Title
Gas Dynamics in Galaxy Clusters

Author
McCourt, Michael Kingsley

Publication Date
2014

Peer reviewed|Thesis/dissertation

Gas Dynamics in Galaxy Clusters

By

Michael Kingsley McCourt, Jr.

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Astrophysics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Eliot Quataert, Chair
Professor Chung-Pei Ma
Professor Christopher McKee
Professor William Holzapfel

Spring 2014

Gas Dynamics in Galaxy Clusters

Abstract

Gas Dynamics in Galaxy Clusters

by

Michael Kingsley McCourt, Jr.

Doctor of Philosophy in Astrophysics

University of California, Berkeley

Professor Eliot Quataert, Chair

Galaxy clusters are the most massive structures in the universe and, in the hierarchical pattern of cosmological structure formation, the largest objects in the universe form last. Galaxy clusters are thus interesting objects for a number of reasons. Three examples relevant to this thesis are:

1. **Constraining the properties of dark energy:** Due to the hierarchical nature of structure formation, the largest objects in the universe form last. The cluster mass function is thus sensitive to the entire expansion history of the universe and can be used to constrain the properties of dark energy. This constraint complements others derived from the CMB or from Type Ia supernovae and provides an important, independent confirmation of such methods. In particular, clusters provide detailed information about the equation of state parameter w because they sample a large redshift range $z \sim 0 - 1$.

2. **Probing galaxy formation:** Clusters contain the most massive galaxies in the universe, and the most massive black holes; because clusters form so late, we can still witness the assembly of these objects in the nearby universe. Clusters thus provide a more detailed view of galaxy formation than is possible in studies of lower-mass objects. An important example comes from x-ray studies of clusters, which unexpectedly found that star formation in massive galaxies in clusters is closely correlated with the properties of the hot, virialized gas in their halos. This correlation persists despite the enormous separation in temperature, in dynamical time-scales, and in length-scales between the virialized gas in the halo and the star-forming regions in the galaxy. This remains a challenge to interpret theoretically.

3. **Developing our knowledge of dilute plasmas:** The masses and sizes of galaxy clusters imply that the plasma which permeates them is both very hot ($\sim 10^8\,\mathrm{K}$) and very dilute ($\sim 10^{-2}\,\mathrm{cm}^{-3}$). This plasma is collisional enough to be considered a fluid,

but collisionless enough to develop significant anisotropies with respect to the local magnetic field. This interesting regime is one of the frontiers in theoretical studies of fluid dynamics. Unlike other astrophysical environments of similar collisionality (e. g. accretion disk coronae), galaxy clusters are optically thin and subtend large angles on the sky. Thus, they are easily observed in the x-ray (to constrain thermal processes) and in the radio (to constrain non-thermal processes) and provide a wonderful environment to develop our understanding of dilute plasmas.

This thesis studies the dynamics of the hot gas in galaxy clusters, which touches on all three of the above topics.

Chapter 2 shows that galaxy clusters are likely to be unstable to a new, vigorous form of convection. As a dynamical process which involves thermodynamic and magnetic properties of the gas, this convection bears directly on our understanding of the physics of dilute plasmas. Furthermore, by moving metals and thermal energy through the cluster, convection may change the cooling rate of the gas and thus significantly impact the process of galaxy formation. Cluster convection also impacts the use of clusters as cosmological probes. Convection may drive turbulence in clusters with mean Mach numbers of order-unity. This changes the force balance in clusters, decreasing the thermal energy of a cluster of a given mass. Current methods for using clusters to constrain dark energy rely on observational probes of the thermal energy as a proxy for total mass. The accuracy of these methods depends on how vigorous cluster convection is.

Chapter 3 studies thermal instability in galaxy clusters. I argue that clusters are all likely to be thermally unstable, but that this instability only grows to large amplitude in a subset of systems. Later studies have applied this result to galaxy formation in clusters and shown that one can reproduce some features of the well-known non-self-similarity at the high mass end of the galaxy luminosity function.

Chapters 4 and 5 extends my work on convection (and, eventually, thermal instability) to consider the cosmological context of galaxy formation. This work aims to remove any arbitrary initial and boundary conditions from my simulations and is an important step toward a self-consistent model for the plasma physics in clusters.

Contents

Acknowledgments

This work was supported in part by NASA Grant NNX10AC95G, NSF-DOE Grant PHY-0812811, Chandra theory grant TM2-13004X, and the David and Lucile Packard Foundation.

Computing time was provided by the National Science Foundation through the Teragrid resources located at the National Center for Atmospheric Research and the Pittsburgh Supercomputing Center, by the *Henyey* cluster at UC Berkeley, supported by NSF AST Grant 0905801, and through XSEDE under grant numbers AST090038 and AST090039.

I am grateful to Mark Voit and Megan Donahue for interesting and helpful conversations as I completed the work in Chapter 3, and to Steve Balbus for clarifying my discussion of the previous literature in that chapter. I thank the anonymous referee for helping to clarify my discussion, particularly in section 3.4. Mike McDonald helpfully suggested several clusters I left out from figure 3.11 in an early draft of that paper. I am also thankful for the hospitality of the Kavli Institute for Theoretical Physics (KITP) at UC Santa Barbara, where I performed some of this work. KITP is supported in part by NSF grant number PHY05-51164.

I am grateful to Chung-Pei Ma, James McBride, and Claude-André Faucher-Giguère for helpful discussions about the structure of dark-matter halos and their mass accretion histories. These conversations shaped the work in Chapter 4. I also thank Mark Voit for pointing out applications of the $T(M)$ relationship, and I thank the anonymous referee for suggesting several improvements to the original manuscript.

I thank my dissertation committee for helpful conversations which shaped this thesis.

I made my figures using the open-source program Tioga. This research has made use of NASA's Astrophysics Data System. This dissertation was typeset using the ucastrothesis LaTeX template.

Chapter 1

Introduction

This is a thesis about the hot gas in galaxy clusters. Galaxy clusters are the largest objects in the universe and consist primarily of dark matter (84%), hot gas (14%), and stars (2%). Though dark matter represents the bulk of the mass in clusters, it is invisible and its behavior is comparatively well-understood. The hot gas, on the other hand, is amazingly complex and encodes valuable information about the formation of clusters and about structure in the universe. As I discuss later on, understanding the dynamics of this gas is also essential for efforts to use galaxy clusters to better constrain dark energy and to understand the future of the universe (e. g. will the universe last forever?).

The physics underlying the research in this thesis is not explicitly tied to galaxy clusters, and many of the processes I describe do in fact extend to lower mass systems like groups and individual galaxies. However, the hot spatially extended halo gas is much more easily detected in galaxy clusters, so its dynamics is better constrained in these objects. X-ray observations of the Perseus cluster (e. g. Fabian et al. 2011) exemplify the rich phenomenology of the intracluster medium: the gas contains bubbles, ripples, spiral features, and consists of multiple roughly co-spatial phases at different densities and temperatures.

The hot gas in galaxy clusters also presents an opportunity to advance plasma physics by testing it in new environments. The gas in clusters is so hot (several hundred million degrees) that it glows in x-rays. It is also extremely dilute, roughly twenty billion times more tenuous than the best laboratory vacuums on earth. However, clusters are also *extremely* large, some twenty-one orders of magnitude larger than tokamaks and other laboratory plasma devices. This is a relatively unstudied regime in fluid dynamics, and Chapter 2 describes how convection operates very differently in galaxy clusters than it does in other environments such as a pot of boiling water, earth's oceans and atmosphere, or the interior of the sun. Chapters 2 and 5 show several simulations of this convection.

Studying the hot gas in galaxy clusters also yields valuable information about how structures form in the universe. Chatper 3 describes the filaments of cold, atomic gas observed in some clusters and presents a theory for their formation mechanism. Later studies have applied this result to argue that the hot gas in clusters plays an essential role in producing the stars, massive galaxies, and super-massive black holes within them.

Finally, Chapters 4 and 5 focus on the growth of clusters over cosmic time. I show how the hot gas in clusters records their formation histories and I use this result to build self-consistent models for the gas profiles in galaxy clusters.

Initially it may seem esoteric to study a dilute gas which floats billions of light-years from Earth, but these applications make it an exciting subject!

Outline

Following this introduction, my thesis is split into four chapters.

In Chapter 2, I study the effects of anisotropic thermal conduction on low-collisionality, astrophysical plasmas using two and three-dimensional magnetohydrodynamic simulations. Balbus (2000) and Quataert (2008) showed that dilute plasmas are buoyantly unstable for either sign of the temperature gradient: the heat-flux-driven buoyancy instability (HBI) operates when the temperature increases with radius while the magnetothermal instability (MTI) operates in the opposite limit. In contrast to previous results, I show that, in the presence of a sustained temperature gradient, the MTI drives strong turbulence and operates as an efficient magnetic dynamo (akin to standard, adiabatic convection). Together, the turbulent and magnetic energies contribute up to $\sim$10% of the pressure support in the plasma. In addition, the MTI drives a large convective heat flux, $\sim 1.5\% \times \rho c_s^3$. These findings are robust even in the presence of an external source of strong turbulence. My results on the nonlinear saturation of the HBI are consistent with previous studies but I explain physically why the HBI saturates quiescently by re-orienting the magnetic field (suppressing the conductive heat flux through the plasma), while the MTI saturates by generating sustained turbulence. I also systematically study how an external source of turbulence affects the saturation of the HBI: such turbulence can disrupt the HBI only on scales where the shearing rate of the turbulence is faster than the growth rate of the HBI. In particular, my results provide a simple mapping between the level of turbulence in a plasma and the effective isotropic thermal conductivity. I discuss the astrophysical implications of these findings, with a particular focus on the intracluster medium of galaxy clusters. The work in Chapter 2 has previously been published in McCourt et al. (2011a).

In Chapter 3, I study the interplay among cooling, heating, conduction, and magnetic fields in gravitationally-stratified plasmas using simplified, plane-parallel numerical simulations. Since the physical heating mechanism remains uncertain in massive halos such as groups or clusters, I adopt a simple, phenomenological prescription which enforces global thermal equilibrium and prevents a cooling-flow. The plasma remains susceptible to *local* thermal instability, however, and cooling drives an inward flow of material. For physically plausible heating mechanisms in clusters, the thermal stability of the plasma is independent of its convective stability. I find that the ratio of the cooling timescale to the dynamical timescale $t_{\rm cool}/t_{\rm ff}$ controls the non-linear evolution and saturation of the thermal instability: when $t_{\rm cool}/t_{\rm ff} \lesssim 1$, the plasma develops extended multi-phase structure, whereas when $t_{\rm cool}/t_{\rm ff} \gtrsim 1$ it does not. When thermal conduction is anisotropic with respect to the magnetic

field, the criterion for multi-phase gas is essentially independent of the thermal conductivity of the plasma. This criterion for local thermal instability to produce multi-phase structure is an extension of the cold vs. hot accretion modes in galaxy formation that applies at all radii in hot halos, not just to the virial shock. I show that this criterion is consistent with data on multi-phase gas in the ACCEPT sample of clusters; in addition, when $t_{\rm cool}/t_{\rm ff} \gtrsim 1$, the net cooling rate to low temperatures and the mass flux to small radii are suppressed enough relative to models without heating to be qualitatively consistent with star formation rates and x-ray line emission in groups and clusters. The work in Chapter 3 has previously been published in McCourt et al. (2012).

In Chapter 4, I present a spherically symmetric model for the origin and evolution of the temperature profiles in the hot plasma filling galaxy groups and clusters. I find that the gas in clusters is generically not isothermal, and that the temperature declines with radius at large distances from the cluster center (outside the core- and scale radii). This temperature profile is determined by the accretion history of the halo, and is not quantitatively well-described by a polytropic model. I explain quantitatively how the large-scale temperature gradient persists in spite of thermal conduction and convection. These results are a consequence of the cosmological assembly of clusters and cannot be reproduced with non-cosmological simulations of isolated halos. I show that the variation in halo assembly histories produces a $\sim 10\%$ scatter in temperature at fixed mass. On top of this scatter, conduction decreases the temperature of the gas near the scale radius in massive clusters, which may bias hydrostatic mass estimates inferred from x-ray and SZ observations. As an example application of these model profiles, I use mixing-length theory to estimate the turbulent pressure support created by the magnetothermal instability (MTI): in agreement with my earlier MHD simulations, I find that the convection produced by the MTI can provide $\sim 5\%$ non-thermal pressure support near r_{500}. The magnitude of this turbulent pressure support is likely to be non-monotonic in halo mass, peaking in $\sim 10^{14.5}\,M_\odot$ halos. The work in Chapter 4 has previously been published in McCourt et al. (2013).

Finally, in Chapter 5, I describe my current research trying to quantify the effect of convection in the outer parts of galaxy clusters. I use the results from Chapter 4 to implement a novel "semi-"cosmological simulation method and I show my preliminary results. Chapter 5 is the only chapter in this thesis which describes ongoing, unpublished work. I expect to submit it for publication during the summer of 2014.

Chapter 2

Nonlinear Saturation of Buoyancy Instabilities

2.1 Introduction

The thermodynamics of a plasma can have dramatic and sometimes unexpected implications for its dynamical evolution. For example, thermal conduction can reduce the accretion rate in spherical accretion flows by as much as two to three orders of magnitude relative to the Bondi value (Johnson & Quataert 2007; Shcherbakov & Baganoff 2010). More relevant to this paper, Balbus (2000) and Quataert (2008) demonstrated that the convective dynamics of conducting plasmas are completely different from those of an adiabatic fluid. This paper focuses on the nonlinear evolution and saturation of this convection.

When anisotropic conduction is rapid compared to the dynamical response of a plasma, the temperature gradient, rather than the entropy gradient, determines the plasma's convective stability. The convective instability in this limit is known as the heat-flux-driven buoyancy instability (HBI) when the temperature increases with height ($\boldsymbol{g} \cdot \nabla T < 0$) or the magnetothermal instability (MTI) when the temperature decreases with height ($\boldsymbol{g} \cdot \nabla T > 0$). We summarize the linear physics of these instabilities in section 2.2. Parrish & Stone (2005, 2007) and Parrish & Quataert (2008) studied the nonlinear development of the MTI and HBI using numerical simulations. These instabilities couple the magnetic structure of the plasma to its thermal properties and potentially have important implications for galaxy clusters (Parrish et al. 2008, 2009; Bogdanović et al. 2009; Sharma et al. 2009a; Parrish et al. 2010; Bogdanović et al. 2011; Ruszkowski & Oh 2010), hot accretion flows onto compact objects (Sharma et al. 2008), and the interiors and surface layers of white dwarfs and neutron stars (Chang et al. 2010).

In this paper, we revisit the nonlinear behavior of the HBI and MTI. We focus on the physics of their saturation, but also include a lengthy discussion of possible astrophysical implications in section 2.6. Our analysis is idealized (we use a plane-parallel approximation and neglect radiative cooling), but we are able to understand the nonlinear behavior of the HBI and MTI, and therefore their astrophysical implications, more thoroughly than in

previous papers. Our results for the saturation of the HBI are similar to those of Parrish & Quataert (2008), but with an improved understanding of the saturation mechanism. Our results for the MTI, however, differ from previous results, significantly changing the predicted astrophysical implications of the MTI. We show in section 2.4.2 that this is because the development of the MTI in many previous simulations has been hindered by the finite size of the simulation domain.

The structure of this paper is as follows. We describe the linear physics of the MTI and HBI in section 2.2. We describe our computational setup in section 2.3 and the results of our numerical simulations in section 2.4. In order to better understand how the saturation of the HBI and MTI may change in a more realistic astrophysical environment, we study the interaction between these instabilities and an external source of turbulence or fluid motion (section 2.5). Finally, in section 2.6, we summarize our results and discuss the astrophysical implications of our work.

2.2 Background

2.2.1 Equations and Assumptions

We assume that the plasma is an ideal gas with an adiabatic index $\gamma = 5/3$ and model it using the magneto-hydrodynamic (MHD) equations, neglecting all dissipative processes except thermal conduction. In Cartesian coordinates, the equations for the conservation of mass, momentum and magnetic flux, and for the evolution of internal energy are

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \boldsymbol{v}) = 0, \tag{2.1a}$$

$$\begin{aligned} \frac{\partial}{\partial t}(\rho \boldsymbol{v}) + \nabla \cdot \left[\rho \boldsymbol{v} \otimes \boldsymbol{v} + \left(P + \frac{B^2}{8\pi}\right)\mathbf{I} + \frac{\boldsymbol{B} \otimes \boldsymbol{B}}{4\pi}\right] \\ = \rho(\boldsymbol{g} + \boldsymbol{f}), \end{aligned} \tag{2.1b}$$

$$\frac{\partial \boldsymbol{B}}{\partial t} = \nabla \times (\boldsymbol{v} \times \boldsymbol{B}), \tag{2.1c}$$

$$\rho T \frac{ds}{dt} = -\nabla \cdot \boldsymbol{Q}_{\mathrm{cond}}, \tag{2.1d}$$

where ρ is the mass density, $\boldsymbol{v}$ is the fluid velocity, $\otimes$ denotes a tensor product, P is the pressure, $\boldsymbol{B}$ is the magnetic field, $\mathbf{I}$ is the unit matrix, $\boldsymbol{g}$ is the gravitational field, $\boldsymbol{f}$ is an externally imposed force, T is the temperature,

$$s = \frac{1}{\gamma - 1}\frac{k_{\mathrm{B}}}{m_{\mathrm{H}}} \ln\left(\frac{P}{\rho^{\gamma}}\right) \tag{2.2}$$

is the entropy per unit mass, $d/dt = \partial/\partial t + \boldsymbol{v} \cdot \nabla$ is the Lagrangian time derivative, and $\boldsymbol{Q}_{\mathrm{cond}}$ is the conductive heat flux.

We ignore the ion contribution to the conductive heat flux, which is smaller than the electron contribution by a factor of $(m_{\rm i}/m_{\rm e})^{1/2} \approx 42$. We assume that the electrons have mean free paths much longer than their gyro-radii (as is the case in the intracluster medium (ICM) of galaxy clusters); the electrons therefore move almost entirely along magnetic field lines. Consequently, the thermal conductivity of the plasma is strongly anisotropic. The conductive heat flux in this limit becomes

$$\boldsymbol{Q}_{\rm cond} = -\kappa_{\rm e}\,\hat{\boldsymbol{b}}\,(\hat{\boldsymbol{b}}\cdot\nabla T), \tag{2.3}$$

where $\hat{\boldsymbol{b}} = \boldsymbol{B}/B$ is a unit vector in the direction of the magnetic field and $\kappa_{\rm e}$ is the thermal conductivity of free electrons (Braginskii 1965). While $\kappa_{\rm e}$ depends sensitively on temperature (Spitzer 1962), we take it to be constant in our calculations to simplify the interpretation of our results. This approximation does not affect our conclusions, because the physics of the buoyancy instabilities is independent of the conductivity in the limit that the thermal diffusion time across the spatial scales of interest is short compared to the dynamical time.

2.2.2 The Physics of Buoyancy Instabilities in Dilute Plasmas

The dissipative term in equation (2.1d) shows that fluid displacements in a plasma are not in general adiabatic. The standard analysis of buoyancy instabilities therefore does not apply, and the convective and mixing properties of a conducting plasma can be very different from those of an adiabatic fluid. In this paper, we focus on the limit in which thermal conduction is much faster than the dynamical time in the plasma. This applies on scales $\lesssim 7\,(\lambda H)^{1/2}$, where λ is the electron mean free path and H is the plasma scale height; this "rapid conduction limit" encompasses scales $\lesssim 0.3r$ in the cores of clusters, and scales $\lesssim r$ in the outer parts. This includes most scales of interest in the ICM. In this limit, the magnitude of the temperature gradient and the local orientation of the magnetic field control the convective stability of the plasma (Balbus 2000; Quataert 2008).

Although a single dispersion relation describes the linear stability of plasmas in this limit, it is easiest to understand the physics by separately considering cases where the temperature increases or decreases with height. Balbus (2000) first considered the case where the temperature decreases with height and identified the magnetothermal instability, or MTI. We show a schematic of this instability in the top row of Figure 2.1. The first panel of this figure shows a plasma in hydrostatic and thermal equilibrium, with a weak horizontal magnetic field. We apply a small, plane wave perturbation and, as the plasma evolves, the field lines follow the fluid displacements. Efficient conduction along these field lines keeps the displacements isothermal. Since the temperature falls with height, upwardly displaced fluid elements are warmer than their new surroundings; they expand and continue to rise. Similarly, downwardly displaced fluid elements sink. An order of magnitude calculation shows that displacements grow exponentially, with the growth rate (or e-folding rate) $p \sim |g\,\partial \ln T/\partial z|^{-1/2}$.

Quataert (2008) investigated the limit in which the temperature of the plasma increases with height. The instability in this case is known as the heat-flux-driven buoyancy instability,

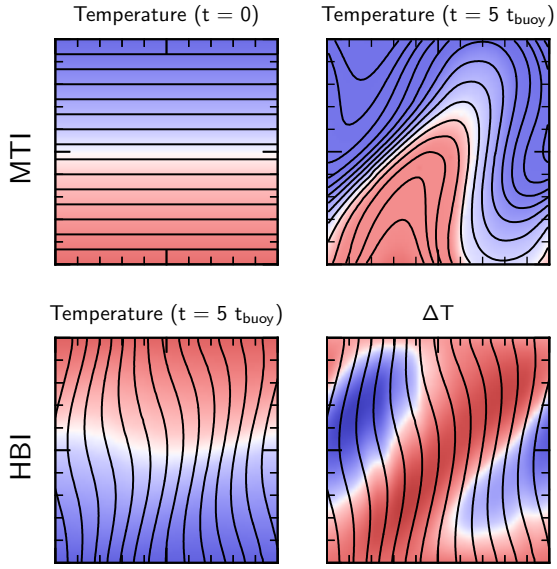

Figure 2.1: Illustration of the linear development of buoyancy instabilities in dilute plasmas from two-dimensional numerical simulations. Color shows the temperature, increasing from blue to red, while black lines trace the magnetic field. **Top Row:** Quasi-linear evolution of the MTI. *Left panel:* Initial equilibrium state, with $\hat{b}_z = 0$ and the conductive heat flux $\boldsymbol{Q} = 0$ everywhere. *Right panel:* The plasma at $t = 5\,t_{\rm buoy}$ given an initial perturbation with $\boldsymbol{k} = 2\pi/L\;\hat{\boldsymbol{x}}$. Flux freezing and rapid conduction along field lines ensure that Lagrangian fluid displacements are nearly isothermal. If the initial state has a positive temperature gradient, upwardly displaced fluid elements are warmer than their surroundings; they will expand and continue to rise. Similarly, downwardly displaced fluid elements are cooler than their surroundings and will sink. **Bottom Row:** Quasi-linear evolution of the HBI. The initial equilibrium state has $\hat{b}_z = 1$ and a net downward conductive heat flux, with $\nabla \cdot \boldsymbol{Q} = 0$ everywhere. *Left panel:* The plasma at $t = 5\,t_{\rm buoy}$ given an initial perturbation with $k_x = k_z$. The perturbation alters the geometry of the magnetic field, so that field lines are not parallel and there can be conductive heating and cooling of the plasma. *Right panel:* Temperature difference at $t = 5\,t_{\rm buoy}$ relative to the initial condition. Magnetic field lines converge in upwardly displaced fluid elements, leading to an increased temperature. These fluid elements will then expand and continue to rise. Similarly, downwardly displaced fluid elements are conductively cooled, contract and continue to sink.

or HBI. We sketch the growth of the HBI in the bottom row of Figure 2.1. Here, the initial equilibrium state has vertical magnetic field lines and a constant heat flux; the latter is required for the plasma to be in thermal equilibrium. However, perturbations to the magnetic field divert the heat flux and conductively heat or cool pockets of the plasma. If the temperature increases with height, upwardly displaced fluid elements become warmer than their surroundings and therefore experience a destabilizing buoyant response. These perturbations grow with the same growth rate $|g\,\partial \ln T/\partial z|^{-1/2}$.

Quataert (2008) performed a WKB analysis on equations (2.1a)–(2.1d) in the Boussinesq limit and obtained the following dispersion relation for plasma in a constant gravitational field $\boldsymbol{g} = -g\,\hat{\boldsymbol{z}}$, threaded by a constant magnetic field with any orientation in the $\hat{\boldsymbol{x}}-\hat{\boldsymbol{z}}$ plane:

$$\begin{aligned}(p^2+\omega_{\rm A}^2)(p+\omega_\kappa) \;+\; p\,N^2(1-\hat{k}_z^2) = \\ \omega_\kappa\left(g\frac{\partial \ln T}{\partial z}\right)\left[(2\hat{b}_z^2-1)(1-\hat{k}_z^2)-2\hat{b}_x\hat{b}_z\hat{k}_x\hat{k}_z\right].\end{aligned} \tag{2.4}$$

Here, $p=-i\omega$ is the local growth rate of the mode (the time dependence of the perturbation is e^{pt}), $\omega_{\rm A}=\boldsymbol{k}\cdot\boldsymbol{v}_{\rm A}$ is the Alfvén crossing frequency,

$$N^2 = \frac{\gamma-1}{\gamma}\,\frac{\mu\,m_{\rm H}}{k_{\rm B}}\,g\,\frac{\partial s}{\partial z} \tag{2.5}$$

describes the buoyant response of an adiabatic plasma, $\hat{\boldsymbol{k}}=\boldsymbol{k}/k$ is the direction of the wave vector, and

$$\omega_\kappa = \frac{\gamma-1}{\gamma}\,\frac{\kappa_{\rm e}\,T}{P}\,\left(\hat{\boldsymbol{b}}\cdot\boldsymbol{k}\right)^2 \tag{2.6}$$

is inversely proportional to the conduction time across the wavelength of the mode. It is convenient to identify $\omega_{\rm buoy}=|g\,\partial\ln T/\partial z|^{1/2}$ as a characteristic frequency for the buoyancy instabilities and we will also use $t_{\rm buoy}=\omega_{\rm buoy}^{-1}$ and $t_{\rm ad}=N^{-1}$ in our analysis.

In the limit that conduction is rapid compared to any dynamical response ($\omega_\kappa \gg \omega_{\rm A}$, N, $\omega_{\rm buoy}$), equation (2.4) simplifies to[1]

$$\begin{aligned}\left(p^2+\omega_{\rm A}^2\right) = {\rm sgn}(\partial T/\partial z)\,\omega_{\rm buoy}^2\times \\ \left[(2\hat{b}_z^2-1)(1-\hat{k}_z^2)-2\hat{b}_x\hat{b}_z\hat{k}_x\hat{k}_z\right].\end{aligned} \tag{2.7}$$

As we stated in section 2.2.1, the growth rate in the rapid conduction limit is independent of the thermal conductivity.

Equation (2.7) shows that magnetic tension can suppress the MTI and HBI; if $\omega_{\rm A}^2 > \omega_{\rm buoy}^2$, p^2 must be negative, and the plasma is stable to small perturbations. In this paper, we focus on the relatively weak field limit in which magnetic tension does not suppress the instabilities,

[1] In going from the cubic equation (2.4) to the quadratic (2.7), we have ignored a solution for p. This solution is exponentially damped on the conduction timescale and is not relevant to our analysis.

and we take $\omega_A \ll \omega_{buoy}$ in equation (2.7). The dominant role of the magnetic field in our analysis is to enforce anisotropic electron heat transport.

For any magnetic field direction $\hat{\boldsymbol{b}}$, the term in square brackets in equation (2.7) can be positive or negative; thus, there are always linearly unstable modes, irrespective of the thermal state of the plasma (excluding the singular case of $\partial T/\partial z = 0$).

We finally note that, although we neglected cooling in equation (2.1d), it adds a term to equation (2.7) which is insignificant when the cooling time is much longer than ω_κ^{-1} (Balbus & Reynolds 2008, 2010). This is not necessarily the case near the centers of cool-core clusters, and we plan to explore the combined effects of cooling and buoyancy instabilities in future work. For now, we assume that the buoyancy instabilities develop independently of cooling; this allows us to understand the nonlinear development and saturation of the instabilities themselves and therefore to assess their possible implications for the thermal balance of the plasma.

2.3 Numerical Method

2.3.1 Problem Setup and Integration

We consider the evolution of a volume of plasma initially in hydrostatic and thermal equilibrium, but subject to either the HBI or MTI. We seed our simulations with Gaussian-random velocity perturbations with a flat spatial power spectrum and a standard deviation of $10^{-4}\, c_s$. The small amplitude of these initial perturbations ensures that the instabilities start out in a linear phase and permits us to compare our results with the predictions of equation (2.7). Astrophysical perturbations are unlikely to be this subsonic, however, and the instabilities in our simulations take much longer to saturate then one would expect from the larger perturbations found in a more realistic scenario. We return to this point in section 2.6.

We highlight the distinctness of the HBI and MTI from adiabatic convection by choosing initial conditions with $\partial s/\partial z > 0$ whenever possible, so that the plasma would be absolutely stable if it were adiabatic. Our results do not rely critically on this choice, however, because the plasma is not adiabatic and its evolution is independent of its entropy gradient to lowest order in $\omega_{buoy}/\omega_\kappa$. Our results are much easier to interpret when ω_{buoy} doesn't vary across the simulation domain, and we prioritize this constraint on the initial condition over the sign of its entropy gradient (as we discuss in more detail below).

We solve equations (2.1a)–(2.1c), along with a conservative form of equation (2.1d), using the conservative MHD code Athena (Gardiner & Stone 2008; Stone et al. 2008) with the anisotropic conduction algorithm described in Parrish & Stone (2005) and Sharma & Hammett (2007). In particular, we use the monotonized central difference limiter on transverse heat fluxes to ensure stability. This conduction algorithm is sub-cycled with respect to the main integrator with a time step $\Delta t \propto (\Delta x)^2$; our simulations are therefore more computationally expensive than adiabatic MHD calculations. We draw most of our conclusions from

simulations performed on uniform Cartesian grids of $(64)^3$ and $(128)^3$. We also performed a large number of two-dimensional simulations on grids of $(64)^2$, $(128)^2$ and $(256)^2$. We found that the kinetic energy generated in our local HBI simulations converges by a resolution of $(64)^3$; our global HBI simulations require roughly 60 grid cells per scale height to give a converged kinetic energy. The MTI is somewhat more sensitive to resolution; we tested the convergence of these simulations by comparing the results of an identical simulation performed at $(128)^3$ and $(256)^3$.

We perform our calculations in the plane-parallel approximation with uniform gravity $\boldsymbol{g} = -g\,\hat{\boldsymbol{z}}$. We fix the temperature at the upper and lower boundaries of our computational domain, and we extrapolate the pressure into the upper and lower ghost cells to ensure that hydrostatic equilibrium holds at the boundary. For all other plasma variables, we apply reflecting boundary conditions in the direction parallel to gravity and periodic boundary conditions in the orthogonal directions. As we discuss in section 2.4, the choice of fixed temperatures at the upper and lower boundaries has an important effect on the non-linear evolution. Simulations with Neumann boundary conditions in which the temperature at the boundaries is free to adjust would give somewhat different results (see, e.g., Parrish & Stone 2005). Our choice of Dirichlet boundary conditions is largely motivated by the fact that many galaxy clusters in the local universe are observed to have non-negligible temperature gradients (Piffaretti et al. 2005).

We perform both local (with the size of the simulation domain L much smaller than the scale height H) and global ($L \gtrsim H$) simulations of the MTI and HBI. The local simulations separate the development of the instability from any large-scale response of the plasma, allowing us to study the dynamics in great detail. However, because the response of the plasma on larger scales can influence the nonlinear evolution and saturation of the instabilities, we also carry out global simulations.

2.3.2 Local Simulations

In sections 2.3–2.5, we work in units with $k_{\rm B} = \mu\, m_{\rm p} = 1$. As noted previously, we restrict our analysis to plasmas with rapid conduction; we find that setting $\kappa_{\rm e} = 10 \times \rho\, \omega_{\rm buoy}\, L^2$ puts our local simulations safely in this limit. This corresponds to a thermal diffusion time across the box of $\sim 0.1\, \omega_{\rm buoy}^{-1}$. We initialize our local HBI simulations with a linear temperature gradient:

$$T(z) = T_0\,(1 + z/H)\,, \tag{2.8a}$$

$$\rho(z) = \rho_0\,(1 + z/H)^{-3}\,, \tag{2.8b}$$

$$P(z) = \rho(z)\,T(z)\,. \tag{2.8c}$$

We choose to set $\rho_0 = 1$ and $g = 1 = 2\,T_0/H$. Unless otherwise noted, we take $H = 2$ ($T_0 = 1$) in our local simulations and we evolve a volume of plasma from $z = 0$ to $z = L = 0.1$, i.e., over $\sim 5\%$ of a scale-height.

We use the setup of Parrish & Stone (2005) for our local MTI simulations:

$$T(z) = T_0 \left(1 - z/H\right), \tag{2.9a}$$
$$\rho(z) = \rho_0 \left(1 - z/H\right)^2, \tag{2.9b}$$
$$P(z) = \rho(z)\, T(z), \tag{2.9c}$$

with $g = 1$ and $H = 3$. The MTI induces large vertical displacements in the plasma, and our reflecting boundary conditions clearly influence its evolution; we attempt to minimize the effects of the boundaries by sandwiching the unstable volume of plasma between two buoyantly neutral layers. Parrish & Stone (2005) describe this setup in more detail.[2]

Both of the above atmospheres have positive entropy gradients and therefore would be stable in the absence of anisotropic conduction. The results we describe in this paper are entirely due to non-adiabatic processes.

We show in section 2.4.1 that local and global HBI simulations give very similar results; we therefore use the simpler, local simulations for most of our analysis of the HBI. By contrast, the large-scale vertical motions induced by the MTI make its evolution inherently global; as we describe in section 2.4.2, local simulations do not give a converged result independent of L/H. We therefore study the saturated state of the MTI using global simulations, which require an alternate set of initial conditions.

2.3.3 Global Simulations

The physical properties of the plasma in a global simulation can vary by an order of magnitude or more across the simulation domain. This complicates our study of the HBI and MTI because the instabilities can be in different stages of evolution at different spatial locations. Similarly, the kinetic and magnetic energies generated by the instabilities in a global simulation can be functions of height, obscuring their dependence on other parameters. While one must confront these problems when studying buoyancy instabilities in an astrophysical context, we avoid them by choosing initial conditions with a buoyancy time that is constant with height. We specify

$$g(z) = g_0\, e^{-z/S}, \tag{2.10a}$$
$$T(z) = \exp\left[\pm \frac{S\omega_{\rm buoy}^2}{g_0}\left(e^{z/S} - 1\right)\right]. \tag{2.10b}$$

The positive and negative signs in equation (2.10b) produce atmospheres unstable to the HBI and MTI, respectively. Note that $T(z = 0) = 1$ as in our local simulations. We take $g_0 = 1$, $S = 3$ and $\omega_{\rm buoy}^2 = 1/2$ and numerically solve for $\rho(z)$ so that the initial atmosphere is in hydrostatic equilibrium. The atmosphere defined by equation (2.10b) does not have a single, well-defined scale height. We therefore define H so that $L/H = \ln\left(T_{\rm max}/T_{\rm min}\right)$; this

[2]Note that, although Parrish & Stone (2005) describe these extra layers as buoyantly stable, the effect of the isotropic conductivity is to make them buoyantly neutral.

definition for H is analogous to that used in our local simulations and L/H reflects the total free energy available to the instabilities. Because H is a defined, rather than fundamental, property of the atmosphere, it is not independent of the size of the simulation domain. We carry out simulations with domain sizes $L = 0.5$ and 2.0, which have $L/H = 0.27$ and 1.4, respectively. The temperature profile defined by equation 2.10b becomes unrealistically steep when $z \gtrsim S$; thus, we only consider models with $L \lesssim S$. We do include simulations with $L/H < 1$ and $L/H > 1$.

Simulations with a larger simulation domain size inherently permit larger vertical displacements, and the boundary conditions do not influence the evolution of the plasma as strongly as they do in a local simulation. We find that the neutrally stable layers described in section 2.3.2 don't alter the results of our global simulations, and thus we do not include them in our setup.

We chose the conductivity in our local simulations so that $\kappa_{\rm e} = 10\,\rho\,\omega_{\rm buoy}L^2$, but the corresponding constraint on the time step becomes impractical for our global simulations. Instead, we adjust $\kappa_{\rm e}$ so that the ratio $\kappa_{\rm e}/L$ is the same as it is in the local simulations. This keeps the ratio of the conduction time to the sound crossing time constant, and is appropriate if the turbulence driven by the MTI or HBI reaches a terminal speed less than or of order the sound speed. We carried out simulations with different values of the conductivity and verified that our results are insensitive to factor of few changes in $\kappa_{\rm e}$.

The atmosphere we use for our global MTI simulations has a negative entropy gradient. One might worry that the results of these simulations—which aim to focus on the MTI—would be biased by the presence of adiabatic convection. This is, however, only a pedagogical inconvenience; $\omega_\kappa > N$ on all relevant scales in the simulation, so the adiabatic limit to equation (2.4) is not important, and the evolution of the plasma is nearly independent of its entropy gradient. To confirm this we carried out $L = 0.27\,H$ simulations with both sets of initial conditions; they give very similar results. Additionally, we performed one simulation of adiabatic convection using this setup, and the behavior of this simulation is entirely different from the MTI, especially at late times.

2.3.4 Turbulence

The pure HBI/MTI simulations described in the previous sections are physically very instructive but astrophysically somewhat idealized. In order to better understand the astrophysical role of the HBI and MTI, we also study their interaction with other sources of turbulence and fluid motion. We assume that these take the form of isotropic turbulence; we include such turbulence via the externally imposed force field $\boldsymbol{f}$ in equation (2.1b). Following Lemaster & Stone (2008), we compute $\boldsymbol{f}$ in momentum space from scales $k_0 = \{4, 6, 8\} \times 2\pi/L$, down to $k_{\rm max} = 2\,k_0$, with an injected energy spectrum $\dot{E}_k^{\rm (inj)} \propto k^{-3}$. We then randomize the phases, perform a Helmholtz decomposition of the field, discard the compressive component, transform the field into configuration space, and normalize it to a specified energy injection rate. We find that the turbulence sets up a nonlinear cascade that is not very sensitive to either the driving scale or the injected spectrum of the turbulence.

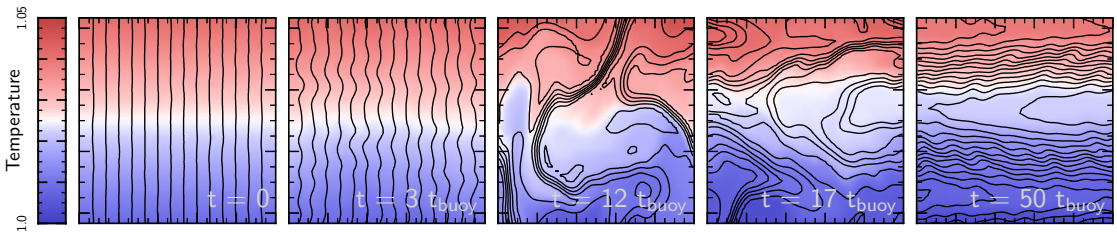

Figure 2.2: Evolution of the HBI with an initially vertical magnetic field in a local, 2D simulation (simulation h1 in Table 2.1). Color shows temperature and black lines show magnetic field lines. A small velocity perturbation to the initial state seeds exponentially growing modes which dramatically reorient the magnetic field to be predominantly horizontal. The induced velocities are always highly subsonic and, after $t \sim 20\,t_{\rm buoy}$, are also almost entirely horizontal. Once the plasma reaches its saturated state, it is buoyantly stable to vertical displacements. The plasma does not resist horizontal displacements, but the saturated state is nearly symmetric to these displacements and they do not change its character.

This prescription for turbulence is statistically uniform in both space and time and therefore provides a controlled environment in which to study the interaction between turbulence and buoyancy instabilities. This may not, however, be a good approximation to real astrophysical turbulence; we discuss the implications of our choice in sections 2.5 and 2.6.

2.4 Nonlinear Saturation

2.4.1 Saturation of the HBI

Parrish & Quataert (2008) described the nonlinear saturation of the HBI, but they did not explicitly test the dependence of their results on the size of the computational domain. This dependence turns out to be crucial for the MTI (see § 2.4.2), but we show here that the size of the domain has little effect on the saturation of the HBI. Nonetheless, we describe the nonlinear behavior of the HBI in reasonable detail, expanding on the physical interpretation given in previous papers (although our HBI results do not differ significantly from those of previous authors). The saturation of the HBI is the simplest process we consider in this paper and serves as a useful comparison for our new results.

We study the saturation of the HBI using 2D and 3D simulations spanning a range of domain sizes L/H. Table 2.1 lists all of the simulations presented in this section.

Figure 2.2 shows snapshots of the evolution of temperature and magnetic field lines in a local, 2D HBI simulation. We chose this simulation to simplify the field-line visualization, but the results in Figure 2.2 apply equally to our local and global 3D simulations. We initialized this simulation in an unstable equilibrium state with vertical magnetic field lines ($\hat{b}_z = 1$). As described in section 2.3.1, we seed this initial condition with small velocity perturbations; the HBI causes these perturbations to grow in the first three panels of Figure 2.2. The evolution

Table 2.1: Parameters for the HBI simulations (§ 2.4.1).

Name	D	res	L/H	κ
h1	2	64	0.05	7.07
h2	3	64	0.05	7.07
h3	3	128	0.75	0.47
h4	3	128	1.40	0.35

All simulations are performed on square Cartesian grids of size L. D is the dimensionality of the simulation, res is the number of grid cells along a side, and κ is the conductivity (in units of $k_{\rm B}/\mu m \times \rho\,\omega_{\rm buoy} L^2$). All of these simulations are local (eq. 2.8), except for the one with $L/H = 1.4$, which uses the global setup (eq. 2.10). We initialized all of these simulations with weak horizontal magnetic fields ($B/\sqrt{4\pi} = 10^{-6}$).

becomes nonlinear in the third panel, when the velocity perturbations reach $\sim 4\%$ of the sound speed. Afterwards, the instability begins to saturate and the plasma slowly settles into a new equilibrium state. The last panel in Figure 2.2 shows that this saturated state is highly anisotropic: the magnetic field lines are almost entirely orthogonal to gravity. Flux conservation implies that the fluid motions must also be anisotropic, with most of the kinetic energy in horizontal motions at late times (see Fig. 2.3, discussed below). These horizontal motions are very subsonic: in all of our simulations, the velocities generated by the HBI are significantly less than 1% of the sound speed in the saturated state.

Because the fluid velocities remain small, the linear dispersion relation (eq. 2.7) captures much of the evolution of the HBI, even at late times. For any magnetic field orientation, the fastest growing modes are the ones with $\boldsymbol{k}$ along the axis $\hat{\boldsymbol{b}} \times (\hat{\boldsymbol{b}} \times \boldsymbol{g})$; these modes have the growth rate

$$p_{\rm max} = |\omega_{\rm buoy}\, \hat{b}_z|, \tag{2.11}$$

which decreases as the field lines become horizontal. Additionally, when $\hat{b}_z^2 < 1/2$, only modes with $\hat{k}_z^2 > 1 - 4(\hat{b}_z^2 - \hat{b}_z^4)$ are unstable. Since the HBI saturates by making the field lines horizontal ($\hat{b}_z \to 0$), both the maximum growth rate of the instability and the volume of phase space for unstable modes decrease as the HBI develops. This strongly limits the growth of the perturbations, and helps explain why the instability saturates relatively quiescently.

As argued by Parrish & Quataert (2008), the HBI saturates when its maximum growth rate $p_{\rm max}$ vanishes, so that there are no longer unstable modes. While this is clearly a sufficient condition for the plasma to reach a new stable equilibrium, it is by no means necessary: the instability could, e.g., saturate via nonlinear effects, but in practice this is not the case (at least for simulations without an additional source of turbulence). Equation (2.11) for $p_{\rm max}$ shows that the HBI could saturate by making either $\partial T/\partial z$ or $\hat{b}_z$ vanish; intuitively, the HBI is powered by a conductive heat flux, which it must extinguish in order to stop growing. Erasing the temperature gradient might seem like the more natural saturation channel, since the conduction time across the domain is much shorter than the time it

takes the HBI to develop and saturate. In an astrophysical setting, however, the large-scale temperature field is often controlled by cooling, accretion or other processes apart from the HBI. We therefore impose the overall temperature gradient in our simulations by fixing the temperature at the top and bottom of the domain, so that $\omega_{\rm buoy}$ is roughly independent of time and saturation requires $\hat{b}_z = 0$.

Since the HBI saturates by making the magnetic field lines horizontal, we take the $\hat{b}_z \to 0$ limit in equation (2.7) to understand the late-time behavior of the plasma:

$$\omega = \pm\, \omega_{\rm buoy} \left(1 - \hat{k}_z^2\right)^{1/2}. \tag{2.12}$$

We have assumed here that the magnetic field is weak enough for magnetic tension to be negligible on the scales of interest.[3] Equation (2.12) shows that the saturated state of the HBI is buoyantly stable, but that there is a family modes with $\hat{k}_z = 1$ which feel no restoring force. This simply reflects the fact that the plasma is stably stratified and resists vertical displacements. Displacements orthogonal to gravity are unaffected by buoyancy, however, and appear as zero-frequency modes in the dispersion relation.

Figure 2.3 demonstrates this asymmetry between vertical and horizontal displacements in the late time evolution of the plasma. This figure shows the kinetic energy in horizontal and vertical motions as a function of time. During the initial, linear growth of the instability ($t \lesssim 10\, t_{\rm buoy}$), buoyantly unstable fluid elements accelerate toward the stable equilibrium, and the kinetic energy is approximately evenly split among vertical and horizontal motions. As the instability saturates, however, the plasma becomes buoyantly stable and traps the vertical motions in decaying oscillations (internal gravity waves). The horizontal motions keep going, however, and retain their kinetic energy for the duration of the simulation.[4] This difference in the response of the plasma to vertical and horizontal motions accounts for the anisotropy of the velocity field in the saturated state of the HBI.

