Hydrodynamic Models of AGN Feedback in Cooling Core Clusters

John C. Vernaleo



Advisor: Christopher S. Reynolds

May 2008

Abstract

Title of Dissertation:	Hydrodynamic Models of AGN Feedback in Cooling Core Clusters
	John C. Vernaleo, Doctor of Philosophy, 2008
Dissertation directed by:	Professor Christopher S. Reynolds Department of Astronomy

X-ray observations show that the Intra Cluster Medium (ICM) in many galaxy clusters is cooling at a rapid rate, often to the point that it should have radiated away all of its energy in less than the age of the cluster. There is however a very clear lack of enough cool end products of this gas in the centers of the clusters. Energetic arguments indicate that Active Galactic Nuclei (AGN) should be capable of heating the inner regions of clusters enough to offset the radiative cooling; truncating massive galaxy formation and solving the cooling flow problem.

We present three sets of high resolution, ideal hydrodynamic simulations with the ZEUS code to test this AGN heating paradigm. For the first set of simulations, we study the dependence of the interaction between the AGN jets and the ICM on the parameters of the jets themselves. We present a parameter survey of two-dimensional (axisymmetric) models of back-to-back jets injected into a cluster atmosphere. We follow the passive evolution of the resulting structures. These simulations fall into roughly two classes, cocoon-bounded and non-cocoon bounded. We find that the cocoon-bounded sources inject significantly more entropy into the core regions of the ICM atmosphere, even though the efficiency with which the energy is thermalized is independent of the morphological class. In all cases, a large fraction of the energy injected by the jet ends up as gravitational potential energy due to the expansion of the atmosphere. For the second set, we present three-dimensional simulations of jetted AGN that act in response to cooling-mediated accretion of an ICM atmosphere. We find that our models are incapable of producing a long term balance of heating and cooling; catastrophic cooling can be delayed by the jet action but inevitably takes hold. At the heart of the failure of these models is the formation of a low density channel through which the jet can freely flow, carrying its energy out of the cooling core.

Finally, we present a set of simulations with both feedback and precessing jets. The addition of jet precession is not sufficient to couple the jets to the ICM energetically although it can deposit a large amount of energy in sound waves. These sound waves are lost to the system in ideal hydrodynamics, but ultimately may provide a powerful heating mechanism for clusters cores by AGN when additional physical effects are taken into account.

Hydrodynamic Models of AGN Feedback in Cooling Core Clusters

by

John C. Vernaleo

Dissertation submitted to the Faculty of the Graduate School of the University of Maryland at College Park in partial fulfillment of the requirements for the degree of Doctor of Philosophy 2008

Advisory Committee:

Professor Christopher S. Reynolds, chair Professor William Dorland Professor M. Coleman Miller Professor Daniel Proga Professor Derek C. Richardson Professor Massimo Ricotti

© John C. Vernaleo 2008

ii

Preface

Portions of the research described in this work have been published elsewhere. Part of Chapter 2 appeared in a modified form in Vernaleo and Reynolds (2006a). Chapter 3 was published in a modified form in Vernaleo and Reynolds (2007b). Chapter 4 was published in a modified form in Vernaleo and Reynolds (2006a). Preliminary versions of the material in Chapter 5 were presented at "Heating vs. Cooling in Galaxies and Clusters of Galaxies" in Garching, Germany, August 2006 (Vernaleo and Reynolds 2007a), the 209th meeting of the AAS, January 2007 (Vernaleo and Reynolds 2006b), and "Galaxy and Black Hole Evolution: Towards a Unified View" in Tucson, Arizona, November 2007. For every author, musician, or philosopher that ever inspired me. Sometimes it is too easy to forget that science and creativity shouldn't stay too far apart.

Acknowledgements

First, I would like to thank my advisor, Chris Reynolds. He was able to provide me funding for a long time so I could concentrate on my work, sent me to some great conferences, and provided an incalculable amount of scientific and writing advice. Chris also gave me enough freedom to proceed with things the way I wanted to, never forcing me to use his preferred set of tools for things. This occasionally led to a duplication of effort on my part, but I learned a lot more this way than I could have any other way and I appreciate it.

Derek Richardson provided useful computer advice throughout, despite the fact that he neither uses FORTRAN nor works on clusters. Cole Miller provided many useful discussions. The computer staff, Bill Sebok, Mark Wolfire, and John Ohlmacher have been immensely helpful, managing to keep our computers working despite my best efforts.

There is no way I could have finished without many of my fellow graduate students. Some provided useful advice and help. Some were just friends and good to complain to. Either way, I hesitate to name anyone since I will certainly leave someone out. This is even worse since too many of my friends have already graduated. As a first order attempt at it, I would include, Mia Bovill, Nicholas Chapman, Kayhan Gültekin, Matthew Knight, Mike Koss, Rachel Kuzio de Naray, Randall Perrine, Rob Piontek, Rahul Shetty, and Kevin Walsh.

I would like to thank the developers of ZEUS-3D and the original developers of ZEUS-MP at NCSA for providing such high quality codes to work with. Without such public codes, none of this work could have been possible.

Portions of this work were funded by the Chandra Theory and Modeling Program under grants TM4-5007X and TM7-8009X and National Science Foundation under grant AST0205990.

I would like to thank my parents for their support throughout all of my schooling (and something like 26 years of school is a lot to ask for). And my brother Dennis for providing a constant source of distractions.

And lastly, of course, my wife, Beth, for going through grad. school with me. I'm not sure if I could imagine grad. school without her, but I know I wouldn't want to.

Contents

Li	List of Tables ix		
\mathbf{Li}	st of	Figures	x
1	Intr	coduction	1
	1.1	The Cooling Flow Problem and Clusters of Galaxies	2
	1.2	The AGN Feedback Paradigm	5
	1.3	Previous Numerical Models	7
	1.4	Persistent Questions	10
	1.5	This Work	11
2	The	e ZEUS Code	13
	2.1	Introduction	13
	2.2	Basics of Grid Based Hydrodynamics	13
	2.3	Methods used in ZEUS	15
	2.4	ZEUS-MP	16
		2.4.1 Code Archaeology \ldots	18
		2.4.2 Modifications	19
		2.4.3 The ZEUS-MP Build/Setup Process	22
		2.4.4 Performance \ldots	25
	2.5	Conclusions	28
3	Ene	ergetic Impact of Jet-Inflated Cocoons in Relaxed Galaxy Clus-	
ters: Two-Dimensional Models		: Two-Dimensional Models	31
	3.1	Introduction	31
	3.2	Simulation Setup	32
		3.2.1 Modeling the Jets	34
		3.2.2 Comparison to Real Units	38
	3.3	Analysis and Results	40
		3.3.1 Morphological Classification	41
		3.3.2 Mass Distribution and Energetics	49
		3.3.3 Entropy Evolution and Thermalization Efficiencies	53

	3.4	Discussion
	3.5	Conclusion
4	The	e Failure of Simple Feedback Models 63
	4.1	Introduction
	4.2	Basic Setup
	4.3	Specific Models
		4.3.1 Radiative Cooling
		4.3.2 Single AGN Outburst Models
		4.3.3 Feedback Models
		4.3.4 Delayed Feedback Models
		4.3.5 Feedback with Rotation
	4.4	Discussion
	4.5	Conclusion
5	Pre	cessing Jet Models 93
-	5.1	Introduction
	5.2	Setup
	5.3	Analysis and Results
	0.0	5.3.1 Basic Properties 98
		5.3.2 The Cooling Flow 90
		5.3.3 Simulated X-ray Maps 101
		5.3.4 Tracing the Energy 102
	5.4	Discussion 106
	5.5	Conclusion 119
	0.0	
6	Con	clusions and Future Directions 113
	6.1	Conclusions
	6.2	Future Directions
Bibliography 118		

List of Tables

3.1	Parameters for 2D Simulations	39
4.1	List of Feedback Simulations	66
5.1	List of Precessing Simulations	95

List of Figures

1.1	Cluster cooling times
1.2	Perseus Cluster
1.3	Bubbles in viscous background
2.1	Energy conservation
2.2	KH test
2.3	KH test, 3D
2.4	ZEUS-MP run time benchmarks
2.5	ZEUS-MP speed-up tests
3.1	Close up of active jet density
3.2	Close up of active jet entropy 36
3.3	Close up of jet density
3.4	Close up of jet entropy
3.5	Cocoon bounded source, entropy
3.6	Cocoon bounded source, density
3.7	Non-cocoon bounded source, entropy 44
3.8	Non-cocoon bounded source, density 45
3.9	Unresolved source, entropy
3.10	Unresolved source, density
3.11	Density contours for cocoon
3.12	Density contours for non-cocoon
3.13	Morphology vs. jet parameters
3.14	Morphology vs. jet parameters, density-velocity plane
3.15	M and E vs. time for cocoon
3.16	M and E vs. time for non-cocoon
3.17	Change in entropy
3.18	Efficiency of internal energy change
3.19	Efficiency of kinetic energy change
3.20	Efficiency of potential energy change
4.1	Mass accretion with pure cooling

4.2	Radial temperature dependence for pure cooling
4.3	Entropy plots for single jet
4.4	Density plots for single jet
4.5	Close up of active jet density
4.6	Close up of active jet entropy
4.7	Mass accretion with single burst
4.8	Radial temperature dependence, run B
4.9	Radial temperature dependence, run C
4.10	Mass accretion for feedback
4.11	Channel formation $\ldots \ldots 80$
4.12	Mass accretion with low- η delayed feedback $\ldots \ldots \ldots$
4.13	Mass accretion with low- η long-delayed feedback
4.14	Mass accretion with very high- η feedback
4.15	Mass accretion with rotating cluster
4.16	Dense, cold disk
5.1	Geometry for precessing jet
5.2	Density and pressure slices for precessing jet
5.3	Mass accretion for precessing jet
5.4	Mass accretion for non-precessing jet
5.5	X-ray Brightness, run I
5.6	X-ray Brightness, run II
5.7	Total energy change for fixed jet
5.8	Total energy change for precessing jet
5.9	Run I, ΔE
5.10	Run II, ΔE
5.11	Projected X-ray brightness

Chapter 1

Introduction

Supermassive back holes have a major impact on the universe, often far beyond the range of their gravitational influence. This appears to be especially true in the context of both structure formation and galaxy formation. The cooling flow problem in galaxy clusters is a possible manifestation of this very effect. In this dissertation, we study, through numerical models, the ability of supermassive black holes, in the guise of jetted Active Galactic Nuclei (AGN), to solve the cooling flow problem, and therefore to help regulate structure formation and galaxy formation.

Observational evidence from X-ray studies shows that the Intra Cluster Medium (ICM) is cooling at a rate that would be unsustainable over cosmological time frames in many clusters. Left unchecked, this would drastically alter the current picture of structure formation and the formation of massive galaxies. It is already known that AGN have an impact on the ICM, but the details of how they alter the overall energetics (and therefore the inferred cooling flow) are harder to pin down. Our simulations will attempt to address this. In the following sections we will review the cooling flow problem and the observations that lead to it (Section 1.1), the potential role of AGN in cluster evolution (Section 1.2), previous numerical work on AGN in cooling flow clusters (Section 1.3), and the remaining questions in this growing picture (Section 1.4).

1.1 The Cooling Flow Problem and Clusters of Galaxies

From very early observation of galaxy clusters, it has been clear that galaxies do not contain most of the mass in a galaxy cluster (Zwicky 1937). In the current standard cosmology (Spergel et al. 2007), clusters are assumed to be around 80% dark matter, 14% ICM, and the rest galaxies and their components. This means that the ICM contains most of the baryons in a galaxy cluster by a large margin.

The ICM has been studied extensively by every X-ray observatory from the beginning of X-ray astronomy and continuing to the modern satellites. Initially observed by Felten et al. (1966) in the Coma cluster and by the *Uhura* X-ray observatory in the Virgo, Coma, and Perseus clusters (Gursky et al. 1972), the ICM is bright in X-rays with a temperature of a few keV. This is thermal emission from an optically-thin, collisionally ionized plasma.

From these X-ray observations, we are able to calculate the temperature and density of the gas. Combined with a known cooling function, Λ , these can be used to calculate the cooling time,

$$t_{cool} = \frac{\frac{5}{2}nkT}{n^2\Lambda},\tag{1.1}$$

of the gas. There is an extra nkT in the numerator of Equation 1.1 (beyond the usual internal energy density $\frac{3}{2}nkT$) because the gas is assumed to cool isobarically. Figure 1.1 show the cooling time vs. radius for a set of relaxed clusters. Many of these times are at least less than the age of the cluster, and certainly less than a Hubble time in about 70% of clusters (Edge et al. 1992; Peres et al. 1998; Sanderson et al. 2006; Vikhlinin et al. 2006). In the inner regions of the cluster, these times



Figure 1.1: Calculated cooling times vs. radius for a set of relaxed clusters observed with the *Chandra* observatory. Originally from Peterson and Fabian (2006).

can be shorter than one billion years.

This means that for a sizable fraction of clusters, in the absence of some heating, the inner regions should have had enough time to radiate away all of their thermal energy. In reality, gas clearly cannot radiate away all of its thermal energy. Instead, the simple interpretation is that the gas should cool and lose pressure support. In a relaxed cluster, as the gas loses pressure support, it will flow in towards the center, becoming denser (with hot gas from further out in the cluster replacing it). Due to the strong density dependence of the thermal bremsstrahlung radiation, this should cause the gas to radiate its energy away even faster. This inferred inward flow of material is know as a cooling flow.

Observationally, there should be clear evidence for this. In many clusters inferred cooling rates of hundreds of solar masses per year are possible (see Fabian (1994) for a compilation of these numbers for a large sample of clusters). Recent estimates lower that value for many of the clusters (for example Sanders et al. (2008)), but even tens of solar masses per year should be easily observable. As it turns out, there is both an absence of evidence for large quantities of cool gas and actual evidence for the absence of such gas. This lack of observed cool gas is the *Cooling Flow Problem* (Fabian 1994).

Star formation is seen in the central giant elliptical (cD) galaxy (Hicks and Mushotzky 2005; O'Dea et al. 2004), but it is not sufficient to account for the cooling gas. Not only are there not enough cool remnants such as stars (and in these circumstances, stars count as cool remnants), but based on *XMM-Newton* dispersive spectroscopy, there does not appear to be any gas below about $T_{virial}/3$ (Peterson et al. 2001; Tamura et al. 2001). Beyond just the stellar component, there is a deficit of galaxies vs. dark matter halos on both the high and low end of the luminosity function. The low end (which is not relevant for central cluster galaxies) is explained by star formation feedback (Dekel and Silk 1986; Larson 1974). For the high end mass deficit, star formation feedback is not efficient enough to explain the cutoff (Benson et al. 2003).

The standard framework for galaxy formation states that baryonic matter falls into the developing dark matter halos and the resulting accretion shocks raises it to approximately the virial temperature of the halo. The "cold" baryonic components of the galaxy then form via radiative cooling of this shocked gas. In the absence of any feedback processes (i.e., with gravitational collapse followed by radiative cooling alone), the galaxy mass function would have to essentially follow the dark matter halo mass function — in essence, the baryons that are within the turn-around radius of the developing dark matter halo are trapped and fated to eventually form the baryonic galaxy at the halo's center. The fact that the cD galaxies do not follow the dark matter halo of the cluster and have a size cutoff requires some form of feedback to explain.

1.2 The AGN Feedback Paradigm

It has long been known that a large fraction of galaxy clusters host a radio loud AGN in their center (Burns and Owen 1977). Radio loud AGN (which account for approximately 15% of all AGN (Krolik 1999)) have powerful, radio-emitting jets and are referred to as radio-galaxies. In many clusters of galaxies, the ICM bears the marks of repeated episodes of AGN jet activity. First seen in *Einstein* and *Röntgensatellit* observations (Böhringer et al. 1995, 1993; Carilli et al. 1994; Feigelson et al. 1987; Heinz et al. 1998), *Chandra* has studied numerous examples of ICM bubbles (Blanton et al. 2001; Fabian et al. 2000; McNamara et al. 2000; Young et al. 2002), ghost bubbles (Choi et al. 2004; Fabian et al. 2000; Heinz et al. 2002; McNamara et al. 2001), ripples (Fabian et al. 2003, 2005b), shells (Fabian et al. 2000), and filaments (Hatch et al. 2005) that are clearly associated with a central radio-galaxy.

In the accepted picture, AGN are the large scale manifestation of accreting supermassive black holes in the center of galaxies. Although the exact details of jet launching are still uncertain, the general idea that material flows onto the supermassive black hole from an accretion disk and back-to-back jets are launched from the poles of the system is well supported. These relativistic jets (and possibly any wideangle winds that also result as a byproduct of the accretion) allow the black hole to influence its environment on scales much larger than possible by gravitational influence alone. One additional component to this is that a given supermassive black hole is not always active. Most galaxies are thought to have a black hole in the center (and if not, certainly nearly every cD galaxy does), yet all galaxies are not AGN. Therefore AGN must have some duty cycle, where accretion is only sufficient to cause the launching of AGN jets for parts of the black hole's lifetime.

As discussed above, jetted AGN can have dramatic effects on the ICM of a cluster. AGN are among the most energetic objects in the universe, and they should release enough energy to easily offset losses from radiative cooling in the cluster (although the sample of Bîrzan et al. (2004) shows that they are close to the line for having sufficient energy to balance cooling and not well over the line as initially thought). They are also a good candidate for offsetting the radiative losses that cause a cooling flow since they are often situated in the center of clusters, right where radiative losses present the largest problem. The AGN feedback model states that as the ICM cools and accretes onto the central object, it eventually fuels the central AGN. The action of the central AGN then heats the ICM gas, offsetting cooling, and presumable halting the cooling flow, and starving the fueling mechanism of the AGN.

An especially good example of how this should work is seen in the deep *Chandra* observation of Perseus (Fabian et al. 2005b) seen in Figure 1.2. In this observation, there are several sets of radio bubbles of varying ages, indicating separate periods of AGN activity. The inflating of the bubbles is one way for the AGN jets to heat the system. There are ripples, sound waves, and weak shock waves in the image as well. These all provide opportunities for the AGN to heat the ICM through the jets.



Figure 1.2: Some of the clearest examples of bubbles, ripples, and complex interactions between an AGN and the ICM can be seen in this *Chandra* observation of the Perseus Cluster from Fabian et al. (2005b). Figure taken from Peterson and Fabian (2006).

1.3 Previous Numerical Models

A large amount of theoretical work has been performed on the effects of AGN on cooling cluster cores. This has included analytic work such as Begelman and Cioffi (1989) on the formation of cocoons from jets and numerical models. With ever increasing computer power, most of the recent work has focused on hydrodynamic models of the AGN/ICM interaction.

There is a large body of numerical work that models various aspects of these interactions. At the first level of complexity are the spherically symmetric models. These include Ciotti and Ostriker (2001); Conroy and Ostriker (2007); and the double heating model of Ruszkowski and Begelman (2002) which included heating by thermal conduction from the outer regions of the cluster. Generally, the spherically symmetric models have been fairly successful at balancing cooling, but they are not realistic enough to fully explain AGN heating. Jets are inherently not spherically symmetric. In one dimension, the only relevant issue is if the total energy can balance, not the spatial distribution of the energy. Additionally, the growth rates of the hydrodynamical instabilities are not the same in one dimension as they are in three dimensions. Due to this, one dimensional simulations can produce stable structures that would be Rayleigh-Taylor (RT) unstable in a multi-dimensional simulation (such as a low-density bubble suspended in a dense medium).

