

HOT GASEOUS HALOS AROUND GALAXIES

by

Michael E. Anderson

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Astronomy and Astrophysics)
in The University of Michigan
2013

Doctoral Committee:

Professor Joel N. Bregman, Chair
Associate Professor Eric F. Bell
Professor August Evrard
Associate Professor Oleg Y. Gnedin
Assistant Professor Christopher J. Miller
Associate Professor Jon M. Miller
Associate Professor Mateusz Ruszkowski

Copyright © Michael E. Anderson 2013
All Rights Reserved

To my parents, with love

ACKNOWLEDGMENTS

It is a pleasure to acknowledge the many people and institutions upon whom I have relied during my graduate studies at Michigan. It has been a wonderful experience working with and learning from my advisor, Joel Bregman. I greatly admire his combination of uncannily prescient leaps of big-picture intuition paired with scrupulous attention to details, and to the extent I have learned to think more like a scientist during my PhD program, Joel deserves the credit for this. He has also been unwaveringly supportive of me throughout this process. I really appreciate his willingness to take me on as an unexpected PhD student, his availability, assistance, and forbearance with technical issues large and small, and his insights into the wider culture of academia.

The rest of my thesis committee – Eric Bell, Gus Evrard, Oleg Gnedin, Chris Miller, Jon Miller, and Mateusz Ruskowski – have also helped me in a number of ways. Joel, Eric, and Oleg have been incredibly supportive of Astrocoffee, and I have learned an enormous amount from their lively debates. Jon went out of his way to welcome me to the program and to guide me through the thicket of relevant departmental and Rackham requirements. And everyone on my committee has contributed a number of useful suggestions throughout this process, and I really appreciate their attention and their contributions.

The department in general has also fostered a very supportive environment. I have benefitted greatly from the large number of opportunities to give talks to the department, and from the genuine interest the faculty take in the personal and academic well-being of the graduate students. The graduate chair, Nuria Calvet, also allowed me to pursue a certificate in Science, Technology, and Public Policy at the Ford School while working on my dissertation. This certificate may not have con-

tributed directly to my thesis, but I have learned a lot about the connection between science and society, which I hope and expect to serve me well in the rest of my career.

Graduate school is also supposed to be a horizontal learning process, and my peers have been a formidable resource during my time here. I would especially like to thank Nate Crockett, Edmund Hodges-Kluck, Ashley King, Matt Miller, Sasha Muratov, and Colin Slater for all their helpful conversations, lessons, and advice, which improved this thesis work immeasurably. I would also gratefully like to acknowledge Xinyu Dai and Jon Miller, who co-authored some of the chapters in this work, as well as Hajira Chaudhry, Elena Gallo, Brendan Miller, and Ezekiel Silverstein, with whom I am collaborating on other work which does not appear in this thesis.

Chapters 2, 3, 4, and 5 of this thesis are based on publications in the *Astrophysical Journal*; the references and copyright information are noted at the beginning of each chapter. Chapter 6 is based on a publication to appear in the *Astrophysical Journal* (Anderson and Bregman 2013, in preparation). This thesis work has been partially funded by Chandra grant GO9-0089X and by NASA ADAP grant NNX11AJ55G. I am also grateful for an NSF Graduate Research Fellowship and a Rackham One-Term Dissertation Fellowship, which have also generously funded me for parts of my PhD program.

Finally, I would like to thank my friends, my girlfriend, and my family, for their love and support throughout this entire process.

CONTENTS

DEDICATION	ii
ACKNOWLEDGMENTS	iii
LIST OF FIGURES	ix
LIST OF TABLES	xiii
ABSTRACT	xiv
CHAPTER	
1 Introduction	1
1.1 References	8
2 Do Hot Haloes Around Galaxies Contain the Missing Baryons?	11
2.1 Abstract	11
2.2 Introduction	12
2.3 On Milky Way Parameters	14
2.4 Constraints on the Milky Way Halo	16
2.4.1 LMC Pulsar Dispersion Measure	16
2.4.2 Other Constraints	18
2.5 OVII Absorption in Other Galaxies	20
2.6 X-Ray Surface Brightness	24
2.7 Energetics for Gas-Rich Late-Type Galaxies	26
2.8 The Effect of Flatter Density Profiles	30
2.9 Discussion and Conclusion	33
2.10 Acknowledgements	36
2.11 References	37

3	Detection of a Hot Gaseous Halo Around the Giant Spiral Galaxy NGC 1961	49
	3.1 Abstract	49
	3.2 Introduction	50
	3.3 Observation	52
	3.4 Flat-fielding	53
	3.4.1 Using Background Frames	53
	3.4.2 In-field Subtraction, Modeling the Background	54
	3.4.3 In-field Subtraction, Conjugate Technique	55
	3.4.4 Point Sources	56
	3.5 Radial Surface Brightness Profile	59
	3.5.1 Parametric Fitting with the β -model	59
	3.6 Spectral Fitting	61
	3.7 Halo Mass	63
	3.7.1 Flattened Profiles	64
	3.8 Implications and Conclusion	65
	3.8.1 Generalizing the Conjugate Subtraction Technique	65
	3.8.2 Halo Faintness and the Baryon Budget of NGC 1961	66
	3.8.3 Halo Cooling Rates and Implications for Galaxy Formation	68
	3.9 Acknowledgements	71
	3.10 References	71
4	XMM-Newton Detects a Hot Gaseous Halo in the Fastest Rotating Spiral Galaxy UGC 12591	79
	4.1 Introduction	79
	4.2 Observation and Data Reduction	80
	4.3 Spectral Analysis	81
	4.4 Spatial Analysis	82
	4.5 Discussion	88
	4.5.1 Baryon Mass Components in UGC 12591	88

4.5.2	Cooling of the hot gas	89
4.6	References	90
5	Extended Hot Halos Around Isolated Galaxies Observed in the	
	<i>ROSAT</i> All-Sky Survey	91
5.1	Abstract	91
5.2	Introduction	92
5.3	Sample	95
5.4	Constructing Stacked Images	96
5.5	Analysis	102
5.5.1	Parameterizing the surface brightness profile	102
5.5.2	Computing the PSF	105
5.6	Simulations	107
5.6.1	Recovering total counts and fit parameters	111
5.6.2	Recovering counts within 50 kpc	111
5.6.3	Identifying extended emission	114
5.6.4	Effects of changing the PSF	118
5.6.5	Including a combination of point source and extended components	119
5.6.6	Summary of simulations	120
5.7	Statistical significance of detections	121
5.8	Results	125
5.8.1	Average Luminosity	127
5.8.2	AGN emission	133
5.8.3	X-ray binaries	134
5.8.4	The L_X - L_K relation	138
5.8.5	Hot Gas Mass	138
5.8.6	Comparison with previous work	145
5.9	Conclusions	146
5.10	Acknowledgements	148
5.11	References	151

6	Modeling X-ray Emission Around Galaxies	155
6.1	Abstract	155
6.2	Introduction	156
6.3	The model	160
6.3.1	Background	160
6.3.2	Source	161
6.3.3	Point sources	164
6.4	Simulating images	165
6.4.1	Simulating Point Sources	166
6.4.2	Detecting Point Sources	167
6.5	Example: Chandra Deep Field South	168
6.6	The Likelihood Function	174
6.6.1	Measuring f_{unvig} for observation 8595	179
6.7	Simulated Images with an Extended Source	179
6.7.1	Source at Large Off-Axis Angle	180
6.7.2	Source at Aimpoint	183
6.8	Application to Real Observation: NGC 720	184
6.9	Conclusion	190
6.10	References	192
7	Conclusion	196
7.1	Late-type Galaxies	196
7.2	Early-type Galaxies	200
7.3	Systematic Uncertainties	201
7.4	Missing Baryons From Galaxies?	206
7.5	Future Prospects	208
7.6	References	209

LIST OF FIGURES

Figure

2.1	Integrated X-Ray emission measures for a halo of hot gas obeying an NFW density profile	41
2.2	Comparison of the scatter of galaxies around the baryonic Tully-Fisher relation	42
2.3	Comparison of galaxy residuals about the baryonic Tully-Fisher relation as a function of their stellar mass fraction	43
2.4	SDSS images of NGC 4138 and NGC 4183	44
2.5	Integrated X-Ray emission measures for a halo of hot gas with a power-law density profile	48
3.1	An ESO DSS image of NGC 1961, with the layout of our four ACIS-I observations overlaid	74
3.2	Our four Chandra ACIS-I observations of NGC 1961	75
3.3	Background-subtracted radial surface brightness profile of NGC 1961 .	76
3.4	Background-subtracted radial surface brightness profiles for all four observations of NGC 1961, after background smoothing	77
3.5	Log-log plot of radial surface brightness profiles for all four observations of NGC 1961	78
4.1	X-ray surface brightness profile of UGC 12591 in the XMM-Newton PN observation	82
4.2	X-ray surface brightness profile of UGC 12591 in the XMM-Newton M1 observation	83

4.3	X-ray surface brightness profile of UGC 12591 in the XMM-Newton M2 observation	84
4.4	Background-subtracted, vignetting-corrected XMM-Newton PN and MOS radial surface brightness profiles of the 0.6-1.4 keV emission around UGC 12591	86
5.1	Histograms of K_S absolute magnitudes and distances for our sample of galaxies in the 2MASS Isolated Galaxies catalog	96
5.2	Stacked images of our samples of 2MIG galaxies	99
5.3	These plots contain the same information as the images in Figure 5.2, expressed as surface brightness profiles instead of images	100
5.4	Mean surface brightness profiles of the primary stacked PSFs	103
5.5	Profiles showing the normalized strength of the PSF as a function of projected radius	108
5.6	Probability distribution functions of the ratio of the measured value of $N \equiv N_\beta + N_\delta$ to the true value of N for simulations of each galaxy sample	112
5.7	Probability distribution functions (pdfs) for the measured value of the slope β of the extended component of the simulated images, compared to the true value of β for those images	113
5.8	Probability distribution functions (pdfs) of the ratio of the measured value of $N \equiv N_\beta + N_\delta$ to the true value of N for simulations of each galaxy sample	115
5.9	Probability distribution functions (pdfs) for (a) the values of N_β and N_δ integrated out to 500 kpc and (b) the ratios of the values of N_β and N_δ integrated out to 50 kpc compared to the true values	116
5.10	Same as Figure 5.9, but for the “extra counts” simulations	117
5.11	Comparison of pdfs of recovered counts within 50 kpc after variations in the redshift distribution/aperture size of point sources we stack to produce the empirical PSFs	119

5.12	Comparison of pdfs of recovered counts within 50 kpc to the true number of counts within 50 kpc for simulations where $N_\beta = N_\delta$	120
5.13	Probability distribution functions (pdfs) for the total number of source counts in the stacked 2MIG images, within a radius of (a) 50 kpc and (b) 500 kpc from the central galaxies	126
5.14	The pdfs for the fraction of source counts within 50 kpc of the galaxy, attributed to the extended component for all four subsamples	129
5.15	Estimates of the conversion factor from counts to energy per unit area for our data, assuming various spectral models for the emission	130
5.16	Probability distribution functions (pdfs) for the average 0.5–2.0 keV X-ray luminosity within radii of (a) 50 kpc and (b) 300 kpc of the galaxy, for each of our four subsamples of 2MIG galaxies	132
5.17	Probability distribution functions (pdfs) for the average 0.5–2.0 keV X-ray luminosity of the <i>extended component</i> of the emission within radii of (a) 50 kpc and (b) 300 kpc of the galaxy for each of our four subsamples of 2MIG galaxies	132
5.18	Hardness ratio of source counts within 50 kpc of the galaxy, for the full sample and all four samples of 2MIG galaxies, along with hardness ratios predicted by various spectral models	135
5.19	L_X - L_K relation for all the luminous, the luminous late-type, and the luminous early-type galaxies	139
5.20	Probability distribution functions (pdfs) for the hot gas mass within 50 kpc of the galaxy	141
5.21	Ratio of the hot gas mass within 200 or 300 kpc to the hot gas mass within 50 kpc, assuming the hot gas follows a β -model	144
6.1	The amount of nongaussianity in the psf within an ecf of (a) 75%, (b) 90%, and (c) 95%	163
6.2	Illustrations of various steps in our analysis pipeline, using Chandra obsid 8595 (the CDF-S) as our example	170

6.3	Comparison of point sources detected in observation 8595 to the full CDF-S	171
6.4	Comparison of point sources detected with two different methods	172
6.5	Simulated images with mock point sources added and then masked with <i>wavdetect</i>	173
6.6	Comparison of a few example likelihood functions in recovering the shape of the background	176
6.7	Simulated images with mock point sources added at random, and an extended source placed near the bottom of the ACIS-I array that contains 4000 source counts	181
6.8	Posterior probability distribution functions (pdfs) for the image parameters	182
6.9	Simulated images with mock point sources added at random, and an extended source placed near aimpoint of the ACIS-I array that contains 4000 source counts	183
6.10	Posterior probability distribution functions (pdfs) for the image parameters	184
6.11	1σ regions corresponding to the estimated mass in the hot halo around NGC 720, from three different analyses of the galaxy	188
6.12	Comparison of predicted surface brightness profiles around NGC 720 to the observed surface brightness profile	189
7.1	A schematic diagram showing example hot halo density profiles and possible constraints	210

LIST OF TABLES

Table

2.1	Constraints on the Milky Way Hot Halo	45
2.2	OVII Sightline Analysis	46
2.3	Halo Emissivity Analysis	47
2.4	Comparison of Two Nearby Disk Galaxies	47
5.1	PSF fitting parameters	106
5.2	Properties of simulations	110
5.3	Significance of detections	123
5.4	Summary of measurements and derived quantities	150
6.1	Likelihood Function Comparison	177
6.2	Measurements of f_{unvig} with various likelihood functions	179
7.1	Summary of Hot Halo Properties	213

ABSTRACT

The modern theory of galaxy formation makes a generic prediction that all massive galaxies, elliptical or spiral, ought to be surrounded by hot gaseous halos. Elliptical galaxies are indeed generally surrounded by hot gas, but for many reasons spiral galaxies offer a better laboratory for testing the theory and its implications for missing baryons. Unfortunately, hot halos around spiral galaxies have eluded detection for decades. In this thesis, we present the first two detections of hot halos around spiral galaxies (NGC1961 and UGC12591), as well as a number of results characterizing the hot gas and its implications. Both the hot halos we detect have similar properties, and it is likely that neither contains more than a third of the missing baryons from either galaxy. We also compile constraints on the mass of hot halos around the Milky Way and other massive spirals, based on upper limits in emission and absorption. For the Galaxy the most robust constraint is based on the dispersion measure of pulsars in the Magellanic Clouds, which argues against a Galactic hot halo containing more than a quarter of the missing baryons from the Galaxy. In addition, we perform a stacking analysis of ROSAT All-Sky Survey images of isolated galaxies, in which we detect X-ray emission from L^* spirals and ellipticals at high confidence, and find moderate evidence that the emission is extended on scales of tens of kpc. We infer hot halo masses consistent with our other studies, and are able to strongly exclude models with more massive hot halos that would have contained the missing baryons from these galaxies. Finally, we introduce significant improvements for the spatial analysis of extended X-ray emission, which we use to trace the emission around the isolated elliptical NGC 720 to below 1/10 of the background. We are able to settle a discrepancy between two different studies of the hot halo around this galaxy, suggesting this galaxy does not have most of its missing baryons in its hot halo. We

conclude that the missing baryons must either be in a cooler phase, or outside the virial radius.

CHAPTER 1

Introduction

The modern theory of galaxy formation traces its roots to White and Rees (1978), who showed that hot gas plays a critical role in the formation of galaxies. In their original (analytic) model, baryons falling into nascent protogalactic potential wells experienced accretion shocks, which heated them to the virial temperature of the system ($T \sim$ a few $\times 10^6$ K for the Milky Way). The result was a hot halo of gas, tracing the dark matter and filling the virial volume out to hundreds of kpc, from which gas slowly cooled from the inside-out, forming a galaxy.

At the same time, it was also appreciated that gas at similar hot temperatures should also be expected around massive galaxies for another, completely independent, reason – supernova-driven winds. These winds were first proposed by Mathews and Baker (1971) to explain the dearth of interstellar gas in massive elliptical galaxies. In their model, the interstellar medium (ISM) was heated as stellar winds collided and shocked to a temperature corresponding to the velocity dispersion of the stars, and then supernovae added additional energy to the ISM and produced galaxy-scale outflows. This model seemed to successfully explain the observed deficit of cold gas in ellipticals (Faber and Gallagher 1976), and was soon incorporated into three-phase models of the ISM in spiral galaxies as well (e.g. McKee and Ostriker 1977).

Ever since these initial pioneering studies, the field of galaxy formation has essentially been the story of the complex interplay between accretion and outflow (feedback) on galactic scales¹. While both accretion and feedback are multiphase phe-

¹Mergers are of course important as well, and they are related to accretion; technically accretion and mergers are the two extreme cases of the broader phenomenon of subhalos falling into larger halos. However, mergers primarily seem to affect the stellar properties of galaxies, while this thesis is more concerned with the gas dynamics and with the overall baryon fraction.

nomena, each spanning more than five orders of magnitude in gas temperature (from molecular gas at $T \sim 10$ K to hot gas at $T > 10^6$ K), both processes were originally studied in their hot phases. The rest of this introduction will review the history of our understanding of galactic accretion and feedback, with emphasis on the hot phases of each process.

It was immediately appreciated that, while the soft X-rays were the natural energy range in which to search for evidence of hot halos and hot winds, detectability would be an issue. For example, Mathews and Baker (1971) predicted a characteristic 4-keV luminosity of $\sim 1 \times 10^{36}$ erg s $^{-1}$ for steady-state supernova-driven winds integrated over the entirety of a massive elliptical galaxy, and pointed out that this translates into such a low surface brightness as to be unobservable with the existing generation of X-ray telescopes (or, indeed, even with their modern successors). Nevertheless, searches have been ongoing for decades, with some notable successes and failures.