Figure 2.4 shows the evolution of the rms magnetic field angle (left panel) and magnetic energy (right panel) in 3D HBI simulations for three different values of the size of the computational domain L relative to the scale height H. The zero frequency modes discussed above also dominate the evolution of the magnetic field at late times, after the motions become nonlinear ($t \gtrsim 10\, t_{\rm buoy}$). The horizontal displacements stretch out the field lines, amplifying and reorienting them. Quantitatively, we expect that $\hat{b}_z \sim \lambda/\xi \propto t^{-1}$, where λ is a characteristic scale for the modes in the saturated state, ξ is the magnitude of the horizontal displacements, and we have assumed that the velocity is constant with time. The left panel of Figure 2.4 shows that the dependence in the simulations is quite close to this, with $\hat{b}_z \propto t^{-0.85}$.[5] Stretching the field lines in this manner amplifies the field strength by an

[3] Note that equation 2.12 only strictly applies when the magnetic field is *exactly* horizontal. More generally, there will still be unstable modes with growth rate given by equation 2.11; this growth rate is very slow in the saturated state of the HBI, however, and these modes don't change the dynamics of the plasma.

[4] These horizontal motions are not susceptible to the Kelvin-Helmoltz instability because vertical mixing is stabilized by buoyancy when the Richardson number $\sim \omega_{\rm buoy}^2/(dv/dz)^2 > 0.25$ (Turner 1979).

[5] The slight difference relative to the simple predictions of flux freezing given the velocity field in Figure 2.3

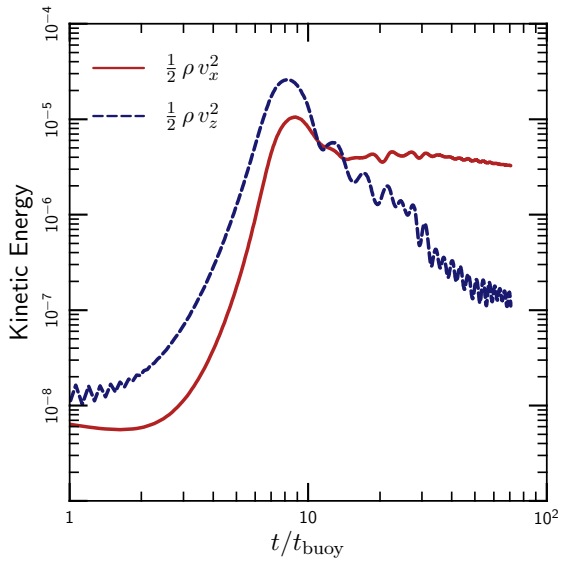

Figure 2.3: Evolution of the vertical and horizontal kinetic energy in a local, 2D HBI simulation (simulation h1 in Table 2.1). The units are such that the thermal pressure $P \approx 1$ and the initial magnetic energy is $B^2/8\pi = 10^{-12}$. After a period of exponential growth in which the x and z motions are in approximate equipartition, the HBI saturates and the kinetic energy ceases to grow. At this point, the energy in the vertical motion is in the form of stable oscillations, which decay non-linearly. The horizontal motions are unhindered, however, and persist for the entire duration of the simulation. These horizontal motions are responsible for the asymmetry of the magnetic field shown in Figure 2.4.

amount $\delta B \propto \xi$; if the velocity is constant with time, we expect $B^2 \propto t^2$. The right panel of Figure 2.4 shows that this time dependence is approximately true for our two larger HBI simulations; the amplification is slightly slower in the very local calculation with $L/H = 0.05$. There is no indication that the magnetic field amplification has saturated at late times in the HBI simulations. We suspect that the amplification would continue until the magnetic and kinetic energy densities reach approximate equipartition, but we would have to run the simulation for a very long time to verify this.

One of the important results in Figure 2.4 is that the saturated state of the HBI is nearly independent of the size of the computational domain L/H. Since the late time evolution of the plasma is driven only by horizontal displacements, the key dynamics all occur at approximately the same height in the atmosphere. The saturation of the HBI is thus essentially local in nature and should not be sensitive to the global thermal state of the plasma or the details of the computational setup.

The dramatic reorienting of the magnetic field caused by the HBI severely suppresses the conductive heat flux through the plasma. The conductive flux is proportional to $\langle \hat{b}_z^2 \rangle$, which

may be due to the finite resolution of our simulations, which prevents us from resolving the field line direction when $\hat{b}_z \lesssim 10 \times dz/L$.

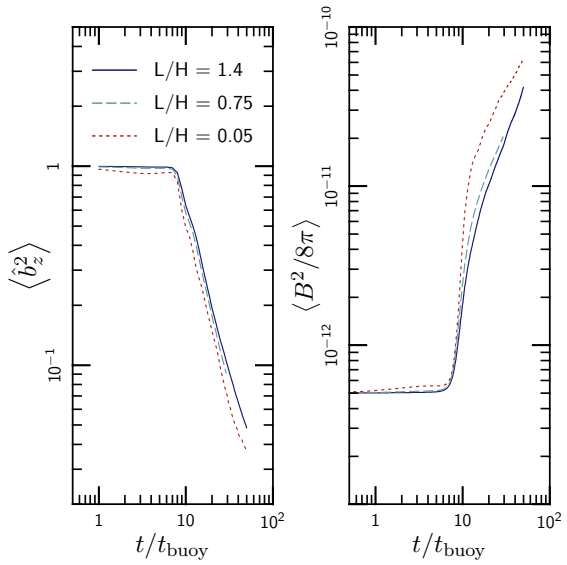

Figure 2.4: Evolution of the orientation (left) and energy (right) of the magnetic field in 3D HBI simulations for three different values of the size of simulation domain relative to the temperature scale-height (L/H) (simulations h2–h4 in Table 2.1). Units are such that the thermal pressure P is $\simeq 1$. The results are nearly independent of size of the simulation domain. As described in § 2.4.1, the HBI saturates by shutting itself off; the linear exponential growth ends at $t \sim 10\, t_{\rm buoy}$, and most of the evolution of the magnetic field happens afterward. This evolution is driven by the horizontal motions shown in Figure 2.3, which both amplify and reorient the magnetic field. After a brief period of exponential growth, the field amplification is roughly linear in time. By contrast, the field amplification by the MTI is exponential in time (see Fig. 2.8).

decreases in time $\propto (t/t_{\rm buoy})^{-1.7}$. The saturated state of the HBI is also buoyantly stable and resists any vertical mixing of the plasma. As a result, the convective energy fluxes in our HBI simulations are very small, $\sim 10^{-6}\,\rho c_s^3$. The effect of the HBI therefore is to strongly insulate the plasma against both conductive and convective energy transport. This can dramatically affect the thermal evolution of the plasma (Parrish et al. 2009; Bogdanović et al. 2009).

The fact that the growth rate of the HBI depends on the local orientation of the magnetic field, as well as the thermal structure of the plasma, makes it very different from adiabatic convection. This dependence on the magnetic field structure provides a saturation channel in which the kinetic energies are very small compared to the thermal energy (e.g., in Fig. 2.3, $\rho v^2/nkT \sim 10^{-5}$). Critically, these highly subsonic motions occur in simulations in which the boundary conditions allow for the presence of a sustained, order unity, temperature gradient ($d\ln T/d\ln z \sim 1$). In an adiabatic simulation, the analogous sustained entropy gradient would generate convective motions with $\rho v^2 \sim nkT$. This does not occur in an HBI-unstable plasma. Thus, although a plasma with a positive temperature gradient is in general buoyantly unstable, the effect of the HBI is to peacefully stabilize the plasma within a few buoyancy times by suppressing the conductive heat flux through the plasma. The resulting, stably-stratified plasma then resists vertical mixing and, in the absence of strong external forcing, we expect the fluid velocities and magnetic field lines to be primarily horizontal. In section 2.5, we perturb this state with externally driven, isotropic turbulence and test the strength of the stabilizing force.

2.4.2 Saturation of the MTI

Figure 2.5 shows the evolution of one of our local, 2D MTI simulations. As in the HBI simulation shown in Figure 2.2, we initialized this simulation in an unstable equilibrium state (a weak horizontal magnetic field) and seeded it with the small velocity perturbations described in section 2.3.1. The MTI and HBI stem from very similar physics, and as a result have very similar linear dynamics. The nonlinear behavior of the two instabilities is entirely different, however. While the HBI saturates relatively quiescently by driving the plasma to a buoyantly stable and highly anisotropic state, the MTI generates vigorous, sustained convection that tends to isotropize both the magnetic and velocity fields.

As we did for the HBI, we study the saturation of the MTI using 2D and 3D simulations spanning a range of domain sizes L/H. Table 2.2 summarizes the simulations presented in this section.

Since the linear dispersion relation successfully describes the nonlinear evolution and saturation of the HBI, it is a good place to begin our discussion of the MTI. The MTI is described by equation (2.7) when $\partial T/\partial z < 0$. The linear evolution of the MTI is the opposite of that of the HBI: the MTI operates when the temperature decreases with height, its fastest growing modes are the ones with wave vectors $\boldsymbol{k}$ parallel to $\hat{\boldsymbol{b}}$, and the force that destabilizes the MTI is exactly that which stabilizes the HBI in its saturated state. Equation (2.7) shows that the maximum growth rate of the MTI goes to zero when $\hat{b}_z = 1$. By analogy with the HBI, it thus seems reasonable to expect that the MTI also saturates quiescently, by making

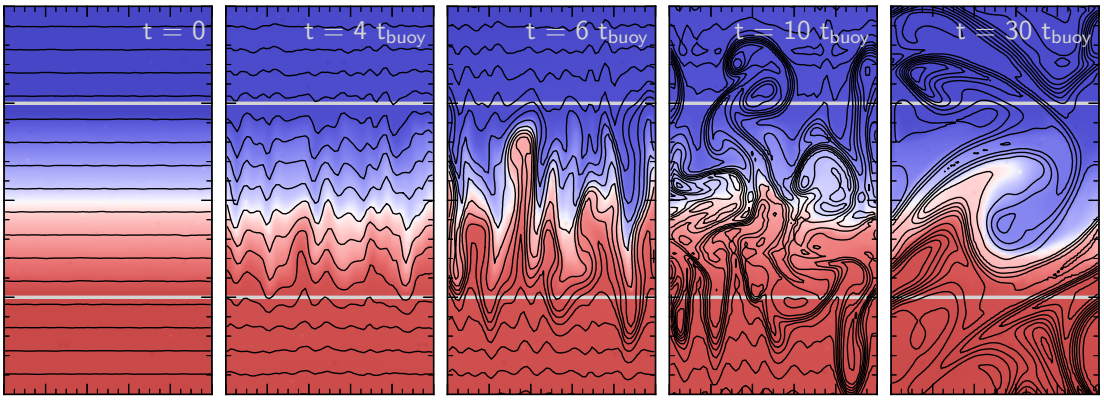

Figure 2.5: Evolution of the MTI with an initially horizontal magnetic field in a local, 2D simulation (simulation m1 in Table 2.2). Gray horizontal lines show the transition to the buoyantly neutral layers described in § 2.3.2; the color scale is identical to that in Figure 2.2. Initial perturbations grow by the mechanism described in § 2.2.2 (Fig. 2.1); rising and sinking plumes rake out the field lines until, by $t = 6\,t_{\rm buoy}$, they are mostly vertical. This configuration is, however, nonlinearly unstable to horizontal displacements, which generate a horizontal magnetic field and thus continually seed the MTI (see Fig. 2.6). The result is vigorous, sustained convection in marked contrast to the saturation of the HBI in Fig. 2.2. In this local simulation, buoyant plumes accelerate until they reach the neutrally stable layers. The boundaries prematurely stop the growth of the MTI, and the local simulation under predicts the kinetic energy generated by the MTI (see Fig. 2.7).

Table 2.2: Parameters for the MTI simulations (§ 2.4.2).

Name	D	res	L/H	κ	Field Configuration
m1	2	64	0.033	7.07	horizontal
m2	2	64	0.033	7.07	vertical
m3	3	64	0.033	7.07	horizontal
m4	3	64	0.033	7.07	vertical
m5*	3	128	0.500	0.31	horizontal
m6*	3	128	1.400	0.35	horizontal

The definitions of L, D, and κ are the same as in Table 2.1. All simulations use the local setup (eq. 2.9), except for the one with L/H = 1.4, which is global (eq. 2.10). Each of these simulations was initialized with a weak magnetic field $B/\sqrt{4\pi} = 10^{-4}$ with the orientation indicated in the table. *We also repeated simulations m5 and m6 with initial field strengths $B/\sqrt{4\pi} = 10^{-4}, 0.0014, 0.0245$

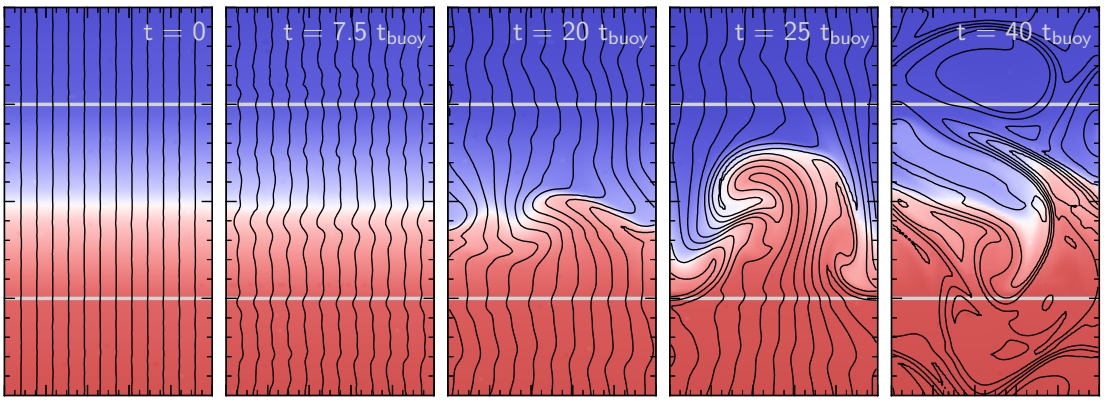

Figure 2.6: Evolution of the MTI in a plasma with an initially vertical magnetic field in a local, 2D simulation (simulation m2 in Table 2.2). This configuration is linearly stable according to equation (2.7), but there are zero-frequency, $k_x = 0$, modes which do not have a restoring force. Physically, these correspond to horizontal motions which do not feel gravity/buoyancy. Because of this zero frequency mode, small random initial perturbations add a horizontal component to the magnetic field, eventually rendering the plasma unstable to the MTI. This creates a feedback loop, allowing the MTI to generate vigorous, sustained convection; the HBI does not have this same feedback loop and so does not generate sustained turbulence (Fig. 2.2). At late times, the results of this simulation with an initially vertical magnetic field are very similar to Figure 2.5 which starts with a horizontal magnetic field.

the field lines vertical.

The first three panels of Figure 2.5 show that this is nearly what happens. As the perturbations grow exponentially, the buoyantly rising and sinking blobs rake out the field lines, making them largely vertical. The growth rate of the MTI goes to zero when the field lines become vertical; since the velocities are still small at this point in the evolution ($\sim 10^{-2} c_s$), one might expect the MTI to operate like the HBI and quiescently settle into this stable equilibrium state. Instead, however, the MTI drives sustained turbulence for as long as the temperature gradient persists. The plasma never becomes buoyantly stable, and the magnetic field and fluid velocities are nearly isotropic at late times.

We can understand this evolution using the same approach we employed for the HBI. Although the plasma in our MTI simulations never reaches a state in which the MTI growth rate is zero, examining the properties of this state is very instructive. The equilibrium state of the MTI with $\hat{b}_z = 1$ (i.e., a vertical field) has precisely the same dispersion relation as the saturated state of the HBI, given by equation (2.12). There are again zero frequency (neutrally stable) modes of the dispersion relation which correspond to horizontal perturbations to the equilibrium state of the MTI; these experience no restoring force, because the restoring force is buoyant in nature and unaffected by horizontal displacements.[6] Critically, however, these zero frequency perturbations now add a horizontal component to the magnetic field, pulling the plasma out of the equilibrium state and rendering it unstable to the MTI.

Figure 2.6 vividly illustrates this process. This figure shows a simulation that starts with the linearly *stable* equilibrium state of the MTI (a vertical magnetic field), seeded with the same highly subsonic velocity perturbations as before. The compressive component of the perturbation rapidly damps because the Mach number is small, and buoyancy traps the vertical components of the perturbation in small-amplitude oscillations, as predicted by equation (2.12). The incompressive, horizontal displacements propagate freely, however, and by $t = 7.5\, t_{\rm buoy}$, they have noticeably changed the local orientation of the magnetic field. The plasma is no longer in its stable equilibrium state, and by $t = 20\, t_{\rm buoy}$, it is clear that this process has excited the MTI. At late times, the simulations initialized with horizontal (linearly unstable; Fig. 2.5) and vertical (linearly stable; Fig. 2.6) magnetic fields are qualitatively indistinguishable. Thus, although equation (2.7) shows that a plasma with $\partial T/\partial z < 0$ is linearly stable if the magnetic field is vertical, that configuration is nonlinearly unstable.

The nonlinear instability of the $\hat{b}_z = 1$ state of the MTI precludes the magnetic saturation channel. The zero frequency horizontal motions generate a horizontal magnetic field component from the vertical magnetic field, seeding the instability and closing the dynamo loop. This continuously drives the MTI and generates sustained turbulence. Without a

[6] This conclusion is more subtle than the analogous argument for the HBI, because the equilibrium state has a nonzero heat flux. If the horizontal displacements aren't incompressive, the field lines could pinch together, heating and destabilizing parts of the plasma. These perturbations do not appear in eq (2.12) because we have taken the Boussinesq limit. It is in principle possible that such compressive perturbations contribute to destabilizing the $\hat{b}_z = 1$ MTI state in our simulations, but all of our analysis is consistent with the neutrally stable zero frequency perturbations being the critical ingredient.

linear means to saturate, the MTI grows until nonlinear effects can compete with the linear instability, which requires $v \sim c_s$.

Figure 2.7 shows the Mach number in 3D MTI simulations as a function of time, for different sizes of the computational domain L/H. The MTI buoyantly accelerates rising or sinking fluid elements. For simulations with $L \ll H$, the velocities generated by the MTI are artificially suppressed because the small size of the computational domain prematurely stops the buoyant acceleration (Fig. 2.5); the results in this Figure 2.7 are reasonably consistent with mixing length estimate of $v \sim \sqrt{L}$. By contrast, for simulations with $L \gtrsim H$, the Mach numbers approach ~ 1 so that the MTI taps into the full buoyant force associated with the unstable temperature gradient. This is only true, of course, because our boundary conditions fix the temperature at the top and bottom of the computational domain. If the temperatures were free to vary, the MTI could saturate by making the plasma isothermal. Which of these saturation mechanisms is realized in a given astrophysical system will depend on the heating and cooling mechanisms that regulate the temperature profile of the plasma.

Figure 2.8 shows the evolution of the magnetic field orientation (left panel) and energy (right panel) as a function of time in our 3D MTI simulations (for two different values of L/H). The sustained, vigorous turbulence generated by the MTI in the nonlinear regime ($t \gtrsim 10\, t_{\rm buoy}$) rapidly amplifies the magnetic field. The magnetic and kinetic energies reach approximate equipartition in our simulations, with $B^2/8\pi \sim 0.1\, \rho v^2$. The kinetic and magnetic energies generated by the MTI together contribute $\sim$ 5–10% of the pressure support in its saturated state.

The evolution of the magnetic field geometry, shown in the left panel of Figure 2.8, nicely illustrates the transition of the MTI from the linear to nonlinear regime. During the linear phase of the instability ($t \lesssim 10\, t_{\rm buoy}$), the plasma accelerates toward the equilibrium state with $\hat{b}_z = 1$. After the evolution becomes nonlinear ($t \gtrsim 10\, t_{\rm buoy}$), however, the MTI drives sustained turbulence, which nearly isotropizes the magnetic field.

While the HBI works to insulate the plasma against vertical energy transport, the MTI enhances it. Figure 2.8 shows that the conductive flux through the plasma is slightly greater than $\sim 1/3$ of the field free value $\kappa_{\rm e} \nabla T$ (because $\langle \hat{b}_z^2 \rangle \sim 0.4$). Moreover, the MTI leads to large fluid velocities and correlated temperature and velocity perturbations—hot pockets of plasma rise, while cool pockets sink. These imply that the MTI drives an efficient outwards convective heat flux. Figure 2.9 shows that this flux can be $\sim 1.5\%$ of ρc_s^3, consistent with the Mach numbers of ~ 0.2 in Figure 2.7. This convective flux is probably not large enough to influence the thermodynamics of the ICM, but it could be important in other environments where the MTI can operate, such as the interiors of white dwarfs and neutron stars (Chang et al. 2010).

2.5 Interaction with Other Sources of Turbulence

Thus far, we have described the development and saturation of buoyancy instabilities only in the idealized case of an otherwise quiescent plasma. In a more astrophysically realistic

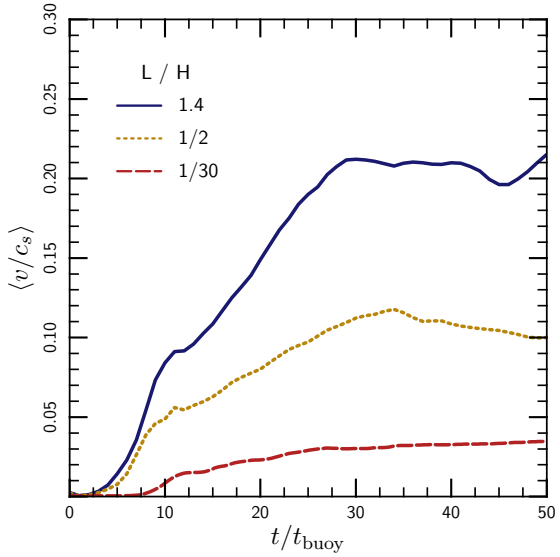

Figure 2.7: Volume-averaged Mach numbers of the turbulence generated by the MTI in 3D simulations, for three different values of the size of simulation domain relative to the temperature scale-height L/H. The MTI buoyantly accelerates the unstable fluid elements; if the size of the simulation domain is smaller than a scale-height, the boundaries suppress the growth of the instability. Local simulations with $L/H = 1/30$ therefore strongly under predict the strength of the turbulence generated by the MTI. In global simulations with $L \gtrsim H$, the MTI leads to turbulence with average Mach numbers of ~ 0.2; the velocity distribution extends up to ~ 5 times the mean.

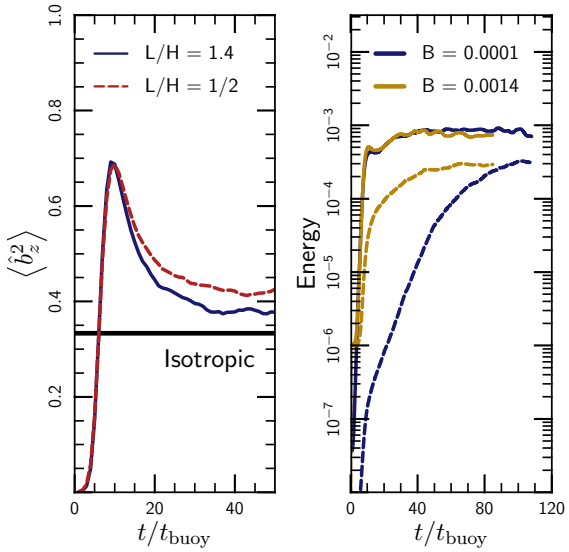

Figure 2.8: **Left panel**: Evolution of the magnetic field orientation in 3D MTI simulations, for two different values of the size of simulation domain relative to the temperature scale-height L/H. During the linear phase of evolution ($t \lesssim 10\,t_{\rm buoy}$), the MTI drives the plasma toward a nearly vertical magnetic field, i.e., $\hat{b}_z \sim 1$. When the instability becomes nonlinear, however, (near $t \sim 10\,t_{\rm buoy}$) the evolution changes. Unlike the HBI, the plasma never settles into an equilibrium state; instead the MTI drives vigorous turbulence. This turbulence amplifies and nearly isotropizes the magnetic field. **Right panel**: Evolution of the magnetic (dashed line) and kinetic (solid line) energy in local ($L/H = 1/2$) 3D MTI simulations, with different initial field strengths. These local simulations have a positive entropy gradient, but under-predict the magnetic and kinetic energy produced by the MTI. The magnetic energy in the saturated state approaches $\sim$ 10% $\times\, \rho v^2$.

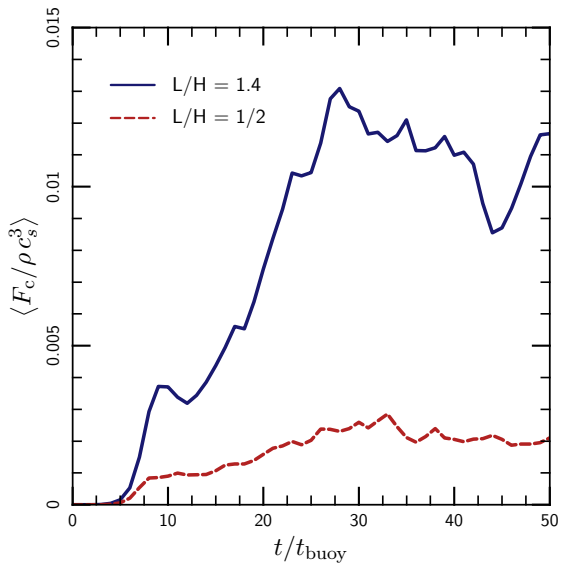

Figure 2.9: Convective energy fluxes generated by the MTI. The large turbulent velocities associated with the MTI (Fig. 2.7) lead to efficient convective transport of energy, at a reasonable fraction of the maximal value, ρc_s^3. In these MTI simulations, the conductive energy flux is ~ 0.4 times the field-free value and always dominates the convective energy flux for galaxy cluster conditions.

scenario, however, other processes and sources of turbulence may also act on the plasma and the resulting dynamics can be more complicated. For example, the evolution of the HBI won't simply proceed until the growth rate is everywhere zero. Instead, we expect the saturated state to involve a statistical balance among the various forces; this balance depends on the buoyant properties of the plasma and provides a test of our understanding of the nonlinear behavior of the HBI and MTI. Furthermore, any change in the saturated state of the plasma due to the interaction between the HBI/MTI and other sources of turbulence could change the astrophysical implications of these instabilities.

We choose to explore the interaction between buoyancy instabilities and other sources of turbulence using the idealized, isotropic turbulence model described in section 2.3.4. While this model glosses over the details of what generates the turbulence, we hope that it captures the essential physics of the problem, allowing us to study the effect of turbulence without unnecessarily restricting our analysis to specific applications. We intend to specialize to specific sources of turbulence in future work, but our present analysis should apply in the ICM, accretion disks, and anywhere else the assumptions summarized in section 2.2.2 apply.

In order to characterize the turbulence, we define a timescale for it to influence the plasma. We define this "distortion time" in terms of the spatial velocity spectrum: $t_{\rm dist}(\ell) = \ell/\delta v(\ell)$, where

$$\delta v(\ell) \equiv \left[\int \left(\delta \boldsymbol{v}(\boldsymbol{k})\right)^2 \, \delta\left(|\boldsymbol{k}| - 2\pi/\ell\right) \, \frac{d^3k}{(2\pi)^3} \right]^{1/2}, \tag{2.13}$$

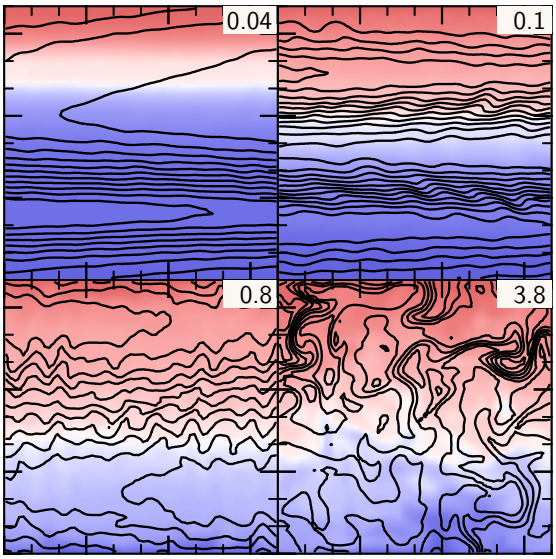

Figure 2.10: Snapshots of the saturated states of our 2D HBI simulations with externally driven turbulence. Colors show the temperature (increasing from blue to red), and black lines show magnetic field lines. Each panel is labeled with the dimensionless strength of the turbulence in the simulation, $t_{\rm buoy}/t_{\rm dist}$, defined in § 2.5. When $t_{\rm buoy} \lesssim t_{\rm dist}$, as in the top two panels, the HBI dominates the evolution of the plasma. When $t_{\rm buoy} \gtrsim t_{\rm dist}$, the turbulence can isotropize the magnetic field, but it does so in a scale-dependent way with the large scales retaining memory of the horizontal field imposed by the HBI.

and $\delta v(\boldsymbol{k})$ is the Fourier transform of the velocity field. We expect the relevant parameter describing the importance of the turbulence to be the ratio of the timescales $t_{\rm buoy}/t_{\rm dist}$. This represents a dimensionless strength of the turbulence; if $t_{\rm buoy}/t_{\rm dist} \gtrsim 1$, the turbulence displaces fluid elements faster than buoyancy can restore them, and we expect the velocities and magnetic field lines to become isotropic. In the opposite limit, buoyancy still plays an important role in the evolution of the plasma. The scale dependence of the distortion time makes the ratio $t_{\rm buoy}/t_{\rm dist}$ a function of scale. We define this ratio at the scale where the velocity spectrum of the injected turbulence peaks. This is roughly consistent with the driving scale of the turbulence and typically represents the scale with the most energy.

In the following sections, we study the transition from a state dominated by buoyancy to one dominated by isotropic turbulence using a number of simulations of the HBI and MTI, with turbulence in the range $0.1 \lesssim t_{\rm buoy}/t_{\rm dist} \lesssim 10$.

2.5.1 Effect of Turbulence on the HBI

Figure 2.10 shows representative snapshots of the temperature and magnetic field lines in saturated states of our HBI simulations with turbulence; the strength of the injected turbulence increases from the top left panel through the bottom right (the labels correspond

Table 2.3: Parameter study for the HBI simulations with turbulence (§ 2.5.1).

D (2)	res (64)	L (0.1)	H (2.0)	B_0 (10^{-6})	$\tilde{k}_0$
—	—	—	—	—	2
—	—	—	1.0	—	2
—	—	—	3.0	—	2
—	—	—	—	—	4
—	—	1.0	—	—	4
—	—	0.3	—	—	4
3	—	—	—	—	4
3	—	—	—	10^{-3}	4
3	—	1.0	—	3×10^{-4}	4
3	128	1.0	—	3×10^{-4}	4
—	—	—	—	—	6
—	—	—	—	—	8
—	256	—	—	—	8
3	—	—	—	—	8

The simulations were initialized with our local setup (eq. 2.8) and an initial magnetic field strength B_0. Each simulation was performed on uniform Cartesian grids of side L, resolution res and dimension D. We varied the size of the simulation domain L (scaling the conductivity as described in § 2.3.3), the plasma scale height H and the initial magnetic field strength B_0; the fiducial values for these parameters are included in the table header (— indicates the fiducial value). For each entry in the table, we performed simulations with both initially horizontal and vertical magnetic fields. As described in the § 2.5.1, these simulations include isotropic turbulence injected at the scale $k_0 = 2\pi/L \times \tilde{k}_0$ and with a range of turbulent energy injection rates to give $0.1 \lesssim t_{\rm buoy}/t_{\rm dist} \lesssim 10$.

to the values of $t_{\rm buoy}/t_{\rm dist}$). When $t_{\rm dist}$ is long compared to the buoyancy time, as in the first panel of Figure 2.10, the turbulence is weak; the HBI therefore dominates and the evolution of the plasma is similar to that described in section 2.4.1. This saturated state of the HBI feels a buoyant restoring force which resists vertical displacements ξ_z with a force per unit mass $f_{\rm buoy} = \omega_{\rm buoy}^2 \xi_z$. As we increase the strength of the applied turbulence in the following panels of Figure 2.10, the vertical displacements grow, and the field deviates more strongly from the $\hat{b}_z = 0$ equilibrium state of the HBI. When $t_{\rm dist}$ is short compared to the buoyancy time, as in the last panel of Figure 2.10, turbulence can displace the fluid elements faster than buoyancy can restore them. The turbulence then dominates the evolution of the plasma, tangling and isotropizing the field lines.

Figure 2.10 also shows the length-scale dependence of the transition from an HBI to a turbulence dominated state. It is clear in the second and third panels that the HBI has globally rearranged the field lines, but that the turbulence is increasingly efficient at smaller scales. This is a consequence of the fact that turbulence typically perturbs the plasma in a scale-dependent way, while the buoyant restoring force of the HBI does not. If the turbulence follows a Kolmogorov cascade, the force $\sim \omega^2/k \propto k^{1/3}$ increases with decreasing scale, so the turbulence will always win on sufficiently small length-scales. It is therefore somewhat ambiguous whether turbulence or the HBI dominates a certain configuration, as the answer will typically depend on scale. As mentioned earlier, we skirt this issue by defining $t_{\rm dist}$ at the scale where the velocity spectrum of the injected turbulence peaks. When assessing the astrophysical importance of the HBI, it is important to keep this scale in mind. If the scale where the turbulent energy spectrum peaks is smaller than the temperature gradient length scale, the HBI may still insulate the plasma against conduction, even if the field lines are isotropized on smaller scales.

In order to quantify the transition from an HBI-dominated configuration to one dominated by turbulence, we measure the mean orientation of the magnetic field via the volume average of $\hat{b}_z^2$. The saturated value for this quantity approaches zero when the HBI dominates, and $1/D$ when the magnetic field is isotropic, where D is the number of dimensions in the simulation.

Figure 2.11 shows the saturated field angle as a function $t_{\rm buoy}/t_{\rm dist}$. The points in this figure represent simulations with different driving scales k_0, different dimensionality, and different buoyancy times $t_{\rm buoy}$ (Table 2.3 summarizes our parameter study). Symmetry of the coordinate axes requires that $b_z^2 = 1/2$ in 2D or $1/3$ in 3D if the magnetic field is isotropic. To include both our 2D and 3D simulations on the same plot, we shift our 3D values of $\langle b_z^2 \rangle$ by a factor of $3/2$. To within the scatter shown in Figure 2.11, we find that the saturated value of $\langle b_z^2 \rangle$ depends only on the ratio $t_{\rm buoy}/t_{\rm dist}$: for $t_{\rm buoy} \gtrsim t_{\rm dist}$ the turbulence is strong and the field becomes relatively isotropic while for $t_{\rm buoy} \lesssim t_{\rm dist}$ the isotropic turbulence is weak and the HBI drives the magnetic field to become relatively horizontal. The fact that the transition between these two states occurs around $t_{\rm buoy} \sim t_{\rm dist}$ suggests that our definition of $t_{\rm dist}$, though somewhat arbitrary, is reasonable.

The bulk of the simulations in Figure 2.11 are 2D, and we do not have any 3D simulations in the very weak turbulence limit. These simulations are computationally expensive, both

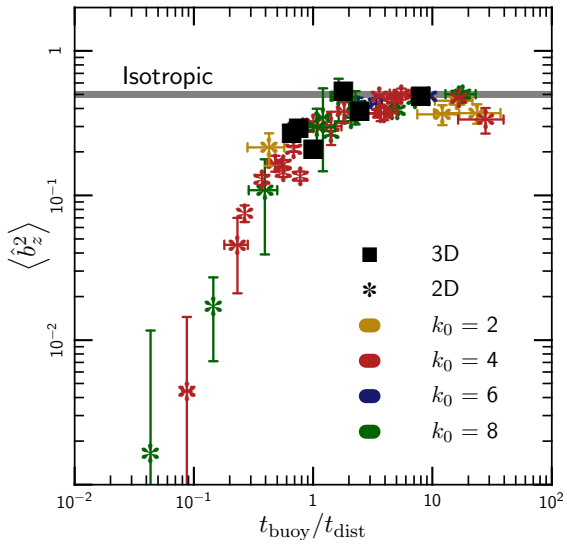

Figure 2.11: Saturated magnetic field orientation as a function of the strength of the externally driven turbulence, $t_{\rm buoy}/t_{\rm dist}$, for local HBI unstable atmospheres with $t_{\rm buoy} = 1$, $\sqrt{2}$ and $\sqrt{3}$. The thick gray line is an isotropic magnetic field in 2D. Colored points represent simulations with turbulence driven on different scales. Squares mark 3D calculations; the values of $\langle b_z^2 \rangle$ for the 3D simulations have been shifted by a factor 3/2 since isotropy implies $b_z^2 = 1/2$ in 2D but $b_z^2 = 1/3$ in 3D. Error bars represent 1σ statistical fluctuations in b_z and $t_{\rm dist}$.

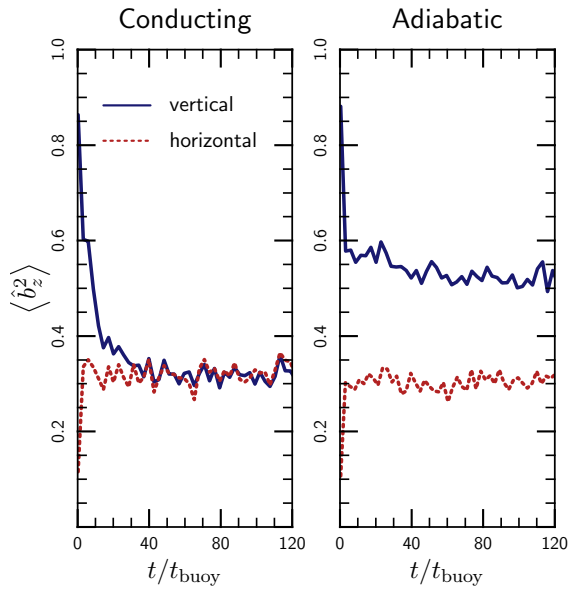

Figure 2.12: Magnetic field orientation as a function of time in 2D simulations with turbulence and the HBI for different initial magnetic field orientations; $t_{\rm buoy}/t_{\rm dist} = 0.8$. *Left panel:* Simulations with anisotropic thermal conduction. *Right panel:* Adiabatic simulations with no conduction (and thus no HBI). Different curves represent simulations with initially vertical and horizontal magnetic field lines, respectively. Conducting simulations eventually reach the same saturated state independent of the initial magnetic field direction. Adiabatic simulations are very similar to the conducting ones if the field is initially horizontal, highlighting the fact that in the saturated state of the HBI the plasma is buoyantly stable and behaves dynamically like an adiabatic fluid.

because of conduction and because we have to run for a long time for turbulence and the HBI to reach a statistical steady state; using 2D simulations allowed us to explore a larger fraction of the interesting parameter space. While the development of turbulence is very different in two and three dimensions, the HBI is essentially two-dimensional in nature. Moreover, the key dynamics governing the interaction between the HBI and the turbulence are dominated by the energy-containing scale of the turbulence—the precise power-spectrum of the fluctuations (which differs in 2D and 3D) is less critical. Scaling for dimension, we find that the saturated states of our 2D and 3D simulations are nearly identical. We thus believe that results in Figure 2.11 in the weak turbulence limit are a good description of the magnetic field structure in 3D systems as well.

These results on the interaction between the HBI and other sources of turbulence support an analogy between the saturated state of the HBI and ordinary, adiabatic stable stratification. The most important effect of the HBI in the saturated states of our simulations is to inhibit mixing in the direction of gravity. Ordinary stable stratification also inhibits vertical mixing and, as suggested by Sharma et al. (2009b), our parameter $t_{\rm buoy}/t_{\rm dist}$ is analogous to

the Richardson number used in the hydrodynamics literature.[7]. To further expand on this analogy, Figure 2.12 shows a comparison of anisotropically conducting (left panel) and adiabatic (right panel) simulations with equal values of $t_{\rm buoy}/t_{\rm dist}$. For the adiabatic simulations, we define $t_{\rm buoy} = t_{\rm ad}$, i.e., using the entropy gradient rather than the temperature gradient. The left panel of Figure 2.12 shows simulations with the same injected turbulence, but initialized with vertical or horizontal magnetic field lines. In these simulations, the saturated state is independent of the initial condition. That is, the interaction between turbulence and the HBI leads to a well defined magnetic field orientation that is independent of the initial field direction.

By contrast, in the adiabatic simulations (right panel of Fig. 2.12), the final magnetic field orientation depends on the initial field direction. For an initially vertical field in an adiabatic plasma, the turbulence slowly isotropizes the magnetic field direction. However, the adiabatic simulations with initially horizontal field lines reach a saturated state that is very similar to that of the HBI simulations. In the adiabatic simulations, the magnetic field is essentially passive, but it traces the fluid displacements. The stable stratification competes with the turbulence and sets a typical scale for vertical displacements in the saturated state. This scale, in turn, determines the magnetic field geometry. The magnetic field plays no dynamical role in this process. The fact that the anisotropically conducting simulations reach the same statistical steady state highlights that the saturated state of the HBI behaves dynamically very much like an adiabatic, stably stratified, plasma.

2.5.2 Effect of Turbulence on the MTI

To complete our analysis, we study how externally imposed isotropic turbulence affects the saturation of the MTI. Figure 2.13 shows the volume averaged magnetic field orientation as a function of time in MTI simulations with additional turbulence, for different values of the strength of the turbulence $t_{\rm buoy}/t_{\rm dist}$. These are local simulations (with domain sizes $L/H = 0.5$; initialized using eq. (2.9)) which have the pedagogical advantage of a positive entropy gradient, but which under-predict the kinetic energy generated by the MTI. We drive the turbulence at relatively small scales ($kL/2\pi = 8$) so that subsonic turbulence can still satisfy $t_{\rm buoy}/t_{\rm dist} > 1$. Both our driven turbulence and the MTI tend to isotropize the magnetic field, so it is not a priori clear whether the field orientation is a good indication of the importance of turbulence relative to the MTI. Although Figure 2.13 shows that there is no strong dependence of the saturated field orientation on the strength of the turbulence, the time dependence of the field orientation clearly shows the effects of the turbulence on the MTI.

In general, the evolution of the MTI proceeds through two stages: there is a linear phase, where the plasma accelerates toward its nominal stable state (which has a vertical magnetic field), and a nonlinear transition to the saturated state, where the strong turbulence generated by the MTI isotropizes the velocities and field lines. In the absence of additional

[7]The Richardson number is defined as $Ri \equiv gh/v^2$, where h is a characteristic vertical scale. Setting this equal to the scale where the turbulent spectrum peaks yields $Ri = (t_{\rm buoy}/t_{\rm dist})^2$

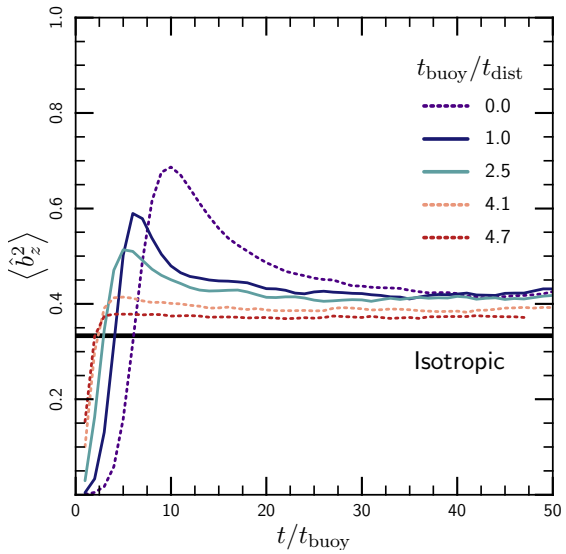

Figure 2.13: Magnetic field orientation as a function of time in 3D simulations with turbulence and the MTI, for different values of the strength of the turbulence $t_{\rm buoy}/t_{\rm dist}$. In all of the simulations, the magnetic field is relatively isotropic in the saturated state. However, the early-time 'overshoot' towards a vertical magnetic field due to the MTI is suppressed in the presence of strong turbulence with $t_{\rm buoy} \gtrsim t_{\rm dist}$. These relatively local simulations ($L = 0.5H$) underestimate the kinetic energy generated by the MTI (Fig. 2.7) and therefore the strength of the turbulence required to influence it.

turbulence, the linear phase is characterized by field lines that are primarily in the direction of gravity (Fig. 2.5). Figure 2.13 shows that additional sources of strong (rapidly shearing) turbulence suppress this linear phase of the MTI. Indeed, the evolution of the field angle with time in our strongest turbulence simulations ($t_{\rm buoy}/t_{\rm dist} = 4.7$) is quite similar to what we find in simulations of an adiabatic plasma in which the MTI is not present.

It would, however, be incorrect to conclude from Figure 2.13 that the MTI is unimportant if there are other strong sources of turbulence in the plasma. The fundamental reason for this is that the growth of the MTI does not depend significantly on scale, while the effects of the other sources of turbulence do. Figure 2.14 shows velocity spectra for the simulations in Figure 2.13; for comparison we also show the velocity spectra in adiabatic simulations, which correspond to the power spectra produced solely by the injected turbulence. Figure 2.13 demonstrates that even in simulations with very strong imposed turbulence ($t_{\rm buoy}/t_{\rm dist} = 4.7$) there is still significant excess power on the largest scales in the computational domain ($kL/2\pi \lesssim 10$). This large-scale power is due to the MTI. Moreover, the turbulent energy due to the MTI dominates the total turbulent kinetic energy in the plasma. These results highlight that 'strong' turbulence is a scale-dependent statement. Suppressing the MTI requires having $t_{\rm buoy}/t_{\rm dist} \gg 1$ on *all scales*, up to the temperature/pressure scale-height of the plasma. Because the MTI itself generates nearly sonic velocities, this suppression would

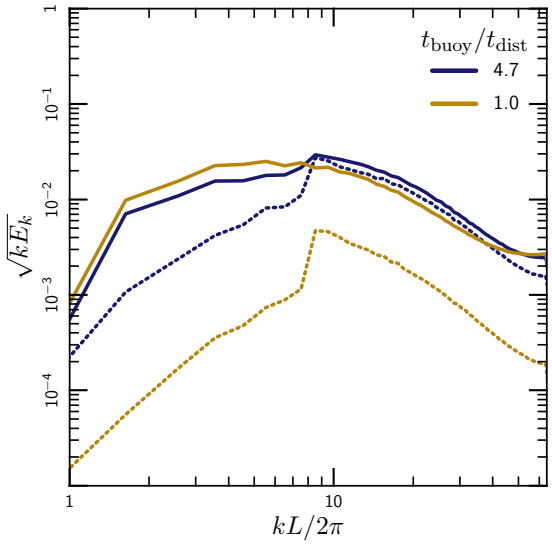

Figure 2.14: Velocity spectra in simulations with turbulence and the MTI (solid lines), for different values of the strength of the injected turbulence $t_{\rm buoy}/t_{\rm dist}$. The turbulence is driven at $kL/2\pi = 8$. For comparison, dashed lines show the power spectra in adiabatic simulations, in which the injected turbulence is the only source of power. Even when $t_{\rm buoy}/t_{\rm dist} \gg 1$ at the driving scale of the turbulence, so that one might expect the MTI to be suppressed, the MTI produces significant turbulent energy on larger scales $kL/2\pi \sim 2 - 10$. Suppressing the MTI on all scales in a plasma would thus require $t_{\rm buoy}/t_{\rm dist} \gg 1$ on all scales smaller than the scale-height.

require close to supersonic turbulence. In practice, it is therefore unlikely that additional sources of turbulence can fully suppress the MTI in most astrophysical environments where it is likely to occur (e.g., accretion disks and galaxy clusters).