Early hydrodynamic models of the putative jet-cocoon/shock structure in Cygnus-A were presented by Clarke et al. (1997). More recent numerical investigations have explored the buoyant evolution of the cocoon after its supersonic expansion phase (Basson and Alexander 2003; Brüggen and Kaiser 2001, 2002; Churazov et al. 2001; Dalla Vecchia et al. 2004; Omma and Binney 2004; Omma et al. 2004; Reynolds et al. 2002; Robinson et al. 2004; Zanni et al. 2005). Various extensions to ideal hydrodynamics have also been studied including the effects of magnetic fields (Jones and De Young 2005; Robinson et al. 2004), plasma transport processes (Ruszkowski et al. 2004), AGN convection (Chandran and Rasera 2007), and realistic (cosmological) background motions (Heinz et al. 2006)

Reynolds et al. (2005b) performed simulations of rising bubbles in a viscous background. Figure 1.3 shows that the addition of physics beyond ideal hydrodynamics





(a) Inviscid Bubble

(b) Viscous Bubble

Figure 1.3: Synthetic X-ray surface brightness maps of the late stages of the buoyant evolution of a bubble in cases both with and without viscosity from Reynolds et al. (2005b). The introduction of viscosity keeps the bubble intact for longer than it would otherwise.

can alter the evolution of a bubble in a cluster. Without viscosity, the bubbles which start at the center of a simulated cluster are shredded by instabilities as they rise. With the addition of viscosity, the bubbles maintain their integrity for much longer times. This is potentially important as the Perseus observations discussed earlier show evidence for multiple generations of bubbles which implies long-lived bubbles.

Much of the numerical work discussed above focused on isotropically inflated bubbles that start in the center of a cluster and buoyantly rise. The problem with these bubble models is that they require making some strong assumptions about the jet/ICM interactions which we ultimately do not believe are warranted. For starters, inflating bubbles does work on the system. Also, to get a perfect spherical bubble assumes that no forces work to shred the bubble during inflation and that inflation happens at a timescale much shorter than the buoyant rise time. Most importantly, an isotropically inflated bubble rules out any shocks or sound-waves that may be generated by the interaction of the jet and the ICM. For these reasons, we believe that models including the jet/ICM interaction are necessary to fully capture the important behavior of these systems. Reynolds et al. (2002) performed two-dimensional, axisymmetric simulations of back-to-back jets inflating cocoons in a simulated cluster. By including the jet inflation phase, this allowed for a more realistic late stage evolution of the bubbles. These simulations served as the starting point for the simulations presented here in Chapter 3.

1.4 Persistent Questions

Despite the attractiveness of the AGN feedback picture, there are still some open issues. Although AGN are energetic enough to balance the cooling, the exact mechanism for getting energy from the AGN into the ICM is not obvious. The expansion of a cocoon (and displacement of ICM gas) can do pdV work and provide heating. Shocks provide an opportunity to heat the ICM. It is not immediately clear if these bulk hydrodynamic processes are enough to balance cooling losses in the ICM. Besides the amount of heating, the location of heating is also important. The core regions of clusters are the area with the need for additional heating. Radio lobes can extend for nearly a megaparsec. Heating at that distance can do nothing to balance the cooling flow. For this reason it is important to see not only how efficiently a given process can heat the ICM, but how quickly it can heat it. The dissipation of waves excited by the jets provide other heating opportunities, but these require physics beyond ideal hydrodynamics. Before looking to more complex plasma physics (which may be necessary) we must fully understand the ability of heating by jets when limited to ideal hydrodynamics.

Even assuming the ICM can be efficiently heated by a jetted AGN, there is the issue of regulation. A true feedback mechanism is needed. While the AGN must be able to heat the cluster gas, it must not over heat the gas. The supermassive black hole powering the AGN must somehow "know" about the state of the ICM gas and power the AGN based on this. An additional constraint on this is that whatever sets the feedback must work on clusters of a variety of scales (both spatial and temperature). So, fine tuning will not be possible if a single form of feedback is to work across all cooling flow clusters. One could argue that different feedback mechanisms work on different size or temperature clusters. Since clusters evolve over cosmological times, this would lead to the odd situation where different feedback mechanisms operate on the same cluster at different times. The simplest explanation seems to be that there is a single way to couple the temperature of the core ICM gas to a central supermassive black hole (or its accretion disk).

1.5 This Work

In this dissertation, we will examine several sets of hydrodynamic models of jetted AGN embedded in cooling core clusters to evaluate their role in heating the cluster center as well as the potential for a closed feedback loop between the cooling gas and the jets from the central AGN. In Chapter 2, we will discuss the hydrodynamic code that we employ for all of our simulations; ZEUS. This will include a general review of the algorithms used in the ZEUS family of codes and a specific discussion of the ZEUS-MP variant.

Chapter 3 will present a set of axisymmetric models of jets in clusters. Due to the

modest CPU requirements for axisymmetric simulations, a large number of models will be presented to allow for an examination of the parameter space available for AGN jets. This will provide our most detailed look at the interactive between the AGN jet and the ICM. Chapter 4 will describe a set of three dimensional models including a feedback mechanism between the jet power and the state of the ICM gas. This will allow us to evaluate the potential of jetted AGN to heat the center of a cluster with an attempt at a closed feedback loop. Chapter 5 will expand on the previous models to include jet precession. We will examine the impact of these jets on the cluster temperature structure and show the differences between precessing jets and fixed jet models. Finally the conclusions in Chapter 6 will summarize the results as well as discuss some further possibilities for balancing ICM cooling with jetted AGN. We will discuss other physical effects which need to be included to ultimately determine the role of jetted AGN in cooling core clusters.

Chapter 2

The ZEUS Code

2.1 Introduction

In this chapter we will discuss the primary tool used in this work; the ZEUS family of hydro-codes. As a code base with a long history in the astrophysics community, there are many versions of the code. Simulations in Chapter 3 used the ZEUS-3D code. Simulations in Chapters 4 and 5 were done using the 1.5 branch of the ZEUS-MP code. We will discuss general features of the code and algorithms which are held in common amongst all version as well as details specific to ZEUS-MP v1.5, the version of the code modified and distributed specifically for this work. A sample use case as well as performance details will also be given for ZEUS-MP.

2.2 Basics of Grid Based Hydrodynamics

The standard equations of hydrodynamics are

$$\frac{D\rho}{Dt} + \rho \nabla \cdot v = 0, \qquad (2.1)$$

$$\rho \frac{Dv}{Dt} = -\nabla P - \rho \nabla \Phi, \qquad (2.2)$$

$$\rho \frac{D}{Dt} \left(\frac{e}{\rho} \right) = -P\nabla \cdot v - \Lambda, \qquad (2.3)$$

where

$$\frac{D}{Dt} \equiv \frac{\partial}{\partial t} + v \cdot \nabla.$$
(2.4)

In Equation 2.2, Λ represents a cooling term. The specific cooling term we use is be given in Equation 4.1 and will be discussed in Section 4.3.1.

For astrophysical cases, the gravitational potential given by the Poisson equation,

$$\nabla^2 \Phi = 4\pi G\rho, \tag{2.5}$$

may also be needed. Equation 2.1 represents conservation of mass, Equation 2.2 provides conservation of momentum, and in Equation 2.3 we see the conservation of internal energy. The system of equations above (Equations 2.1 to 2.3) is closed by the equation of state,

$$P = (\gamma - 1)e, \tag{2.6}$$

with $\gamma = 5/3$ (the value for an ideal gas) in our simulations. As long as collisions between particles are frequent enough that we can treat the material as continuous (which is generally the case for astrophysical plasmas), we can use these equations as a basis to model the system (although if other effects such as magnetic fields are present they must be modified).

In Eulerian hydrodynamics, the problem domain is split into a grid and the fluid elements are then advanced on this grid. This is in contrast to Smoothed Particle Hydrodynamics (SPH) which use a Lagrangian method where the coordinates follow the fluid. In ZEUS (a Eulerian code), the grid is fixed (although it does not have to be uniform). This has the advantage of simplicity over Adaptive Mesh Refinement (AMR) codes where the grid can be adjusted to improve resolution in areas of interest. In ZEUS, there are actually two separate grids used, an a and a b grid, where x1b(i) = x1a(i + 1/2). This makes it easier to deal with quantities that are more naturally defined on the face of a cell (the a grid) such as velocity and quantities that are more naturally defined on the center of a cell (the b grid) like density and internal energy. There is also a buffer of two grid cells (called ghost zones) surrounding the edge of the computational grid to allow for second order differencing of values even at the grid boundaries.

Starting with the fluid variables on this fixed grid, ZEUS solves the equations explicitly using an operator split, finite difference method. An operator split method solves the Partial Differential equation,

$$\frac{\partial x}{\partial t} = L(x) \tag{2.7}$$

as the sum of the finite difference equations $L(x) = L_1(x) + L_2(x) + \dots$ The solution in ZEUS is further broken into a source step and a transport step. The source step takes account any source or sink terms in the Euler equations (e.g., pressure gradients, gravitational forces, and viscous stresses). The transport step then solves for the advection of the fluid variables. This splitting has the useful side effect that if a new physical process can be formulated as a single source term, it can be added on to the end of the source step in a very simple manner.

2.3 Methods used in ZEUS

In the standard mode of operation (using van Leer advection as we do), ZEUS is second-order accurate. Artificial viscosity is used to smooth shocks onto a finite number of cells. All simulations were run as ideal hydrodynamic cases, neglecting magnetic fields and self gravity. ZEUS was originally described in Stone and Norman (1992a,b) and these remain the authoritative sources about the code and algorithms. As we have not used Magnetohydrodynamics (MHD) in this work, we will only discuss the hydrodynamic methods and leave the discuss of the MHD algorithm to the above references.

Stability is determined by the usual Courant-Friedrichs-Lewy (CFL) condition,

$$\Delta t \le \min\left(\frac{\Delta x}{|v_x| + c_s}\right). \tag{2.8}$$

Any other physical processes included (such as radiative cooling) have a limiting time step of their own which must be included. As the jets in our simulations are typically very supersonic, their speed provides the limiting timestep during most of the simulations.

In its standard configuration, ZEUS is not a strictly energy conserving code. We do not, however, believe that this has a major impact on our results. Figure 2.1 shows (for one of our jet/cluster simulations) that during the majority of time during the simulation, total energy is conserved to a sufficient degree for our needs. At very early times, the jet injects hot material onto the grid, which is clearly not conservative. At very late times, material (and sound waves carrying energy) leave the outer boundary of the grid. During the time in between, the energy change remains flat with only minor variations (at the 10^{-5} level), showing that energy is mostly conserved.

2.4 ZEUS-MP

For large, three-dimensional simulations, the single processor ZEUS-3D is not sufficient. To take advantage of multiple processors, we had to use a parallel version of the ZEUS code, ZEUS-MP (Norman 2000). Our version is based on the initial National Center for Supercomputing Applications (NCSA) release, version 1.0b (which



Figure 2.1: Fractional change in total energy within the simulated domain as a function of time. All forms of energy (kinetic, internal, and gravitational) for all of the gas in the computational domain are considered. Shown here is the change in total energy compared with the initial time ΔE divided by the initial total energy E_0 .

is a direct descendant of ZEUS-3D). Our modifications, documentation, and several supporting scripts have been made publicly available at:

http://www.astro.umd.edu/~vernaleo/zeusmp.html

under the same terms as previous public ZEUS releases (very similar to a BSD-style license but with a non-commercial use clause).

ZEUS-MP uses the Message Passing Interface (MPI) library for all parallel operations. In parallel mode, the main grid is broken up into several blocks. Each processor (or process, although putting more than one process on a single processor or core is terribly inefficient and only useful for testing) is given one block. Although there is a master process (myid = 0 in MPI terms) which takes care of more of the I/O than the other processes, each process must construct its own grid (and the overall block topology) based on the startup files. Each process then checks for any special conditions (stops, output dumps, restart dumps, etc.) and then advances the grid variables by a single timestep. All processes share a single timestep which is the smallest of the individual block timesteps. After advancing, each block uses an MPI_Send command to send the values of the border (ghost) cells to its neighbors. Due to this type of communication, ZEUS-MP is efficient and mostly CPU limited as long as $N_{inner cells} >> N_{border cells}$ (basically low surface area to volume in terms of number of cells, not size of cells). The timesteps are then combined for all blocks and the next step is taken.

We also note that another group released a Fortran 90 version 2.0 of ZEUS-MP (Hayes et al. 2006) at the same time as our initial release. Although the two projects are not connected and we are not involved with version 2.0, due to the common ancestor, bug fixes have been shared between the two versions of ZEUS-MP.

2.4.1 Code Archaeology

Our work with ZEUS illustrates a common occurrence in academic computing; code or software archaeology. Although this is likely common in commercial areas as well, the lack of standard practices, the frequent non-continuous use of various codes, the lack of version control, and the lack of a focus on documentation makes this especially common and problematic in scientific computing.

In the case of the ZEUS code, the algorithms are extremely well described in the literature, and there is some scattered documentation which does not always match the versions of the code available. Even when dealing with only local users of a code, a large amount of knowledge is passed from researcher to researcher and is not committed to any permanent documentation. This is a dangerous and inefficient
situation which we have tried to avoid in our modified release of ZEUS-MP.

Along with a need for documentation, version control software (CVS, Subversion, etc.) is a must. This is not just for the standard reasons (code safety, development ease, etc.). An available revision control (even in the crude form of numbered or dated releases) allows comparisons to be made between versions. Reconstructing how a previous author (or even the current author in the past) got to a certain point is sometimes the best way to understand how something works. When combined with documentation (both in the form of comments in the code and actual manuals), this can greatly increase the productivity of new users of long-lived codes such as ZEUS.

2.4.2 Modifications

Here, we describe our changes and modifications to ZEUS-MP. First, the code had to be ported to compile with a current FORTRAN 77 compiler. As different compilers implement the standard (and the extensions to it) differently, we had to choose certain compilers as our targets. The Intel Compiler was used for the speed of the executables it is supposed to produce for x86 compatible machines. For the sake of portability, we also maintain compatibility with the GNU compilers. None of these changes modify the behavior of the code. Primarily this involved removing multiply defined variables and cleaning up the namelist routines and the namelists themselves. Also, all filenames that differ only by case (common in the FORTRAN 77 build process) were changed to allow building on non-case-sensitive filesystems. A large amount of vestigial debugging and profiling code was removed and replaced with a smaller and more manageable amount of debugging code. Similarly, several incomplete routines that could not be called in their current form but were not necessary for any standard runs were removed. The overall philosophy was to begin to move from a catch-all of source code with a well tested base to a more manageable sized source code with only well tested add-ons to the already well tested base.

To allow for long runs, the restart routines (which did not work for parallel simulations) were completely replaced. The new restart routines work for parallel simulations and write to alternating files to save disk space. A wrapper script written in Perl is provided to correctly pick the most recent restart dump (if present) and start ZEUS-MP using that dump. The wrapper script also performs some basics checks of the integrity of the restart dumps before using them as incomplete or missing restart dumps have proven far more likely in parallel simulations than in single processor simulations.

Despite the fact that ZEUS-MP is written nearly entirely in FORTRAN 77, the C preprocessor is used heavily to allow for conditional inclusion of code. Some portions of this preprocessor code needed to be fixed; many preprocessor directives were not properly matched up or were improperly nested, resulting in preprocessor flags that only worked for one of the two possible values. Also, some preprocessor directives did not enclose all of the code relevant to a given option. This made it impossible to completely turn certain options off in the code. To facilitate portability, preprocessor directives were initially included for a large number a machines and architectures. Due to the near total dominance of UNIX-like systems (GNU/Linux and Mac OS X primarily) in high performance and scientific computing, these system specific directives were no longer needed and were removed for all extinct operating systems and platforms. The post-processor was modified to be useful for large numbers of output files and will eventually be completely replaced with a more general postprocessor.

Several new problem specific routines were added. A routine to update boundary values during a run was added. A number of other changes and bug-fixes were also made, primarily involving geometry specific bugs. Also, the build process was updated, relying on improved makefiles and a custom Perl script.

To ensure that ZEUS-MP is portable, we have tested and run benchmarks on several different systems. We are aided by the fact that ZEUS uses NCSA's hdf4 format as its primary output format. By using a portable output format, we do not have to worry about endian issues with our output files. The one place we break this is in the restart dumps which are produced as unformatted FORTRAN binary data. We have run and used ZEUS-MP on a variety of GNU/Linux distributions (both 2.4.x and 2.6.x kernels) on both workstations and Beowulf clusters. AMD Athlon processors, AMD Opteron processors (in 32-bit mode only for benchmarks although tests and science runs have also been successfully done in pure 64-bit mode), and Intel Pentium 4 processors have primarily been used. We have also compiled and run ZEUS-MP on Apple's Mac OS X (darwin) on the G4 PowerPC processor and on Intel Core Duo processors. We have also attempted to run ZEUS-MP on a Sun Ultra80 (UltraSPARC II processor) running Solaris 8.0. As we were not able to successfully use any MPI library, we were not able to run ZEUS-MP. This is the one weakness in our portability. If a working MPI library (preferably lam-mpi) and the hdf4 library cannot be compiled, ZEUS-MP cannot run on a system. This should not be a problem on any modern, Unix-like system. We have also had very limited success compiling and running ZEUS-MP with the Cygwin Unix-like environment for Microsoft Windows, although this has not proven stable or well-behaved enough for any serious use.

2.4.3 The ZEUS-MP Build/Setup Process

We will now describe a simple application with ZEUS-MP to illustrate the use of the code and to provide the results of a sample test problem. To simplify coding, handle certain limits of the FORTRAN 77 language, and to ensure maximum performance, the executable, zeusmp.x, is built specifically for each application with a limited set of parameters that can be set at run time instead of compile time. The sample problem we will look at is a simple simulation of the Kelvin-Helmholtz (KH) instability. All files for this test are included in our source code release. This is an instability that arises when fluids of different densities and different velocities share a boundary layer. This is the simplest type of test to illustrate how to setup a problem that does not require adding any additional physics to the code.

First, the preprocessor definitions must be set. This must be done in a file called **src/zeusmp.def**. The file needed for this test is below.

