not alter star formation, black hole accretion, AGN feedback, etc.). We then compare the results of these reruns with the original simulation and with each other to find out how much they depend on the underlying numerical methods as opposed to their direct, more controlled dependence on the sub-resolution models and their parameters.

For these tests, we alter the numerical implementation of the hydrodynamics in two distinct ways: Firstly, we alter the hydrodynamical solver that is used to calculate the basic hydrodynamical forces between the gas particles, as described in section 4.2.2. We compare the pressure-entropy (PS) SPH formalism used in the original simulation to a run using pressure-energy (PE) SPH, and one using a meshless-finite-mass (MFM) solver. Secondly, for each of the three different solvers, we use two different schemes to limit the size of time-steps of particles near shocks and feedback events, as described in section 4.2.3: the time-step limiter (TSL) that we used in the original simulations, and the wake-up scheme (WU). Together, this leaves us with six runs to compare (see Table 4.2 for an overview).

We start our comparison by looking at the radial metallicity profiles of the gas at the




Figure 4.4: Radial gas metallicity profiles at the end of the simulation, for the different runs, as well as in the initial condition (see legend). Also shown are observed values from O’Sullivan et al. (2007) for NGC 7796 (circles) and NGC 57 (squares). Changing the hydrodynamical solver and/or the time-step-limiting scheme has a strong effect on the profiles. If the TSL is kept, switching from PS-SPH to either PE-SPH or MFM greatly reduces the the reach of metal-enriching outflows, essentially shutting them off completely in the case of PE-SPH. If WU is used instead of TSL, the CGM is enriched to very large distances (>55 kpc) using PE-SPH or MFM. However, when using PS-SPH, it is enriched to higher values, but to lower distances (only ~ 30 kpc) with WU compared to the TSL run.

end of the simulation for all the runs, shown in Fig. 4.4. Comparing this figure with Fig. 4.1, it is immediately obvious that altering the underlying hydrodynamics scheme has as least as much influence on the resulting metal distribution as changing ϵ_w by a factor of 8, even though the wind feedback is the main driver of the metal distribution and the different hydrodynamical implementations should—in principle—be equivalent. It follows that the latter assumption is untrue, at least for simulations including very high energy feedback like the ones presented here. The various weaknesses and idiosyncrasies of the different methods, which might lead to relatively small differences in idealised test cases, affect the impact of the AGN feedback on the galactic gas distribution drastically. As we shall see, the simulations do not just lead to different results at a given time due to their very dynamic nature combined with small variations caused by the changed details of the implementation. Instead, the differences are qualitative in nature: the AGN accretion and feedback behave in fundamentally different ways in the different runs.

Looking at Fig. 4.4 in more detail, we can see two specific trends: If TSL is used, switching from pressure-entropy SPH to either of the two other solvers strongly reduces the metal-enrichment of the gas outside of the very centre, especially in the PE_TSL run, where the metallicity profile outside of the central ~ 5 kpc is almost identical to that of the eff0025 run in Fig. 4.1, which has essentially no outflows. Furthermore, using WU instead of TSL strongly increases the metallicity at large radii, except for the runs with pressure-entropy SPH, in which it decreases it outside of the central ~ 25 kpc.

To investigate why these differences occur, we plot the ratio of the inflow to the outflow rate of the gas through a shell of 20 kpc radius (OFR/IFR) over time for the various runs in Fig. 4.5. While in the two pressure-entropy SPH runs, OFR/IFR oscillates around 1, which indicates an overall balance of inflowing and outflowing gas (i.e. the existence of a “galactic fountain”), OFR/IFR is below 1 most of the time in the PE_TSL and MFM_TSL runs. Metal-rich gas from the centre can only be efficiently ejected to larger radii if a net-outflow exists, i.e. if $\text{OFR/IFR} > 1$. If pressure-entropy SPH is used, or—to a slightly lesser degree—if the WU switch is used instead of the TSL with the other two hydrodynamical solvers, this condition is often fulfilled, leading to the large-scale metal enrichment visible in Fig. 4.4. Note also the period just before the end, in which the PE_WU run shows an almost constant net-outflow, which coincides with high metallicity at very large radii at the end. In the MFM_TSL run, net-outflow is quite common in the first two thirds of the simulation, but absent after that, while in the PE_TSL run (with the lowest metallicity outside of the very centre), the net-flow is always directed inwards after the first half Gyr.

We visualise the behaviour of the AGN-driven outflows further in figures 4.6 and 4.7, which show maps of the gas temperature and metallicity in a 100×100 kpc wide region around the centre for the six different runs after ~ 1.5 Gyr (Fig. 4.6) and after ~ 4.5 Gyr (the end time of the simulations; Fig. 4.7). After ~ 1.5 Gyr, we can see how hot, metal-rich gas is flowing from the central region out to larger radii, where it then mixes into the CGM in all runs except PE_TSL, i.e. in all runs which have periods with significant $\text{OFR/IFR} > 1$. The outflows in the MFM-runs are significantly more collimated than those in the SPH-runs, creating very distinct mushroom-cloud shapes (compare e.g. MFM_TSL in the upper right of Fig. 4.6 with PS_WU in the lower left). Note that, while all runs




Figure 4.5: Ratio of outflow to inflow rate (OFR/IFR) through a shell at 20 kpc distance from the centre over time for the different models (see legend). Specifically, the running mean with a ~ 30 Myr (10 snapshot) step size is plotted. The horizontal black line shows a 1:1 ratio of OFR to IFR, i.e. zero net-flow of gas. In both the PS_TSL and PS_WU runs, OFR/IFR oscillates around 1, i.e. outflow and inflow are in approximate equilibrium if averaged over time. In the other four runs, OFR/IFR is generally negative, indicating net-inflow. The MFM_TSL run shows some significant net-outflows in the first ~ 3 Gyr, while in the PE_TSL run, the net-flow is always inwards with only one exception very early on (around ~ 0.5 Gyr). Using WU instead of TSL (which has little effect for PS-SPH) increases OFR/IFR in the PE-SPH and MFM runs, leading to more (and stronger) net-outflows, especially for PE_WU, and particularly in the last ~ 500 Myr.

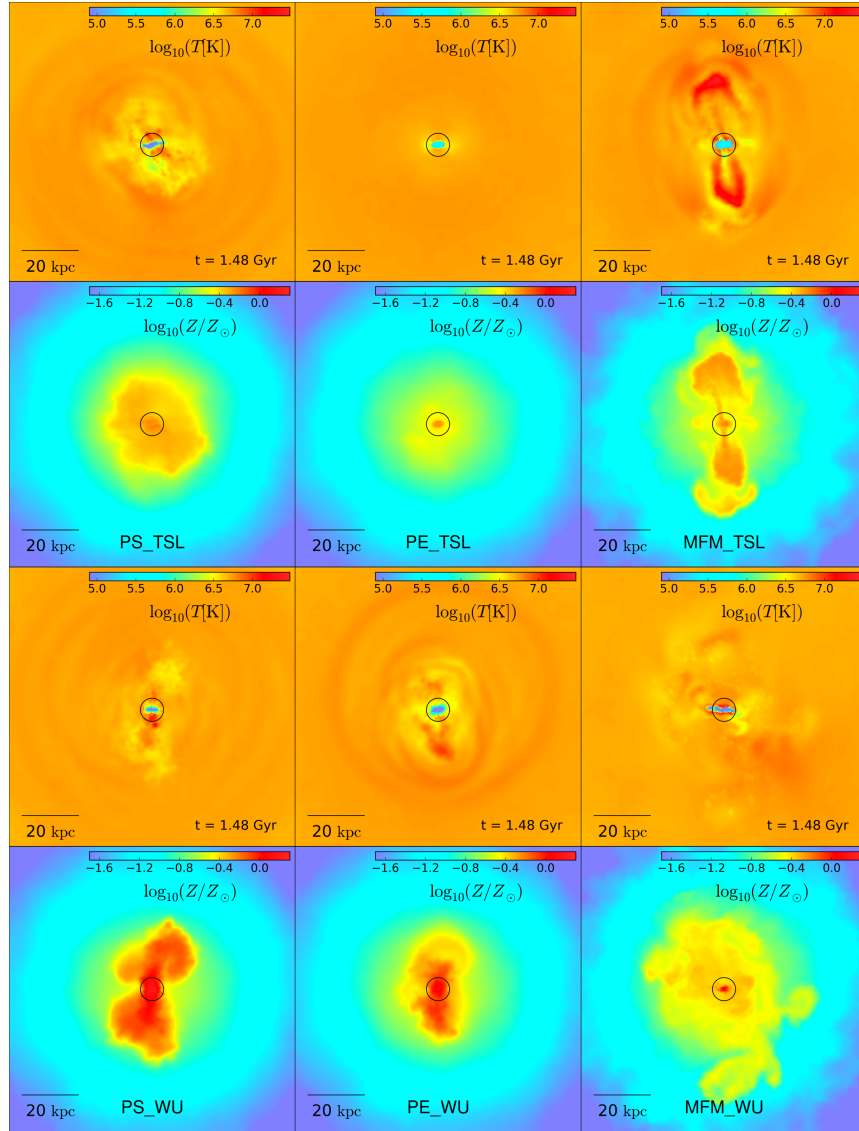


Figure 4.6: Maps of the density-weighted average temperature (first and third row) and metallicity (second and fourth row) of the gas in a 3 kpc thick, 100x100 kpc wide region around the galactic centre. The columns show, from left to right, the states after ~ 1.5 Gyr for PS-SPH, PE-SPH, and MFM, with runs using TSL in the top two rows, and with WU in the bottom two. The black circle marks 1% of the virial radius. Hot, metal-rich outflows (or their remnants) can be seen in all runs except PE_TSL, which shows a warm, metal-rich disc in the centre, surrounded by mostly metal-depleted gas.

except for PE_TSL have far-reaching metal-rich outflows until ~ 1.5 Gyr in the simulation, some of the runs (PE_TSL, MFM_WU) are in between outflow events at the plotted time. Only the mixing of the gas into the CGM is visible for these runs, not the outflow itself.