2.6 Discussion

The motion of electrons along, but not across, magnetic field lines in dilute, magnetized plasmas produces efficient, anisotropic transport of heat. Such plasmas are therefore non-adiabatic, and the standard analysis of buoyancy (or convective) instabilities does not necessarily apply. Quantitatively, conduction plays an essential role on scales less than $7\,(\lambda H)^{1/2}$, where λ is the electron mean free path and H is the plasma scale height. In this "rapid conduction limit," the temperature gradient, rather than the entropy gradient, dictates the stability of the plasma, and the plasma is unstable for either sign of the temperature gradient (Balbus 2000; Quataert 2008). The convective instability in this limit is known as the HBI (MTI) when the temperature increases (decreases) with height.

Parrish & Stone (2005, 2007) and Parrish & Quataert (2008) extended the original linear analysis of the MTI and HBI into the nonlinear regime using numerical simulations. In this paper, we have reconsidered the nonlinear saturation of the HBI and MTI. Our work

adds to previous investigations because we have identified a key difference between the two instabilities and are able to understand the nonlinear behavior of the MTI more completely. This paper therefore represents a significant change in our understanding of the possible astrophysical implications of the MTI (but not the HBI). We have also studied the effect of an external source of turbulence on both the MTI and HBI. We conclude that other sources of turbulence in a plasma can change the saturation of the HBI, but that it is much harder to disrupt the MTI. Below we summarize our results and discuss their astrophysical implications, focusing on the intracluster medium (ICM) of galaxy clusters.

2.6.1 HBI

The HBI occurs whenever the temperature increases with height in an anisotropically conducting plasma. Plasmas in the rapid conduction limit are in general linearly unstable, unless the magnetic field lines are horizontal (eq. 2.7). Horizontal field lines represent a fixed point in the evolution of the plasma: if the plasma somehow reaches such a state (and is not perturbed away by another process), it will remain there forever. A horizontal magnetic field is therefore a natural saturation channel for the HBI.

This "magnetic" saturation mechanism can be understood using the linear dispersion relation of the plasma (§ 2.4.1). Starting from a linearly unstable state, the HBI induces both horizontal and vertical motions in the plasma. As the HBI develops, however, the vertical motions become trapped in internal gravity waves. These waves decay, leaving only the horizontal motions at late times; thus, the fluid velocities are very anisotropic in the saturated state. Since the horizontal motions don't incur a buoyant response, the horizontal displacements can be very large. These motions stretch out the magnetic field lines, amplifying and reorienting them, and drive the plasma towards its stable equilibrium with horizontal magnetic field lines. These horizontal motions therefore drive the nonlinear evolution and saturation of the HBI.

Since the saturation of the HBI is dominated by horizontal displacements, the key dynamics all occur at approximately the same height in the atmosphere. The saturation of the HBI is thus essentially local in nature. We have demonstrated this explicitly by carrying out simulations with different domain sizes relative to the plasma scale-height; the results of these simulations are very similar (Fig. 2.4).

The growth rate of the HBI decreases dramatically as the instability progresses. The HBI therefore saturates quiescently, and the velocities in the saturated state are very subsonic (Fig. 2.3). The saturation is driven by horizontal motions with nearly constant velocities, so the nonlinear magnetic field amplification is approximately linear, rather than exponential, with time. These findings are consistent with those of Parrish & Quataert (2008). We note that the HBI is very much unlike adiabatic convection, which can only saturate by changing the thermal state of the plasma and therefore generates vigorous turbulence. The key difference between the HBI and adiabatic convection is that the source of free energy for the HBI is a conductive heat flux through the plasma, not merely the existence of a temperature gradient. Since the heat flux can be suppressed by rearranging the magnetic

field, the HBI has a magnetic saturation channel that is not available to adiabatic convection.

The astrophysical implications of the HBI follow immediately from the nature of its saturated state. By reorienting the magnetic field lines, the HBI dramatically reduces the conductive heat flux through the plasma. The HBI should operate in the innermost ∼100–200 kpc in the intracluster medium of cool-core galaxy clusters, where the observed temperature increases outward. As noted in Parrish & Quataert (2008), this is precisely where the cooling time of the ICM is shorter than its age; the HBI removes thermal conduction as a source of energy for the cores, potentially exacerbating the cooling flow problem (Parrish et al. 2009).

Our results demonstrate that the saturated state of the HBI is buoyantly stable. This may not seem surprising, because it is exactly what one would expect if the ICM were adiabatic. The ICM is not, however, adiabatic, and thermal conduction would render it buoyantly neutral to vertical displacements if the magnetic field lines were tangled. The saturated state of the HBI is buoyantly stable only because of the nearly horizontal magnetic field lines (that are perpendicular to the temperature gradient). The HBI therefore inhibits vertical mixing and allows for the existence of weakly damped internal gravity waves in the plasma.

Sharma et al. (2009b) noted that the stable stratification associated with the saturated state of the HBI competes with other sources of turbulence in a well-defined way. This competition can be understood using a modified Richardson number $t_{\rm buoy}/t_{\rm dist}$, where $t_{\rm buoy}$ is the timescale for the HBI to grow and $t_{\rm dist}$ is a characteristic "distortion time," or "eddy turnover time," of the turbulence. When $t_{\rm buoy} \gtrsim t_{\rm dist}$, the turbulence can isotropize the plasma and remove all traces of the HBI (Fig. 2.11). When $t_{\rm buoy} \lesssim t_{\rm dist}$, the saturated state of the plasma represents a statistical balance between turbulence and the HBI, with the magnetic field becoming more horizontal, and the plasma more HBI-dominated, for smaller values of $t_{\rm buoy}/t_{\rm dist}$. The strength of other sources of turbulence is therefore crucial for understanding the astrophysical implications of the HBI.

Figure 2.11 provides a very simple mapping between the properties of the turbulence and the magnetic field geometry in the plasma. Given the strength of the turbulence and the thermal state of the plasma, this figure provides a recipe for determining the mean geometry of the magnetic field, and therefore the effective conductivity of the plasma. This can be used to interpret observational results or to construct semi-analytic models of anisotropic conduction for use in cosmological simulations.

Turbulence in the ICM is currently poorly constrained, and thus it is difficult to determine precisely how important the HBI is for the evolution of clusters. Reasonable estimates suggest that $t_{\rm buoy}/t_{\rm dist} \sim 1$, but more detailed simulations of clusters are required to determine this ratio more precisely. Future observations of the ICM with space-based x-ray calorimeters will place observational constraints on the level of turbulence. In addition, Faraday rotation measurements of the ICM will measure the orientation of the magnetic fields in clusters and constrain the role of the HBI (Bogdanović et al. 2011) (Pfrommer & Dursi (2010) describe another mechanism to measure the magnetic orientation in the ICM). Even if the turbulence is strong ($t_{\rm buoy}/t_{\rm dist} \gtrsim 1$), the driving scale and filling factor of the turbulence may allow

the HBI to dominate on some scales or at some locations (see Fig. 2.10 and associated discussion in § 2.5.1). As suggested by Parrish et al. (2010) and Ruszkowski & Oh (2010), the interaction between turbulence and the HBI might be part of a feedback loop for the thermal evolution of the ICM.

2.6.2 MTI

The nonlinear evolution of the MTI is more complex than that of the HBI. Just like the HBI, the MTI has linearly stable equilibria, but they are transposed: the linearly stable equilibrium states of the MTI have vertical field lines. We have shown, however, that the linearly stable equilibrium states turn out to be nonlinearly unstable; i.e., they are unstable to perturbations with a finite amplitude. This nonlinear instability arises because neutrally buoyant, horizontal displacements add a horizontal component to the magnetic field. This takes the plasma out of its linearly stable state and re-seeds the instability (Fig. 2.6 and § 2.4.2). As a result, the linearly stable equilibrium states of the MTI do not represent fixed points in the evolution of the instability, and the MTI cannot saturate simply by reorienting the magnetic field.

This difference eliminates the quiescent, magnetic saturation channel for the MTI and dramatically changes its evolution. Without a linear means to saturate, the instability grows until nonlinear effects dominate, which occurs when $v \sim c_s$. Unlike the HBI, the MTI therefore drives strong turbulence and operates as an efficient magnetic dynamo, much more akin to adiabatic convection. The astrophysical implications of the MTI are therefore entirely different from those of the HBI. The kinetic and magnetic energy generated by the MTI can contribute a significant (up to ten percent) non-thermal pressure support to the plasma in the saturated state. This is consistent with observational constraints on non-thermal pressure support in the ICM near the virial radius (George et al. 2009). This non-thermal pressure support may have consequences for mass estimates of clusters, which often rely on the assumption of hydrostatic equilibrium with thermal pressure support. Note, in particular, that the MTI is predicted to be present at precisely the same radii ($\gtrsim$ the scale radius) to which x-ray and SZ mass measurements are most sensitive.

Because the MTI operates by buoyantly accelerating fluid elements until they approach the sound speed, the results of numerical simulations of the MTI are sensitive to the size of the computational domain. The boundaries of the domain can artificially suppress this acceleration, and simulations with sizes smaller than a scale height under-predict the kinetic energy generated by the MTI (this was the case in the original MTI simulations of Parrish & Stone 2005, 2007). The nonlinear development of the MTI is therefore quite sensitive to the global thermal state of the plasma, and an understanding of the MTI requires more careful numerical simulations than are needed for the HBI.

The saturated state of the MTI corresponds to a largely isotropic magnetic field, with a slight but persistent vertical (or radial) bias; this bias is robust even in the presence of other sources of strong turbulence (Fig. 2.13). We find that the magnetic energy generated by the MTI saturates at about 30% of the kinetic energy (Fig. 2.8). However, it may be difficult

to observationally distinguish the turbulence generated by the MTI from that generated by other processes.

The large velocities generated by the MTI, along with correlations between the temperature and velocity perturbations, imply that the MTI drives a large convective heat flux, $\sim 1.5\% \times \rho c_s^3$ (Fig. 2.9). While this convective flux is probably too small to alter the thermal evolution of the ICM, it could be important in higher density astrophysical plasmas, where the electron mean free path is smaller and conduction isn't as efficient.

While the MTI cannot saturate by reorienting the magnetic field, it can saturate by making the plasma isothermal. Simulations with Neumann boundary conditions in which the temperature at the boundaries is free to adjust find that this is the case; the atmosphere becomes isothermal before the MTI has a chance to develop (Parrish et al. 2008). We fixed the temperature at the top and bottom boundaries of our simulations; this is partially motivated by the fact that many galaxy clusters in the local universe are observed to have non-negligible temperature gradients.

Sharma et al. (2008) carried out numerical simulations of the MTI in spherical accretion flows and found nearly radial magnetic fields, with modest turbulence. This quasi-linear saturation of the MTI might seem to contradict the results presented in this paper. Note, however, that in the Bondi inflow studied by Sharma et al. (2008), the plasma undergoes at most $\sim 5 - 10$ MTI growth times before flowing in. After 10 growth times, our simulations also show approximately radial field lines and modest turbulence (Fig. 2.8). Moreover, the simulations of Sharma et al. (2008) covered a very large dynamic range in radius and may have lacked the resolution to see the full nonlinear development of the MTI.

The MTI growth time in the outer parts of the ICM is about 1 Gyr. Although our typical MTI simulations take ∼10–20 growth times to saturate, this does not preclude the importance of the MTI in galaxy clusters. Figure 2.8 shows that there is a long, linear ramp-up phase where the instability grows from the tiny perturbations we apply to the nonlinear state. Astrophysical perturbations are unlikely to be this subsonic. Figure 2.13 shows that the MTI can saturate in ∼2–5 growth times when subjected to larger perturbations, suggesting that the MTI is likely to become nonlinear in the outer parts of clusters. Cosmological simulations will be required to fully understand the implications of the MTI for galaxy clusters.

Chapter 3

Thermal Instability in Gravitationally-Stratified Plasmas

3.1 Introduction

While the formation of dark matter halos can be understood via gravitational interactions alone, the combined effects of cooling and gravity are essential to galaxy formation (Rees & Ostriker 1977; Silk 1977; White & Rees 1978). This interplay remains poorly understood, however, because the dense plasma in many high-mass halos is predicted to cool and accrete far more rapidly than is observed (Peterson & Fabian 2006). As a result, theoretical models and numerical simulations routinely over-predict the amount of cooling and star formation in massive galaxies (e. g. Saro et al. 2006); this discrepancy is an example of the well-known "cooling-flow problem." Some studies avoid the cooling-flow problem by focusing only on very hot halos (e. g. Sijacki & Springel 2006) or by "pre-heating" the gas to very high entropies (e. g. Oh & Benson 2003; McCarthy et al. 2004); this solution cannot work in general, however, because such hot systems are not representative of the cluster population (Cavagnolo et al. 2008). In particular, the central cooling time in many clusters is shorter than the Hubble time. Significant heating ("feedback") is required even at low redshift to suppress cooling in high-mass halos (Benson et al. 2003), and thus to explain the observed cutoff in the galaxy luminosity function (Cole et al. 2001; Kochanek et al. 2001).

Detailed x-ray observations of groups and clusters also highlight the need for significant heating of the intracluster plasma. Though these objects contain large amounts of radiating plasma (Fabian 1994), their x-ray spectra indicate a paucity of material cooling below $\sim 1/3$ of the maximum temperature (Peterson & Fabian 2006). This demonstrates that most of the plasma radiates without actually cooling; i. e., an energy source heats the plasma at a rate similar to its cooling rate.

Although heating dramatically suppresses cooling in groups and clusters, there is clear evidence for *some* cool gas in these systems. Studying this cold material can provide an important window into the heating mechanisms in groups and clusters and may help us understand how the balance between heating and cooling is maintained. The existence

of a cold phase can be inferred from star formation (O'Dea et al. 2010), but it has also been directly imaged in a number of cases, revealing filamentary nebulae located tens of kiloparsecs from the center of the potential (e. g. Fabian et al. 2008; McDonald et al. 2010, 2011a). Despite more than five decades of study, the origin of these dramatic filaments has yet to be conclusively established: they have been interpreted as the remnant from an enormous, central explosion (Lynds & Sandage 1963; Lynds 1970), mass dropout from a cooling catastrophe (Fabian & Nulsen 1977; Cowie et al. 1980; Nulsen 1986), debris from a high-speed merger of two gas-rich galaxies (Holtzman et al. 1992), or material dredged from the central galaxy by rising bubbles inflated by its AGN (Fabian et al. 2003, 2008). Studies have shown, however, that the cold gas is highly correlated with short central cooling times in the hot intracluster plasma (e. g. Hu et al. 1985; Heckman et al. 1989; Cavagnolo et al. 2008; Rafferty et al. 2008, clearly illusrated in Voit et al. 2008), suggesting that its origin involves cooling of the intracluster medium (ICM).

In this paper, we investigate the possibility that the cold phase forms as a consequence of *local* thermal instability in a *globally stable* atmosphere. Though many authors (e. g. Fabian & Nulsen 1977; Nulsen 1986) have previously proposed that the filaments form via thermal instability, this idea has typically been analyzed in the context of a cooling-flow background. Subsequent analytic and numerical studies (e. g. Malagoli et al. 1987; Balbus 1988; Balbus & Soker 1989; Hattori & Habe 1990; Malagoli et al. 1990; Joung et al. 2011) showed, however, that the linear thermal instability is ineffective at amplifying perturbations in a cooling flow and concluded that it is unlikely unlikely to produce the cool filaments seen in many clusters. By contrast, the thermal instability is not suppressed in a globally stable atmosphere (Defouw 1970; Balbus 1986), which is now believed to be a better approximation to the thermal state of the ICM. Quantitatively studying the thermal instability in this context has proven difficult because of the cooling-flow problem: studies that include cooling and gravity generally find that the plasma is globally thermally unstable, and that the entire cluster core collapses monolithically.[1]

We avoid the cooling-flow problem in this paper using a new strategy. Rather than attempting ab initio calculations of heating in clusters, we start from the observational fact that the ICM does not cool catastrophically. We therefore implement a phenomenological heating model that enforces approximate thermal equilibrium when averaged over large scales. We use this model to study the formation of multi-phase structure and we compare our results with archival data for groups and clusters. We find that the thermal stability of the plasma does not depend on its convective stability (see section 3.4.2). Instead, we find that the ratio of the thermal instability timescale $t_{\rm TI}$ to the dynamical (or "free-fall") timescale $t_{\rm ff}$ governs the non-linear saturation of the local thermal instability: the plasma develops extended, multi-phase structure only where this ratio falls below a critical threshold (§3.5.1). This

[1]One-dimensional models of the ICM with simplified heating prescriptions can be stable or quasi-stable with episodes of heating and cooling (e. g. Guo & Oh 2008; Ciotti & Ostriker 2001); however, creating a realistic, stable model in two or three dimensions is significantly more challenging. Moreover, just as convection cannot be modeled in one dimension, the dynamics of cool, over-dense gas sinking through the hot atmosphere is absent in one-dimensional models.

conclusion is not sensitive to significant perturbations about our idealized feedback prescription (§3.5.4) and is unchanged even in the presence of very rapid thermal conduction (§3.7.3) (Though the threshold may depend somewhat on the geometry and initial conditions of the system; see section 3.8).

This paper is the first in a series; here we present our model of local thermal instability and demonstrate its properties and implications using linear theory and non-linear simulations. The aim of this paper is to develop an understanding of the essential physics of the problem and we therefore study stratified plasmas using idealized, plane-parallel calculations. In our companion paper (Sharma et al. 2012b; hereafter Paper II), we present more realistic calculations of groups and clusters with spherical geometries and NFW halos. In both papers, we focus our analysis on the transition of material from the hot phase to the cold phase; we are not yet able to quantitatively predict any precise properties (such as sizes or luminosities) of the cold filaments produced via thermal instability. We discuss in section 3.8 how our results can nonetheless be tested observationally.

Because we put in by hand that hot halos are in approximate global thermal equilibrium, our model provides no direct insight into how this balance is maintained. This is both a weakness and a strength of our current approach: though our setup is necessarily phenomenological, our results are not tied to any particular heating mechanism. Thus, we expect that our conclusions should apply to a wide range of systems, ranging in mass from galaxies to galaxy clusters. We return to this point in sections 3.5 and 3.8 and we study more physically motivated heating models in Paper II. Our present aim is not to identify a plausible solution to the cooling-flow problem, but rather to understand what implications a stabilizing heat source has for the local thermal stability and dynamics of the ICM.

The structure of this paper is as follows. In section 3.2, we describe our model for the plasma, including our phenomenological heating prescription. We describe our numerical method in section 3.3, linear theory results in section 3.4, and our primary numerical results in section 3.5. Section 3.6 provides a physical interpretation of the numerical results. For simplicity, we initially ignore magnetic fields and thermal conduction in this paper; section 3.7 shows results including these effects. Finally, in section 3.8, we speculate on the astrophysical implications of our model and compare our results with observational data from the ACCEPT catalog (Cavagnolo et al. 2009).

3.2 Plasma Model

In this section, we describe our model for the cooling, heating and dynamics of the plasma in a dark matter halo. Due to the wealth of observations of the ICM, we explicitly motivate our model for galaxy clusters, and some of the details we present in this section may not apply to galaxies. Nonetheless, our analysis is fairly general and we expect that some of our basic conclusions also hold massive galaxies (see Paper II for more details).

We model the plasma as an ideal gas, sitting in the fixed gravitational potential of the halo and subject to both optically-thin radiative cooling and heating by a stabilizing feedback

mechanism. In the interest of simplicity, we initially ignore both thermal conduction and the dynamical effect of the magnetic field; these effects are important in the ICM (see, e. g. McCourt et al. 2011b and references therein), but do not change our qualitative conclusions. We generalize our results to conducting, magnetized plasmas in section 3.7.

The equations for the conservation of mass and momentum in the plasma, and for the evolution of its internal energy are:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \boldsymbol{v}) = 0, \tag{3.1a}$$

$$\frac{\partial}{\partial t}(\rho \boldsymbol{v}) + \nabla \cdot (\rho \boldsymbol{v} \otimes \boldsymbol{v} + P \mathbf{g}) = \rho \boldsymbol{g}, \tag{3.1b}$$

$$\rho T \frac{ds}{dt} = \mathcal{H} - \mathcal{L}, \tag{3.1c}$$

where ρ is the mass density, $\boldsymbol{v}$ is the fluid velocity, $\otimes$ denotes a tensor product, P is the pressure, $\mathbf{g}$ is the unit matrix, $\boldsymbol{g}$ is the gravitational field, T is the temperature,

$$s = \frac{1}{\gamma - 1} \frac{k_\mathrm{B}}{\mu m_\mathrm{H}} \ln\left(\frac{P}{\rho^\gamma}\right) \tag{3.2}$$

is the entropy per unit mass, and $d/dt = \partial/\partial t + \boldsymbol{v} \cdot \nabla$ is the Lagrangian (or convective) time derivative. In equation 3.2, k_B is Boltzmann's constant and μm_H is the mean mass of the particles contributing to thermal pressure in the plasma. The functions $\mathcal{H}$ and $\mathcal{L}$ describe heating and cooling of the plasma, respectively; we explain our prescriptions for these processes in the following sections.

3.2.1 Feedback

The physical origin of heating in clusters remains uncertain, but it is simple to understand why our model requires the heating function $\mathcal{H}$. Equation 3.1c shows that the timescale for the ICM to cool in a cluster is $\sim nT\,|\mathcal{L} - \mathcal{H}|^{-1}$; if $\mathcal{H} = 0$, this timescale near the centers of many clusters can be orders of magnitude shorter than the Hubble time (this is the aforementioned cooling-flow problem). The continued existence of the ICM in these clusters therefore strongly suggests that it is very nearly in thermal equilibrium, with $\mathcal{H}$ approximately equal to $\mathcal{L}$ when averaged over sufficient length- or time-scales.[2] Nonetheless, the multi-phase structure seen in many clusters (e. g. McDonald et al. 2010, 2011a) suggests that the thermal instability also operates. For the purposes of this paper, we call this behavior globally stable, but locally thermally unstable.

[2] An alternative is a "cooling-flow" model, where the cooling gas flows inward and is replenished by continued accretion (see Fabian 1994). This is not a viable alternative to heating, however, as these models over-predict the rate of gas cooling to low temperatures and the star formation rates in clusters by a factor of 10–1000; furthermore, the resulting density profiles are strongly disfavored by x-ray observations (McNamara & Nulsen 2007). We therefore do not consider cooling-flow models in this work.

The processes maintaining global thermal stability in clusters are not fully understood. The condition of global stability with local instability constrains the possible heating mechanisms, however, and suggests a phenomenological model for heating in the ICM. This is a model in which $\mathcal{H} \approx \mathcal{L}$ on average, but not $\mathcal{H} = \mathcal{L}$ identically. We adopt a specific implementation of this feedback model which simply fixes thermal equilibrium at all radii in our model halos. We set

$$\mathcal{H} = \langle \mathcal{L} \rangle, \tag{3.3}$$

where $\langle \cdots \rangle$ denotes a spatial average at a given radius. Thus, heating in our simplified model is a function only of r and t. By construction, this heating function ensures global thermal equilibrium at all radii in the plasma (precluding a cooling catastrophe), but permits the thermal instability to grow on smaller scales. It thus captures what we believe is the essential physics for the formation of multi-phase structure and meets our observationally-motivated requirements for the thermal stability of the ICM.

Equation 3.3 can be roughly motivated by positing a causal relationship between cooling on small scales and heating on large scales. Accretion onto a central AGN induced by cooling at larger radii is a promising mechanism for this "feedback" (Pizzolato & Soker 2005, 2010), and feedback from star formation could play a similar role in lower mass halos. Our specific heating implementation instantaneously balances cooling in every radial shell—in detail, this behaviour is non-local, acausal, and unphysical. Equation 3.3 is intended to mimic the end result of very effective feedback, but does not directly model the feedback process. Finding a physically-motivated heating mechanism that also leads to global stability is an important goal in the theory of the ICM, but it is outside the scope of our present study.

Our heating model is necessarily idealized, and it is important to separate tautological results (put in by hand) from the results which more generally reflect the global stability and local instability of the plasma. Though the subtleties of feedback are likely to strongly affect the evolution of the plasma, we find that our qualitative conclusions are not sensitive to the precise form of our heating function. We demonstrate this in section 3.5.4 by applying spatial and temporal variations to equation 3.3. Moreover, the simulations in Paper II reach similar conclusions using a very different setup. Thus we believe that the results derived using equation 3.3 capture some of the essential (and robust) dynamics of local thermal instability in globally stable systems. It is, however, difficult to prove this conclusively given current uncertainties in the heating of the ICM.

In our heating model, spatial variations between heating and cooling drive thermal instability; a more realistic model would likely introduce temporal, in addition to spatial, variations. We show in §3.5.4 that our conclusions do not change unless these temporal fluctuations around the thermal equilibrium are very large ($\sim 300\%$). We choose to begin our study using equation 3.3 because it is analytically tractable and lends itself to a thorough investigation.

Equation 3.3 is appropriate for a heating process which distributes energy per unit volume, such as the dissipation of MHD waves. Other processes like photoelectric heating distribute energy per unit mass. Since it is not yet known how feedback energy is thermalized

in the ICM, we generalize equation 3.3 to other processes:

$$\mathcal{H} = n^{\alpha} \frac{\langle \mathcal{L} \rangle}{\langle n^{\alpha} \rangle}. \tag{3.4}$$

Here, $\alpha = 0$ corresponds to volumetric heating and $\alpha = 1$ corresponds to mass-weighted heating. We show in section 3.4.1 that the thermal instability takes ~ 3 times longer to develop in plasmas with $\alpha = 1$ than in plasmas with $\alpha = 0$; after scaling the timescales by this factor, however, we find very similar evolution for plasmas with volumetric and mass-weighted heating (see figure 3.2, below).

3.2.2 Cooling

In the idealized spirit of this paper, we adopt a simple cooling function $\mathcal{L}$ dominated by thermal Bremsstrahlung

$$\mathcal{L}_{\rm B} = n^2 \Lambda(T) = \Lambda_0 n^2 T^{1/2}, \tag{3.5}$$

where $n = \rho/\mu m_{\rm H}$ is the number density of particles in the plasma and we have introduced the standard notation $\Lambda(T)$ for consistency with other work. Our conclusions are not sensitive to the shape of the cooling function as long as the plasma remains locally thermally unstable (§3.4); this is the case in the ICM for temperatures above $\sim 10^4$K.

In an unstratified plasma, thermally unstable clumps of cool gas collapse to the Field length in the cold phase (the length-scale below which thermal conduction suppresses local thermal instability; Field 1965, Begelman & McKee 1990). Resolving the realistic Field length in the cold phase of the ICM is numerically impractical, so we introduce a temperature floor at which we truncate the cooling function (see Sharma et al. 2010, §2.2 for a discussion of this approximation; also see §3.5.3 of this paper). We use the modified cooling function

$$\mathcal{L} = \mathcal{L}_{\rm B}\, \Theta_{\rm H}(T - T_{\rm floor}), \tag{3.6}$$

where $\Theta_{\rm H}$ is the Heaviside function, and $T_{\rm floor}$ effectively becomes the temperature of the cold phase.

The microphysical processes heating and cooling the cold phase in the ICM are likely to be very complicated (see Ferland et al. 2009) and we do not consider them here. Our use of a temperature floor amounts to the reasonable assumption that, once a thermally-unstable fluid element cools below $T_{\rm floor}$, it is unlikely to enter back into the hot phase. This simplification, along with our omission of line-cooling from equation 3.5, prevents us from studying the evolution of the cold material in detail. This is not a major limitation, however, because we are primarily interested in the transition of material from the hot phase to the cold phase. Following the internal structure of the cold clumps would be crucial for calculating the emission from filaments or for studying the intermittency in the accreted mass flux, but these applications are beyond the scope of our present study.

The simplified cooling function used here (eq. 3.6) prevents the gas from cooling below $T_{\rm floor}$ and therefore artificially lowers the gas density in thermally unstable clumps or filaments. We have confirmed, however, that the quantitative results in this paper are insensitive to the numeric value of $T_{\rm floor}$, provided it is much lower than the initial (or virial) temperature of the plasma. In Paper II we use a realistic cooling function that includes both Bremsstrahlung and line emission, and which does not implement a temperature floor. The results from this more realistic model agree with our conclusions here.

3.3 Numerical Model

We solve equations 3.1a–3.1c using the conservative MHD code ATHENA, modified to implement equations 3.4 and 3.6 via a semi-implicit, operator-split method (Sharma et al. 2010). Specifically, we evolve the thermal energy per unit volume $E = n\,k_{\rm B}T/(\gamma-1)$ using

$$\delta E^{(n)} = \left(\mathcal{H}^{(n)} - \mathcal{L}^{(n)}\right)\,\delta t \tag{3.7a}$$

$$E^{(n+1)} = \begin{cases} E^{(n)} + \delta E^{(n)} & \delta E^{(n)} > 0 \\ E^{(n)}/\left(1 + |\delta E^{(n)}/E^{(n)}|\right) & \delta E^{(n)} < 0 \end{cases}, \tag{3.7b}$$

where δ indicates a finite approximation to a differential, and $f^{(n)}$ denotes the function f during the $n^{\rm th}$ time-step of the simulation. This method explicitly prevents the temperature from becoming negative, even in the extreme case that the cooling time becomes shorter than the simulation time-step (although equation 3.7 is no longer accurate in this limit). Equation 3.7 is asymmetric and is only accurate to first order in $\delta t/t_{\rm cool}$. In order to test the sensitivity of our simulations to these shortcomings, we have also run simulations using a fully explicit, sub-cycled method. The two methods yield very similar results. We use equation 3.7 because it is faster than an explicit method and because it does not alter our results.

We perform most of our calculations on 2D Cartesian grids of resolution $(300)^2$ or 3D Cartesian grids of resolution $(128)^3$. We show a resolution study in section 3.5.3. In the remaining sections, as in our simulations, we work in units with $k_{\rm B} = \mu m_{\rm p} = 1$.

We perform our calculations in the plane-parallel approximation, with $\boldsymbol{g} = -g(z)\,\hat{\boldsymbol{z}}$. We therefore use the words 'height' and 'radius' interchangeably in the following sections. We make our setup symmetric about the $z = 0$ plane, with

$$g = g_0 \frac{z/a}{\left[1 + (z/a)^2\right]^{1/2}}. \tag{3.8}$$

Thus, g is nearly constant outside $|z| = a$, with a smooth transition through zero at the center. This setup enables us to place the computational boundaries far from the center, where most of the cooling and feedback take place (see figure 3.1). To further diminish the influence of the boundaries, we end our simulations before one cooling time transpires at

the boundary. We use reflecting boundary conditions in the direction parallel to gravity and periodic boundary conditions in the orthogonal directions.

We set the softening radius $a = 0.1\,H$, where H is the plasma scale-height (defined below). We turn off cooling and heating within $|z| \leq a$ because the physics at small radii is particularly uncertain and our feedback prescription (equation 3.4) may not be a good approximation to what happens there. We allow cold material to accumulate in the center $|z| \leq a$, but we otherwise ignore this region in our analysis. We have also performed simulations in which we do not turn off cooling in the center and have confirmed that it does not change our conclusions at larger radii $z \gtrsim H$.

We initialize the ICM in hydrostatic equilibrium, with a constant temperature T_0 and with the density profile

$$\rho(z) = \rho_0 \, \exp\left[-\frac{a}{H}\left(\left[1 + (z/a)^2\right]^{1/2} - 1\right)\right], \tag{3.9}$$

where the scale-height $H = T_0/g_0$. For computational convenience, we set $\rho_0 = T_0 = g_0 = 1$ and we take $T_{\rm floor} = 1/20$. This roughly corresponds to a virialized halo, in which the thermal and gravitational energy in the plasma are approximately equal. The atmosphere defined by equation 3.9 is buoyantly stable, with $\partial s/\partial z > 0$. To test the sensitivity of our results to stratification, we also use the buoyantly neutral atmosphere defined by

$$T(z) = T_0\left[1 - \frac{\gamma - 1}{\gamma}\frac{a}{H}\left(\left[1 + (z/a)^2\right]^{1/2} - 1\right)\right], \tag{3.10a}$$

$$\rho(z) = \rho_0\left(\frac{T}{T_0}\right)^{1/(\gamma-1)}. \tag{3.10b}$$

We refer to the conditions defined by equations 3.9 and 3.10 as isothermal and isentropic, respectively. Note that our use of the entropy gradient to determine convective stability is only appropriate because we have neglected conduction. When thermal conduction is efficient, the temperature gradient and the magnetic field orientation determine the convective stability of the plasma (Balbus 2000; Quataert 2008). We describe this in more detail in section 3.7.

We seed thermal instability in our model atmospheres by applying an isobaric perturbation with a flat spectrum ranging from $k = 2\pi/L$ to $k = 40\pi/L$, where L is the size of the simulation domain. The cutoff at high k makes the perturbation independent of resolution and permits a detailed convergence study. Unless otherwise noted, the modes of this perturbation have Gaussian-random amplitudes with an RMS value of 10^{-2}.

We define the free-fall time and the cooling time as follows:

$$t_{\rm ff} = \left(\frac{2z}{g_0}\right)^{1/2} \tag{3.11a}$$

$$t_{\rm cool} = \frac{3}{2}\frac{T^{1/2}}{n\Lambda_0}. \tag{3.11b}$$

Since these timescales are functions of height in our simulations, we quote them in the plane $z = H$ to give single values. When our analysis depends on the ratio $t_{\text{TI}}/t_{\text{ff}}$, we restrict it to this plane. We perform simulations with different initial values of the ratio $t_{\text{cool}}/t_{\text{ff}}$ by changing the parameter Λ_0; this permits direct and unambiguous comparison among our simulations because each is initialized identically. In reality, of course, Λ_0 is set by fundamental physics, and different values of $t_{\text{cool}}/t_{\text{ff}}$ correspond to clusters with different ICM entropies or densities.

3.4 Linear Theory Results

Equations 3.1–3.6 completely specify our model. In the rest of this paper, we study the properties of this model and apply it to astrophysical systems. In this section, we describe the linear stability of our model and derive the timescale for the formation of multi-phase structure in the plasma. We discuss the well-known linear results in some detail because they inform our interpretation of the non-linear behavior described later. In addition, the interpretation of these linear results has generated some confusion in the literature, leading to conflicting claims about the thermal stability of gas in hot halos.

3.4.1 Linear Stability

We define the net cooling rate $\Theta = \mathcal{L} - \mathcal{H}$ and assume that the plasma is initially in thermal equilibrium with $\Theta = 0$ everywhere. The derivative $(\partial\Theta/\partial T)_P$ describes how the net cooling responds to a linear, Eulerian perturbation. If this derivative is negative, a decrease in temperature at a fixed location in the plasma leads to an increase in the net cooling rate; thus, the temperature decreases further and the perturbed fluid element runs away to low temperatures. Similarly, an increase in temperature causes the fluid element to run away to high temperatures. The plasma is therefore unstable to small temperature fluctuations when $(\partial\Theta/\partial T)_P < 0$. A similar line of reasoning demonstrates that the plasma is thermally stable if $(\partial\Theta/\partial T)_P > 0$. Following Field (1965), we derive this result by linearizing and perturbing equations 3.1a–3.1c. This analysis yields the linear growth rate of the perturbations and will assist our interpretation of the non-linear results presented later.

We Fourier transform equations 3.1a–3.1c and perform a standard WKB analysis. We seek solutions with growth times much longer than the sound-crossing time and therefore make the Boussinesq approximation, which filters out sound waves (Balbus 2000, 2001).[3]

[3]We show later that the most important modes have wavelengths of order the Field length, $\sim 10\,\text{kpc}$. These modes have a sound-crossing time less than 10% of the cooling time; thus, they cool almost isobarically.

Under these approximations, the dynamical equations become

$$\boldsymbol{k} \cdot \delta \boldsymbol{v} = 0 \tag{3.12a}$$

$$-i\,\omega k^2 \delta \boldsymbol{v} = -\frac{\delta n}{n}\left[k^2 \boldsymbol{g} - \boldsymbol{k}(\boldsymbol{k}\cdot\boldsymbol{g})\right] \tag{3.12b}$$

$$-i\,\omega \delta s + \delta v_z \frac{\partial s}{\partial z} = -\frac{\delta \Theta}{nT}. \tag{3.12c}$$

In deriving equation 3.12b, we have crossed the momentum equation with $\boldsymbol{k}$ twice and used equation 3.12a to eliminate the compressive component of the velocity. This is consistent with the Boussinesq approximation and simplifies the algebra later on. Additionally, in the Boussinesq limit,

$$\delta s = -\frac{\gamma}{\gamma - 1}\frac{\delta n}{n} \tag{3.13}$$

and

$$\delta \Theta = -T\left(\frac{\partial \Theta}{\partial T}\right)_P \frac{\delta n}{n}. \tag{3.14}$$

In deriving equation 3.14, we have used the thermodynamic identity

$$\left(\frac{\partial \ln X}{\partial \ln T}\right)_P = \left(\frac{\partial \ln X}{\partial \ln T}\right)_n - \left(\frac{\partial \ln X}{\partial \ln n}\right)_T \tag{3.15}$$

for any state function $X(n, T)$. Note that, although the net cooling rate Θ varies explicitly with position, this dependence does not enter into equation 3.14 because $\delta\Theta$ represents an Eulerian perturbation at a fixed point in space. Although heating in our model has explicit radial and temporal dependencies, it experiences no first-order change under an Eulerian perturbation. We therefore ignore changes to the heating in this linear analysis. This should not give the impression that heating is immaterial to the linear results; on the contrary, these results presume an initial equilibrium state with a stabilizing heat source. The growth of the thermal instability is very different in the absence of such heating (Balbus 1988; Balbus & Soker 1989).

Combining equations 3.12–3.14, we find that the linear dispersion relation for the plasma is

$$\omega^2 - i\,\frac{T}{\mathcal{L}}\left(\frac{\partial \Theta}{\partial T}\right)_P \omega\omega_{\rm cool} - N^2\left(1 - \hat{k}_z^2\right) = 0 \tag{3.16}$$

where

$$\omega_{\rm cool} = \frac{\gamma - 1}{\gamma}\frac{\mathcal{L}}{nT} = (\gamma\, t_{\rm cool})^{-1} \tag{3.17}$$

is the cooling rate,

$$N = \sqrt{\frac{\gamma - 1}{\gamma}\, g \frac{\partial s}{\partial z}} \tag{3.18}$$

is the frequency for internal gravity waves, and $\hat{\boldsymbol{k}} = \boldsymbol{k}/k$ is the direction of the wave vector of the perturbation. As noted previously, we have neglected conduction and thus equations 3.16 and 3.18 only apply on relatively large length scales, $\gtrsim$ the Field length (see section 3.7). Equation 3.16 implies that perturbations grow exponentially in amplitude $\sim e^{p_{\rm TI} t}$, with

$$p_{\rm TI} = -\frac{\gamma - 1}{\gamma} \frac{1}{n} \left(\frac{\partial \Theta}{\partial T} \right)_P \tag{3.19a}$$

$$= \frac{\gamma - 1}{\gamma} \left(2 - \frac{\partial \ln \Lambda}{\partial \ln T} - \alpha \right) \frac{\mathcal{L}}{nT} \tag{3.19b}$$

$$= \left(\frac{3}{2} - \alpha \right) \omega_{\rm cool}. \tag{3.19c}$$

The three forms of equation 3.19 are equivalent and are useful in different contexts. In equation 3.19c, we have specialized to Bremsstrahlung cooling. In this case, plasmas with $\alpha < 3/2$ are locally thermally unstable (with $p_{\rm TI} > 0$), even though our model is (by construction) globally stable against a cooling catastrophe.

3.4.2 Local Stability, Global Stability, and Convection

Following Field (1965) and Defouw (1970), we showed in the previous section that the ICM is likely to be locally thermally unstable, and we propose that thermal instability may produce at least some of the multi-phase structure in galaxy clusters. At first, our analysis may appear inconsistent with other claims (such as can be found in, e. g. Balbus & Soker 1989 and Binney et al. 2009) about the importance of local thermal instability in galaxy and cluster halos. We review this apparent contradiction here and show that there is no inconsistency.

Balbus (1988) and Balbus & Soker (1989) extensively studied thermal instability using Lagrangian techniques and discovered that it is significantly stabilized in a cooling-flow. In a globally stable atmosphere, however, perturbations *do* grow exponentially (Defouw 1970; Balbus 1986). Since it is now thought that clusters are globally thermally stable and that the ICM persists for many cooling times, we expect the thermal instability to undergo many e-foldings and to become highly non-linear in clusters (though this does not always imply a large amplitude; see §3.6). Thus, assumptions about the global stability of the ICM also dictate conclusions about its local thermal stability and one must be careful to choose an appropriate background model.

Even though we expect perturbations to grow exponentially in clusters, they do not necessarily grow monotonically: equation 3.16 shows that a thermally unstable perturbation

oscillates as it grows if the cooling time is longer than the buoyancy time. This overstability represents a driven gravity wave (see Defouw 1970). Since the thermal instability in this case is not purely condensational, its identification with multi-phase gas becomes somewhat unclear (Malagoli et al. 1987; Binney et al. 2009). However, the growth rate of the thermal instability is essentially unaffected by buoyancy (equation 3.16), and thus perturbations are also likely to become highly non-linear in this limit. Earlier studies of the thermal overstability in stratified plasmas have either focused entirely on the linear evolution of perturbations (Defouw 1970; Malagoli et al. 1987; Binney et al. 2009) or have studied them in the context of a cooling-flow (Hattori & Habe 1990; Malagoli et al. 1990; Joung et al. 2011), in which the thermal instability is suppressed (Balbus & Soker 1989).

For the reasons listed above, we argue that earlier studies cannot directly predict the astrophysical implications of thermal instability in cluster halos. The astrophysical implications of the thermal instability depend on how the linear growth saturates in a globally stable environment. This motivates our present study. A series of previous investigations are very similar to ours (Nulsen 1986; Pizzolato & Soker 2005; Soker 2006; Pizzolato & Soker 2010), but focus on the survival of preexisting cold filaments rather than their formation via thermal instability. Our investigation compliments these studies and produces the initial conditions they require.

The saturation of the thermal instability involves the sinking of cool over-densities; in this respect, it bears some similarity to convection. This connection between thermal and convective stability was first recognized by Defouw (1970) and was significantly sharpened by Balbus & Soker (1989). Specifically, Balbus & Soker (1989) showed that thermal instability necessarily implies convective instability if the heating and cooling are state functions of the plasma. Heating in galaxy groups and clusters is very unlikely to be a state function of the ICM plasma, however. As a concrete example of spatially dependent heating, consider heating by turbulence (induced by, e. g. buoyant bubbles created by star formation or an AGN). The heating rate in this case is set by the rate at which turbulent energy is transferred to small scales, and thus by the turbulence properties as a function of position. In this case, i. e. when the heating depends explicitly on position, there is no one-to-one relationship between convective and thermal stability (as noted by Balbus & Soker 1989). Fundamentally, buoyancy determines convective stability, while heating and cooling determine thermal stability; these processes are not related in a globally stable atmosphere, and the thermal stability of an atmosphere is independent of its convective stability.

3.5 Simulation Results

We extend our analysis into the non-linear regime using the numerical setup described in section 3.3. We have run a large suite of 2D and 3D simulations, summarized in Table 3.1. We focus our analysis on the presence of multi-phase structure (§3.5.1) and on the accreted mass flux (§3.5.2), both of which can be compared with observations of groups and clusters.