#define	XYZ	
#define	I_ZONES	50
#define	J_ZONES	50
#define	K_ZONES	100
#define	MAX_OF_IN_JN_KN	100
#define	SMALL_NO	1.000d-99
#define	LARGE_NO	1.000d+99
#define	MPI_USED	
#define	NO_CPU_TIMER	
#define	RES_DUMPS	
#define	HDF_DUMPS	
#define	TSL_DUMPS	
#define	VISCOSITY avisc	
#define	PROBLEM kh	
#define	SPECIAL writeout	
#define	RESTART empty	
#define	USERDUMP textdmp	
#define	BNDYUPDATE empty	

The most important thing to note is that the number of zones in this file is for

each processor, not for the entire simulation. The next input file, exe/zmp_inp has the number of cells for the entire simulations (as well as defining the grid spacing and boundaries and contains the MPI layout (mpitop). Values in this file may be changed at run time, but in the case of the grid values, it is limited by the per process values set at compile time in src/zeusmp.def.

```
$mpitop ntiles(1)=2,ntiles(2)=2,ntiles(3)=1,periodic=3*.true. /$
$rescon irestart=0,dtdump=0.1,resfile='resaa000000.000' /$
$pcon nlim=5000000, tlim=20.0, cpulim=1.0E8, mbatch=1 /$
$hycon qcon=2.0,qlin=0.0,dtrat=1.0d-5,courno=0.5,iord=2 /$
$iib niis(1)=4 /$
$oib nois(1)=4 /$
$ijb nijs(1)=4 /$
$ojb nojs(1)=4 /$
$ikb niks(1)=4 /$
$okb noks(1)=4 /$
$ggen1 nbl=100,x1min=-0.5,x1max=0.5,igrid=1,x1rat=1.0,lgrid=.true.
                                                                     /$
$ggen2 nbl=100,x2min=-0.5,x2max=0.5,igrid=1,x2rat=1.0,lgrid=.true. /$
$ggen3 nbl=100,x3min=-0.5,x3max=0.5,igrid=1,x3rat=1.0,lgrid=.true.
                                                                     /$
$grvcon g=1.0e0, ptmass=0e0, x1ptm=0.0, x2ptm=0.0, x3ptm=0.0 /$
$eqos gamma=1.4 /$
$gcon x1fac=0.0, x2fac=0.0, x3fac=0.0, ia=3, ja=3, ka=3, igcon=0 /$
$iocon thdf=0.0, dthdf=0.1, dttsl=0.1 /$
```

As the namelist file, exe/zmp_inp, is one of the least human-readable files in ZEUS and yet contains most of the important user defined settings, we will go through it line by line. The namelist mpitop contains the number of processors to be used in each direction, and if any of them are periodic boundaries. The namelist rescon handles the restart files and is usually only modified by the wrapper script that runs the code. Stopping conditions are set it pcon. The hycon namelist allows for tuning the hydrodynamic solver and should generally not be modified. The next four namelists give the inner (i) and outer (o) boundary conditions for the i, j, and k directions. For this problem, we have used 4 which represents a periodic boundary in all directions. The ggen namelists control the generation of the grid in each direction. This includes the total number of cells (for the entire problem, not per

processor), the first and last value of the grid, and how the grid spacing is supposed to grow (for non-uniformly spaced grids). It is possible to have multiple regions for each direction with different spacing. The final region is marked by lgrid=.true. although in our example, each direction has a single, uniform grid. Next, grvcon and eqos set details about the gravity and the equation of state for the gas. Finally, gcon sets rarely changed details about the grid and iocon sets the frequency of outputs.

The actual grid values have to be set at compile time in a problem file indicated by the PROBLEM macro (kh in this case). The problem file, src/kh.F is given below.

```
#include "zeusmp.def"
       subroutine kh
       implicit NONE
#include "mod_names.h"
#include "param.h"
#include "field.h"
#include "bndry.h"
#include "grid.h"
#include "root.h"
#include "scratch.h"
#include "mpicom.h"
       integer i,j,k
       if(myid.eq.0)then
          write(*,*)"K-H problem setup."
       endif
       do k=ks-2,ke+2
          do j=js-2, je+2
              do i=is-2,ie+2
                 e(i, j, k) = 2.5
                 v2(i,j,k)=0.0
                 v3(i,j,k)=0.0
                 if(x2a(j).le.-0.25)then
                    d(i, j, k) = 1.0
                    v1(i, j, k) = -0.5
                 endif
                 if(x2a(j).ge.0.25)then
```

```
d(i,j,k)=1.0
v1(i,j,k)=-0.5
endif
if((x2a(j).gt.-0.25).and.(x2a(j).lt.0.25))then
        d(i,j,k)=10.0
        v1(i,j,k)=0.5
        endif
        enddo
enddo
enddo
return
end
```

The physical setup of the problem is that if $-0.25 < x^2 < 0.25$, $\rho = 10$ and $v_1 = 0.5$. Otherwise, $\rho = 1$ and $v_1 = -0.5$. Since the same code must work for each processor, grid values must be set based on the value of the grid (x1(i)) and not the grid count (i).

Compilation is handled by a Makefile which is generated for a given system by the **src/configure.pl** script (based on either a predefined set of values or reasonable guesses). The code is then run by the **exe/mpiwrap.pl** script which takes care of much of the complication of possible restart files and MPI commands. The results of this problem are shown in the density slices of Figure 2.2 and three-dimensional volume renderings in Figure 2.3. After a fairly short amount of time, the three layers have nearly completely mixed.

2.4.4 Performance

To test performance and scaling, a simulation was run with a 100^3 Cartesian grid for 200 time steps in pure hydrodynamics mode with radiative cooling added. This simulation was run as one to eight processes (not necessarily processors). For single processor machines, this means multiple processes on one processor (and presumably a large performance penalty). For the cluster tests this is not a problem. The



Figure 2.2: Density slices for KH test problem.

timing results are shown in Figure 2.4. There are two surprises in these results. The first is that at eight processors, code compiled with the GNU compilers seems to outperform the code from the Intel compiler (when used on AMD processors). This is surprising as the Intel compiler is widely considered to produce the fastest code for x86 processors. The second surprising feature is that the G4, while it performs far below anything else in the tests, does not appear to have any penalty for multiple processes on a single processor.

The results from the scaling test are shown in Figure 2.5. While at no point beyond two processors do we get the theoretical linear increase, we continue to get a decent speedup (with a constant slope in most cases) up to at least eight processors,



Figure 2.3: Results for KH test problem showing 3-dimensional features in density.

with the code produced with the GNU compiler getting much closer to linear than the Intel compiler.



Figure 2.4: ZEUS-MP run time benchmarks for a variety of processors on a 100^3 cell grid.

2.5 Conclusions

Having given a brief description of our main numerical tool, we may now move on to using the ZEUS code for astrophysics in the remaining chapters. The ability to use an existing, high-quality, and well tested code had greatly facilitated this work. We have made our modifications public in an attempt to help keep the availability of such codes high.

All simulations were done on the University of Maryland, Astronomy Department's GNU/Linux Beowulf cluster, borg, or the University of Maryland's High Performance Computing cluster, deepthought. Run times were of the order of one to two months for four processor simulations.



Figure 2.5: ZEUS-MP speed-up tests for a variety of processors on a 100^3 cell grid along with the theoretical performance.

Chapter 3

Energetic Impact of Jet-Inflated Cocoons in Relaxed Galaxy Clusters: Two-Dimensional Models

3.1 Introduction

In this chapter we will discuss a series of two-dimensional models of back-to-back jets in a cluster center. Collectively, the existing body of simulation work (see Chapter 1 for details and references) has allowed us to explain many of the observed features of AGN/ICM interactions, investigate how AGN induced flows mix metals within the ICM core, and study entropy injection and heating of the ICM. However, most of the previous sets of simulation were performed for a limited set of jet parameters and hence it is unclear how to generalize the results to the population as a whole. Also, many of the past simulations have modeled only the buoyant phase of bubbles rather than the self-consistent jet-inflated cocoons. In this chapter, we present a moderately large set of (axisymmetric) high-resolution hydrodynamic simulations of jet/ICM interactions that survey a wide range of jet powers and jet velocities. We study the morphology of the radio-galaxy as well as the injection of energy and entropy into the ICM as a function of the jet properties. In Section 3.2, we discuss the details of our setup. In Section 3.3, we present the analysis of our simulations and their classifications. In Section 3.4 we discuss the results, followed by conclusions in Section 3.5.

3.2 Simulation Setup

Our goal is to model a relaxed cluster and its interaction with a central radio galaxy that produces back-to-back jets. To do this, we start out with an initial cluster that is spherically symmetric, static, and isothermal, with the (adiabatic) sound speed $c_s = 1$ everywhere. The initial density of the cluster gas is assumed to follow a β -model profile,

$$\rho(r) = \frac{1}{\left[1 + \left(\frac{r}{r_0}\right)^2\right]^{3/4}},\tag{3.1}$$

where the core radius is $r_0 = 2.0$ in code units.

A static gravitational potential,

$$\Phi = \frac{c_s^2}{\gamma} \ln(\rho), \qquad (3.2)$$

is set to make the gas initially in hydrostatic equilibrium and is assumed to remain fixed in time. This is equivalent to a gravitational potential which is set entirely by a stationary dark matter distribution that dominates the system, so the self gravity of the dynamically evolving gas is ignored. The intracluster gas, while containing a large fraction of the baryons in the cluster, is not a significant contributor to the overall cluster mass (around 14% of the total cluster mass, see Spergel et al. (2007) for the most recent cosmological values).

We use spherical polar coordinates (r, θ, ϕ) and impose symmetry in the ϕ direction $(\partial/\partial \phi = 0)$. The radial coordinate varied between 0.1 and 30.0 in scaled code units. The comparison between code units and physical units is given in Section 3.2.2. The values of θ were allowed to vary between 0 and π (i.e., we are modeling the full sphere). To avoid numerical difficulties, a very small region around $\theta = 0$ and $\theta = \pi$ was excluded from the grid with a reflecting boundary condition. In our geometry and setup, this has the effect of a small cone-shaped region excluded down the center of each jet axis. Under the assumptions of hydrodynamics this should have no effect on the results (although in the more general MHD case, this could be problematic for the field geometry). Since we have included the full range of angles, this allows us to model back-to-back jets. Thus, we can model interaction between the backflows from the jets rather than the commonly employed technique of imposing reflecting boundary conditions on the $\theta = \pi/2$ plane (e.g., Chapter 4 and 5). We included second-order differences in grid spacing in the θ direction for each hemisphere (i.e., the difference between 1.002 and 1/0.998; see below), giving us two realizations of the jet/cocoon structure per simulations. In no cases were there major differences between the two sides, confirming that the details of the gridding do not affect our results. A small circle around r = 0 was excluded from the computational domain to avoid the coordinate singularity.

All simulations were run with a $n_r \times n_{\theta} = 1200 \times 600$ grid. In the radial direction, a ratioed grid was used, with successive cells increasing by a factor of 1.003. The θ direction also used a ratioed grid increasing by a factor of 0.998 from 0 to 0.5π and then decreasing by a factor of 1.002 from 0.5π to π . This provides the greatest number of zones near the center of the cluster and near the two jet axes while providing the second-order differences in the θ direction mentioned above.

Clearly jets are three-dimensional structures. However, there is a long history of two-dimensional modeling of jets that supports the usefulness of two-dimensional jets (Carvalho and O'Dea 2002; Churazov et al. 2001; Mizuta et al. 2001; Reynolds et al. 2002). There are, however, several aspects of the dynamics that cannot be captured in two dimensions. There is no way to have realistic random motions in the background. This lack of random motions produces lobes that are far more regular and symmetric than those found in any real source. The two-dimensional assumption also enhances mixing between the high entropy material and ambient material. This is because non-axisymmetric KH and RT modes are clearly not possible in an axisymmetric system. Without non-axisymmetric modes, all of the mixing must be done by the axisymmetric modes, which are more efficient at mixing (Reynolds et al. 2002). Fortunately, this will be most problematic only at late times when there has been enough time for significant mixing to occur.

The hydrodynamic evolution of the jet/ICM system was followed using the ZEUS-3D code (Stone and Norman 1992a,b) in its axisymmetric mode.. For further details on ZEUS, see Chapter 2. Radiative cooling, represented by the Λ -term in Equation 2.3, is neglected in this chapter order to concentrate on the hydrodynamic evolution and the thermodynamic response of the ICM to the jets.

3.2.1 Modeling the Jets

Two, back-to-back jets were injected into the center of the model cluster. This was achieved with an inflow boundary condition on the inner radial boundary. All of the jets were highly supersonic with respect to the ambient sound speed and moderately to highly supersonic with respect to their internal sound speeds. Direct injection allows us to keep the details of jet acceleration and central engines outside of the grid, so we are able to solve the equations of hydrodynamics consistently within our computational grid.

The injected jets are initially conical with half-opening angles of 15°, and are injected in pressure balance with the ICM core. The drop of internal pressure associated with subsequent jet expansion leads to external, pressure driven recollimation and internal shocks, ultimately resulting in some new, self-consistently determined opening angle. Figures 3.1 and 3.2 show the internal structure of the jet near the end of the active phase and Figures 3.3 and 3.4 shows the internal structure of the jet during the passive evolution in one of our simulations. Given these (fixed) assumptions, the properties of a given jet are parameterized by the kinetic luminosity $L_{\rm kin}$ and internal Mach number $\mathcal{M}_{\rm jet}$ of the injected jet matter. These are related to the injection density ρ , internal energy density e, and velocity v by

$$\mathcal{M}_{jet} = \frac{v_j}{c_s} = v_j \sqrt{\frac{\rho}{\gamma(\gamma-1)e}},\tag{3.3}$$

$$L_{kin} = \frac{1}{2}\rho_j v_j^3 \pi r_{in}^2.$$
 (3.4)

where $r_{\rm in}$ is the radius of the "injection nozzle" for the jet. The jet parameters for the base run ($L_{\rm kin} = 26.47$ and $\mathcal{M}_{\rm jet} = 10.55$; $\rho_j = 0.01$ and $v_j = 105.5$) are the same as the run presented in Reynolds et al. (2002).

There has been significant analytic work on the stability of AGN jets which can explain why the jets in our simulations are able to propagate out to many times their radius. First we note that most of the jet instabilities are seeded by perturbations on the boundary layer between the jet and the background. The background in our simulations is smooth, so while there is still noise to seed perturbations, there should be less to drive it than in real jets. Since we restrict ourselves to non-magnetized simulations, the KH instability is the most relevant for the jets. Hardee (1979) showed that the lowest order mode of this (the "pinch" mode) grows slower than the jet propagation speed for high Mach number jets, so most of our jets should be



Figure 3.1: Density map of jet at t = 1.0 (run 1). Only the region around a single jet is shown.



Figure 3.2: Entropy map of jet at t = 1.0 (run 1). Only the region around a single jet is shown.



Figure 3.3: Density map of jet at t = 2.0 (run 1). Only the region around a single jet is shown.



Figure 3.4: Entropy map of jet at t = 2.0 (run 1). Only the region around a single jet is shown.

stable to this instability. Hardee (1982) showed that the next mode (the "helical" instability) decreases in wavelength if the jet expands (which is allowed by our opening angle) for reasonable jet parameters for AGN jets, showing that once again, our jets should be stable to these type of instabilities for reasonable times (compared to the jet lifetimes in our simulations) and can propagate long distances.

A total of 26 models were run, where only the kinetic luminosity and/or the Mach number of the jets were varied. In two of these simulations, the jet lifetimes were also varied. The full list of simulations is given in Table 3.1. Internal Mach numbers in the range 3.7–21.1 and kinetic luminosity in the range 6.6–68.8 were explored. The choice of parameters used was based on raising or lowering the Mach number and/or kinetic luminosity by a factor of two from one of the previous runs.

3.2.2 Comparison to Real Units

Throughout this chapter, quantities are quoted in dimensionless code units unless explicitly stated otherwise. Since we are solving the equations of ideal hydrodynamics (Equations 2.1-2.4) with no additional physics added, there is no unique set of physical scales to our problem. Each simulation may therefore be compared to several different sets of physical scales. Following Reynolds et al. (2002), we quote two scalings for our simulations (although there is actually an entire threeparameter family of scalings). For our rich cluster scaling, we set the core radius $r_0 = 100$ kpc, meaning one code unit in r is equal to 50 kpc. Such clusters are hot, with sound speed $c_s = 1000 \text{ km s}^{-1}$. This, along with the typical central number density $n_0 = 0.01 \text{ cm}^{-1}$, gives a time unit of 50 Myr and a kinetic luminosity unit of $L_{kin} = 3.5 \times 10^{44} \text{erg s}^{-1}$. We can also consider a poor cluster scaling, $r_0 = 50$ kpc, $c_s = 500 \text{ km s}^{-1}$, and $n_0 = 0.1 \text{ cm}^{-1}$. These give a length unit of 25 kpc, a time unit of 10 Myr and a kinetic luminosity unit of $L_{kin} = 4.4 \times 10^{42} \text{ erg s}^{-1}$.

Run	Jet Active Time	L_{kin}	\mathcal{M}_{jet}	Morphology
1	1.0	26.47	10.55	Cocoon
2	1.0	52.94	21.1	Cocoon
3	1.0	26.43	15.89	Cocoon
4	1.0	26.31	7.39	Cocoon
5	1.0	68.97	10.46	Cocoon
6	1.0	23.32	10.54	Cocoon
7	1.0	52.94	10.55	Cocoon
8	0.5	26.47	10.55	Cocoon
9	1.5	26.47	10.55	Cocoon
10	1.0	8.91	6.47	Cocoon
11	1.0	17.91	9.48	Cocoon
12	1.0	19.98	15.0	Cocoon
13	1.0	31.09	20.02	Cocoon
14	1.0	13.24	21.1	Non-Cocoon
15	1.0	13.23	10.55	Non-Cocoon
16	1.0	26.47	21.1	Non-Cocoon
17	1.0	6.62	10.55	Non-Cocoon
18	1.0	9.96	8.99	Non-Cocoon
19	1.0	21.01	13.0	Non-Cocoon
20	1.0	30.1	18.02	Non-Cocoon
21	1.0	24.98	19.0	Non-Cocoon
22	1.0	52.99	5.28	Unresolved
23	1.0	26.44	3.71	Unresolved
24	1.0	26.43	5.27	Unresolved
25	1.0	10.05	5.01	Unresolved
26	1.0	12.15	6.54	Unresolved

Table 3.1: Parameters defining the 26 simulations presented in this chapter. Also given are the (visual) morphological classification of the resulting structure, as judged at time t = 2.0. See text for more details.

3.3 Analysis and Results

The primary diagnostic for distinguishing injected jet material from ambient and shocked-ambient material is the specific entropy index,

$$\sigma = \frac{P}{\rho^{\gamma}}.\tag{3.5}$$

The ambient material (both undisturbed and shock heated) has a relatively low specific entropy index (which ranges from 0.6 to 9) and the injected jet material has a specific entropy of injected material does vary with jet power, ranging from near 1 for the weakest jets to 10^7 for the most powerful. With only two ways to change entropy (increase with shocks and increase or decrease through numerical mixing), specific entropy generally provides a powerful way to distinguish injected (and shocked) material from the background gas. A cutoff of $\sigma = 10$ was set to separate jet material from ambient material. This is close to the highest ambient value so we are not likely to artificially include ambient matter as part of the cocoon. We use the same entropy-index cutoff for all simulations to allow for direct comparisons to be made. In the case of the weakest jets, this entropy index cutoff cannot pick out the pre-shocked jet material. Once the jet material passes through the terminal shock, however, it is always clearly distinguished from ambient material by a $\sigma = 10$ cutoff. For a further discussion of the entropy cutoff diagnostic, see Reynolds et al. (2002).

3.3.1 Morphological Classification

Based on the entropy maps, the simulations were visually placed into three classes. Simulations with a distinct, well-defined, monolithic cocoon of jet-originating plasma were labeled as "cocoon bounded sources" (or simply "cocoon sources"). We consider a cocoon to be a jet-inflated region of high entropy material with a coherent structure that is well separated from the low entropy background and relatively safe from mixing (until late times). Simulations that showed faint, wispy plumes of jet-originating material instead of well defined cocoons were labeled "non-cocoon bounded sources" ("non-cocoon sources"). The final group were those simulations that do not appear sufficiently resolved for an accurate classification. In the discussion below, we abbreviate the term "jet-originating material" to "jet material".