At the end of the simulation time (Fig. 4.7), the situation looks different. In the PE_TSL run, the lack of any AGN-driven outflows has led to a strong depletion of metals from the CGM gas, as more and more it has been accreted into the central region. The MFM_TSL run, which still showed strong outflows at the earlier time of Fig. 4.6, now also has a mostly metal-depleted CGM. The remnants of weaker, less far-reaching outflows are still visible in the metallicity map out to about 20 kpc, including a small outflow almost parallel to the central disc plane. In all other runs, the metal-enriching outflows have continued until the end of the simulation, further increasing the metallicity of the CGM, and (except in PS_WU) increasing the radius out to which metal-enrichment occurs.

There are strong differences in the shape of the metal-rich outflow though. In the PS_TSL and the PS_WU runs, the metal-rich hot gas is distributed more or less isotropically around the central region (with relatively small anisotropies introduced by more recent outflow events). In contrast, the metal-distribution in the PE_WU run is clearly bipolar, with two clouds of metal-rich gas above and below the central disc. The metal-rich gas is mixed into the CGM to some degree, but outside of about 10 kpc from the centre, very little metal-rich gas reaches the central plane. This bipolarity is even more pronounced in the MFM_WU run. While the metal-rich gas was still quite isotropically distributed after ~ 1.5 Gyr (Fig. 4.6), it is now almost exclusively within a very collimated, hot, bipolar outflow that is slightly inclined to the plane of the central cold gas disc. Most of the metal-rich gas only starts to fan out into a mushroom cloud at very large distances from the centre (>60 kpc, see the lower right corner of the lower right panel of Fig. 4.7).

Why does exchanging TSL for WU have such a large impact on the outflows in the PE-SPH and the MFM runs? In the former, the change replaces essentially no outflows with massive outflows out to more than 50 kpc from the centre, and in the latter, it causes a less extreme, but still very strong enhancement in the outflow rate and reach. The very likely reason for this effect is the missing immediate activation of gas particles receiving feedback energy and momentum in the WU scheme, which is present in the TSL scheme (compare the last paragraph of section 4.2.3). Without this mechanism, gas particles that are given energy and momentum from the AGN feedback (or the SN feedback, for that matter), can travel far away from the central regions (as they now have large velocities) before their current time-step ends and they start interacting with their neighbouring gas particles again. At these distances, the gas density is much lower, which reduces the decelerating pressure forces on the outflows, allowing them to reach much farther out into the CGM. As more gas can escape the central region, this also enhances the total outflow rates.

While this argument explains the differences between the TSL and WU runs in PE-SPH and MFM, and leads to the straightforward conclusion that TSL should be used instead of WU in simulations that include feedback mechanisms, it does nothing to explain why the change from TSL to WU has a completely different effect on the PS-SPH runs: PS_WU does not show farther reaching outflows than PS_TSL, but the opposite, and the net-outflow rate is about the same in both runs. Of course, the PS_TSL also differs strongly

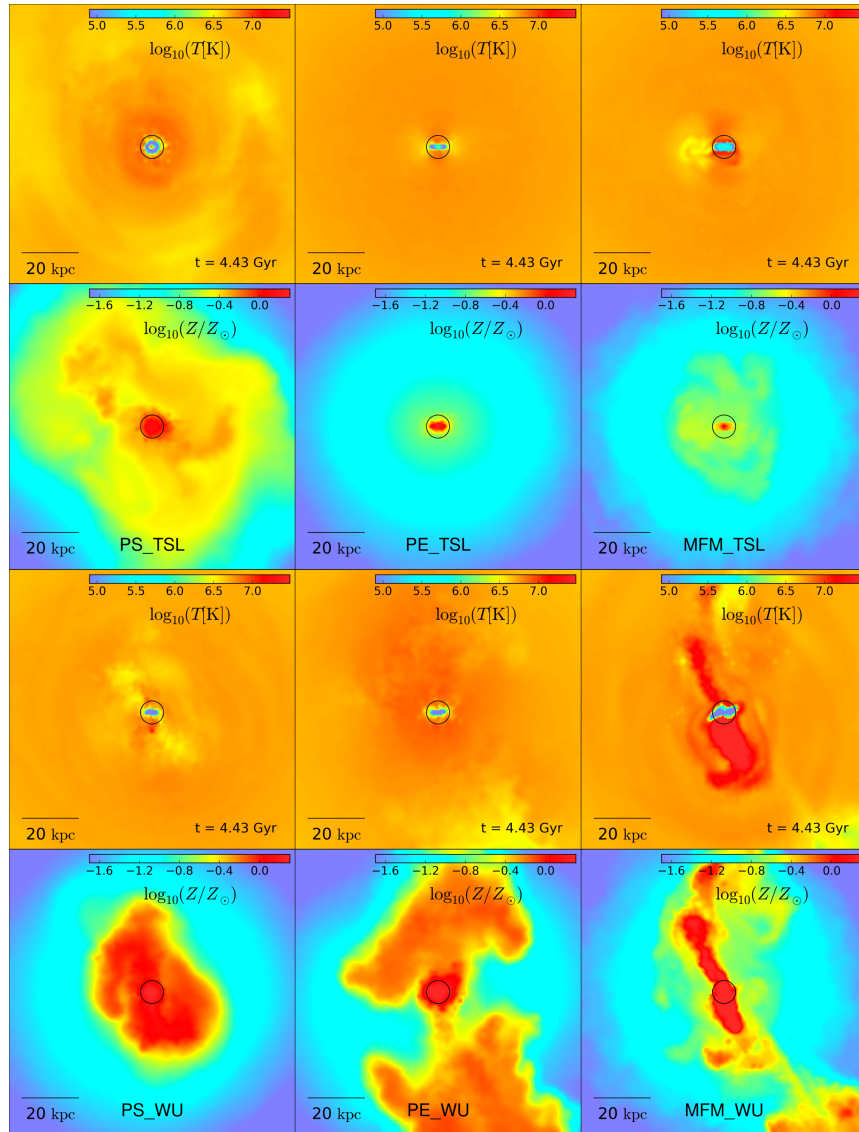


Figure 4.7: Maps of the density-weighted average temperature (first and third row) and metallicity (second and fourth row) of the gas in a 3 kpc thick, 100x100 kpc wide region around the galactic centre. The columns show, from left to right, the states at the end of the simulation for PS-SPH, PE-SPH, and MFM, with runs using TSL in the top two rows, and with WU in the bottom two. The black circle marks 1% of the virial radius. At the end of the simulation, the CGM is filled with metal-enriched gas out to various radii in all runs except PE_TSL and MFM_TSL. In the PE_WU and MFM_WU runs, the hot, metal-rich outflows are further reaching and more collimated than in the runs with PS-SPH. Furthermore, the metals are much less mixed in the MFM_WU run than in the other runs showing large-scale metal-enrichment.

from the two other TSL runs, as it shows consistent, strong, far-reaching outflows capable of enriching the CGM out to ~ 50 kpc.

To explain these discrepancies, we focus on the region from which the outflows originate in more detail. In Fig. 4.8, we show the gas surface densities (first and third row), and temperatures (second and fourth row) in the central 10×10 kpc for the six different simulations at the end of the simulation time. In all runs, a dense central gas disc has formed that is colder than the gas remaining in the hot halo. However, the properties of these central gas discs differ strongly between the different implementations. In the PE_TSL and MFM_TSL runs, the central disc is massive, dense, and mostly heated to a few 10^5 - 10^6 K by radiative AGN feedback. It is densest in the very centre. Switching TSL for WU, the PE_WU and MFM_WU runs now show sparser, more structured discs, which are overwhelmingly cold (a few 10^3 K). In the case of MFM_WU (which shows the strongest, most collimated outflows), a hole is blown into the innermost kiloparsec of the disc. This supports our explanation from above. In the WU cases, the gas that is affected by the feedback escapes the central disc before interacting with it, leaving it less dense and cold. In the TSL cases, where the feedback-affected gas interacts with its surroundings immediately, much less gas is able to escape, leading to a much denser disc, which is kept at high temperatures by the feedback energy.