Equation 3.19c shows that the growth rate of the thermal instability is a factor of 3

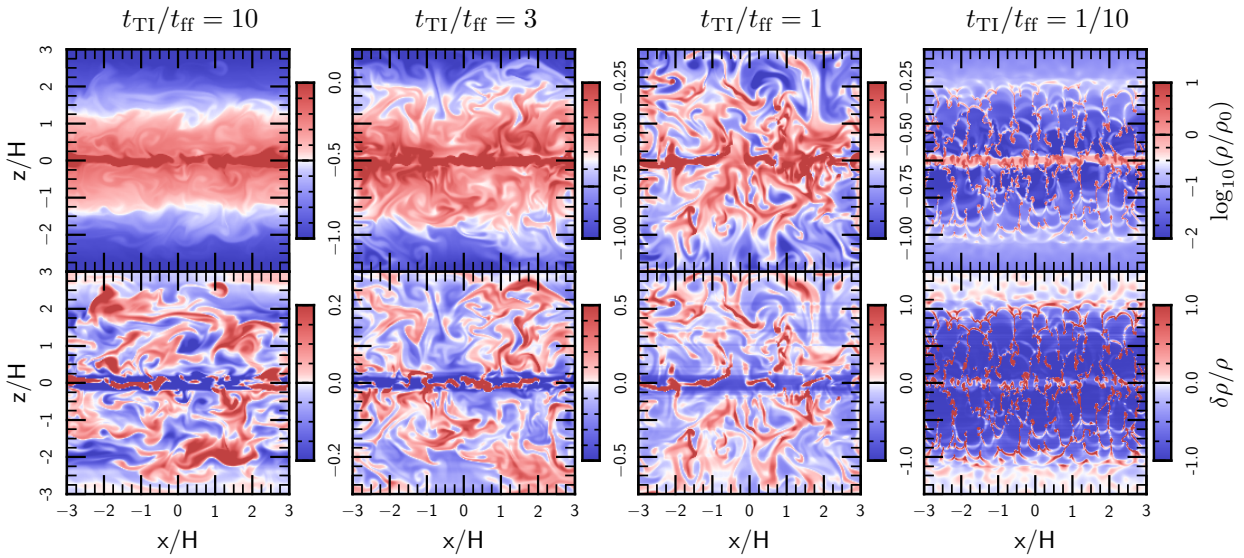

Figure 3.1: Snapshots of the density (top) and fractional density inhomogeneity $\delta\rho/\rho \equiv (\rho - \langle\rho\rangle)/\langle\rho\rangle$ (bottom) at the time $t = 10\,t_{\rm TI}(z = H)$ in our simulations. Note that our simulations are symmetric about the plane $z = 0$; this enables us to put the boundaries far from the center, where most of the cooling and feedback take place. Gravity points down in the top half of the domain and up in the bottom half (for the remainder of the paper, we primarily show images of the top of the domain). We applied the heating function in equation 3.4 to ensure global thermal stability, distributing the energy per unit volume ($\alpha = 0$). From left to right, these simulations have initial values of $t_{\rm TI}/t_{\rm ff} = 10, 3, 1$ and 0.1. These simulations demonstrate that cooling and heating drive internal gravity waves when $t_{\rm TI}/t_{\rm ff} \gtrsim 1$. The amplitude of these waves increases with the cooling rate and approaches the size of the simulation domain when $t_{\rm TI} \sim t_{\rm ff}$. When $t_{\rm TI}/t_{\rm ff} < 1$, the thermally unstable gas collapses into dense clumps, which then rain down into the center of the potential. For clarity, we have restricted the color bar on plots with $t_{\rm TI}/t_{\rm ff} = 1/10$. In this simulation, $(\delta\rho/\rho)_{\rm max} \sim 20$ is set by the (arbitrary) temperature floor we impose. While the simulations have been run for 10 thermal instability times at $z = H$, gas near the boundaries has not yet had time to cool. The initial perturbations are still visible near the boundaries in the lower-rightmost plot. This figure also clearly shows the accumulation of cool material in the center of our simulation domain. We describe this process in more detail in §3.6. Animated versions of the figures in this paper can be found at: http://astro.berkeley.edu/∼mkmcc/research/thermal_instability/movies.html.

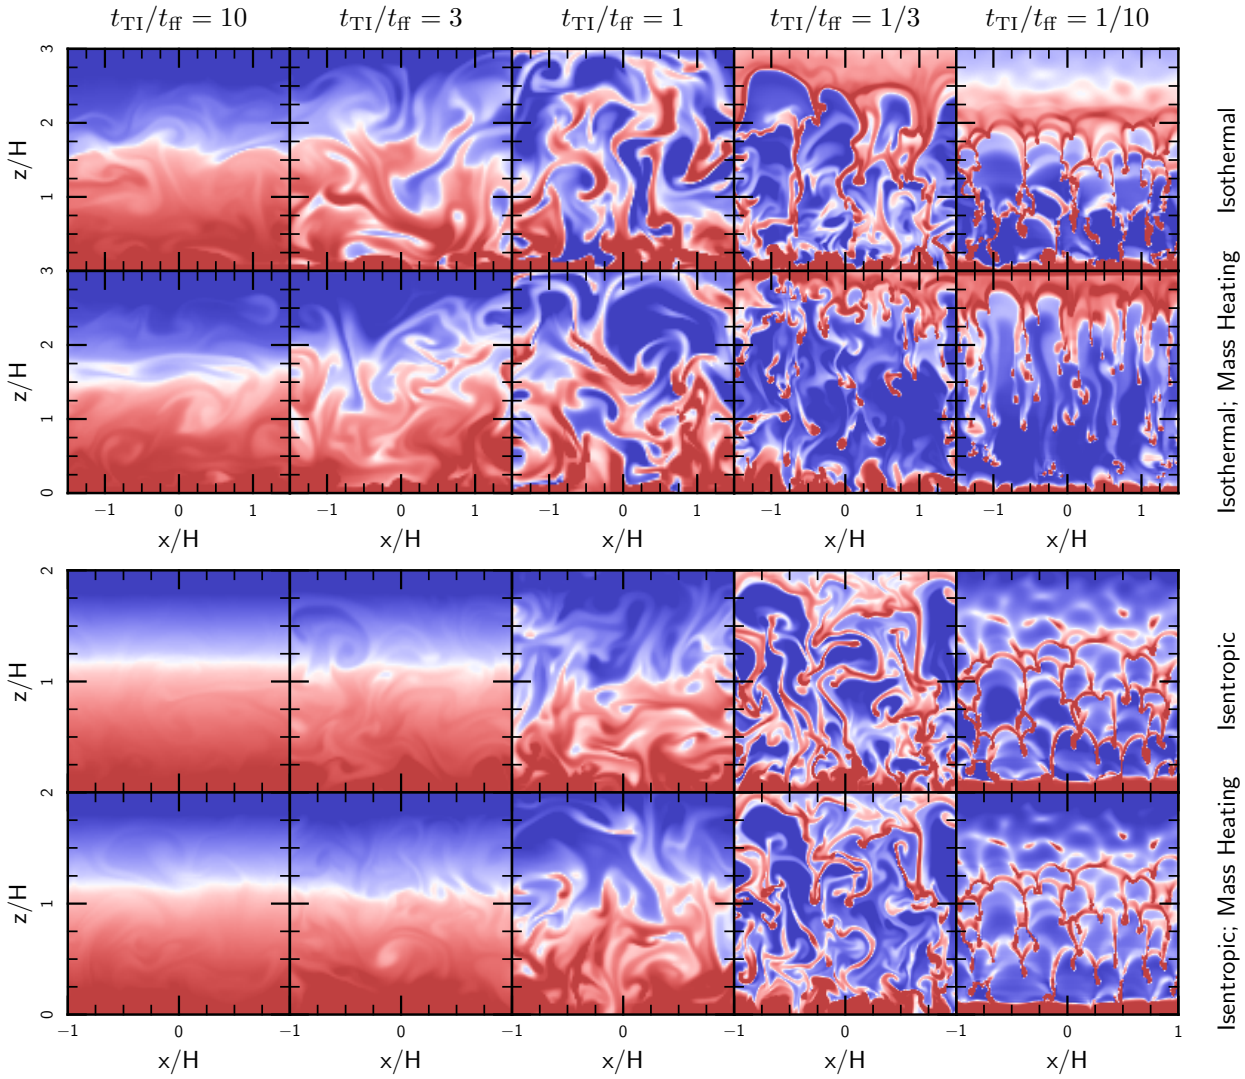

Figure 3.2: Snapshots of the density at the time $t = 10\,t_{TI}$, for different values of the time-scale ratio t_{TI}/t_{ff}. The top two rows show our simulations with isothermal initial conditions and volume and mass-weighted heating, while the bottom two rows show isentropic initial conditions with volume and mass-weighted heating. These results show that the non-linear behavior of the thermal instability is relatively independent of the initial stratification and the details of the heating. Note that the color scale was chosen to show the features in the gas and varies from plot to plot; the ranges for a given value of t_{TI}/t_{ff} are similar to those shown in figure 3.1.

Table 3.1: Parameters for simulations without conduction (§3.5).

Initial Condition	L/H	α	$t_{\rm TI}/t_{\rm ff}$
Isothermal	3	0	$1.57\,\Lambda_0^{-1}$
		1	$4.70\,\Lambda_0^{-1}$
Isentropic	2	0	$1.25\,\Lambda_0^{-1}$
		1	$3.74\,\Lambda_0^{-1}$

We performed all simulations on square Cartesian grids of resolution $(300)^2$ or $(128)^3$ and physical size $2L$ (the scale-height H is defined in §3.3). We also performed simulations at other resolutions as part of a convergence study (§3.5.3). The cooling constant Λ_0 is a free parameter in our model, which we choose to obtain the desired $t_{\rm TI}/t_{\rm ff}$. Each combination of the listed parameters was simulated with initial values of $\log_{10}(t_{\rm TI}/t_{\rm ff})$ at $z = H$ spanning between -1 and 1 with increments of 1/4. The top row represents our fiducial setup; we also performed 3D simulations using this setup with $\log_{10}(t_{\rm TI}/t_{\rm ff}) = -1$, -0.75, -0.5, 0, and 1.

smaller in plasmas with heating per unit mass than in plasmas with heating per unit volume. More generally, the timescale depends on the uncertain parameter α and cannot be directly applied to (or inferred from) observations. Nonetheless, it is convenient to use $t_{\rm TI} = p_{\rm TI}^{-1}$ to normalize time when considering the physics of the thermal instability with different values of α. We also use the cooling time $t_{\rm cool} \equiv E/\mathcal{L} = (\gamma\,\omega_{\rm cool})^{-1}$ when we compare our results with observations. These two timescales differ only by an uncertain factor of order unity.

3.5.1 Multi-phase Structure

We performed simulations with the ratio of time-scales $t_{\rm TI}/t_{\rm ff}$ ranging from 0.1 to 10 (measured at $z = H$) and ran each for ten growth times, until $t = 10\,t_{\rm TI}$. Figure 3.1 shows representative snapshots of the density at the end of our fiducial simulations with volumetric heating and isothermal initial conditions. Our simulations show that plasmas with cooling times shorter than the dynamical time ($t_{\rm TI} \ll t_{\rm ff}$) develop spatially extended multi-phase structure, whereas plasmas with cooling times longer than the dynamical time ($t_{\rm TI} \gtrsim t_{\rm ff}$) do not. Thus, the ratio $t_{\rm TI}/t_{\rm ff}$ controls the non-linear saturation of the thermal instability in stratified plasmas; this observationally-testable prediction is the primary result of our study. This conclusion does not depend strongly on either the initial stratification of the plasma or on our choice of heating per unit volume. To demonstrate this, figure 3.2 shows variations of our fiducial simulations with isentropic initial conditions and with mass-weighted heating. In all four cases, the saturated state transitions from single-phase to multi-phase when the ratio of time-scales $t_{\rm TI}/t_{\rm ff}$ becomes less than one. Below, we describe the plasma properties in these two limits and the physics of the transition between them.

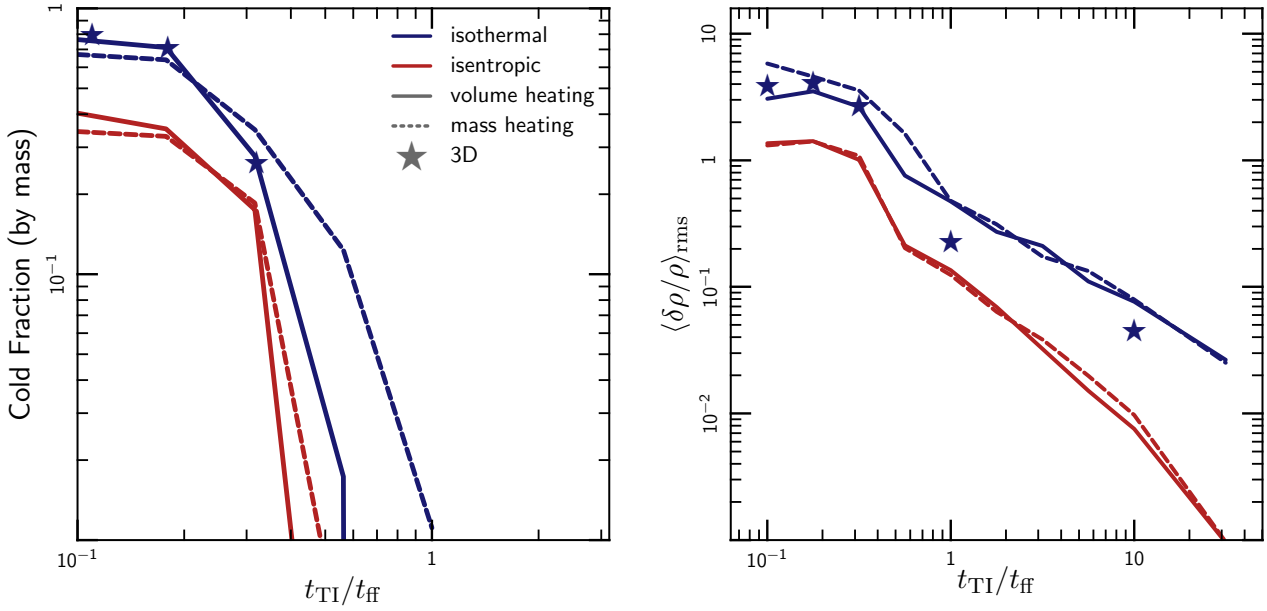

Figure 3.3: *(Left:)* Mass fraction of cold material (with $T \leq T_{\rm initial}/3$) as a function of the timescale ratio $t_{\rm TI}/t_{\rm ff}$. The mass in cold material drops off sharply when $t_{\rm TI} \sim t_{\rm ff}$, and there is no extended multi-phase structure in the weak-cooling limit. All quantities in these plots represent averages from $z = 0.9$–$1.1\,H$ and from $t = 9$–$10\,t_{\rm TI}$. *(Right:)* Fractional density inhomogeneity $\delta\rho/\rho$ as a function of the timescale ratio $t_{\rm TI}/t_{\rm ff}$. Blue (red) lines indicate isothermal (isentropic) initial conditions, solid (dashed) lines indicate volumetric (mass-weighted) heating. The blue stars represent 3D simulations using our fiducial setup (isothermal initial condition with volumetric heating); the remaining simulations are 2D.

The evolution of plasmas with short cooling times $t_{\rm TI} \ll t_{\rm ff}$ is straightforward: in this limit, the thermal instability develops and saturates before the plasma can buoyantly respond. The initial perturbations therefore collapse into dense clumps essentially in-situ, and the ICM develops a highly inhomogeneous, multi-phase structure wherever $t_{\rm TI}(z) < t$. The clumps of cold gas then rain down onto the central galaxy on the (much longer) free-fall time, while bubbles of heated gas rise outwards. The rightmost panels of figures 3.1 and 3.2 illustrate this behavior. The result is a hotter atmosphere (in which $t_{\rm TI}/t_{\rm ff} > 1$), filled with clumps of cold gas. We show in Paper II that this end state resembles the observed properties of some cool-core groups and clusters.

The saturation of the thermal instability is fundamentally different when the cooling time is long compared to the dynamical time. In this limit, gravity and buoyancy influence the linear evolution of the perturbations (though the growth rate changes only by a factor of two). Nonlinearly, however, buoyancy provides a critical saturation channel for the thermal instability that prevents the formation of multi-phase gas. This conclusion is qualitatively similar to that reached by Balbus & Soker (1989); however, the physics is very different in our case because the background atmosphere remains statistically in thermal equilibrium for many cooling timescales. As initial perturbations cool and grow, they sink in the gravitational potential and mix with gas at lower radii. The cooling thus drives a slow, inward flow of material; the associated mass flux is, however, significantly smaller than is predicted by models without heating. We return to this point in the following sections. Rather than creating strong density inhomogeneities, cooling in this limit excites internal gravity waves with an amplitude that depends on the timescale ratio $t_{\rm TI}/t_{\rm ff}$. These waves represent the overstability highlighted by Balbus & Soker (1989) and by Binney et al. (2009); we discuss their saturation below.

Our results depend *crucially* on the existence of a globally-stabilizing heating mechanism; if heating were not present, the atmospheres shown in figures 3.1 and 3.2 would collapse monolithically. This globally unstable case has been studied extensively by Balbus & Soker (1989). Consistent with their analysis, we find that atmospheres with small initial density inhomogeneities do not form multi-phase gas, regardless of $t_{\rm TI}/t_{\rm ff}$ (see Paper II for a more detailed discussion).

Figures 3.1 and 3.2 show that, assuming the existence of a globally stabilizing heating mechanism, plasmas with short cooling times $t_{\rm TI} \ll t_{\rm ff}$ develop spatially extended multi-phase structure, while plasmas with long cooling times do not. The left panel of figure 3.3 demonstrates this result more quantitatively. Here, we plot the mass fraction of cold gas (with $T \leq 1/3\, T_0$) at late times in the plane $z = H$ as a function of $t_{\rm TI}/t_{\rm ff}$. (Recall that, in our units, $T_0 \sim T_{\rm virial}$.) This figure shows that the fraction of cold gas drops precipitously around $t_{\rm TI}/t_{\rm ff} \sim 1$ and that there is essentially no multi-phase gas at large radii in simulations with $t_{\rm TI}/t_{\rm ff} > 1$.

The right panel of figure 3.3 quantifies the dependence of the saturated density fluctuations on the time-scale ratio $t_{\rm TI}/t_{\rm ff}$ and hints at the physics of the transition between the

two limits. Here, we plot the root-mean-square (RMS) average of the density perturbations

$$\frac{\delta\rho}{\rho} \equiv \frac{\rho - \langle\rho\rangle}{\langle\rho\rangle} \tag{3.20}$$

as a function of $t_{\rm TI}/t_{\rm ff}$ in the plane $z = H$ as in equation 3.3, $\langle\cdots\rangle$ indicates a spatial average at a given radius. In the short cooling time limit, the plasma develops multi-phase structure with large density perturbations $\delta\rho/\rho \gtrsim 1$. By contrast, in the long cooling time limit, the density perturbations saturate at much lower values $\delta\rho/\rho \ll 1$. For plasmas with stable background stratification (e. g. our isothermal initial conditions), $\delta\rho/\rho$ in this limit represents the amplitude of the gravity waves driven by cooling. Note that, while the mass fraction of cold gas drops off sharply around $t_{\rm TI}/t_{\rm ff} \sim 1$, figure 3.3 shows that the mean density fluctuation is a smooth function of this parameter, even in the weak cooling limit. We work to understand this quantitatively in section 3.6.

We emphasize that the difference in the evolution of plasmas with long and short cooling times does not simply result from the thermal instability taking longer to develop in simulations with weak cooling. We have run each simulation for a fixed number of growth times $t_{\rm TI}$ and, if gravity were not present, the results of our simulations with rapid and slow cooling would be nearly identical (we have confirmed this numerically). In fact, even in our simulations with long cooling times, the density contrast $\delta\rho/\rho$ becomes large near the center of the potential, where gravity is weak (eq. 3.8) and the thermal instability has time to develop. The development of multi-phase structure depends on both gravity and cooling and therefore on the parameter $t_{\rm TI}/t_{\rm ff}$, rather than simply on the cooling time alone.

3.5.2 Accreted Mass Flux

Figure 3.4 shows the instantaneous, mean mass flux through the plane $z = H$ as a function of time in three of our 3D, fiducial simulations. The mass fluxes are normalized to the values predicted by cooling-flow models without heating, $\dot{M}_{\rm CF} = \rho H/t_{\rm TI}$. It is clear that the mass flux is strongly suppressed relative to the cooling-flow solution whenever $t_{\rm TI}/t_{\rm ff} \gtrsim 1$. This is not a trivial consequence of our feedback heating mechanism, because for $t_{\rm TI} \lesssim t_{\rm ff}$, $\dot{M}$ approaches $\dot{M}_{\rm CF}$. Rather, the suppression of $\dot{M}$ for $t_{\rm TI} > t_{\rm ff}$ is also due to the non-linear saturation of the thermal instability (described below). The mass fluxes we find for $t_{\rm TI} \gtrsim t_{\rm ff}$ are $\lesssim 1\%$ of the cooling-flow estimates and are therefore reasonably consistent with observational limits for cooling in the ICM (Peterson & Fabian 2006). In Paper II we show that this suppression is even stronger in spherical potentials and we explore its dependence on the details of our heating model.

In the rapid cooling limit, we find that gas heated at small radii, where the cooling time is shorter, rises up through the plane $z = H$ and initially drives an outflow. As the thermal instability progresses, however, this outflow reverses and a strong accretion flow develops (although the accreted material is all in the cold phase, rather than the hot phase; see the left panel of figure 3.3). The accretion rate approaches the cooling-flow value and eventually depletes the atmosphere of its gas. Thus, even our idealized feedback model (eq. 3.4) cannot

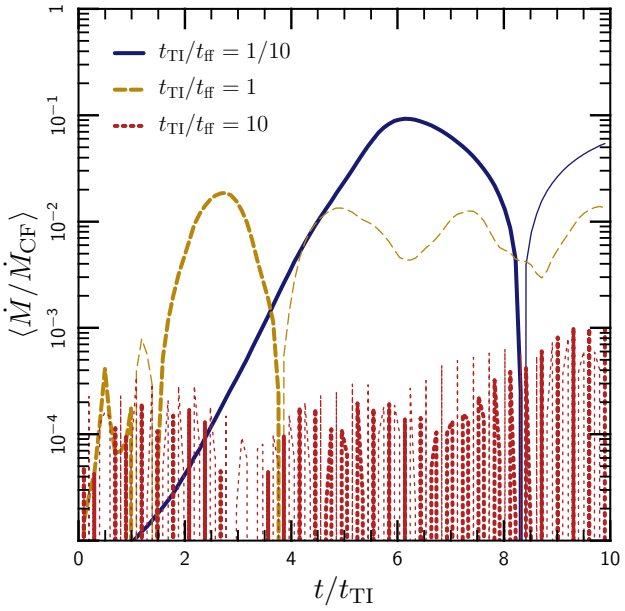

Figure 3.4: Mass flux (averaged from $z = 0.9$–$1.1\,H$) as a function of time in 3D simulations, normalized to the cooling-flow flux for that atmosphere. The mass flux is severely suppressed when $t_{\rm TI} \gtrsim t_{\rm ff}$ (see eq. 3.27). The mass flux is not suppressed as strongly when $t_{\rm TI} < t_{\rm ff}$, but it is highly variable. Thick lines indicate a positive mass flux (i. e., an outflow), while thin lines indicate a negative mass flux (i. e., an inflow). Note that gravity waves dominate the instantaneous mass flux when $t_{\rm TI}/t_{\rm ff} = 10$; the time-averaged accretion rate is much smaller than suggested by this plot.

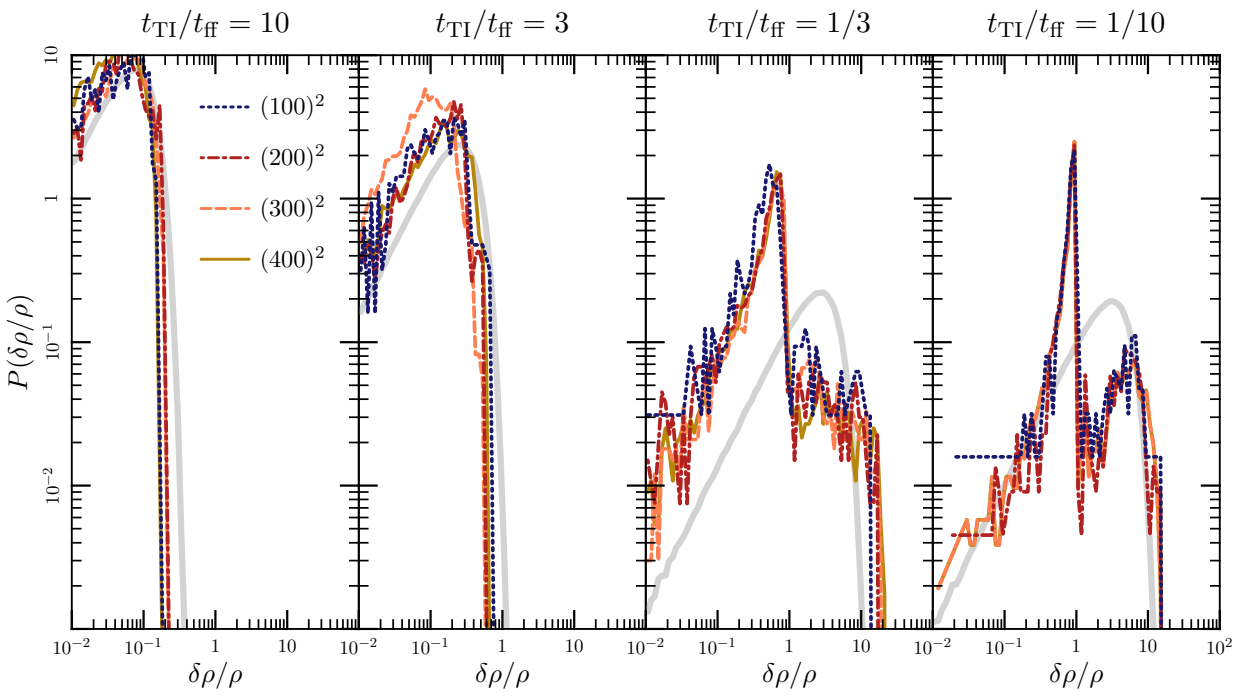

Figure 3.5: Convergence of the probability distribution function for density fluctuations $P(\delta\rho/\rho)$. Colored lines show simulations at different resolutions, and the thick gray line shows the best-fit Gaussian distribution. This figure shows that the density inhomogeneities are reasonably converged in our simulations, apart from the obvious fact that one can resolve finer structure, and therefore higher $\delta\rho/\rho$, at higher resolution. Note that in the limit of rapid cooling, the properties of the high density regions are determined in part by the temperature floor we apply, which determines the density of cold clumps that can be in pressure equilibrium with the surrounding hot plasma.

suppress a cooling catastrophe when $t_{\rm TI}/t_{\rm ff} \ll 1$. We discuss the implications of this result in section 3.8 and, more thoroughly, in Paper II.

Figure 3.4 is only meant to be suggestive, as several subtleties in our analysis complicate a precise interpretation of the accreted mass flux. For instance, we use only the initial value $\dot{M}_{\rm CF}$, though this quantity changes dramatically over the course of some of our simulations. Additionally, in simulations with $t_{\rm TI} \ll t_{\rm ff}$, a more appropriate normalization for the mass flux might be $\rho H/t_{\rm ff}$, since the gas is not likely to flow in faster than its free-fall rate. Paper II presents a much more realistic and thorough analysis of mass accretion rates.

3.5.3 Resolution Study

We test the numerical convergence of our results with 2D calculations on grids of resolution $(100)^2$, $(200)^2$, $(300)^2$, and $(400)^2$ for the full range of $t_{\rm TI}/t_{\rm ff}$. Additionally, we performed 3D calculations of resolution $(128)^3$ and $(256)^3$ for atmospheres with $t_{\rm TI}/t_{\rm ff} = 0.1$, 1, and 10.

Figure 3.5 shows the distribution of density perturbations in our 2D simulations; this quantity has no apparent trend with resolution. Similarly, our 3D simulations are nearly identical at resolutions of $(128)^3$ and $(256)^3$. Though figure 3.5 only demonstrates convergence of an integrated quantity, our simulations also "look" very similar at different resolution: for example, in the rapid-cooling limit, the clumps of cold gas have similar shapes and sizes, and they appear in the same locations.

We find rapid convergence in our simulations, even without including thermal conduction. By contrast, Sharma et al. (2010) found that convergence requires resolving the Field length (§3.7.2) in the cold phase of the ICM. The temperature floor we apply (eq. 3.6) implies that the Field length is not defined for the cold phase in our simulations, and therefore that it is not relevant for convergence. Because the cold phase in our simulations does not cool, it can become pressure-supported at a finite size and resist further collapse. Convergence is somewhat less restrictive in our simulations than in those studied by Sharma et al. (2010).

We performed both 2D and 3D simulations and have confirmed that they give similar results. Many of the plots in this paper show the results of 2D simulations, since they are less expensive and permit a much larger parameter study. Because 2D simulations contain fewer grid cells than 3D simulations, however, integrated quantities derived from 2D calculations are noisier. Thus, we chose to include only 3D simulations in figures 3.4 and 3.8.

While our 2D and 3D simulations produce similar results, they are fundamentally different from one-dimensional simulations. Spatial variations between heating and cooling drive the local thermal instability in our model; hence, the development of multi-phase structure in our simulations is an inherently multi-dimensional effect. Additionally, the symmetry of a one-dimensional model prevents over-dense material from sinking and removes an important saturation channel from the thermal instability (§3.6). Much of the physics we describe in this paper is therefore absent in one-dimensional treatments of the ICM such as those described in Ciotti & Ostriker (2001) and Guo & Oh (2008).

3.5.4 Sensitivity to the Heating Function

An important test of our model is the sensitivity of our conclusions to the details of the (unknown) heating function. We study this dependence by adding random heating fluctuations of the form

$$\mathcal{H} \to \mathcal{H}\,(1+\delta)\,, \tag{3.21}$$

where $\delta(\boldsymbol{x}, t)$ is a Gaussian-random field with a white-noise spatial power spectrum and a temporal autocorrelation function $R_{\delta\delta}(\tau) = e^{-\tau/t_{\rm corr}}$. Thus, δ introduces both spatial and temporal imbalances between heating and cooling, which persist for the coherence time $\sim t_{\rm corr}$. These fluctuations detune our feedback model while still preserving average thermal equilibrium, and are intended to mimic the temporal and spatial differences between heating and cooling which might arise in a more realistic feedback scenario. More importantly, including these fluctuations allows us to distinguish between results that are a consequence

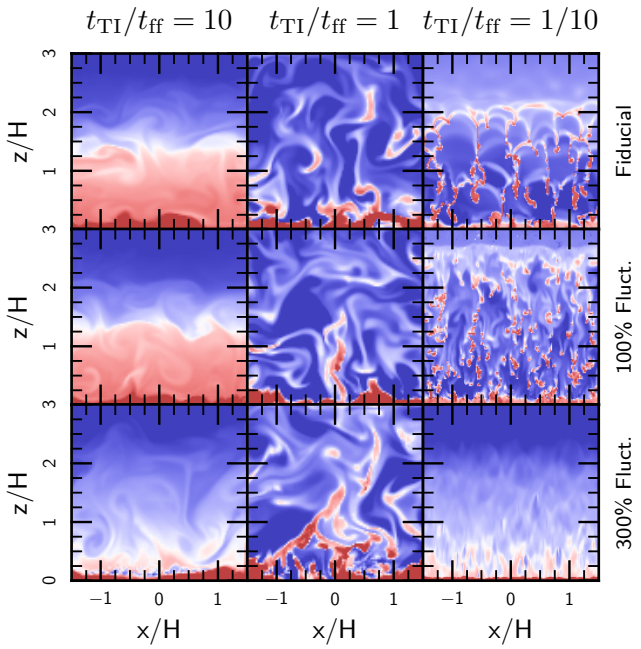

Figure 3.6: Comparison of the plasma density in simulations with our fiducial heating function (equation 3.4) to simulations where we have added significant, random fluctuations to the heating function $\mathcal{H}$ (see eq. 3.21). These are white noise fluctuations with a temporal correlation $t_{\rm corr} = t_{\rm TI}$; simulations with longer correlation times $t_{\rm corr} = 10\,t_{\rm TI}$ give similar results (see figure 3.7). From top to bottom, the panels show the density at $t = 10\,t_{\rm TI}$ in our fiducial simulations, simulations with 100% fluctuations in heating, and simulations with 300% fluctuations in heating. Fluctuations of 300% produce a cooling flow, but 100% fluctuations do not and instead produce results similar to our fiducial model. In all panels, color represents the log of the density, which ranges from 10^{-2} (blue) to 10 (red).

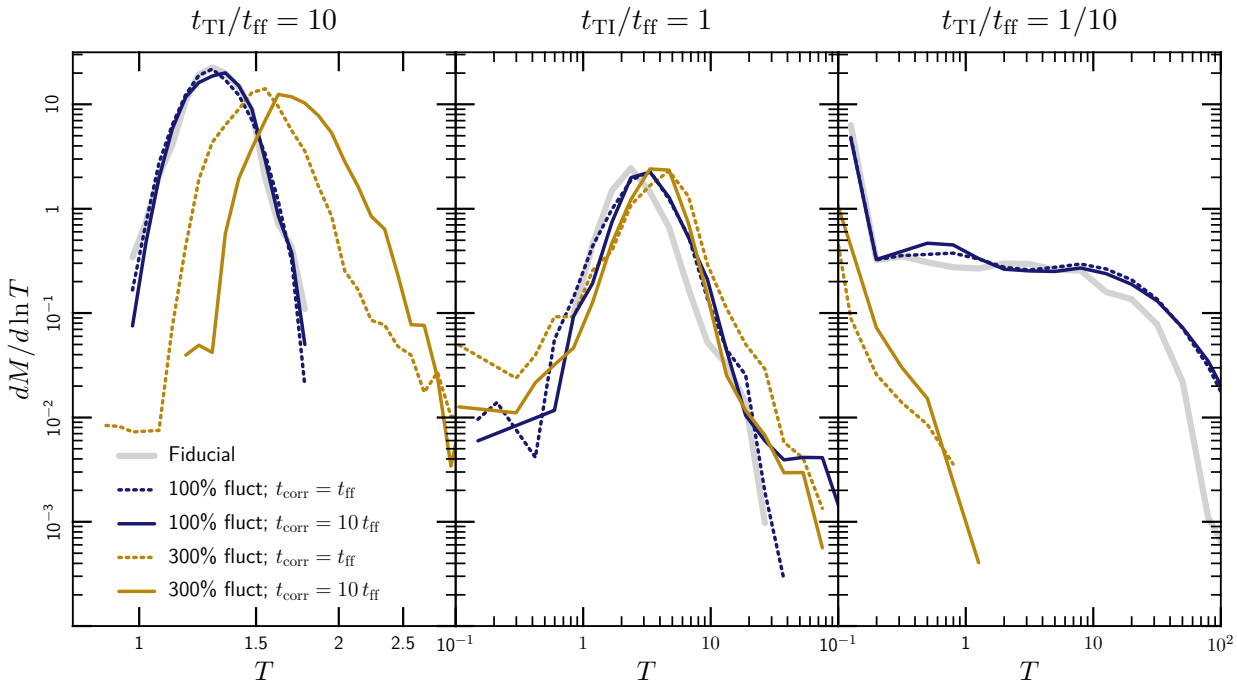

Figure 3.7: Gas mass as a function of temperature in simulations with different types of fluctuations about thermal equilibrium, measured at $t = 10t_{\rm TI}$ and $z \sim H$. These simulations are for isothermal initial conditions in which the initial temperature $T = 1$. The fluctuations are of the form $\mathcal{H} \to \mathcal{H}(1+\delta)$, where $\delta(\boldsymbol{x}, t)$ has a white-noise spatial spectrum and a temporal coherence time $t_{\rm corr}$ (see §3.5.4 for details). The gas properties are not sensitive to strong fluctuations in heating of up to 100% in amplitude. Stronger fluctuations of 300% generate significant cold material when $t_{\rm TI}/t_{\rm ff} \gg 1$; the cold material sinks to small radii (see figure 3.6), leading to a modest heating of the gas that remains at $z \sim H$. When $t_{\rm TI}/t_{\rm ff} \ll 1$, fluctuations of 300% break our ansatz of approximate thermal equilibrium and induce a cooling catastrophe. The correlation time $t_{\rm corr}$ has only a modest influence on these results.

of the exact (and, in detail, unphysical) balance in equation 3.4 and results that are more robust and are primarily a consequence of global thermal stability.

We carried out 2D simulations with $t_{\rm corr} = \{0.1,\, 1,\, 10\} \times t_{\rm ff}$ and with the fluctuations normalized to root-mean-square (RMS) amplitudes of 50%, 100% and 300%. (Note that we quote the RMS, or '1σ' amplitude of the fluctuations; the peak values are considerably higher.) figure 3.6 shows images of the density fluctuations for these simulations and figure 3.7 shows the temperature distribution function for different values of $t_{\rm corr}$ and the fluctuation amplitude. These figures demonstrate that our conclusions about the development of the thermal instability are essentially unaffected by order-unity fluctuations, over at least 10 cooling times. This important result implies that, as long as the plasma is in approximate global thermal equilibrium on reasonable time-scales $\sim t_{\rm TI}$ and length-scales $\sim H$, the development and saturation of local thermal instability will proceed approximately as shown in figures 3.1–3.4. We think that the existence of an approximate thermal equilibrium, rather than the specific details of our heating function (eq. 3.4), determines how the thermal instability develops and saturates. This conclusion is bolstered by Paper II, which finds very similar results using an entirely different heating function.

Figs 3.6 and 3.7 show that extremely strong heating fluctuations with RMS amplitudes of 300% spoil the thermal equilibrium of the plasma and induce a cooling catastrophe; even our extremely optimistic feedback model cannot withstand arbitrarily large heating perturbations. Though the feedback mechanism is not yet understood in clusters, this places a constraint on the heating: it should not differ persistently from the local cooling rate by more than a factor of several. Figure 3.7 shows that this conclusion is essentially independent of the coherence time $t_{\rm corr}$ of the heating.

3.6 Interpretation of the Non-Linear Saturation

In this section, we show that the linearized dynamical equations provide valuable insight into the non-linear saturation of the thermal instability and its astrophysical implications. As in section 3.5, we focus on the development of multi-phase structure (§3.6.2) and on the accreted mass flux (§3.6.3), which have been extensively studied observationally. Our basic procedure is to estimate a saturation amplitude for the linear instability. Because we use linearized equations, the interpretation in this section only strictly holds in the weak cooling limit ($t_{\rm TI} \gg t_{\rm ff}$), so that the density perturbations remain relatively small.

Figure 3.8 illustrates the development and saturation of the thermal instability. The density inhomogeneity $\delta\rho/\rho$ (or any other quantity linear in the perturbation) initially grows exponentially according to the dispersion relation (eq. 3.16), but eventually freezes out at a finite amplitude. This amplitude, along with the relations 3.12a–3.12c, then approximately determines the state of the plasma at late times. In the following sections, we estimate this amplitude and show that we can reproduce elements of the non-linear saturation shown in figure 3.3. Figure 3.8 shows that the difference between atmospheres which develop multi-phase gas and ones which do not is *fundamentally a non-linear effect*. The linear growth

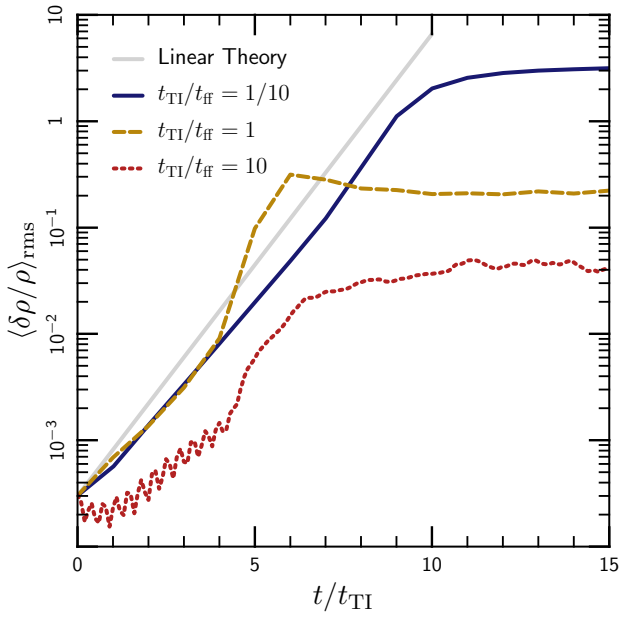

Figure 3.8: Evolution of the density fluctuation $\delta\rho/\rho$ as a function of time in our simulations with isothermal initial conditions. The plotted quantity is an RMS average from $z = 0.9H$–$1.1H$. The density inhomogeneity grows from the initial perturbation until the characteristic infall time becomes comparable to the local cooling time. At this point, the density contrast saturates at approximately the value given by equation 3.25b. This figure shows the results from 3D simulations, but the results from 2D simulations are similar.

rate of the perturbations is largely independent of the time-scale ratio $t_{\rm TI}/t_{\rm ff}$, and it is the saturation which determines the degree of inhomogeneity at late times.

3.6.1 Saturation Amplitudes

In the limit that the plasma is buoyantly neutral ($N = 0$), we can estimate the saturation amplitude by inspecting the linearized, Lagrangian form of the momentum equation (eq. 3.1b):

$$\frac{dv_z}{dt} = -\frac{\delta n}{n} g. \tag{3.22}$$

The characteristic inflow (or outflow) time for a perturbed fluid element is $t_{\rm sink} \sim H/\delta v_z$. Initially, δv_z is small and this inflow time is long compared to the growth time of the thermal instability. As the perturbation grows, however, $\delta n/n$ increases and the fluid element accelerates according to equation 3.22. The inflow time $t_{\rm sink}$ thus becomes shorter as the instability develops. We assume that the growth ceases when the inflow time is comparable to, or slightly shorter than, the growth time of the thermal instability. In this case, the thermal instability saturates when the velocity satisfies

$$\delta v_z \sim \frac{H}{t_{\rm TI}}. \tag{3.23}$$

Thus, non-linear saturation occurs when a fluid element flows to smaller radii after one cooling time, as seems intuitively reasonable.

The physical picture of a sinking fluid element does not apply in a stably stratified atmosphere, since the fluid element does not flow monotonically inwards, but instead oscillates with the gravity wave frequency N. The velocity associated with this oscillation dwarfs the mean, inward velocity. These waves are sourced by cooling, however, and we assume that they reach a steady-state in which the dissipation rate due to non-linear mode coupling equals the driving rate due to the thermal instability $\sim t_{\rm TI}^{-1}$. Thus the instability saturates when the dissipation time $t_{\rm diss} \sim H/\delta v \sim t_{\rm TI}$, where we have assumed strong turbulence and used the fact that the waves are driven on large scales, $\sim H$ (as suggested by the bottom panels in figure 3.1). Though this saturation mechanism is very different from that described above for buoyantly neutral plasmas, it implies an equivalent saturation amplitude. We therefore assume that equation 3.23 describes the late-time evolution of the perturbations in all of our simulations. We show in the following sections how the behavior described in section 3.5 can be understood in terms of this saturation amplitude.

3.6.2 Multi-Phase Structure

Inserting our ansatz for the saturation amplitudes (equation 3.23) into the momentum equation (3.12b) and using the dispersion relation (equation 3.16) to replace ω, we express

the density inhomogeneity $\delta n/n$ at late times in terms of other properties of the plasma:

$$\frac{\delta n}{n} \propto \left(\frac{t_{\rm ff}}{t_{\rm TI}}\right)^2 \left[1 \pm \sqrt{1 - 4\hat{\boldsymbol{k}}_\perp^2 \left(\frac{t_{\rm TI}}{t_{\rm buoy}}\right)^2}\right] \tag{3.24}$$

where $t_{\rm buoy} \equiv N^{-1}$ and $\hat{\boldsymbol{k}}_\perp^2 \equiv (1 - \hat{k}_z^2)$ is the squared horizontal component of the direction of the wave vector (typically ~ 1).

Equation 3.24 has the asymptotic forms

$$\frac{\delta n}{n} \propto \left(\frac{t_{\rm ff}}{t_{\rm TI}}\right)^2 \qquad t_{\rm TI} \ll t_{\rm buoy} \tag{3.25a}$$

$$\frac{\delta n}{n} \propto \left(\frac{t_{\rm ff}}{t_{\rm TI}}\right)\left(\frac{t_{\rm ff}}{t_{\rm buoy}}\right) \qquad t_{\rm TI} \gg t_{\rm buoy}. \tag{3.25b}$$

Equation 3.25 shows that weakly stratified plasmas with $t_{\rm TI} \ll t_{\rm buoy}$ develop smaller density inhomogeneities than plasmas with $t_{\rm TI} \gg t_{\rm buoy}$ (in the limit that $t_{\rm TI} \gtrsim t_{\rm ff}$). This difference arises because because plasmas with $t_{\rm TI} \gg t_{\rm buoy}$ can sustain internal gravity waves, while atmospheres with $t_{\rm TI} \ll t_{\rm buoy}$ cannot.

Somewhat surprisingly, the right panel of figure 3.3 shows that the measured dependence of $\delta n/n$ on the time-scale ratio $t_{\rm TI}/t_{\rm ff}$ is in good agreement with equation 3.25b for both isothermal and isentropic initial conditions, even though $t_{\rm buoy} \to \infty$ in an isentropic atmosphere. This is because equation 3.25a applies only when the atmosphere is very nearly buoyantly neutral (at least when $t_{\rm TI}$ is long compared to the dynamical time, as must be for equation 3.25 to be valid). While this is the case initially in our isentropic atmospheres, the entropy gradient evolves somewhat with time and equation 3.25a ceases to describe the plasma after only a few cooling times. The saturated values of $t_{\rm buoy}$ differ in our simulations with isentropic and isothermal initial conditions; equation 3.25b suggests that this may explain the systematic offset between these simulations shown in figure 3.3. For the longest cooling times, the evolution of the background profile is smallest; the slight steepening of $\delta\rho/\rho$ for our isentropic simulations in this limit may represent an intermediate case between equations 3.25a and 3.25b.