An example of a typical cocoon source can be seen in Figures 3.5 and 3.6. Figures 3.7 and 3.8 show a typical non-cocoon source. A poorly resolved simulation can be seen in Figures 3.9 and 3.10. We also observe that it is possible to separate the cocoon-bounded from the non-cocoon bounded sources by use of density contours. Figures 3.11 and 3.12 show density contours for a cocoon bounded and non-cocoon bounded source respectively. Simulations in the cocoon category appear to have the high entropy jet material in a distinct region separated from the background by a solid density contour and remain that way until fairly late times. Non-cocoon sources do not have such distinctions, and the jet material seems mixed in with the background even early on.

These classifications were performed at the same time for all simulations, t = 2.0. At this time, the structures have evolved passively for the same amount of time that they were driven by the jet. However, our classification does not appear to depend sensitively on the choice of time, although there are some limitations. In the very



Figure 3.5: Entropy map of central regions of typical cocoon bounded source (run 1) at t=2.0.

earliest stages of the active jet phase (t < 0.2-0.4), the jets are still pushing out from the inner boundary of the simulation (i.e., the morphology is strongly determined by the artificial aspects of the simulation) and classification is not meaningful. At late times (t > 10 - 15), mixing has disrupted all of our simulated sources and, again, eliminated any morphological distinction.

Using this classification, we can study how morphology depends on $L_{\rm jet}$ and



Figure 3.6: Density map of central regions of typical cocoon bounded source (run 1) at t=2.5.

 \mathcal{M}_{jet} . Our results are shown in Figure. 3.13. We see three separate regions on this parameter space. At the high kinetic luminosities and lower Mach numbers, the runs are unresolved. From this point on, the unresolved sources were excluded from all analysis. Low luminosity sources seem to fall in the non-cocoon category, almost regardless of Mach number (except that at high Mach number they can have a higher kinetic luminosity). The middle region contains the cocoon bounded sources.



Figure 3.7: Entropy map of central regions of typical non-cocoon bounded source (run 21) t=2.0.

In Figure. 3.14, we show the parameter space again, on a density-velocity plot. Although the overlap between the morphologies is more apparent in this parameter space, we can see that there appears to be a continuum of morphologies with the fastest and lightest jet producing unresolved sources at one end and the slow heavy jets producing non-cocoon sources. Again, the cocoon sources occupy the middle ground between the other two.



Figure 3.8: Density map of central regions of typical non-cocoon bounded source (run 21) t=2.5.

As revealed by Figure. 3.13, the line separating cocoon bounded and non-cocoon bounded sources on the (L_j, \mathcal{M}_{jet}) plane is approximately linear (i.e., $L_j \propto \mathcal{M}_{jet}$). This can be understood from the analytic models of cocoon expansion from Begelman and Cioffi (1989). Equation (4) of Begelman and Cioffi (1989) states that the pressure of a cocoon is given by

$$p_c \sim \frac{(L_{\rm kin} v_j \rho_a A_h)^{1/2}}{A_c}$$
 (3.6)



Figure 3.9: Entropy map of central regions of typical unresolved source (run 23) at t=2.0.

where A_h is the area of the working surface over which the jet deposits its momentum, A_c is the cross-sectional area of the cocoon, and ρ_a is the density of the ambient ISM/ICM. Well defined cocoons can only be sustained if p_c exceeds the pressure of the ambient medium, p_a . Noting that the axial ratio of the cocoons are approximately constant (i.e., $A_h/A_c \sim \text{constant}$), and that the density and pressure of our model ambient ICM atmosphere outside of the ICM core is the same in all



Figure 3.10: Density map of central regions of typical unresolved source (run 23) at t=2.5.

simulations and drops off as $r^{-3/2}$, the condition that $p_c > p_a$ can be written as the condition

$$(L_{\rm kin}v_j)^{1/2} > \Upsilon_{\rm crit}, \tag{3.7}$$

where $\Upsilon_{\rm crit}$ is a weak function of r ($\Upsilon \propto r^{1/4}$) and hence time (as the jet/cocoon structure grows and reaches larger r with time). Since the jets are injected with fixed initial pressure, it can be shown that initial jet velocity depends on the kinetic



Figure 3.11: Density contours for the inner half of a cocoon bounded source (run 1) at t=2.0.

luminosity and initial internal Mach number according to $v_j \propto L_{\rm kin}/\mathcal{M}_{\rm jet}^2$. Thus, we can use Equation 3.7 to see that the line separating cocoon from non-cocoon sources (i.e., where the cocoon just comes into pressure balance with the ambient ICM) can be written as $L_{\rm jet} \propto \Upsilon \mathcal{M}_{\rm jet}$. Further discussion of the physical origin of the cocoon is reserved for the Discussion (Section 3.4).



Figure 3.12: Density contours for the inner half of a non-cocoon bounded source (run 21) at t=2.0.

3.3.2 Mass Distribution and Energetics

To examine the mass distribution and energetics during the simulations, several quantities were calculated at each output time. The total mass of jet material at each time is given by:

$$M_{cocoon} = \int_C \rho dV, \tag{3.8}$$

where C is the region with $\sigma > 10$. The internal, kinetic, and potential energies were calculated for both the jet and the ambient material,

$$E_{int}^{C,A} = \frac{1}{\gamma - 1} \int_{C,A} P dV,$$
 (3.9)

$$E_{kin}^{C,A} = \frac{1}{2} \int_{C,A} \rho v^2 dV, \qquad (3.10)$$



Figure 3.13: Morphology vs. jet parameters. Lines separating the different regions were done by eye. Although the lines between the different regions are not perfect, there appear to be regions in which different classes are clearly excluded.

$$E_{pot}^{C,A} = -\int_{C,A} \rho \Phi dV. \tag{3.11}$$

where A is the region with $\sigma \leq 10$. In all cases, the integral is taken over only the region of interest (i.e., cocoon or ambient only, as determined by the entropy cutoff). We then subtract the initial energies in order to derive the change of energy $\Delta E_{int}^{C,A}$, $\Delta E_{kin}^{C,A}$, and $\Delta E_{pot}^{C,A}$.

The results of these calculations for a representative cocoon bounded simulation are shown in Figure 3.15. Figure 3.16 shows similar plots for a non-cocoon bounded simulation. These figures are similar to Figure (4) of Reynolds et al. (2002) and primarily differ due to our increased outer radius (and the fact that plots are continued until later times). The top-left panel in each figure shows the total mass of the cocoon material compared to the amount of injected material. In both cases, this



Figure 3.14: Morphology vs. jet parameters. On the density-velocity plane, we see the previously distinct regions as more of a continuum moving from one type of source to the next.

increases until the jet turns off and then begins to decrease. For the cocoon-bounded source, this decrease is mostly steady for the entire lifetime of the simulation. For the non-cocoon source, the drop is much greater early on, so that about halfway through the simulation most of the high entropy material has mixed with the background, and thus the decay rate has lessened. This enhanced mixing is a direct consequence of the lack of a well-defined contact discontinuity in the non-cocoon sources. Note that the cocoon mass does not correspond to the cocoon mass plot in Reynolds et al. (2002) as we are interested in the total cocoon mass, not the change in cocoon mass.

Also shown in Figures 3.15 and 3.16 (top-right panel) is the length of the region containing jet material. In the cocoon bounded sources, the length of the cocoon evolves in a smooth manner, gradually decelerating as the cocoon comes into pressure balance with the ambient ICM. The discontinuity at late times occurs well into the phase in which the cocoon has transformed into two buoyantly rising plumes and corresponds to the complete mixing/dispersal of the leading part of the plume. The evolution of the length of the non-cocoon bounded sources, on the other hand, suffers a sharp break at the time that the jet is shut off. This is presumably because in the non-cocoon cases, there is no large amount of hot shocked gas to continue to drive the cocoon growth after the jet has ceased activity; there is only the buoyant rise of the hot material to increase the length of the cocoon region. Also, as Equations 3.6 and 3.7 show, the non-cocoon cases are below the line where the cocoon pressure exceeds the ambient pressure, so not only is it unable to form a real cocoon, but in the absence of an active jet, it is no longer able to expand in the length as quickly as the cocoon sources.

Figures 3.15 and 3.16 (bottom panels) show the change in energies for the jet material and the ICM material respectively. For the jet material, there is little difference in the time-dependence of the energies for the two classes with one exception. In the non-cocoon case, the changes in internal and kinetic energy peak at a comparable value although the kinetic drops to almost zero very soon after. For the cocoon case, the change in kinetic energy never reaches a comparable level to the change in internal energy. Physically, this is close to the heart of the difference between the cocoon and non-cocoon sources. In each case, a set amount of energy is injected by the jet (along with a set amount of mass as seen in the potential change which is not very dramatic). The evolution of the system then determines how this energy is split between the internal energy and the kinetic energy (and eventually how much is transferred to the ICM). In the case of the cocoon, the hot jet and shocked material is kept separate from the background and goes almost entirely into the internal energy of the shocked jet gas. For the non-cocoon case, the energy is split nearly evenly between the internal and the kinetic energy. This means that there is no longer enough internal energy available to inflate a cocoon of shocked gas; more of the energy goes to the kinetic energy of the mostly bare jet (with wispy areas of hot, shocked gas around it). This is consistent with the analytic estimate of Cioffi and Blondin (1992).

The time-dependence of the ICM energies show a greater difference between the two cases. The potential energy suffers a greater change (within a few times the jet lifetime) in the cocoon case due to the well-defined shell of ICM that is lifted up by the expanding cocoon, while the non-cocoon case has a much more gradual change in the potential energy. The internal energy in the cocoon case drops back down to the initial value around the same time the potential peaks, while the non-cocoon case has what looks like a slow trade-off between the two. Finally, the kinetic energy changes are comparable in the two cases, as there appears to be a similar amount of disturbance in the ICM regardless of the nature of the inflated structure.

3.3.3 Entropy Evolution and Thermalization Efficiencies

We now examine the effect of the jets on the entropy of the background gas. First, we define the specific entropy,

$$\Delta S = S_1 - S_0 = \log \sigma_1 - \log \sigma_0, \qquad (3.12)$$

where σ is the specific entropy index from Equation 3.5.

Figure 3.17 shows the entropy difference (compared to the initial condition) at the final output of the simulations. At this point, the systems were allowed to evolve passively for 19 times the active lifetime of the jet. For each simulation, the total entropy at each radius was calculated and then averaged over the angular



Figure 3.15: Mass and energy vs. time for a cocoon bounded simulation (run 1). Starting from the top left and moving clockwise, the panels are: mass of the cocoon material divided by the injected mass, the maximum radial extent of the high entropy material, the change in internal, kinetic, and potential energies for the ambient, low entropy material compared to the initial state, and the same energy changes for the cocoon material.

coordinate. The split between the different morphologies is evident in this figure. Within the core, the cocoon sources have a higher ΔS overall. Further discussion of this issue is reserved for Section 3.4.

As a final look at the effect of the jets on the cluster energetics, we show the efficiencies for conversion of the injected energy into various forms within the ICM as measured at the final time for both the cocoon and non-cocoon sources. This is defined as the change in each type of energy for the ICM material divided by the total amount of energy injected by the jet at the final output time of the simulation. Figures 3.18, 3.19 and 3.20 show the efficiency for conversion of injected energy into


Figure 3.16: Mass and energy vs. time for a non-cocoon bounded simulation (run 21). Starting from the top left and moving clockwise, the panels are: mass of the cocoon material divided by the injected mass, the maximum radial extent of the high entropy material, the change in internal, kinetic, and potential energies for the ambient, low entropy material compared to the initial state, and the same energy changes for the cocoon material.

ICM internal energy, kinetic energy, and potential energy, respectively. There is no clear segregation between the two types although the cocoon sources do appear to cluster together somewhat. In all cases, a large fraction of the injected energy (50%-80%) ends as gravitational potential energy of the ICM as it expands in response to the AGN heating.



Figure 3.17: Angle-averaged change in entropy. See discussion in text for full description.

3.4 Discussion

In the (resolved) simulations, we see a split in the morphology of the jet-created structures. Although the presence or absence of a cocoon does not correspond in a simple way to the Mach number or kinetic luminosity of the jet (Figure 3.13), it does correspond in a clear way with the jet velocity and jet density (Figure 3.14). The inverse of this is that the forms we see (cocoon or non-cocoon) tells us something about both the jet interaction with the ICM and (certain) aspects of the jets themselves.

The cocoon itself is formed by shocks. The supersonic jet initially flows freely through the ICM. After traveling some distance (usually short on the scale of these



Figure 3.18: Efficiency for change in internal energy vs.kinetic luminosity for ambient medium.

simulations), the jet will build up enough material in front of it that it can no longer flow freely. The material at the head of the jet is shock heated and expands. Up to this point, the evolution is the same for both cocoon and non-cocoon sources. As this shocked material expands, the evolution diverges.

The hot, shocked material expands into the surrounding medium. At the boundaries of the shocked material, fluid instabilities (i.e., the RT and KH instabilities) work to shred the contact discontinuity and destroy the forming bubble. In the cocoon bounded sources, the bubble is inflated before it is shredded and continues to inflate until it reaches pressure equilibrium with the background. While active, the jet deposits energy only at the point where it comes into contact with the shocked ICM and freshly shocked material flows back from there, continuing to inflate the bubble. Once the jet ceases, the bubble floats buoyantly from the cluster center.



Figure 3.19: Efficiency for change in kinetic energy vs. kinetic luminosity for ambient medium.

In the case of non-cocoon sources, the jet is unable to inflate a bubble. If the jet cannot deposit enough energy at the contact point to inflate the bubble, the fluid instabilities will mix the shocked material with the background material, preventing the bubble from ever growing. This leaves areas of hot, shocked gas that can rise, but never produces a single coherent structure.

The clearest distinction in the thermodynamics of cocoon and non-cocoon sources can be seen in the spatial distribution of the entropy enhancement. We find that in the core regions, cocoon sources cause a greater enhancement of the ICM entropy than non-cocoon sources although there is some overlap. It is fairly clear why the cocoon sources have a great entropy enhancement overall; shocks. The cocoon is formed by shocked gas, so jets that are capable of forming cocoons are also generally the ones that shock more ICM and generate stronger shocks (although this is not



Figure 3.20: Efficiency for change in potential energy vs. kinetic luminosity for ambient medium.

the entire story, as the rate of shocked gas vs. mixed gas matters as well). The non-cocoon sources lack (or at least have less of) these strong shocks and therefore do not have the same entropy enhancement in the core. Outside the core (around r = 5), both appear the same and cause a slight negative change in the entropy. This negative change corresponds to lower entropy material that has been pushed upwards by the expansion of the core. For the non-cocoon sources, this change persists to the edge of the simulation. Some, but not all, of the cocoon sources have a spiky, positive change when approaching the edge of the grid; this corresponds to the remnants of the buoyant cocoon that, in some cases, remains intact until the end of the simulation. It is important to note that the comparison of overall entropy change was done at the final output (t = 20). This is long after our classifications were done. At this point, mixing has disturbed the structure enough that it would no longer be possible to separate the simulations into cocoon or non-cocoon categories. The fact that we still see a difference at this late a time shows that the jets have long lasting effects on the energetics of the ICM and that these effects vary based on the jet properties in a way that is more complicated than simply depending on jet power.

The energy efficiencies (Figures 3.18, 3.19, and 3.20) are initially puzzling. The non-cocoon simulations generally seem to coincide with the highest efficiency cocoon simulations for internal and potential energy efficiency. For the kinetic energy (Figure 3.19), the non-cocoon runs seem to have close to both the highest values and the lowest values. This is in contrast with the situation for entropy where the cocoon simulations have a greater impact. It is not as immediately obvious why jets that are good at increasing the entropy are not always as good at increasing the energy and why there can be a spread in efficiencies for jets with similar parameters. To see why, we must go back to Equation 3.5, the definition of the specific entropy index (keeping in mind that the pressure, P, is directly proportional to the internal energy. For a given change in internal energy, the change in entropy can be higher (or lower) by having it take place at a lower (or higher) density. So if the shocks happen higher in the potential well (and hence at a lower density), it is possible to have a greater change in entropy for the same amount of energy input. The spread in kinetic energies shows that it is possible to stir up the gas significantly with both the cocoon producing and the non-cocoon producing jets. Also, the cocoon-producing jets leave more energy available as internal energy which can go towards shocks, while the non-cocoon producing jets have more of their energy in kinetic energy, which does not go to shocks and therefore is not available to increase the entropy. It is also the case that different jets will cause a different amount of material and energy to be lost from the system at the outer edge of the grid. We consider it

to be reasonable to count this as lost energy as anything that has not managed to heat the gas by the time it reaches the outer edge is clearly not going to be able to change the temperature of the system.

Along with the energetic issues, we may also compare these simulations to observations in a broad sense. Classically, radio galaxies are split into two categories: Faranoff-Riley Type I (FR I) and Faranoff-Riley Type II (FR II) (Fanaroff and Riley 1974). Lacking any information on luminosities and radio emission, we cannot directly compare our simulations to these classes of real radio galaxies, but we can note some interesting correlations. FR II galaxies encompass the large, back-toback classical double sided sources. Our cocoon-bounded sources are reminiscent of the classical doubles. The less lobe dominated and often bent FR I sources share a similar structure with our non-cocoon sources. Although our classifications were done based on the passive phase of the system, we can also compare the jets from the active phase to the FR I/FR II divide. As seen in Figures 3.13 and 3.14, slow, heavy, and lower luminosity jets (the non-cocoon cases) should correspond to the FR I sources while fast, light, and more luminous jets should correspond to the FR II type. As FR II galaxies have higher luminosity jets (Ledlow and Owen 1996), this is at least consistent with our simulations. Even without a definite correspondence between the two, we do show that it is possible to produce qualitatively different structures with differences in the jet parameters without any environmental changes.

Two additional simulations (run 8 and run 9) were performed where all the parameters were kept the same as our standard case (run 1), but the active time of the jet was varied. In run 8, the jet had an active time of only $t_j = 0.5$, while run 9 had an active time of $t_j = 1.5$. In both cases, these simulations fall in the same cocoon-bounded category as run 1. This give us some confidence that our results are not simply based on the amount of total energy inject by the jet, but instead depend on the detailed hydrodynamic interactions between the jets and the ambient material.

3.5 Conclusion

We have performed a large number of axisymmetric simulations of AGN jets in cluster atmosphere. By limiting ourselves to pure hydrodynamics with no radiative cooling, we are able to carefully study the interaction of the atmosphere with the jets and the formation of cocoons and bubbles in the cluster. With the continued interest in the probable role of jets and bubbles in the cooling flow problem, it is important that we understand their dynamics as clearly as possible.

Our main results are summarized in Figures 3.13 and 3.17. The morphology versus jet parameters plot shows that we are able to produce two distinct classes of structures, cocoon bound and non-cocoon bound, by varying the initial jet parameters. This distinction does not change with a change in total injected energy (as seen when we vary the jet lifetime) and is a function of both jet parameters (Mach number and kinetic luminosity). We can also visually draw comparisons between our two classes and the FR I and FR II distinction among real radio galaxies. The angle-averaged entropy change versus radius shows that along with the split in morphology, the effect on the energetics of the ICM also depends on the jet parameters. We also see that even short-lived single burst jets are capable of a long-lived enhancement of the entropy of a cluster core. Finally, we show that regardless of morphological type, our jets are efficient (50%-80%) at converting the injected energy into the gravitational energy of the ICM. That is, the jets are efficient at thermalizing their energy overall and making the cluster "puff-up".