While the PE-SPH and the MFM runs look similar to each other (though the disc is significantly larger in the MFM runs), the PS-SPH runs show a very different central region. Both in PS_TSL and in PS_WU, the central disc is more of a cold, dense ring around a central hole. The PS_TSL run behaves essentially like the runs described in chapter 2. The AGN wind feedback blows away all of the gas entering the very centre, creating the large-scale metal-rich outflows, while the AGN radiation feedback heats up and expands the gas in the surrounding disc, which then quickly cools down and contracts again. The ring is cold at the time shown in Fig. 4.8 because the last AGN feedback event happened a few 100 Myr earlier (compare Fig. 4.11), as no gas has fallen into the central accretion region around the SMBH since then, but generalising over the whole simulation time, the behaviour is as described above. In the PS_WU run, the wind feedback still blows a hole into the centre, as in PS_TSL, but the radiative feedback fails to heat up the surrounding disc, which stays cold for the entire time, as it does in the other two WU runs.

The wind feedback always affects gas particles close to the central black hole. Hence, the central hole that is created by the wind feedback in the PS-SPH runs is the reason for the much smaller impact on the outflow-properties from switching TSL for WU, compared to the other runs. As the immediate environment of the black hole is generally hot and of low density, the gas particles of the wind feedback encounter little resistance leaving the central region, even if they are activated immediately. In stark contrast, the neighbourhood of the central black hole is the densest part of the gas distribution in the PE_TSL and MFM_TSL runs, making it much harder for wind particles to escape.

Why is the wind feedback able to create a hole around the SMBH in the PS_TSL run, but not in the PS_TSL and MFM_TSL runs? The answer to this seems to lie in the problems of PS-SPH with energy conservation at shocks. Our wind feedback creates a blast wave by injecting a large amount of momentum and energy into a few spatially concentrated




Figure 4.8: Maps of the gas surface density (first and third row) and the density-weighted average temperature (second and fourth row) in a 10 kpc thick, 10x10 kpc wide region around the galactic centre. The columns show, from left to right, the states at the end of the simulation for PS-SPH, PE-SPH, and MFM, with runs using TSL in the top two rows, and with WU in the bottom two. In the PS-SPH runs, a hole is blown into the central cold gas disc that formed from gas cooling out of the hot halo. This central hole is bigger and hotter in PS-TSL than in PS-WU. PE-TSL and MFM-TSL show a smooth dense central disc (smoother for MFM than for PE), which is heated in the inner few kpc by the AGN radiation feedback. In contrast to PS-TSL, the wind feedback does not create a central hole. In the PE-WU and MFM-WU runs, the central disc is sparser and less dense than in their TSL counterparts, and stays cold despite the radiation feedback. Furthermore, MFM-WU leads to the formation of a significantly larger central disc than the other models.

gas particles. In PS-SPH, errors generated by converting entropy into energy (see section 4.2.2 and Hu et al., 2016, for details) can increase the total energy of a blast wave by a significant amount in the beginning of its evolution. If this additionally created energy is enough to blow out the central gas, when it otherwise would not be (which can be the case especially very early in the simulation, when only a small amount of gas has accreted into the centre), a hole can be created. Once the hole exists, it is easier to generate an outflow in subsequent feedback events, which maintains the hole and, in turn, the outflow.

The differences in the AGN feedback’s ability to affect the central gas do not just result in different outflow and metal-enrichment properties, but also affect the stellar and blackhole mass evolution of the simulated galaxy. In Fig. 4.9, we show the time-evolution of the specific star formation rate (sSFR) for the various runs, compared to an empirical “quiescence limit” by Franx et al. (2008), below which a galaxy can be considered quiescent (i.e. without significant star formation). In the PS_TSL run, the sSFR is fluctuating rapidly over several orders of magnitude, resulting in an overall very low star formation rate and a galaxy that is almost always quiescent. In contrast, all the other runs show much less fluctuation in sSFR. Specifically, in the runs using the WU scheme, the sSFR shows almost no short-term fluctuation and is almost always above the quiescence limit. This is because the radiative AGN feedback is unable to heat the central disc in these runs, resulting in the constant presence of a dense, cold, star-forming gas disc in the galactic centre. PE_TSL and MFM_TSL show slowly increasing sSFRs with small (for MFM_TSL) to intermediate (\sim a factor of two, for PE_TSL) short term scatter. In these runs, the galaxy is still quiescent the majority of the time, but the median SFR is much higher than in PS_TSL. The dense gas disc in the centre of the galaxy is mostly heated to high temperatures by the radiative feedback and thereby kept from forming stars, but a fraction of the dense gas cools and becomes star-forming. The amount of star-forming gas slowly increases as more and more gas is accreted onto the central disc while almost none is removed from it through outflows.

Interestingly, the total stellar mass growth (shown in Fig. 4.10 as the total stellar mass over time for all runs) is similar (and close to zero) for both the PS_TSL and the MFM_TSL runs, even though the SFR behaves very differently. This is due to peaks of high star formation in the PS_TSL run. In contrast, the PE_TSL run, which shows an overall similar sSFR to the MFM_TSL run, has a much larger total stellar mass growth because of the larger short term scatter. Unsurprisingly, the WU runs form the most stars in total.

Similar to how the SFR is much smoother and less bursty in the PE-SPH and MFM runs than in the PS_TSL run, the black hole accretion rate (shown over time in Fig. 4.11 for all simulations) is strongly fluctuating in the PS_TSL run, but much more even in the runs using PE-SPH or MFM. Different from the behaviour of the SFR, the black hole accretion is (on average) higher in the PS-SPH runs than in the other four. Using WU reduces the accretion rate compared to using TSL for the PE-SPH and MFM runs by a factor of about 3 to 5, as the generated outflows remove gas from the central region where it can be accreted by the SMBH. The PS_WU run shows the strangest behaviour of the accretion rate. While it is bursty in the PS_TSL run, it almost never shuts off completely,




Figure 4.9: The specific star formation rate over time for the different runs (see legend). The black horizontal line is the quiescence limit (Franx et al., 2008). While the sSFR is strongly fluctuating in the PS_TSL run, its evolution is much smoother in all other combinations of hydrodynamical solver and time-step-limiting scheme. In PE_TSL and MFM_TSL, the sSFR follows a similar trend: slowly rising, but below the quiescence limit for the majority of the time. The scatter is larger for PE_TSL, leading to a significantly larger overall stellar mass growth compared to MFM_TSL (compare Fig. 4.10). If WU is used instead of TSL, the sSFR is much higher (and with less scatter) for all hydrodynamical solvers, leading to a galaxy that is actively star-forming for most of the simulation time.




Figure 4.10: Total stellar mass of the galaxy over time for the runs given in the legend. The runs using WU all show a significant ($\sim 6\%$ to $\sim 8\%$) stellar mass growth over the simulation time, as does PE_TSL ($\sim 5\%$). In contrast, both PS_TSL and MFM_TSL show almost no net stellar mass growth (note that the total stellar mass can decrease due to feedback-related stellar mass loss, if the SFR is low enough).




Figure 4.11: Running mean with a step size of about 30 Myr (10 snapshots) of the black hole mass accretion rate over time for the runs given in the legend. Accretion is very bursty in the PS-SPH runs (to the extreme of discrete, very high, very short accretion peaks in between long periods of no accretion for PS_WU), but significantly smoother with PE-SPH or MFM. There is little difference in the accretion rate between PE-SPH and MFM, while switching from TSL to WU reduces the average accretion rate by about a factor of 3 to 5.

but in the PS_WU run, accretion onto the black hole only happens in very short, very high bursts, which are then followed by prolonged periods of no accretion.

The most likely explanation for this strange behaviour is that the cold, dense disc around the black hole, which exists much more consistently in the PS_WU run than in the PS_TSL run, leads to all of the gas particles neighbouring the black hole to have short smoothing lengths. Then none of them overlap with the Bondi radius of the SMBH, leading to no accretion. Once some gas particles do come close enough to the centre to be eligible for accretion, the accretion rate is very high (as the gas is cold and dense), and they are immediately either accreted or ejected as wind particles. On the other hand, in the PS_TSL run, the gas in the central disc is regularly heated and disturbed by the radiative AGN feedback, which leads to more gas particles with higher temperatures and lower densities in the centre. These particles have larger smoothing lengths and are therefore more likely to be eligible for accretion at larger distances from the black hole. This leads to the low, but non-zero accretion rate periods in the PS_TSL run.

We can see this behaviour in the left two panels of Fig. 4.12, which shows the cumulative accreted gas mass over time in three different temperature bins for the different runs. In the PS_TSL run, about half of the accreted gas is from the intermediate temperature bin ($1.2 \times 10^4 \text{ K} < T \leq 2 \times 10^5 \text{ K}$), and almost 20% is hot gas ($T > 2 \times 10^5 \text{ K}$), while in the PS_WU run, no hot gas is ever accreted, and about 2/3 of the total accreted gas is cold ($T \leq 1.2 \times 10^4 \text{ K}$) at the time of accretion.