3.6.3 Accreted Mass Flux

We can also use our estimate of the non-linear saturation to understand the inward mass flux induced by the thermal instability. Defining the mass flux $\dot{M} = \langle \delta n \, \delta v_z \rangle$ and using the estimates of δv and $\delta n/n$ from the momentum equation and from equation 3.24, we find

$$\begin{aligned}\dot{M} =& \dot{M}_{\rm CF} \left(\frac{t_{\rm ff}}{t_{\rm TI}}\right)^2 \frac{1}{4\hat{\boldsymbol{k}}_\perp^2} \\ &\times \left\langle {\rm Re}\left(\left[1 + \sqrt{1 - 4\hat{\boldsymbol{k}}_\perp^2 \left(\frac{t_{\rm TI}}{t_{\rm buoy}}\right)^2}\right] e^{i\boldsymbol{k}\cdot\boldsymbol{x}}\right) {\rm Re}\left(e^{i\boldsymbol{k}\cdot\boldsymbol{x}}\right) \right\rangle,\end{aligned} \tag{3.26}$$

where $\dot{M}_{\rm CF} = \rho H / t_{\rm TI}$ is the mass flux expected in the absence of heating (recall that we are in the limit that $t_{\rm TI} \gg t_{\rm buoy}$). This yields

$$\frac{\dot{M}}{\dot{M}_{\rm CF}} \simeq \begin{cases} \frac{1}{4}\left(\frac{t_{\rm ff}}{t_{\rm TI}}\right)^2 - \frac{1}{4}\left(\frac{t_{\rm ff}}{t_{\rm buoy}}\right)^2 & t_{\rm TI} \ll t_{\rm buoy} \\ \frac{1}{8}\left(\frac{t_{\rm ff}}{t_{\rm TI}}\right)^2 & t_{\rm TI} \gg t_{\rm buoy} \end{cases}. \tag{3.27}$$

Equation 3.27 shows that the mass flux is dramatically suppressed when $t_{\rm TI}/t_{\rm ff} \gg 1$, qualitatively consistent with figure 3.4. Furthermore, this suppression is nearly independent of the initial stratification of the plasma, even though stably stratified plasmas show much stronger density inhomogeneities (eq. 3.25). This is because the internal gravity waves that enhance $\delta n/n$ when $t_{\rm TI}/t_{\rm buoy} \gg 1$ do not contribute to the net mass flux. Note that gravity waves dominate the instantaneous mass flux shown in figure 3.4 when $t_{\rm TI}/t_{\rm ff} \gg 1$; the time-averaged accretion rate is smaller than the figure suggests.

3.7 Simulations Including Conduction

The previous sections describe a simplified model of the thermal instability that neglects both conduction and the dynamical effect of the magnetic field. This model nicely isolates the physics of the thermal instability, but astrophysically it is too idealized. For example, conduction is critical for the thermal evolution of the plasma over a wide range of scales in the ICM, and Balbus (2000) and Quataert (2008) have shown that this completely changes the stability and dynamics of the plasma. In this section, we present results including magnetic fields and conduction and show that the conclusions from the previous sections largely apply in this more realistic case.

3.7.1 Setup

Our setup is very similar to that described in sections 3.2 and 3.3. We generalize equations 3.1b and 3.1c to include the effects of thermal conduction and the magnetic field:

$$\frac{\partial}{\partial t}(\rho \boldsymbol{v}) + \nabla \cdot \left[\rho \boldsymbol{v} \otimes \boldsymbol{v} + \left(P + \frac{B^2}{8\pi}\right)\mathbf{I} + \frac{\boldsymbol{B} \otimes \boldsymbol{B}}{4\pi}\right] = \rho \boldsymbol{g}, \tag{3.1b$'$}$$

$$\rho T \frac{ds}{dt} = (\mathcal{H} - \mathcal{L}) - \nabla \cdot \boldsymbol{Q}_{\rm cond}, \tag{3.1c$'$}$$

where $\boldsymbol{B}$ is the magnetic field and $\boldsymbol{Q}_{\rm cond}$ is the conductive heat flux. We evolve the magnetic field using the induction equation:

$$\frac{\partial \boldsymbol{B}}{\partial t} = \nabla \times (\boldsymbol{v} \times \boldsymbol{B}). \tag{1d$'$}$$

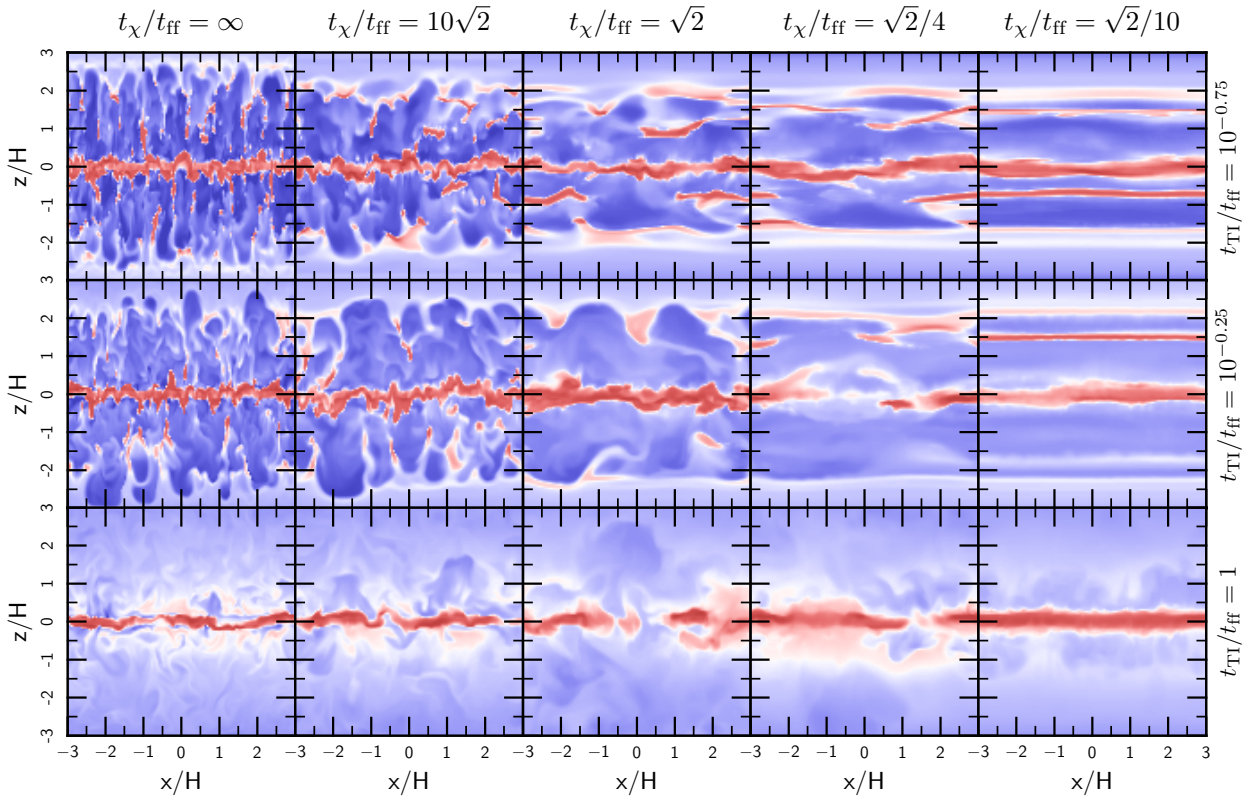

Figure 3.9: Comparison of the gas density in the non-linear state of simulations with different time-scale ratios $t_{\rm TI}/t_{\rm ff}$ and $t_\chi/t_{\rm ff}$, including an initially horizontal magnetic field and anisotropic thermal conduction. We take the conduction time to be the time it takes heat to diffuse across one scale-height: $t_\chi \equiv H^2/\chi$. Rapid conduction dramatically changes the morphology of the cold gas, smearing it out in the direction of the magnetic field. However, conduction does not appreciably change the mass of gas in the cold phase (figure 3.10). The ratio of timescales $t_{\rm TI}/t_{\rm ff}$ still determines whether or not the plasma develops multi-phase structure, even in the limit of rapid conduction. In all frames, the color scale represents the log of the density; blue corresponds to a density of 10^{-2} and red corresponds to a density of 10. The filaments are very straight in simulations where the Field length is longer than the domain size (right-most column) because conduction effectively eliminates all horizontal structure in the initial perturbations; the subsequent evolution is therefore nearly two-dimensional. We show both the top and bottom of the computational domain here to emphasize that the number of filaments/blobs produced by the thermal instability is somewhat stochastic.

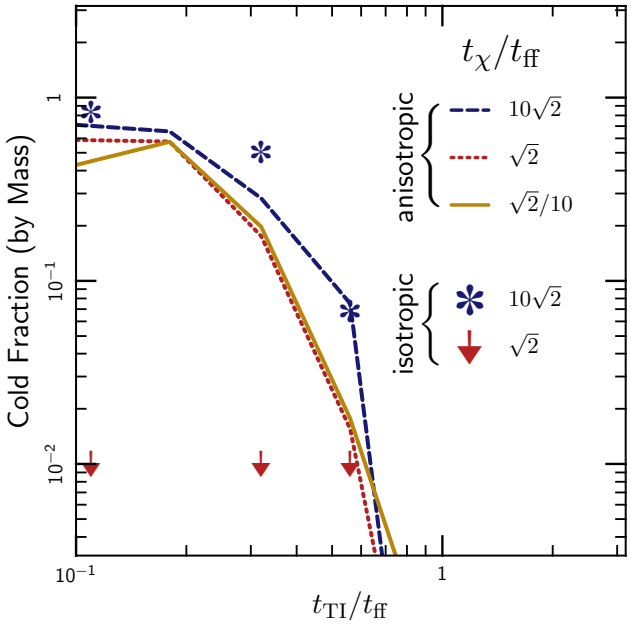

Figure 3.10: Mass fraction of cold material (with $T \leq T_0/3$) as a function of the timescale ratio $t_{\rm TI}/t_{\rm ff}$, for simulations with different conductivities. As in figure 3.3, this mass fraction is determined by averaging from $z = 0.9$–$1.1\,H$ and from $t = 9$–$10\,t_{\rm TI}$. Lines show simulations with anisotropic thermal conduction and points indicate simulations with isotropic conduction. As suggested by figure3.9, anisotropic thermal conduction does not strongly influence the multi-phase structure produced by the thermal instability: though the conductivities in these simulations differ by a factor of 100, the cold mass fractions agree to within about a factor of 2. In particular, the results appear to converge in the rapid-conduction limit. By contrast, we see no multi-phase structure in simulations with isotropic conduction if the Field length is comparable to, or greater than a scale-height (the arrows in this figure indicate upper limits on the mass fractions of cold gas).

We have ignored both (explicit) viscous and magnetic dissipation in equations 3.1b′–1d′. These effects can influence MHD simulations in subtle and unexpected ways (see, e. g. Fromang & Papaloizou 2007; Davis et al. 2010), and so will need to be studied in detail in the future.

The thermal conductivity of the plasma is strongly anisotropic in the ICM and as a result the conductive heat flux is given by (Braginskii 1965)

$$\boldsymbol{Q}_{\rm cond} = -nk_{\rm B}\chi_{\rm e}\,\hat{\boldsymbol{b}}\,(\hat{\boldsymbol{b}}\cdot\nabla T), \tag{3.28}$$

where $\hat{\boldsymbol{b}} = \boldsymbol{B}/B$ is a unit vector in the direction of the magnetic field and $\chi_{\rm e}$ is the thermal diffusivity of free electrons (with units of cm^2/s). While the diffusivity $\chi_{\rm e}$ depends sensitively on temperature (Spitzer 1962), we take it to be constant in this exploratory analysis. This enables us to control the ratio of the conduction time to other timescales in the problem and thus to isolate the physics of cooling and conduction. Note that we still use the heating function defined by equation 3.4; any conductive heating or cooling of the plasma happens on top of the feedback heating.

We initialize the plasma with a weak, horizontal magnetic field. (By 'weak,' we mean that magnetic tension is negligible in our simulations.) Because we impose reflecting boundary conditions at the upper and lower boundaries of the domain (§3.3), the magnetic field remains horizontal there and prohibits a conductive heat flux into the domain.

As before, we solve equations 3.1a and 3.1b′–1d′ using ATHENA with the modifications described in section 3.3. We also implement equation 3.28 via operator splitting, using the anisotropic conduction algorithm described in Parrish & Stone (2005) and Sharma & Hammett (2007). In particular, we use the monotonized central difference limiter on transverse heat fluxes to ensure stability. This conduction algorithm is sub-cycled with respect to the main integrator with a time step $\Delta t \propto (\Delta x)^2$; these simulations are therefore more computationally expensive than adiabatic MHD calculations, especially at high resolution.

3.7.2 Linear Properties

We linearize equations 3.1a and 3.1b′–1d′ and perform a WKB analysis (see Quataert 2008 for more details). Assuming that magnetic tension is negligible, and proceeding as in section 3.4.1, the dispersion relation for the plasma is (cf. Balbus & Reynolds 2010)

$$p^3 - p^2 p_{\rm F} + p\,N^2\hat{\boldsymbol{k}}_\perp^2 - \omega_\chi\,p_{\rm HBI}^2 = 0. \tag{3.29}$$

In the above, $p = -i\omega$ is the growth rate of the perturbation, $p_{\rm F} = p_{\rm TI} - \omega_\chi$ is the growth rate of the thermal instability accounting for conduction (Field 1965),

$$\omega_\chi = \frac{\gamma-1}{\gamma}\chi_{\rm e}\left(\hat{\boldsymbol{b}}\cdot\boldsymbol{k}\right)^2 \tag{3.30}$$

is inversely proportional to the conduction time across the wavelength of the perturbation, and

$$p_{\rm HBI}^2 = g\frac{\partial \ln T}{\partial z} \times \left[(2\hat{b}_z^2 - 1)(1 - \hat{k}_z^2) - 2\hat{b}_x\hat{b}_z\hat{k}_x\hat{k}_z\right] \tag{3.31}$$

is the growth rate of either the magnetothermal instability (MTI; Balbus 2000), or the heat-flux driven buoyancy instability (HBI; Quataert 2008).

The MTI is unlikely to influence the development of multi-phase structure in galaxy clusters, since it operates outside the cool core, where the ratio of timescales $t_{\rm TI}/t_{\rm ff}$ is typically much greater than unity. While the HBI does operate efficiently in cool cores, it behaves like ordinary stable stratification in its saturated state (McCourt et al. 2011b) and the growth time $t_{\rm HBI}$ is analogous to the timescale $t_{\rm buoy}$ used earlier. Thus, we do not expect the HBI or MTI to change our results in any essential way (although this must be studied more carefully in future work). We anticipate that the same will be true for the overstabilities associated with the MTI and HBI (Balbus & Reynolds 2010). In this section, we use simulations with isothermal initial conditions (in which $p_{\rm HBI} \to 0$) so that these instabilities and overstabilities do not operate (at least in our initial conditions). This allows us to focus on the physics of thermal instability.

The conduction frequency ω_χ is a function of scale, while the growth rate of the thermal instability $p_{\rm TI}$ is not; the modified growth rate $p_{\rm F}$ therefore must switch sign at the length-scale

$$\lambda_{\rm F} = |\hat{\boldsymbol{b}}\cdot\hat{\boldsymbol{k}}| \times \left[(2\pi)^2\frac{\gamma-1}{\gamma}\frac{\chi_{\rm e}}{p_{\rm TI}}\right]^{1/2}, \tag{3.32}$$

known as the Field length (Field 1965). Intuitively, the Field length is the distance heat can diffuse in one cooling time; if the wavelength of a perturbation is larger than this distance, conduction cannot stabilize it against cooling and the perturbation grows exponentially.

Conduction suppresses the thermal instability on scales smaller than the Field length, but the Field length in a magnetized medium depends on direction, as well as position. Even if the term in square brackets in equation 3.32 becomes arbitrarily large, the Field length will be small in directions orthogonal to the magnetic field. Because of this anisotropy, the thermal instability can still grow on scales much smaller than $\sqrt{\chi_{\rm e} t_{\rm TI}}$. Sharma et al. (2010) have studied this growth in the absence of gravity; here we generalize their results to stratified plasmas.

The growth rate for the thermal instability in our simulations is

$$p = \frac{1}{2}\left[p_{\rm F} \pm \sqrt{p_{\rm F}^2 - 4N^2\hat{\boldsymbol{k}}_\perp^2}\right] \tag{3.33a}$$

$$= \begin{cases} p_{\rm F} - N^2\hat{\boldsymbol{k}}_\perp^2/p_{\rm F} & p_{\rm F} \gg |N| \\ \frac{1}{2}p_{\rm F} \pm iN\hat{\boldsymbol{k}}_\perp & p_{\rm F} \ll |N| \end{cases}. \tag{3.33b}$$

Equation 3.33 shows that the characteristic growth time of the thermal instability is $p_{\rm F}^{-1}$, regardless of the entropy gradient. The growth rate $p_{\rm F}$ reduces to $p_{\rm TI}$ on large scales; thus, our results with and without conduction are very similar on scales larger than the Field length. Conduction prevents perturbations from growing below the Field length and therefore plays a similar role to the temperature floor in our non-conducting simulations. The primary difference between our conducting and non-conducting simulations is that the Field length is anisotropic in the conducting simulations, and the thermally unstable fluid elements collapse into long filaments, rather than the approximately spherical clumps shown in figures 3.1 and 3.2.

3.7.3 Numerical Results

Figure 3.9 shows 2D slices of 3D simulations with different values of the cooling constant Λ_0 and the conductivity $\chi_{\rm e}$. We use only 3D simulations in this section because, just as an over-dense fluid element cannot sink in one dimension, the dynamics of a sinking magnetized filament changes in going from two to three dimensions. These simulations all use our isothermal initial condition and initially have weak, horizontal magnetic field lines in the plane of the figure. Rapid conduction smears out the cold clumps into filaments of length $\sim \lambda_{\rm F}$, but does not otherwise alter the growth of the thermal instability. Specifically, figure 3.9 demonstrates that, even in the limit of very rapid conduction, the ratio of time-scales $t_{\rm TI}/t_{\rm ff}$ determines whether the plasma develops multi-phase structure. This result depends *critically* on the anisotropic nature of thermal conduction. In the rightmost panels of figure 3.9, the Field length is larger than the entire simulation domain; if conduction were isotropic, the entire atmosphere would become nearly isothermal and the thermal instability would be suppressed. The insulating effect of the magnetic field permits large temperature gradients orthogonal to the magnetic field and thus the formation of multi-phase structure (Sharma et al. 2010).

Figure 3.10 quantifies the effect of conduction on the thermal instability: we show the mass fraction of cold gas (as in the left panel of figure 3.3) for 3D simulations with different thermal conductivities. In simulations with anisotropic thermal conduction, this mass fraction is almost independent of the conductivity, and it appears to converge in the limit that the conductivity becomes large. This behavior is consistent with figure 3.9. Together, these results imply that anisotropic conduction alters the morphology of the gas in the cold phase, but not the presence, absence, or amount of multi-phase structure.

We have also run a number of simulations with isotropic thermal conduction. These simulations use the same setup as before, but with the conductive heat flux $\boldsymbol{Q}_{\rm cond} = -nk_{\rm B}\chi_{\rm e}\nabla T$, where (as before) $\chi_{\rm e}$ is a constant, free parameter. In order to prevent conduction from changing the total energy content of the plasma, we set $\chi_{\rm e} = 0$ at the upper and lower boundaries of the computational domain so that there is no conductive heat flux into the domain. Figure 3.10 shows that, while anisotropic conduction does not strongly influence the amount of cold gas produced by the thermal instability, isotropic conduction can quench it entirely: we see no multi-phase structure in our simulations with isotropic conduction

Table 3.2: Clusters used in figure 3.11.

Extended H_α	No Extended H_α
Abell 133	Abell 85
Abell 478	Abell 644
Abell 496	Abell 744
Abell 780	Abell 1650
Abell 1795	Abell 2029
Abell 1991	Abell 2142
Abell 2597	Abell 4059
Sersic 159-03	
Centaurus	

We use the surveys of McDonald et al. (2010, 2011a,b) to determine whether a cluster shows multi-phase gas, and we use the data in the ACCEPT catalog to estimate $t_{\rm TI}/t_{\rm ff}$ for the hot ICM. Our label 'extended H_α' signifies that the H_α emission can be resolved and is known to exist outside the BCG; these are the Type I systems from McDonald et al.

whenever the Field length is comparable to, or larger than, the pressure scale-height. These conclusions also apply to other properties of the plasma quantified in section 3.5, e. g. the accreted mass flux: anisotropic thermal conduction has little effect on this quantity, while isotropic thermal conduction can strongly suppress it.

Voit et al. (2008) suggested that thermal instability produces multi-phase gas in clusters when the Field length is comparable to, or smaller than, the size of the cool core, but that conduction suppresses the formation of multi-phase structure for larger Field lengths. Coincidentally, in typical cool-core clusters, this criterion is quantitatively similar to our criterion on the ratio $t_{\rm TI}/t_{\rm ff}$.[4] However, because the ICM is magnetized, thermal conduction is extremely anisotropic; the results of this section demonstrate that even very rapid thermal conduction cannot suppress local thermal instability. Thermal conduction only stabilizes modes parallel to the magnetic field, and multi-phase structure continues to develop via perturbations that are roughly orthogonal to the local magnetic field.

3.8 Discussion

Observational limits from x-ray spectroscopy (Peterson & Fabian 2006) and from the shape of the galaxy luminosity function (Benson et al. 2003) indicate that the diffuse plasma in galaxy groups and clusters does not cool as quickly as it radiates. These observations imply that some heating process offsets radiative cooling and that the gas remains in approximate

[4]This comparison makes use of the result from Paper II that the threshold for multi-phase gas in spherical systems is closer to $t_{\rm TI}/t_{\rm ff} \sim 10$.

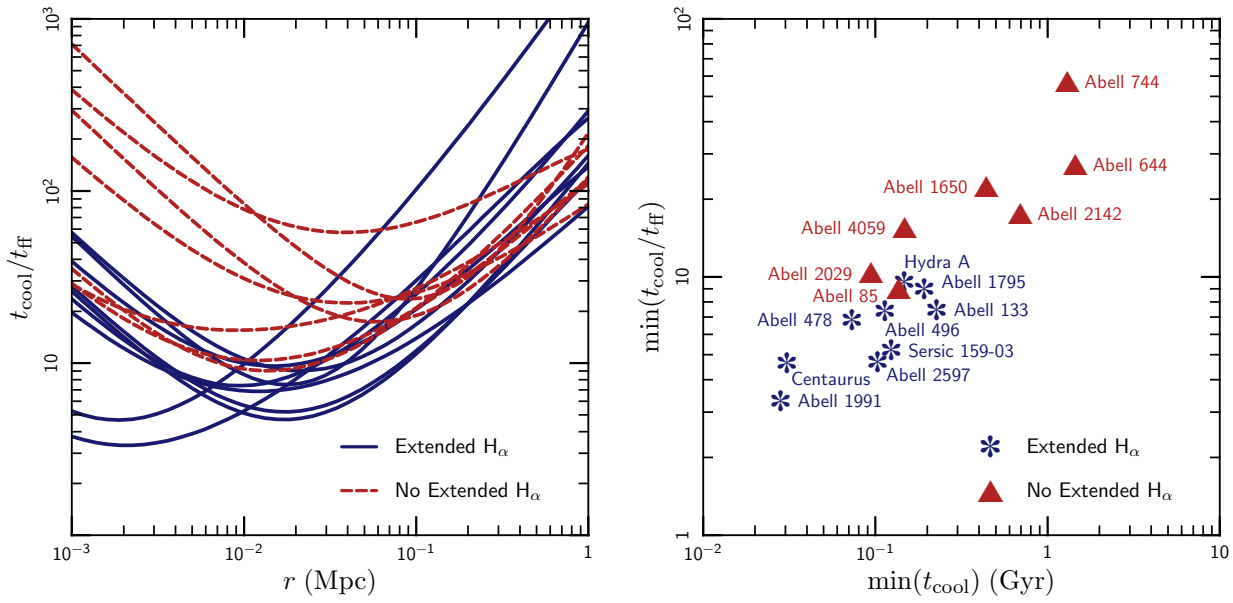

Figure 3.11: (*Left:*) The timescale ratio $t_{\rm cool}/t_{\rm ff}$ as a function of radius for clusters in both the ACCEPT catalog and the McDonald et al. (2010) survey. Solid blue lines show clusters with filaments and dashed red lines show clusters that lack detected extended H_α emission. Clusters with filaments have systematically lower values of $t_{\rm cool}/t_{\rm ff}$. Furthermore, this ratio is smallest between ∼10–50 kpc, where most filaments are found. (*Right:*) The same clusters in the $t_{\rm cool}$–$t_{\rm cool}/t_{\rm ff}$ plane. The coloring is the same as in the left panel. The ratio $t_{\rm cool}/t_{\rm ff}$ appears to be a slightly better predictor of multi-phase structure than $t_{\rm cool}$ alone. Table 3.2 lists the clusters plotted in this figure.

thermal equilibrium, at least when averaged over length-scales comparable to the scale-height or time-scales comparable to the cooling time. The nature of this "feedback" heating is not yet fully understood, although it appears to involve heating by a central AGN (e. g. Bîrzan et al. 2004; McNamara & Nulsen 2007). More fully understanding the mechanism(s) that regulate the heating to so closely match cooling remains a major challenge in theories of galaxy formation.

Though heating strongly suppresses cooling in galaxy groups and clusters, star formation and multi-phase gas provide clear evidence for cold gas in many cluster cores. Observational indicators of this cold gas strongly correlate with the cooling time of the ambient hot ICM (e. g. Voit et al. 2008; Rafferty et al. 2008; Cavagnolo et al. 2008, 2009), motivating a model in which thermal instability in the hot ICM produces much of the cold gas in cluster cores (as has been suggested many times in the past, e. g. Fabian & Nulsen 1977; Cowie et al. 1980; Nulsen 1986; Loewenstein et al. 1991). Theoretically studying local thermal instability in the ICM has proven difficult, however, because of the cooling-flow problem: studies that include both cooling and gravity typically find that the plasma is *globally* thermally unstable, and that the entire cluster core collapses monolithically. This difficulty has led several authors to conclude that the thermal instability does not produce multi-phase structure in stably-stratified systems at all.

We avoid the cooling-flow problem in this paper by adopting a phenomenological heating model that enforces thermal equilibrium when averaged over large scales (§3.2.1). Our heating model is approximate, over-simplified, and wrong in detail. However, our results are insensitive to large temporal and spatial fluctuations about the average heating (§3.5.4). Moreover, in Paper II we obtain similar results using a more physical feedback heating prescription. We therefore believe that our conclusions about the saturation of the local thermal instability are reasonably robust.

In the current paradigm in which clusters are approximately in global thermal equilibrium, heating of the ICM is very likely to depend explicitly on position in the cluster. In this case, the thermal stability of the plasma is independent of its convective stability (§3.4.2). Fundamentally, buoyancy drives convection, while heating and cooling drive thermal instability. These two processes are formally related only under the restrictive assumption that heating is a state function of the plasma; more generally, they are unrelated and the thermal stability of a plasma is independent of its convective stability (see Balbus & Soker 1989).

Nonetheless, it remains true that the competition between buoyancy and thermal instability determines the net effect of cooling on a stratified plasma. We parametrize the relative importance of these effects using the dimensionless ratio $t_{\rm TI}/t_{\rm ff}$ (the ratio of the thermal instability growth time $t_{\rm TI}$ to the local dynamical, or free-fall, time $t_{\rm ff}$). When this ratio is small, thermal instability dominates and the plasma develops significant multi-phase structure; when the ratio $t_{\rm TI}/t_{\rm ff}$ is large, buoyancy dominates and the plasma remains in a single, hot phase (§3.5). This dependence of the saturation on $t_{\rm TI}/t_{\rm ff}$ is true for *both* stably stratified and neutrally stratified plasmas (figure 3.3), even though the effect of buoyancy is very different in these two cases.

More quantitatively, we find (figure 3.3) that the saturated density inhomogeneities pro-

duced by the thermal instability approximately obey the relation

$$\frac{\delta\rho}{\rho} \sim \left(\frac{t_{\rm TI}}{t_{\rm ff}}\right)^{-1}; \tag{3.34}$$

this scaling can be understood analytically by assuming a saturation amplitude for the thermal instability in which the characteristic fluid velocities approach $v_{\rm sat} \sim H/t_{\rm cool}$ (§3.6). Thus, by assuming that some heating mechanism prevents cooling catastrophes in clusters, we find that the ICM breaks up into multiple phases via local thermal instability (with $\delta\rho/\rho \gtrsim 1$) only if the dimensionless ratio of timescales $t_{\rm TI}/t_{\rm ff} \lesssim 1$; specifically, there is almost no cold gas at large radii when $t_{\rm TI} \gtrsim t_{\rm ff}$ (figure 3.3). This finding is one of the primary results of our analysis. We note that the linear growth of the thermal instability is largely independent of the ratio $t_{\rm TI}/t_{\rm ff}$. Thus the difference between atmospheres which develop multi-phase gas and those which do not is fundamentally due to how the *non-linear* saturation of the thermal instability depends on the atmosphere's properties.

The calculations in Paper II show that the criterion for multi-phase structure is actually somewhat less stringent in spherical systems, $t_{\rm TI}/t_{\rm ff} \lesssim 10$. This difference stems from the fact that fluid elements are compressed as they move inwards in a spherical system; this compression enhances the density perturbations and accelerates the growth of the thermal instability.

Our criterion for multi-phase structure is not sensitive to large variations about our idealized heating prescription (§3.5.4 and figure 3.6). Furthermore, the multi-phase structure that develops via thermal instability is largely independent of the magnitude of the thermal conduction, even on scales much smaller than the Field length (provided conduction is anisotropic, as is the case in galaxy groups and clusters; §3.7). Anisotropic thermal conduction changes the morphology of the cold gas produced via thermal instability (blobs $\rightarrow$ filaments), but not the presence or amount of cold gas. There is thus a very strong, qualitative difference between isotropic and anisotropic thermal conduction (figure 3.10), which cannot be captured by simply multiplying the heat flux by a suppression factor (as is often done, e. g. Zakamska & Narayan 2003; Voit et al. 2008; Guo & Oh 2008).

Our heating prescription imposes global thermal equilibrium and reduces the accreted mass flux in our model halos relative to cooling-flow values. This reduction is not inevitable, however, because thermal instability can produce cooling-flow-like inflow rates when $t_{\rm TI} \lesssim t_{\rm ff}$. In a globally stable system, the thermal instability thus plays an important role in regulating gas inflow rates. We study the connections among thermal instability, mass inflow and feedback more fully in Paper II.

We argue that a locally stable heating mechanism (such as the one proposed in Kunz et al. 2011) is not required to explain the reduced star formation and cooling rates in clusters. Instead, global stability arising from approximate thermal equilibrium, together with the physics of local thermal instability in stratified plasmas, is sufficient to reproduce the low net cooling rates in clusters. Moreover, the correlation of H_α filaments and star formation in clusters with the cooling time in the hot ICM strongly suggests that the plasma is in fact locally thermally unstable. We have shown that the suppression of accretion rates is

not sensitive to thermal conduction or to significant variations about our specific feedback prescription (§3.5.4 and figure 3.6), and we explore this further in Paper II.

We now compare our model predictions with observational results. The thermal instability time $t_{\rm TI} \equiv p_{\rm TI}^{-1}$ (eq. 3.19) depends on the unknown way in which feedback energy is thermalized (eq. 3.4) and cannot be directly inferred from observations. We therefore use the cooling time when comparing our results with observations. These two timescales differ by an unknown factor of order unity. The ratio of timescales $t_{\rm cool}/t_{\rm ff}$ can be reexpressed in more familiar terms using

$$\frac{t_{\rm cool}}{t_{\rm ff}} \sim 3\frac{(K/10\ {\rm keV\ cm^2})^{3/2}}{T_7^{1/2}\,\Lambda_{-23}\,(t_{\rm ff}/30\ {\rm Myr})}, \tag{3.35}$$

where $K = k_{\rm B}T/n^{2/3}$ is a measure of the plasma entropy, T_7 is the temperature in units of $10^7\,{\rm K}$ and Λ_{-23} is the cooling function (eq. 3.5) in units of 10^{-23} erg cm^3 s^{-1}. Thus, the plasma in clusters and galactic halos should show extended multi-phase structure wherever

$$K \lesssim (20\ {\rm keV\ cm^2})\left[T_7^{1/2}\,\Lambda_{-23}\left(\frac{t_{\rm ff}}{30\ {\rm Myr}}\right)\right]^{2/3}, \tag{3.36}$$

where we have used the threshold from Paper II for multi-phase structure in spherical systems: $t_{\rm TI}/t_{\rm ff} \sim 10$. This criterion is consistent with observations that clusters with central entropies below 30 keV cm^2 preferentially show signs of cold gas such as star formation and $\rm H_\alpha$ emission (Voit et al. 2008; Cavagnolo et al. 2008, 2009). Note, however, the relatively strong dependence of this criterion on $t_{\rm ff}$ and on Λ; this is because the entropy K is not the fundamental parameter governing the saturation of the thermal instability in a stratified system.

The $\rm H_\alpha$ survey conducted by McDonald et al. (2010, 2011a) permits another test of our criterion. McDonald et al. provide lists of groups and clusters with and without extended $\rm H_\alpha$ emission. We test our criterion by estimating the time-scale ratio $t_{\rm cool}/t_{\rm ff}$ for these systems using data from the ACCEPT catalog (Cavagnolo et al. 2009). We fit the entropy profiles of clusters in the ACCEPT catalog using $K(r) = K_0 + K_1(r/100{\rm kpc})^a$ and we fit the pressure profiles $P(r)$ using the form provided in Arnaud et al. (2010). Of the groups and clusters in both the $\rm H_\alpha$ surveys and in ACCEPT, sixteen give reasonable fits (listed in Table 3.2).[5] From our fits to $K(r)$ and $P(r)$, we calculate $n(r)$ and $T(r)$ and estimate $t_{\rm cool}(r)$ using the fit to the cooling function provided by Tozzi & Norman (2001) with 1/3 solar metallicity. We estimate $g(r)$ from the pressure and density profiles by assuming spherical symmetry and hydrostatic equilibrium; from this we calculate $t_{\rm ff}(r)$.

The left panel of figure 3.11 shows $t_{\rm cool}/t_{\rm ff}$ as a function of radius for these sixteen groups and clusters. As predicted by our analysis, the clusters with short cooling times $t_{\rm cool}/t_{\rm ff} \lesssim 10$

[5]Unfortunately, we were unable to fit several clusters with well-known filament systems, including Perseus, Abell 2052 and M87. The pressure gradients in Perseus and Abell 2052 are positive at some radii and cannot be fit by the universal profile. Similarly, the pressure profile for M87 deviates from the broken power-law universal profile.

show extended filaments, while the clusters with long cooling times $t_{\rm cool}/t_{\rm ff} \gtrsim 10$ do not. Additionally, most of the filaments are found at radii around 10–50 kpc, where the ratio $t_{\rm TI}/t_{\rm ff}$ is the smallest. Although this evidence is not conclusive, these data support our hypothesis that the filaments condense from the ICM due to the local thermal instability.

We emphasize that both cooling and gravity influence the development of the thermal instability in the ICM. A short cooling time (or low K) is not sufficient for the formation of filaments; rather, the ratio $t_{\rm TI}/t_{\rm ff}$ is the relevant parameter. The right panel of figure 3.11 shows clusters in the $(t_{\rm cool})$–$(t_{\rm cool}/t_{\rm ff})$ plane. More data are needed to conclusively test our model, but these results are consistent with our interpretation that the ratio of timescales $t_{\rm cool}/t_{\rm ff}$ is a better predictor of multi-phase gas in hot halos than $t_{\rm cool}$ alone.

These simple comparisons support a model in which local thermal instability produces at least some of the H_α filaments seen in clusters. This does not, however, imply that thermal instability alone can explain all of the observed properties of multi-phase gas in clusters and/or galaxies. On the contrary, processes such as conductive condensation of hot gas to cool gas ("non-radiative cooling;" Fabian et al. 2002; Soker et al. 2004) and the inflow of cold gas through the virial radius (Kereš & Hernquist 2009) may also be important (in higher and lower mass halos, respectively). Furthermore, other processes in the ICM such as merger shocks, galaxy wakes and buoyant radio bubbles may influence the evolution of the filaments.

Accretion of the cold gas formed via thermal instability likely plays an important role in the evolution of brightest cluster galaxies and their central AGN (Pizzolato & Soker 2005, 2010). In addition, the high-velocity clouds surrounding the Milky Way may also be manifestations of the thermal instability (Maller & Bullock 2004; Sommer-Larsen 2006; Kaufmann et al. 2006; Peek et al. 2008); this process could provide an important source of unenriched gas to maintain metallicity gradients (Jones et al. 2010) and continued star formation (Bauermeister et al. 2010) in the Milky Way and other galaxies.

Chapter 4

Temperature Profiles in the ICM

4.1 Introduction

X-ray observations of the hot, diffuse gas in galaxy groups and clusters suggest negative temperature gradients at large radii (Leccardi & Molendi 2008; George et al. 2009; Simionescu et al. 2011). This observation is somewhat surprising because negative temperature gradients are susceptible to a convective instability known as the magnetothermal instability, or MTI (Balbus 2001). Without a clear source of free energy to maintain the convection, it seems unusual that clusters should so uniformly be found in unstable states. This observation becomes even more surprising in clusters more massive than $\sim 10^{14.5} M_{\odot}$, where the timescale for heat to diffuse through the ICM can be much shorter than the Hubble time. Left to their own devices, both conduction and convection tend to erase temperature gradients, and one might expect them to make the gas isothermal. In this paper, we study how the assembly of clusters creates large-scale temperature gradients and maintains them in spite of convection and thermal conduction.

Our subject is not purely academic. Systematic trends in temperature gradients with mass and redshift may influence the conversion of observable quantities (such as the x-ray surface brightness) to thermodynamic quantities (gas density, pressure, etc.). Understanding the origin of the temperature gradients in clusters would also enable us to calculate how the turbulence produced by the MTI depends on halo mass or redshift. These trends create systematic variations in non-thermal pressure support and might affect current efforts to use the cluster mass function to constrain cosmology (Allen et al. 2008; Shaw et al. 2010; Allen et al. 2011). Thus, while the temperature profiles in galaxy clusters present an interesting puzzle in their own right, understanding the processes which control them may also find useful application in cluster cosmology.

Several studies, including Dolag et al. (2004), Burns et al. (2010), and Ruszkowski et al. (2011) have begun to address the effects of conduction and the MTI in the ICM using cosmological simulations. These results are computationally expensive, however, and can be difficult to interpret. For example, the MTI is expected to be a sub-dominant, but significant, source of turbulence in clusters (cf. Lau et al. 2009 and Parrish et al. 2012b). Using cos-

mological simulations to study its trends with mass and redshift would require very careful calibration of other sources of turbulence. Thus, we feel that a simplified treatment which affords an intuitive understanding of the results remains useful.

The temperature and entropy profiles in the ICM are related by hydrostatic equilibrium. Accordingly, this paper closely follows earlier work by Tozzi & Norman (2001) and Voit et al. (2003), who study the entropy profiles in clusters. However, the processes controlling temperature gradients in clusters are slightly more subtle than those which determine their entropy profiles. Radial variations in entropy tend to be much larger than those in temperature, so small differences in the entropy profiles translate to much larger differences in temperature profiles. For the same reason, thermal conduction has a more pronounced effect on the temperature profile than on entropy and we must include it in our analysis. Thus, despite significant similarities to both Tozzi & Norman (2001) and Voit et al. (2003), our models represent a generalization of these earlier studies and we use them to explore different astrophysical applications. Komatsu & Seljak (2001) also present analytic models for temperature profiles in clusters. As we describe below, however, our method differs significantly from theirs because we do not assume a polytropic model for the gas; instead we directly calculate its thermodynamic state.

We describe our general methodology in the next section, followed by two sets of models. We begin in section 4.3 with an idealized but intuitive model. This section illustrates the key process that sets temperature gradients in clusters, but is too idealized to be directly compared with real clusters. We generalize this model in section 4.4 to more accurately reflect the formation histories and gravitational potentials of clusters. We also study the influence of thermal conduction on our results. We close in § 4.5 with a brief summary of our conclusions and a description of our future plans to apply them to observations and to more realistic simulations.

4.2 Method

Following Tozzi & Norman (2001) and Voit et al. (2003) (hereafter V03), we model the cosmological assembly of a galaxy cluster and use the evolving properties of its accretion shock to calculate the thermal state of the gas in its ICM. Then, by assuming the ICM is in hydrostatic equilibrium, we determine its temperature and pressure profiles. We focus our attention on how the accretion history of a halo influences its temperature profile, and on how this temperature profile is later modified by thermal conduction. As discussed in section 4.1, we build on previous work that has focused on cluster entropy profiles.

Since the dynamical timescale in clusters is typically shorter than the timescales for either accretion or thermal conduction, we approximate the dynamics of the gas by assuming that it is in hydrostatic equilibrium at all times. We note that this assumption of strict hydrostatic equilibrium within the virial shock is not quantitatively justified: cosmological simulations of cluster formation (e. g. Rasia et al. 2006; Nagai et al. 2007; Lau et al. 2009; Vazza et al. 2011) consistently find ∼10–20% turbulent pressure support driven by mergers near the

virial radius. Our assumption of hydrostatic equilibrium represents a first approximation and provides a simple model with no free parameters.

Our assumption of spherical symmetry is also an approximation: galaxy cluster halos are tri-axial and grow partly by accreting smaller sub-halos. The extent to which these properties influence the temperature profiles in clusters is not entirely clear; including them in our analysis would require cosmological simulations, however, and would likely obscure our results. Instead, we model clusters as spherically symmetric and we assume that they grow primarily via smooth accretion. Though these approximations are not entirely realistic, they are partially justified in V03, who compare models based on smooth accretion with more realistic ones based on hierarchical structure formation: the differences are modest for high-mass halos and mostly reflect changes in normalization due to clumping factors.

Each shell of material accreted by the cluster contains both dark and baryonic matter. We do not solve for the evolution of the dark matter; instead, we assume that the dark matter virializes rapidly and we parametrize it with a time-dependent fit to the potential (we neglect the gravity due to the baryons). We consider both simplified, isothermal potentials (§ 4.3) and more realistic fits to cosmological simulations (§ 4.4).

The baryons enter the cluster through a shock with a density $\rho_{\rm b,i}$ and velocity $v_{\rm i}$ determined by the accretion rate and by the depth of the potential:

$$\rho_{\rm b,i}\, v_{\rm i} = \frac{f_{\rm b}}{4\pi}\frac{1}{r_{\rm sh}^2}\frac{dM_{\rm sh}}{dt}$$
$$v_{\rm i}^2 = \frac{2GM_{\rm sh}}{r_{\rm sh}}(1-\xi).$$

In the above, $r_{\rm sh}$ is the radius of the accretion shock, $M_{\rm sh}$ is the total mass contained within it, and $f_{\rm b} \approx 0.17$ is the cosmic baryon fraction. The parameter $\xi = r_{\rm sh}/r_{\rm ta}$ parametrizes the strength of the virial shock, and $r_{\rm ta}$ is the "turnaround" radius, at which the kinetic energy of the shell vanishes (Gunn & Gott 1972).[1] We assume that the IGM is cold enough that the gas enters the cluster through a strong shock (with an upstream Mach number $M \gg 1$). Thus, the post-shock density and pressure are given by:

$$P_{\rm sh} = \frac{3}{4}\rho_{\rm b,i}\left(\frac{4}{3}v_{\rm i}\right)^2$$
$$\rho_{\rm sh} = 4\rho_{\rm b,i},$$

where the factor of 4/3 in front of the velocity transforms the infall velocity from the frame of the cluster to the frame of the shock. Thus, the post-shock entropy ($K \equiv k_{\rm B}Tn^{-2/3}$) is:

$$K_{\rm sh} = \frac{1}{3}\left(\frac{f_{\rm b}}{4\pi}\right)^{-2/3}\left[\frac{G^4M_{\rm sh}^2(\mu m_{\rm p})^5}{\left(\frac{\partial \ln M_{\rm sh}}{\partial t}\right)^2}(1-\xi)^4\right]^{1/3}. \quad (4.1)$$

[1] Note that identifying ξ with $r_{\rm sh}/r_{\rm ta}$ assumes that the baryons do not separate from dark matter before reaching the virial shock. If this were not the case, the baryons would feel a time-dependent potential due to the dark matter. Tozzi & Norman (2001) find that this may introduce a ∼10% correction to the infall velocity.

Equation 4.1, above, is identical to equation 8 in V03. As emphasized in V03, apart from an overall normalization $\propto (M/f_{\rm b})^{2/3}$, the entropy profile depends only on the accretion history and on the shock strength ξ.

The jump conditions above represent the conservation of mass, momentum, and energy and thus reflect the *total* pressure behind the shock. In general, however, the electron and ion pressures may differ. The difference may be significant in clusters because the shock preferentially heats ions and because the timescale for electrons and ions to equilibrate is long (Fox & Loeb 1997). We don't distinguish between electron and ion temperatures in our calculation because the temperature difference does not effect hydrostatic equilibrium and thus should not influence our solutions. Moreover, the simulations by Rudd & Nagai (2009) show that this temperature difference is modest within the virial radius of the cluster.

Our assumption of a strong virial shock may not be valid for the innermost shells of material, which accreted when the IGM was hotter, and when the gravitational potential of the halo was shallower, than they are today. Hence, our model will not accurately reproduce the gas profiles near the centers of clusters (Tozzi & Norman 2001). Many other processes, including cooling, heating, and the formation of the central galaxy also affect the structure of the core, however (Voit 2011). We instead focus on the gas at larger radii, near the virial radius. Both the large mass and the long cooling time of this gas enable us to ignore galaxy formation processes at smaller radii (e. g. Voit & Ponman 2003; Voit 2011).

After undergoing the virial shock, the gas entropy evolves due to thermal conduction. Thermal conduction is highly anisotropic in the ICM because the electron mean free path is much longer than its gyroradius. Anisotropic conduction renders the ICM unstable to the MTI (Balbus 2001), however, which generates turbulence and may isotropize the magnetic field (McCourt et al. 2011b; Parrish et al. 2012b). We therefore parametrize thermal conduction through an effective isotropic conductivity. We introduce the effective conductivity $\kappa_{\rm eff} \equiv f_{\rm Sp}\kappa_{\rm e}$, where $f_{\rm Sp} \sim 1/3$ is a suppression factor due to the magnetic field.[2] Thus, the evolution of the entropy is determined by the following energy equation:

$$\frac{d}{dt}\ln K = \frac{2}{3P}\nabla\cdot\left(\frac{\kappa_{\rm eff}}{k_{\rm B}}\nabla T\right). \tag{4.2}$$

In the simplifying case that $\kappa_{\rm eff} \to 0$, the entropy of each shell is a constant in time and equal to $K_{\rm sh}$, given by equation 4.1. In addition to the conductivity κ, we also use the thermal diffusion coefficient $\chi_{\rm e} \equiv \kappa_{\rm e}/(n_{\rm e}k_{\rm B})$, which has units of $\rm cm^2/s$.