Chapter 4

The Failure of Simple Feedback Models

4.1 Introduction

In the previous chapter, we discussed a series of two-dimensional axisymmetric models of jets in clusters. Although this allowed us to examine the heating and energy distribution possible by hydrodynamic jets, it did not provide us with a way to directly address the cooling flow problem. By building on these models and relaxing several of the limiting assumptions in them, we can make a more direct attack on the cooling flow problem.

In this chapter we perform high-resolution, three-dimensional ideal hydrodynamic simulations of AGN feedback in a relaxed cooling cluster. Unlike many of the previous three-dimensional investigations (with the notable exceptions of Basson and Alexander (2003); Omma and Binney (2004); Omma et al. (2004)), we simulate the injection of a supersonic jet into the ICM atmosphere rather than starting with an initial condition of a pre-inflated, static radio cocoon. We consider this an important issue — only through a direct modeling of the jet can we hope to be able to capture the jet-induced ICM shock heating as well as the complex internal dynamics of the cocoon. While the effect of the jet on the ICM is treated from first principles (through the evolution of the equations of hydrodynamics), the enormous range of scales between the ICM core and the gravitational sphere of influence of the black hole prevents us from treating the formation of the jet or the AGN fueling properly.

In this work, we introduce several different feedback prescriptions by which the AGN jet power is related to the ICM properties in the innermost region. Our ultimate goal is to attempt to construct a model system in which AGN heating and ICM cooling are, in the long term, in balance. We find that our ideal hydrodynamic models *fail* to achieve this balance. Assuming that the AGN feedback picture is indeed correct, we conclude that our models must be failing to capture some crucial, but as of yet undetermined aspect of the system.

In Section 4.2, the basic setup for our simulations is described. Section 4.3 presents our results for the different feedback scenarios considered. We discuss our results along with possible ways to rescue the AGN feedback hypothesis in Section 4.4.

4.2 Basic Setup

We aim to model the ICM of a cluster and its interaction with a central radio galaxy. Many aspects of these simulations are similar to the axisymmetric models in Chapter 3. Our basic picture is a spherical cluster consisting initially of stationary dark matter and gas. The gas is initially in hydrostatic equilibrium, but unlike in Chapter 3 is cooling through optically thin thermal emission. As the gas cools, some of it will flow across the inner radial boundary of the simulation and is removed from the simulation domain. The amount of gas to cross this boundary is used as our primary diagnostic of the cooling flow.

Initially, the cluster is static, spherically symmetric, and isothermal. The gas is set up with a β -model profile (Equation 3.1) and a potential given by Equation 3.2 as in our two-dimensional models in Chapter 3.

Spherical polar (r, θ, ϕ) coordinates are used. All simulations were run on a $200 \times 200 \times 100$ cell grid, with enhanced resolution near the center of the cluster and near the jet axis. This corresponds to a physical grid with a inner radius at r = 10 kpc and an outer radius at r = 1000 kpc. The angular coordinate was only allowed to vary from $\theta = 0$ to $\pi/2$. This effectively only covers half the cluster and therefore only allows for one jet. Reflecting boundaries were used on the $\theta = \frac{\pi}{2}$ plane to mimic the effect of the missing half of the cluster. This aids in maintaining a high-resolution while keeping a reasonable number of grid cells.

In ideal (adiabatic) hydrodynamics, the equations of hydrodynamics can be written in dimensionless form and, hence, the results of a single simulation may be scaled to a whole family of problems with different mass, size and timescales (e.g., see Reynolds et al. (2002)). Since we must add radiative cooling to our simulations (see Section 4.3.1), we are forced into a set of physical units. The values for the rich cluster of Reynolds et al. (2002) were used. This gives us a core radius of $r_0 = 100$ kpc, a code length unit of r = 50 kpc, time units of 50 Myr, and a sound speed of $c_s = 1000$ km s⁻¹. The resulting total radiative luminosity of the model ICM is 1.22×10^{45} erg s⁻¹.

Name	Feedback	Efficiency	Delay	Rotation
Run A	none	N/A	N/A	N/A
Run B	strong jet	N/A	N/A	N/A
$\operatorname{Run} C$	weak jet	N/A	N/A	N/A
Run D	simple feedback	0.0001	N/A	N/A
$\operatorname{Run} E$	simple feedback	0.00001	N/A	N/A
$\operatorname{Run} F$	delayed feedback	0.0001	short	N/A
Run G	delayed feedback	0.00001	short	N/A
$\operatorname{Run} H$	delayed feedback	0.0001	long	N/A
Run I	delayed feedback	0.00001	long	N/A
Run J	delayed feedback	0.01	long	N/A
$\operatorname{Run}\mathrm{K}$	delayed feedback	0.1	long	N/A
$\operatorname{Run} L$	delayed feedback	0.00001	long	Solid body rotation
Run M	delayed feedback	0.00001	long	Fraction of grav. rotation

Table 4.1: List of simulations and included effects.

4.3 Specific Models

A set of 13 three-dimensional simulations were performed. Table 4.1 lists the simulations along with some relevant parameters. In the following sections we describe each simulation and present its results.

An important diagnostic used to compare the different models was the mass accretion rate across the inner boundary of the simulated grid. This was measured in units of solar masses per year. To calculate the mass flow, the amount of mass in each cell on the inner boundary with a negative (inward) velocity in the radial direction was summed. This is probably not completely accurate at the highest densities and mass accretion rates, as the outflow boundaries are not perfectly efficient and sound waves may be reflected off the boundary. However, this should not change the results since these inaccuracies do not occur until the mass flow and density reach physically unrealistic values.

4.3.1 Radiative Cooling

The driving force behind the cooling flow is the thermal bremsstrahlung and line radiation that removes thermal energy from the ICM core. This is modeled with the optically thin cooling law,

$$\Lambda = [C_1 (k_B T)^{\alpha} + C_2 (k_B T)^{\beta} + C_3] 0.704 (\rho/m_p)^2 \times 10^{-22} \text{ ergs cm}^{-3} s^{-1}, \qquad (4.1)$$

which is the same law used by Ruszkowski and Begelman (2002) with the coefficients $C_1 = 8.6 \times 10^{-3}$, $C_2 = 5.8 \times 10^{-2}$, $C_3 = 6.4 \times 10^{-2}$, $\alpha = -1.7$, and $\beta = 0.5$ (Sutherland and Dopita 1993). This cooling term is proportional to ρ^2 , which produces the eventual catastrophic cooling as the center of the cluster becomes denser as it cools. Below a minimum temperature (kT = 0.1 keV), cooling was manually truncated. Even if we allowed material to cool below this limit, our spatial resolution would be insufficient to follow the resulting structures, and additional cooling processes not captured by Equation 4.1 would become applicable.

The cooling represents an extra term in the energy equation (Equation 2.3) and is implemented as an explicit source term. As with all physical processes, there is a maximum allowed time step for numerical stability associated with the cooling, but in our case it is always above the normal hydrodynamic CFL condition (Equation 2.8) and need never be implemented.

Due to the density dependence of the radiative cooling, it is only relevant in the inner regions of the cluster. This is important for three reasons. First, as the core cools, the inner regions loses pressure support. In the absence of heating, this causes the inward flow of material, some part of which could fuel the accretion on the central compact object. Second, as the inner regions become denser due to the sagging, the cooling increases. This causes the eventual runaway cooling whose absence represents part of the cooling flow problem. Finally, since the outer region does not cool appreciably with a Hubble time or so, it represents a large reservoir of hot material that could potentially be exploited to balance cooling (although this does not seem possible in purely hydrodynamic models).

To establish a control case, our first simulation (run A) followed the pure radiative collapse of our model ICM atmosphere (i.e., a spherically symmetric homogeneous cooling flow). The evolution of this system was very simple. The cluster cools, primarily in the inner regions. As it cools, the inner regions become denser. As the gas becomes denser, it cools quicker, and the process runs away (until numerical issues with the very cool dense gas force us to terminate the simulation). The mass accretion rate for pure cooling can be seen in Figure 4.1. With no mechanism to stop or slow the cooling, the mass accretion show a featureless, approximately exponential increase. This unbounded cooling and mass accretion is our first example of a catastrophic cooling.

To diagnose the time taken for the system to undergo a cooling catastrophe, we measure the time taken for each of our simulations to exceed a mass accretion rate of $5000 M_{\odot} yr^{-1}$ (an arbitrary "large" mass accretion rate). For the pure cooling flow run (run A), this threshold mass accretion rate is crossed at 244.1 Myr. Material first fell below our imposed lower temperature limit at 248.9 Myr. The radial dependence of temperature for run A is shown in Figure 4.2. Initially the cluster is isothermal. The inner regions cool first, while the outer regions barely change in temperature. As catastrophic cooling occurs, the temperature gradient in the inner region becomes progressively steeper.



Figure 4.1: Mass accretion rate for pure cooling (run A).

4.3.2 Single AGN Outburst Models

The onset of the cooling catastrophe in the homogeneous cooling flow is well known and unsurprising. The (almost) standard hypothesis and the focus of this work is the idea that jet activity by a central AGN can heat the ICM core and prevent this catastrophic cooling. As an initial exploration of jet heating models, we follow the evolution and effects of a period of jet activity in which the jet has a fixed and constant power for a pre-defined duration. In run B, the jet is active for 50 Myr, after which it is completely shut off and the resulting ICM allowed to evolve passively. These cases are essentially three-dimensional generalizations of the axisymmetric simulations of Reynolds et al. (2002) and share many characteristics with (but are higher resolution than) Basson and Alexander (2003). The axisymmetric simulations



Figure 4.2: Radial temperature dependence for pure cooling (run A).

of Zanni et al. (2005) provide another example of this type of single burst jet heating. This can be viewed as a zeroth-order approximation to AGN feedback.

In detail, the action of the AGN is modeled as a jet of low-density material injected from the inner radial edge of the simulation grid using an inflow boundary condition. The jet has a density of 1/100 of the initial inner ICM density and (at the inner boundary) is in pressure equilibrium with the initial inner ICM. Therefore, the internal sound speed of the injected jet material is 10 times that in the initial ICM atmosphere. The jet has an opening angle of 15° . In the case of a single jet outburst of run B, a Mach number of around 10.5 compared to the background material was used, resulting in a total kinetic luminosity of $3.6 \times 10^{46} \,\mathrm{erg \, s^{-1}}$. We also modeled a weaker jet with a kinetic luminosity of $8.9 \times 10^{45} \,\mathrm{erg \, s^{-1}}$ in run C (which was otherwise identical to run B).

Plots of entropy are shown at various times for run B in Figure 4.3. Density slices for the same run and the same output times are in Figure 4.4. As discussed in Chapter 3, plots of entropy are a good tool for differentiating shocked jet material from background material. The density plots a useful in showing the shock produced by the jet impact with the ICM. It is clear from the density plots that th shock will touch a much larger fraction of the ICM than the jet or bubble will touch. This includes regions which are 90° away from the jet axis. Initially, the jet carves a path through the ICM and terminates in a shock. It can be seen that the jet channel is surrounded by a backflow of shocked (high-entropy) jet material. This shocked material is over-pressurized compared to the background, and expands laterally into a cocoon structure. The cocoon is separated from the background by a contact discontinuity. RT and KH instabilities work to shred this cocoon and mix in background material. After the jet has been turned off, the cocoon is left to evolve passively. Buoyancy forces cause it to rise, leaving the cluster core. Once the cocoon completely detaches itself from the core, it takes on the appearance of a rising bubble, which spreads out and fades somewhat into the cluster background (much like observed ghost bubbles) until it leaves the computational grid.

Figures 4.5 and 4.6 show a close up view of the jet just before the end of the active phase for run B. It is useful to compare these to the two dimensional jets of Chapter 3 seen in Figures 3.1 and 3.2. The three dimensional jets appear slightly "fatter", that is, the width to length ration is greater. This is due to the different growth rates of the hydrodynamic instabilities in two and three dimensions. There is also less internal structure visible, although how much of that is due to the difference in resolution and how much is due to the differences between two and three dimensional models is unclear.

Figure 4.7 shows mass accretion for run B and run C. Although these differ



Figure 4.3: Entropy plots for single jet (run B) at t=50, 125, 250, and 375 Myr. The semi-circle in each plot represents the outer edge of the simulated grid at 1000 kpc.



Figure 4.4: Density plots for single jet (run B) at t=50, 125, 250, and 375 Myr. The semi-circle in each plot represents the outer edge of the simulated grid at 1000 kpc.



Figure 4.5: Density map of jet at t = 1.0 Myr (run B).

slightly in the timing, they qualitatively show the same behavior. Initially, there is moderate mass flow (a few hundred solar masses per year). Ironically, the jet activity is actually responsible for this initial enhanced period of inflow; some matter in the innermost region of the ICM core becomes caught in the backflow that results from the onset of the of the jet activity and is swept across the inner boundary. After the initial spike, the mass accretion rate drops off to very low values while the jet is on and remains very low for over 100 Myr after the jet stops due to the shock heating of



Figure 4.6: Entropy map of jet at t = 1.0 Myr (run B).

the ICM core. The cooling flow then starts to re-establish, and the mass flow begins to increase (eventually) catastrophically as it does in the pure cooling case. The time for catastrophic cooling has been delayed somewhat from the pure cooling case. For run B, \dot{M} does not reach 5000 M_{\odot}yr⁻¹ until 306 Myr (more than 60 Myr later than the pure cooling model). Material falls below the minimum cooling temperature at 305 Myr (i.e., at essentially the same time that the mass accretion rate crosses our "catastrophe" threshold). The weaker jet, run C, reaches the catastrophic point at



Figure 4.7: Mass accretion rate for single burst simulations (runs B and C).

285 Myr, which is only slightly delayed from the pure cooling case. The temperature falls below our floor at 274 Myr (also slightly later than the pure cooling case).

Figures 4.8 and 4.9 show the radial temperature dependence for runs B and C at a late time (300 Myr). The temperature is show off the jet access to avoid the relic bubble and just look at the temperature in the ICM. In both cases, the central temperature has dropped dramatically by this time (although the central temperature is slightly higher for run B). In both cases the shock excited by the jet can be clearly seen around 500 kpc and 400 kpc respectively. The shock also separate the cooling inner regions from the outer regions which do not cool appreciably.



Figure 4.8: Radial temperature dependence for a single burst jet (run B) at 300 Myr away from the jet axis ($\theta = 70^{\circ}$).

4.3.3 Feedback Models

The radio-loud AGN feedback hypothesis would argue that the central AGN acts such as to maintain a long-term balance of jet heating and radiative cooling of the ICM. Furthermore, this must occur in systems with widely varying radiative luminosities. Clearly, this mechanism requires that the AGN power be somehow regulated by conditions within the ICM so that it does not underheat or overheat the ICM core (see Binney and Tabor (1995) for early work on AGN feedback on cooling flows). In this work, we explore a set of models in which the kinetic luminosity of the jets is connected to the mass accretion rate in the cooling flow. Unlike the other effects (e.g., cooling and jet propagation), feedback is the only one that does not directly represent fundamental physics of the cluster gas. Instead, it tells us



Figure 4.9: Radial temperature dependence for a single burst jet (run C) at 300 Myr away from the jet axis ($\theta = 70^{\circ}$).

something about fueling in the central engine of the AGN. This provides us with the most freedom in how to implement it. Due to the large range of scales involved, it is not possible to directly model feedback. This would require resolving gas from the outer edges of the cluster (nearly a megaparsec) all the way down to the inner accretion disc (sub-parsec scales). A direct assault on this problem would not be practical at the current time even for an AMR code. Instead, we must use some simplified model for feedback.

In our first set of feedback simulations (runs D and E), the jet is injected with a kinetic energy proportional to the instantaneous mass accretion rate across the inner radial boundary of our simulation domain. This is achieved by modulating the injection velocity to be

$$v_{jet} = \left(\frac{2\eta \dot{M}c^2}{A\rho}\right)^{\frac{1}{3}},\tag{4.2}$$

where \dot{M} is the mass flow rate across the inner radial boundary of the simulation, c is the speed of light, ρ is the density of the injected jet material, A is the area of the jet "nozzle" on the inner boundary, and η is the efficiency with which the rest mass energy of the ICM cooling flow is converted into jet power. The appropriate choice for η is far from clear. In the extreme case that all of the matter that flows across the inner boundary was to accrete onto the central supermassive black hole, η would be the jet-production efficiency of the actual black hole accretion disk itself. For an efficient disk, we would have $\eta \sim 0.1$, although higher efficiencies are possible if the black hole is spinning (Novikov and Thorne 1973) or if magnetic coupling within the radius of marginal stability becomes important (Armitage and Reynolds 2003; Gammie 1999). However, as has been extensively discussed, supermassive black holes in the centers of cD galaxies accreting at a modest rate from the hot interstellar medium may well have a significantly smaller efficiency ($\eta \sim 10^{-3}$). Much of the remaining energy is advected across the event horizon or driving a slow, uncollimated wind from the disk. Furthermore, it is plausible that only a small fraction of the mass that flows across the inner boundary of our simulation domain (located at 10 kpc) actually enters the sphere of influence of the massive black hole, with some of the remaining going to fuel low-level star formation.

Motivated by this discussion, we explored two cases of simple feedback with efficiencies of $\eta = 10^{-4}$ (run D) and $\eta = 10^{-5}$ (run E). The resulting mass accretion for run E (lower efficiency) is given is Figure 4.10. The higher efficiency (run D) shows the same pattern. Initially, the mass accretion rate cycles up and down in response to the jet power (partly due to the backflow noted above), as we would expect if feedback works to stop cooling. However, by 150 Myr, this breaks down.



Figure 4.10: Mass accretion rate for feedback model (run E).

From that point on, the mass accretion rate increases, and the increasingly powerful jet does nothing to stop it. The time for catastrophic cooling and for material to fall below the minimum cooling temperature is not significantly affected by this type of feedback. Both runs have very similar catastrophic cooling times with 246.7 Myr for run D and 245.43 Myr for run E. This is essentially the same as the time for the pure cooling model.

The reason for this failure of feedback to prevent the cooling catastrophe can be seen in Figure 4.11. After the jet has been active for a while (even at very low power) it clears a low-density channel in the ICM. As cooling and ICM accretion proceeds along equatorial latitudes (i.e., in the plane perpendicular to the jet axis), the increasingly powerful jet flows freely down the pre-cleared channel. This prevents the jet from depositing energy near the cluster core. Instead the kinetic energy is



Figure 4.11: Temperature Map (top left), pressure map (top right), entropy map (bottom left), and density map (bottom right) for immediate feedback jet. The temperature map only shows the highest temperatures to pick out primarily jet material. The thin, low-density channel can be easily seen in temperature and pressure (and to a lesser extent entropy). Only the inner 254 kpc of the simulation is shown.

carried to the head of the cocoon reasonably unimpeded. With no energy deposition near the core, cooling proceeds on the catastrophic course in the equatorial regions almost as if there were no jet.