Fig. 4.12 also shows a distinct difference in the temperature of the accreted gas between the PE_TSL and the MFM_TSL runs, specifically a much higher fraction of hot accreted gas in the latter (almost 50% compared to $\sim 15\%$). The MFM_TSL run, in which the AGN feedback is able to create some outflows (especially in the first 2/3 of the simulation, in which the hot accreted fraction is the highest), has much more hot gas close to the SMBH than the PE_TSL run. This suggests that, in MFM, the gas cools more slowly in a blast wave (as the central gas is affected by the wind feedback) than it does in PE-SPH, which also allows for more, and stronger outflows. In the PE_WU and MFM_WU runs, the accreted gas is on average much colder than in the corresponding TSL runs, which reflects the inability of the feedback in these runs to heat the central gas, as most feedback-affected gas is immediately escaping to larger radii.

4.5 Summary & conclusions

In this chapter, we studied the importance of the underlying hydrodynamical methods used in simulations with a powerful sub-resolution AGN feedback model for the effects of this feedback on these simulations. For this, we reran the fiducial simulation described in chapter 2 (i.e. an isolated, old, massive elliptical galaxy evolved with an SPH code including models for cooling, star formation, stellar feedback, and AGN feedback in the form of fast winds and radiation), while changing some of the underpinnings of the hydrodynamical code.

For context, we started with a simple parameter study. Keeping everything else con-




Figure 4.12: Cumulative mass of gas accreted by the central black hole until a given time in different temperature bins over time, normalised by the total gas mass accreted onto the black hole during this time. The red solid lines show the accreted gas mass fractions for gas hotter than 2×10^5 K, the cyan dashed line the fractions for warm gas within $1.2 \times 10^4 \text{ K} < T \leq 2 \times 10^5$ K, and the blue dashed-dotted line for gas colder than 1.2×10^4 K. The top row shows the runs using TSL, the bottom row those using WU, while the columns show, from left to right, the runs with PS-SPH, PE-SPH, and MFM. In the PS_TSL run, most accreted gas is warm, and a significant fraction is hot, while in the PS_WU run, most of the gas is accreted while it is cold, and no gas is accreted hot. The same trend of the accreted gas being on average colder in the WU runs compared to the TSL runs is true for the PE-SPH and MFM runs, as well. The MFM_TSL run shows a much higher fraction of hot accreted gas than the PE_TSL run.

stant, we varied the free parameter of the AGN wind feedback efficiency ϵ_w over a factor of 8 (see section 4.3), and compared the reach of metal-enriching AGN-powered outflows dependent on ϵ_w . We found that ϵ_w has to be above $\sim 0.25\%$ for the AGN wind feedback to generate any large-scale metal-enriching outflows, and that further increasing the efficiency can increase the reach of these outflows, but only if the increase is large enough.

With the initial motivation to test the stability and reliability of these results, we then compared the results of six simulations, in which we varied the numerical solver used to compute the hydrodynamical forces, and the methods used to limit the time-steps of the simulation particles, but kept everything else (initial condition, sub-resolution models and their parameters, etc.) constant (see section 4.4). For these runs, we took the highest wind efficiency used in the preceding parameter study, $\epsilon_w = 0.02$, to maximise the effect of the AGN wind feedback. The three different hydrodynamical solvers used (see section 4.2.2) are pressure-entropy SPH (PS-SPH; the formalism used in chapter 2 and the parameter study), pressure-energy SPH (PE-SPH), and meshless-finite-mass (MFM). These are combined with two different methods for limiting time-steps of gas particles (see section 4.2.3): the time-step-limiter (TSL), and the wake-up scheme (WU), resulting in six simulations in total.

Our results show that varying the hydrodynamics in this way, while keeping everything else constant, has an even larger impact on the (AGN-driven) metal distribution than changing the wind efficiency parameter. Furthermore, these differences are not just of a spurious, quantitative kind (i.e. different outcomes of the same kind of evolution generated from small changes due to the highly dynamical nature of the process), but based on fundamental, qualitative differences in how the AGN feedback affects its surrounding gas under the different hydrodynamical schemes. We see several trends that can be explained by various flaws of the different methods used:

- In runs using the WU scheme instead of TSL, the AGN generates stronger outflows if the hydrodynamical solver used is either PE-SPH or MFM. This is a consequence of the missing instantaneous activation of gas particles directly affected by feedback in WU, which is present in TSL. This missing activation leads to the escape of almost all heated gas from the centre before it can distribute its heat, resulting in a permanently cold central gas disc, and a correspondingly high star formation rate (with essentially no periods of quiescence) in all of the runs using WU.
- If PS-SPH is used, the AGN feedback blows a hole of significant size into the centre of the gas disc surrounding the black hole. The feedback is then able to generate powerful outflows (even using TSL) while in the PE-SPH and MFM runs (using TSL), the centre stays dense, and much fewer strong outflows are generated. A likely reason for this difference is the energy-conservation problem that PS-SPH has in shocks (see Hu et al., 2016).
- While, if TSL is used, the feedback in the PE-SPH and MFM runs is not able to enrich the gas out to large radii until the end of the simulation time, it keeps the central dense gas disc mostly hot, thereby preventing star formation. As more gas

flows into the centre, and the density and mass of the central disc slowly rise, the SFR slowly rises over time, but still stays below the Franx et al. (2008) quiescence limit for the majority of the time. Even though the general behaviour looks similar, the total stellar mass grows much more in the PE-SPH TSL run than in the MFM TSL run, because the star formation rate fluctuates more and to higher values.

- Due to the differences in the central gas properties (i.e. hole or no hole), the accretion rate of the black hole is very bursty in the PS-SPH runs, but much smoother (and overall lower) in the PE-SPH and MFM runs. As using the WU scheme increases the outflow rate—and thereby lowers the central gas mass—greatly with these solvers, the black hole growth is even lower in the PE-SPH and MFM runs using WU than in the corresponding TSL runs.
- The gas accreted by the black hole is overall much colder at the time of accretion in the WU runs than in the TSL runs, as all hot gas is instantly removed from the central region in the former. More interestingly, the accretion temperature of the gas is significantly higher on average in the MFM TSL run than in the PE-SPH TSL run. Hotter accretion temperatures coincide with stronger outflows for the MFM TSL run, and heat tends to disperse less quickly into the surrounding gas in MFM than in PE-SPH, leading to more collimated outflows.

The first, and simplest, conclusion from the presented tests is that the TSL formalism is clearly better suited for simulations with sub-resolution feedback models than the WU scheme, as the latter’s missing mechanism to activate feedback-affected particles instantly is crucial. This is of course generalisable to all similar schemes dedicated to limiting time-steps of gas particles. Furthermore, in simulations with strong shocks, e.g. all simulations including AGN feedback of any sort, pressure-energy SPH is preferable to the pressure-entropy variant, as the latter’s energy conservation issues, even if they are small, can lead to drastically different outcomes for almost all aspects of the simulations results. Pressure-energy SPH and MFM show much more similar results, which is promising for the applicability of both of these types of simulations, but even between these, there are several significant differences, especially in how the gas behaves in shocks in regards to cooling and collimation. These detailed differences also result in large disparities in the results in our simulations, both for the large-scale metal-enrichment in the first half of the simulation time, and for the total stellar mass growth.

It is also not immediately clear which of these results is preferable, as there is no “correct” solution to compare to, and any constraints by observations might be matched with either hydrodynamical method using different sub-resolution models (or even just varying the free parameters in them). Hopkins (2015) suggests that MFM is more accurate than either version of SPH, and backs up this claim by a large amount of tests with known analytical solutions in which MFM outperforms SPH. However, in his test simulations of an isolated disc galaxy using self-consistent sub-resolution models for star formation and stellar feedback, the different methods’ results agree well at first order, suggesting that the sub-resolution models determine the outcome to a much larger degree than the

hydrodynamical method used. Several other code comparison studies arrived at the same conclusion (e.g. Scannapieco et al., 2012; Kim et al., 2016; Sembolini et al., 2016b). While the large influence of the chosen sub-resolution models is certainly also true for simulations of massive systems with AGN feedback (changing the AGN feedback model can vastly change the outcome of those simulations), the hydrodynamical model used clearly has a very large impact as well, and not just on the second order.