As mentioned above, anisotropic thermal conduction in the ICM drives a convective instability known as the MTI. This convection carries a heat flux which should technically be added to equation 4.2. However, the convective heat flux is smaller than the conductive flux

[2]Though this approximation is suitable for our purposes, we note in passing that the use of an isotropic conductivity significantly alters other processes in the ICM, such as convection (Balbus 2001; Quataert 2008) and thermal instability (Sharma et al. 2010; McCourt et al. 2012; Sharma et al. 2012a), and is thus not suitable for more detailed dynamical studies. Interestingly, a suppression factor $f_{\rm Sp} \sim 1/3$ turns out to be appropriate even if the magnetic field is tangled on very small scales (Narayan & Medvedev 2001).

by a factor of

$$\frac{Q_{\rm conv}}{Q_{\rm cond}} \sim \sqrt{\frac{m_{\rm e}}{m_{\rm i}}} \left(\frac{H}{\lambda_{\rm e}}\right) \alpha^3 \left(H \frac{d\ln T}{dr}\right)^{3/2},$$

where $m_{\rm e}$ and $m_{\rm i}$ are the electron and ion masses, H is the pressure scale-height, $\lambda_{\rm e}$ is the electron mean-free-path, and α is the mixing-length parameter of the convection. This ratio is small ($\sim 10^{-2}$) in the ICM, enabling us to ignore the convective heat flux in equation 4.2 (cf. Parrish et al. 2008).

Our spherically symmetric model is most easily represented in Lagrangian coordinates with the enclosed mass as the independent variable. Thus, the continuity equation and the equation for hydrostatic equilibrium take the form:

$$\mu m_{\rm p} \frac{dr}{dM_{\rm b}} = \frac{1}{4\pi r^2} \left(\frac{K}{P}\right)^{3/5}, \tag{4.3}$$

$$\frac{dP}{dM_{\rm b}} = -\frac{g}{4\pi r^2}, \tag{4.4}$$

where $M_{\rm b}$ is the *baryonic* mass contained within the radius r, and we adopt the entropy K and pressure P as our primary thermodynamic variables. In this initial study, we ignore sources of non-thermal pressure support (such as cosmic rays, magnetic fields, or turbulence). However, we discuss in section 4.5 our plans to self-consistently include turbulence generated by the MTI and to compare these results with cosmological simulations.

Equations 4.1–4.4 completely specify our model except for the mass accretion history $M_{\rm sh}(t)$ and the gravitational field $g(r,t)$ of the halo. We assume that these are determined by cosmology and are unaffected by the baryonic formation of the cluster. We describe a simple, idealized model for the halo formation in section 4.3 and a more realistic model in section 4.4.

4.3 Simplified Adiabatic Models

Before studying the consequences of thermal conduction, it is useful to understand the 'baseline' temperature profile generated by the halo's mass accretion history. Thus, we begin with adiabatic models which ignore thermal conduction. We further isolate the influence of the accretion history by assuming the gravitational potential is isothermal and that the accretion rate of the cluster is independent of time: $M_{\rm sh}(t) = M_0 \times t/t_0$. As we show later, this prescription contains only a single free parameter and is perhaps the simplest nontrivial model of the process we wish to study. Though the results in this section cannot be directly applied to clusters, they highlight some of the key physics determining the temperature gradients in clusters, and will assist in our interpretation of the more detailed models in section 4.4.

The next step in our model is to determine the shock strength ξ. The most logical choice would be to calculate $\xi(M_{\rm b})$ so that the cluster was in hydrostatic equilibrium at every

epoch; in fact, this is formally required to use equation 4.1 for the post-shock entropy. In this section, however, we make the simplifying assumption that $\xi(M_{\rm b})$ is a constant. Thus, the models presented in this section are not entirely self-consistent. Our goal in this section is only to obtain an intuitive understanding of how the accretion rate of a halo influences its temperature gradient. We present more detailed models, with more accurate results, in section 4.4.

Since we have assumed a solution for $\xi(M_{\rm b})$ and that the evolution of the gas is adiabatic, the entropy profile $K(M_{\rm b})$ is uniquely determined at all times by equation 4.1. Thus, it suffices to solve hydrostatic equilibrium only at the present epoch – this solution cannot depend on the state of the cluster at earlier times. Hence, we need not track the evolution of the cluster, and the equations determining the state of the gas reduce to ordinary differential equations.

4.3.1 Method

Before solving the equations of our model, we recast them in a more versatile dimensionless form. Since we have a system of ordinary differential equations with the shock as one boundary, we de-dimensionalize the equations in this section using the properties at the *shock radius*. (Note that this differs from the usual convention of using the virial radius of the underlying dark-matter potential.) We introduce the constants $R_{\rm sh}$ and M_0, which represent the shock radius of the cluster at the present time t_0 and the total mass enclosed within it. (In an isothermal potential, $R_{\rm sh}$ and M_0 each differ from the virial radius and mass by a factor of 2ξ.) We also define a dynamical time $t_{\rm dyn} \equiv (GM_0/R_{\rm sh}^3)^{-1/2}$, along with the spatial coordinate $x \equiv r/R_{\rm sh}$ and the Lagrangian mass coordinate $y \equiv M_{\rm b}/(f_{\rm b}M_0)$.

We introduce the dimensionless gas variables P_1, ρ_1, and K_1 via:

$$\rho_{\rm b} \equiv \frac{f_{\rm b}}{4\pi}\frac{M_0}{R_{\rm sh}^3} \times \rho_1 \tag{4.5a}$$

$$P \equiv \frac{f_{\rm b}}{4\pi}\frac{GM_0^2}{R_{\rm sh}^4} \times P_1 \tag{4.5b}$$

$$K_1 \equiv P_1/\rho_1^{5/3}. \tag{4.5c}$$

Thus, the equations for mass conservation and hydrostatic equilibrium become:

$$\frac{dx}{dy} = \frac{(1-\xi)^{4/5}}{3^{3/5}}\left(\frac{t_{\rm dyn}}{t_0}\right)^{-2/5}\frac{y^{4/5}}{x^2 P_1^{3/5}} \tag{4.6a}$$

$$\frac{dP_1}{dy} = -\frac{1}{x^3}, \tag{4.6b}$$

which we solve subject to the boundary conditions at the shock:

$$x(y=1) = 1 \tag{4.7a}$$

$$P_1(y=1) = \frac{4\sqrt{2}}{3}(1-\xi)^{1/2}\left(\frac{t_{\rm dyn}}{t_0}\right) \tag{4.7b}$$

and at the center:

$$x(y = 0) = 0. \tag{4.7c}$$

With these definitions, the solution is independent of the parameters $f_{\rm b}$, μ, $R_{\rm sh}$, and M_0. Furthermore, the system is over-determined with three boundary conditions and two equations. The shock radius ξ is therefore an eigenvalue which must be chosen to meet the inner boundary condition in equation 4.7c.

Thus, the *only* free parameter in our system of equations is the ratio $t_{\rm dyn}/t_0$, and we expect to find a one-dimensional family of models. Since the average density of any dark matter halo $\bar{\rho} \sim 200\rho_{\rm crit}$ is independent of mass in hierarchical structure formation, we do not expect the dynamical timescale $t_{\rm dyn}$ to vary strongly among clusters at any cosmological epoch. The ratio $t_{\rm dyn}/t_0$ thus measures the age (or, equivalently, the assembly rate) of the halo.

We note that the following solution to equations 4.6 and 4.7:

$$P_1 = \frac{1}{2x^2}, \quad x = y, \quad \xi = \frac{1}{4} \tag{4.8}$$

exists when the assembly rate satisfies:

$$\frac{t_{\rm dyn}}{t_0} = \sqrt{\frac{3}{32}}.$$

We describe the physical significance of this special, isothermal solution in section 4.3.2.

For other values of the assembly rate $t_{\rm dyn}/t_0$, we solve equations 4.6 and 4.7 numerically using a predictor-corrector method, and we solve the eigenvalue problem for ξ with a bisection search. We avoid the singularity in equations 4.6a and 4.6b by solving them on a logarithmic grid. Since we cannot apply the boundary condition equation 4.7c in the logarithmic coordinates, we obtain an approximate boundary condition at a finite radius by expanding the equation for hydrostatic equilibrium near $x = 0$. Assuming that the temperature remains finite, this equation becomes:

$$T \to T_{\rm vir} \left(\frac{d \ln \rho^{-1/2}}{d \ln r} \right)^{-1}.$$

Combined with the fact that $K \propto T\rho^{-2/3} \propto y^{4/3}$, this implies that $\rho \propto r^{-2}$, and thus that

$$T \to T_{\rm vir}, \text{ and} \tag{4.9}$$

$$y \to \left(\frac{T_{\rm vir}}{T_{\rm sh}} \right)^{1/2} \left(\frac{3\rho_{\rm sh}}{\bar{\rho}} \right)^{1/3} x \tag{4.10}$$

as $x \to 0$. We use equation 4.10 as a boundary condition for the eigenvalue problem and equation 4.9 to check the accuracy of our integration. The analytic solution in equation 4.8 also permits a more pedantic test of our method.

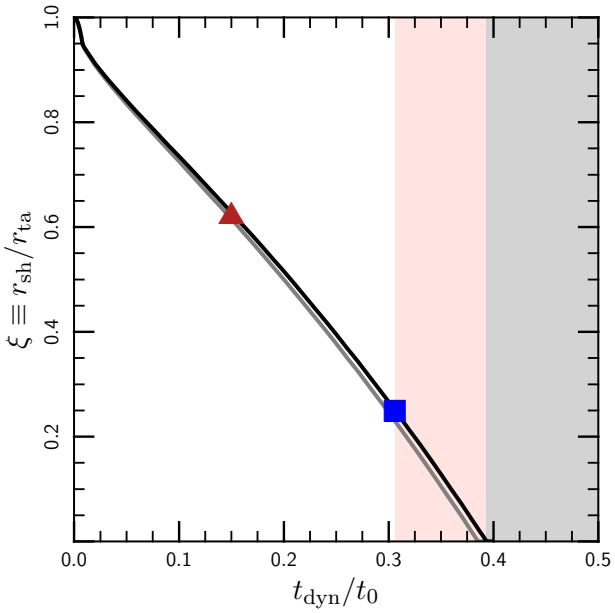

Figure 4.1: Shock radius ξ as a function of the assembly parameter $t_{\rm dyn}/t_0$ for the simplified models described in section 4.3. The timescale $t_0 = M/\dot{M}$ is the age of the halo. Points further to the left on this plot correspond to clusters which form slowly and points further to the right correspond to clusters which form rapidly. The red triangle represents the approximate value of $t_{\rm dyn}/t_0$ expected for cluster halos. The blue square marks the isothermal solution in equation 4.8. Clusters in the white region of the plot have negative temperature gradients, clusters in the pink region have positive temperature gradients, and clusters in the gray region cannot exist in our steady-state model. Real clusters are expected to have moderate, negative temperature gradients based on this analysis (see figure 4.3). The black line illustrates the calculation with an isothermal potential, and the gray line shows results with a more realistic NFW potential.

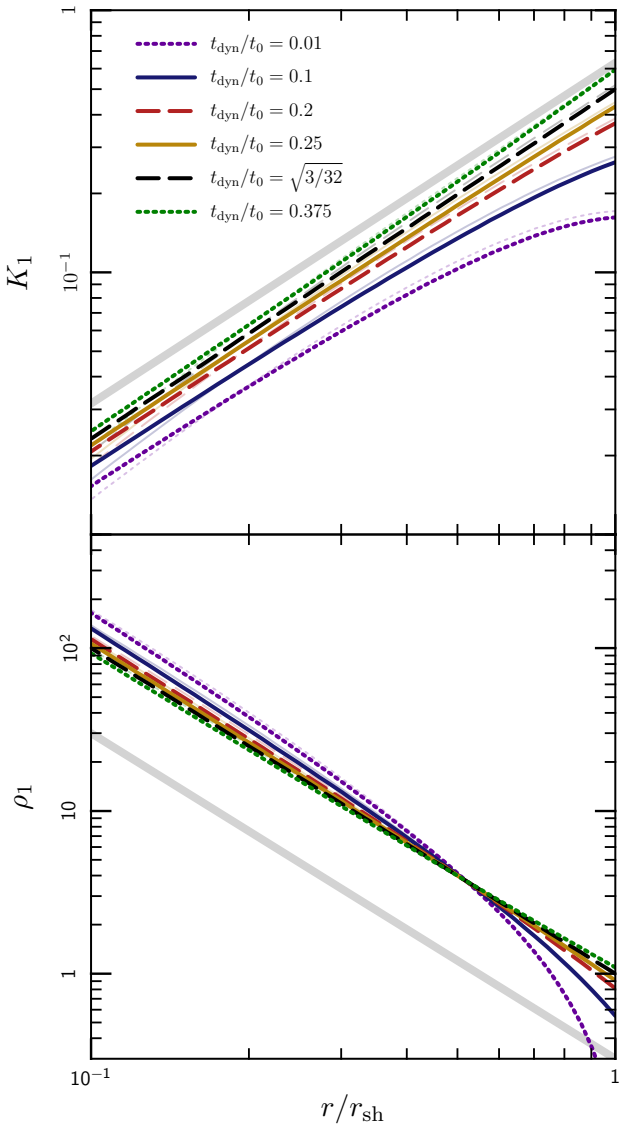

Figure 4.2: Profiles of density and entropy for representative models from fig. 4.1 with different values of the assembly parameter $t_{\rm dyn}/t_0$. Thick gray lines show typical power-law slopes derived from x-ray observations (Croston et al. 2008; Cavagnolo et al. 2009) (the normalization is arbitrary). Massive clusters should mostly lie between the solid blue and long-dashed red lines. Our solutions are approximately, but not exactly, power-laws. The deviations from power-law behavior are dictated by the outer boundary condition on the pressure and determine the temperature profile (fig. 4.3). This boundary condition depends on the ram pressure behind the shock, and thus on the speed with which the cluster formed (parametrized by $t_{\rm dyn}/t_0$ in this model). As in fig. 4.1, thick curves show models with simplified, isothermal potentials and thin curves show more realistic models with NFW potentials.

4.3.2 Results

Figure 4.1 shows the solution for the dimensionless shock radius $\xi(t_{\rm dyn}/t_0)$ from our numerical calculations. In the limit of very slow accretion (i. e. as $t_{\rm dyn}/t_0 \to 0$), the shock is comparatively weak and the dimensionless shock radius $\xi \to 1$. As the assembly rate $t_{\rm dyn}/t_0$ increases, the shock radius ξ decreases monotonically. Thus, the virial shock moves inwards as the accretion rate increases. This result seems intuitively reasonable, as a higher accretion rate implies a higher ram pressure behind the shock. A stronger shock (or smaller ξ) is thus required to hold back the infalling material and to keep the ICM in hydrostatic equilibrium.

Interestingly, figure 4.1 indicates a maximum assembly rate around $t_{\rm dyn}/t_0 \approx 0.39$ at which $\xi \to 0$. Beyond this point, thermal pressure alone cannot hold back the accretion shock and hydrostatic equilibrium becomes impossible. This represents an extremely rapid accretion rate, however, with the halo forming over only ~ 2.5 dynamical times; our quasi-static model for the ICM breaks down in this limit. A fully dynamical calculation (e. g. a simulation) with such a high accretion rate would likely produce a model ICM with significant time dependence and turbulent pressure support, but with a finite shock radius.

Figure 4.1 shows the isothermal solution (eq. 4.8) as a blue square and a point representative of a massive cluster [with $t_{\rm dyn} = 0.1 H_0^{-1}$ and $t_0 \sim (2/3) H_0^{-1}$] as a red triangle. Our model predicts that $\xi \sim 0.6$ for this fiducial cluster, similar to what has been expected in the past (e. g. Rees & Ostriker 1977). Note that the isothermal solution requires a much faster assembly than is typical for galaxy clusters; this simple model thus suggests that clusters should not be isothermal.

Figure 4.1 directly illustrates the effect of the accretion rate on the location of the virial shock. The gas properties in the ICM must match onto jump conditions at the shock; thus, by moving the shock radius, the accretion rate also influences the thermodynamic structure of the ICM. Figure 4.2 demonstrates this by showing profiles of the gas density and entropy for models with different assembly rates $t_{\rm dyn}/t_0$. The profiles are nearly isothermal, with $\rho \sim r^{-2}$ and $K \sim r^{4/3}$, and are broadly consistent with determinations from x-ray data (Croston et al. 2008; Cavagnolo et al. 2009).

The small deviations from power-laws in the density and entropy profiles lead to significant temperature gradients, however. We show this explicitly in figure 4.3, where we plot the temperature profiles for the models from figure 4.2. The temperature profiles are nearly linear, and gradients of either sign are possible, depending on the assembly rate $t_{\rm dyn}/t_0$. This result is not surprising, since the central temperature must equal the virial temperature of the halo (eq. 4.9) and the temperature at the shock is dictated by the jump conditions. A roughly linear interpolation between these boundary conditions seems reasonable given the simplicity of the model. The isothermal solution with $t_{\rm dyn}/t_0 = \sqrt{3/32}$ divides models with negative and positive temperature gradients. In reality, most clusters satisfy $t_{\rm dyn}/t_0 \sim 0.1 - 0.2$; in this case our model predicts temperature profiles which decrease by a factor of ~ 2 from the center to the shock radius. This result is in line with recent observations of the gas temperature near the virial radius (George et al. 2009; Simionescu et al. 2011).

Thus, while we solve the system of equations 4.6a and 4.6b and boundary conditions 4.7a–

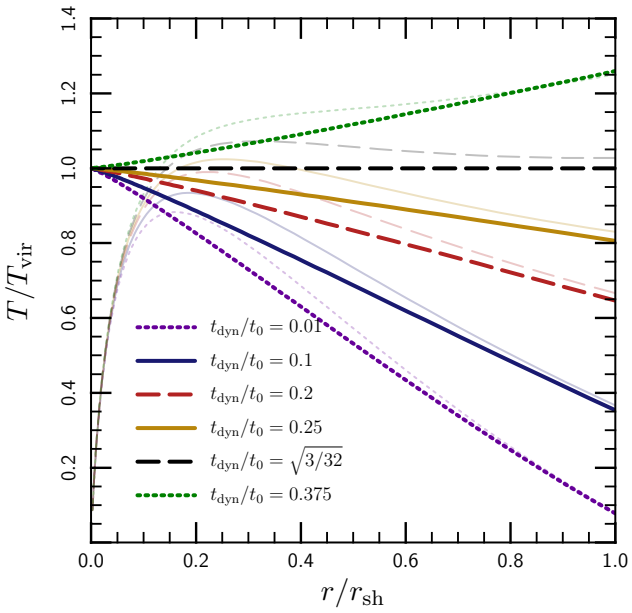

Figure 4.3: Temperature profiles for representative models from fig. 4.1 with different values of the assembly parameter $t_{\rm dyn}/t_0$. The black line corresponds to the isothermal solution (eq. 4.8); massive clusters should lie between the blue and red lines. Our model implies temperature profiles which decrease by a factor of $\sim$ 2 from the center to the virial radius, in line with x-ray observations. Thin, light lines show the results of calculations with NFW potentials, rather than isothermal ones (see § 4.4 for details). The temperature gradients are similar outside the scale radius of the halo.

4.7c for the shock radius and for the structure of the ICM, we find that the boundary conditions essentially dictate the temperature profile. The temperature must reach the virial temperature of the potential at the center of the halo, and it must match onto the jump conditions at the shock. The solution to hydrostatic equilibrium then implies a nearly linear interpolation between these two boundary conditions.

As presented here, the temperature profiles in figure 4.3 may seem specific to our assumption of an isothermal potential. To demonstrate that this is not the case, we have included calculations with NFW potentials in figure 4.3 (shown as thin, light lines). Although the different potential has a dramatic effect on the temperature profile within the scale radius, the trend between the assembly rate and the overall temperature gradient at large radii is similar. We include these lines only for illustration, but present much more detailed models with NFW potentials in the following section 4.4.

The simple model presented in this section suggests that the assembly rate of the halo (or, equivalently, the ram pressure behind the accretion shock) dictates the large-scale temperature gradient in the ICM. The assembly rates of massive clusters are such that they should have moderate, negative temperature gradients outside the scale radius. This is one of our primary findings. In what follows, we show that this result holds true even when we relax the simplifying assumptions in this section. We also study how thermal conduction modifies this 'baseline' temperature profile.

4.4 Conduction and Realistic Assembly Histories

Our models with isothermal potentials and linear accretion histories are especially transparent. The models are very idealized, however, and it is not clear how accurately they carry over to real clusters. In this section, we generalize our results to include more realistic potentials and accretion histories; we also include thermal conduction in our calculation and study its influence on the temperature profiles in clusters.

4.4.1 Method

Coordinates and Assumptions

Introducing realistic potentials and accretion histories into our model necessitates a few changes to our method. Though the unit system introduced in section 4.3 is ideal for our model equations, the virial mass of the halo (as opposed to the mass enclosed by the virial shock) is an eigenvalue of the problem and cannot be specified ahead of time. Since we want to study the variation in temperature profiles at fixed *virial* mass, we must alter the equations and boundary conditions slightly.

We adopt a unit system based on the virial mass $M_{\mathrm{vir},0}$ and radius $r_{\mathrm{vir},0}$ of the halo at redshift zero. Following our approach in the last section, we introduce the dimensionless

variables P_1, ρ_1, and K_1 via:

$$\rho_{\rm b} \equiv \frac{f_{\rm b}}{4\pi}\frac{M_{\rm vir,\,0}}{r^3_{\rm vir,\,0}} \times \rho_1 \tag{4.11a}$$

$$P \equiv \frac{f_{\rm b}}{4\pi}\frac{GM^2_{\rm vir,\,0}}{r^4_{\rm vir,\,0}} \times P_1 \tag{4.11b}$$

$$K_1 \equiv \frac{P_1}{\rho_1^{5/3}}. \tag{4.11c}$$

We also define the spatial coordinate $x \equiv r/r_{\rm vir}$, the Lagrangian coordinate $y \equiv M_{\rm b}/(f_{\rm b}M_{\rm vir})$, and the dynamical time $t_{\rm dyn} \equiv (GM_{\rm vir}/r^3_{\rm vir})^{-1/2}$. Note that our definitions of x, y, and $t_{\rm dyn}$ use instantaneous values of $r_{\rm vir}$ and $M_{\rm vir}$, while our definitions of ρ_1, P_1, and K_1 are normalized to $r_{\rm vir,\,0}$ and $M_{\rm vir,\,0}$. Thermodynamic quantities in our calculation (e. g. K_1) are thus directly comparable at different redshifts, while coordinates (e. g. x) are not.

We define $m \equiv M_{\rm vir}(t)/M_{\rm vir,\,0}$, which tracks the formation of the halo and functions as a time coordinate. We take the virial radius to be r_{200}, the radius within which the mean density of the halo is 200 times the critical density of the universe. Thus, $t_{\rm dyn} = (10H)^{-1}$, where H is the Hubble parameter.

Since the virial radius does not directly enter into our model (§ 4.3), our choice defining the virial radius is essentially arbitrary. We chose r_{200} because it is straightforward to compute and facilitates comparison with much of the existing literature. We continue to use the notation "$r_{\rm vir}$" over the notation "r_{200}" in order to de-emphasize this arbitrary definition, however.

Dark Matter

As discussed in section 4.2, we do not solve for the evolution of the dark matter; instead, we assume that the dark matter evolves independently of the baryons and we parametrize it using fits to halos from cosmological n-body simulations. In particular, we assume that the dark matter follows an NFW distribution at all times (Navarro et al. 1997), with a constant concentration parameter $c = 5$, as is appropriate for actively forming, massive halos (Zhao et al. 2009). The only free parameter in this model for the dark matter is the mass accretion history of the halo.

We use fits to halo mass accretion histories of the form

$$m(z) = [(1+z)^b \exp(-z)]^\gamma, \tag{4.12}$$

derived by McBride et al. (2009) from the Millennium simulation (Springel et al. 2005). This fit is calibrated to the "friends of friends" mass $M_{\rm FoF}$ (Davis et al. 1985), which is similar to our choice of M_{200} for the virial mass (White 2001). Figure 4.4 shows the parameter space for the exponents γ and b, along with the definitions of Type I, II, III, and IV accretion

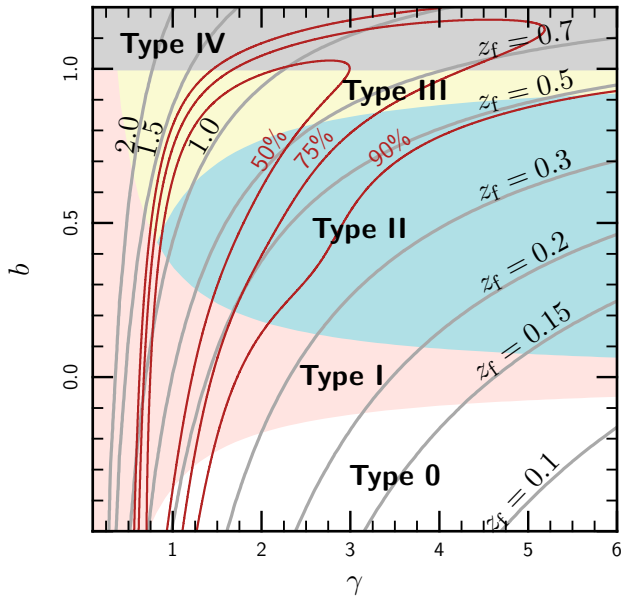

Figure 4.4: Parameter space for dark matter accretion histories based on the Millennium Simulation. The exponents γ and b are defined in equation 4.12. Colors demarcate the different accretion "types" from McBride et al. (2009) and gray lines show contours of the formation redshift z_f such that $M_{vir}(z_f) = 0.5\, M_{vir}(z = 0)$. At redshift zero, Type III halos are accreting very slowly and Type 0 halos are accreting rapidly. The red curves mark contours of the probability density for $10^{15} M_\odot$ clusters to have a given accretion history (taken from the Appendix of McBride et al. 2009). The contours for $10^{14} M_\odot$ halos are very similar to those shown here.

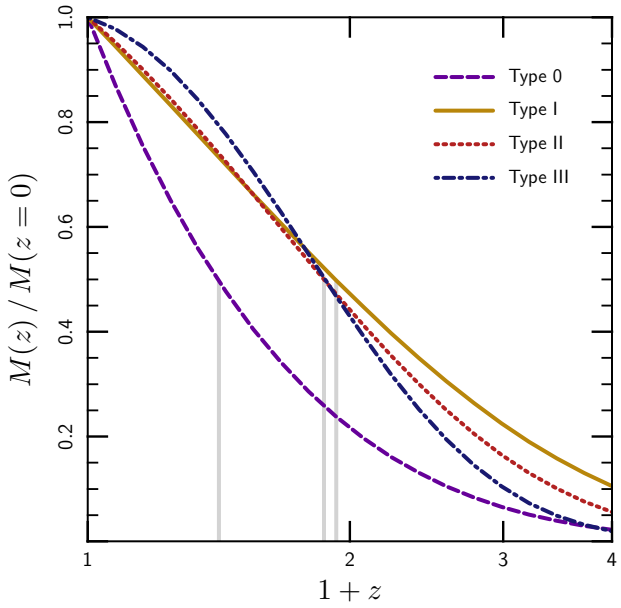

Figure 4.5: Example accretion histories from equation 4.12 illustrating each of the four "types." The examples correspond to $\{\gamma, b\} = \{0.75, -1.5\}$ (Type 0), $\{0.75, 0.0\}$ (Type I), $\{1.25, 0.5\}$ (Type II), and $\{2.25, 0.9\}$ (Type III). These examples are also used in figure 4.6. Note that Type 0 halos have larger accretion rates, and that Type III halos have lower accretion rates, than intermediate halo types at low redshift. Vertical, gray lines indicate the "formation redshift" for which $M(z_{\rm f}) = 0.5\,M(z = 0)$. Though the formation redshift does not have a one-to-one correspondence with accretion type, Type III halos tend to form early and Type 0 halos tend to form late (cf. figure 4.4).

histories from McBride et al. (2009),[3] and figure 4.5 shows an example of each "Type" of accretion history. For reference, the red curves in figure 4.4 show contours of the probability density function for halo accretion parameters from the appendix of McBride et al. (2009). We also show contours of the "formation redshift" $z_{\rm f}$ at which the cluster reaches half its present mass; most of the halos formed at a redshift between 0.5 and 1.

Note that a Type I accretion history is roughly exponential in redshift, which is typical in ΛCDM (Wechsler et al. 2002). At fixed mass, Type 0 halos are younger, while Type II and Type III halos are older, than Type I halos. A Type IV accretion history corresponds to mass loss at late times; this would imply a negative ram pressure at the virial shock in our quasi-equilibrium model. These cases should be studied with fully dynamical, cosmological simulations. Fortunately, at the high masses we wish to study, a relatively small fraction of the total halos exhibit Type IV accretion histories. Moreover, these are systems which have recently undergone major mergers; they are likely to be morphologically disturbed and may

[3] McBride et al. (2009) do not distinguish between Type II halos with negative and positive values of the exponent b. Since this distinction is important in our application, we denote Type II halos with $b < 0$ as Type 0.

be excluded from cosmological samples.

In order to calculate the strength of the virial shock, we require an estimate for the turnaround radius $r_{\rm ta}$ (§ 4.2). Unfortunately, fits to $r_{\rm ta}$ from n-body simulations do not seem to be available. Therefore, we simply use the virial theorem to estimate that the turnaround radius is twice the virial radius. The shock radius is then given in terms of the virial radius by $r_{\rm sh} = 2\xi r_{\rm vir}$. We note that this approximation may over-estimate the turn-around radius (cf. Diemand et al. 2007), causing us to predict temperature profiles which are too shallow. This is one of the primary sources of uncertainty in our models.

The quasi-equilibrium model for the dark matter described in this section greatly simplifies our method by eliminating the need to solve for the dark matter dynamics (e. g. by using an n-body simulation or by solving the Jeans equations). This model is ambiguous outside the virial radius, however. Since the virial shock typically lies exterior to the virial radius, it is not a priori clear what value of $M_{\rm sh}$ to use in equation 4.1 for the post-shock entropy. We proceed by presuming that the gas and dark matter first separate at the virial shock; thus, the gravitating mass $M_{\rm sh}$ in equation 4.1 corresponds to the mass $M_{\rm vir}$. After the shock, the gas remains in hydrostatic equilibrium at $r_{\rm sh}$ while the dark matter continues to collapse and virializes at the virial radius $r_{\rm vir} < r_{\rm sh}$. Therefore, when we solve for hydrostatic equilibrium in the post-shocked gas, we assume that the dark matter has relaxed and we extrapolate the NFW profile between the virial- and shock radii. This inconsistency in our treatment of the gravitating mass is an unavoidable consequence of applying a quasi-equilibrium model for the dark matter outside the virial radius.

Gas Equations

We solve for the state of the gas using a method very similar to that described in section 4.3, but we now solve for the shock radius $\xi(m)$ self-consistently. We discretize the halo formation into the accretion of a finite number of shells and, for each shell m_i, we solve the full eigenvalue problem for its shock strength $\xi(m_i)$. Thus, we simultaneously build up solutions for the shock radius and for the temperature profile as functions of time. As discussed above, we assume that the gas and dark matter first separate at the virial shock. The baryonic accretion rate at the shock radius $r_{\rm sh}$ is thus also proportional to equation 4.12. This is qualitatively consistent with the findings of Faucher-Giguère et al. (2011), who show that the baryonic accretion closely tracks the dark matter accretion history in simulations of high-mass halos.

The equations for mass conservation and hydrostatic equilibrium are:

$$\frac{dx}{dy} = \frac{h^2}{x^2}\left(\frac{K_1(y)}{P_1(y)}\right)^{3/5}, \tag{4.13a}$$

$$\frac{dP_1}{dy} = -\frac{m^{2/3}h^{8/3}}{x^3}\frac{1}{x}\frac{\log(1+cx) - cx/(1+cx)}{\log(1+c) - c/(1+c)}, \tag{4.13b}$$

with the boundary conditions:

$$P_1(y=1) = \frac{4\sqrt{2}}{3} m^{2/3} h^{8/3} \eta \sqrt{\frac{1-\xi}{(2\xi)^5}} \tag{4.14a}$$

$$x(y=1) = 2\xi \tag{4.14b}$$

$$x(y=0) = 0. \tag{4.14c}$$

In the above, $\eta \equiv t_{\rm dyn}(\partial \ln M_{\rm vir}/\partial t)$ measures the accretion rate of the halo, and $h \equiv H/H_0 = [\Omega_{\rm m}(1+z)^3 + \Omega_\Lambda]^{1/2}$ is the Hubble parameter in units of H_0. In equations 4.13 and 4.14, η, h, and ξ are all evaluated at the epoch of the most recently accreted shell. It is straightforward to show from equation 4.12 that $\dot{m} = m\gamma(1+z-b)H$. Thus, the dimensionless accretion rate η is given by $\eta = \gamma(1+z-b)/10$.

The shock entropy $K_1^{\rm sh}$ takes the form:

$$K_1^{\rm sh}(y) = \frac{1}{3}\left[\frac{y}{\eta[z(y)] \times h[z(y)]}\right]^{2/3} [1-\xi(y)]^{4/3}, \tag{4.15}$$

where $z(y)$ represents the redshift at which the baryonic shell y accreted and $\xi(y)$ represents the shock radius of the cluster at that redshift.

Inserting the conductivity appropriate for a fully-ionized hydrogen plasma (Spitzer 1962) into equation 4.2 yields:

$$\begin{aligned}\frac{d\ln K_1}{dm} = 0.70 f_{\rm Sp} \frac{4\pi}{f_{\rm b}} \frac{\mu^{7/2}}{hm\eta} \left(\frac{M_{\rm vir,0}}{10^{15} M_\odot}\right) \\ \times K_1^{5/2} \left[\frac{2}{x}\frac{\partial T_1}{\partial x} + \frac{5}{2T_1}\left(\frac{\partial T_1}{\partial x}\right)^2 + \frac{\partial^2 T_1}{\partial x^2}\right].\end{aligned} \tag{4.16}$$

(Recall that the coordinate m tracks the formation of the cluster and thus functions as a time coordinate). We integrate this equation between accretion events with an explicit, sub-cycled method. The conductive heat flux must vanish at the origin by spherical symmetry; thus, we adopt $\partial T/\partial r = 0$ as the inner boundary condition for equation 4.16. The precise boundary condition on the heat flux at the shock is uncertain because it depends on the physics of collisionless shocks in the presence of strong thermal conduction. We proceed by assuming that electrons do not diffuse across the shock into the upstream flow; in our model, the shock thus serves as an insulating boundary. This assumption is convenient because thermal conduction does not modify the structure of the shock. We have also tried calculations in which we keep the heat flux constant at the shock; the results with this alternative boundary condition were very similar to those we present here.

Note that equations 4.13–4.15 are independent of halo mass, while equation 4.16 is not. Conduction thus introduces non-self-similar behavior and may influence mass-observable relations. We quantify this departure from self-similarity in the following section.

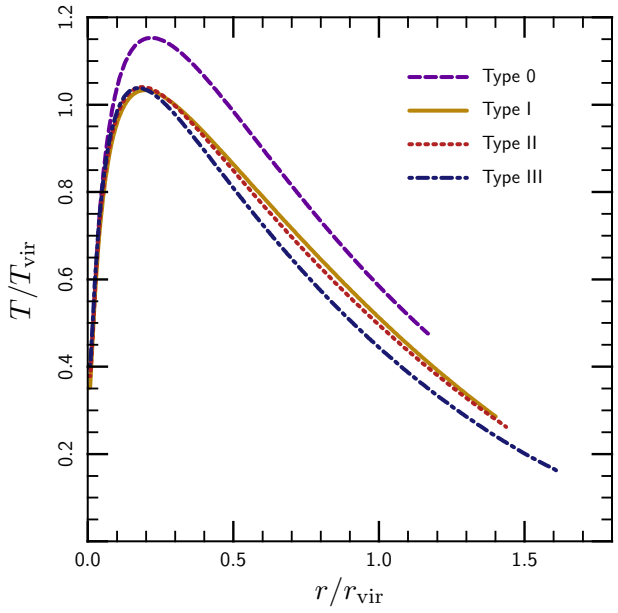

Figure 4.6: Representative temperature profiles at redshift $z = 0$ generated by each type of accretion history from McBride et al. (2009) (see figures 4.4 and 4.5). The profiles are calculated using the Lagrangian method from section 4.4, but with no thermal conduction (the temperature profiles are thus independent of the halo mass). Note the qualitative agreement with the simple models from section 4.3. In addition, the location of the shock radius depends fairly sensitively on the accretion history of the halo.

The post-shock entropy (eq. 4.15) and the outer boundary condition on the pressure (eq. 4.14a) depend on the accretion history through the dimensionless accretion rate η (cf. § 4.3). Thus, the diversity in accretion histories may generate scatter in the ICM properties at fixed halo mass. In order to estimate the statistics in ICM properties, we generate an ensemble of accretion histories for each halo mass $M_{\rm vir,\,0}$, with the exponents γ and b drawn from the distribution in the appendix of McBride et al. (2009). This ensemble yields information about the statistics of the cluster population and the extent to which the variation in accretion histories creates scatter in the temperature profile and mass-observable relations. We present these results in the next section.

4.4.2 Results

Before studying the effect of thermal conduction, we present adiabatic models in which we set the effective conductivity to zero. These profiles facilitate comparison with the simpler models in section 4.3 and also provide a basis for understanding the models with conduction.

Figure 4.6 shows representative, adiabatic solutions for the redshift-zero temperature profiles resulting from several different accretion histories. (Because we have neglected thermal

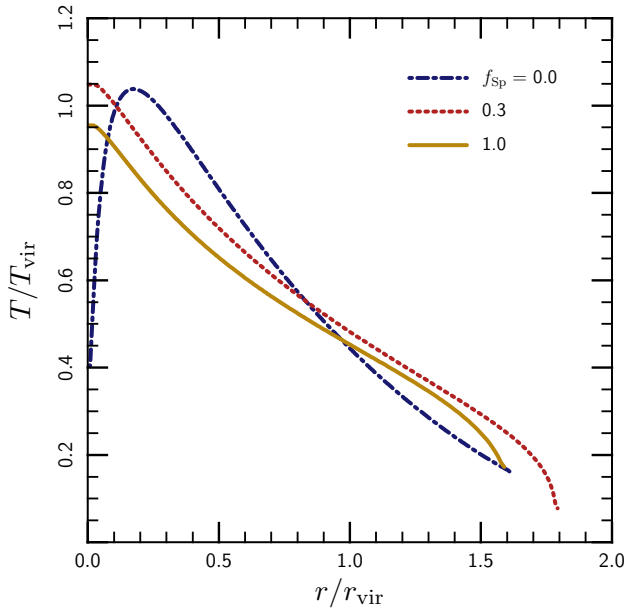

Figure 4.7: Influence of the effective conductivity on cluster temperature profiles. These curves show temperature profiles for massive, $10^{15} M_{\odot}$ halos with Type III accretion histories at redshift $z = 0$. This choice of mass, accretion type, and redshift maximizes the effect of thermal conduction; nonetheless, the change to the temperature profile at large radii is modest. As discussed in the text, the marked effect of thermal conduction on the temperatures at small radii is not a firm prediction, since our models neglect both radiative cooling and feedback heating, which are important in this region (Voit 2011).

conduction, these solutions are independent of halo mass.) In all cases, the temperature approaches the virial temperature near the scale radius of the halo, and falls by a factor of ~ 2 by the virial radius. We note that the temperature profiles in figure 4.6 are not monotonic with density and thus cannot be described with polytropic models. Interestingly, however, a polytrope with an index $\gamma \sim 1.2$ (as assumed by v03) provides a good fit between the scale radius and $\sim 0.5\, r_{\rm vir}$. These profiles are qualitatively very similar to the NFW models in figure 4.3 with $t_{\rm dyn}/t_0 \sim 0.15$, as is reasonable for clusters (§ 4.3.2). The location of the accretion shock is also consistent with our interpretation in section 4.3: Type III halos, which experience slower accretion at late times, have larger shock radii than the more rapidly accreting Type 0 halos. Thus, the intuition we developed in section 4.3 likely holds even for the more complex models in this section.

Ignoring thermal conduction is not a well-motivated approximation, however: as discussed in section 4.1, the timescale for heat to diffuse through massive clusters ($r_{\rm vir}^2/\chi_{\rm e} \sim 1\,$Gyr) is shorter than the typical age of the ICM ($\sim 5\,$Gyr). Consequently, non-cosmological simulations of isolated halos (e. g. Parrish et al. 2008) show that the ICM becomes almost completely isothermal after $\sim 2\,$Gyr. By analogy, one might therefore expect conduction to significantly modify the temperature profile shown in figures 4.3 and 4.6.

Figure 4.7 compares the temperature profiles of clusters with different effective conductivities. In order to maximize the influence of thermal conduction, we show $10^{15} M_\odot$ clusters (which are hotter, and thus more conductive than lower mass clusters), with Type III accretion histories (which formed comparatively early and thus provide more time for conduction to operate). As expected, thermal conduction smooths out the temperature profile in the ICM; the effect, however, is substantially weaker than has been found in non-cosmological simulations (see above). The profiles we obtain do *not* become isothermal, in qualitative agreement with x-ray observations of the ICM (George et al. 2009; Simionescu et al. 2011), and also with cosmological simulations of clusters which include thermal conduction (Dolag et al. 2004; Burns et al. 2010). The disagreement with simulations of isolated halos suggests that cosmological accretion and the continued formation of clusters is essential not only for understanding the origin of the large-scale temperature gradient in the ICM (§ 4.3), but also for how this gradient persists in spite of thermal conduction.

Two important effects differentiate cosmological calculations of thermal conduction from non-cosmological simulations of isolated halos. First, even though the present-day timescale for thermal conduction is shorter than the mean age of massive clusters, the gas near the accretion shock is always young and thus unaffected by conduction. The jump conditions at the virial shock therefore still determine the outer temperature, even when we take thermal conduction into account. The second effect differentiating cosmological and non-cosmological calculations is that the ICM at higher redshift had a lower temperature and thus a much lower conductivity than it has at redshift zero. Even at small radii, the age of the ICM ($\sim 5\,$Gyr) thus over-estimates the timescale over which thermal diffusion operates. These two effects, which are not present in non-cosmological simulations, strongly limit the influence of thermal conduction on the large-scale temperature profile of the ICM.

We illustrate these points in figure 4.8, where for each model from figure 4.7 we plot

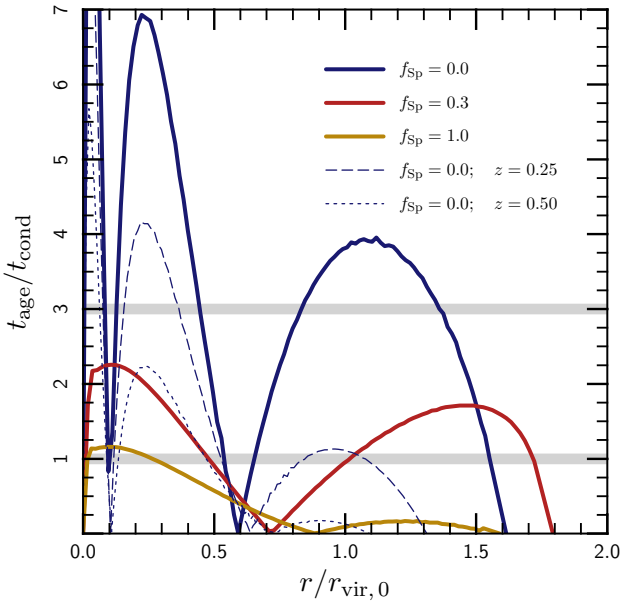

Figure 4.8: Ratio of the age of each shell in the ICM to its field-free conduction timescale (see text for definitions). The models are the same as in figure 4.7. This ratio is defined using the field-free conductivity (i. e. for $f_{\rm Sp} = 1$). Models with finite suppression factors should satisfy $t_{\rm age}/t_{\rm cond} \lesssim f_{\rm Sp}^{-1}$; the thick gray lines show this limit for $f_{\rm Sp} = 1$ and $1/3$. Adiabatic models predict $t_{\rm cond} \lesssim (0.2 - 0.3) t_{\rm age}$ at intermediate radii; as a result, conduction alters the temperature profile and maintains $t_{\rm age}/t_{\rm cond} \lesssim f_{\rm Sp}^{-1}$.

the ratio of the age of each shell in the ICM to its field-free conduction timescale. We define the conduction timescale as $t_{\rm cond} \equiv |d \ln K/dt|^{-1}$, calculated using equation 4.2 with $f_{\rm Sp} = 1$. We define the age of the shell $t_{\rm age}$ as the time over which conduction operates at its present-day efficiency; we approximate this as the minimum of the time since the shell accreted and $|d \ln m/dt|^{-1}$, the timescale over which the ICM conductivity changes appreciably. Conduction should limit the temperature profile of the ICM to everywhere satisfy $t_{\rm cond}/t_{\rm age} \lesssim f_{\rm Sp}^{-1}$. Thus, by looking at the adiabatic models in figure 4.8 (blue, with $f_{\rm Sp} = 0$), we can estimate how large of an effect conduction would have if we were to include it.

Figure 4.8 shows that the ratio $t_{\rm age}/t_{\rm cond}$ falls steeply near the virial shock, again reflecting that the recently accreted gas is too young to be influenced strongly by conduction. At smaller radii, however, $t_{\rm age}/t_{\rm cond} \sim 5$ in the adiabatic model, suggesting that thermal conduction plays a significant role in determining the temperature profile in the ICM. The thin blue curves in figure 4.8 show that this conclusion depends strongly on redshift, however. The flattening of the temperature profile shown in figure 4.7 is thus a fairly recent phenomenon and is far less pronounced at redshift $z \gtrsim 0.2$. Recall also that figures 4.7 and 4.8 show results for massive, $10^{15} M_\odot$ clusters with Type III accretion histories. Clusters with lower masses or different accretion histories are likely to be even less strongly influenced by conduction (see figure 4.10, below).