The first major success was the detection (Forman et al. 1979) of extended X-ray coronae around massive elliptical galaxies. This study observed the core of the Virgo Cluster with the Einstein Observatory's Imaging Proportional Counter in the 0.5-3.0 keV band. M87, the brightest cluster galaxy (BGC), lies at the center of the cluster emission, but the massive ellipticals M86, NGC 4438, M34, and NGC 4388 are detected as well and the emission from M86 is identified as thermal bremsstrahlung. Follow-up observations of 55 massive ellipticals led to the conclusion in Forman et al. 1985 that "hot gaseous coronae are a common and perhaps ubiquitous feature" of these galaxies.

Forman et al. (1985) found that the observed X-ray luminosity of these hot halos scales as the square of the B-band luminosity of the associated galaxies, $L_X \propto L_B^2$. They show that this scaling, as well as the normalization of the relation, are both consistent with the hot halo having a mass equal to the mass lost from the old AGB stars in these galaxies, exactly as predicted by Faber and Gallagher (1976). However, Forman et al. (1985) pointed out that the total amount of mass loss expected from the stars in these galaxies is much higher than this observed hot halo mass, since younger stars lose mass very efficiently through stellar winds and supernovae. When

accounting for this additional mass loss, they point out that the observed hot halos are less massive than expected by up to a factor of 50.

Alternatively, Nulsen et al. (1984) showed that a similar relation ($L_X \propto L_B^{5/3}$) is also produced through hot-mode accretion onto massive ellipticals, along the lines of the White and Rees (1978) model. So, while most authors favored the stellar mass-loss explanation for hot halos around ellipticals, accretion was also considered a potentially viable solution. Decades later, two new types of observations tipped the scale decisively in favor of the mass-loss scenario. First, XMM spectroscopy failed to detect line emission around BCGs at the strengths expected if hot gas was cooling to temperatures below ~ 1 keV (e.g. Peterson et al. 2001, Molendi and Pizzolato 2001, Peterson et al. 2003). This is obviously inconsistent with a model where the stars in the galaxy condense out of the hot accreted gas, which was the prediction of “cooling flow” models in clusters of galaxies. Significant cooling flows have also not been detected around isolated ellipticals. Secondly, the hot halos around massive ellipticals are observed to have super-solar metallicities (Mathews and Brighenti 2003), which requires the gas to have been processed through stellar evolution. This is a natural prediction of mass-loss models, but is difficult to explain if the gas is supposed to have an extragalactic origin.

While the mass-loss model is generally favored, the debate about the origin of hot halos around ellipticals is not entirely settled. AGN feedback is a plausible mechanism for preventing the gas from cooling (Tucker and David 1997, Mathews and Brighenti 2003), which can explain the absence of strong cooling lines. And, as Crain et al. (2013) point out, most of the soft X-ray emission from these hot halos is line emission, so enriched gas has a much higher emission measure than pristine gas. Therefore it is possible to have a component of enriched gas provided by stellar mass loss that “screens out” a more extended and more massive component of pristine, accreted gas.

The situation for hot halos around spiral galaxies is less clear. The White and Rees (1978) model for hot halos formed from accretion shocks depends on the balance between the dynamical timescale for the infalling gas (set by the mass of the galaxy + dark matter halo) and the cooling time for the gas (set primarily by its metallicity,

and assumed to be low for newly infalling gas). Thus hot halos should form around sufficiently massive galaxies, regardless of morphology, and so this model generically predicts hot halos around spiral galaxies. The mass-loss model, on the other hand, depends on having a pressure-supported distribution of stars to shock-heat the gas, and spiral galaxies are generally rotation-supported instead. Thus one would not expect much of a hot halo to be produced around spiral galaxies by this mechanism. However, in addition to velocity dispersion heating, the other ingredient in the Mathews and Baker (1971) model for hot feedback is supernova feedback, and this source is abundant in spiral galaxies due to the enhanced star formation rates. So some level of hot gas is still expected around spiral galaxies due to heating from supernovae.

The first search for hot gas around spiral galaxies was undertaken by Bregman and Glassgold (1982). They looked at the nearby massive edge-on spirals NGC 3628 and NGC 4244 with the Einstein IPC, using a similar observing strategy as Forman et al. (1979). However, in contrast to the situation with the elliptical galaxies, Bregman and Glassgold found no evidence of diffuse soft X-ray emission around these spirals down to extremely stringent upper limits (2×10^{38} erg s⁻¹). Subsequently, pointed observations with the ROSAT PSPC were undertaken by Benson et al. (2000) for the same purpose; they found no extended soft X-ray emission around the massive spirals NGC 2841, NGC 4594, and NGC 5529 down to upper limits of below 10^{41} erg s⁻¹. Finally, another search was conducted with the Chandra X-ray observatory (Pedersen et al. 2006) around the massive spirals NGC 5746 and NGC 5170. These authors originally reported a detection of a hot halo around NGC 5746, but the detection disappeared upon reanalysis with an updated Chandra calibration (Rasmussen et al. 2009); their 3σ upper limit on the hot halo luminosity around NGC 5746 is 4×10^{39} erg s⁻¹.

For most models of galaxy formation, and most numerical simulations, the non-detection of hot halos around massive spirals was a surprise. Bregman and Glassgold computed that their nondetection suggests that $< 0.1\%$ of the energy supplied to the ISM by supernovae escapes from the galaxy in the soft X-rays; since SNe shocks heat gas to X-ray temperatures, this very low fraction was unexpected. Today this

is understood as a result of the porosity of the ISM: where the ISM is dense enough for SNe shocks to radiate strongly, they cool quickly and the optical depth to soft X-rays is high; where the ISM is rarefied the hot gas flows freely but radiates very little. Indeed, in actively starbursting spirals, significant amounts of X-ray emission are visible throughout the galaxies and in coronae extending 10-20 kpc above and below the disk (Fabbiano and Trinchieri 1984, Fabbiano and Trinchieri 1987, Bregman and Pildis 1994, Wang et al. 1995, Dahlem et al. 1998, Strickland et al. 2004ab, Tüllmann et al. 2006ab, Li and Wang 2013). This emission often traces H α emission and arises at the interfaces between different phases of the ISM, and is thought to be powered by a complex combination of cooling radiation, shocks, and charge exchange (Heckman et al. 1990, Strickland et al. 2004ab).

The other major implication of these nondetections is the problem of the missing baryons from galaxies. Ever since the concordance model of cosmology was established and the approximate values of Ω_m and Ω_b became fairly settled quantities (current values are 0.28 and 0.05, respectively; Hinshaw et al. 2013), it was realized that the observed baryon fraction in galaxies is much lower than the baryon fraction of the Universe (Silk 2003, McGaugh 2005). This is related to, but conceptually slightly separate from, the classical “missing baryons problem”, and is known as the problem of missing baryons from galaxies.

Before the nondetection of hot halos around massive spirals was widely appreciated, these hot halos were considered the obvious reservoir for the missing baryons from galaxies, as they were already invoked in the theory of White and Rees (1978) as a natural product of the galaxy formation process and as the reservoir for continuing accretion and fuel for star formation over the lifetime of the galaxy. This theory was elaborated by White and Frenk (1991), who predicted soft X-ray luminosities in excess of 10^{41} erg s $^{-1}$ for L* galaxies accreting at $1M_{\odot}$ yr $^{-1}$, and luminosities more than an order of magnitude larger for the giant spirals. Essentially the same calculation was repeated again by Fukugita and Peebles (2006), and they found a similar value for the total predicted luminosity (though they focus on the luminosity at radii beyond 10 kpc, which they predict to be $\sim 10^{40}$ erg s $^{-1}$). However, these predictions

are in tension with the observations described above, so theorists developed other solutions to the problem.

One factor which helps to reduce the tension with the X-ray observations is the realization that a significant fraction of intergalactic gas does not undergo accretion shocks, but is instead able to advect onto the galaxy directly (or into a cold disk rotating around the galaxy). This “cold mode” of accretion (Birnboim and Dekel 2003, Katz et al. 2003, Kereš et al. 2005) therefore does not produce a hot halo. While cold accretion might be expected to occur preferentially (or entirely) in low-mass halos where the cooling time of the gas is shorter than the dynamical time, it has also been proposed that cold accretion is the dominant mode of accretion across all galaxy masses (Dekel et al. 2009, Brooks et al. 2009, Oser et al. 2010), which obviates the tension between theory and the observed non-detection of hot halos around massive spirals.

However, this still does not explain where the missing baryons from galaxies are. In fact, if a significant amount of gas was accreted in the cold mode at $z \gtrsim 2$, then it should have settled into the disk by today – this actually makes the missing baryons problem worse! In simulations with cold flow accretion onto massive spirals, powerful stellar feedback is therefore necessary to eject this material from the galaxy. This feedback has been an essential ingredient in simulations of galaxy formation for decades (Dekel and Silk 1986, Begelman and Fabian 1990, Slavin et al. 1993), and its efficiency per unit accreted mass has been steadily increased in simulations as the overcooling problem (Somerville and Primack 1999, Cole et al. 2000, Springel and Hernquist 2003) and the missing baryons problem manifested themselves (Brook et al. 2011, 2012, Agertz et al. 2013).

Another solution is to claim that the hot halo still contains the missing baryons from the galaxy, but that the density profile is much flatter than initially predicted. Since the dominant emission processes (thermal bremsstrahlung and/or CIE line emission, depending on temperature and metallicity) are both two-body processes, the X-ray emissivity goes as the square of the gas density. So the observed surface brightness is quite sensitive to the shape of the density profile, and density profiles that

decline with radius more gradually than isothermal profiles can effectively “hide” vast amounts of hot gas at large radii, where the gas shines at below the X-ray surface brightness detection threshold of modern telescopes. These profiles are produced in simulations by injecting entropy into the gas from a combination of stellar and/or AGN feedback (Toft et al. 2002, Maller and Bullock 2004, Sommer-Larsen 2006, Kaufmann et al. 2009, Guedes et al. 2011, Feldmann et al. 2013, Fang et al. 2013).

Finally, the third class of solutions is to invoke pre-heating (Mo and Mao 2002, Oh and Benson 2003, Benson and Madau 2003, Mo et al. 2005) to increase the entropy of the gas before it can accrete onto the galaxy (i.e., at $z \gtrsim 2$). In linear theory, if the gas is driven beyond the turnaround radius (twice the virial radius), then it will never fall into the galaxy, and this allows galaxies to form without most of their baryons, obviating the need for massive hot halos or powerful stellar feedback. Possibilities for sources of pre-heating include AGN feedback (Scannapieco and Oh 2004), accretion shocks (Dopita et al. 2011), and pop III stars (Madau et al. 2001); at present, observational constraints on the possible details of pre-heating mechanisms are not very secure, especially on the galactic scale. On cluster scales, there is stronger evidence of some sort of pre-heating, based both on the evolution of the cluster luminosity function (Kaiser 1991, Evrard and Henry 1991, Tozzi and Norman 2001) and on individual cluster entropy profiles (Tozzi and Norman 2001, Nagai et al. 2007, Cavagnolo et al. 2009). Finally, it is worth pointing out that many simulations smuggle-in a form of pre-heating by initializing their simulations with a significantly depleted galaxy. The most striking example of this is probably Mo, Mao, and White (1998), an extremely influential paper on disk formation that parameterizes the baryon fraction of the disk as m_d , with a default value of 0.05 - less than a third of the cosmological value of 0.17.

In this thesis, we present five studies of the properties of hot halos around nearby galaxies, with an emphasis on the late-type galaxies. In Chapter 2, we infer constraints on the amount of gas present in typical hot halos around L^* galaxies based on nondetections in both emission and absorption. We also examine in detail the amount of hot gas around the Milky Way using a variety of different constraints.

In Chapter 3, we announce the first detection of a hot gaseous halo around a spiral galaxy, for the giant spiral NGC 1961. We are able to detect this emission out to about 40-50 kpc, and infer a hot gas mass. In Chapter 4, we announce another detection of a hot gaseous halo; this one is around the giant spiral UGC 12591 and has similar properties to the halo around NGC 1961. In Chapter 5, we constrain the properties of hot halos around less massive galaxies than these giant spirals, by stacking observations from the ROSAT All-Sky Survey, and are able to estimate properties of hot halos around L^* spiral and elliptical galaxies. Finally, in Chapter 6, we present refinements to the method of flat-fielding X-ray images, which enables much more accurate measurements of very extended and faint X-ray emission. We apply this method to the elliptical galaxy NGC 720 and are able to constrain the properties of its hot halo much more accurately than previous work. Finally, in the conclusion (Chapter 7), we discuss the evolution of this field over the last five years, and prospects for the future.

1.1 References

- Agertz, O. et al. 2013, ApJ, 770, 25
- Begelman, M. C. and Fabian, A. C. 1990, MNRAS, 244, 26
- Benson, A. J. et al. 2000, MNRAS, 314, 557
- Benson, A. J. and Madau, P. 2003, MNRAS, 344, 835
- Birnboim, Y. and Dekel, A. 2003, MNRAS, 345, 349
- Bregman, J. N. and Glassgold, A. E. 1982, ApJ, 263, 564
- Bregman, J. N. and Pildis, R. A. 1994, ApJ, 420, 570
- Brook, C. B. et al. 2011, MNRAS, 415, 1051
- Brook, C. B. et al. 2012, MNRAS, 419, 771
- Brooks, A. et al. 2009, 694, 396
- Cavagnolo, K. W. et al. 2009, ApJS, 182, 12
- Cole, S. et al. 2000, MNRAS, 319, 168
- Crain, R. A. et al. 2013, MNRAS, 432, 3005
- Dahlem, M., Weaver, K. A., and Heckman, T. M. 1998, ApJ, 118, 401

Dekel, A. et al. 2009, *Nature*, 457, 451

Dekel, A. and Silk, J. 1986, *ApJ*, 303, 39

Dopita, M. A. et al. 2011, *Ap&SS*, 335, 345

Evrard, A. E. and Henry, J. P. 1991, *ApJ*, 383, 95

Fabbiano, G. and Trinchieri, G. 1984, *ApJ*, 286, 491

Fabbiano, G. and Trinchieri, G. 1987, *ApJ*, 315, 46

Faber, S. M. and Gallagher, J. S. 1976, *ApJ*, 204, 365

Fang, T., Bullock, J., and Boylan-Kolchin, M. 2013, *ApJ*, 762, 20

Feldmann, R., Hooper, D., and Gnedin, N. Y. 2013, *ApJ*, 763, 21

Forman, W. et al. 1979, *ApJL*, 234, 27

Forman, W. Jones, C., and Tucker, W. 1985, *ApJ*, 293, 102

Fukugita, M. and Peebles, P. J. E. 2006, *ApJ*, 639, 590

Guedes, J. et al. 2011, *ApJ*, 742, 76

Heckman, T., Armus, L., and Miley, G. K. 1990, *ApJS*, 74, 833

Hinshaw, G. et al. 2013, *ApJS*, in press

Kaiser, N. 1991, *ApJ*, 383, 104

Katz, N. et al. 2003, *ASSL*, 281, 185

Kaufmann, T. et al. 2009, *MNRAS*, 396, 191

Kereš, D. et al. 2005, *MNRAS*, 363, 2 Li, J.-T. and Wang, Q. D. 2013, *MNRAS*, 428, 2085

Madau, P., Ferrara, A., and Rees, M. J. 2001, *ApJ*, 555, 92

Maller, A. H. and Bullock, J. S. 2004, *MNRAS*, 355, 694

Mathews, W. G. and Baker, J. C. 1971, *ApJ*, 170, 241

Mathews, W. G. and Brighenti, F. 2003, *ARA&A*, 41, 191

McGaugh, S. S. 2005, *ApJ*, 632, 859

McKee, C. F. and Ostriker, J. P. 1977, *ApJ*, 218, 148

Mo, H. J., Mao, S., and White, S. D. M. 1998, *MNRAS*, 295, 319

Mo, H. J. and Mao, S. 2002, *MNRAS*, 333, 768

Mo, H. J. et al. 2005, *MNRAS*, 363, 1155

Molendi, S. and Pizzolato, F. 2001, *ApJ*, 560, 194

Nagai, D., Kravtsov, A. V., and Vikhlinin, A. 2007, ApJ, 668, 1
Nulsen, P. et al. 1984, MNRAS, 208, 185
Oh, S. P. and Benson, A. J. 2003, MNRAS, 342, 664
Oser, L. et al. 2010, ApJ, 725, 2312
Pedersen, K. et al. 2006, NewA, 11, 465
Peterson, J. R. et al. 2001, A&AL, 365, 104
Peterson, J. R. et al. 2003, ApJ, 590, 207
Rasmussen, J. et al. 2009, ApJ, 697, 79
Scannapieco, E. and Oh, S. P. 2004, ApJ, 608, 62
Silk, J. 2003, MNRAS, 343, 249
Slavin, J. D. et al. 1993, ApJ, 407, 83
Somerville, R. S. and Primack, J. 1999, MNRAS, 310, 1087 Sommer-Larsen, J. 2006,
ApJL, 644, 1
Springel, V. and Hernquist, L. 2003, MNRAS, 339, 289
Strickland, D. K. 2004a, ApJS, 151, 193
Strickland, D. K. 2004b, ApJ, 606, 289
Toft, S. et al. 2002, MNRAS, 335, 799
Tozzi, P. and Norman, C. 2001, 546, 63
Tucker, W. and David, L. P. 1997, ApJ, 484, 602
Tüllmann, R. et al. 2006a, A&A, 448, 43
Tüllmann, R. et al. 2006b, A&A, 457, 779
Wang, Q. D. et al. 1995, ApJ, 439, 176
White, S. D. M. and Frenk, C. S. 1991, ApJ, 379, 52
White, S. D. M. and Rees, M. J. 1978, ApJ, 183, 341

CHAPTER 2

Do Hot Haloes Around Galaxies Contain the Missing Baryons?