In general, it is therefore advisable to be very careful when interpreting the results of simulations of complex systems, especially those involving high energy processes on unresolved scales, which need to be implemented in some approximated form as a sub-resolution model. The outcomes of these simulations are not only sensitive to the specific implementation of the unresolved physical processes, but also to the far more fundamental methods used to determine the basic hydrodynamic forces between the fluid elements. It is very hard to predict how a change in these methods might affect the simulation, e.g. if it would make a process more or less effective, and by how much, even if these methods were checked and compared against each other in simpler test cases before. This makes it essentially impossible to determine when any given method is accurate enough to properly represent the actual effects of a high-energy physical phenomenon like AGN feedback on its surrounding gas without extensive testing of the method in the circumstances it is actually used in, which is prohibitively computationally expensive. Hence, as long as it is not possible to resolve these kinds of physical processes down to scales at which their effects are properly constrained (so that their actual implementation is sufficiently similar to a test case with a known analytical result), conclusions drawn from simulations including these high-energy effects should always be carefully caveated, taken to apply to the specific simulation setup, and not extrapolated to general explanations of how the investigated processes work in the real universe.

Summary

In this thesis, we investigated several aspects of the evolution of the multiphase gas in and around early-type galaxies and how it is influenced by the supermassive black holes at the centres of these systems. For this purpose we performed multiple hydrodynamical simulations using the code SPHGal (an improved version of GADGET-3) with various different settings and sub-resolution models.

In chapter 2, we simulated the effects of kinetic and radiative feedback from the active galactic nucleus on the cooling gas from the hot halo of an idealised, massive, quiescent elliptical galaxy. We showed that AGN feedback is able and necessary to keep the galaxy quiescent over the simulation time of 4.35 Gyr. The inclusion of both radiative and kinetic feedback is essential to prevent the gas from condensing into a consistent cold, dense phase in the galactic centre, while keeping other properties of the galaxy (such as the SMBH mass) consistent with observations as well. The AGN feedback also generates large-scale outflows of metal-rich gas, which keep the circumgalactic medium of the galaxy enriched with metals out to a radius of about 30 kpc.

In chapter 3, we simulated the evolution of a circumnuclear, molecular gas disc embedded in the deep gravitational potential of a massive early-type galaxy with a central supermassive black hole. In close comparison to the observed properties of similar CNDs, particularly the one in NGC 4429 (Davis et al., 2018), we showed that the fast rotation of the gas (which is caused by the deep potentials of the stars and SMBH) is unable to sufficiently stabilise the simulated gas disc against fragmentation into very dense ($\rho > 10^{-20} \text{ g cm}^{-3}$) clumps that form stars at a high rate. Due to this, the star formation in the simulated CND is too high compared to NGC 4429 and observations in similar systems, where it is suppressed with respect to the level of star-forming spiral galaxies. Several other tested mechanisms—including a stronger UV radiation field, cosmic ray ionisation, a new, more effective implementation of supernova feedback, and the inclusion of black hole accretion and feedback—also fail to suppress the star formation sufficiently to match the observations. We concluded that, to achieve the necessary star formation suppression, some other, non-included physical mechanism must be at play, e.g. magnetic fields. However, we found that SMBH accretion and feedback strongly affect the central $\sim 100 \text{ pc}$ of the CND, which are kept largely devoid of gas by these mechanisms. Furthermore, the overall growth of the SMBH is correlated to the overall star formation, as supernovae drive gas towards the centre, where it can be accreted by the SMBH, though only if the AGN feedback is ignored.

Finally, in chapter 4, we investigated the effect of changing the underlying numerics

on the evolution of the galaxy by repeating our simulations from chapter 2 while changing first the AGN wind feedback efficiency, and then the implementation of the hydrodynamic force calculations. We found that changing the wind efficiency considerably affects the effectiveness of the AGN feedback at driving outflows, such that significantly lowering it prevents any outflow from occurring, while strongly raising it increases the reach of the outflows. Changing the hydrodynamical solver has an even larger effect on the simulations: Pressure-entropy SPH results in much stronger AGN-driven outflows than either pressure-energy SPH or MFM, due to energy conservation problems inherent to the method. The inclusion of a time-step-limiting scheme that takes into account gas particles that were affected by feedback is similarly important to prevent overly energetic AGN feedback. The differences in the results between pressure-energy SPH and MFM are less severe, but still significant, which led us to conclude that the results of any hydrodynamical simulation that includes very energetic events of the sorts of AGN feedback need to be interpreted carefully, as they may depend to a large degree on the specifics of their fundamental implementation.

Bibliography

- Abbott B. P., et al., 2016, *Physical Review Letters*, 116, 061102
- Agertz O., et al., 2007, *MNRAS*, 380, 963
- Alatalo K., et al., 2011, *ApJ*, 735, 88
- Alatalo K., et al., 2013, *MNRAS*, 432, 1796
- Alatalo K., et al., 2015, *ApJ*, 798, 31
- Anderson M. E., Gaspari M., White S. D. M., Wang W., Dai X., 2015, *MNRAS*, 449, 3806
- Anglés-Alcázar D., Özel F., Davé R., 2013, *ApJ*, 770, 5
- Anglés-Alcázar D., Davé R., Faucher-Giguère C.-A., Özel F., Hopkins P. F., 2017, *MNRAS*, 464, 2840
- Antonucci R., 1993, *Annual Review of Astronomy and Astrophysics*, 31, 473
- Arnaud K. A., 1996, in Jacoby G. H., Barnes J., eds, *Astronomical Society of the Pacific Conference Series Vol. 101, Astronomical Data Analysis Software and Systems V*. p. 17
- Aumer M., White S. D. M., Naab T., Scannapieco C., 2013, *MNRAS*, 434, 3142
- Balkowski C., Bottinelli L., Gouguenheim L., Heidmann J., 1972, *A & A*, 21, 303
- Barnabè M., Czoske O., Koopmans L. V. E., Treu T., Bolton A. S., 2011, *MNRAS*, 415, 2215
- Barnes J., Hut P., 1986, *Nature*, 324, 446
- Beckmann V., Shrader C. R., 2012, *Active Galactic Nuclei*
- Bender R., Burstein D., Faber S. M., 1992, *ApJ*, 399, 462
- Bieri R., Dubois Y., Rosdahl J., Wagner A., Silk J., Mamon G. A., 2017, *MNRAS*, 464, 1854
- Binney J., 1977, *ApJ*, 215, 483

- Binney J., Tremaine S., 2008, *Galactic Dynamics: Second Edition*. Princeton University Press
- Birzan L., Rafferty D. A., McNamara B. R., Wise M. W., Nulsen P. E. J., 2004, *ApJ*, 607, 800
- Bocquet S., Saro A., Dolag K., Mohr J. J., 2016, *MNRAS*, 456, 2361
- Boizelle B. D., Barth A. J., Darling J., Baker A. J., Buote D. A., Ho L. C., Walsh J. L., 2017, *ApJ*, 845, 170
- Bondi H., 1952, *MNRAS*, 112, 195
- Bondi H., Hoyle F., 1944, *MNRAS*, 104, 273
- Booth C. M., Schaye J., 2009, *MNRAS*, 398, 53
- Boroson B., Kim D.-W., Fabbiano G., 2011, *ApJ*, 729, 12
- Bower R. G., Benson A. J., Malbon R., Helly J. C., Frenk C. S., Baugh C. M., Cole S., Lacey C. G., 2006, *MNRAS*, 370, 645
- Brighenti F., Mathews W. G., 2006, *ApJ*, 643, 120
- Brusa M., et al., 2015, *A & A*, 578, A11
- Cano-Díaz M., Maiolino R., Marconi A., Netzer H., Shemmer O., Cresci G., 2012, *A & A*, 537, L8
- Cappellari M., et al., 2011, *MNRAS*, 413, 813
- Carniani S., et al., 2016, *A & A*, 591, A28
- Casares J., Jonker P. G., 2014, *Space Science Reviews*, 183, 223
- Cavaliere A., Fusco-Femiano R., 1976, *A & A*, 49, 137
- Chandrasekhar S., 1935, *MNRAS*, 95, 207
- Chang Y.-Y., van der Wel A., da Cunha E., Rix H.-W., 2015, *ApJ Supplement Series*, 219, 8
- Cheung E., et al., 2016, *Nature*, 533, 504
- Choi E., Ostriker J. P., Naab T., Johansson P. H., 2012, *ApJ*, 754, 125
- Choi E., Naab T., Ostriker J. P., Johansson P. H., Moster B. P., 2014, *MNRAS*, 442, 440
- Choi E., Ostriker J. P., Naab T., Oser L., Moster B. P., 2015, *MNRAS*, 449, 4105