The red and yellow curves in figure 4.8 show that when we include thermal conduction with a given suppression factor $f_{\rm Sp}$, the temperature gradient adjusts so as to keep the ratio $t_{\rm cond}/t_{\rm age}$ below $f_{\rm Sp}^{-1}$. The ratio $t_{\rm cond}/t_{\rm age}$ is proportional to the gradient term:

$$\left[\frac{2}{x}\frac{\partial T_1}{\partial x} + \frac{5}{2T_1}\left(\frac{\partial T_1}{\partial x}\right)^2 + \frac{\partial^2 T_1}{\partial x^2}\right].$$

Thus, thermal conduction can either make the ICM isothermal (so that the heat flux vanishes), or make $T \sim r^{-2/7}$ (so that the heat flux is constant). We find that our model clusters initially take the second approach: the temperature profile adjusts so as to have a nearly constant heat flux as a function of radius. This is not possible near the origin, however, where spherical symmetry requires the heat flux to vanish. Thus, over a longer timescale, the center of the ICM cools and the ICM begins to become isothermal. The yellow curve in figure 4.7 shows the beginning of this process, but there is not enough time before redshift $z = 0$ for a significant fraction of the ICM to become completely isothermal, even in cases where conduction is most effective.

Though conduction may have considerable effect on the temperature profiles of massive clusters by redshift zero, its influence on their entropy profiles is less pronounced. We show this in figure 4.9, which compares entropy profiles for massive, $10^{15} M_\odot$ clusters with and without conduction. Between the scale radius and the virial radius, both models agree fairly well with a power-law fit derived from x-ray observations (e. g. Cavagnolo et al. 2009). Interestingly, thermal conduction smooths out some of the non-power-law behavior introduced by the accretion history, leading to an entropy profile more similar to the adiabatic self-similar profile.

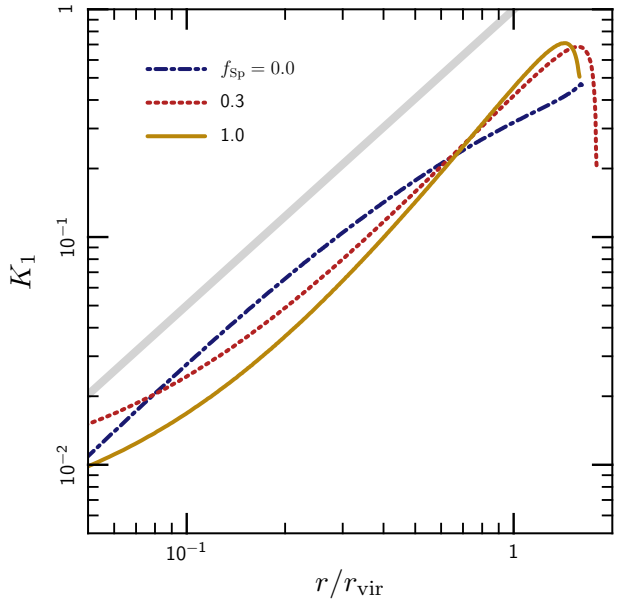

Figure 4.9: Entropy profiles in calculations with different values of the conductivity. As in figure 4.7, this figure shows $10^{15}\,M_\odot$ halos with Type III accretion histories; these properties maximize the effect of thermal conduction. Also shown is the $K \sim r^{1.3}$ power-law derived in v03.

Figures 4.6–4.9 illustrate the effects of cosmological accretion and thermal conduction on the temperature of the ICM in individual clusters. As mentioned in section 4.4.1, in order to study the statistics of temperature profiles as a function of virial mass, we run an ensemble of models with accretion histories drawn from the probability distribution in the appendix of McBride et al. (2009). The upper panels of figure 4.10 show these temperature profiles for three different halo masses with (yellow) and without (blue) thermal conduction. In the models with conduction, we have assumed a conductive suppression factor $f_{\rm Sp} = 0.3$. The temperature profile roughly follows the pattern found in section 4.3, but with fairly significant, ∼10%, scatter. This scatter is caused purely by the variation in accretion histories.[4] We have neglected several processes, including mergers, heating by dynamical friction, radiative cooling, and heating by AGN outflows, all of which may increase the scatter above that shown here.

Figure 4.10 also illustrates the mass-dependence of conduction's role in the ICM. Because the thermal conductivity of a plasma depends sensitively on its temperature as $\kappa \propto T^{5/2}$, higher-mass halos are more strongly influenced by conduction. More quantitatively, the conduction timescale in the ICM falls with mass: $t_{\rm cond} \sim r_{\rm vir}^2/\chi_{\rm e} \sim 1/M_{\rm vir}$. Figure 4.10 shows

[4]Note that we have neglected any correlation between halo concentration and accretion history, which may introduce additional variation in the temperature profiles (e. g. Bullock et al. 2001). The analysis in Zhao et al. (2009) suggests that a constant value for the concentration $c \sim 5$ is appropriate for massive clusters, however.

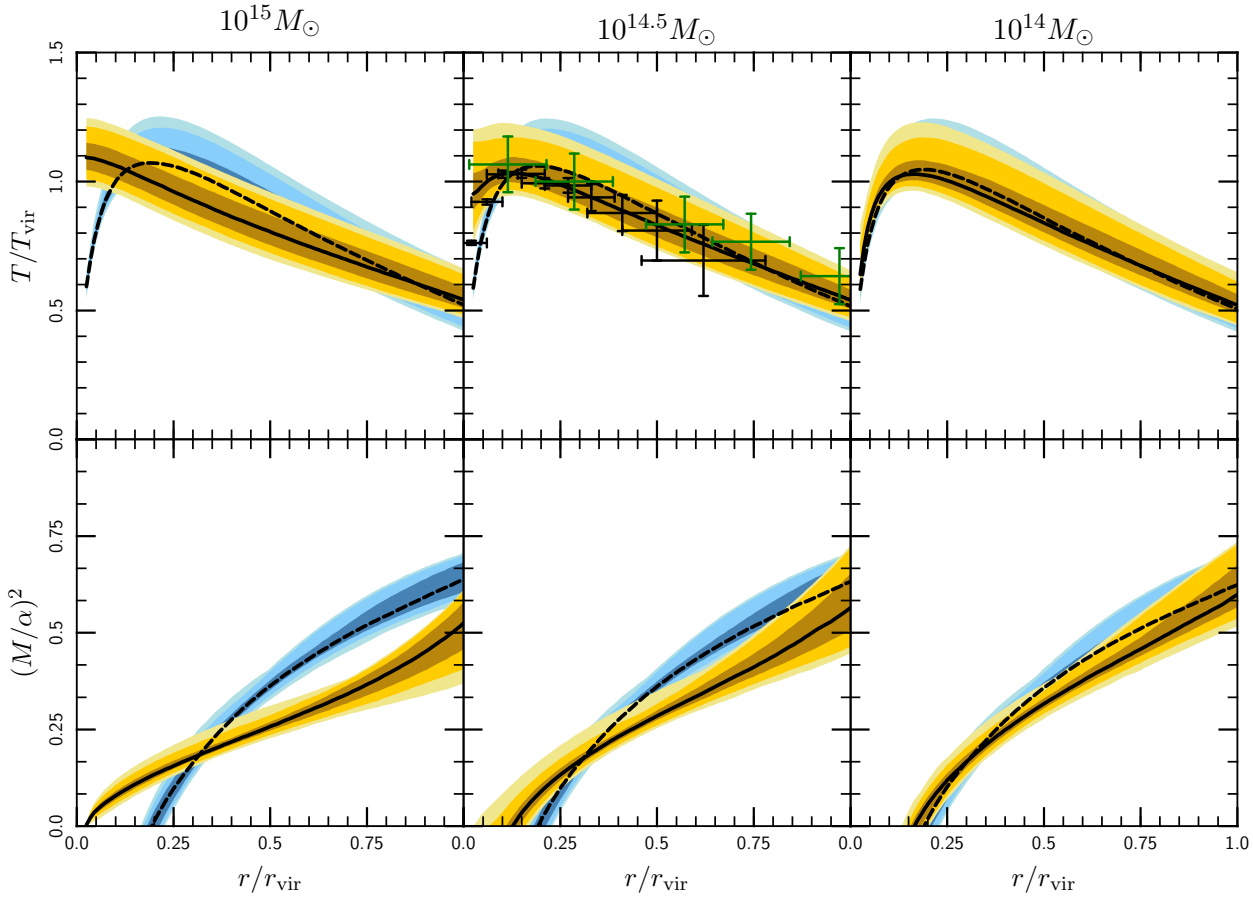

Figure 4.10: *(Top row):* Temperature profiles with (yellow; $f_{\rm Sp} = 0.3$) and without (blue) thermal conduction for different halo masses. Black lines show the median profile from our ensemble of accretion histories based on figure 4.4 and colors show contours enclosing 50%, 90%, and 99% of the models. The points with error bars in the top-right panel show temperature profiles derived from x-ray observations. The black and green points are taken from Leccardi & Molendi (2008) and Simionescu et al. (2011), respectively. *(Bottom row):* Estimated fraction of non-thermal pressure support generated by the MTI. See section 4.4.3 for details.

that the normalization is such that conduction has a minor influence on $10^{14} M_{\odot}$ halos,[5] but that it is significant for more massive, $10^{15} M_{\odot}$ clusters.

Finally, in the upper-right panel of figure 4.10, we have overlaid observational data from Leccardi & Molendi (2008) and Simionescu et al. (2011).[6] The data from Leccardi & Molendi (2008) (shown in black) represent an average temperature profile derived from 48 clusters from the XMM-Newton archive. The points from Simionescu et al. (2011), shown in green, represent SUZAKU observations of the Perseus cluster out to r_{200}. Outside the scale radius, where our method is applicable, our model agrees favorably with the data. Though this agreement is encouraging, we caution that both the normalization and slope of the data points in this plot are sensitive to the assumed virial mass of the halo, which is uncertain in x-ray observations. Furthermore, several uncertainties in our model, including the turnaround radius, the effect of non-smooth accretion, and deviations from spherical symmetry preclude a very quantitative comparison with the data.

4.4.3 Mixing-Length Theory and the MTI

As an example application of our models, we use mixing-length theory to estimate the turbulent pressure support produced by convection in clusters. In dilute, magnetized plasmas such as the ICM in galaxy clusters, convective stability depends on the temperature gradient of the plasma (Balbus 2001, later generalized by Quataert 2008 and Kunz et al. 2011). This convection, known as the MTI, may produce strong turbulence in clusters (McCourt et al. 2011b; Parrish et al. 2012b) and thus may provide enough non-thermal pressure support to bias hydrostatic mass estimates of cluster halos. Parrish et al. (2012b) found that the convective velocities produced by the MTI roughly obey mixing-length theory. This motivates us to use our model temperature profiles to estimate the turbulent pressure support produced by the MTI as a function of cluster mass and redshift. We note, however, that Kunz et al. (2012) have shown that the strength of this turbulence depends on the magnetic field strength in the ICM, an effect which we do not account for in our simple estimates.

Assuming that the convective motions retain their coherence for a fraction α of a pressure scale-height and that magnetic tension does not suppress the convective motions, we expect the instability to drive turbulent convection with Mach numbers of order:

$$M \sim \alpha \left(H \frac{d \ln T}{dr} \right)^{1/2}, \tag{4.17}$$

where $H \equiv (d \ln P / dr)^{-1}$ is the pressure scale-height. The bottom panels of figure 4.10 show this estimate for the turbulent pressure support (proportional to M^2). Assuming (as suggested by the simulations in Parrish et al. 2012b) that the mixing length parameter

[5] This conclusion applies to the large-scale temperature gradient in the ICM. Of course, small-scale features may be strongly influenced by conduction, even in $10^{14} M_{\odot}$ halos (cf. Dolag et al. 2004).

[6] The observation presented in George et al. (2009) provides another useful constraint on the properties of the ICM at large radii. We do not include it in figure 4.10, however, because it is not de-projected and thus not directly comparable to our results.

$\alpha \sim 0.5$, this figure suggests that turbulence driven by the MTI contributes $\sim 5 - 10\%$ of the pressure support outside of r_{500}. Cosmological simulations of cluster formation (e. g. Rasia et al. 2006; Nagai et al. 2007; Vazza et al. 2011) find significant turbulent pressure support due to subsonic, bulk flows driven by mergers near the virial radius. We note that any turbulence produced by the MTI would add to, but would likely be sub-dominant to, that produced by infalling subhalos.

Since the MTI is not suppressed by other sources of turbulence (McCourt et al. 2011b; Parrish et al. 2012b), this pressure support adds to that already present due to turbulence driven by infalling substructure or by galaxy wakes. This analysis suggests that the MTI plays an important role in the dynamics of the ICM, especially at radii $\gtrsim 0.5 r_{\rm vir}$. Thus, the MTI may provide an interesting correction for hydrostatic mass estimates of cluster halos. Unfortunately, both the large scatter in the strength of the MTI and its strong radial dependence (cf. Parrish et al. 2012b) seem to preclude a simple fitting function for the fraction of turbulent pressure support as a function of halo mass.

By applying mixing-length theory in equation 4.17, we have implicitly assumed that the MTI grows rapidly enough to establish convection by redshift zero. We find that $t_{\rm MTI} \lesssim 0.5 t_{\rm age}$ within $r_{\rm vir}$; thus, the Mach numbers in figure 4.10 (and also in Parrish et al. 2012b) are likely to be reasonable estimates.

Another assumption implicit in our use of mixing-length theory is that conduction is rapid enough to sustain the MTI. Thus, the results in the lower panels of figure 4.10 are only valid when the conduction timescale across an unstable mode $\sim (\alpha H)^2/\chi$ is less than the growth time of the MTI. Massive clusters ($M_{\rm vir} \gtrsim 10^{14.5}\ M_\odot$), likely satisfy this ordering of timescales, but lower masses halos ($M_{\rm vir} \lesssim 10^{14}\ M_\odot$) may not. Thus, the mass-dependence of the MTI is not likely to be monotonic. In the most massive halos ($M_{\rm vir} \gtrsim 10^{15}\ M_\odot$), thermal conduction is more efficient and weakens the temperature gradient at intermediate radii (fig. 4.10). In lower mass halos ($M_{\rm vir} \lesssim 10^{14}\ M_\odot$), on the other hand, conduction may not be fast enough to drive the MTI to its full potential. This analysis suggests that $10^{14.5}\ M_\odot$ halos experience the most vigorous convection driven by the MTI.

4.5 Discussion

This paper provides a simplified, spherically symmetric model for the temperature profiles of the hot plasma in galaxy groups and clusters. Our model is similar in spirit to earlier studies of the entropy profiles in clusters (e. g. Tozzi & Norman 2001; Voit et al. 2003), but builds on these earlier studies by focusing on temperature and by including the effects of thermal conduction. Our results agree reasonably well with the profiles derived from x-ray observations (cf. Leccardi & Molendi 2008) and from numerical simulations (e. g. Dolag et al. 2004).

We have shown that the large-scale temperature gradient in the ICM is primarily determined by the accretion history of its halo: while the gas near the center of a cluster reaches the virial temperature of the halo, the temperature at the virial shock is determined by the

ram pressure of the accreting gas. This difference sets the overall shape of the temperature profile. The timescale for thermal conduction ($t_{\rm cond} \sim r_{\rm vir}^2/\chi_{\rm e}$) is somewhat shorter than the age of the ICM. However, the influence of thermal conduction on the global temperature profile in clusters is mitigated for two reasons. The gas near the center of the cluster is less strongly effected by conduction because it was cooler in the relatively recent past. The gas near the virial shock, on the other hand, has only recently accreted and is younger than the conduction timescale. Thus, conduction has a diminished effect both near the center of the cluster and near the outskirts. As a result, it does not dramatically change the mean temperature profile (fig. 4.10). Of course, conduction can have a dramatic effect on small-scale inhomogeneities in the ICM (cf. Dolag et al. 2004); such inhomogeneities cannot be studied in our one-dimensional model.

Our results demonstrate the close relationship between the temperature gradient in clusters and the cosmological evolution of the host halo. This implies that numerical studies of isolated cluster models (e. g. Parrish et al. 2008) cannot correctly predict the evolution of the large-scale temperature profile, though they are very useful for studying other aspects of the ICM, such as the interplay among cooling, feedback, and plasma instabilities within the scale radius of the halo.

One of our motivations for studying the effects of conduction and halo accretion history on temperature gradients in clusters is that the free energy in the non-zero temperature gradient drives an efficient convective instability, the MTI. The results in section 4.4.3 (e. g. fig. 4.10) show that the turbulent pressure support generated by the MTI may be of order ~ 5 percent of the thermal pressure, and that it scales non-monotonically with halo mass. The magnitude of the turbulent pressure support is sensitive to the accretion history and does not seem amenable to a simple fitting formula.

Halo accretion histories have been studied extensively with numerical simulations. We use the fits to the Millennium simulation from McBride et al. (2009) to estimate the scatter in temperature profiles as a function of halo mass; at redshift $z = 0$, this scatter is of order 10%. This scatter likely contributes to the dispersion in cluster mass-observable relations relevant to x-ray and SZ observations. Perhaps more interesting are the effects of thermal conduction and convection, which introduce systematic changes to the temperature profile with mass. In particular, conduction smooths out the temperature profile (and decreases the peak temperature in the halo) by an amount that increases monotonically with halo mass (fig. 4.10); convection, on the other hand, produces turbulent pressure support that is non-monotonic in halo mass, peaking around $10^{14.5} M_\odot$ halos (§ 4.4.3).

Figures 4.6, 4.7, and 4.10 demonstrate that the effect of thermal conduction on a cluster's temperature profile is at least as large as the differences produced by normal variation in accretion histories. Any variation in the factor $f_{\rm Sp}$ (which parameterizes the suppression of the effective radial thermal conductivity relative to the field-free value), if it exists, would create additional scatter in the temperature profiles at fixed mass. Possible effects influencing $f_{\rm Sp}$ include magnetic draping around infalling or orbiting substructure (Dursi & Pfrommer 2008; Pfrommer & Dursi 2010) and the strength of the magnetic field. These processes may also contribute to the scatter in cluster mass-observable relations.

The models presented in this paper provide a simple explanation for the physics that sets the temperature profiles in galaxy groups and clusters at large radii. Our results are consistent with current observational constraints on cluster temperature profiles at large radii (fig. 4.10). They also highlight several processes which may bias hydrostatic mass estimates of clusters, including modifications to the scaling relation $T(M)$ by conduction and by turbulence driven by the MTI. The approximations in this paper, especially the assumptions of smooth accretion and spherical symmetry, preclude precise estimates of the non-thermal pressure support produced by the MTI. These limitations can be addressed using cosmological simulations.

Chapter 5

"Semi-"Cosmological Simulations of Cluster Convection

5.1 Introduction

Conditions in the ICM indicate an electron gyroradius which is smaller than the Coulomb mean free path by some 10 orders of magnitude; even allowing for the significant uncertainty in the inferred strength of the magnetic field, this separation in scales implies that ICM electrons are confined to move only along magnetic field lines. This motion, in turn, implies that the viscosity and thermal conductivity of the plasma are strongly anisotropic with respect to the magnetic field (Braginskii 1965). One of the consequences of this anisotropic transport is that the dominant convective instability switches from the usual Schwarzschild instability (which depends on the entropy gradient of the plasma) to the magnetothermal instability (MTI), which depends on the *temperature* gradient of the plasma (Balbus 2000; Quataert 2008). Thus, galaxy clusters, which were long thought to be convectively stable, are actually unstable to convection.

The analytic theory for the MTI was invented by Balbus (2000) and later refined by Quataert (2008), Balbus & Reynolds (2010), and Kunz et al. (2011). Parrish & Stone (2005, 2007) used idealized, local simulations to confirm the instability numerically and to study its nonlinear saturation. Initially, the MTI seemed likely to drive strong convective turbulence in galaxy clusters, potentially mixing metals throughout the ICM, amplifying the intracluster magnetic field, and producing enough turbulent pressure support to bias hydrostatic mass estimates of clusters (e. g. Quataert 2008).

Parrish et al. (2008) and Sharma et al. (2008) first studied the MTI in global simulations of galaxy clusters and of the galactic center, respectively. Surprisingly, both studies found that the MTI did *not* grow to large amplitude: rather than driving fully-developed turbulence, the instability instead saturated "quasi-linearly," with certain properties such as the final magnetic field strength still determined by the initial condition. Parrish et al. (2008) identified the saturation channel for the MTI in their simulation: thermal conduction rendered the ICM isothermal on a timescale short compared to the dynamical time, effectively

shutting off the MTI before it had a chance to grow.

McCourt et al. (2011b) pointed out that observed clusters do not seem to be isothermal, and speculated that an unknown process maintains the large-scale temperature gradient in clusters. Thus, the isothermal saturation channel might not be available to clusters, forcing MTI to drive convection. Consequently, Parrish et al. (2012b) ran a new set of simulations with different boundary conditions and showed that the MTI may in fact drive vigorous turbulence, as had been thought initially. The implications of the MTI are thus not obvious a priori, and seem to depend on other processes controlling the structure of the ICM.

A significant shortcoming of existing work on the MTI is that the timescale for the MTI to grow is only a factor of $\sim$10 shorter than the age of a cluster. The time-evolution of the temperature gradient and gravitational potential may thus alter the growth of the instability in an important way. To date, however, nearly all numerical studies of the MTI in galaxy clusters have utilized simulations of "isolated" halos which do not evolve in time. More recently, McCourt et al. (2013) studied simplified models of forming clusters and found that thermal conduction has a much smaller influence on the cluster plasma than simulations of isolated halos predict. One might expect that the effect of the MTI is similarly diminished; simulations with an evolving potential and a realistic virial shock are needed to fully assess the importance of the MTI for cluster astrophysics and cosmology.

McCourt et al. (2013) provided a simplified model for the formation of galaxy clusters and showed that the continued assembly of clusters mandates a large-scale temperature gradient similar to those observed. In part, this justifies the speculation in McCourt et al. (2011b) and the choice of boundary conditions in Parrish et al. (2012b). However, the processes identified by McCourt et al. (2013) which moderate the influence of thermal conduction on the ICM might similarly suppress the MTI – these processes rely on the cosmological evolution of the cluster and, unfortunately, were not present in the simulation of Parrish et al. (2012b). In this chapter, I present a set of simulations which aims to determine whether the MTI grows to large amplitude in galaxy clusters, taking into account the competition between the cosmological assembly of the cluster and the growth of the instability in an idealized model.

5.1.1 Plan

Simulating the MTI requires solving the full Braginskii equations, along with the cosmological physics which determines the large-scale structure and evolution of the ICM. This chapter presents idealized simulations which abstract away as much of the cosmological physics as possible, while still keeping the aspects which seem essential for understanding the present-day properties of galaxy clusters.

I have tried as much as possible to extend the spherically symmetric models of McCourt et al. (2013) and Voit et al. (2003) to full 3D numerical simulation. In particular, I treat the baryonic evolution of the cluster as separate from the dark matter dynamics. I model the dark matter using a time-dependent gravitational potential fit to the "NFW" form (Navarro et al. 2004) with a mass accretion history taken from the fits in McBride et al. (2009). I assume

the baryonic accretion rate tracks the dark matter and apply a boundary condition on the gas density at the "turnaround radius" (see § 5.2 for more details). The gas self-consistently finds a virial shock, which determines the thermal state of the intracluster medium. These cosmological approximations closely follow those in McCourt et al. (2013) and in Voit et al. (2003), but can now be applied to multi-dimensional numerical simulations to study the MTI in a cosmologcail context.

These simulations will also enable me to assess the importance of viscosity in the dynamics of the ICM. While Kunz et al. (2011) pointed out that ignoring viscosity is not formally justified in the linear theory, the practical importance of viscosity is not yet known. The two studies to date (Parrish et al. 2012a and Kunz et al. 2012) reach somewhat different conclusions and the results appear to depend on the details of the simulation setup. By self-consistently evolving both the cluster and the MTI, I aim to remove much of the arbitrariness in the simulation setup.

5.2 Method

As discussed in § 5.1.1, I aim to repeat the calculation of McCourt et al. (2013) as closely as possible, but relaxing the assumptions of hydrostatic equilibrium and spherical symmetry in the gas dynamics. In addition to providing a more realistic simulation of the MTI, the simulations in this chapter will also serve to test those idealized, one-dimensional models of galaxy cluster assembly.

I use the non-cosmological MHD code ATHENA, which solves the non-relativistic fluid equations in physical (i. e. not co-moving) coordinates. The code does not solve for dark matter dynamics; instead, I model the dark matter via an external potential as a function of radius and time. At each time-step in the MHD simulation, I find the "turnaround radius" (assumed to be twice the virial radius) and I calculate the density and velocity of the gas there from the accretion history of the cluster. I apply this density and velocity as a boundary condition on the calculation, overwriting whatever values already exist on the grid. This enables me to apply a moving boundary condition within the fixed grid structure of ATHENA. I illustrate this procedure schematically in figure 5.1 and I show an example computational setup in figure 5.2.

I thus model the cosmological context of the cluster via a time-dependent potential and a boundary condition. Conceptually, this separates the gas dynamics of the cluster from the detailed cosmological formation, while keeping the aspects of cosmology which are essential for the problem. It is of course possible that the detailed cosmological evolution of the cluster matters – angular momentum, tri-axiality of the potential, non-smooth accretion could all matter in detail. I plan to test this in the near future by comparing my results with fully cosmological simulations (e. g. Ruszkowski et al. 2011 or Smith et al. 2013).

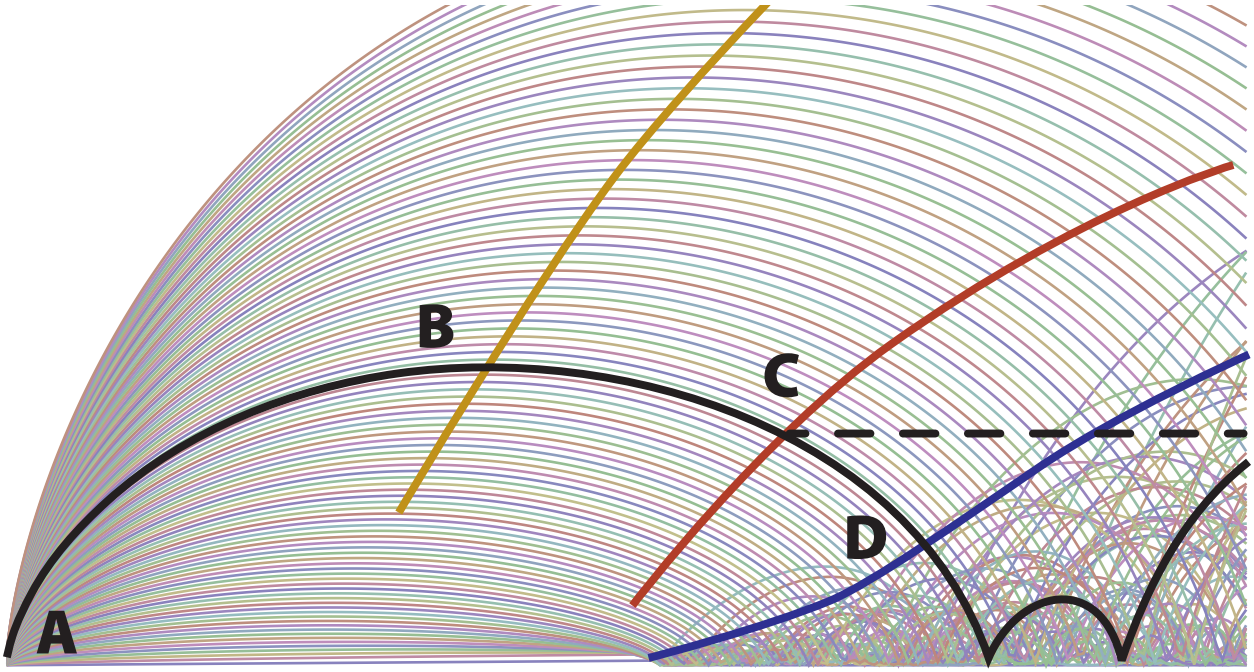

Figure 5.1: Schematic formation of a galaxy cluster. The highlighted shell initially expands from the big bang at point 'A' as $r \sim t^{2/3}$ before turning around at point 'B' due to gravitational instability. As the shell falls in toward the cluster, the gas shocks at point 'C' and sits in hydrostatic equilibrium at roughly constant radius (dashed line). The dark matter falls through the virial shock until it reaches the virial radius at point 'D,' after which shell crossings become important. The dark matter then begins to orbit and makes up the halo of the cluster. The yellow, red, and blue curves show the 'turnaround,' shock, and virial radii as functions of time. In this work, I assume the dark matter dynamics as specified and I solve the gas dynamics within the evolving potential due to the dark matter. I use the NFW fit to the dark matter distribution (Navarro et al. 2004) and fits to the mass accretion histories from McBride et al. (2009). I furthermore assume that the turnaround radius is twice the virial radius at all times, which specifies the dark matter dynamics completely. Within this model, I solve the gas dynamics in 3D including MHD with anisotropic conduction and viscosity to determine the shock radius and the thermal state of the intracluster medium self-consistently.

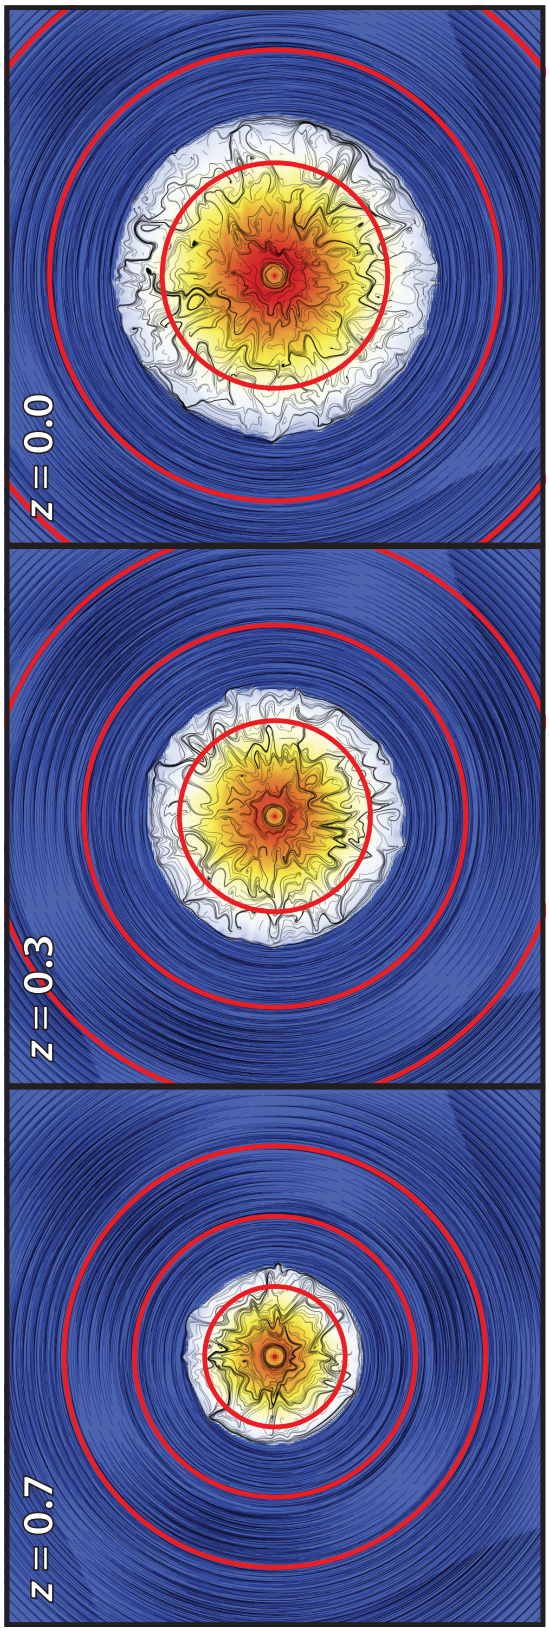

Figure 5.2: Example simulation showing the entire computational domain at several different times. Color shows temperature and black lines trace the magnetic field. The red circles denote r_{200}, $2\,r_{200}$, and $3\,r_{200}$. The middle circle represents the computational boundary (but not the edge of the grid); at each time-step, I overwrite the gas variables outside this sphere using a top-hat collapse model as described in § 5.2.2.

5.2.1 Potential and Mass Accretion History

As described in McCourt et al. (2013), I assume the dark matter follows an NFW (Navarro et al. 2004) distribution at all times with a constant concentration parameter $c = 5$. This choice of a constant concentration parameter is appropriate for a cluster still undergoing active stages of formation (Zhao et al. 2009; Wu et al. 2013).

I adopt the fits to cluster accretion histories from McBride et al. (2009) and assume the mass evolves with redshift as $M(z) \propto [(1+z)^b e^{-z}]^g$. McBride et al. (2009) provide a probability distribution for the parameters g and b and classify accretion histories into distinct "types." For simplicity, I assume a Type I accretion history represents a "typical" cluster, while a Type 0 accretion history represents a young (late-forming) cluster and a Type III accretion history represents an old (early-forming) cluster. This is roughly correct, though "Type" is not one-to-one with age (see figure 4 in McCourt et al. 2013).

5.2.2 Boundary Condition

I derive my outer boundary condition by assuming the dark matter far outside the virial radius follows a simple top-hat collapse model in which shells of dark matter decouple from the hubble flow and "turn around" at a given radius before collapsing to form the cluster (e. g. Gunn & Gott 1972; Padmanabhan 1993). See, e.g. Tozzi & Norman (2001), Voit et al. (2003), or McCourt et al. (2013) for a more detailed description of this boundary condition.

In the frame of the cluster, the radial velocity vanishes at turnaround. By fixing the accretion rate through this radius ($= 4\pi r_{\rm ta}^2 \rho_{\rm ta} \times d r_{\rm ta}/dt$) to the mass accretion history of the cluster, I derive the density at turnaround as a function of time:

$$\frac{\rho_{\rm ta}}{200\rho_{\rm crit}} = \frac{1}{8}\frac{\ln 11}{\ln 6}\left[1 + \frac{3\Omega_{\rm m}}{g(1+z-b)}\frac{(1+z)^3}{\Omega_{\rm m}(1+z)^3+\Omega_\Lambda}\right]^{-1}, \tag{5.1}$$

illustrated in figure 5.3. Physically, this corresponds to the initial density perturbation which collapses to form the cluster.

I have described how I calculate the density and velocity of the gas at turnaround using a top-hat collapse model. I still need to apply a boundary condition on the temperature and the magnetic field. I choose an outer temperature which is low enough to guarantee a strong virial shock with $M \gg 1$. The results are thus independent of the outer temperature. I assume the magnetic field is dynamically weak ($\beta \gg 1$) and is initially arranged in concentric loops about the "z" axis of the cluster. Though this magnetic field configuration is highly idealized, it introduces no free parameters into the setup and it highlights the role of the MTI in my simulations. Previous simulations (Parrish et al. 2008; McCourt et al. 2011b) suggest that the evolution of the MTI is not strongly dependent on the initial geometry of the magnetic field. Prior to turning this chapter into a published paper, I will extend the results presented here to include initially tangled magnetic fields.

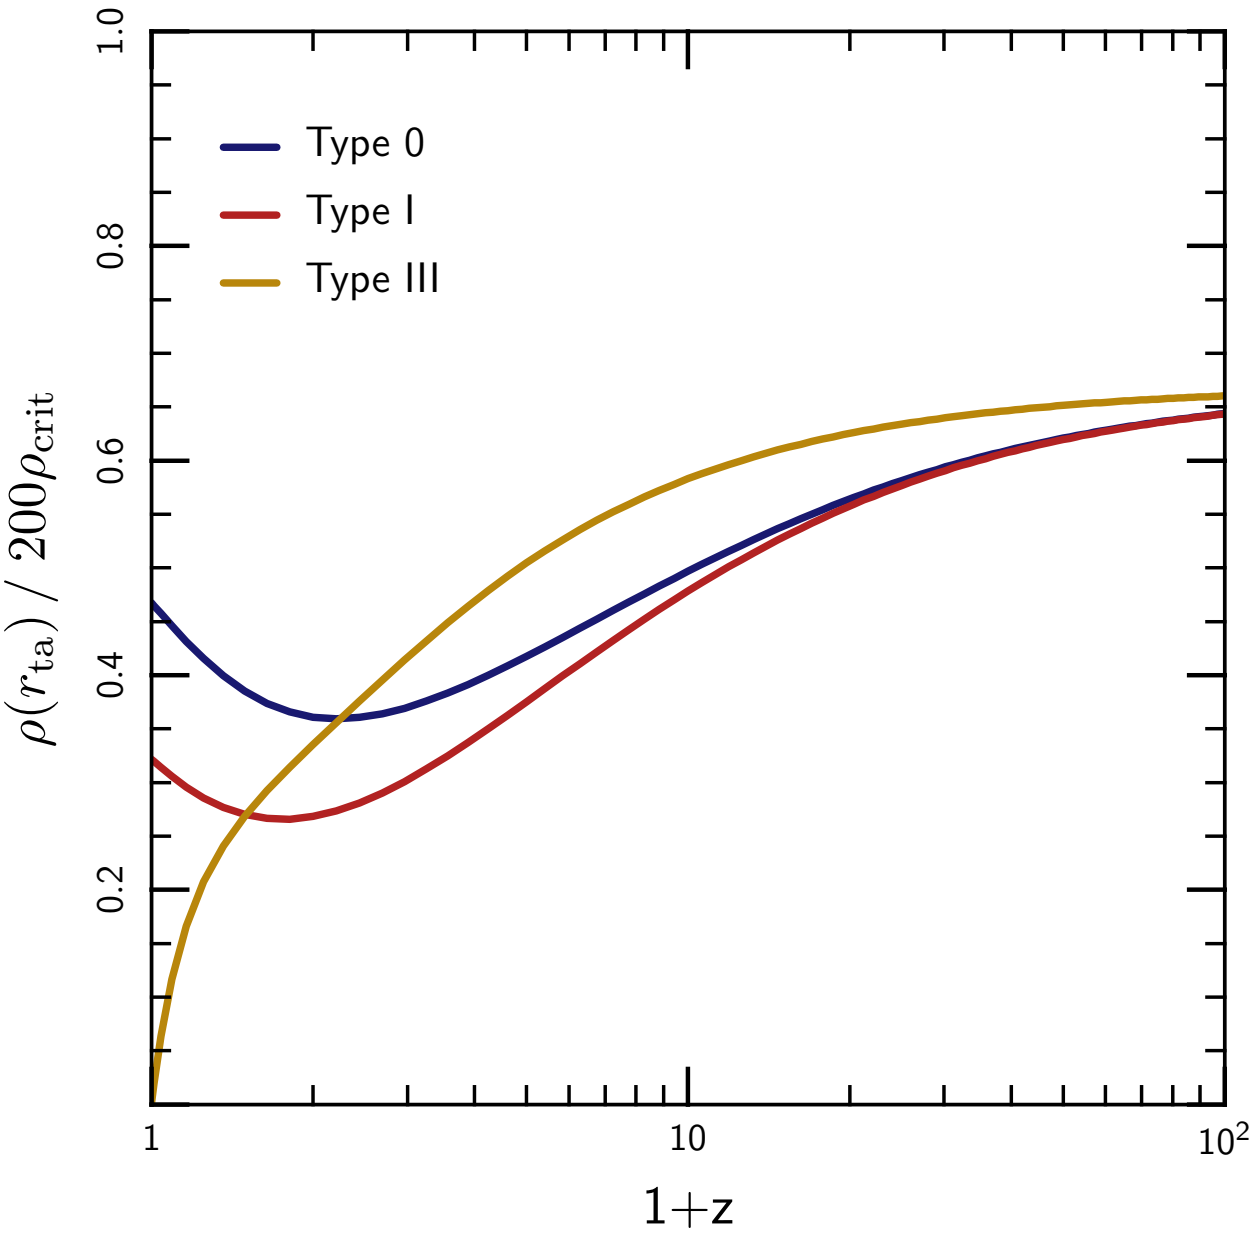

Figure 5.3: Density at turnaround as a function of redshift for three different models for the cosmological growth of cluster mass halos (which roughly bracket the range of accretion histories found in cosmological simulations such as Millenium; see McBride et al. 2009).

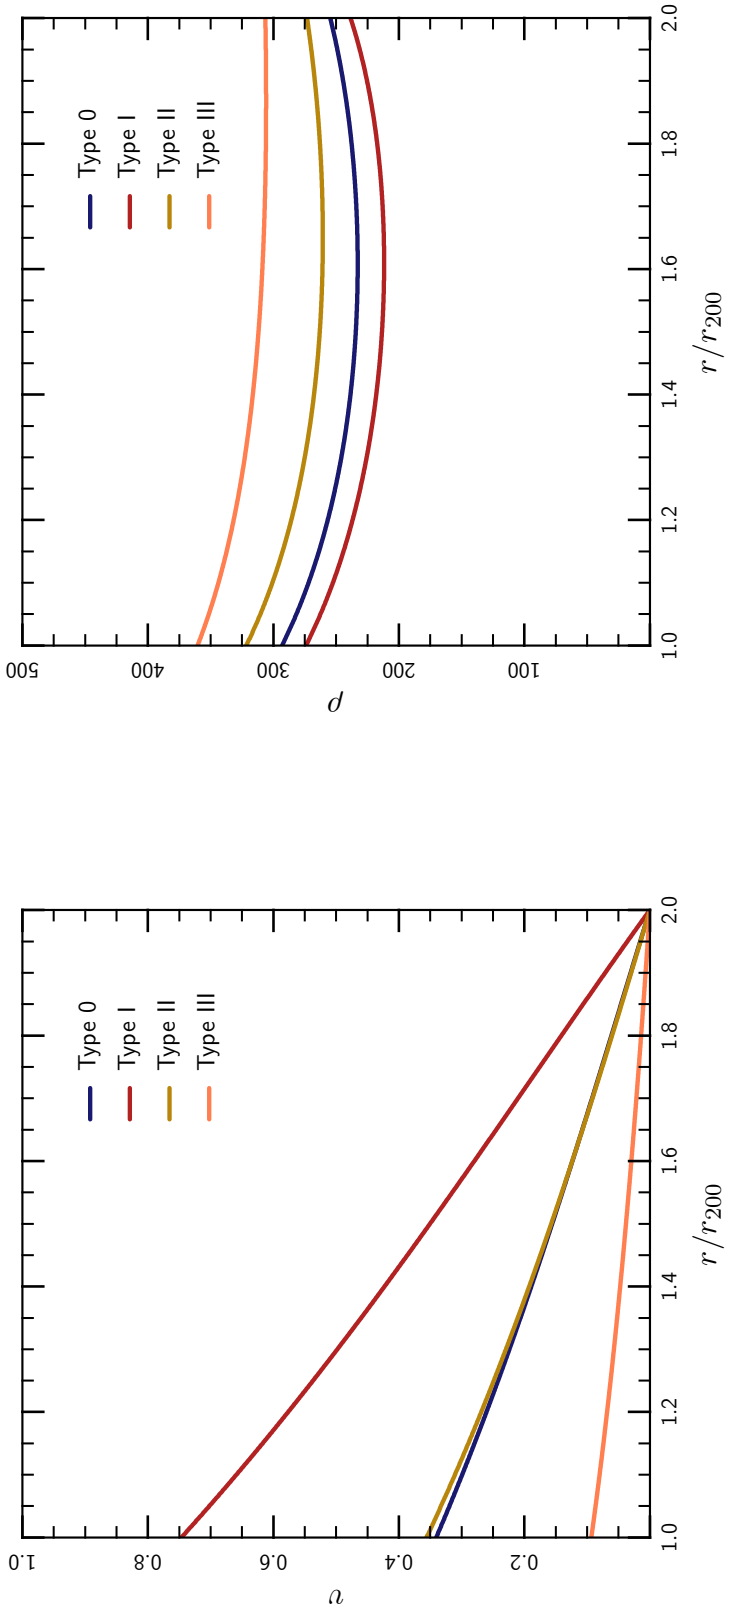

Figure 5.4: Initial condition for density and velocity between the turnaround radius and the virial radius. I treat the gas exterior to the turnaround radius as a boundary condition. I initialize the gas inside the virial radius in hydrostatic equilibrium with the potential of the dark-matter halo. See section 5.2 for details.

5.2.3 Initial Condition

I begin my simulations at redshift $z = 6$, with the density and velocity at turnaround as described above. For the initial condition, I find the density and velocity of the gas inside turnaround by extending the top-hat collapse model. Specifically, I initialize a series of shells filling the volume between $r_{\rm ta}$ and r_{200} and I integrate the equation of motion:

$$\frac{dv}{dt} = -\nabla\phi(r,t) \tag{5.2}$$

back in time to find when the shells turned around. Assuming that shell-crossing is negligible outside r_{200}, the mass enclosed by a given shell is a constant in time and equation 5.2 reduces to an ordinary differential equation. I then use the density at turnaround at that time, along with the change in volume between adjacent shells, to determine the density and velocity of the gas between $r_{\rm ta}$ and r_{200}. Figure 5.4 shows example solutions to this equation.

Since this method breaks down inside r_{200}, I model the gas inside this region with a hydrostatic atmosphere separated from the infalling gas by a shock. Though my choice of a shock radius at r_{200} is somewhat arbitrary, I have confirmed my results are insensitive to this choice provided I start the simulations early enough ($z_{\rm initial} \gtrsim 4$).

Inside the shock, I assume the initial state is a hydrostatic and isentropic core, with

$$\frac{\rho}{\rho_{\rm core}} = \left[1 + A\left(\frac{\phi}{\phi_{\rm core}} - 1\right)\right]^{3/2}$$
$$P = K\left(\frac{\rho}{\rho_{\rm core}}\right)^{5/3}$$

where $A = (6|\phi_{\rm core}|)/(5v_{\rm i,\ core}^2)$, $K = \frac{1}{3}v_{\rm i,core}^2(4\rho_{\rm core})^{-2/3}$, and $\rho_{\rm core}$ and $v_{\rm i,core}$ represent the pre-shock density and velocity in the frame of the cluster.

5.2.4 Numerical Method

I use a method very similar to the one described in Parrish et al. (2012a) to determine the Braginskii conductive heat flux and viscous momentum flux. Since the closure to the fluid equations is not known in this quasi-collisionless limit, however, I modified the method to optionally implement the "Sharma" closure from Sharma et al. (2006), which restricts the pressure anisotropy to lie within the bounds for the firehose and mirror instabilities. In the text below, I refer to these simulations as using the "Sharma" closure, as opposed to the "Braginskii" closure.