Note: This chapter is a reproduction of a paper of the same title, published in the *Astrophysical Journal*. The reference is Anderson, M. E. and Bregman, J. N. 2010, *ApJ*, 714, 320-331. The paper is copyright 2010, the American Astronomical Society, reproduced here under the non-exclusive right of republication granted by the AAS to the author(s) of the paper.

2.1 Abstract

Galaxies are missing most of their baryons, and many models predict these baryons lie in a hot halo around galaxies. We establish observationally motivated constraints on the mass and radii of these haloes using a variety of independent arguments. First, the observed dispersion measure of pulsars in the Large Magellanic Cloud allows us to constrain the hot halo around the Milky Way: if it obeys the standard NFW profile, it must contain less than 4-5% of the missing baryons from the Galaxy. This is similar to other upper limits on the Galactic hot halo, such as the soft X-ray background and the pressure around high velocity clouds. Second, we note that the X-ray surface brightness of hot haloes with NFW profiles around large isolated galaxies is high enough that such emission should be observed, unless their haloes contain less than 10-25% of their missing baryons. Third, we place constraints on the column density of hot haloes using nondetections of OVII absorption along AGN sightlines: in general they must contain less than 70% of the missing baryons or extend to no more than 40 kpc. Flattening the density profile of galactic hot haloes weakens the surface

brightness constraint so that a typical L_* galaxy may hold half its missing baryons in its halo, but the OVII constraint remains unchanged, and around the Milky Way a flattened profile may only hold 6 – 13% of the missing baryons from the Galaxy ($2-4 \times 10^{10} M_\odot$). We also show that AGN and supernovae at low to moderate redshift - the theoretical sources of winds responsible for driving out the missing baryons - do not produce the expected correlations with the baryonic Tully-Fisher relationship and so are insufficient to explain the missing baryons from galaxies. We conclude that most of missing baryons from galaxies do not lie in hot haloes around the galaxies, and that the missing baryons never fell into the potential wells of protogalaxies in the first place. They may have been expelled from the galaxies as part of the process of galaxy formation.

2.2 Introduction

The so-called “missing baryon problem” actually refers to two separate, but related, issues. First is the realization that the mean density of all detected baryons in the local universe accounts for less than half of the cosmological baryon density. The universal baryon fraction is known precisely from, among other sources, the Wilkinson Microwave Anisotropy Probe 5-year data (Dunkley et al. 2009): $f_b \equiv \Omega_b/\Omega_m = 0.171 \pm 0.006$ and $\Omega_b h^2 = 0.0227 \pm 0.0006$. But studies that attempt to count baryons in the low-redshift universe (Fukugita, Hogan, and Peebles 1998, Fukugita and Peebles 2004, Nicastro et al. 2005a) find that only about a tenth of this figure is observed in stars and cool gas in galaxies, while another third of Ω_b is divided between the hot plasma in the intracluster medium of galaxy clusters, and the cool intercluster gas detected in Ly α absorption lines (for more on the latter, see e.g. Penton, Stocke, and Shull 2004).

This result introduces the second missing baryon problem, which is the realization that most galaxies are severely baryon-depleted relative to the cosmological fraction (e.g. Bell et al. 2003). Based on cosmological simulations, the rest of the baryon budget is thought to reside in the “warm-hot intergalactic medium” (WHIM), at temperatures between 10^5 and 10^7 K (Cen and Ostriker 1999, Davé et al. 2001). This

is thought to include most of the baryons associated with galaxies, since the optically-luminous portion of galaxies is severely baryon-depleted relative to the universal fraction (Dai et al. 2010). Observational evidence for the existence of the WHIM has subsequently begun to accrue, primarily through detections with ultraviolet and X-ray telescopes of highly-ionized Oxygen absorption lines (for a recent review, see Bregman 2007).

These observations do not yet constrain the nature and distribution of the WHIM. Essentially all models agree that the WHIM exists as a cosmic filamentary web with some material in hot haloes around galaxies, but the mass fraction as a function of WHIM density is still unknown. Some numerical models (e.g. Davé et al. 2009) find a diffuse WHIM at $z = 0$ with most of its baryonic content lying outside of galactic haloes. However, many simulations either assume (e.g. Bower et al. 2006, Croton et al. 2006) or self-consistently derive (e.g. Cen and Ostriker 2006, Tang et al. 2009, Kim, Wise, and Abel 2009) the result that cosmologically significant reservoirs of baryons are embedded in hot haloes around massive galaxies ($kT \sim$ a few hundred eV). These hot haloes have also been modeled in separate theoretical calculations (Fukugita and Peebles 2006, Dekel and Birnboim 2006, Kaufmann et al. 2009) and have important consequences in the galaxy formation process such as differentiating the “blue cloud” from the “red sequence” (Dekel and Birnboim 2006, Bouché et al. 2009). Observationally, the evidence for hot haloes is still unclear: gas at these temperatures has been observed in X-rays around disk galaxies out to a few kpc beyond the disk (Strickland et al. 2004, Li et al. 2006, Li, Wang, and Hameed 2007), but this gas seems to result from starburst activity in the host galaxy instead of emerging as a byproduct of galaxy formation (Rasmussen 2009), and its inferred mass falls significantly short of theoretical predictions (Tüllmann et al. 2006).

In this paper, we examine the implications of this paradigm in more detail. If most of the baryons associated with galaxies reside in hot haloes around the optically luminous part, there are several predictions that can be tested with existing observations. We seek to answer two questions. First, how much baryonic mass resides in these haloes? And, what can these haloes and the second missing baryon problem

tell us about the process of galaxy formation?

To answer these questions, we present several independent lines of argument. First, we consider observational constraints on the density of hot gas around the Milky Way using the dispersion measure of pulsars in the Large Magellanic Cloud, as well as the ambient pressure around high velocity clouds, and the galactic soft X-ray background. We then extend this argument to other galaxies. We constrain the characteristic density and radius of hot haloes using existing X-ray observations of absorption along quasar sightlines and emission from galactic haloes. We also discuss the energetics of driving a galactic-scale wind, arguing that existing mechanisms for expelling most of the missing baryons seem incomplete. Finally, we discuss flattened density profiles and other prospects for resolving the second missing baryon problem. Before engaging these arguments, however, we will begin by describing the assumptions we use to model hot haloes, which are very similar to those of Fukugita and Peebles (2006).

2.3 On Milky Way Parameters

Much of the following discussion directly relates to the Milky Way or uses scaling relations to extrapolate from the Galaxy to other galaxies. Therefore we begin by discussing the parameters and scaling relations used throughout this paper.

Recent estimates for the virial mass of the Milky Way vary significantly. The consensus figure is roughly $2.5 \times 10^{12} M_{\odot}$ (Sakamoto, Chiba, and Beers 2003, Bellazzini 2004, Loeb et al. 2005, Li and White 2008), although a recent estimate was as low as $1.0^{+0.3}_{-0.2} \times 10^{12} M_{\odot}$ (Xue et al. 2008) while another (Reid et al. 2009) revised upwards the circular velocity by about 15 km s^{-1} , which should increase the mass estimate by more than 10%. Since larger masses for the Galaxy increase the amount of missing baryons, we will adopt a conservative virial mass estimate of $2.0 \times 10^{12} M_{\odot}$.

Regarding the virial radius of the Milky Way, there is scatter around a central value. Shattow and Loeb (2009) use $R_{\text{vir}} = 277$ for a solar distance from the Milky Way center of $d_{\odot} = 8.0 \text{ kpc}$, while the earlier work of Loeb et al. (2005) uses a virial radius of 207 kpc. Klypin, Zhao, and Somerville (2002) favor a virial radius of 258

kpc. We adopt an estimate for the virial radius of 250 kpc.

Another relevant Galactic parameter is amount of missing baryons. There is a surprisingly tight correlation between the total mass of a galaxy and its baryon fraction (McGaugh 2000, Dai et al. 2010), which we parametrize as $f_b = 0.04(M/2 \times 10^{12} M_\odot)^{1/2}$ for $M \lesssim 6 \times 10^{12} M_\odot$. The relation flattens significantly for galaxy groups and clusters, but those are outside the scope of this work. This implies a baryonic mass of $8 \times 10^{10} M_\odot$ for the Galaxy, which broadly comports with the estimate of a stellar mass of $5 \times 10^{10} M_\odot$ for the Galaxy (Binney and Tremaine 2008), with an additional $1 - 2 \times 10^{10} M_\odot$ in gas and dust. We assume the total mass is comprised of three components: the dark matter halo, the observable baryons in the galaxy, and the missing baryons in a hot halo. In this model, the mass of the missing baryons can be related to the observed mass M_0 of the other two components:

$$M_{\text{miss}} = \frac{0.17 - f_b}{0.83 + f_b} M_0 \quad (2.1)$$

For the Milky Way, the missing baryonic mass is therefore $3 \times 10^{11} M_\odot$. As a simple first-order model, we assume some fraction of the missing mass is distributed in a hot virialized halo around the disk. We can model the gaseous halo by assuming the Galactic dark matter halo obeys a Navarro, Frenk, and White (NFW) profile (Navarro, Frenk, and White 1997)). The missing baryons would follow the gravitational potential set by the dark matter, and therefore also obey the same density profile, except that the missing baryons do not extend into the disk, so we occasionally truncate the gaseous halo at an inner radius r_1 (this does not have a significant effect on our analysis). We examine many values for r_1 and for the NFW concentration parameter $C \equiv r_0/R_{\text{vir}}$, but our nominal choice will be $r_1 = 10$ kpc and $C = 12$ (so $r_0 = 20.8$ kpc).

Finally, the gaseous halo is assumed to be approximately isothermal at the virial temperature $kT_{\text{vir}} \sim 300$ eV. This is also the temperature of the hot gas observed around NGC 891 (Tüllmann et al. 2006), a near 'twin' of the Milky Way (van der Kruit 1984), and is similar to the 250 eV gas surrounding the Galactic disk observed

by Snowden et al. (1998). We expect the gas to have $Z \sim 0.2Z_{\odot}$ (using Anders and Grevesse 1989 abundances) following Cen and Ostriker (2006) and Rasumssen et al. (2009), and since similar metallicity is found in the hot intracluster medium even towards the virial radius of distant galaxy clusters (Maughan et al. 2008, Anderson et al. 2009). In computing the electron density of the halo, we take $\mu = 0.6$ and $\mu_e = 1.2$.

2.4 Constraints on the Milky Way Halo

2.4.1 LMC Pulsar Dispersion Measure

The strongest constraint we can derive for the hot halo around the Galaxy comes from the dispersion measure (DM) of pulsars in the Large Magellanic Cloud (LMC). The DM is defined as

$$DM = \int_0^d n(l) dl \quad (2.2)$$

where d is the distance to the pulsar and $n(l)$ is the free electron density along the line of sight. The Galactic free electrons (primarily from the so-called “Warm Ionized Medium”, or “WIM”) have been modeled as an exponential disk by, e.g, Gaensler et al. (2008). Integrating over this disk out to the LMC (and assuming a distance of 50 kpc from the Sun), we estimate the Galactic WIM contributes about $47 \text{ cm}^{-3} \text{ pc}$ to the DM of pulsars in the LMC. In a survey of the Magellanic Clouds, Manchester et al. (2006) measured the DM of 11 pulsars in the direction of the LMC. Three pulsars have DMs below $47 \text{ cm}^{-3} \text{ pc}$ (the values are 26, 28, and $45 \text{ cm}^{-3} \text{ pc}$), suggesting these sources are located within the Galaxy and randomly superimposed in front of the LMC. The next-lowest DMs for LMC pulsars are 66, 69, and $73 \text{ cm}^{-3} \text{ pc}$, which leads to a conservative estimate of $\approx 70 \text{ cm}^{-3} \text{ pc}$ for the minimum DM introduced by the free electrons towards the LMC. This is also consistent with the lowest DMs for pulsars in the same survey towards the Small Magellanic Cloud (SMC): 70 and $76 \text{ cm}^{-3} \text{ pc}$.

Subtracting $47 \text{ cm}^{-3} \text{ pc}$ of Galactic electrons leaves $23 \text{ cm}^{-3} \text{ pc}$ of electrons from other sources. The LMC has an ionized halo that may be responsible for much of this excess DM. One can assess the amount of ionized gas from the high-ionization absorption lines, such as OVI and CIV. The column densities of these ions in the LMC are comparable to the Milky Way (Howk et al. 2002). Considering that the metallicity is lower in the LMC, the inferred electron column is likely to be greater than that of the Milky Way. This implies that only a fraction of the excess DM is due to an ambient medium along the line of sight toward the LMC. However, since we do not know precisely the position of the LMC pulsars within its halo, in order to keep our estimate as conservative as possible we neglect the LMC electrons and assign all $23 \text{ cm}^{-3} \text{ pc}$ of excess DM to the hot halo of the Milky Way.

With an LMC distance of 50 kpc, $23 \text{ cm}^{-3} \text{ pc}$ corresponds to an average electron density between the Sun and the LMC of $\langle n \rangle \approx 5 \times 10^{-4} \text{ cm}^{-3}$. If the entire halo had this mean density out to the virial radius, it would contain $8 \times 10^{11} M_{\odot}$ – all of the missing baryons. This is an unphysical profile, however. We instead consider the standard assumption of an NFW profile for the missing baryons, with a concentration of 12 (Klypin, Zhao, and Somerville 2002), and find a halo obeying this profile and the DM constraint would have a mass of $1.2 \times 10^{10} M_{\odot}$, which is only 4% of the missing baryons from the Galaxy. However, since the Galactic electrons have a scale length of 1.8 kpc (Gaensler et al. 2008), we might reasonably argue that baryons within a few scale lengths are Galactic, not missing baryons. We therefore also consider an NFW profile with an inner truncation radius of 10 kpc, which is five scale lengths and therefore a very conservative estimate. Such a profile would have a mass of $1.5 \times 10^{10} M_{\odot}$, or just 5% of the missing baryons from the Galaxy.

As a consistency check, we compute the cooling time of this halo, using our halo temperature and metallicity with the cooling function of Sutherland and Dopita (1993). Our model assumes a stable halo, but we find the halo is probably not quite stable against cooling over a Hubble time. For $n = 5 \times 10^{-4} \text{ cm}^{-3}$ (the upper limit on the average gas density out to the LMC), the cooling time is $\tau \approx 10 \text{ Gyr}$. For the above NFW profile, then, $6.4 \times 10^9 M_{\odot}$ (40% of the total hot halo gas mass)

will cool within the age of the Galaxy and re-accrete onto the disk. To maintain self-consistency, the constraints above must be treated as upper limits on an even more rarefied halo, or there must be continuous re-heating of the halo gas by a wind or some other mechanism.

2.4.2 Other Constraints

We can compare the constraints on the Galactic hot halo from LMC pulsars to other constraints derived using a variety of methods. For reference, the DM upper limit on the electron density at $r = 50$ kpc is $n_e \leq 7.7 \times 10^{-5} \text{ cm}^{-3}$ for an NFW halo, or $n_e \leq 1.0 \times 10^{-4} \text{ cm}^{-3}$ if we truncate the profile at $r_{\text{in}} = 10$ kpc.

Our first comparison is to the result of Stanimirović et al. (2002), who used 21 cm observations of the Magellanic Streams to measure the radius and velocity width of clumps of neutral gas in the streams. From these measurements, they derive the internal pressure of the clumps, and equate this to the gas pressure of the hot halo (making the assumption that the clumps are confined entirely by gas pressure, and neglecting magnetic and turbulent pressure; this is therefore probably an upper limit on the gas pressure). They assume a halo temperature of 1×10^6 K; for our higher temperature of 3.5×10^6 K, the inferred upper limit on the hot halo electron density is $n(45 \text{ kpc}) \leq 9 \times 10^{-5} \text{ cm}^{-3}$ - consistent with the above constraint.

We also consider high velocity clouds, using UV measurements of the temperature and density of numerous partially-ionized metals and inferred gas pressures (Fox et al. 2005). This analysis suffers from a size-distance degeneracy, however, and so Fox et al. perform their analysis for several potential cloud distances. The lower distance (10 kpc) yields a lower upper limit, and the higher distance (100 kpc) yields a higher upper limit, but they seem to prefer a distance of 50 kpc based on a tentative association of some of these clouds with the Magellanic Stream. Using this distance and converting to our halo temperature, the upper limits on the electron density at this radius are approximately $n_e = 3 - 6 \times 10^{-5} \text{ cm}^{-3}$.

We also estimate density using the emission measure of the hot halo, as measured by Kuntz and Snowden (2000). They decompose the ROSAT soft X-ray background

into components, tentatively identifying the harder emission from an unidentified component as the hot halo of missing baryons. This emission has a temperature of about 3×10^6 K, close to our estimate of the hot halo temperature, and it appears to have a fairly uniform distribution across the sky. There is also a softer background component with a temperature of $T \sim 1 \times 10^6$ K, but this temperature is too cool to remain in a hot halo out to the virial radius, and its spatial distribution is not isotropic: it follows the Galactic disk.

The hotter component of the Galactic X-ray background has a best-fit emission measure of $2.1 \times 10^{-3} \text{ cm}^{-6} \text{ pc}$, for an assumed solar-metallicity gas. This corresponds to an NFW hot halo of total mass $3.1 \times 10^9 M_\odot$ (1.0% of the missing baryons from the Galaxy), or a truncated NFW halo of mass $4.1 \times 10^9 M_\odot$ (1.4% of the missing baryons). Using our assumed metallicity of $0.2Z_\odot$, we recalculated the emissivity using the XSPEC utility with an APEC model (Smith et al. 2001) and all other parameters matched to the observations in Kuntz and Snowden (2000), which makes these halo masses increase by 190%. The corrected halo masses match very closely with the previous constraints: $5.9 \times 10^9 M_\odot$ (2.0% of the missing baryons) for an NFW halo, $7.8 \times 10^9 M_\odot$ (2.6% of the missing baryons) for a truncated NFW halo. This halo has a cooling time larger than 10 Gyr for $r > 15$ kpc (no truncation) or $r > 18$ kpc (with truncation).