4.3.4 Delayed Feedback Models

The failure of the simple feedback model is, at least in part, related to the formation of the low-density channel along the jet axis. We note that Omma et al. (2004) also discussed the importance of such channels within the context of their slow-jet simulations and argued that the effective heating required the channels to fill in between powerful outbursts. Motivated by this, and in an attempt to model more realistic feedback, we performed a set of simulations in which a time delay was introduced between the mass accretion rate and the response of the jet. Clearly the immediate feedback of runs D and E is not physically accurate. Some time delay must be added to account for the material travel time from the inner edge of the grid to the accretion disk and onto the black hole. We note that the time for the relativistic jet to reach from the black hole and to enter the computational domain is negligible.

In runs F and G, a delay of 10 Myr was introduced. This is the sound crossing time of the cluster center (the area from the inner edge to r = 0) and represents the minimum physically reasonable delay. Runs H, I, J, and K all have a delay time of 100 Myr. This is approximately the dynamical time for the central galaxy and is a reasonable "best guess" at a physically plausible time delay. Interestingly, this is also the approximate cooling time observed by *Chandra* in the centralmost regions of cluster cooling cores. We did not test delay times longer than this as they are not physically relevant. If the delay time is longer than the central cooling time, then the AGN would not be able to respond to the cooling gas quickly enough to allow for any type of feedback. The efficiencies of these runs is listed in Table 4.1.

An example of the low-efficiency delayed feedback with a short delay can be seen in Figure 4.12. There is not much difference between this and the immediate feedback. There is also no essential difference in the time to reach catastrophic cooling. The only minor effect we see here is that the lower efficiency run (G) has a slightly lower rate of increase for the accretion, probably due to less efficient channel formation.

For runs H through K, the physically motivated delay of 100 Myr was used. Runs H and I are essentially copies of runs F and G with a longer delay between the mass flow and the jet. As Figure 4.13 shows, this have a more significant effect on the mass flow than the shorter delay had. For the first 100 Myr, there was no jet, so the run proceeded as a pure cooling run. The onset of (weak) jet activity



Figure 4.12: Mass accretion rate for low-efficiency delayed feedback (run G).

caused a sharp dip in the accretion, followed by a spike (up to 600 M_{\odot} year⁻¹) lasting for about 50 Myr (or half of the dynamic time). However, after several smaller spikes and dips, runaway cooling proceeds in the equatorial plane. In these cases, the cooling catastrophe is delayed as compared with the pure cooling model, but only by 40 Myr. The higher efficiency version (run H) behaved similarly only it reached the cooling catastrophe slightly faster (once again due to increased channel formation).

Given the difficulties with the feedback models found above, we also ran two simulations with *very* efficient jet production; runs J and K possessed $\eta = 0.01$ and 0.1 respectively. Run K was performed in the spirit of an *extreme case*, and we note that it is not reasonable that accretion onto the central galaxy and then onto the central AGN are both perfect (with only the efficiency limits forced by relativity).



Figure 4.13: Mass accretion rate for low-efficiency long-delayed feedback (run I).

Even at an efficiency of 0.01 it is hard to come up with a plausible scenario for such nearly perfect accretion.

Even the extreme runs are not able to completely stop cooling. For run J, the temperature falls below $T_{\rm min}$ at 245 Myr. This is even sooner than in the pure cooling case. Run K does not have this until 267 Myr, which is between the times for the other delayed runs and the single jet runs. The extreme efficiency feedback does, however, significantly delay catastrophic cooling. Run J does not experience catastrophic cooling until 345 Myr. The mass accretion behavior is shown if Figure 4.14. Although it has taken longer for the cooling to reach our catastrophic level, the final catastrophe is extremely sharp, with the mass accretion rate exponentiating on timescales of only 1 Myr or so. This is much more dramatic than any of the less efficient runs. Run K holds off catastrophic cooling for 50 Myr beyond run J (398)



Figure 4.14: Mass accretion rate for very high efficient feedback (run J).

Myr), but it too becomes catastrophic rapidly at that point.

4.3.5 Feedback with Rotation

The AGN feedback and jet heating models described above clearly fail to adequately heat the ICM core in a manner that stably offsets radiative cooling. Our final two simulations (runs L and M) explore the effect of an ICM atmosphere that is not initially static. The ICM of real clusters is never, of course, a perfect hydrostatic atmosphere. Clusters have complex merger and formation histories that leave an imprint on the dynamics of the ICM (Burns and Owen 1977). In several clusters observed with *Chandra*, the ghost cavities are not symmetric about the cluster center (e.g., A4059, Heinz et al. (2002); Perseus A, Fabian et al. (2003)), suggesting that they are buoyantly rising in an ICM atmosphere, which itself has velocity structure. While a full exploration of this class of models is beyond the scope of this chapter, our final two simulations model the case of a rotating ICM atmosphere. Due to the assumed reflection symmetry in the plane perpendicular to the jet axis, we are forced to consider only rotations with an axis coincident with the jet axis.

In run L, a delayed feedback model similar to run I is examined in which the initial ICM atmosphere is assumed to undergo solid body rotation. The rotation speed at the outer edge of the grid is set equal to the sound speed of the cluster. This leave the rotation in the core at a very low value. The results from this were nearly indistinguishable (in cooling, not in detailed physical structure) from run I. This is because the rotation in the core was not high enough to mix the ICM effectively or to stop accretion in any other way. Without having gas moving supersonically in the outer parts of the cluster (which is clearly unphysical), there is no way to get solid-body rotation to help offset cooling.

In run M, a rotation law was chosen such that the centrifugal force associated with the rotation is a fixed fraction (10%) of the gravitational force at that location,

$$v_{\phi} = 0.1r \frac{3}{4\gamma} \sin \theta \frac{1}{[1 + (\frac{r}{r_0})^2]^{0.5}}.$$
(4.3)

This results in a rotation law in which the angular velocity increases with decreasing radius until one gets well within the ICM core, at which point the angular velocity smoothly goes to zero at the center. Figure 4.15 from run M shows the first accretion curve that differs significantly from the previous patterns. Initially, the accretion rate goes up and down in response to the jet. At around 240 Myr, gas falls below the cooling limit (similar time to the pure cooling and most feedback runs). At 285 Myr, the accretion passes the 5000 M_{\odot} year⁻¹ limit and continues to climb. Unlike previous setups, it does not grow without bound (or flatten out at some unrealistically high value). Instead at 16,000 M_{\odot} yr⁻¹, the curve reverses and over the course of 50 Myr drops back down to low levels (a few hundred solar masses per year. It then continues at that rate with only small fluctuations for the duration of the simulation.

In this case, the AGN behavior is actually rather incidental. At early times, the slowly rotating inner core of the cluster cools and accretes. Higher angular momentum material flows inwards and, eventually, the ICM core becomes rotationally supported. Given the assumptions of our model, this material will conserve its angular momentum and, irrespective of radiative cooling, will not accrete. The main role of the AGN is to prevent accretion of ICM from the high-latitude regions within the centrifugal barrier. Indeed, the cooling catastrophe has *not* in fact been averted. Rapid cooling into a rotationally supported disk can clearly be seen in the late stages of this simulation (Figure 4.16). Our mass accretion diagnostic (based on mass flow across the inner radial boundary of the simulation) fails to detect this particular manifestation of the cooling catastrophe.

4.4 Discussion

Our work with simple hydrodynamic models shows that it is more difficult than previously assumed to halt a cooling flow through AGN activity. The ultimate hope (as in, for example, the one-dimensional models of Ruszkowski and Begelman (2002)) is that a proper coupling of the AGN kinetic luminosity to the cluster gas will allow a long-term balance to be established in which the AGN heating balances the ICM radiative cooling. In our three-dimensional simulations, we do not see such balance. Instead, the AGN jet seems to effectively clear a low-density channel through the ICM core, after which time the kinetic energy of the jet is deposited well outside of the cooling core. Thus, simulations with delayed feedback do not produce multiple outbursts. Instead they produce one or two long outbursts with



Figure 4.15: Mass accretion rate for rotating cluster with feedback (run M).



Figure 4.16: Dense cold disk formed in the inner regions of run M. Density at 880 Myr is shown. Only the inner 254 kpc of the simulation is shown.

varying jet power before the ICM core catastrophically cools.

Our modeled AGN feedback is not totally ineffectual at heating the ICM. Both the time for catastrophic cooling and the time for gas to fall below our cooling limit are increased by the AGN in most cases. The single jet cases are rather effective (delaying the catastrophe for as long the most effective feedback simulations). However, the instantaneous feedback and the feedback with a time delay of the core sound crossing time have basically no effect of the catastrophe time (no more than a few megayears). The more realistic delays are capable of delaying the cooling catastrophe by several tens of megayears (more in the unrealistically efficient cases). The only exception to this behavior was the rotating cluster modeled in run M. As shown in Figure 4.15, there is a brief period of very high accretion, followed by an extended period of lower accretion. However, as discussed above, catastrophic cooling is still occurring into a rotationally supported disk. The rotational support prevents this gas from falling into the central region, but it is still a large body of cool and cooling gas around the cluster center that contradicts observations.

Due to the computationally expensive nature of our simulations and the multiple parameters present, it was not possible to test every permutation of the model parameters. It is possible that there is some preferred area of the parameter space where the AGN heating balances the radiative cooling without producing unrealistic jets. The potential existence of such a privileged area of the parameter space does not change our conclusion as it would bring up a serious fine-tuning problem. Cooling and AGN occur in a wide range of clusters of varying masses and temperature. For a regulatory mechanism to work, it must be general enough to work in the different clusters. We believe that our simulations sample physically plausible parts of parameter space, thus the failure to model successful AGN feedback is an important result. The two-dimensional simulations of Chapter 3 sample a much larger area of the parameter space, but they still show a disconnect between entropy increase and energy increase. This both give us added confidence that we are not missing an important area of the parameter space and reaffirms the point that in ideal hydrodynamics it is hard to efficiently heat the cluster gas, even with powerful jets.

As discussed in Section 4.1, however, there is significant observational support for the notion that radio galaxy heating can provide a long-term balance against radiative cooling. Thus, it seems clear that our models are not capturing some ingredients that are important in real systems. In the rest of this section, we explore some of these missing ingredients.

We must acknowledge that while we do model the jet, we may well not capture all of the relevant jet dynamics. Real AGN jets have internal dynamics that can only be captured by considering special-relativistic (Marti et al. 1997) and MHD effects (Omma and Binney 2004; Omma et al. 2004). Furthermore, our simulation probably has insufficient resolution to capture the full internal dynamics of a hypersonic jet even within the context of ideal non-relativistic hydrodynamics. These effects may allow the jet momentum to be spread over a larger working surface (i.e., "dentist drill effect"), thereby helping alleviate the problems associated with the jet drilling out of the ICM core. The fact that our simulations produce structures (cocoons and ICM cavities) with aspect ratios comparable to those observed, however, suggests that this is not a major deficiency of our simulations.

Before leaving issues associated with the properties of the jet itself, we note that Omma et al. (2004) and Omma and Binney (2004) simulated the effects of slow ($\sim 10,000 \,\mathrm{km \, s^{-1}}$), heavier jets on a cooling ICM core. In agreement with our conclusions, they also identify the importance of channel formation and the need for the channels to collapse between AGN outbursts. A new phenomenon that appears

to emerge from these simulations is the excitation of large-scale g-modes in the ICM that persist (and slowly dissipate their energy in to the ICM) long after the AGN activity has shut down. Tentative observational evidence for g-modes may exist in the cluster surrounding the radio galaxy 3C 401 (Reynolds et al. 2005a).

Any aspect of the system that alleviates the problem of channel formation will be potentially important in ensuring that the jet energy is thermalized in the ICM core. The initial cluster in our simulations start with a smooth, hydrostatic atmosphere. Real clusters have complex merger histories and turbulent ambient motions. This turbulence could help to distribute the energy from the jet more uniformly, or to oppose the formation of channels. Since the jets are powerful and higher velocity than even mildly supersonic ambient motions for most of their lifetime, it is uncertain how effective turbulence would be in preventing channels and distributing energy. Heinz et al. (2006) explored the possible effects of large scale winds of the distribution of heating. More importantly, real AGN jets may precess over significant angles (for a recent claim of jet precession in Perseus A, see Dunn et al. (2006)). ICM channels may become irrelevant and the kinetic energy thermalized in the ICM core if the jets can indeed precess over large angles on a timescale comparable to or less than the inner ICM cooling time. We will explore this possibility in detail in Chapter 5.

Finally, our results might be pointing to the importance of physics beyond ideal hydrodynamics. A very real possibility is that plasma transport processes may be crucial. Thermal conduction provides a way to move thermal energy across a temperature gradient in the ICM and may be essential in helping the system avoid the cooling catastrophe. There is a large heat reservoir in the outer parts of the cluster that could be exploited to warm the core. Furthermore, thermal conduction and viscosity could be important mechanisms for thermalizing the sound waves and weak shocks driven into the ICM by the radio-galaxy activity. All of these
processes may delay the ICM cooling sufficiently to allow the jet-blown channels to refill, thus allowing a long-term balance to be established. We note that there is growing observational evidence for the action of thermal conduction in real clusters (e.g., Perseus A; Fabian et al. (2005a,b)). Magnetic fields are also certainly present in the ICM, and at some level, an MHD (or even a kinetic theory) treatment may be necessary. This will have an important effect on the stability of the buoyantly rising bubbles and the mixing of the high-entropy radio plasma with the thermal ICM. But magnetic fields may have more profound consequences. Thermal conduction and plasma viscosity will both become very anisotropic due to the suppression of transport processes perpendicular to magnetic field lines, leading to qualitatively new fluid instabilities (the magnetothermal instability Balbus (2000) and the fire-hose instability). Cosmic ray pressure may also be relevant to driving ICM convection (Chandran 2004). In a cooling cluster core with an embedded AGN, one can speculate that these instabilities drive ICM turbulence, which may itself be a crucial ingredient in the dynamics of the system. Of course, this would be above and beyond any turbulence or bulk motions due to the merging history of the cluster.

An important consequence of this work is that it firmly underscores the inadequacy of AGN feedback simulations that model AGN feedback through isotropically inflated bubbles at prescribed locations in the ICM. Models that involve pre-inflated bubbles, or injecting energy isotropically at given locations in the ICM, will fail to capture precisely that aspect of the physics that turn out to be crucial to the failure of our models — the development of a low-density channel along which the AGN outflow can travel unimpeded thereby carrying its kinetic energy out of the cooling core. For this reason, it is vital that jet dynamics be included if a model is to properly address AGN-halted cooling in a cluster core.

4.5 Conclusion

We have performed a set of high-resolution three-dimensional simulations of jetted AGN embedded in the cooling ICM cores of galaxy clusters in which the AGN power reacts in response to the cooling of the cluster gas. We make the first attempt to model regulated AGN feedback in the hope of achieving a long-term balance between radiative cooling and jet heating. However, we fail to construct a model in which AGN heating comes into long-term balance with the radiative cooling. The early time jet activity is extremely effective at clearing a channel through the ICM core. The jet at later times flows freely down this channel and hence deposits its kinetic energy well outside the region which is starting to undergo catastrophic cooling. This essential behavior is robust to changing the efficiency or time delay characterizing the AGN feedback. Hence, we argue that some vital ingredient of AGN feedback is missing from our model; this may include more realistic jets, ambient ICM motions/turbulence, jet-precession, or additional ICM microphysics.Our results also highlight the absolute necessity of following the jet dynamics in any attempt to address radio galaxy heating of the ICM.

Chapter 5

Precessing Jet Models

5.1 Introduction

In the simulations of Chapter 4 we showed that with a fixed jet and a closed feedback loop, an AGN could not efficiently heat the core of a relaxed, cooling cluster. The specific failure mechanism of the jet heating in those simulations was the formation of a low-density channel in the ICM which prevented efficient heating of the cluster center regardless of the amount of energy provided by the jet. In this chapter, we study a set of simulations that remove this limitation by introducing precession in the direction of the jet. This will allow us to determine if there are other factors besides the unnecessary symmetry in our previous simulations which prevented efficient heating.

Recently, Sternberg and Soker (2008) have shown (in a restricted axisymmetric geometry) that it is possible to inflate spherically bubbles with thin jets if the jets are precessing. Kurosawa and Proga (2008) have shown that a precessing accretion disk can produce an outflow which also precesses. Leaving aside the question of the driving force behind the disk (and therefore jet) precession, we have performed a set of simulations where the AGN jet is allowed to precess around a set point.

This allows us to examine the effect of the localization of energy deposition on the heating and cooling of the cluster gas. This addresses the tendency of fast jets to drill channels in the ICM and deposit their energy far from the core. These simulations are limited to ideal hydrodynamics (along with optically thin cooling). This allows us to examine effects that depend only on the varying orientation of the jet, even though there is reason to believe that non-ideal effects may be important in the overall evolution.

Observationally, there is some evidence for precessing jets in cooling core clusters. Specifically, Dunn et al. (2006) present evidence for the precession of the jet in NGC 1275 (3C 84) based on the alignment of pairs of bubbles in the Perseus cluster. They use the position of the north and south jets as well as a pair of inner bubbles, a pair of outer bubbles, a pair of ghost bubbles, and a pair of ancient ghost bubbles (with a combination of X-ray and radio data) to get a precession angle and a timescale for precession. Their best fit is a precession opening angle of 50° and a precession timescale of approximately 7.5×10^7 yr. The presence of these misaligned pairs of relic radio bubbles makes a strong case for jets which precess with a timescale less than the total lifetime of a bubble. Although other causes are possible (e.g., environmental effects altering the trajectory of a rising bubble), they would be less likely to produce pairs of aligned bubbles.

In this chapter, we examine a set of three-dimensional hydrodynamic simulations of precessing jets in cooling core clusters. In addition we have rerun several nonprecessing cases specifically for use as comparisons. In Section 5.2 we will discuss the details of the setup and the simulations performed. Section 5.3 will describe the analysis done on the simulations as well as the results. A discussion of the results will be done in Section 5.4 followed by conclusions in Section 5.5.

Name	Precession	Cooling	Feedback
Run I	No	No	No
Run II	Yes	No	No
Run III	No	Yes	No
Run IV	Yes	Yes	No
$\operatorname{Run} V$	No	Yes	Yes
Run VI	Yes	Yes	Yes

Table 5.1: List of simulations and included effects.

5.2 Setup

As in Chapter 4, our goal is to set up a consistent, quasi-steady state solution for a relaxed cooling cluster being heated by a central, jetted AGN. Our model cluster starts out spherically symmetric, isothermal, and in hydrostatic equilibrium. The details of the simulation setup and the computational grid used are the same as those in Chapter 4 and will not be repeated here. All simulations were performed using the ZEUS-MP MHD code in pure hydrodynamic mode. ZEUS is discussed in detail in Chapter 2.

There are three additional effects that varied in our set of simulations: jet precession, radiative cooling, and feedback. A list of simulations and the included processes in each is given in Table 5.1.