- Choi E., Ostriker J. P., Naab T., Somerville R. S., Hirschmann M., Núñez A., Hu C.-Y., Oser L., 2017, *ApJ*, 844, 31
- Churazov E., Forman W., Jones C., Böhringer H., 2000, *A & A*, 356, 788
- Churazov E., Brüggén M., Kaiser C. R., Böhringer H., Forman W., 2001, *ApJ*, 554, 261
- Churazov E., Sunyaev R., Forman W., Böhringer H., 2002, *MNRAS*, 332, 729
- Churazov E., Sazonov S., Sunyaev R., Forman W., Jones C., Böhringer H., 2005, *MNRAS*, 363, L91
- Cicone C., et al., 2014, *A & A*, 562, A21
- Cielo S., Bieri R., Volonteri M., Wagner A. Y., Dubois Y., 2018, *MNRAS*, 477, 1336
- Ciotti L., Ostriker J. P., 1997, *ApJ Letters*, 487, L105
- Ciotti L., Ostriker J. P., 2001, *ApJ*, 551, 131
- Ciotti L., Ostriker J. P., 2007, *ApJ*, 665, 1038
- Ciotti L., Ostriker J. P., Proga D., 2009, *ApJ*, 699, 89
- Ciotti L., Ostriker J. P., Proga D., 2010, *ApJ*, 717, 708
- Ciotti L., Pellegrini S., Negri A., Ostriker J. P., 2017, *ApJ*, 835, 15
- Clark P. C., Glover S. C. O., Klessen R. S., 2012, *MNRAS*, 420, 745
- Combes F., Young L. M., Bureau M., 2007, *MNRAS*, 377, 1795
- Constantin A., Shields J. C., Ho L. C., Barth A. J., Filippenko A. V., Castillo C. A., 2015, *ApJ*, 814, 149
- Couchman H. M. P., Thomas P. A., Pearce F. R., 1995, *ApJ*, 452, 797
- Crenshaw D. M., Kraemer S. B., George I. M., 2003, *Annual Review of Astronomy and Astrophysics*, 41, 117
- Cresci G., et al., 2015a, *A & A*, 582, A63
- Cresci G., et al., 2015b, *ApJ*, 799, 82
- Crocker A. F., Bureau M., Young L. M., Combes F., 2011, *MNRAS*, 410, 1197
- Croton D. J., et al., 2006, *MNRAS*, 365, 11
- Dai X., Bregman J. N., Kochanek C. S., Rasia E., 2010, *ApJ*, 719, 119

- Danielson A. L. R., Lehmer B. D., Alexander D. M., Brandt W. N., Luo B., Miller N., Xue Y. Q., Stott J. P., 2012, *MNRAS*, 422, 494
- Dasyra K. M., Bostrom A. C., Combes F., Vlahakis N., 2015, *ApJ*, 815, 34
- Davé R., Thompson R., Hopkins P. F., 2016, *MNRAS*, 462, 3265
- Davies R. I., Müller Sánchez F., Genzel R., Tacconi L. J., Hicks E. K. S., Friedrich S., Sternberg A., 2007, *ApJ*, 671, 1388
- Davis T. A., et al., 2014, *MNRAS*, 444, 3427
- Davis T. A., Bureau M., Onishi K., Cappellari M., Iguchi S., Sarzi M., 2017, *MNRAS*, 468, 4675
- Davis T. A., et al., 2018, *MNRAS*, 473, 3818
- Debuhr J., Quataert E., Ma C.-P., 2011, *MNRAS*, 412, 1341
- Debuhr J., Quataert E., Ma C.-P., 2012, *MNRAS*, 420, 2221
- Dekel A., et al., 2009, *Nature*, 457, 451
- Delvecchio I., et al., 2015, *MNRAS*, 449, 373
- Di Matteo T., Springel V., Hernquist L., 2005, *Nature*, 433, 604
- Diehl S., Statler T. S., 2007, *ApJ*, 668, 150
- Downes D., Solomon P. M., 1998, *ApJ*, 507, 615
- Draine B. T., Bertoldi F., 1996, *ApJ*, 468, 269
- Dressler A., 1980, *ApJ*, 236, 351
- Dubois Y., Devriendt J., Slyz A., Teyssier R., 2010, *MNRAS*, 409, 985
- Dubois Y., Devriendt J., Slyz A., Teyssier R., 2012, *MNRAS*, 420, 2662
- Dubois Y., Gavazzi R., Peirani S., Silk J., 2013, *MNRAS*, 433, 3297
- Dubois Y., et al., 2014, *MNRAS*, 444, 1453
- Dubois Y., Peirani S., Pichon C., Devriendt J., Gavazzi R., Welker C., Volonteri M., 2016, *MNRAS*, 463, 3948
- Dunn J. P., et al., 2010, *ApJ*, 709, 611
- Durier F., Dalla Vecchia C., 2012, *MNRAS*, 419, 465

- Ebneter K., Djorgovski S., Davis M., 1988, *Astronomical Journal*, 95, 422
- Eddington A. S., 1916, *MNRAS*, 77, 16
- Eisenreich M., Naab T., Choi E., Ostriker J. P., Emsellem E., 2017, *MNRAS*, 468, 751
- Eke V. R., Navarro J. F., Frenk C. S., 1998, *ApJ*, 503, 569
- Faber S. M., Jackson R. E., 1976, *ApJ*, 204, 668
- Fabian A. C., 2012, *Annual Review of Astronomy and Astrophysics*, 50, 455
- Forman W., Schwarz J., Jones C., Liller W., Fabian A. C., 1979, *ApJ Letters*, 234, L27
- Forman W., et al., 2007, *ApJ*, 665, 1057
- Franx M., van Dokkum P. G., Schreiber N. M. F., Wuyts S., Labbé I., Toft S., 2008, *ApJ*, 688, 770
- Frenk C. S., et al., 1999, *ApJ*, 525, 554
- Gaburov E., Nitadori K., 2011, *MNRAS*, 414, 129
- Gaibler V., Khochfar S., Krause M., Silk J., 2012, *MNRAS*, 425, 438
- Gan Z., Yuan F., Ostriker J. P., Ciotti L., Novak G. S., 2014, *ApJ*, 789, 150
- Genel S., et al., 2014, *MNRAS*, 445, 175
- Genzel R., Pichon C., Eckart A., Gerhard O. E., Ott T., 2000, *MNRAS*, 317, 348
- Ghez A. M., Morris M., Becklin E. E., Tanner A., Kremenek T., 2000, *Nature*, 407, 349
- Girichidis P., et al., 2016, *MNRAS*, 456, 3432
- Glover S. C. O., Clark P. C., 2012, *MNRAS*, 421, 116
- Glover S. C. O., Mac Low M.-M., 2007, *ApJ Supplement Series*, 169, 239
- Gomes J. M., et al., 2016, *A & A*, 588, A68
- González-Alfonso E., et al., 2013, *A & A*, 550, A25
- Górski K. M., Hivon E., 2011, *HEALPix: Hierarchical Equal Area isoLatitude Pixelization of a sphere*, *Astrophysics Source Code Library* (ascl:1107.018)
- Grange Y. G., de Plaa J., Kaastra J. S., Werner N., Verbunt F., Paerels F., de Vries C. P., 2011, *A & A*, 531, A15
- Greene J. E., Ho L. C., 2007, *ApJ*, 667, 131

- Greene J. E., Murphy J. D., Graves G. J., Gunn J. E., Raskutti S., Comerford J. M., Gebhardt K., 2013, *ApJ*, 776, 64
- Habing H. J., 1968, *Bulletin of the Astronomical Institutes of the Netherlands*, 19, 421
- Hayward C. C., Torrey P., Springel V., Hernquist L., Vogelsberger M., 2014, *MNRAS*, 442, 1992
- Heckman T. M., Best P. N., 2014, *Annual Review of Astronomy and Astrophysics*, 52, 589
- Heckman T. M., Kauffmann G., Brinchmann J., Charlot S., Tremonti C., White S. D. M., 2004, *ApJ*, 613, 109
- Hernquist L., 1990, *ApJ*, 356, 359
- Hicks E. K. S., Davies R. I., Malkan M. A., Genzel R., Tacconi L. J., Müller Sánchez F., Sternberg A., 2009, *ApJ*, 696, 448
- Hicks E. K. S., Davies R. I., Maciejewski W., Emsellem E., Malkan M. A., Dumas G., Müller-Sánchez F., Rivers A., 2013, *ApJ*, 768, 107
- Ho L. C., 2009, *ApJ*, 699, 626
- Ho L., Kormendy J., 2000, *Supermassive Black Holes in Active Galactic Nuclei*. p. 2365, doi:10.1888/0333750888/2365
- Hopkins P. F., 2013, *MNRAS*, 428, 2840
- Hopkins P. F., 2015, *MNRAS*, 450, 53
- Hopkins P. F., Quataert E., 2010, *MNRAS*, 407, 1529
- Hopkins P. F., Quataert E., 2011, *MNRAS*, 415, 1027
- Hopkins P. F., Hernquist L., Cox T. J., Di Matteo T., Martini P., Robertson B., Springel V., 2005, *ApJ*, 630, 705
- Hopkins P. F., Cox T. J., Kereš D., Hernquist L., 2008, *ApJ Supplement Series*, 175, 390
- Hopkins P. F., Kereš D., Oñorbe J., Faucher-Giguère C.-A., Quataert E., Murray N., Bullock J. S., 2014, *MNRAS*, 445, 581
- Hopkins P. F., Torrey P., Faucher-Giguère C.-A., Quataert E., Murray N., 2016, *MNRAS*, 458, 816
- Hopkins P. F., et al., 2017, preprint, ([arXiv:1702.06148](https://arxiv.org/abs/1702.06148))
- Hoyle F., Lyttleton R. A., 1939, *Proceedings of the Cambridge Philosophical Society*, 35, 405