We summarize all of these constraints in Table 2.1. Upper limits on the total hot halo mass, electron density at a radius of 50 kpc, and fraction of the missing baryons from the Galaxy in the halo, are computed for each constraint. The ranges on some parameters reflect ranges in the original constraint combined with fits to both a truncated and a non-truncated NFW profile. We also present the upper limits in terms of a flattened, power-law profile, discussed further in section 2.8.

Finally, we also include a lower limit on the Galactic hot halo in addition to the upper limits previously noted. Blitz and Robishaw (2000) measure HI mass for the dwarf spheroidal galaxies in the Local Group, and use these masses to place upper limits on the amount of ram-pressure stripping that these galaxies could have experienced. This leads to an upper limit on the density of the Local Group intragroup

medium, corresponding to $n_e \geq 2.5 \times 10^{-5} \text{ cm}^{-3}$. The Galactic hot halo must be denser than the ambient medium around it, so this figure also serves as a strict lower limit on the density of the hot halo.

Our picture of the Milky Way hot halo is therefore a halo obeying an NFW profile, and containing several $\times 10^9 M_\odot$ of material. It extends out to at least 50 kpc, although its density declines to the IGM density by a radius of 80 kpc or so. Such a halo is fully consistent with all of the constraints we have identified, but it does not contain more than a few percent of the missing baryons from the Galaxy.

2.5 OVII Absorption in Other Galaxies

The previous techniques will not work for constraining the haloes of most other galaxies, since we cannot measure the DM of distant extragalactic pulsars, and even if we could, we cannot reliably separate the dispersion from intergalactic electrons from the dispersion caused by the galaxy's halo of missing baryons. However, the integrated column density of free electrons from the hot haloes still will have a measurable effect in other galaxies through OVII absorption features along quasar sightlines.

At the temperatures of the haloes around large galaxies (a few hundred eV), intergalactic Oxygen exists in the highly ionized OVII or OVIII states, of which the former can be detected through its 21.60 \AA $K\alpha$ absorption line if superimposed in front of a bright background X-ray source. Zero-redshift OVII absorption has been detected several times, and attributed variously to the Local Group IGM (Nicastro et al. 2002, Rasmussen, Kahn, and Paerels 2003) or to a hot halo around the Milky Way (Wang et al. 2005, Fang et al. 2006, Bregman and Lloyd-Davies 2007). In addition to the local detections, there are a few published claims of nonlocal OVII absorption lines (Mathur, Weinberg, and Chen 2003, Nicastro et al. 2005b, Buote et al. 2009), but these have yet to yield any significant constraints (see, e.g., Kaastra et al. 2006).

One successful approach comes from a statistical argument by Fang et al. (2006). They studied 20 active galactic nucleus (AGN) sightlines without any detections of nonlocal OVII absorption and argued based on the number density of L_* galaxies and of Local Group-sized galaxy groups that most of the missing baryons must be

restricted to haloes around the former; if these haloes extended to megaparsec scales, it would be statistically unlikely ($\sim 3.7\sigma$) not to have seen them in OVII absorption.

We extend this result by examining observations along specific sightlines, and broadening the search to galaxies smaller than L_* . We select critical sightlines with the best combination of redshift interval and high signal-to-noise observations discussed by Fang et al. (2006), and count the intervening galaxies along the sightline, focusing on the twelve AGN that fall within the coverage of the Sloan Digital Sky Survey (SDSS) Seventh Data Release (Abazajian et al. 2009). For each of these twelve sightlines, we searched the SDSS archive for foreground galaxies whose virial radii overlapped the sightline. If such a galaxy is present, but no OVII absorption is detected at its redshift, then we can place an upper limit on the OVII column from the galaxy.

We searched within a $4^\circ \times 4^\circ$ region around each AGN for SDSS galaxies with photometric redshifts less than the AGN’s redshift. We also found a known Lyman- α absorption line at each of these photometric redshifts to help verify the photometric redshifts; see Table 2.2. The photometric redshift was used to compute the distance to each galaxy and the angular distance between each galaxy and the AGN sightline. The virial radius for each galaxy was estimated, and galaxies were selected whose virial radii were larger than their distance to the AGN sightline. We inferred the virial radius by computing the SDSS g-band galaxy luminosity and assuming a scaling relation between luminosity and virial radius, normalized to the Milky Way ($L_g = 1 \times 10^{10} L_\odot$, $R_{\text{vir}} = 250$ kpc). We consider two forms for the scaling relation. The first, (scaling relation A), is based on empirical scaling between the baryon fraction (taken to be a proxy for luminosity) and the total mass of a galaxy (McGaugh 2000, Dai et al. 2010): $f_b \propto L/M \propto M^{0.5}$, so $L \propto M^{3/2}$ and $R \propto L^{2/9}$. This seems to represent the data over many orders of magnitude in galaxies with disk components (McGaugh et al. 2005). For elliptical galaxies, the well-established ”tilt” in the fundamental plane (e.g. Tortora et al. 2009) has the baryon fraction scaling in the opposite direction with mass: $M/L \propto L_B^{1/4}$, so $L \propto M^{4/5}$ and $R \propto L^{5/12}$ (scaling relation B). We classify each galaxy as early-type or late-type using its SDSS $u - r$ color, with early-type

galaxies having $u - r > 2.2$ and late-type galaxies having $u - r < 2.2$ (Strateva et al. 2001, Mateus et al. 2006). For the former, we apply scaling relation B, and for the latter we apply scaling relation A to estimate the virial radius.

Using this technique, out of our sample of twelve sightlines we find three with nearby galaxies whose virial radii encompass the sightline: 3C273 ($z_{em} = 0.158$), Q1821+643 ($z_{em} = 0.297$), and Ton 1388 ($z_{em} = 0.177$). For each intervening galaxy, the mass and the baryon fraction are estimated, as well as the expected mass in missing baryons. We assume these baryons are distributed in a halo around the galaxy, extending to the virial radius, and that they obey an NFW density profile (using several values for the concentration parameter between $c=9$ and $c=17$, as in Klypin, Zhao, and Somerville 2002). The expected electron column contributed by each galaxy is computed by integrating along the AGN sightline as it passes through the halo. This gives us $N_{e,NFW}$, the electron column expected from a hot halo in hydrostatic equilibrium around each galaxy.

In Table 2.2, we compare this column to the minimum detectable electron column inferred from Fang et al. (2006). For each sightline, they provide the expected width of a 3σ detection of OVII absorption, from which we compute the OVII column that would produce this line using $N(\text{O VII}) = 3.48 \times 10^{14} \text{ EW}$, where N is in cm^{-2} and EW in $\text{m}\text{\AA}$ (Bregman and Lloyd-Davies 2007). The electron column is estimated from the Oxygen column by adapting the conversion derived by Bregman and Lloyd-Davies (2007):

$$N_e = 2.9 \times 10^{20} \left(\frac{N_{OVII}}{10^{16} \text{ cm}^{-2}} \right) \left(\frac{f}{0.5} \right)^{-1} \left(\frac{Z}{0.2Z_\odot} \right)^{-1} \text{ cm}^{-2}$$

where f is the O VII ion fraction (we use $f = 0.5$) and assuming gas of $Z = 0.2Z_\odot$. This yields $N_{e,th}$, the minimum (threshold) electron column that can be detected at 3σ using OVII absorption. These assumptions are conservative. The fraction of Oxygen in the O VII state is close to unity for temperatures between $3.5 \times 10^5 \text{ K}$ and $1.5 \times 10^6 \text{ K}$, which are the expected temperatures for hydrostatic haloes around

most of our selected galaxies. This is probably a low estimate since hot intracluster medium gas typically has a metallicity of about $Z = 0.2Z_{\odot}$, even towards the virial radius of distant galaxy clusters (Maughan et al. 2008, Anderson et al. 2009). For both f and Z , choosing less conservative values reduces $N_{e,th}$, which would imply even stronger expected detections than those listed in Table 2.2.

Table 2.2 shows that two of these galaxies should have been detected in existing observations. The nondetection of these absorption lines implies real electron column must be much smaller than the $N_{e,NFW}$ value derived above. In light of our conservative assumptions about unknown parameters like f and Z , the most reasonable way to reduce $N_{e,NFW}$ below $N_{e,th}$ is to bring down the mass of the haloes. The radius could also be increased to several times the virial radius, which would reduce the density but increase the path length, but this would also increase the number of galaxies whose haloes intersect with the AGN sightlines. For example, if R is doubled to twice the virial radius, we would expect $4\times$ as many absorbers, yielding about 6 detections of 3σ or more. Reducing the halo radius below the virial radius increases the density, which worsens the pulsar dispersion measure inconsistencies in the previous section and leads to a problem of excess X-ray emission (see next section). Also, our strongest constraints come from two galaxies of 5-10% of L_* , each within about 40 kpc of the sightline. As long as these galaxies have haloes at least 40 kpc in extent, we can put a 3σ upper limit on the mass in the halo of 70% of the missing baryons, so reducing the radius over a reasonable range will not have much of an effect. Thus the only remaining option is to lower the mass, below 70% of the missing baryons at 3σ .

This is a stronger conclusion than Fang et al. (2006) reached in their analysis. They were able to rule out haloes with columns greater than $N_{e,th}$ with radii of 1 Mpc at 3.7σ , but they did not rule out such haloes with $R \sim 100$ kpc. Their statistical analysis predicted only 0-1 L_* galaxies to fall within 100 kpc of a sightline in their entire sample, so their method did not really constrain the parameters of haloes of this size. Indeed, we find only one L_* galaxy, and it lies at an impact parameter of 135 kpc, in agreement with their prediction. However, there are several $0.1L_*$ galaxies

closer than 100 kpc, and these are critical in deriving our constraints. At larger radii, our predictions are also consistent with not observing excess OVII absorption in a sightline that passes within 350 kpc of M31 (Bregman and Lloyd-Davies 2007).

2.6 X-Ray Surface Brightness

The failure to detect extended hot gaseous halos from X-ray emission (e.g. Strickland et al. 2004, Li et al. 2006, Li, Wang, and Hameed 2007) places particularly strong constraints on the fraction of missing baryons present around galaxies. Unlike the other constraints in our paper, this constraint has been known for at least a decade; see, for example, Benson et al. (2000) and Wu, Fabian, and Nulsen (2001). The former paper worked out X-ray luminosities of hot haloes around massive late-type galaxies, and found these should be some of the most luminous X-ray sources in the local Universe. The latter paper explored the unresolved X-ray background, and noted that the predicted contribution to the background from hot haloes around galaxy groups was an order of magnitude larger than the unresolved background. Since these haloes are so large and diffuse, however, measurements of both the luminosity and the unresolved background number density require extrapolation out to large radii and low surface brightness. Therefore, we instead consider this discrepancy in terms of X-ray surface brightness, since this quantity is directly measurable and can easily be used to establish observational constraints on these hot haloes.

We calculate the X-ray emissivity for three different galaxy masses: $3 \times 10^{11} M_{\odot}$, $8 \times 10^{11} M_{\odot}$, and $2 \times 10^{12} M_{\odot}$. For each mass, the hot halo is assumed to obey an NFW profile with a concentration $C=12$ (the profile is not very sensitive to this value), and the hydrostatic equilibrium temperature is computed. The emission from the baryons in the hot halo is modeled with an optically thin emission model (the APEC model) with photoelectric absorption from Galactic neutral Hydrogen of column density $N_H = 3 \times 10^{20} \text{ cm}^{-2}$. For each halo, we use the WebSpec utility to compute the flux F_1 over an energy range of 0.1-6 keV on the ACIS-S instrument. The metallicity was set to $0.2Z_{\odot}$.

We estimate a typical Chandra ACIS background of $1.3 \times 10^{-12} \text{ erg s}^{-1} \text{ arcmin}^{-2}$

from, e.g., Bregman et al. (2009), which is converted into a flux (in $\text{erg s}^{-1} \text{cm}^{-2}$) for each halo by dividing by the telescope effective area A determined using XSPEC. Assuming halo emission can be detected once it reaches a surface brightness of about 30% of the background, this background can be combined with the flux F_1 from above and with the definition of normalization in the APEC model (at $z = 0$, $\text{norm} = n^2 R \times 10^{-14} / 4\pi$) to calculate the minimum detectable emission measure for gas in each halo:

$$(n^2 R)_{\text{min}} = \frac{4\pi \times 0.3 \times (1.3 \times 10^{-12} \text{ erg s}^{-1} \text{ arcmin}^{-2})}{10^{-14} \times F_1 \times (8.5 \times 10^{-8} \text{ arcmin sterad}^{-1}) \times A} \text{ cm}^{-5}$$

The results of this analysis are presented in Table 2.3.

We integrate along columns through the halo to find the expected EM at each projected radius from this galaxy. In Figure 2.1, we compare the expected emission measure to the minimum detectable EM for galaxies of various concentrations at all three masses. While the emission within the disk may be obscured by or confused with the disk, the haloes can be detected out to many times the radius of the thick disk in all three cases. A Milky Way-sized halo, for example, should be detectable in emission out to about 75 kpc, and even a hot halo around a galaxy of total mass $3 \times 10^{11} M_{\odot}$ should be detectable out to twice the radius of the thick disk.

These haloes should be very obvious in X-rays around all reasonably-sized galaxies, and yet the observed X-ray haloes are confined to the immediate vicinity of the thick disk. This is a serious problem for the paradigm of hot haloes full of missing baryons. The missing baryon electron density must be smaller than the model predicts; the upper limit on the mass of the halo to bring the emissivity at the edge of the thick disk down to the detection threshold is between 11% and 24% for the three different masses. In general larger galaxies have stricter upper limits, since they have more mass in their haloes. Therefore, this constraint can be improved further using deep X-ray observations of very massive isolated disk galaxies.

Observations can also constrain the density profile of the hot halo; based on the predicted EM, in the detectable region outside the thick disk our NFW profiles

roughly obey a β model with $\beta \approx 0.87$. This is slightly steeper than the profiles of hot gas in early-type galaxies, which have $\beta \approx 0.4 - 0.6$ (e.g. Brown and Bregman 2001). In Figure 2.5, we also consider a flatter density profile as discussed further in section 2.7. This profile has $\beta \approx 0.36$ in the observable region, shallower than observed gas profiles. However, since observed values of β fall in between these two profiles, it seems likely that real haloes are neither as steep as NFW profiles nor as shallow as the flattened profiles of section 2.7.

2.7 Energetics for Gas-Rich Late-Type Galaxies

Having established some observational constraints on hot haloes around the Milky Way and other large galaxies, we comment briefly on the overall energetics of expelling the missing baryons from galaxies (either into virialized haloes or outside of the entire system) using large-scale galactic winds. These winds are a common component of the current generation of semi-analytic models of galaxy formation, required to reconcile the dominant “cold flows” picture (e.g. Kereš et al. 2005), which assumes galaxies accrete the cosmological fraction of baryons, with the observed severe baryon depletion in nearly every galaxy today. Winds are also observed directly at $z \sim 2-3$ in galaxies undergoing periods of intense star formation and AGN activity (e.g. Pettini et al. 2001, Geach et al. 2005, Law et al. 2007), but appear insufficient to expel the missing baryons (Martin 2005).

Such winds come in two varieties: “energy-driven” and “momentum-driven”. The former carry kinetic energy proportional to the energy injected into the interstellar medium by supernovae and/or an AGN, while the latter (Murray, Quataert, and Thompson 2005) carry momentum proportional to the momentum injected into the ISM by supernova shocks and radiation pressure from stars and AGN. In both cases, however, the constant of proportionality (the “efficiency” of the coupling) has to be large in order to drive enough of a wind to expel the missing baryons. For example, recent simulations by Dutton and van den Bosch (2009) require 25% efficiency for energy-driven winds or 100% efficiency for momentum-driven winds to expel the missing baryons, and they neglect cooling of the ejected baryons, which can result

in the baryons re-accreting onto the disk in a galactic “fountain” and can therefore increase the required ejection energy by several times. Oppenheimer and Davé (2008) use momentum-driven winds with an equivalent efficiency of 200% with respect to the energy available to energy-driven winds. Such high efficiencies are an attractive feature of momentum-driven winds, but we note that the theory still needs to be fully developed and observationally supported. In particular, it is not yet clear that momentum-driven winds are stable enough to drive most of the baryons from a galaxy instead of fragmenting into a less efficient multiphase wind, although further research may illuminate this question.

Regardless, while requiring such a large efficiency may already be problematic, driving these winds also leads to some challenging conclusions. There are two major sources of energy/momentum available for this purpose: supernovae and accretion onto a supermassive black hole (SMBH). In addition, one or both of these can be triggered by major mergers, but mergers primarily increase the amount of energy in these two reservoirs rather than adding a new component. (Mergers also do release gravitational energy, but we are not aware of any way to couple efficiently this energy to the baryons in galaxies.) The maximum energy or momentum available in a galaxy from supernovae should be proportional to the stellar mass of the galaxy, and the maximum energy/momentum available from the AGN should be proportional to the mass of the black hole. Since the black hole mass is proportional to the galactic bulge mass (Marconi and Hunt 2003), which is dominated by stars, this would seem to suggest that the total energy available to eject baryons is related to the stellar mass of the galaxy, which would lead naturally to a baryonic Tully-Fisher relation (McGaugh 2005).

However, this picture also leads to the prediction that variations in the energy sources should affect the baryon fractions of galaxies, and this is not observed. The observations show that, at least for disk galaxies, the baryon fraction seems to depend almost solely on the circular velocity at large radius (a proxy for total mass; McGaugh 2005, Stark, McGaugh, and Swaters 2009). But if the baryon fractions are caused by supernova- and AGN-driven winds, then galaxies with fewer supernovae or smaller

SMBHs should have higher baryon fractions than galaxies with more supernovae or larger SMBHs. So, for example, galaxies with no central black hole such as M33 (Gebhardt et al. 2001) should have higher baryon fractions than galaxies with the same total mass but a large SMBH, when in reality M33 lies on the BTF.