In simulations with a precessing jet, the basic setup is the same as the fixed jet runs only the velocity of the jet is varied with time so the jet no longer points at a fixed location (directly down the x1 axis). The jet axis is offset from the $\theta = 0$ axis by a fixed angle and rotates around the axis with a constant angular velocity. In simulations with feedback, the jet injection velocity is then also modified by the mass accretion rate across the inner boundary as described below. The angle between the jet axis and $\theta = 0$ (the half opening angle) is set at 45°. The geometry of the precessing jet can be seen in Figure 5.1. Although this is at the extreme end of what one would expect for precessing jets, if it is incapable of halting cooling, it is unlikely that less dramatic precession will fare any better. The period of precession is 100 Myr (close to the dynamical time of the central galaxy). As in our previous simulations, there is a reflecting boundary on the $\theta = \pi/2$ plane (we only simulated half of a sphere and one jet). Although this is an unphysical choice, we believe that a reflecting boundary is the most conservative option. A second, counter jet could produce a backflow which would interact with the backflow from the first jet (and the cooling flow) in a limited area while the reflecting boundary does it over the entire plane. So a real jet can only do at most as much work on its counterpart (and the center) as our model jet (and likely much less) and therefore we do not feel that the reflecting boundary could be responsible for preventing a balance with cooling.

It is well known that relaxed clusters are bright in the X-ray wavelengths due to thermal bremsstrahlung radiation. This is exactly the driving force behind the cooling flow problem. To model this, we have included optically thin cooling in the ZEUS code as in Equation 4.1 from Chapter 4.

The final addition piece selectively included in our simulations is a feedback prescription. In principle, the feedback needed in a cooling cluster can be fairly simple: the AGN power (in the jets) needs to related to some state of the cooling gas in the cluster core in some way that allows the jet to balance cooling. In reality, there are innumerable complications possible. Even before worrying about accretion and jet launching from the central AGN, the cooling gas has the possibility of star formation or any other cool end product between when it reaches the central galaxy and when it reaches the central accretion disk.

We follow the same feedback scheme as Chapter 4, which is the simplest feedback scheme that could still capture as much of the physics as possibly while having very



Figure 5.1: Geometry for precessing jet. $\alpha = 45^{\circ}$ and β varies with time.

few tunable parameters. As the gas cools, the inner most (and likely coldest) gas can leave the grid across the inner radial boundary. The code keeps track of how much gas flows (irreversibly) across this boundary. The quantity of gas is then used to set the kinetic energy of the jet by way of the velocity of the jet material and Equation 4.2. As in Chapter 4, the only truly free parameter in this is η , the efficiency. The maximum theoretical efficiency for the supermassive black hole to convert accreting material into energy is $\eta = 0.1$. We cannot account for where the gas goes after it has left our grid (but before it reaches the central black hole), so our efficiency must be less than this. We have chosen an efficiency of $\eta = 10^{-4}$. See Chapter 4 for a further exploration of different efficiencies and other variations to the feedback scenarios of this type.

To account for the time required for the gas to reach the central black hole and accrete, a delay of 100 Myr was added between when the gas leaves the grid and when that value of \dot{M} is used to calculate a jet velocity. This time was set to match the dynamical time of the central galaxy.

5.3 Analysis and Results

5.3.1 Basic Properties

The most direct analysis tool we have is to look at slices through the various hydrodynamic variables. For a single burst precessing jet, see Figures 5.2(a) and 5.2(b). There are several features that stand out. First we note that when viewing a slice, two distinct lobes may be seen, surrounded by a single (nearly spherical) outer shock wave. Unlike the case for a non-precessing jet (see Figure 4.4), there is not a single contact discontinuity. Instead, the geometry of the contact discontinuity is much more complicated, with significant disruptions due to the KH and RT instabilities. There is also a strong shock clearly visible in both the density and pressure maps around the lobe with the most recent jet activity (the right lobe). This shock appears very strong, most likely because due to the precession, the jet impacts denser material at later times than it does in the fixed jet case (where it only impacts dense material initially).



(a) Density slice



(b) Pressure slice

Figure 5.2: Slice through density (top) and pressure (bottom) for single burst precessing jet (run II) at t=4.0.

5.3.2 The Cooling Flow

The most basic quantitative tool to investigate the state of the cooling flow is the mass accretion rate across the inner radial boundary. Although this can tell us little about the actual energy deposition or temperature structure of the cluster gas, it does show if the simulation reaches a stable solution during its run time. For a



Figure 5.3: Mass accretion for a precessing jet with delayed feedback.

precessing jet with delayed feedback (run VI), the mass accretion rate can be seen in Figure 5.3. We can compare this to the similar plot for delayed feedback without precession (run V), seen in Figure 5.4. It is clear from the values of \dot{M} that the a precessing jet (with feedback) does no better than a fixed-angle jet (with feedback) at stabilizing cooling. We will return to this point in Section 5.4. This inability to halt the cooling was seen in all precessing jet simulations we ran (runs II, IV, and VI).



Figure 5.4: Mass accretion for non-precessing jet with delayed feedback.

5.3.3 Simulated X-ray Maps

X-ray maps were generated to compare the output more closely to observations. The X-ray emissivity was assumed to be proportional to $\rho^2 T^{1/2}$. This was then integrated along a line of sight perpendicular to the $\theta = 0$ axis to produce the maps shown in Figures 5.5 and 5.6. These maps clearly show that the rapid cooling is proceeding in the equatorial regions of the cluster. Even with our very wide angle jets, as long as there is a region where the gas can "get around" the jets, it will, and this allows runaway cooling to proceed. It is interesting to note that the X-ray bright regions around the base of the jet (where the majority of the cooling is taking place) are quite asymmetric in the case of the precessing jet unlike the fixed angle jet.



Figure 5.5: X-ray brightness at t=2.0 for run I.

5.3.4 Tracing the Energy

Although we know that the jets are not balancing cooling, we do know that they are putting a large amount of energy into the system. It is likely that much of the energy resides in sound waves excited by the jets. To understand the excitation of sound waves, we must trace the spatial distribution of the injected energy. To evaluate this, the total energy in each cell was calculated as

$$E_{tot} = \frac{1}{2}\rho v^2 + e_{int} - \rho\Phi, \qquad (5.1)$$

and then the difference between this at two consecutive output times $(E_{tott} - E_{tott-1})$ was taken. The results for several times can be seen in Figure 5.7 and 5.8. This



Figure 5.6: X-ray brightness at t=2.0 for run II.

shows that there are regions of fast moving, high energy material excited by the jets. We use this particular quantity because it can pick out waves. In a complicated geometry (both for the system and the waves), there is no perfect way to isolate waves from a somewhat messy background, but this quantity allows us to come close to uniquely picking them out. Ignoring the highest value regions (those with a value near 1 and an orange color) which come from recently injected jet material, the precessing jet case has more energy flux in the waves, and the waves present are generally higher energy than the waves in the non-precessing case.

Possibly more important than the strength of the waves is the fact that the precessing jets also have these waves spread out over a larger region. While these



Figure 5.7: Change in total energy for run V at t=2.5, 4.5, 6.5, and 7.5.

do not heat the background gas under the assumption of ideal hydrodynamics, they do represent energy that has been added to the system and could be a major source of heating in a more realistic setting (e.g., with the addition of thermal conduction, realistic viscosity, or any other process which can dissipate and thermalize sound waves). They also represent the reason these jets fail to halt the cooling flow. Even though energy is injected to balance cooling, the energy is converted into waves and under our assumptions waves cannot deposit energy into the gas.

In Figures 5.7 and 5.8, we compare two runs with feedback (runs V and VI). Although this is fine for a qualitative comparison, it presents a problem for directly comparing the results. Since the power of the jets depend on the state of the ICM at



Figure 5.8: Change in total energy for run VI at t=2.5, 4.5, 6.5, and 8.5.

some time, we cannot pick a snapshot where the jet in each simulation has injected the same amount of energy into the cluster (unless we look at different times, but then the ICM will have cooled by different amounts).

To directly compare the energies of the precessing and non-precessing jets, we will look at the fixed time, single burst simulations: run I and II. Starting with the change in total energy (Equation 5.1), we sum over the total change in energy at a radius r. We do this for a region on the jet axis ($\theta \leq 60^{\circ}$ which contains the entire jet) and a region off the jet axis ($\theta > 60^{\circ}$). Figures 5.9(a) and 5.9(b) show this for the fixed jet (run I). Figures 5.10(a) and 5.10(b) show this for the precessing jet (run II). If, as discussed above, this total energy difference preferentially picks out energy contained in waves, then these show that the precessing runs have more energy present in waves than the fixed jet runs. This is true both close to the jet axis (top panels) and away from the jet axis (bottom panels). It is also important to notice that the near-axis energy are not on the same scale for the two runs since the fixed jet run would not show any detail on the scale used for the precessing jet run (emphasizing the large difference between the two).

Figures 5.11(a) and 5.11(b) show the projected X-ray brightness for run I and II respectively through the line at $\theta = 70^{\circ}$ (far away from the jet) at t = 3.0. The first thing that we notice is that the peak value for the precessing jet is 3.5 times greater than the peak for the fixed jet (which had a peak X-ray brightness of roughly $10\times$ greater than the background value for the wave). Outside of the highest values, the rest are similar between the two (so we do not just have a different distribution of the same overall energy). This is additional evidence that the precessing jet is capable of pushing more energy (in the form of waves) away from the jet axis. In pure hydrodynamics this energy is not available for heating, but the presence of it is interesting in itself and may point to the ultimate origin of the balance between heating and cooling. This is also interesting as the previous search for waves used a difference in total energy which is very clearly not an observable quantity. This is much closer to an observable quantity of real systems.

5.4 Discussion

The first and most obvious thing to note is that simply allowing the jet to precess around a fixed point does not give the jet the ability to balance cooling in any self regulated way. This can be seen in the mass accretion plots (Figures 5.3 and 5.4). They also do not delay the onset of catastrophic cooling beyond what fixed jets do



Figure 5.9: ΔE at t = 2.0 for the inner region (on the jet axis) and the outer region (off the jet axis) for run I.



Figure 5.10: ΔE at t = 2.0 for the inner region (on the jet axis) and the outer region (off the jet axis) for run II. For the top panel, note the increase in vertical scale compared to Figure 5.9(a).



Figure 5.11: Projected X-ray brightness through $\theta = 70^{\circ}$ at t = 3.0.

 $(\sim 40 \,\mathrm{Myr}).$

The interesting thing is that we have broken the symmetry present in Chapter 4 and thereby were able to prevent the formation of low density channels in the ICM, but were still not able to offset cooling. Clearly, the formation of low density channels is not the only failure mode for feedback in these types of simulations.

This does not mean that the precessing jets have the same impact on the ICM as the fixed jets. Rather, the precessing jets excite waves in the ICM beyond what a fixed angle jet can do. These waves are stronger in the precessing jet case and cover larger area, especially in the equatorial plane where most of the cooling takes place. In ideal hydrodynamics, these waves represent energy lost to the system which cannot help balance radiative losses. In reality however, waves do dissipate as they travel through a plasma, providing a possible way for AGN jets to heat the cluster core.

Physically, it is not hard to understand why precessing jets should be better at producing waves in the core of the cluster than a fixed jet. The fixed axis jet produces an outer shock wave when it first impacts the ICM in much the same way as the precessing jet. After that, the jet will impact only with material at the head of the jet which moves out radially into progressively less dense parts of the ICM. If the jet were to pulse it is possible that it could produce large amounts of waves at the head of the jet, but this will have left the cluster core after a fraction of the cooling time.

The precessing jet on the other hand does not appear as a steady flow to a given region of the ICM even if the actual linear velocity and power of the jet remains the same. It gets to impact denser material for much longer and will actually stop hitting a given region of the ICM. This avoids the underdense channels and keeps wave production closer to the core (and increases wave production over all). Although only a limited number of precessing jet models were run, they are fairly extreme examples which should easily capture all the relevant effects. With a 45° precession angle, our jets are at the outer edge of the physically plausible range for jet precession angles. With only hydrodynamic effects available, a larger precession angle will invariably lead to a greater effect on the atmosphere. Therefore, our large angle case puts an upper limit what jets with a smaller precession angles could do. Along the same lines, it seems obvious that if a jet had a 90° precession angle (that is it covered the entire half-circle) over a time frame less than the cooling time, it could easily balance cooling. This would effectively be the same thing as isotropic (or one-dimensional) models. While models like this have shown some success (Kaiser and Binney (2003) for example) they are clearly not the same as jet models; instead they are effectively isotropic heating models. While it is not impossible that there is an isotropic component to AGN based heating, there is clear observational evidence for the existence of jets, so this type of numerical solution is not realistic enough to pursue.

We are left with several possibilities. First, there is the possibility that our simulations (both fixed and precessing jets) are simply not capturing some (hydrodynamic) aspect of the system that is vital for our model to properly match up with reality (or even that our resolution is insufficient to capture the relevant details). Especially in light of the results of Chapter 3 on the ability of jets to alter the entropy of the background, we do not believe this is the case. More likely is the case where the missing piece is some additional real physical effect that can efficiently dissipate the wave energy (magnetic fields, thermal conduction, etc.) that is left out of our simulations by design is possible and will be discussed further in Chapter 6.

5.5 Conclusion

Extending our previous work in pure hydrodynamic models of AGN jets embedded in cooling core cluster, we have examined the effects of jet precession on the temperature structure of the ICM. We have shown that even in the limiting case of an large rotation angle, fast, hydrodynamic jets are unable to balance cooling when combined with an \dot{M} based feedback loops. Not only are these precessing jets unable to halt catastrophic cooling, they are not able to significantly delay the onset of catastrophic cooling.

We do however find an important difference between the failure mode of the precessing jet models and the fixed jet models. The fixed jet models fail due to a low density channel which prevents the jet from interacting with the cluster core after a short time. The precessing jets are able to easily avoid that problem. Instead, the precessing jets appear to convert more of their kinetic energy into sound waves. These waves are not confined to regions near the jet axis and occur mainly near the cluster core. In ideal hydrodynamics, these waves cannot heat the background gas and therefore are unable to balance the energy budget, but they do hint to the ability of precessing AGN jets to have a major impact on the cluster core's thermal evolution if additional physics beyond ideal hydrodynamics come into play.

Chapter 6

Conclusions and Future Directions

6.1 Conclusions

To examine the action of jetted AGN on the ICM in a cooling flow cluster, we have performed three sets of simulations. The first of these was a set of axisymmetric jet models to study varying jet parameters. The second was a set of three dimensional models with an attempt at a feedback loop between the jet and the ICM. The third was set of models with precessing jets added to the feedback models.

In our parameter survey of two-dimensional, axisymmetric jet models (Chapter 3), we have shown that by varying the jet parameters, it is possible to produce two distinct populations of systems, cocoon bounded and non-cocoon bounded. The total injected energy (via the jet lifetime) did not have as strong an effect on the resulting morphology. In the cocoon bounded systems, the jet inflates a high entropy bubble which remains distinct for long times as it rises in the atmosphere. In the non-cocoon cases, the inflated structure lacks a sharp boundary and mixes with the background from early on. There is a possible correspondence between these two classes and the Faranoff-Riley types, although we cannot definitively say so. We have also shown that the jets are capable of converting a large fraction (50% - 80%) of their energy into the gravitational energy of the ICM causing it to "puff-up".

In our two sets of feedback simulations, we have shown that it is hard to couple fast jets to the cooling ICM under the assumptions of ideal hydrodynamics. The first set of simulations (Chapter 4) focuses on the feedback prescription. We set up a jet in a cooling cluster where the jet power was determined by \dot{M} , the amount of mass to flow across the inner radial boundary. Although it was possible to delay catastrophic cooling (~ 50 Myr) we were not able to set up a stable feedback loop or even a long term offsetting of the cooling. This held true over a large range of efficiencies for the feedback. We also introduced (and varied) a delay between the mass accretion and its effect on the jet power. Much like the efficiency, this did not ultimately aid in preventing the cooling catastrophe. In all cases, the jet bores a low-density channel in the ICM and subsequently stops heating the core of the cluster. After that point, heating proceeds uninhibited in the equatorial regions and rapidly reaches catastrophic rates while the increasing powerful jet only encounters gas in the outer regions of the cluster.

Finally, a set of simulations was run to test the effect of jet precession on the cooling flow models in Chapter 5. Although the jets were able to avoid the low density channels which plagued our previous feedback simulations, they were still not able to effectively heat the cluster core. Even with a large precession angle, cooling proceeded uninhibited in the equatorial regions which were untouched by the jet. This showed that even in the absence of low-density channels, it is very hard to efficiently couple an AGN jet to the core regions of the ICM when restricted to ideal hydrodynamics.

Although none of our hydrodynamic feedback models were able to heat the cluster center efficiently enough to offset cooling, the precessing jets appear to transfer a large amount of their energy into sound waves. With only ideal hydrodynamics, any energy in sound waves is lost to the system. The sound waves were present in the fixed jet simulations, but are stronger in the precessing jet simulations, especially off the jet axis (precisely where most of the cooling takes place). If some additional physical process could extra energy from the sound waves and use it to heat the cluster core, it is possible that they could be a major contributor to the energy balance of the cluster core.

6.2 Future Directions

While our work here puts some restrictions on how AGN jet feedback could work, it is not necessarily fatal to the AGN jet feedback picture. There are several possible ways to extend these models beyond ideal hydrodynamics which could allow AGN jets to successfully set up a feedback loop and balance cooling.

The most promising aspect of these models is the sound waves of Chapter 5. A large amount of the energy injected by the jets is transported out of the core region (but in the nearly isotropic way that seem vital for a successful heating mechanism) in sound waves. If some dissipative process is active in the ICM, it is possible that these waves could heat the cluster core. Thermal conduction and viscosity are both capable of efficiently dissipating wave energy. Viscosity has been studied in the case of rising bubbles in Reynolds et al. (2005b).

Conduction has two effects which can help the energy balance. First, the hot outer regions of the cluster contain a large amount of thermal energy which could be brought into the core through thermal conduction. Second, conduction can dissipate the energy in sound waves. Conduction is given by

$$\mathbf{F}_{cond} = -\kappa \nabla T,\tag{6.1}$$

with

$$\kappa = \kappa_s = \frac{1.84 \times 10^{-5} T^{-5/2}}{\ln \lambda} \,\mathrm{ergs}^{-1} \mathrm{cm}^{-2} \mathrm{K}^{-1}, \tag{6.2}$$

where κ_s is the Spitzer value for an unmagnetized plasma from Spitzer (1965) and depends only on the temperature. We have done some initial work on conduction in feedback models and have implemented a thermal conduction module in ZEUS-MP. Unfortunately, numerical instabilities have prevented us from using it in the relevant simulations unless κ is lowered by several orders of magnitude from the Spitzer value (thereby rendering thermal conduction irrelevant for the evolution of the system). We do not believe that there are any physical instabilities at work here, so a solution to this obstacle is needed to pursue dissipation of sound waves through thermal conduction in these cluster simulations.

Magnetic fields provide another potential avenue for extending these models. Although field strengths are not completely certain in clusters, there is clear evidence for their presence in at least the μG range (see Pollack et al. (2005) for example). While not very strong, these are enough to be dynamically relevant. Unlike the dissipative processes, it is not immediately clear how magnetic fields could help heat the ICM core. The most obvious relevant effect they have is that magnetic fields alter the hydrodynamic instabilities (and introduce new instabilities not present in ideal hydrodynamics). Since instabilities have such an important effect on the evolution of the radio bubbles inflated by the jets, anything that modifies these instabilities could be an important ingredient in these systems.