- Hu C.-Y., Naab T., Walch S., Moster B. P., Oser L., 2014, MNRAS, 443, 1173
- Hu C.-Y., Naab T., Walch S., Glover S. C. O., Clark P. C., 2016, MNRAS, 458, 3528
- Hu C.-Y., Naab T., Glover S. C. O., Walch S., Clark P. C., 2017, MNRAS, 471, 2151
- Hubble E. P., 1926, ApJ, 64
- Humphrey P. J., Buote D. A., Gastaldello F., Zappacosta L., Bullock J. S., Brighenti F., Mathews W. G., 2006, ApJ, 646, 899
- Illustris Collaboration 2014, Illustris Simulation, <http://www.illustris-project.org/media/>
- Iverson R. J., et al., 2010, A & A, 518, L35
- Iwamoto K., Brachwitz F., Nomoto K., Kishimoto N., Umeda H., Hix W. R., Thielemann F.-K., 1999, ApJ Supplement Series, 125, 439
- Janka H.-T., 2012, Annual Review of Nuclear and Particle Science, 62, 407
- Johansson P. H., Naab T., Ostriker J. P., 2009, ApJ Letters, 697, L38
- Jones C., Forman W., 1984, ApJ, 276, 38
- Kauffmann G., Heckman T. M., 2009, MNRAS, 397, 135
- Kauffmann G., et al., 2003a, MNRAS, 341, 33
- Kauffmann G., et al., 2003b, MNRAS, 346, 1055
- Kehrig C., et al., 2012, A & A, 540, A11
- Kennicutt Jr. R. C., 1998, ApJ, 498, 541
- Kereš D., Katz N., Weinberg D. H., Davé R., 2005, MNRAS, 363, 2
- Kereš D., Vogelsberger M., Sijacki D., Springel V., Hernquist L., 2012, MNRAS, 425, 2027
- Khandai N., Di Matteo T., Croft R., Wilkins S., Feng Y., Tucker E., DeGraf C., Liu M.-S., 2015, MNRAS, 450, 1349
- Kim D.-W., Fabbiano G., 2013, ApJ, 776, 116
- Kim D.-W., Fabbiano G., 2015, ApJ, 812, 127
- Kim C.-G., Ostriker E. C., 2015, ApJ, 802, 99
- Kim J.-h., Wise J. H., Alvarez M. A., Abel T., 2011, ApJ, 738, 54

- Kim J.-h., et al., 2014, *ApJ Supplement Series*, 210, 14
- Kim J.-h., et al., 2016, *ApJ*, 833, 202
- Knapp G. R., Turner E. L., Cunniffe P. E., 1985, *Astronomical Journal*, 90, 454
- Konami S., Matsushita K., Nagino R., Tamagawa T., 2014, *ApJ*, 783, 8
- Kormendy J., Ho L. C., 2013, *Annual Review of Astronomy and Astrophysics*, 51, 511
- Kormendy J., Richstone D., 1995, *Annual Review of Astronomy and Astrophysics*, 33, 581
- Kravtsov A. V., Vikhlinin A. A., Meshcheryakov A. V., 2018, *Astronomy Letters*, 44, 8
- Laganá T. F., Martinet N., Durret F., Lima Neto G. B., Maughan B., Zhang Y.-Y., 2013, *A & A*, 555, A66
- Lanson N., Vila J.-P., 2008a, *SIAM Journal on Numerical Analysis*, 46, 1912
- Lanson N., Vila J.-P., 2008b, *SIAM Journal on Numerical Analysis*, 46, 1935
- Le Brun A. M. C., McCarthy I. G., Schaye J., Ponman T. J., 2014, *MNRAS*, 441, 1270
- Loenen A. F., et al., 2010, *A & A*, 521, L2
- Makino N., Sasaki S., Suto Y., 1998, *ApJ*, 497, 555
- Malkan M. A., Jensen L. D., Rodriguez D. R., Spinoglio L., Rush B., 2017, *ApJ*, 846, 102
- Maoz D., 2007, *MNRAS*, 377, 1696
- Maoz D., Mannucci F., 2012, *Publications of the Astronomical Society of Australia*, 29, 447
- Martig M., Bournaud F., Teyssier R., Dekel A., 2009, *ApJ*, 707, 250
- Martig M., et al., 2013, *MNRAS*, 432, 1914
- Martizzi D., Teyssier R., Moore B., 2012, *MNRAS*, 420, 2859
- Mathews W. G., Brighenti F., 2003, *Annual Review of Astronomy and Astrophysics*, 41, 191
- McCarthy I. G., et al., 2010, *MNRAS*, 406, 822
- McCarthy I. G., Schaye J., Bower R. G., Ponman T. J., Booth C. M., Dalla Vecchia C., Springel V., 2011, *MNRAS*, 412, 1965
- McDermid R. M., et al., 2015, *MNRAS*, 448, 3484

- McNamara B. R., et al., 2000, *ApJ Letters*, 534, L135
- Medling A. M., et al., 2014, *ApJ*, 784, 70
- Micic M., Glover S. C. O., Federrath C., Klessen R. S., 2012, *MNRAS*, 421, 2531
- Minkowski R., Osterbrock D., 1959, *ApJ*, 129, 583
- Mitchell N. L., McCarthy I. G., Bower R. G., Theuns T., Crain R. A., 2009, *MNRAS*, 395, 180
- Mo H., van den Bosch F. C., White S., 2010, *Galaxy Formation and Evolution*. Cambridge University Press, Cambridge
- Moe M., Arav N., Bautista M. A., Korista K. T., 2009, *ApJ*, 706, 525
- Morganti R., Oosterloo T., Oonk J. B. R., Frieswijk W., Tadhunter C., 2015, *A & A*, 580, A1
- Moster B. P., Macciò A. V., Somerville R. S., Naab T., Cox T. J., 2011, *MNRAS*, 415, 3750
- Moster B. P., Naab T., White S. D. M., 2013, *MNRAS*, 428, 3121
- Mulchaey J. S., Jeltema T. E., 2010, *ApJ Letters*, 715, L1
- Naab T., Ostriker J. P., 2017, *Annual Review of Astronomy and Astrophysics*, 55, 59
- Nagai D., Kravtsov A. V., Vikhlinin A., 2007, *ApJ*, 668, 1
- Navarro J. F., Frenk C. S., White S. D. M., 1996, *ApJ*, 462, 563
- Nelson R. P., Langer W. D., 1997, *ApJ*, 482, 796
- Nelson D., Vogelsberger M., Genel S., Sijacki D., Kereš D., Springel V., Hernquist L., 2013, *MNRAS*, 429, 3353
- Nelson D., et al., 2018, *MNRAS*, 475, 624
- Novak G. S., Ostriker J. P., Ciotti L., 2011, *ApJ*, 737, 26
- Núñez A., Ostriker J. P., Naab T., Oser L., Hu C.-Y., Choi E., 2017, *ApJ*, 836, 204
- Nyland K., et al., 2016, *MNRAS*, 458, 2221
- O'Sullivan E., Sanderson A. J. R., Ponman T. J., 2007, *MNRAS*, 380, 1409
- Omnia H., Binney J., Bryan G., Slyz A., 2004, *MNRAS*, 348, 1105

- Onishi K., Iguchi S., Davis T. A., Bureau M., Cappellari M., Sarzi M., Blitz L., 2017, MNRAS, 468, 4663
- Opik E., 1922, ApJ, 55
- Oppenheimer J. R., Volkoff G. M., 1939, Physical Review, 55, 374
- Oser L., Ostriker J. P., Naab T., Johansson P. H., Burkert A., 2010, ApJ, 725, 2312
- Osterbrock D. E., 1960, ApJ, 132, 325
- Ostriker E. C., McKee C. F., Leroy A. K., 2010a, ApJ, 721, 975
- Ostriker J. P., Choi E., Ciotti L., Novak G. S., Proga D., 2010b, ApJ, 722, 642
- Padovani P., et al., 2017, Astronomy and Astrophysics Review, 25, 2
- Page M. J., et al., 2012, Nature, 485, 213
- Pakmor R., Edelmann P., Röpke F. K., Hillebrandt W., 2012, MNRAS, 424, 2222
- Peng Y.-j., et al., 2010, ApJ, 721, 193
- Pinto C., et al., 2014, A & A, 572, L8
- Planck Collaboration et al., 2013, A & A, 557, A52
- Planck Collaboration et al., 2014, A & A, 571, A16
- Planck Collaboration et al., 2016, A & A, 594, A1
- Price D. J., 2012, Journal of Computational Physics, 231, 759
- Prieto M. A., Maciejewski W., Reunanen J., 2005, Astronomical Journal, 130, 1472
- Proga D., Kallman T. R., 2004, ApJ, 616, 688
- Puchwein E., Springel V., 2013, MNRAS, 428, 2966
- Qu Y., et al., 2017, MNRAS, 464, 1659
- Read J. I., Hayfield T., Agertz O., 2010, MNRAS, 405, 1513
- Rees M. J., Ostriker J. P., 1977, MNRAS, 179, 541
- Remillard R. A., McClintock J. E., 2006, Annual Review of Astronomy and Astrophysics, 44, 49
- Renzini A., Peng Y.-j., 2015, ApJ Letters, 801, L29
- Renzini A., Ciotti L., D'Ercole A., Pellegrini S., 1993, ApJ, 419, 52