If the supermassive black hole (SMBH) in a galaxy is responsible for a significant amount of the baryon loss, there should be a relationship between the baryon fraction and the galaxy type. Equivalently, in a fit to the baryon Tully-Fisher relationship, galaxies with relatively larger SMBHs (earlier-type galaxies) would lie below the normal scaling relation if AGN activity from a SMBH expels gas. To search for this effect, we examine the sample of galaxies from McGaugh (2005) and use the galaxy type as a proxy for the relative importance of a SMBH. When the galaxy type was not determined to sufficient precision, the galaxy was not used for this comparison. For the 60 galaxies with good galaxy types, there is no trend of the residual with galaxy type (Figure 2.2). We estimate that we could have detected a trend of about 20% or larger in the residual between Sb and Sd, suggesting AGN activity is responsible for removing a mass equivalent to less than 20% of the remaining baryons. Since present-day baryon fractions are typically 5-30% of the cosmological value, this corresponds to AGN removing no more than 1-5% of the original baryonic content from galaxies over their lifetime.

We apply a similar analysis to examine the effect of supernova energy on baryon fraction. We computed the ratio of stellar mass to total baryon mass in the sample of gas-rich disk galaxies of Stark, McGaugh, and Swaters (2009). To estimate the total baryon mass, we assumed an NFW halo potential and applied the relation $M_{200} = 2.3 \times 10^5 v_f^3 h^{-1} M_\odot$ for v_f in km s^{-1} from Navarro (1998). The stellar mass is computed in Stark et al. for each galaxy from stellar population modeling. In Figure 2.3, we compare this ratio to the circular velocity at large radius v_f for each galaxy. The figure shows these galaxies are all significantly depleted in stellar fraction compared to average disk galaxies, so there should be significantly less supernova energy available to remove baryons in these galaxies. We examined whether the residuals about the BTF are correlated with the degree to which the galaxies are gas-

dominated. No such correlation was found, using both parametric and non-parametric tests, which suggests that supernovae from stars in these galaxies do not appear to be responsible for removing most of the baryons.

An alternative explanation of the lack of BTF correlation with galaxy type or gas fraction can be inferred from the result of Navarro and Steinmetz (2000). They noted that the remarkably low scatter in the I-band Tully-Fisher relation can be explained through galaxies with different disk mass fractions scattering along the Tully-Fisher relation instead of perpendicular to the curve. The idea is that the disk mass actually contributes to the measured rotation speed in addition to the halo mass, so if the disk mass increases, the rotation speed will also increase, regardless of the halo mass. This explanation does not apply directly to our argument, however. Neither the gas fraction nor the galaxy type correlates with the disk mass, so there is no reason to expect the scatter between these quantities and the rotation speed to follow the BTF. However, if supernovae or AGN accretion has expelled a dynamically significant mass of baryons from one of these galaxies, we would expect the rotation speed of the galaxy to decrease near the disk, in an analogous effect to the result of Navarro and Steinmetz. To minimize this concern, the rotation speeds used in our analysis are inferred from 21 cm emission at large radii, outside of the disk, where the rotation curve has flattened (for details, see McGaugh 2005 and Stark et al. 2009). At these radii (tens of kpc), the dark matter dominates the gravitational potential, and so the efficiency of expelling baryons should not affect the measured rotation velocity. Thus we do not believe correlations between baryon fraction and rotation velocity affect our claim that AGN accretion and supernovae are not primarily responsible for expelling the missing baryons.

A simple way to understand the argument of this section is to consider two galaxies: NGC 4183 and NGC 4138. Figure 2.4 shows images of these two galaxies taken from the Sloan Digital Sky Survey, and in Table 2.4 we present measured or inferred parameters for the baryons in these galaxies, based on the results in NED and in Stark, McGaugh, and Swaters (2009) as described above.

These two galaxies are similar in mass, but very different in composition. NGC

4138 is a bulge-dominated lenticular galaxy with 91% of its baryons in stars. NGC 4183 is a late-type disk galaxy with only 31% of its baryons in stars (the other 69% are in gas). Our sample contains more extreme cases of both gas-rich and gas-poor galaxies than these, but these two present a useful contrast in both gas fraction and galaxy type, with little difference in total mass. Since NGC 4138 has such a high stellar fraction and such a prominent bulge, it must have been able to use its AGN and its supernovae to produce a significant galactic wind and drive out more of its baryons than NGC 4183, and yet the latter has a smaller baryon fraction than the former.

More quantitatively, if we assume a negligible SMBH in NGC 4183 and assume its $2.1 \times 10^9 M_\odot$ of stars has produced 1.8×10^7 supernovae¹, we can estimate the average mass each supernova needs to eject to account for the missing baryons. From eq. 2.1, we calculate NGC 4183 is missing $7.9 \times 10^{10} M_\odot$ of baryons, so each supernova needs to heat and eject $4400 M_\odot$ from the galaxy. It already requires 180% of the supernova energy (i.e., 1.8×10^{51} erg) to heat $4400 M_\odot$ to the virial temperature of NGC 4183, and a comparable amount of energy to drive this material up into a halo from the gravitational potential of the disk, for a total supernova energetic efficiency of 350% (or more, if cooling is included). These extreme efficiencies may pose a problem even for momentum-driven winds, and this leads to our suggestion that the classical picture of AGN- and supernova-driven winds is incomplete for explaining the missing baryon problem.

2.8 The Effect of Flatter Density Profiles

There exists one significant modification to the standard picture of hot haloes that must be considered. Significant pre-heating of the baryons can lead to excess entropy in the halo gas, which causes it to deviate from the dark matter NFW profile and obey a flatter density profile instead (e.g. Mo and Mao 2002, Kauffmann et al. 2009). We

¹This is estimated by integrating over a Chabrier (2003) initial mass function, letting stars with $M > 8 M_\odot$ explode in a core-collapse supernova, and applying a 20% correction (Mannucci et al. 2005) to include type Ia supernovae

approximate the relaxed high-entropy profile in Figure 6 of Kaufmann et al. (2009) with a power-law:

$$\rho(r) = \rho_0 \left(\frac{r}{r_{\text{vir}}} \right)^{-p} \quad (2.3)$$

with $p = 0.9$. We recomputed the Galactic constraints for this flattened density profile (see Table 2.1). The DM constraint is significantly relaxed: the halo can hold 58% of the missing baryons before reaching the LMC DM limit of $23 \text{ cm}^{-3} \text{ pc}$. This limit is dramatically weaker compared to the NFW upper limit of about 5%. On the other hand, the pressure confinement constraints based on high velocity clouds and Magellanic Stream clouds place an upper limit of 6 – 13% of the missing baryons ($2 - 4 \times 10^{10} M_{\odot}$), which still fails to resolve the second missing baryon problem. A flattened profile containing all the missing baryons also violates the EM measurement of Kuntz and Snowden (2000). They measured an EM of $2.1 \times 10^{-3} \text{ cm}^{-6} \text{ pc}$ (which becomes approximately $7.7 \times 10^{-3} \text{ cm}^{-6} \text{ pc}$ for our assumed metallicity), while the flattened halo would have an EM of $1.9 \times 10^{-2} \text{ cm}^{-6} \text{ pc}$. Thus the picture we described in Section 2.3.2 remains - the Galaxy has a limited hot halo of some form that declines to the IGM density, but it does not contain most of the missing baryons from the Galaxy.

The extragalactic emission measure constraint (X-ray surface brightness) is also significantly affected by a flatter profile. Kaufmann et al. claim their profile lowers the X-ray luminosity of the inner region of the halo by a factor of ~ 70 . We computed the emission measure for the flattened power-law profile (Figure 2.5) and indeed found the EM lowered by a factor of 70-300 at the edge of the thick disk. However, the emissivity of the gas per unit mass increases, by about 75%, because the flattened profile is hotter than the NFW profile. Finally, since the slope of the profile is flatter, the emission can remain significant out to larger radii. The result of these combined effects is that the emission from within the disk is lowered significantly, but the emission remains flatter, so at larger radii it is still visible. As Figure 2.5 shows, smaller galaxy haloes are not detected, but a Milky Way-sized galaxy would still

have a halo detectable out to 70 kpc or so, and a larger galaxy would have a hot halo visible to even larger radii. Searching for X-ray emission at significant fractions of the virial radius around large, isolated disk galaxies is therefore a good way to constrain these flattened profiles.

This claim differs from the conclusions in Kauffmann et al. (2009). They argue that the surface brightness of their high-entropy halo is consistent with observations, while we argue the emission from such a halo should be detectable tens of kpc away from the galaxy disk. This distinction is primarily due to different assumptions about the baryon fraction. They initialize their models 10 Gyr ago with a baryon fraction $f_b = 0.1$, while we assume galaxies form with the cosmological baryon fraction $f_b = 0.17$. This gives our halo a density 1.7 times larger than theirs, which yields a factor of three in the emission measure. Correcting for this effect reconciles our respective predictions with each other, although we disagree on the interpretation. To understand the second missing baryon problem, however, one must begin with the cosmological baryon fraction, and then explain the fate of all the baryons originally associated with the galaxy over cosmic time; thus we argue that flattened density profiles do not hide all of the missing baryons from galaxies.

The flattened profiles are still detectable through their OVII absorption as well. At large fractions of the virial radius, a flattened profile has a higher density compared to an NFW profile, so sightlines that only pass through a small section of the hot halo produce larger absorption lines if the profile is flatter. Sightlines that pass closer to the central galaxy produce roughly equivalent sightlines to the NFW halo. We computed the absorption lines expected for our seven absorbers, and displayed the results in Table 2.2 above. In the former regime, the small ($< 1\sigma$) detections predicted for NFW haloes increased moderately if the profile was flatter and the $1 - 2\sigma$ detections became roughly 3σ detections. In the latter regime, the two large $> 4\sigma$ detections both decreased slightly. Since a 3σ absorption line could be observable, the flattened density profiles do not seem to solve the OVII absorption issue, and actually make the problem worse.

In general, flattened profiles can contain a larger fraction of the missing baryons

from galaxies than NFW profiles. For the Milky Way, a flattened profile can meet the LMC pulsar DM constraint with up to 58% of the missing baryons in the hot halo instead of 5%, although the pressure constraints bring this figure down to 6-13%. The X-ray emissivity argument constrains the mass of these haloes around other large galaxies to about 59% of the missing baryons from these galaxies. This is a factor of three more mass than the NFW profile can hold without violating surface brightness limits. The OVII argument also limits the mass of these haloes in the same way it limits the mass of NFW haloes, with an independent upper limit of about 70% on the fraction of the missing baryons that can reside in the halo. Flattened haloes are therefore an improvement, but they do not seem to be a perfect solution.

If the halo is driven from the galactic disk by a superwind, then the energetics problem we highlighted still applies. The flattened profile puts some additional constraints on the wind, however. In Kaufmann et al. (2009), the high entropy that keeps the profile flat is assumed to be injected at redshift 2 or 3. This pushes back the energetics problem 10 Gyr or so, but still does not quite resolve it. On the other hand, if the halo has existed in its present form for 10 Gyr, perhaps the gas simply never fell into the disk in the first place. We will discuss this idea further in the conclusion.

2.9 Discussion and Conclusion

The above arguments place separate and independent constraints on the baryons around galaxies. These arguments lead us to two general conclusions:

First, the majority of the missing baryons from spiral galaxies do not lie in hot haloes around these galaxies. For the Milky Way, the combination of observed constraints suggests that the common assumption of a hot halo of missing baryons obeying an NFW density profile with a concentration of 12 can only hold 2-3% of the baryons missing from the Galaxy. The emission measure of theoretical haloes places an upper limit of 11-24% on the fraction of the missing baryons that can reside in similar hot haloes around other large galaxies.

The analysis of O VII sightlines constrains hot haloes around a wider mass range

of galaxies. The free electrons in these haloes should be detected through O VII absorption in the spectra of several strong background X-ray sources. This absorption has not been observed, so the column must be several times smaller than the prediction. This would occur if the halo has a radius of $\lesssim 40$ kpc or if its mass is reduced by factor of 3-4. The former case leads to further difficulties. The density would be 0.01-0.1 cm^{-3} , so the halo would cool quickly and should be far too bright in X-rays. The latter case is possible, but such a rarefied halo would not account for the missing baryons from galaxies.

The proposal of hot haloes obeying flattened density profiles instead of NFW profiles was also considered. Around the Milky Way, a typical flattened profile could hold 6 – 13% of the missing baryons without violating any of our constraints. The emission measure of hot haloes with a flattened profile around other large disk galaxies sets an upper limit on the mass of about 59% of the missing baryons from these galaxies. The OVII constraint is approximately unchanged, which also limits the mass and extent of flattened density profiles.

Collectively, these arguments rule out the existence of a large reservoir of hot gas around large galaxies that could account for the missing baryons. The most likely scenario is the presence of a small gaseous halo around L_* galaxies containing $10^9 - 10^{10} M_\odot$ and extending to a radius of 50-100 kpc. This picture also satisfies pulsar DM observations, measurements of the pressure around high-velocity clouds, and X-ray surface brightness constraints. If such a limited halo exists, it is not the primary reservoir of the missing baryons from galaxies.

Our second general conclusion is that galaxies do not expel their baryons primarily by galactic winds driven by supernovae or AGN activity. We infer this from the lack of correlation between missing baryon fraction in galaxies and the stellar fraction or bulge mass in the galaxy. Since the latter observables are taken to be proxies for the energy available in a galaxy from supernovae or from a supermassive black hole, the lack of correlation implies these energy sources do not explain the observed missing baryon fractions.

Additionally, there are galaxies where most of the baryons are still gaseous yet

they have the same missing baryon fraction as their counterparts where the baryons are nearly all in stars. To explain the baryon depletion of some of these star-poor, gas-rich galaxies, the energy required is several times that expected from galactic winds powered by supernovae.

It is beyond the scope of this paper to model the remaining mechanisms for resolving the second missing baryon problem, but we can comment on the possible solutions. Based on our conclusions, a straightforward way to account for the missing baryons from galaxies is if the baryons never fell into galaxies in the first place. This would require pre-heating before or just as the galaxy is forming. The most obvious source of non-gravitational energy at these times is early supernovae, whose contribution to heating would be much more effective in the epoch of protogalaxies (when potential wells are shallow) than it is today. These early supernovae could even come from the hypothetical “Population III” stars, which have the advantage of a top-heavy IMF, so this process would not convert much of the missing baryonic mass into low-mass stars.

If the missing baryons never accreted onto protogalaxies because of early supernova heating, there are several observable consequences in addition to resolution of the second missing baryon problem. These early supernovae pre-enrich gas in addition to pre-heating it, and this helps to explain phenomena like the G-dwarf problem and the “floor” in iron abundance noted in section 2.4. For more discussion of this possibility, see Benson and Madau (2003). Due to the top-heavy IMF, they would have short lifetimes, so few of these early unbound stars would survive today, and the strict limits on intracluster light (e.g Krick and Bernstein 2007) would not be violated. A top-heavy IMF would also increase the efficiency per unit mass of metal production, which has further implications for the metallicity in galaxy clusters, as discussed in Bregman, Anderson, and Dai (2010). There are also observable constraints on this process. For example, if supernova heating is to prevent the baryons from accreting onto galaxies, the supernovae must occur before the galaxies form, i.e. $z \gtrsim 6$. There is also an important period at $z \sim 8$ when Compton cooling can efficiently cool the baryons (Cen 2003). Depending on the details of the supernovae

and the cosmology, the bulk of the early star formation and baryon expulsion might occur either before or after this period. More detailed observation and theoretical work can further constrain these possibilities.

The remarkably tight correlation for galaxies between dark matter halo mass and baryon fraction seems to offer a powerful clue for understanding galaxy formation. The correlation persists over variations in star/gas mass ratio. The mechanism and details of this connection remain unclear, however, and still need to be understood. Further observations could also help to constrain more tightly the mass and extent of hot haloes around galaxies. Direct imaging around large galaxies in soft X-rays could strengthen the surface brightness constraints. Measurements of extragalactic pulsars or additional OVII lines could also provide useful constraints on the electron column density around galaxies. Detailed examination of simulations could also be helpful, predicting, for example, the fraction of dark matter in the Universe that falls into galactic haloes for comparison to the behavior of the baryons. Finally, more modeling of galaxy formation could greatly assist in producing testable predictions that would differentiate between various solutions to this second missing baryon problem.

2.10 Acknowledgements

We would like to thank the anonymous referee for a thoughtful and helpful report that greatly improved this manuscript. We wish to thank M. Ruzkowski, S. McGaugh, O. Gnedin, E. Bell, and A. Muratov for helpful comments and discussions that related to this work. This research has made use of the Sloan Digital Sky Survey. Funding for the SDSS and SDSS-II has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, the U.S. Department of Energy, the National Aeronautics and Space Administration, the Japanese Monbukagakusho, the Max Planck Society, and the Higher Education Funding Council for England. The SDSS Web Site is <http://www.sdss.org/>. This research has made use of NASA's Astrophysics Data System. This research has made use of the NASA/IPAC Extragalactic Database (NED) which is operated by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and

Space Administration. We gratefully acknowledge support from NASA in support of this research. MEA also acknowledges support from the NSF in the form of a Graduate Research Fellowship.