Magnetic fields also modify the action of the dissipative processes such as thermal conduction. When magnetic fields are present, Equation 6.1 becomes

$$\mathbf{F}_{cond} = -\kappa' \hat{\mathbf{b}} (\hat{\mathbf{b}} \cdot \nabla T), \tag{6.3}$$

with

$$\kappa' = \kappa_s/3 \tag{6.4}$$

along the field lines. This almost completely prevents conduction across field lines while still allowing it (although reduced somewhat) along the field lines. Due to the dependence of field geometry (which is likely complicated), there is a large amount of work possible in this particular area before a definitive answer can be reached about the effects of magnetic fields and conduction in these models.

With an active numerical community using increasingly powerful codes and building on the tremendous amount of high quality X-ray data on clusters, it seems possible that a complete picture of the cooling flow problem will emerge soon. While work like this shows that it is possible the feedback from AGN may not be the entire answer as once thought, it still seems certain that they will play an important role.

Bibliography

- Armitage, P. J. and Reynolds, C. S. (2003). The variability of accretion on to Schwarzschild black holes from turbulent magnetized discs. MNRAS, 341:1041– 1050.
- Bîrzan, L., Rafferty, D. A., McNamara, B. R., Wise, M. W., and Nulsen, P. E. J. (2004). A Systematic Study of Radio-induced X-Ray Cavities in Clusters, Groups, and Galaxies. ApJ, 607:800–809.
- Balbus, S. A. (2000). Stability, Instability, and "Backward" Transport in Stratified Fluids. ApJ, 534:420–427.
- Basson, J. F. and Alexander, P. (2003). The long-term effect of radio sources on the intracluster medium. *MNRAS*, 339:353–359.
- Begelman, M. C. and Cioffi, D. F. (1989). Overpressured cocoons in extragalactic radio sources. ApJ, 345:L21–L24.
- Benson, A. J., Bower, R. G., Frenk, C. S., Lacey, C. G., Baugh, C. M., and Cole,S. (2003). What Shapes the Luminosity Function of Galaxies? *ApJ*, 599:38–49.
- Binney, J. and Tabor, G. (1995). Evolving Cooling Flows. MNRAS, 276:663-+.
- Blanton, E. L., Sarazin, C. L., McNamara, B. R., and Wise, M. W. (2001). Chandra Observation of the Radio Source/X-Ray Gas Interaction in the Cooling Flow Cluster Abell 2052. ApJ, 558:L15–L18.
- Böhringer, H., Nulsen, P. E. J., Braun, R., and Fabian, A. C. (1995). The interaction

of the radio halo of M87 with the cooling intracluster medium of the Virgo cluster. MNRAS, 274:L67–L71.

- Böhringer, H., Voges, W., Fabian, A. C., Edge, A. C., and Neumann, D. M. (1993). A ROSAT HRI study of the interaction of the X-ray-emitting gas and radio lobes of NGC 1275. MNRAS, 264:L25–L28.
- Brüggen, M. and Kaiser, C. R. (2001). Buoyant radio plasma in clusters of galaxies. MNRAS, 325:676–684.
- Brüggen, M. and Kaiser, C. R. (2002). Hot bubbles from active galactic nuclei as a heat source in cooling-flow clusters. *Nature*, 418:301–303.
- Burns, J. O. and Owen, F. N. (1977). A statistical investigation of radio sources in the directions of Zwicky clusters of galaxies. ApJ, 217:34–44.
- Carilli, C. L., Perley, R. A., and Harris, D. E. (1994). Observations of Interaction Between Cluster Gas and the Radio Lobes of Cygnus-A. MNRAS, 270:173–+.
- Carvalho, J. C. and O'Dea, C. P. (2002). Evolution of Global Properties of Powerful Radio Sources. II. Hydrodynamical Simulations in a Declining Density Atmosphere and Source Energetics. ApJS, 141:371–414.
- Chandran, B. D. G. (2004). Convection in Galaxy-Cluster Plasmas Driven by Active Galactic Nuclei and Cosmic-Ray Buoyancy. *ApJ*, 616:169–177.
- Chandran, B. D. G. and Rasera, Y. (2007). Convection and AGN Feedback in Clusters of Galaxies. ApJ, 671:1413–1433.
- Choi, Y., Reynolds, C. S., Heinz, S., Rosenberg, J. L., Perlman, E. S., and Yang, J. (2004). Observations of A4059 with Chandra, Hubble Space Telescope, and the Very Large Array: Unraveling a Complex Cluster/Radio Galaxy Interaction. *ApJ*, 606:185–195.
- Churazov, E., Brüggen, M., Kaiser, C. R., Böhringer, H., and Forman, W. (2001). Evolution of Buoyant Bubbles in M87. ApJ, 554:261–273.

- Cioffi, D. F. and Blondin, J. M. (1992). The evolution of cocoons surrounding light, extragalactic jets. *ApJ*, 392:458–464.
- Ciotti, L. and Ostriker, J. P. (2001). Cooling Flows and Quasars. II. Detailed Models of Feedback-modulated Accretion Flows. *ApJ*, 551:131–152.
- Clarke, D. A., Harris, D. E., and Carilli, C. L. (1997). Formation of cavities in the X-ray emitting cluster gas of Cygnus A. MNRAS, 284:981–993.
- Conroy, C. and Ostriker, J. P. (2007). Thermal Balance in the Intracluster Medium: Is AGN Feedback Necessary? *ArXiv e-prints*, 712.
- Dalla Vecchia, C., Bower, R. G., Theuns, T., Balogh, M. L., Mazzotta, P., and Frenk, C. S. (2004). Quenching cluster cooling flows with recurrent hot plasma bubbles. *MNRAS*, 355:995–1004.
- Dekel, A. and Silk, J. (1986). The origin of dwarf galaxies, cold dark matter, and biased galaxy formation. ApJ, 303:39–55.
- Dunn, R. J. H., Fabian, A. C., and Sanders, J. S. (2006). Precession of the supermassive black hole in NGC 1275 (3C 84)? MNRAS, 366:758–766.
- Edge, A. C., Stewart, G. C., and Fabian, A. C. (1992). Properties of cooling flows in a flux-limited sample of clusters of galaxies. *MNRAS*, 258:177–188.
- Fabian, A. C. (1994). Cooling Flows in Clusters of Galaxies. ARA&A, 32:277–318.
- Fabian, A. C., Reynolds, C. S., Taylor, G. B., and Dunn, R. J. H. (2005a). On viscosity, conduction and sound waves in the intracluster medium. *MNRAS*, 363:891–896.
- Fabian, A. C., Sanders, J. S., Allen, S. W., Crawford, C. S., Iwasawa, K., Johnstone,
 R. M., Schmidt, R. W., and Taylor, G. B. (2003). A deep Chandra observation of the Perseus cluster: shocks and ripples. *MNRAS*, 344:L43–L47.
- Fabian, A. C., Sanders, J. S., Ettori, S., Taylor, G. B., Allen, S. W., Crawford,C. S., Iwasawa, K., Johnstone, R. M., and Ogle, P. M. (2000). Chandra imaging

of the complex X-ray core of the Perseus cluster. MNRAS, 318:L65–L68.

- Fabian, A. C., Sanders, J. S., Taylor, G. B., and Allen, S. W. (2005b). A deep Chandra observation of the Centaurus cluster: bubbles, filaments and edges. MNRAS, 360:L20–L24.
- Fanaroff, B. L. and Riley, J. M. (1974). The morphology of extragalactic radio sources of high and low luminosity. MNRAS, 167:31P–36P.
- Feigelson, E. D., Wood, P. A. D., Schreier, E. J., Harris, D. E., and Reid, M. J. (1987). X-rays from the radio halo of M87. ApJ, 312:101–110.
- Felten, J. E., Gould, R. J., Stein, W. A., and Woolf, N. J. (1966). X-Rays from the Coma Cluster of Galaxies. ApJ, 146:955–958.
- Gammie, C. F. (1999). Efficiency of Magnetized Thin Accretion Disks in the Kerr Metric. ApJ, 522:L57–L60.
- Gursky, H., Solinger, A., Kellogg, E. M., Murray, S., Tananbaum, H., Giacconi, R., and Cavaliere, A. (1972). X-Ray Emission from Rich Clusters of Galaxies. ApJ, 173:L99+.
- Hardee, P. E. (1979). On the configuration and propagation of jets in extragalactic radio sources. *ApJ*, 234:47–55.
- Hardee, P. E. (1982). Helical and pinching instability of supersonic expanding jets in extragalactic radio sources. ApJ, 257:509–526.
- Hatch, N. A., Crawford, C. S., Fabian, A. C., Johnstone, R. M., and Sanders, J. S. (2005). Feedback from Brightest Cluster Galaxies. In Bulletin of the American Astronomical Society, volume 37 of Bulletin of the American Astronomical Society, pages 1239–+.
- Hayes, J. C., Norman, M. L., Fiedler, R. A., Bordner, J. O., Li, P. S., Clark, S. E., ud-Doula, A., and Mac Low, M.-M. (2006). Simulating Radiating and Magnetized Flows in Multiple Dimensions with ZEUS-MP. ApJS, 165:188–228.

- Heinz, S., Brüggen, M., Young, A., and Levesque, E. (2006). The answer is blowing in the wind: simulating the interaction of jets with dynamic cluster atmospheres. MNRAS, 373:L65–L69.
- Heinz, S., Choi, Y., Reynolds, C. S., and Begelman, M. C. (2002). Chandra ACIS-S Observations of Abell 4059: Signs of Dramatic Interaction between a Radio Galaxy and a Galaxy Cluster. ApJ, 569:L79–L82.
- Heinz, S., Reynolds, C. S., and Begelman, M. C. (1998). X-Ray Signatures of Evolving Radio Galaxies. ApJ, 501:126-+.
- Hicks, A. K. and Mushotzky, R. (2005). Star Formation Rates in Cooling Flow Clusters: A UV Pilot Study with Archival XMM-Newton Optical Monitor Data. *ApJ*, 635:L9–L12.
- Jones, T. W. and De Young, D. S. (2005). Magnetohydrodynamic Simulations of Relic Radio Bubbles in Clusters. ApJ, 624:586–605.
- Kaiser, C. R. and Binney, J. (2003). Simple models of cooling flows. MNRAS, 338:837–845.
- Krolik, J. H. (1999). Active galactic nuclei : from the central black hole to the galactic environment. Active galactic nuclei : from the central black hole to the galactic environment /Julian H. Krolik. Princeton, N. J. : Princeton University Press, c1999.
- Kurosawa, R. and Proga, D. (2008). Three-dimensional Simulations of Inflows Irradiated by a Precessing Accretion Disk in Active Galactic Nuclei: Formation of Outflows. ApJ, 674:97–110.
- Larson, R. B. (1974). Effects of supernovac on the early evolution of galaxies. MNRAS, 169:229–246.
- Ledlow, M. J. and Owen, F. N. (1996). 20 CM VLA Survey of Abell Clusters of Galaxies. VI. Radio/Optical Luminosity Functions. AJ, 112:9–+.

- Marti, J. M. A., Mueller, E., Font, J. A., Ibanez, J. M. A., and Marquina, A. (1997). Morphology and Dynamics of Relativistic Jets. ApJ, 479:151-+.
- McNamara, B. R., Wise, M., Nulsen, P. E. J., David, L. P., Sarazin, C. L., Bautz, M., Markevitch, M., Vikhlinin, A., Forman, W. R., Jones, C., and Harris, D. E. (2000). Chandra X-Ray Observations of the Hydra A Cluster: An Interaction between the Radio Source and the X-Ray-emitting Gas. ApJ, 534:L135–L138.
- McNamara, B. R., Wise, M. W., Nulsen, P. E. J., David, L. P., Carilli, C. L., Sarazin,
 C. L., O'Dea, C. P., Houck, J., Donahue, M., Baum, S., Voit, M., O'Connell,
 R. W., and Koekemoer, A. (2001). Discovery of Ghost Cavities in the X-Ray
 Atmosphere of Abell 2597. ApJ, 562:L149–L152.
- Mizuta, A., Yamada, S., and Takabe, H. (2001). Numerical Study of AGN Jet Propagation with Two Dimensional Relativistic Hydrodynamic Code. Journal of Korean Astronomical Society, 34:329–331.
- Norman, M. L. (2000). Introducing ZEUS-MP: A 3D, Parallel, Multiphysics Code for Astrophysical Fluid Dynamics. In *Revista Mexicana de Astronomia y Astrofísica Conference Series*, pages 66–71.
- Novikov, I. D. and Thorne, K. S. (1973). Astrophysics of Black Holes. In DeWitt, C. and DeWitt, B., editors, *Black Holes*, pages 343–450. Gordon and Breach, Paris.
- O'Dea, C. P., Baum, S. A., Mack, J., Koekemoer, A. M., and Laor, A. (2004). Hubble Space Telescope STIS Far-Ultraviolet Observations of the Central Nebulae in the Cooling-Core Clusters A1795 and A2597. ApJ, 612:131–151.
- Omma, H. and Binney, J. (2004). Structural stability of cooling flows. *MNRAS*, 350:L13–L16.
- Omma, H., Binney, J., Bryan, G., and Slyz, A. (2004). Heating cooling flows with jets. MNRAS, 348:1105–1119.
- Peres, C. B., Fabian, A. C., Edge, A. C., Allen, S. W., Johnstone, R. M., and White,

D. A. (1998). A ROSAT study of the cores of clusters of galaxies - I. Cooling flows in an X-ray flux-limited sample. *MNRAS*, 298:416–432.

- Peterson, J. R. and Fabian, A. C. (2006). X-ray spectroscopy of cooling clusters. *Phys. Rep.*, 427:1–39.
- Peterson, J. R., Paerels, F. B. S., Kaastra, J. S., Arnaud, M., Reiprich, T. H., Fabian, A. C., Mushotzky, R. F., Jernigan, J. G., and Sakelliou, I. (2001). X-ray imaging-spectroscopy of Abell 1835. A&A, 365:L104–L109.
- Pollack, L. K., Taylor, G. B., and Allen, S. W. (2005). VLA polarimetry observations of PKS2322-123: estimating magnetic fields in the Abell 2597 cluster. MNRAS, 359:1229–1236.
- Reynolds, C. S., Brenneman, L. W., and Stocke, J. T. (2005a). The interaction of 3C 401 with the surrounding intracluster medium. *MNRAS*, 357:381–386.
- Reynolds, C. S., Heinz, S., and Begelman, M. C. (2002). The hydrodynamics of dead radio galaxies. MNRAS, 332:271–282.
- Reynolds, C. S., McKernan, B., Fabian, A. C., Stone, J. M., and Vernaleo, J. C. (2005b). Buoyant radio lobes in a viscous intracluster medium. *MNRAS*, 357:242– 250.
- Robinson, K., Dursi, L. J., Ricker, P. M., Rosner, R., Calder, A. C., Zingale, M.,
 Truran, J. W., Linde, T., Caceres, A., Fryxell, B., Olson, K., Riley, K., Siegel,
 A., and Vladimirova, N. (2004). Morphology of Rising Hydrodynamic and Magnetohydrodynamic Bubbles from Numerical Simulations. ApJ, 601:621–643.
- Ruszkowski, M. and Begelman, M. C. (2002). Heating, Conduction, and Minimum Temperatures in Cooling Flows. ApJ, 581:223–228.
- Ruszkowski, M., Brüggen, M., and Begelman, M. C. (2004). Cluster Heating by Viscous Dissipation of Sound Waves. ApJ, 611:158–163.

Sanders, J. S., Fabian, A. C., Allen, S. W., Morris, R. G., Graham, J., and John-

stone, R. M. (2008). Cool X-ray emitting gas in the core of the Centaurus cluster of galaxies. *MNRAS*, 385:1186–1200.

- Sanderson, A. J. R., Ponman, T. J., and O'Sullivan, E. (2006). A statistically selected Chandra sample of 20 galaxy clusters - I. Temperature and cooling time profiles. *MNRAS*, 372:1496–1508.
- Spergel, D. N., Bean, R., Doré, O., Nolta, M. R., Bennett, C. L., Dunkley, J., Hinshaw, G., Jarosik, N., Komatsu, E., Page, L., Peiris, H. V., Verde, L., Halpern, M., Hill, R. S., Kogut, A., Limon, M., Meyer, S. S., Odegard, N., Tucker, G. S., Weiland, J. L., Wollack, E., and Wright, E. L. (2007). Three-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Implications for Cosmology. ApJS, 170:377–408.
- Spitzer, L. (1965). Physics of fully ionized gases. Interscience Tracts on Physics and Astronomy, New York: Interscience Publication, 1965, 2nd rev. ed.
- Sternberg, A. and Soker, N. (2008). Inflating fat bubbles in clusters of galaxies by precessing massive slow jets. MNRAS, 384:1327–1336.
- Stone, J. M. and Norman, M. L. (1992a). ZEUS-2D: A radiation magnetohydrodynamics code for astrophysical flows in two space dimensions. I - The hydrodynamic algorithms and tests. ApJS, 80:753–790.
- Stone, J. M. and Norman, M. L. (1992b). ZEUS-2D: A Radiation Magnetohydrodynamics Code for Astrophysical Flows in Two Space Dimensions. II. The Magnetohydrodynamic Algorithms and Tests. ApJS, 80:791–+.
- Sutherland, R. S. and Dopita, M. A. (1993). Cooling functions for low-density astrophysical plasmas. ApJS, 88:253–327.
- Tamura, T., Kaastra, J. S., Peterson, J. R., Paerels, F. B. S., Mittaz, J. P. D., Trudolyubov, S. P., Stewart, G., Fabian, A. C., Mushotzky, R. F., Lumb, D. H., and Ikebe, Y. (2001). X-ray spectroscopy of the cluster of galaxies Abell 1795
with XMM-Newton. A & A, 365:L87–L92.

- Vernaleo, J. C. and Reynolds, C. (2007a). AGN Heating of Cooling Flow Clusters: Problems with 3D Hydrodynamic Models. In Böhringer, H., Pratt, G. W., Finoguenov, A., and Schuecker, P., editors, *Heating versus Cooling in Galaxies* and Clusters of Galaxies, pages 248–+.
- Vernaleo, J. C. and Reynolds, C. S. (2006a). AGN Feedback and Cooling Flows: Problems with Simple Hydrodynamic Models. ApJ, 645:83–94.
- Vernaleo, J. C. and Reynolds, C. S. (2006b). Hydrodynamic Models of AGN Feedback in Cooling Core Clusters. In Bulletin of the American Astronomical Society, volume 38 of Bulletin of the American Astronomical Society, pages 1062–+.
- Vernaleo, J. C. and Reynolds, C. S. (2007b). Energetic Impact of Jet-Inflated Cocoons in Relaxed Galaxy Clusters. ApJ, 671:171–180.
- Vikhlinin, A., Kravtsov, A., Forman, W., Jones, C., Markevitch, M., Murray, S. S., and Van Speybroeck, L. (2006). Chandra Sample of Nearby Relaxed Galaxy Clusters: Mass, Gas Fraction, and Mass-Temperature Relation. ApJ, 640:691– 709.
- Young, A. J., Wilson, A. S., and Mundell, C. G. (2002). Chandra Imaging of the X-Ray Core of the Virgo Cluster. ApJ, 579:560–570.
- Zanni, C., Murante, G., Bodo, G., Massaglia, S., Rossi, P., and Ferrari, A. (2005).
 Heating groups and clusters of galaxies: The role of AGN jets. A&A, 429:399–415.
- Zwicky, F. (1937). On the Masses of Nebulae and of Clusters of Nebulae. ApJ, 86:217–+.