- Richardson M. L. A., Scannapieco E., Devriendt J., Slyz A., Thacker R. J., Dubois Y., Wurster J., Silk J., 2016, *ApJ*, 825, 83
- Rodriguez-Gomez V., et al., 2016, *MNRAS*, 458, 2371
- Röttgers B., Arth A., 2018, preprint, ([arXiv:1803.03652](https://arxiv.org/abs/1803.03652))
- Sage L. J., Wrobel J. M., 1989, *ApJ*, 344, 204
- Saglia R. P., et al., 2016, *ApJ*, 818, 47
- Saitoh T. R., Makino J., 2009, *ApJ Letters*, 697, L99
- Sarzi M., et al., 2013, *MNRAS*, 432, 1845
- Sazonov S. Y., Ostriker J. P., Ciotti L., Sunyaev R. A., 2005, *MNRAS*, 358, 168
- Scannapieco C., Tissera P. B., White S. D. M., Springel V., 2005, *MNRAS*, 364, 552
- Scannapieco C., Tissera P. B., White S. D. M., Springel V., 2006, *MNRAS*, 371, 1125
- Scannapieco C., et al., 2012, *MNRAS*, 423, 1726
- Schaller M., Dalla Vecchia C., Schaye J., Bower R. G., Theuns T., Crain R. A., Furlong M., McCarthy I. G., 2015, *MNRAS*, 454, 2277
- Schartmann M., Mould J., Wada K., Burkert A., Durré M., Behrendt M., Davies R. I., Burtscher L., 2018, *MNRAS*, 473, 953
- Schaye J., et al., 2015, *MNRAS*, 446, 521
- Schechter P., 1976, *ApJ*, 203, 297
- Schmidt M., 1959, *ApJ*, 129, 243
- Sedov L. I., 1959, *Similarity and Dimensional Methods in Mechanics*. Academic Press, New York
- Sembolini F., et al., 2016a, *MNRAS*, 457, 4063
- Sembolini F., et al., 2016b, *MNRAS*, 459, 2973
- Serra P., et al., 2012, *MNRAS*, 422, 1835
- Shakura N. I., Sunyaev R. A., 1973, *A & A*, 24, 337
- Shin M.-S., Ostriker J. P., Ciotti L., 2010, *ApJ*, 711, 268
- Sijacki D., Springel V., Di Matteo T., Hernquist L., 2007, *MNRAS*, 380, 877

- Sijacki D., Vogelsberger M., Kereš D., Springel V., Hernquist L., 2012, *MNRAS*, 424, 2999
- Sijacki D., Vogelsberger M., Genel S., Springel V., Torrey P., Snyder G. F., Nelson D., Hernquist L., 2015, *MNRAS*, 452, 575
- Silk J., 1977, *ApJ*, 211, 638
- Silk J., 2013, *ApJ*, 772, 112
- Silk J., Rees M. J., 1998, *A & A*, 331, L1
- Smith R. J., Lucey J. R., Price J., Hudson M. J., Phillipps S., 2012, *MNRAS*, 419, 3167
- Smith A., Bromm V., Loeb A., 2017, *Astronomy and Geophysics*, 58, 3.22
- Soltan A., 1982, *MNRAS*, 200, 115
- Somerville R. S., Davé R., 2015, *Annual Review of Astronomy and Astrophysics*, 53, 51
- Springel V., 2005, *MNRAS*, 364, 1105
- Springel V., 2010, *MNRAS*, 401, 791
- Springel V., Hernquist L., 2003, *MNRAS*, 339, 289
- Springel V., Di Matteo T., Hernquist L., 2005a, *MNRAS*, 361, 776
- Springel V., et al., 2005b, *Nature*, 435, 629
- Steinborn L. K., Dolag K., Hirschmann M., Prieto M. A., Remus R.-S., 2015, *MNRAS*, 448, 1504
- Tabatabaei F. S., Minguez P., Prieto M. A., Fernández-Ontiveros J. A., 2018, *Nature Astronomy*, 2, 83
- Tamura Y., Nakanishi K., Kohno K., Kawabe R., 2007, in Combes F., Palouš J., eds, *IAU Symposium Vol. 235, Galaxy Evolution across the Hubble Time*. pp 430–430, doi:10.1017/S174392130601057X
- Telfer R. C., Zheng W., Kriss G. A., Davidsen A. F., 2002, *ApJ*, 565, 773
- Teyssier R., 2002, *A & A*, 385, 337
- Thomas D., Maraston C., Bender R., Mendes de Oliveira C., 2005, *ApJ*, 621, 673
- Toomre A., 1964, *ApJ*, 139, 1217
- Tumlinson J., Peebles M. S., Werk J. K., 2017, *Annual Review of Astronomy and Astrophysics*, 55, 389

- Urry C. M., Padovani P., 1995, *Publications of the Astronomical Society of the Pacific*, 107, 803
- Valiante R., Agarwal B., Habouzit M., Pezzulli E., 2017, *Publications of the Astronomical Society of Australia*, 34, e031
- Vogelsberger M., Sijacki D., Kereš D., Springel V., Hernquist L., 2012, *MNRAS*, 425, 3024
- Vogelsberger M., et al., 2014, *MNRAS*, 444, 1518
- Voit G. M., Kay S. T., Bryan G. L., 2005, *MNRAS*, 364, 909
- Wada K., 2001, *ApJ Letters*, 559, L41
- Wada K., Schartmann M., Meijerink R., 2016, *ApJ Letters*, 828, L19
- Walch S., et al., 2015, *MNRAS*, 454, 238
- Weinberger R., et al., 2017, *MNRAS*, 465, 3291
- Werk J. K., Prochaska J. X., Thom C., Tumlinson J., Tripp T. M., O’Meara J. M., Peeples M. S., 2013, *ApJ Supplement Series*, 204, 17
- Wiklind T., Combes F., Henkel C., 1995, *A & A*, 297, 643
- Williams R. J., Quadri R. F., Franx M., van Dokkum P., Toft S., Kriek M., Labbé I., 2010, *ApJ*, 713, 738
- Willman B., Strader J., 2012, *Astronomical Journal*, 144, 76
- Woosley S. E., Weaver T. A., 1995, *ApJ Supplement Series*, 101, 181
- Wylezalek D., Zakamska N. L., 2016, *MNRAS*, 461, 3724
- Young L. M., et al., 2011, *MNRAS*, 414, 940
- Young L. M., et al., 2014, *MNRAS*, 444, 3408
- Zubovas K., Nayakshin S., King A., Wilkinson M., 2013, *MNRAS*, 433, 3079
- de Zeeuw P. T., et al., 2002, *MNRAS*, 329, 513
- van der Tak F. F. S., Weiß A., Liu L., Güsten R., 2016, *A & A*, 593, A43

Acknowledgements

First, I want to thank my supervisor Thorsten Naab. His guidance and support—from when I started working in this field during my Masters throughout all of my PhD until now, when I am about to finish it—helped me navigate the multitude of problems (both physical and technical) that make up numerical astrophysics. His advice kept me focussed on the important ones.

I also thank the people in my PhD committee—Simon White, Eiichiro Komatsu, and Benedetta Ciardi—for their guidance and for reminding me that things not working out like we would want them to can still lead to important results. I thank my collaborators, Ena Choi, Jerry Ostriker, and Eric Emsellem for their help and advice with the paper that became the basis of chapter 2.

I am grateful to Chia-Yu Hu, Bernhard Röttgers, Andreas Schmidt, and Ludwig Oser, whose insight in the various numerical codes and models I used in my work was very helpful. In particular, I thank Chia-Yu for implementing many integral parts of the code I worked with and explaining them to me, Bernhard for creating and letting me use and modify the analysis suite Pygad, and Ludwig and Andreas for their general help with coding and deciphering GADGET.

Furthermore, I wish to thank Ricard Ardevol, Miranda Jarvis, Matteo Frigo, Andreas Schmidt, Bernhard Röttgers, Philipp Busch, Jens Stücker, Anabele Pardi, Aoife Boyle, Vlas Sokolov, Timo Halbesma, Tim-Eric Rathjen, Francesco Rizzuto, Magnus Lindström, and all my other amazing colleagues—many of whom became my friends—for interesting, useful, and entertaining discussions, as well as occasional distractions from the routine of scientific work. Special thanks goes out to Ricard, for being a great flatmate beside being a great friend.

Last but not least, I thank my family, whose continued, unwavering support made my entire academic career possible in the first place.