2.11 References

- Abazajian, K. N. et al. 2009, *ApJS*, 182, 543
- Anders, E. and Grevesse, N. 1989, *Geochim. Cosmochim. Acta* 53, 197
- Anderson, M. E. et al. 2009, *ApJ* 698, 317
- Bahcall, J. N. et al. 1993, *ApJS* 87, 1
- Bell, E. F. et al. 2003, *ApJ* 585, 117L
- Bellazzini, M. 2004, *MNRAS* 347, 119
- Benson, A. J. et al. 2000, *MNRAS* 314, 557
- Benson, A. J. and Madau, P. 2003, *MNRAS* 344, 835
- Binney, J. and Tremaine, S. 2008, *Galactic Dynamics, 2nd Ed.* Princeton, Princeton U.P.
- Blitz, L. and Robishaw, T. 2000, *ApJ* 541, 675
- Bouché, N. et al. 2009, preprint (astro-ph/0912.1858)
- Bower, R. G. et al. 2006, *MNRAS* 370, 645
- Bregman, J. N. 2007, *ARA&A* 45, 221
- Bregman, J. N. and Lloyd-Davies, E. J. 2007, *ApJ* 669, 990
- Bregman, J. N. et al. 2009, *ApJ* 699, 1765
- Bregman, J. N., Anderson, M. E., and Dai, X. 2010, submitted to *ApJ*
- Brown, B. A. and Bregman, J. N. 2001, *ApJ* 547, 154
- Buote, D. et al. 2009, *ApJ* 695, 1351
- Cen, R. and Ostriker, J. P. 1999, *ApJ* 514, 1
- Cen, R. 2003, *ApJ* 591, 12
- Cen, R. and Ostriker, J. P. 2006, *ApJ* 650, 560
- Chabrier, G. 2003, *PASP* 115, 763
- Croton, D. J. et al. 2006, *MNRAS* 365, 11
- Dai, X. et al. 2010, submitted to *ApJ*

Davé, R. et al. 2001, ApJ 552, 473
Davé, R. 2009, preprint (astro-ph/0901.3149)
Dekel, A. and Birnboim, Y. 2006, MNRAS 368, 2
Dunkley, J. et al. 2009, ApJS 180, 306
Dutton, A. A. and van den Bosch, F. C. 2009, MNRAS 396, 141
Fang, T. et al. 2006, ApJ 644, 174
Fox, A. J. et al. 2005, ApJ 630, 332
Fukugita, M., Hogan, C. J., and Peebles, P. J. E. 1998, ApJ 503, 518
Fukugita, M. and Peebles, P. J. E. 2004, ApJ 616, 643
Fukugita, M. and Peebles, P. J. E. 2006, ApJ 639, 590
Gaensler, B. M. et al. 2008, PASA 25, 184
Geach, J. E. et al. 2005, MNRAS 363, 1398
Gebhardt, K. et al. 2001, AJ 122, 2469
Howk, J. C. et al. 2002, ApJ 569, 214
Kaastra, J. S. et al. 2006, ApJ 652, 189
Kaufmann, T. et al. 2009, MNRAS 396, 191
Kereš, D. et al. 2005, MNRAS 363, 2
Kim, J. H., Wise, J. H., and Abel, T. 2009, ApJ 694, 123L
Klypin, A., Zhao, H. S., and Somerville, R. S. 2002, ApJ 573, 597
Krick, J. E. and Bernstein, R. A. 2007, AJ 134, 466
Kuntz, K. D. and Snowden, S. L. 2000, ApJ 543, 195
Law, D. R. et al. 2007, ApJ 656, 1
Li, Y. S. and White, S. D. M. 2008, MNRAS 384, 1459
Li, Z. et al. 2006, MNRAS 371, 147L
Li, Z., Wang, Q. D., and Hameed, S. 2007, MNRAS 376, 960L
Loeb, A. et al. 2005, ApJ 633, 894L
Manchester, R. N. et al. 2006, ApJ 649, 235
Mannucci, F. et al. 2005, A&A 433, 807
Marconi, A. and Hunt, L. K. 2003, ApJ 589, 21L
Martin, C. 2005, ApJ 621, 227

Mateus, A. et al. 2006, MNRAS 370, 721
Mathur, S., Weinberg, D. H, and Chen, X. 2003, ApJ 582, 82
Maughan, B. J. et al. 2008, ApJS 174, 117
McGaugh, S. S. et al. 2000, ApJ 533, L99
McGaugh, S. S. 2005, ApJ 632, 859
Mo, H. J. and Mao, S. 2002, MNRAS 333, 768
Murray, N., Quataert, E., and Thompson, T. A. 2005, ApJ 618, 569
Navarro, J. F., Frenk, C. S., and White, S. D. M. 1997, ApJ 490, 493
Navarro, J. F. 1998, preprint (astro-ph/9807084)
Navarro, J. F. and Steinmetz, M. 2000, ApJ 538, 477
Nicastro, F. et al. 2002, ApJ 573, 157
Nicastro, F. et al. 2005a, Nature 433, 495
Nicastro, F. et al. 2005b, ApJ 629, 700
Oppenheimer, B. D. and Davé, R. 2008, MNRAS 387, 577
Penton, S. V., Stocke, J. T., and Shull, J. M. 2004, ApJS 152, 29
Pettini, M. et al. 2001, ApJ 554, 981
Rasmussen, A., Kahn, S. M., and Paerels, F. 2003, ASSL Conf. Proc. 281, 109
Rasmussen, J. et al. 2009, ApJ 697, 79
Reid, M. et al. 2009, ApJ 700, 137
Sakamoto, T., Chiba, M, and Beers, T.C. 2003, A&A 397, 899
Sembach, K. R. et al. 2001, ApJ 561, 573
Sembach, K. R. et al. 2004, ApJS 155, 351
Shattow, G. and Loeb, A. 2009, MNRAS 392, L21
Smith, R. K. et al. 2001, ApJ 556, L91
Snowden, S. et al. 1998, ApJ 493, 715
Stanimirović, S. et al. 2002, ApJ 576, 773
Stark, D. V., McGaugh, S. S., and Swaters, R. A. 2009, AJ 138, 392
Strateva, I. et al. 2001, AJ 122, 1861
Strickland, D. K. et al. 2004, ApJS 151, 193
Sutherland, R. R. and Dopita, M. A. 1993, ApJS 88, 253

Tang, S. et al. 2009, MNRAS 392, 77
Tortora, C. et al. 2009, MNRAS 396, 1132
Tüllmann, R. et al. 2006, A&A 457, 779
van der Kruit, P. C. 1984, A&A 140, 470
Wang, Q. D. et al. 2005, ApJ 635, 386
Wu, K. K. S, Fabian, A. C., and Nulsen, P. E. J. 2001, MNRAS 324, 95
Xue, X. X. et al. 2008, ApJ 684, 1143




Figure 2.1 Integrated X-Ray emission measures for a halo of hot gas obeying an NFW profile in hydrostatic equilibrium. The three colors denote the total mass of the host galaxy and its dark matter halo; each line represents the EM for an NFW profile of concentration parameter $C=12$. The three thick horizontal bars represent the approximate radial extent of the thick disk of a galaxy of the given mass; within these radii, the bright emission should be obscured and unobservable. The three thin horizontal lines are estimates of the minimum emission measures detectable with *Chandra* ACIS-S for each halo. These haloes are all observable out to many times the radius of the thick disk, and should be detected unless the halo mass is reduced to 11-24% of the mass of missing baryons. The slope of the surface brightness between the thick disk and the detection limit follows a β -model with $\beta \approx 0.87$ in each case.




Figure 2.2 Comparison of the scatter of galaxies around the baryonic Tully-Fisher relation (BTFR), based on the sample from McGaugh (2005). We use the galaxy type as a proxy for the size of the supermassive black hole (SMBH). For the 60 galaxies with good galaxy types, there is no trend of the residual around the BTFR with galaxy type. As discussed in section 2.6, we conclude that, at least for galaxies of type Sb and later, AGN activity associated with SMBHs does not play a dominant role in the removal of baryons from those galaxies.




Figure 2.3 Comparison of galaxy residuals about the baryonic Tully-Fisher relation as a function of their stellar mass fraction for the gas-dominated sample of Stark et al. (2009). For several of these galaxies, the stars comprise less than 25% of the baryons, yet they have the same baryon fraction as the more star-rich galaxies. As discussed in section 2.6, we conclude that supernovae from the stars in these galaxies are not responsible for ejecting most of the missing baryons from the galaxies.



Figure 2.4 SDSS images of NGC 4138 (left) and NGC 4183 (right). NGC 4138 is at a distance of 15.4 ± 2.3 Mpc, NGC 4183 is at a distance of 17.3 ± 0.9 Mpc. Both images have a scale of $0.3961''/\text{pixel}$. For more information on and discussion of these galaxies, please see section 2.6.

Table 2.1 Constraints on the Milky Way Hot Halo

Method	NFW profile			Flattened profile			Reference
	Hot halo mass ($10^9 M_\odot$)	$n_e(50 \text{ kpc})$ (10^{-5} cm^{-3})	frac	Hot halo mass ($10^9 M_\odot$)	$n_e(50 \text{ kpc})$ (10^{-5} cm^{-3})	frac	
LMC Pulsar DM	< 12 – 15	< 7.7 – 10	< 0.04 – 0.05	< 170	< 27	< 0.58	§3.1
Mag. Stream HI	< 10 – 11	< 7	< 0.03 – 0.04	< 53	< 8	< 0.18	Stanimirović et al. (2002)
HVCs	< 5 – 9	< 3 – 6	< 0.02 – 0.03	< 19 – 39	< 3 – 6	< 0.06 – 0.13	Fox et al. (2005)
Galactic XRB	< 5.9 – 7.8	< 4 – 5	< 0.02 – 0.03	< 110	< 17	< 0.38	Kuntz & Snowden (2000)

Parameters determined from fitting profiles to the constraints noted in section 2.4.2. *frac* is the ratio of the mass of the hot halo to the mass of the missing baryons from the Galaxy ($3 \times 10^{11} M_\odot$).

Table 2.2 OVII Sightline Analysis

Sightline	photo-z	Ly α EW (m \AA)	b (kpc)	L $_g$ (L $_{\odot}$)	u - r	N $_{e,th}$ (10 19 cm $^{-2}$)	N $_{e,NFW}$ (10 19 cm $^{-2}$)	(S/N) $_{NFW}$	N $_{e,pl}$ (10 19 cm $^{-2}$)	(S/N) $_{pl}$
3C273	0.11 \pm 0.04	138 ^a	41	9 \times 10 ⁸	2.00	9.1	13.8 $^{+20.4}_{-10.2}$	4.6	13.8 $^{+22.3}_{-12.2}$	4.5
Ton 1388	0.09 \pm 0.02	121 ^b	28	9 \times 10 ⁷	2.77	12.0	0.2 $^{+1.3}_{-0.2}$	0.1	0.7 $^{+2.8}_{-0.7}$	0.2
Ton 1388	0.17 \pm 0.01	233 ^b	39	4 \times 10 ⁹	3.21	12.0	20.5 $^{+12.2}_{-8.4}$	5.1	16.9 $^{+23.2}_{-13.7}$	4.2
Ton 1388	0.06 \pm 0.03	79 ^b	43	5 \times 10 ⁸	1.76	12.0	9.2 $^{+17.9}_{-7.8}$	2.3	11.1 $^{+19.6}_{-10.7}$	2.8
Ton 1388	0.13 \pm 0.01	132 ^b	135	1.2 \times 10 ¹⁰	3.07	12.0	6.2 $^{+3.5}_{-2.5}$	1.5	11.8 $^{+16.1}_{-6.1}$	2.9
Q1821+643	0.21 \pm 0.05	560 ^c	75	1 \times 10 ⁹	3.40	10.0	0.9 $^{+3.0}_{-0.9}$	0.3	2.7 $^{+7.1}_{-2.7}$	0.8
Q1821+643	0.18 \pm 0.10	350 ^c	96	4 \times 10 ⁹	1.81	10.0	6.4 $^{+12.3}_{-5.6}$	1.9	11.2 $^{+19.9}_{-7.3}$	3.4

The impact parameter b is the distance between the galaxy and the relevant AGN sightline. Each galaxy was matched to Lyman α absorption detected at or very near its photometric redshift; the equivalent width is noted. The neutral Hydrogen column densities implied by the Ly- α lines are $\sim 10^{13} - 10^{14}$ cm $^{-2}$ - significantly lower than the electron column densities, as would be expected from highly ionized plasma. The Lyman α references are: (a) Sembach et al. 2001, (b) Sembach et al. 2004, (c) Bahcall et al. 1993. The $u - r$ color is used to distinguish between early-type ($u - r > 2.2$) and late-type ($u - r < 2.2$) galaxies when computing $N_{e,NFW}$. EW_{th} is the equivalent width of a 3σ OVII detection for the *Chandra* observation of this AGN in Fang et al. (2006), and $N_{e,th}$ is the electron column corresponding to this EW. $N_{e,NFW}$ is the electron column we expect from the hot halo around this galaxy if it contains the missing baryons (the uncertainty reflects 1σ uncertainties propagated from the redshift, an assumed 5 kpc uncertainty in the impact parameter, the g-band photometry, scaling relation, and different choices for C in the NFW profile). $(S/N)_{NFW}$ is the strength of the detection we would expect for this halo if it extended to the virial radius. $N_{e,pl}$ and SN_{pl} are the column density and signal-to-noise of a flatter, power-law profile that will be discussed in section 2.7.

Table 2.3 Halo Emissivity Analysis

Mass (M_\odot)	kT (keV)	F_1 (erg s $^{-1}$ cm $^{-2}$)	A (cm 2)	$(n^2R)_{\min}$ (cm $^{-6}$ kpc)
2×10^{12}	300	4.06×10^{-10}	217	2.1×10^{-5}
8×10^{11}	160	2.31×10^{-10}	137	5.9×10^{-5}
3×10^{11}	85	5.12×10^{-11}	62	5.9×10^{-4}

Numbers used for derivation of emission measure $(n^2R)_{\min}$ of halo gas in hydrostatic equilibrium that would be detected by *Chandra* ACIS-S at 30% of background. Mass is the total mass of the galaxy and dark matter halo, and kT is the temperature of the halo gas. F_1 is the estimated flux from an APEC model (see text). And A is the effective area of the detector for each observation.

Table 2.4 Comparison of Two Nearby Disk Galaxies

Name	i (degrees)	v_f (km s $^{-1}$)	M_{200} (log M_\odot)	M_b (log M_\odot)	f_b (M_b/M_{200})	M_g (log M_\odot)	M_s (log M_\odot)	Type
NGC 4138	53	148 ± 4	12.02	10.34	0.021	9.30	10.30	S0
NGC 4183	82	111 ± 2	11.64	9.85	0.016	9.69	9.33	SAcd

Data taken from Stark, McGaugh, and Swaters (2009) and from the NASA Extragalactic Database, as described above.

M_{200} is computed from v_f as described above as well, assuming $h = 0.72$.




Figure 2.5 As in Figure 2.1, the integrated X-Ray emission measures for a halo of hot gas, but in this plot the gas obeys a power-law profile: $n(r) \propto r^{-0.9}$. These hot haloes have a lower emission measure, but a higher emissivity per gram due to their increased temperature. The largest should be detected unless the halo mass is reduced to 59% of the mass of missing baryons, the smaller haloes have weaker constraints. The slope of the surface brightness between the thick disk and the detection limit follows a β -model with $\beta \approx 0.36$ in each case.

CHAPTER 3

Detection of a Hot Gaseous Halo Around the Giant Spiral Galaxy NGC 1961

Note: This chapter is a reproduction of a paper of the same title, published in the *Astrophysical Journal*. The reference is Anderson, M. E. and Bregman, J. N. 2011, *ApJ*, 1, 22-32. The paper is copyright 2011, the American Astronomical Society, reproduced here under the non-exclusive right of republication granted by the AAS to the author(s) of the paper.

3.1 Abstract

Hot gaseous halos are predicted around all large galaxies and are critically important for our understanding of galaxy formation, but they have never been detected at distances beyond a few kpc around a spiral galaxy. We used the Chandra ACIS-I instrument to search for diffuse X-ray emission around an ideal candidate galaxy: the isolated giant spiral NGC 1961. We observed four quadrants around the galaxy for 30 ks each, carefully subtracting background and point source emission, and found diffuse emission that appears to extend to 40-50 kpc. We fit β -models to the emission, and estimate a hot halo mass within 50 kpc of $5 \times 10^9 M_{\odot}$. When this profile is extrapolated to 500 kpc (the approximate virial radius), the implied hot halo mass is $1 - 3 \times 10^{11} M_{\odot}$. These mass estimates assume a gas metallicity of $Z = 0.5Z_{\odot}$. This galaxy's hot halo is a large reservoir of gas, but falls significantly below observational upper limits set by previous searches, and suggests that NGC 1961 is missing 75% of its baryons relative to the cosmic mean, which would tentatively place it below an extrapolation of the baryon Tully-Fisher relationship of less massive galaxies. The

cooling rate of the gas is no more than $0.4 M_{\odot}/\text{year}$, more than an order of magnitude below the gas consumption rate through star formation. We discuss the implications of this halo for galaxy formation models.

3.2 Introduction

Hot gaseous halos around galaxies have been an important prediction of galaxy formation models since White and Rees (1978). Theory predicts these hot halos form as matter accretes onto the dark matter halo and the baryons shock to the virial temperature (White and Frenk 1991, also see the review by Benson 2010). Depending on the details of the assumed pre-heating, heating from galactic feedback, and cooling rates, these hot halos are often predicted to contain as much or more baryonic mass as the galaxies within the halos (Sommer-Larsen 2006, Fukugita and Peebles 2006), making them cosmologically important as reservoirs of the “missing baryons” from galaxies (although also see Anderson and Bregman 2010). The hot halo is also thought to produce the galactic color-magnitude bimodality (Dekel and Birnboim 2006) and to help explain galactic “down-sizing” in the star formation history (Bower et al. 2006, De Lucia et al. 2006).