Quantitatively, I re-express the Braginskii stress tensor in terms of the pressure anisotropy $\Delta P \equiv P_\perp - P_\parallel$:

$$\mathbf{T}^{\rm (vis)} = \Delta P\left(\frac{1}{3}\mathbf{g} - \hat{\boldsymbol{b}}\otimes\hat{\boldsymbol{b}}\right), \tag{5.3}$$

where $P_\parallel$ is the gas pressure along magnetic field lines, and $P_\perp$ is the pressure in orthogonal directions. In a coordinate system with magnetic field lines along the $\hat{\boldsymbol{z}}$ axis, this stress tensor takes the form:

$$\mathbf{T} = \mathbf{T}^{(\mathrm{vis})} + P\mathbf{g} = \begin{pmatrix} P_\perp & 0 & 0 \\ 0 & P_\perp & 0 \\ 0 & 0 & P_\parallel \end{pmatrix} \tag{5.4}$$

Individual particles contain an adiabatic invariant $v_\perp^2/B$ which is approximately conserved. Thus $P_\perp/B \sim$ constant, and changes in the magnetic field strength induce a finite pressure anisotropy ΔP. One can understand the Braginskii closure used in Parrish et al. (2012a) loosely by balancing this driving with relaxation by collisions:

$$\left(\frac{\Delta P}{P}\right)_{\mathrm{brag}} \sim \frac{1}{\nu_{\mathrm{ii}}} \frac{d \ln B}{dt}, \tag{5.5}$$

where $\nu_{\mathrm{ii}} \sim c_{\mathrm{s}}^2/\nu$ is the ion-ion collision frequency.

However, a non-zero ΔP represents a source of free energy in the plasma which drives velocity-space instabilities; because these instabilities are not captured in the fluid limit, their implications for my simulations are somewhat uncertain. One possibility, suggested by Sharma & Hammett (2007) and recently bolstered by particle-in-cell (PIC) simulations (Kunz et al. 2014; Riquelme et al. 2014) is that these instabilities induce anomalous pitch-angle scattering, effectively making the plasma more collisional. Thus, I follow Sharma & Hammett (2007) and limit ΔP to be within the bounds for the firehose, mirror, and ion-cyclotron instabilities in some of my simulations. I show below that the choice of closure has only a modest effect on the MTI.

5.3 Results

Figure 5.5 compares the results of adiabatic simulations with the simplified, one-dimensional models from chapter 4. The entropy profile shows that the simulations reproduce both the strength and location of the virial shock as determined in the one-dimensional calculations. More impressive is the agreement between the temperature profiles: as I discussed in the previous chapter, this temperature profile depends on the entire accretion history of the galaxy cluster. This test suggests that my simulation method works surprisingly well; below I use it to study the nonlinear evolution of the MTI in galaxy clusters.

Figure 5.6 compares simulations with isotropic and anisotropic thermal conduction. These are two-dimensional slices through the midplane of full, three-dimensional simulations. Black lines show the magnetic field, and color shows the temperature fluctuation $\delta T/T = (T - \langle T \rangle)/T$, where $\langle \ldots \rangle$ denotes an average on spherical shells. Even at the relatively low mass of $10^{14}\,M_\odot$, isotropic conduction strongly suppresses temperature perturbations; as expected, this suppression becomes stronger in hotter, more massive clusters. Anisotropic conduction permits temperature gradients orthogonal to the magnetic field lines.

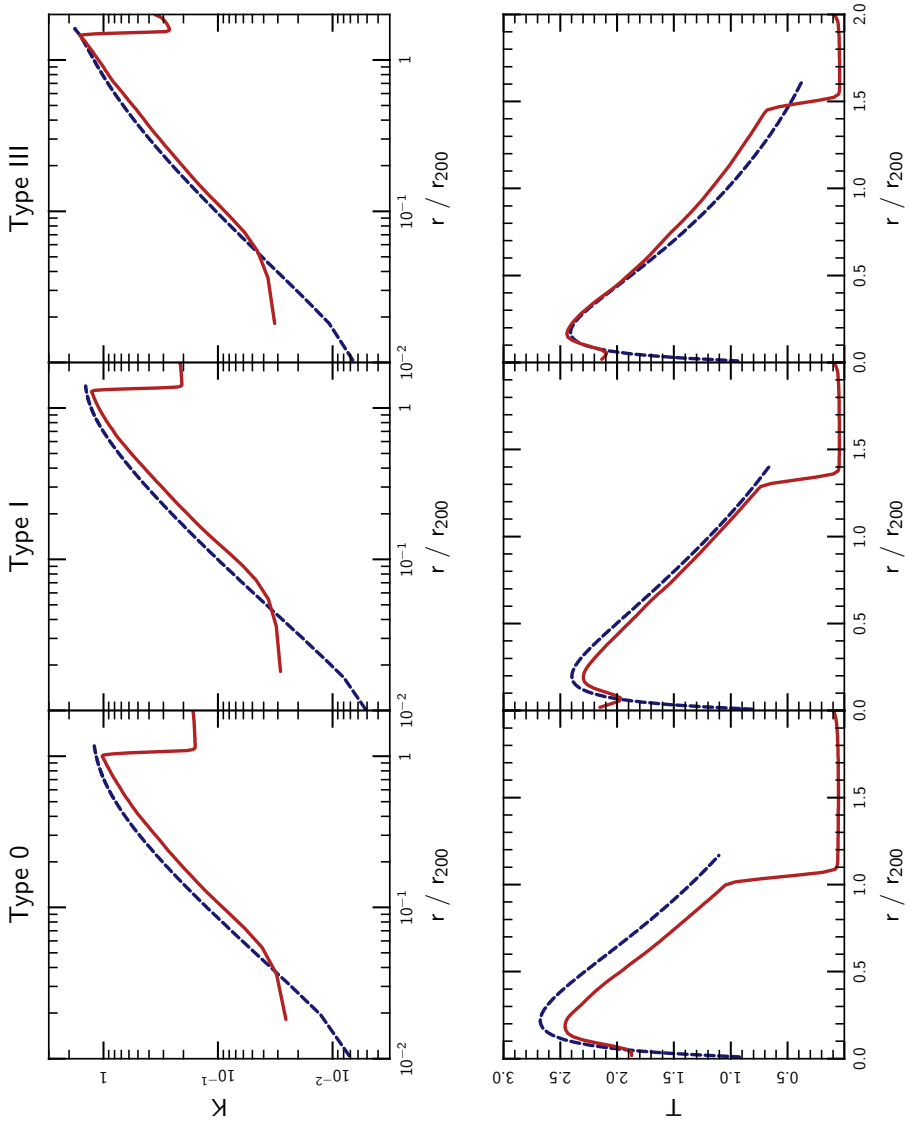

Figure 5.5: Comparison of my three dimensional adiabatic simulations (solid red lines) with the simplified analytic one-dimensional models of cluster profiles from chapter 4 (dashed blue lines). I show profiles of entropy ($K \equiv P/\rho^{5/3}$; top) and temperature (bottom) for three example accretion histories (Type I-III). Units are such that the Hubble constant, Newton's gravitational constant, and the mass of the cluster $H_0 = G = M_0 = 1$. The virial shocks in my simulations have very similar radii and strengths as predicted in the simplified one-dimensional models. The temperature profiles, which are sensitive to the evolution of the cluster as a whole, also agree very well.

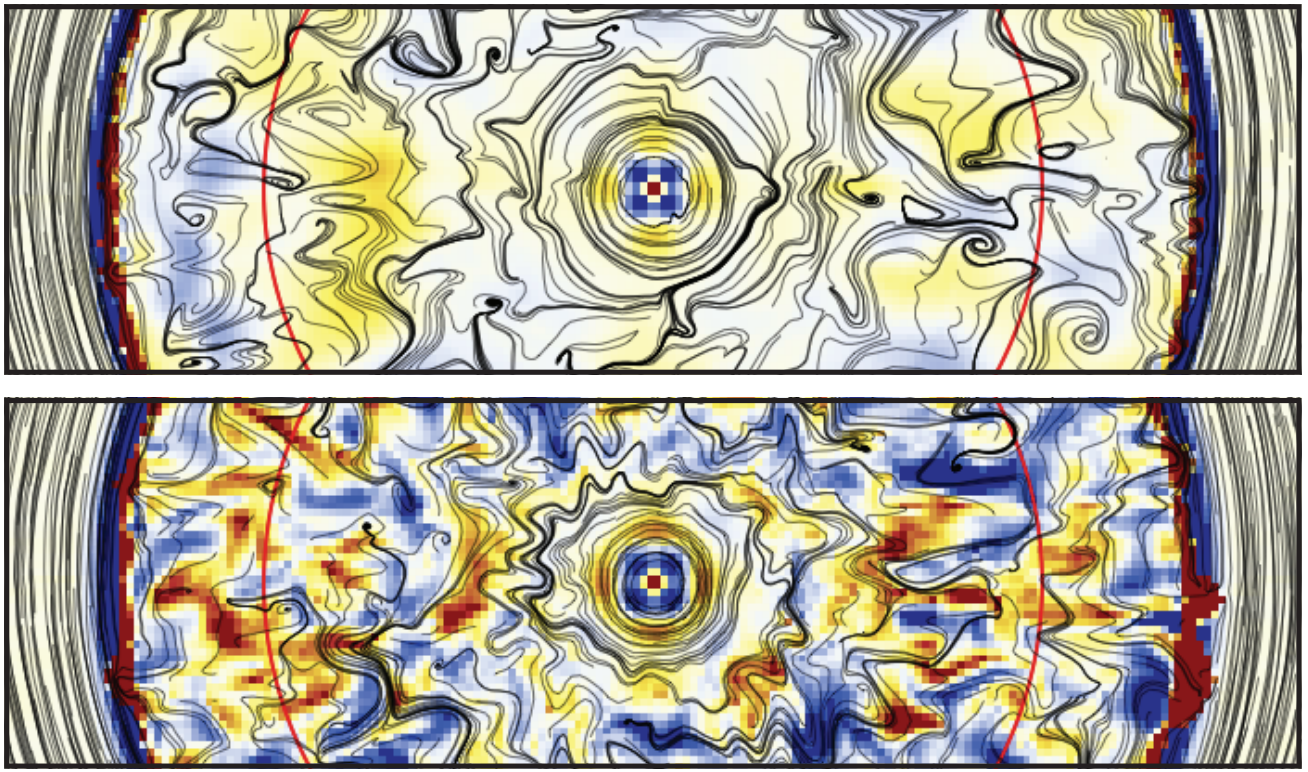

Figure 5.6: Comparison of simulations with isotropic (top) and anisotropic (bottom) thermal conduction. This figure shows 2D slices through the centers of 3D simulations. Black lines show the projected magnetic field, and the solid red circle shows the virial radius r_{200}. Color shows the temperature fluctuation $\delta T/T$; the scale ranges from -0.05 (deep blue) to +0.05 (deep red). Even at the relatively low mass of $10^{14}\,M_\odot$, isotropic conduction strongly suppresses temperature perturbations. Anisotropic conduction, on the other hand, permits temperature gradients orthogonal to the local magnetic field; moreover, these fluctuations are amplified by the MTI (see figure 5.7, below).

Rather than diffusing them away, thermal conduction actually amplifies temperature fluctuations in these simulations; this is the action of the MTI. See figure 5.7, below.

Figure 5.7 shows temperature fluctuations as a function of halo mass for simulations with both isotropic and anisotropic thermal conduction. From top to bottom, this figure shows simulations of 10^{14}, 3×10^{14}, and $10^{15}\,M_\odot$ clusters. Within each panel, the left half shows a simulation with anisotropic conduction and the right half shows the corresponding simulation with isotropic conduction. As in figure 5.6, black lines show the projected magnetic field, and color shows the temperature fluctuation $\delta T/T$. Simulations with isotropic conduction behave intuitively: increasing the halo mass increases the effective conductivity of the ICM, which reduces the amplitude of the temperature fluctuations. In the simulations with anisotropic conduction, however, this trend reverses: the amplitude of the fluctuations *increases* with increasing halo mass. This is the result of amplification due to the MTI, which becomes more vigorous with increasing conductivity. Note that conduction is strong enough in $10^{15}\,M_\odot$ clusters to significantly change the dynamics of the ICM. In the isotropic simulation, the conductive heat flux has inflated the ICM at large radii, pushing the virial shock out to the computational boundary at $2 \times r_{200}$; this does not occur in my simulation with anisotropic conduction because the initially azimuthal field inhibits the radial transport of energy on large scales.

Figure 5.8 compares an inviscid simulation with a simulation which includes anisotropic viscosity using the Braginskii closure. Viscosity smooths out the flow on small scales and appears to organize the magnetic field lines into more coherent, large-scale plumes. Interestingly, however, this has only a small effect on volume-integrated quantities such as the turbulent Mach number or the convective heat flux; see figure 5.10, below. The simulations with Braginskii viscosity exhibit strong bunching of the magnetic field lines; this may be a form of the firehose instability. Figure 5.9 shows that this effect essentially disappears from simulations using the Sharma closure.

Figure 5.10 quantifies the growth of the MTI in my semi-cosmological cluster simulations. The left panel shows the turbulent Mach number, defined as $\sqrt{\langle [(v_r - \langle v_r \rangle)/c_\mathrm{s}]^2 \rangle}$, where $\langle \ldots \rangle$ represents an average over a spherical shell. The mean Mach number can reach $\gtrsim 0.3$ near the virial radius in massive clusters, and appears to be only weakly influenced on viscosity. The right panel shows the average $\langle (\delta T)(\delta v_r)/c_s^3 \rangle$; this quantity is similar to the convective heat flux, but corresponds more directly to the linear growth of the MTI. As in the left panel, viscosity has a negligible effect on the MTI in $10^{14}\,M_\odot$ clusters, but becomes more prominent by $10^{15}\,M_\odot$.

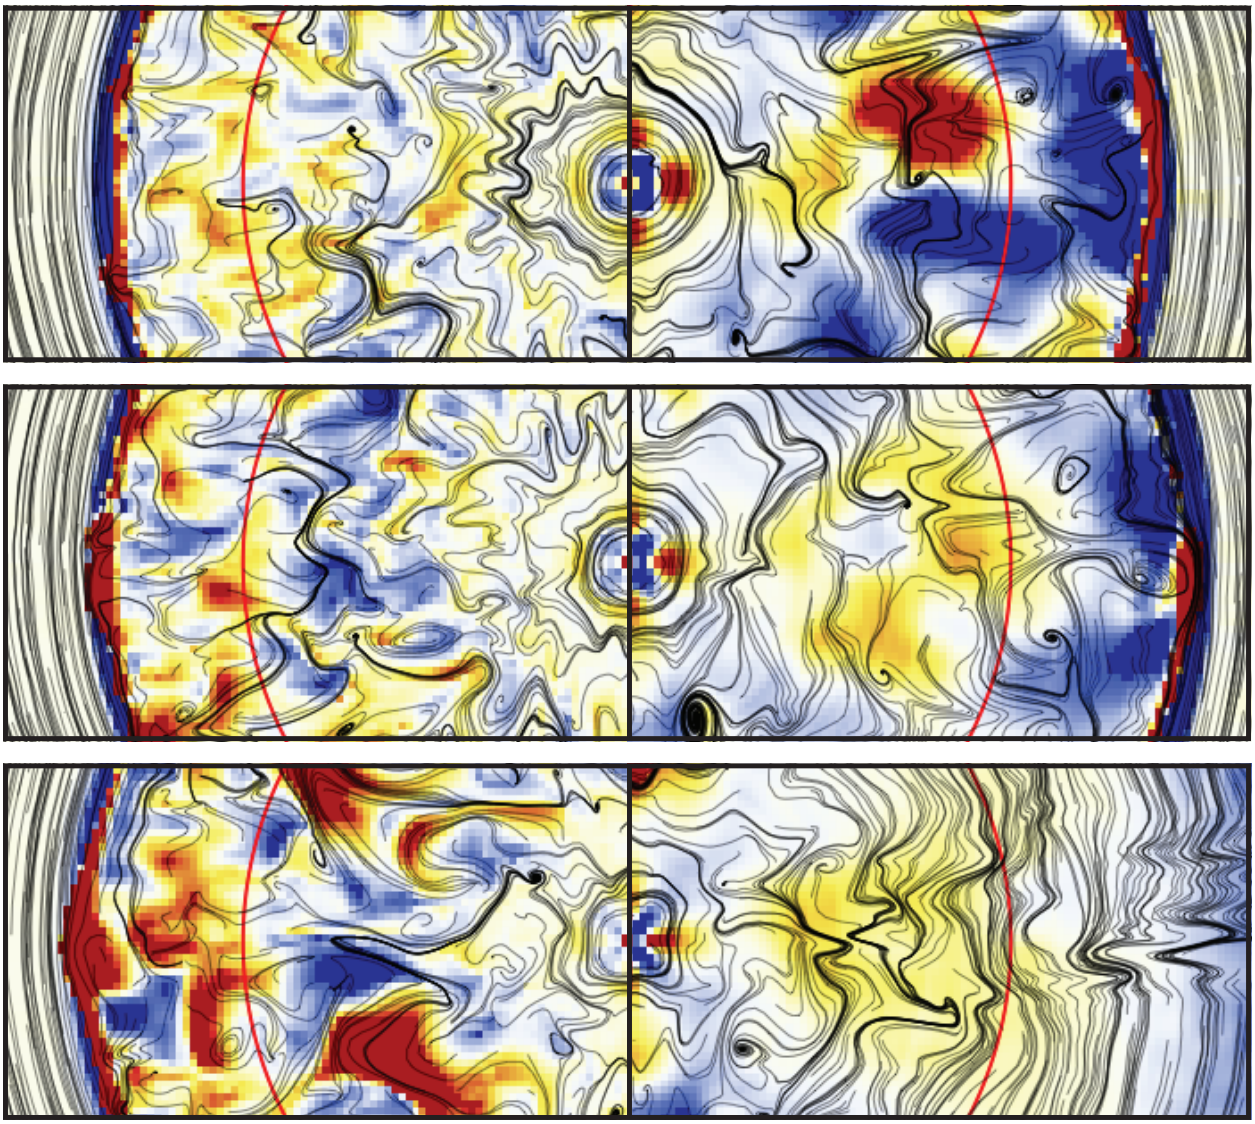

Figure 5.7: (Caption on next page)

Figure 5.7: (Previous page.) Temperature fluctuations as a function of halo mass. From top to bottom, this figure shows simulations of 3×10^{13}, 10^{14}, and $3 \times 10^{14}\,M_\odot$ clusters. Within each panel, the left half shows a simulation with anisotropic conduction and the right half shows the corresponding simulation with isotropic conduction. As in figure 5.6, black lines show the projected magnetic field, and color shows the temperature fluctuation $\delta T/T$. The color scale on the left ranges from -0.05 (deep blue) to +0.05 (deep red); the color scale on the right is a factor of five smaller. Note that the color scale saturates in the top-right and bottom-left panels; the maximum temperature fluctuations in those simulations significantly exceed the ranges plotted here. Simulations with isotropic conduction behave intuitively: increasing the halo mass increases the effective conductivity of the ICM, which reduces the amplitude of the temperature fluctuations. In the simulations with anisotropic conduction, however, this trend reverses: the amplitude of the fluctuations *increases* with increasing halo mass. This is the result of amplification due to the MTI, which becomes more vigorous with increasing conductivity. Note that conduction is strong enough in $10^{15}\,M_\odot$ clusters to significantly change the large-scale dynamics of the ICM. In the isotropic simulation, the conductive heat flux has inflated the ICM at large radii, pushing the virial shock out to the computational boundary at $2 \times r_{200}$; this does not occur in my simulation with anisotropic conduction because the initially azimuthal field inhibits the radial transport of energy on large scales.

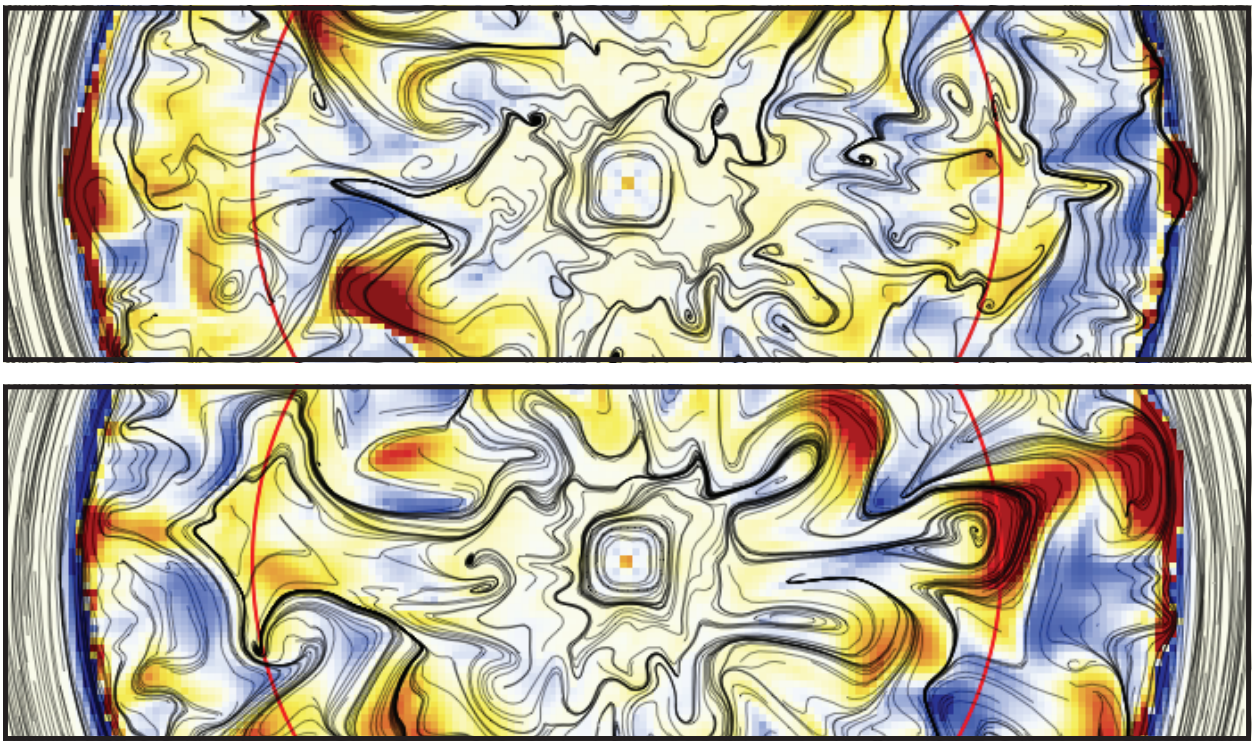

Figure 5.8: Comparison of viscous and inviscid simulations. This figure shows simulations of $10^{15}\,M_\odot$ clusters to maximize the importance of viscosity. As in figure 5.6, black lines show the projected magnetic field, and color shows the temperature fluctuation $\delta T/T$. Here, the color scale ranges from -0.15 (deep blue) to +0.15 (deep red). Viscosity smooths out the flow on small scales and appears to organize the magnetic field lines into more coherent, large-scale plumes. Interestingly, this has only a small effect on volume-integrated quantities such as the turbulent Mach number or the convective heat flux; see figure 5.10, below.

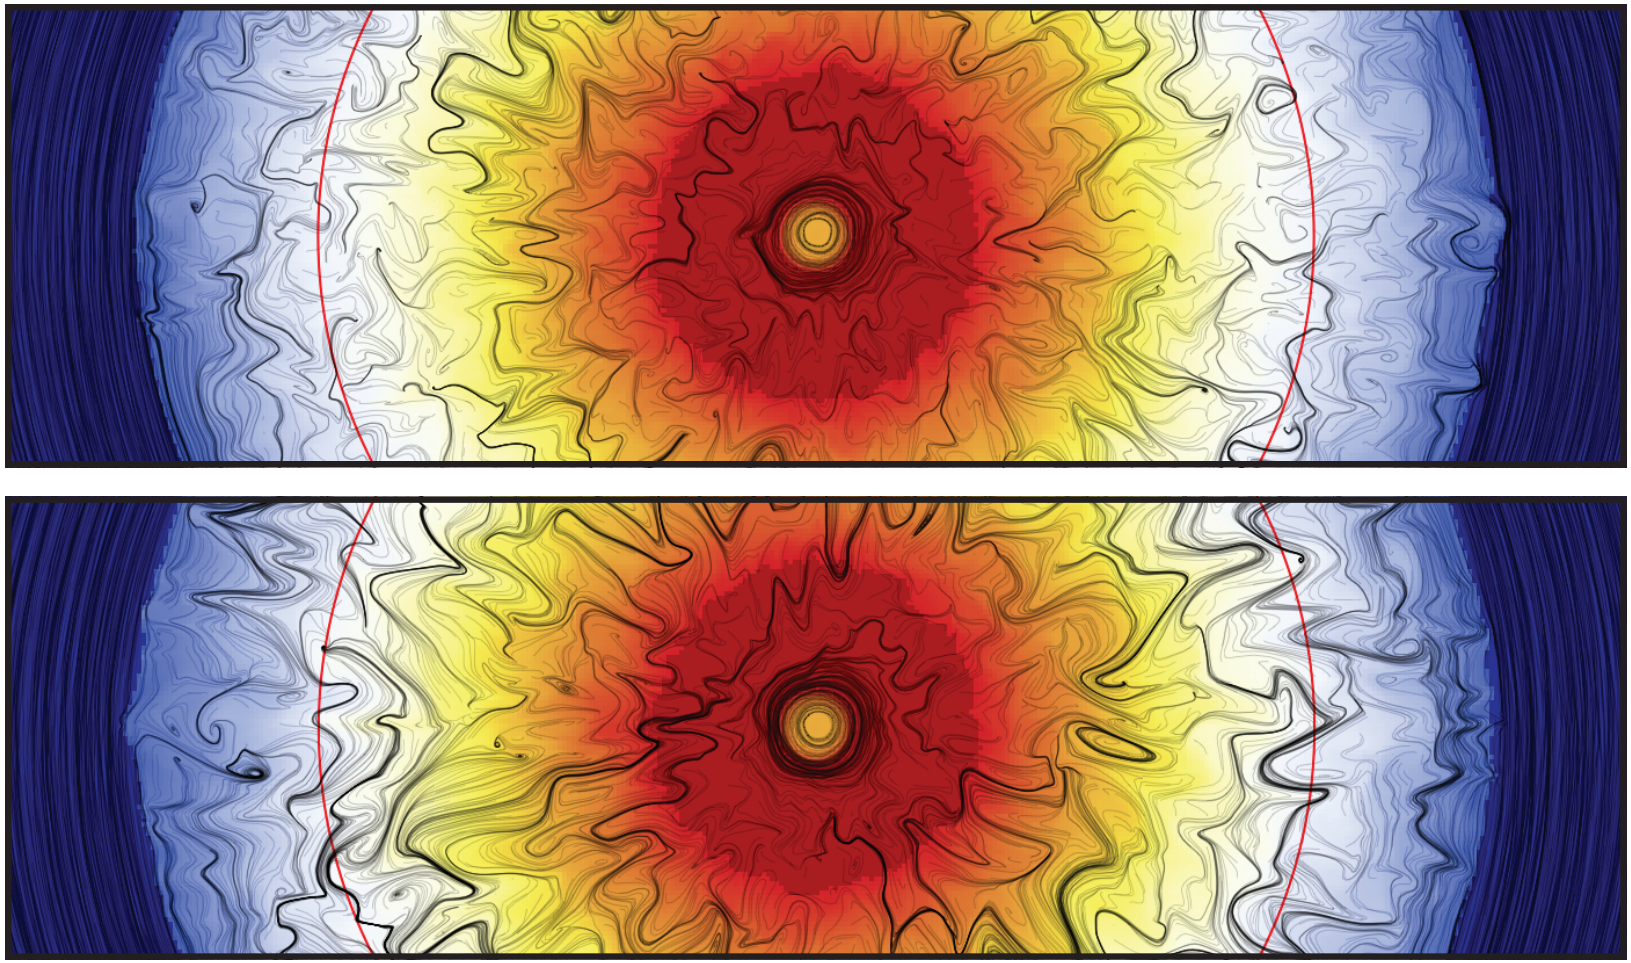

Figure 5.9: Comparison of viscous simulations with the Braginskii closure (bottom) and the "Sharma" closure (top), which limits the pressure anisotropy to within the bounds of the firehose and mirror instabilities (Sharma et al. 2006). The Sharma closure eliminates the strong bunching of the magnetic field lines seen in simulations using the Braginskii closure; this bunching may be due to the firehose instability. The two simulations are otherwise very similar.

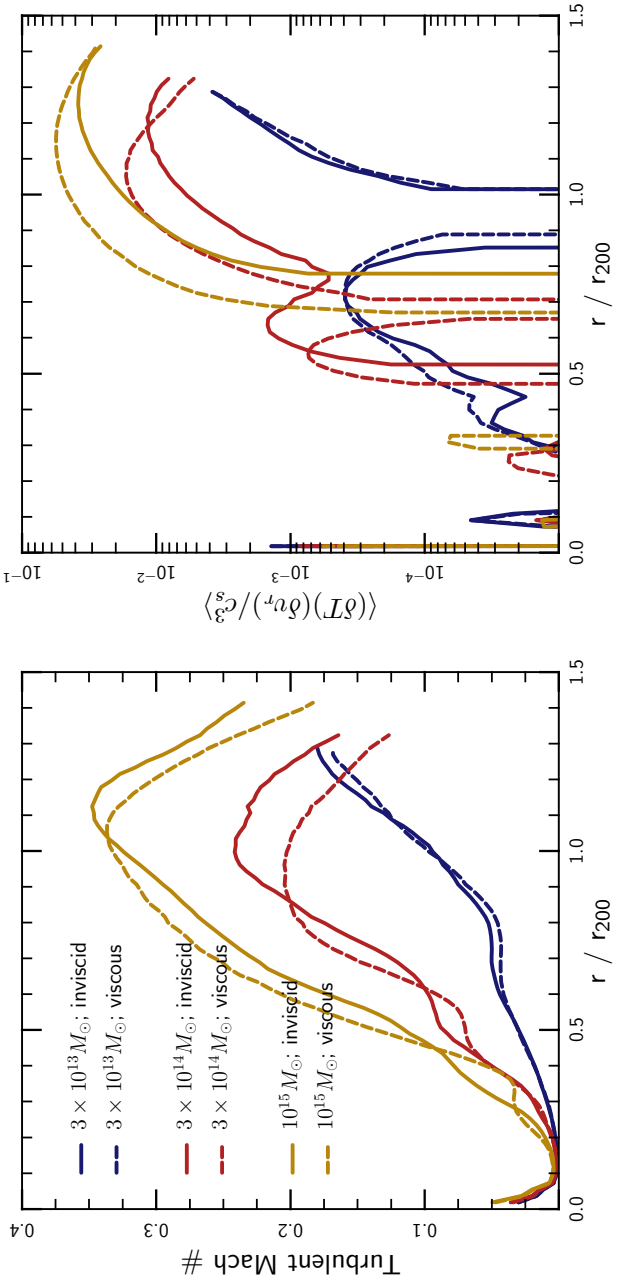

Figure 5.10: Quantification of the effect of viscosity in the idealized cluster simulations. The left panel shows the turbulent Mach number, defined as $\sqrt{[(v_r - \langle v_r \rangle)/c_s]^2}$, where $\langle \ldots \rangle$ represents an average over a spherical shell. Viscosity has a negligible effect on the evolution of the MTI in $10^{14}\,M_\odot$ clusters, but reduces the strength of the turbulence by $\sim 15\%$ in $10^{15}\,M_\odot$ clusters. The right panel shows the average $\langle (\delta T)(\delta v_r)/c_s^3 \rangle$; this quantity is similar to the convective heat flux, but corresponds more directly to the linear growth of the MTI. As in the left panel, viscosity has a negligible effect on the MTI in $10^{14}\,M_\odot$ clusters, but becomes more prominent by $10^{15}\,M_\odot$. Here, however, the effect of viscosity is to *increase* the flux by a factor of ~ 2.

Bibliography

Allen, S. W., Evrard, A. E., & Mantz, A. B. 2011, ARA&A, 49, 409
Allen, S. W., Rapetti, D. A., Schmidt, R. W., et al. 2008, MNRAS, 383, 879
Arnaud, M., Pratt, G. W., Piffaretti, R., et al. 2010, A&A, 517, A92
Balbus, S. A. 1986, ApJ, 303, L79
—. 1988, ApJ, 328, 395
—. 2000, ApJ, 534, 420
—. 2001, ApJ, 562, 909
Balbus, S. A., & Reynolds, C. S. 2008, ApJ, 681, L65
—. 2010, ApJ, 720, L97
Balbus, S. A., & Soker, N. 1989, ApJ, 341, 611
Bauermeister, A., Blitz, L., & Ma, C. 2010, ApJ, 717, 323
Begelman, M. C., & McKee, C. F. 1990, ApJ, 358, 375
Benson, A. J., Bower, R. G., Frenk, C. S., et al. 2003, ApJ, 599, 38
Binney, J., Nipoti, C., & Fraternali, F. 2009, MNRAS, 397, 1804
Bîrzan, L., Rafferty, D. A., McNamara, B. R., Wise, M. W., & Nulsen, P. E. J. 2004, ApJ, 607, 800
Bogdanović, T., Reynolds, C. S., Balbus, S. A., & Parrish, I. J. 2009, ApJ, 704, 211
Bogdanović, T., Reynolds, C. S., & Massey, R. 2011, ApJ, 731, 7
Braginskii, S. I. 1965, Reviews of Plasma Physics, 1, 205
Bullock, J. S., Kolatt, T. S., Sigad, Y., et al. 2001, MNRAS, 321, 559
Burns, J. O., Skillman, S. W., & O'Shea, B. W. 2010, ApJ, 721, 1105
Cavagnolo, K. W., Donahue, M., Voit, G. M., & Sun, M. 2008, ApJ, 683, L107
—. 2009, ApJS, 182, 12
Chang, P., Strubbe, L. E., Menou, K., & Quataert, E. 2010, MNRAS, 1003
Ciotti, L., & Ostriker, J. P. 2001, ApJ, 551, 131
Cole, S., Norberg, P., Baugh, C. M., et al. 2001, MNRAS, 326, 255
Cowie, L. L., Fabian, A. C., & Nulsen, P. E. J. 1980, MNRAS, 191, 399
Croston, J. H., Pratt, G. W., Böhringer, H., et al. 2008, A&A, 487, 431
Davis, M., Efstathiou, G., Frenk, C. S., & White , S. D. M. 1985, ApJ, 292, 371
Davis, S. W., Stone, J. M., & Pessah, M. E. 2010, ApJ, 713, 52
Defouw, R. J. 1970, ApJ, 160, 659
Diemand, J., Kuhlen, M., & Madau, P. 2007, ApJ, 667, 859

Dolag, K., Jubelgas, M., Springel, V., Borgani, S., & Rasia, E. 2004, ApJ, 606, L97
Dursi, L. J., & Pfrommer, C. 2008, ApJ, 677, 993
Fabian, A. C. 1994, ARA&A, 32, 277
Fabian, A. C., Allen, S. W., Crawford, C. S., et al. 2002, MNRAS, 332, L50
Fabian, A. C., Johnstone, R. M., Sanders, J. S., et al. 2008, Nature, 454, 968
Fabian, A. C., & Nulsen, P. E. J. 1977, MNRAS, 180, 479
Fabian, A. C., Sanders, J. S., Crawford, C. S., et al. 2003, MNRAS, 344, L48
Fabian, A. C., Sanders, J. S., Allen, S. W., et al. 2011, Monthly Notices of the Royal Astronomical Society, 418, 2154
Faucher-Giguère, C.-A., Kereš, D., & Ma, C.-P. 2011, MNRAS, 417, 2982
Ferland, G. J., Fabian, A. C., Hatch, N. A., et al. 2009, MNRAS, 392, 1475
Field, G. B. 1965, ApJ, 142, 531
Fox, D. C., & Loeb, A. 1997, ApJ, 491, 459
Fromang, S., & Papaloizou, J. 2007, A&A, 476, 1113
Gardiner, T. A., & Stone, J. M. 2008, Journal of Computational Physics, 227, 4123
George, M. R., Fabian, A. C., Sanders, J. S., Young, A. J., & Russell, H. R. 2009, MNRAS, 395, 657
Gunn, J. E., & Gott, III, J. R. 1972, ApJ, 176, 1
Guo, F., & Oh, S. P. 2008, MNRAS, 384, 251
Hattori, M., & Habe, A. 1990, MNRAS, 242, 399
Heckman, T. M., Baum, S. A., van Breugel, W. J. M., & McCarthy, P. 1989, ApJ, 338, 48
Holtzman, J. A., Faber, S. M., Shaya, E. J., et al. 1992, AJ, 103, 691
Hu, E. M., Cowie, L. L., & Wang, Z. 1985, ApJS, 59, 447
Johnson, B. M., & Quataert, E. 2007, ApJ, 660, 1273
Jones, T., Ellis, R. S., Jullo, E., & Richard, J. 2010, ArXiv e-prints: astro-ph.CO/1010.1538
Joung, M. R., Bryan, G. L., & Putman, M. E. 2011, ArXiv e-prints
Kaufmann, T., Mayer, L., Wadsley, J., Stadel, J., & Moore, B. 2006, MNRAS, 370, 1612
Kereš, D., & Hernquist, L. 2009, ApJ, 700, L1
Kochanek, C. S., Pahre, M. A., Falco, E. E., et al. 2001, ApJ, 560, 566
Komatsu, E., & Seljak, U. 2001, MNRAS, 327, 1353
Kunz, M. W., Bogdanović, T., Reynolds, C. S., & Stone, J. M. 2012, ApJ, 754, 122
Kunz, M. W., Schekochihin, A. A., Cowley, S. C., Binney, J. J., & Sanders, J. S. 2011, MNRAS, 410, 2446
Kunz, M. W., Schekochihin, A. A., & Stone, J. M. 2014, ArXiv e-prints, arXiv:1402.0010 [astro-ph.HE]
Lau, E. T., Kravtsov, A. V., & Nagai, D. 2009, ApJ, 705, 1129
Leccardi, A., & Molendi, S. 2008, A&A, 486, 359
Lemaster, M. N., & Stone, J. M. 2008, ApJ, 682, L97
Loewenstein, M., Zweibel, E. G., & Begelman, M. C. 1991, ApJ, 377, 392
Lynds, C. R., & Sandage, A. R. 1963, ApJ, 137, 1005
Lynds, R. 1970, ApJ, 159, L151
Malagoli, A., Rosner, R., & Bodo, G. 1987, ApJ, 319, 632

Malagoli, A., Rosner, R., & Fryxell, B. 1990, MNRAS, 247, 367
Maller, A. H., & Bullock, J. S. 2004, MNRAS, 355, 694
McBride, J., Fakhouri, O., & Ma, C.-P. 2009, MNRAS, 398, 1858
McCarthy, I. G., Balogh, M. L., Babul, A., Poole, G. B., & Horner, D. J. 2004, ApJ, 613, 811
McCourt, M., Parrish, I. J., Sharma, P., & Quataert, E. 2011a, MNRAS, 413, 1295
—. 2011b, MNRAS, 413, 1295
McCourt, M., Quataert, E., & Parrish, I. J. 2013, MNRAS, 432, 404
McCourt, M., Sharma, P., Quataert, E., & Parrish, I. J. 2012, MNRAS, 419, 3319
McDonald, M., Veilleux, S., & Mushotzky, R. 2011a, ApJ, 731, 33
McDonald, M., Veilleux, S., Rupke, D. S. N., & Mushotzky, R. 2010, ApJ, 721, 1262
McDonald, M., Veilleux, S., Rupke, D. S. N., Mushotzky, R., & Reynolds, C. 2011b, ApJ, 734, 95
McNamara, B. R., & Nulsen, P. E. J. 2007, ARA&A, 45, 117
Nagai, D., Kravtsov, A. V., & Vikhlinin, A. 2007, ApJ, 668, 1
Narayan, R., & Medvedev, M. V. 2001, ApJ, 562, L129
Navarro, J. F., Frenk, C. S., & White, S. D. M. 1997, ApJ, 490, 493
Navarro, J. F., Hayashi, E., Power, C., et al. 2004, MNRAS, 349, 1039
Nulsen, P. E. J. 1986, MNRAS, 221, 377
O'Dea, K. P., Quillen, A. C., O'Dea, C. P., et al. 2010, ApJ, 719, 1619
Oh, S. P., & Benson, A. J. 2003, MNRAS, 342, 664
Padmanabhan, T. 1993, Structure Formation in the Universe (Cambridge Press)
Parrish, I. J., McCourt, M., Quataert, E., & Sharma, P. 2012a, MNRAS, 422, 704
—. 2012b, MNRAS, 419, L29
Parrish, I. J., & Quataert, E. 2008, ApJ, 677, L9
Parrish, I. J., Quataert, E., & Sharma, P. 2009, ApJ, 703, 96
—. 2010, ApJ, 712, L194
Parrish, I. J., & Stone, J. M. 2005, ApJ, 633, 334
—. 2007, ApJ, 664, 135
Parrish, I. J., Stone, J. M., & Lemaster, N. 2008, ApJ, 688, 905
Peek, J. E. G., Putman, M. E., & Sommer-Larsen, J. 2008, ApJ, 674, 227
Peterson, J. R., & Fabian, A. C. 2006, Phys. Rep., 427, 1
Pfrommer, C., & Dursi, J. L. 2010, Nature Physics, 6, 520
Piffaretti, R., Jetzer, P., Kaastra, J. S., & Tamura, T. 2005, A&A, 433, 101
Pizzolato, F., & Soker, N. 2005, ApJ, 632, 821
—. 2010, MNRAS, 408, 961
Quataert, E. 2008, ApJ, 673, 758
Rafferty, D. A., McNamara, B. R., & Nulsen, P. E. J. 2008, ApJ, 687, 899
Rasia, E., Ettori, S., Moscardini, L., et al. 2006, MNRAS, 369, 2013
Rees, M. J., & Ostriker, J. P. 1977, MNRAS, 179, 541
Riquelme, M., Quataert, E., & Verscharen, D. 2014, ArXiv e-prints, arXiv:1402.0014 [astro-ph.HE]

Rudd, D. H., & Nagai, D. 2009, ApJ, 701, L16
Ruszkowski, M., Lee, D., Brüggen, M., Parrish, I., & Oh, S. P. 2011, ApJ, 740, 81
Ruszkowski, M., & Oh, S. P. 2010, ApJ, 713, 1332
Saro, A., Borgani, S., Tornatore, L., et al. 2006, MNRAS, 373, 397
Sharma, P., Chandran, B. D. G., Quataert, E., & Parrish, I. J. 2009a, ApJ, 699, 348
Sharma, P., Chandran, B. D. G., Quataert, E., & Parrish, I. J. 2009b, in American Institute of Physics Conference Series, Vol. 1201, American Institute of Physics Conference Series, ed. S. Heinz & E. Wilcots, 363
Sharma, P., & Hammett, G. W. 2007, Journal of Computational Physics, 227, 123
Sharma, P., Hammett, G. W., Quataert, E., & Stone, J. M. 2006, ApJ, 637, 952
Sharma, P., McCourt, M., Parrish, I. J., & Quataert, E. 2012a, MNRAS, 427, 1219
Sharma, P., McCourt, M., Quataert, E., & Parrish, I. J. 2012b, MNRAS, 420, 3174
Sharma, P., Parrish, I. J., & Quataert, E. 2010, ApJ, 720, 652
Sharma, P., Quataert, E., & Stone, J. M. 2008, MNRAS, 389, 1815
Shaw, L. D., Nagai, D., Bhattacharya, S., & Lau, E. T. 2010, ApJ, 725, 1452
Shcherbakov, R. V., & Baganoff, F. K. 2010, ApJ, 716, 504
Sijacki, D., & Springel, V. 2006, MNRAS, 366, 397
Silk, J. 1977, ApJ, 211, 638
Simionescu, A., Allen, S. W., Mantz, A., et al. 2011, Science, 331, 1576
Smith, B. D., O'Shea, B. W., Voit, G. M., Ventimiglia, D., & Skillman, S. W. 2013, ArXiv e-prints, arXiv:1306.5748 [astro-ph.CO]
Soker, N. 2006, New A, 12, 38
Soker, N., Blanton, E. L., & Sarazin, C. L. 2004, A&A, 422, 445
Sommer-Larsen, J. 2006, ApJ, 644, L1
Spitzer, L. 1962, Physics of Fully Ionized Gases (John Wiley & Sons Inc.)
Springel, V., White, S. D. M., Jenkins, A., et al. 2005, Nature, 435, 629
Stone, J. M., Gardiner, T. A., Teuben, P., Hawley, J. F., & Simon, J. B. 2008, ApJS, 178, 137
Tozzi, P., & Norman, C. 2001, ApJ, 546, 63
Turner, J. 1979, Buoyancy Effects in Fluids, Cambridge Monographs on Mechanics (Cambridge University Press)
Vazza, F., Brunetti, G., Gheller, C., Brunino, R., & Brüggen, M. 2011, A&A, 529, A17
Voit, G. M. 2011, ApJ, 740, 28
Voit, G. M., Balogh, M. L., Bower, R. G., Lacey, C. G., & Bryan, G. L. 2003, ApJ, 593, 272
Voit, G. M., Cavagnolo, K. W., Donahue, M., et al. 2008, ApJ, 681, L5
Voit, G. M., & Ponman, T. J. 2003, ApJ, 594, L75
Wechsler, R. H., Bullock, J. S., Primack, J. R., Kravtsov, A. V., & Dekel, A. 2002, ApJ, 568, 52
White, M. 2001, A&A, 367, 27
White, S. D. M., & Rees, M. J. 1978, MNRAS, 183, 341
Wu, H.-Y., Hahn, O., Wechsler, R. H., Mao, Y.-Y., & Behroozi, P. S. 2013, ApJ, 763, 70
Zakamska, N. L., & Narayan, R. 2003, ApJ, 582, 162

Zhao, D. H., Jing, Y. P., Mo, H. J., & Börner, G. 2009, ApJ, 707, 354