Hot halos have been extensively observed in soft X-rays (roughly 0.5 - 2 keV) around early-type galaxies (Forman, Jones, and Tucker 1985, O’Sullivan, Forbes, and Ponman 2001, Mulchaey and Jeltama 2010). The halos are typically luminous ($L_{X,0.5-2 \text{ keV}} \sim 10^{39} - 10^{41}$ for non-BCG ellipticals), mass-dependent (for most definitions of L_X and L_K , $L_X \propto L_K^2$) and are often visible out to many tens of kpc. But these halos are difficult to connect to the formation of the galaxies because coronal gas can also be produced in the mergers and associated star formation that occurred when the galaxy became elliptical (Read and Ponman 1998), and because it is difficult to disentangle halo gas with the intergroup medium (IGM) in which most large ellipticals reside (Dressler 1980).

In contrast, hot halos around quiescent disk galaxies should be much more direct tracers of the galaxy formation process. While the morphology-density relation makes it difficult to disentangle elliptical galaxies from their dense environments, it also

ensures a large supply of isolated spiral galaxies in low-density environments. Late-type disks are destroyed by strong mergers (e.g. Robertson et al. 2006), and it is easy to identify and exclude starbursting galaxies, so it should be straightforward to search for hot halos around quiescent isolated spirals, and to connect these halos to models of galaxy formation.

Unfortunately the search for extended soft X-ray emission around isolated spirals has so far been unsuccessful. There are several detections of emission a few kpc above the disk (Strickland et al. 2004a, Li et al. 2006, Tüllmann et al. 2006, Rasmussen et al. 2009, Owen and Warwick 2009, Yamasaki et al. 2009), but these observations are linked to the star formation in the galaxy and probably represent galactic fountains. In terms of more extended emission, Li, Wang, and Hameed (2007) observe gas around the Sombrero galaxy out to 20 kpc, but this galaxy is significantly bulge-dominated, and the extended emission has been linked to a galactic bulge-driven wind. Finally, Pedersen et al. (2006) claimed to detect extended hot halo emission around NGC 5746, but this emission disappeared after subsequent re-analysis with newer calibration files (Rasmussen et al. 2009).

A recent paper (Crain et al. 2010b) attributes these detections of extended emission to galactic coronae, instead of the standard explanation of the emission as a fountain or a wind originating from within the galaxy. This interpretation is in disagreement with the standard understanding of galactic fountains in spiral galaxies, but regardless of interpretation it still is true that no hot halo has been detected around a disk galaxy at a radius of more than a few kpc.

In this paper, we present an analysis of *Chandra* ACIS-I observations of the environs of the extremely massive spiral galaxy NGC 1961, in which we detect X-ray emission out to at least 40 kpc and attribute the emission to a hot halo. The outline of the paper is as follows. In §2, we discuss the properties of NGC 1961 and the details of our observation. In §3, we discuss the reduction of the data and explain various approaches to flat-fielding we adopted in our analysis. In §4 we present the spatial analysis, and in §5 we present the spectral results. Finally, in §6 we discuss the derived properties of the hot halo and in §7 we place them in the context of galaxy

formation.

3.3 Observation

This galaxy is one of the most massive spiral galaxies known (Rubin et al. 1979), with a maximum (deprojected) HI circular velocity of 402 km s^{-1} at 34 kpc (Haan et al. 2008). These authors fit for an inclination angle of 42.6° , close to the LEDA value of 47° . A recent paper (Combes et al. 2009) makes a plausibility argument for $i \approx 65^\circ$, which would reduce the circular velocity by 25%. We use the more standard inclination in this paper, but if the higher inclination is correct, the total mass of the galaxy would be $0.75^{-3} = 2.4$ times smaller. While this is a significant difference, it does not change our final conclusion.

If the relation $M_{\text{dyn}} \propto V_{\text{max}}^3$ holds at these velocities, we can compare NGC 1961 to the Milky Way ($v_{\text{circ}} = 220 \text{ km s}^{-1}$) and infer that NGC 1961 has six times the dynamical mass of the Milky Way. Similarly, its 2MASS K-band magnitude is -26.0, which for an assumed mass-to-light ratio of 0.6 (Bell and de Jong 2001) corresponds to a stellar mass of $3 \times 10^{11} M_\odot$ (which is also six times the stellar mass of the Milky Way). Extrapolating from the $L_X - L_K$ relation for elliptical galaxies, we therefore expect NGC 1961 to have an unusually bright X-ray halo ($L_{0.5-2 \text{ keV}} \approx 1 \times 10^{41} \text{ erg s}^{-1}$ for the diffuse emission), making this galaxy an ideal target for identifying extended X-ray emission.

The virial radius of the Milky Way is $\sim 250 \text{ kpc}$ (Shattow and Loeb 2009, Klypin, Zhao, and Somerville 2002), so by extension the virial radius of NGC 1961 would be around 450 kpc^1 . Within this radius, NGC 1961 has several, much smaller, companions (Gottesman et al. 2002, Haan et al. 2008), including three dwarfs ($M_{\text{HI}} < 10^9 M_\odot$) at 120, 140, and 160 kpc, and several slightly larger galaxies at 200-500 kpc distances. It is therefore the dominant galaxy in a small group, but no IGM emission is observed. We adopt for this galaxy a distance of 56 Mpc (NASA Extragalactic Database average), which matches independent measurements of distances

¹We use the form for the virial radius defined in terms of the critical density ($\rho_{\text{vir}} \approx 200\rho_c$).

to other galaxies in the group (Gottesman et al. 2002), and is probably uncertain to 10%. At this distance, 1 arcminute corresponds to 16 kpc.

Our observing strategy was to use a 2×2 mosaic with the ACIS-I array on the Chandra X-ray Observatory, which allowed us to sample the extended emission out to about 260 kpc ($17'$) - roughly $2/3$ of the virial radius (see Figure 3.1). The observations (obs ids 10528-10531) were approved for 35 ks each, and ranged from 31.75-33.25 ks of good time. We also observed two background fields (obs ids 10532,10533) for 10.14 and 10.02 ks. All observations were taken in VFAINT mode with ACIS-I.

The data were processed using CIAO, version 4.1.2, and the latest calibration files. Data taken during flares were excised, but the observations were remarkably clean, with only about 0.4 ks of bad time to remove from the total integration time for each observation. A 0.6-6 keV image was produced for point source detection (using the WAVDETECT algorithm in CIAO). For the rest of the analysis, we use a 0.6-2.0 keV image for observations of the hot halo, and a 2.0-6.0 keV image for constraints on contamination from emission from X-ray binaries, as discussed in section 3.4.4.

3.4 Flat-fielding

This diffuse X-ray emission is very faint, even from such a large galaxy, which is why it has posed such an observational challenge. Therefore the most difficult and important part of the data reduction is the flat-fielding procedure. We tried several different methods of flat-fielding and background subtraction, described below. We eventually developed our own method of flat-fielding the observations (section 3.4.3), which should be generalizable to future observations of faint diffuse emission around nearby galaxies (section 3.7.1).

3.4.1 Using Background Frames

Our original intent was to use images taken a degree off-axis from the galaxy as background images which could be subtracted pixel-by-pixel from the halo images, accomplishing both flat-fielding and background subtraction at once. We took two background pointings for this purpose. The former (obs id 10532) was pointed at

blank sky 17.3 degrees away from the galaxy for 10.14 ks, and was taken 29 days before the science pointings. The latter (obs id 10533) was pointed at blank sky 0.9 degrees away from the galaxy for 10.02 ks, 7.5 days before the science pointings began.

Unfortunately these background images proved unsuitable for background subtraction because the count rate across the image was different than the count rates in the source images. After processing and point source subtraction, the mean 0.6-2 keV count rate in the background images is $0.31 \text{ count s}^{-1}$ compared to $0.36 \text{ count s}^{-1}$ in the source images. This 17% discrepancy is comparable to the total signal from hot halo emission integrated out to 500 kpc (see section 3.7), and probably stems from variations either in the unresolved X-ray background of the images or in variations in the solar X-ray flux between the background observations and the source observations. In either case, the discrepancy was large enough that we were unable to use the background images for background subtraction. Similarly, the standard background images are even more spatially and temporally separate from our source images and could not be used either.

3.4.2 In-field Subtraction, Modeling the Background

We therefore decided to use in-field background subtraction instead. In-field subtraction is not ideal for our project, since we are trying to measure diffuse emission that, in principle, could fill the entire field of view. However, the diffuse emission is only detectable out to a few arcminutes from the galaxy, and that even over most of that range the background is larger than the signal. Thus it is possible to compute the background in-field, and use this for background subtraction.

We first attempted to model the background. We assumed the 0.6 - 2.0 keV background consisted of two components: a vignettted component from the diffuse X-ray background (which we assumed would be proportional to the effective throughput at each pixel, computed using the exposure map), and an unvignettted component from particles in the Solar wind. In each observation, we sampled the background at various points across the detector (away from the direction of the galaxy), and fit the

background to a linear combination of these two components.

This procedure worked adequately for two of the four observations (10528 and 10530), but the other two had large-scale variations in the background across the image or other un-modeled effects and therefore did not yield reasonable fits to the two background components.

3.4.3 In-field Subtraction, Conjugate Technique

We found more success with in-field conjugate subtraction. The idea is to assume the vignettted background is azimuthally symmetric over large scales around the aimpoint of the image (located on the ACIS-I3 chip at approximately $(x,y) = (974,969)$ in detector coordinates²). We selected “conjugate” background regions in the image of the same shape and size as the source, but at the opposite position angle from the aimpoint of the image. We then subtracted each conjugate region from the corresponding source region (see Figure 3.2).

As shown in Figure 3.2, we faced different geometrical combinations of source, conjugate, and aimpoint for each observation. For observation 10528, the galaxy is just off the edge of the boundary between the I0 and I1 chips, the aimpoint is on the opposite side of the center of the ACIS-I array, and so the conjugate point is 170 arcseconds beyond the edge of the I3/I4 chips. We therefore had to choose annuli at the same distance from the aimpoint, but at angles of 120° and 240° instead of 180° to find conjugate regions on the detector for the source photons from the inner 170 arcseconds (45 kpc) around the galaxy in this observation. For observation 10529, the galaxy is just off the edge of the I0/I2 chips, the aimpoint is again on the opposite side of the center of the array, and the conjugate point is 130 arcseconds beyond the edge of the I3/I1 chips. We therefore had to choose annuli at the same distance from the aimpoint, but at angles of 120° and 240° instead of 180° to find conjugate regions on the detector for the source photons from the inner 130 arcseconds (35 kpc) around the galaxy in this observation. For observation 10530, the galaxy is 200 arcsec off the edge of the I1/I3 chips, but the aimpoint is now on the same side of the center of the

²Chandra Proposers’ Observatory Guide, version 13.0, <http://cxc.harvard.edu/proposer/POG/html/index.html>

array, so the conjugate point is close to the edge of the I0/I2 chips. But since the galaxy is off the detector, we do not have a measurement of the source for the inner 200 arcseconds (50 kpc) around the galaxy in this observation. For observation 10531, the galaxy is on the I3 chip, near the I2 chip, 270 arcseconds from the aimpoint. The conjugate point is also on the detector, on the boundary between the I0/I1 chips. So for this observation, we have measurements of the source and background emission out to 270 arcseconds (75 kpc); beyond this radius the background and source annuli begin to overlap. We excluded the data beyond this radius for observation 10531 for the rest of the analysis, but did also repeat the analysis with these points included, and found that they have no effect on the results since the halo emission has disappeared by 50 or 60 kpc.

This yielded similar backgrounds to the results of the modeling for 10528 and 10530, but the results were much better for 10529 and 10531, so we adopted this approach for the rest of the analysis. To help verify the reliability of this technique, we also tested the conjugate technique 90° on either side of the source, and obtained zero signal.

3.4.4 Point Sources

It is critically important to be sure we are measuring the hot diffuse emission and not a collection of X-ray binaries in and around the galaxy, whose surface density also falls off with radius like the halo gas. The first step to ensuring a clean measurement is the automated point source removal using WAVDETECT, described above, which removed six point sources within the inner 50 arcseconds, with the faintest point source having a luminosity of $L_{0.6-2 \text{ keV}} \sim 3 \times 10^{38} \text{ erg s}^{-1}$ if at the assumed distance of 56 Mpc. One of these point sources falls on the galactic nucleus, which hosts a low-luminosity AGN (Roberts and Warwick 2000) for which we measured a luminosity of $1 \times 10^{40} \text{ erg s}^{-1}$.

To estimate the contribution of point source emission to the surface brightness profile, we extracted and reduced an image of the 2-6 keV emission, using the identical procedure as we used for the 0.6-2 keV images. We expect no contribution from

the < 1 keV gas in this higher-energy band, so all the emission should come from point sources in the galaxy or the background. Using the in-field conjugate subtraction technique, we subtracted the background emission and derived radial surface brightness profiles for the 2-6 keV galactic emission. We attributed all this emission to unresolved point sources.

Irwin, Athey, and Bregman (2003) found that the integrated emission from low-mass X-ray binaries (which dominate the point source emission over most of the region in our analysis) has a universal spectrum that can be fit with a power-law distribution with a slope $\Gamma = 1.56$. Using this slope and accounting for absorption, we find that each unit of 2-6 keV emission corresponds to 1.7 units of 0.6-2 keV emission. We scaled the 2-6 keV galactic emission by a factor of 1.7 and subtracted this from the 0.6-2 keV emission to remove the point source contribution. We stopped the unresolved point source removal after reaching the average background 2-6 keV surface brightness for each observation; this occurred at 32 kpc for observation 10528, 30 kpc for observation 10529, and 25 kpc for observation 10531 (observation 10530 has the galaxy 50 kpc off the edge of the detector, so there is no contribution from galactic point sources. The total amount of emission due to unresolved point sources was $\approx 1.5 \times 10^{40}$ erg s $^{-1}$, about 10-20% of the total emission.

We attempted to verify this result by assuming the six point sources detected with WAVDETECT represent a complete sample down to 3×10^{38} erg s $^{-1}$. Using the *Chandra* point source number counts of Kim et al. (2007), we expect 0.6 background point sources above our flux limit in the (50 arcsecond radius) aperture, so one of these six sources is likely actually an unrelated background object (and one is the central low-luminosity AGN). We applied the luminosity function of Grimm, Gilfanov, and Sunyaev (2002), which was calibrated using Galactic LMXBs and HMXBs, to the five non-nuclear point sources. We adopted a power-law function for $N(> L)$ with a slope of $\alpha = -0.3$. The ratio of the total luminosity (with a lower-luminosity cut-off at 10^{36} erg s $^{-1}$ to the luminosity above 3×10^{38} erg s $^{-1}$ is 5.0, and the total 0.6-2 keV luminosity of our five point sources is 7.0×10^{39} erg s $^{-1}$, so we expect a point source luminosity of 3.5×10^{40} erg s $^{-1}$. This figure is about twice as high as the point

source luminosity we infer from the 2-6 keV flux, although the likely presence of a background point source is probably causing us to overestimate the background using this method.

If we assume that all five detected non-nuclear point sources within 50 arcseconds (13 kpc) are associated with the galaxy, then our method of scaling from the high-energy emission corresponds to a flat power-law slope of $\alpha = -0.15$ for $N(> L)$. This slope is 1.4σ away from the best-fit slope to the LMXB luminosity function in Grimm, Gilfanov, and Sunyaev (2002). It does seem more likely for NGC 1961 to have a different point source luminosity function than the Milky Way instead of a different X-ray spectrum for its point sources. Supporting this result, a flattening of the point source luminosity function below a few $\times 10^{37}$ erg s^{-1} is also observed in Centaurus A (Voss et al. 2009).

We also included a correction for X-ray emission from stars at a level fainter than 10^{36} erg s^{-1} . This emission seems to scale with total stellar mass (inferred from the K-band luminosity), at least for old stellar populations (Revnivtsev et al. 2008). While NGC 1961 is a late-type spiral and therefore not likely to be dominated by an old stellar population, the Revnivtsev et al. scaling is the best available at present for accounting for this emission. We took the K-band radial surface brightness profile for NGC 1961 from 2MASS (Jarrett et al. 2003), re-binned it to match our X-ray annuli, and assumed a K-band mass-to-light ratio of 0.6 to convert into stellar mass (Bell and de Jong 2001). We then applied the Revnivtsev et al. conversion between stellar mass and 0.5-2.0 keV X-ray luminosity to estimate the stellar X-ray emission. Finally, we multiplied the predicted X-ray emission by 73% to account for absorption by the Galactic Hydrogen column in the direction of NGC 1961. We find that the stellar emission is never a significant fraction of the total X-ray emission, even at very small radii; the total 0.5-2.0 keV emission from stars is 2.6×10^{39} erg s^{-1} , or about 5% of the total luminosity in the central region. The K-band half-light radius is 35 arcsec (9.6 kpc), and the K-band surface brightness reaches the 2MASS 1σ background at about 78 arcsec (21 kpc) – a much more concentrated profile than the X-ray emission.

3.5 Radial Surface Brightness Profile

Our primary interest is in deriving a radial surface brightness profile for the hot gas around NGC 1961. The emission is very faint, so we have to make a few critical assumptions in order to make any progress in parameterizing the surface brightness profile. Most importantly, we expect the halo to be roughly spherical, so we compute the surface brightness profile in circular annuli around the galaxy. Starting with just this assumption, in Figure 3.3 we present background-subtracted radial surface brightness profiles for our four images of the halo around NGC 1961. We chose annuli such that each source annulus contains at least 20 photons, although we also tried annuli of constant radius and found no significant differences in the shape of the profile.

Inspecting Figure 3.3, emission seems visible out to about 160 arcsec, corresponding to about 40 kpc. This is the key result of our paper, and it does not depend on any further assumptions or statistical techniques (and we find this excess in multiple observations/quadrants). Some of the interior emission is due to X-ray binaries and a low-luminosity AGN, but we attribute the bulk of the emission at 40 kpc to diffuse gas around the galaxy. We discuss in section 3.8.3 the possibility that this emission results from some internal galactic process (i.e. a supernova-driven wind or galactic fountain), but we conclude this possibility is unlikely.

In the rest of this section, we will attempt to quantify the level of emission above the background out to 40 kpc, by introducing parametric fits to the surface brightness profile and then by smoothing the background.

3.5.1 Parametric Fitting with the β -model

Our assumed parametric form for the surface brightness profile is the class of models known as the β -models. These models parameterize the surface brightness as

$$S(r) = S_0 \left[1 + \left(\frac{r}{r_0} \right)^2 \right]^{0.5-3\beta}$$

which, if the gas is isothermal and has a constant metallicity, corresponds to a

density distribution of the form

$$n(r) = n_0 \left[1 + \left(\frac{r}{r_0} \right)^2 \right]^{-1.5\beta}$$

These models provide good fits to hot gas around elliptical galaxies (Forman, Jones, and Tucker 1985) as well as the hot gas in galaxy groups and clusters (Sarazin 1986). We quantified fits to the β -models using χ^2 -minimization and we binned the data to have at least 20 photons per radial bin. Using the radial surface brightness profiles, we attempt to exclude specific choices of β -model, and the best-fit models are the profiles which can be excluded at the lowest confidence. We want to be able to exclude at less than 95% confidence for the model to be considered statistically acceptable.

We fit the surface brightness profiles to the β -profiles and solved for S_0 , r_0 , and β using χ^2 minimization. We include annuli extending out to 370 arcsec in our fit, since this radius appears to enclose all the excess emission visible in the surface brightness profile (see Figure 3.3). However, varying this radius does not affect the result much. We also require a core radius of at least 1 kpc. It would be better if we did not have to constrain the core radius at all, but we have no observations within the core radius so an observational constraint is difficult. A 1 kpc core radius is very small for a hot gaseous halo around a galaxy of this size, so our constraint is at least still somewhat conservative.

We fit all four profiles simultaneously and to find a single set of parameters that worked for all four observations. The best-fit parameters were $S_0 = 9.77 \times 10^{-8}$ count $\text{s}^{-1} \text{ cm}^{-2} \text{ arcsec}^{-2}$, $r_0 = 1.00$ kpc, $\beta = 0.47$, and the full range of acceptable fits is narrowly clustered around these values (see Figure 3.3).

While we had at least 20 source photons in each radial bin, in the inner annuli where the galaxy is brighter than the background there are fewer than 20 background photons per bin. Additionally, from the size of the radial background variations at large radii, we can infer there are small-scale spatial variations in the background (probably due to unresolved point sources). We attempted to correct for these two

effects by fitting a second-order polynomial to the surface brightness profile of the background. Quadratics were the lowest-order polynomials with acceptable fits to the background surface brightness profile (reduced $\chi^2 = 1.42, 1.70, 1.08,$ and $1.02,$ respectively, for 33, 33, 26, and 47 degrees of freedom). Observation 10529 has the least well-behaved background at small radii, and this is partially responsible for the lower values at small radii for this observation.

As noted above, our main conclusion – that there is extended coronal emission around NGC 1961 out to at least 40 kpc – does not depend on this smoothing techniques, and we can still get an acceptable fit to the data without any smoothing. Smoothing the background allows us to remove a principal source of error in our analysis, however, yielding a wider and more reliable range of acceptable fits. The data with smoothed background, as well as the acceptable fits, are shown in Figure 3.4. We take the fits to the smoothed data (unacceptable at less than 95% confidence) as the fiducial range for the rest of the paper. Note that this range encompasses the entire range of acceptable fits to the unsmoothed data.

We present the smoothed surface brightness profile in log-log space in Figure 3.5. In this figure, we have subtracted out the estimated contribution from X-ray binaries and from stars, as discussed in section 3.4.4. For ease of visualization in log-log space, for this figure we have not subtracted out the smoothed background; rather, we indicate the level of the smoothed background for comparison. Again, we clearly detect emission above the background out to 40-50 kpc, in multiple quadrants, and this emission is more extended than the emission from stars and X-ray binaries.

The parameters for the joint fit with the highest enclosed mass are ($S_0 = 3.85 \times 10^{-8}$ count s $^{-1}$ cm $^{-2}$ arcsec $^{-2}$, $\beta = 0.41,$ $r_0 = 1.00$ kpc), and the parameters for the fit with the lowest enclosed mass are ($S_0 = 1.38 \times 10^{-8}$ count s $^{-1}$ cm $^{-2}$ arcsec $^{-2}$, $\beta = 0.54,$ $r_0 = 4.04$ kpc).

3.6 Spectral Fitting

We also examined the spectrum of the source photons, to verify that the emission is consistent with strong metal emission lines atop a thermal bremsstrahlung continuum

as expected (and not, for example, unresolved point sources), and to measure the temperature of the hot emitting gas.

We examined the spectrum of observation 10531. We made a 0.5-6 keV image and removed point sources using WAVDETECT, then selected the inner 40 kpc (160 arcsec) a source aperture and a conjugate circle of the same size as a background region. This is the radius at which the signal approximately matches the background, and therefore should yield the spectrum with the optimal signal/noise ratio.

Using XSPEC version 12.6.0, we fit various models to the data. The hot halo is expected to match an APEC model, and the X-ray binaries are expected to follow a $\Gamma = 1.56$ powerlaw, and both components should experience Galactic photoabsorption, so we fit the signal with an (APEC + powerlaw) \times PHABS model. To account for the background (Galactic and extragalactic) and instrumental features, we followed Humphrey et al. (2011) and used an (absorbed) $\Gamma = 1.41$ powerlaw for the extragalactic background and two (unabsorbed) APEC models with Solar abundance and $kT = 0.07$ keV and $kT = 0.20$ keV for the Galactic background. We also included two zero-width Gaussian features at 1.77 and 2.2 keV to account for Si and Au emission lines from the instrument itself. The normalizations of all these components were allowed to float, but we fixed the hot halo metallicity at $Z = 0.5Z_{\odot}$ and the Galactic column at $8.28 \times 10^{20} \text{ cm}^{-2}$ (Dickey and Lockman 1990). The best-fit model has a temperature for the hot halo component of $kT = 0.60_{-0.09}^{+0.10}$ keV and a normalization of $1.07_{-0.27}^{+0.36} \times 10^{-4}$. The $\Gamma = 1.56$ powerlaw component (corresponding to XRB emission from NGC 1961) had a normalization of $1.71_{-1.71}^{+0.79} \times 10^{-5}$. The χ^2 was 119.5 for 122 degrees of freedom, corresponding to a reduced χ^2 of 0.98.

We tried a fit with an additional PHABS component to account for absorption within the disk of NGC 1961, but this component did not improve the fit at all. Letting the photoabsorbing column density float instead of fixing it at 8.28×10^{20} , we are not able to constrain meaningfully the value. Similarly, we also tried a fit with the hot halo metallicity left floating (instead of fixed at $0.5Z_{\odot}$), and found the same halo temperature but no constraint on the metallicity. We were able to estimate the neutral Hydrogen column in NGC 1961 using the 21-cm observations

of Haan et al. (2008). They find a clumpy distribution of column density ranging between $2 - 4 \times 10^{20}$ throughout the inner 20 kpc, with a sharp dropoff at larger radii. We chose 3×10^{20} as a rough average column for the inner 20 kpc. A fit with this additional photoabsorption component on the APEC and the powerlaw models is equally acceptable and yields essentially the same temperature and normalization. Based on the Haan et al. (2008) column densities, however, we do apply a $\sim 10\%$ upwards correction to the X-ray surface brightness in Figures 3.3 and 3.4 and in the rest of the spatial analysis for bins within 20 kpc to account for the additional photoabsorption, although this has almost no effect on the final profile.

In our aperture, we have approximately 767 ± 65 source photons out of a total of 2510 photons in the 0.5-6 keV energy range. Overall, the results of the spectral fitting are in agreement with our theoretical expectations for the source of the emission. In particular, we expect a halo temperature under 1 keV, and the APEC component to have a total emissivity about 5-10 times as large as the powerlaw component. Both these expectations are met, and additionally the APEC normalization from the spectral fit is consistent with the APEC normalization we infer from spatial fitting in the next section.

3.7 Halo Mass

We estimate the halo mass from the results of the spatial fit (S_0 , r_0 , and β) after removing individual and unresolved point sources. We integrate the β model out to 50 kpc, after which radius the signal is down to below 20% of the background. The enclosed source photon count rate is between 0.011-0.012 count s^{-1} . We then use the PIMMS utility (assuming an APEC emission model with $kT = 0.60$ keV and $Z = 0.5Z_{\odot}$) to convert to an unabsorbed flux ($9.08 - 10.52 \times 10^{-14}$ erg $cm^{-2} s^{-1}$) and an unabsorbed luminosity ($L_{0.6-2} = 3.4 - 3.9 \times 10^{40}$ erg s^{-1}). We examined APEC models in XSPEC using the above parameters, and found that the acceptable fits all correspond to a normalization between $8 - 9 \times 10^{-5}$. This normalization can be converted into a central electron density ($n_0 = 1.2 - 2.8 \times 10^{-2} cm^{-3}$) and an enclosed gas mass within 50 kpc of $4.9 - 5.2 \times 10^9 M_{\odot}$ after applying a 30% upwards correction

to account for the cosmological Helium fraction).

If we extrapolate this integration out to 500 kpc (the approximate virial radius of this galaxy), the count rate increases to 0.018-0.036 count s⁻¹, corresponding to an unabsorbed flux of $1.50 - 3.06 \times 10^{-13}$ erg cm⁻² s⁻¹) and an unabsorbed luminosity of $L_{0.6-2} = 5.6 - 11.5 \times 10^{40}$ erg s⁻¹. The enclosed gas mass within 500 kpc is $1.4 - 2.6 \times 10^{11} M_{\odot}$. The best-fit profile corresponds to an enclosed gas mass of $2.3 \times 10^{11} M_{\odot}$.

3.7.1 Flattened Profiles

We also examined joint and individual fits with two-component β -models. One component is the $\beta \sim 0.5$ profile described above, and another is a much flatter profile ($\beta \sim 0.35$) with much a larger core radius ($r_0 \sim 100$ kpc). This joint profile is motivated by theoretical work that suggests hot halos might have higher entropy than expected, which would flatten the profiles and allow them to contain more material (Crain et al. 2010a, Kaufmann et al. 2009, Guedes et al. 2011). In a previous paper (Anderson and Bregman 2010) we used a fiducial slope of $\beta = 0.35$ for flattened profiles, so we constrain the flattened profile to have that slope in our fit. We also fix the core radius at 50 kpc to reduce the parameter space, but our conclusion is not very sensitive to this parameter. We choose logarithmically-spaced values for the normalization of the flattened component, and for each normalization, we search for values of β , r_0 , and S_0 corresponding to the non-flattened ($\beta \sim 0.5$) profile, such that we minimize the total χ^2 of the fit to the data of the sum of the flattened and non-flattened profiles. We fit the data out to radii of 500 arcseconds so as to incorporate constraints from beyond the region of visible emission.

The resulting minimum χ^2 as a function of the normalization of the flattened profile component, and the strongest constraint (ruled out at 95%) is a maximum normalization of $S_0 = 1.5 \times 10^{-10}$ count s⁻¹ cm⁻² arcsec⁻². This corresponds to a total count rate out to 500 kpc of 6.5×10^{-2} counts s⁻¹, or an unabsorbed flux of 5.2×10^{-13} erg s⁻¹ cm⁻², or a 0.6-2 keV luminosity of 2.0×10^{41} erg s⁻¹. The normalization is 1.1×10^{-3} , corresponding to a hot halo mass in the flattened component of 7.4×10^{11}

M_{\odot} .

3.8 Implications and Conclusion

A full treatment of the implications of this measurement and this unusual galaxy is best reserved for future work, so here we comment briefly on a few salient points.

3.8.1 Generalizing the Conjugate Subtraction Technique

Given the historical observational difficulties in detecting diffuse emission at large radii around spiral galaxies, and our own experience and difficulties with this observation, we point out a few lessons learned that may be of use for future observers in this field. First, it is inadvisable to place the galaxy too close to the aimpoint of the telescope. There are numerous instrumental effects (vignetting, point source detectability, decreasing throughput, etc) that all vary with distance from the aimpoint; while corrections can be applied for each, it is very difficult to disentangle instrumental effects from diffuse emission at or below the surface brightness of the background. Second, it is also inadvisable to place the galaxy at the very edge of the detector (as we did for observations 10528 and 10529). While this configuration allows for conjugate subtraction to circumvent many of the radially varying instrumental effects, the throughput and vignetting effects are so large at the edge of the detector that it becomes difficult to see the galaxy. The optimal configuration is probably similar to our observation 10531, where the galaxy is placed a few arcminutes from the edge of the detector. This allows for conjugate subtraction while still retaining enough throughput to detect faint emission. The primary shortcoming of this configuration is that the conjugate annuli begin to overlap sooner than they would if the galaxy is at the edge of the detector. For observation 10531, we can only probe the inner 4.5 arcminutes of the halo, but the halo falls below the background at a radius of about 3 arcminutes, so we can still see all the useful emission.

As described in section 3.5.1, we also found that a quadratic fit to the soft X-ray background as a function of radius is generally statistically acceptable. Using quadratic fits to the background for conjugate subtraction, our technique is able to

detect emission down to about 20% of the background for very faint and diffuse emission (see Figure 3.4).

3.8.2 Halo Faintness and the Baryon Budget of NGC 1961

As mentioned in section 3.2, it has been known that hot halos around spiral galaxies are underluminous in comparison to hot halos around ellipticals, but now that we have a detection of one, we can finally quantify this. Benson et al. (2000) examined a sample of three large spirals, of which the largest (NGC 4594, the Sombrero galaxy) is similar in mass but much more bulge-dominated than NGC 1961. They predict a bolometric hot halo luminosity within the inner 76 kpc of $6.9 \pm 0.5 \times 10^{41}$ erg s⁻¹ for a metal abundance of $0.3Z_{\odot}$ (which would be 50% higher for our assumed abundance of $0.5Z_{\odot}$). This is not observed, but they do establish an observational upper limit of $4.4 \pm 2.8 \times 10^{40}$ erg s⁻¹ for the emission at radii between 16 and 30 kpc. Our detection falls below this limit. We find an absorbed 0.6-2 keV luminosity (which is about half the bolometric luminosity for 0.6 keV gas) within the 16-30 kpc annulus of $5.2 - 6.5 \times 10^{39}$ erg s⁻¹, which is 75% below their upper limit and only 2% as luminous as the theoretical prediction. The hypothetical flattened component would be even fainter within this annulus, with a luminosity of only 2.3×10^{39} erg s⁻¹.

The baryon fraction of this galaxy is dominated by the stars: the 2MASS K_S-band total magnitude of 7.73 corresponds to an absolute K_S magnitude of -26.0. Assuming a mass-to-light ratio of 0.6 (Bell and de Jong 2001), we infer a stellar mass of $3.1 \times 10^{11} M_{\odot}$. The cold gas component can be approximated by the HI mass, which is $4.7 \times 10^{10} M_{\odot}$ (Haan et al. 2008). We find a hot halo mass of $1.4 - 2.6 \times 10^{11} M_{\odot}$. So the baryon budget of this galaxy is approximately 6:1:3-5 stars : cold gas : hot gas (with the flattened profile the ratio is 6:1:14). Within 50 kpc, the baryon budget is approximately 60:10:1 stars : cold gas : hot gas. So, out to 50 kpc, most of the gas in NGC 1961 is cold, but if we are justified in integrating to the virial radius, the majority of the gas in the galaxy is in the hot phase, as predicted by theory for large galaxies.

It is important to highlight the uncertainties in our measurement of the hot halo

mass. Our β -model fitting has a 95%-confidence uncertainty of a factor of three in the halo mass. A 10% uncertainty in the distance to the galaxy would introduce a roughly 20% additional uncertainty in the estimated mass, and our choice of 500 kpc for an outer radius of integration probably introduces another 10% error relative to the true (unknown) virial radius. The assumptions of isothermality and azimuthal symmetry may not be true at larger radii as well. By far the biggest source of uncertainty, however, is the metallicity of the gas. At the virial temperature of this galaxy, X-ray emissivity is nearly linearly proportional to the gas metallicity, so if the gas is actually enriched to $1Z_{\odot}$ instead of our assumed value of $0.5Z_{\odot}$, the true halo mass is half the value we infer, and if the gas is actually $0.25Z_{\odot}$, the true halo mass is twice the value we infer. A metallicity gradient in the hot halo would make an accurate mass estimate even more difficult. Obtaining a measurement of the metallicity of the hot halo gas around a spiral galaxy is critically important for accurately measuring the mass (as well as constraining the formation history of the halo).

For our value of $Z = 0.5Z_{\odot}$, we can convert the baryon budget into a baryon fraction with an estimate of the total dark halo mass. The most recent measurement of the circular velocity finds an inclination-corrected HI velocity at 34 kpc of 402 km s^{-1} (Haan et al. 2008), which is comparable to older measurements for this galaxy (e.g. Rubin et al. 1979). If this measurement is correct and the dark matter follows an NFW profile, the expected virial mass is $M_{200} = 2.3 \times 10^5 v_f^3 h^{-1} M_{\odot}$ (Navarro 1998). So NGC 1961 would have an inferred mass of $2.1 \times 10^{13} M_{\odot}$ - the mass of a medium-sized galaxy group. Even if the halo is not precisely NFW, the total mass seems unlikely to differ by more than 50% or so.

Since NGC 1961 is at the center of a poor group, the other galaxies in its group should also be added to the baryon budget. We added the HI masses for the six other candidate group members (Haan et al. 2008) and found a total HI mass of $7.5 \times 10^9 M_{\odot}$, which is less than a sixth of the HI mass of NGC 1961. We estimated the stellar mass of the six group members using their H-band absolute magnitudes from NED and a mass-to-light ratio of 0.6; this yields a total stellar mass for the group members of $3 \times 10^{10} M_{\odot}$, less than a tenth of the stellar mass of NGC 1961.

NGC 1961 is therefore by far the dominant reservoir of baryons in its group, but we do include the small baryonic contributions from the other galaxies when we compute the baryon budget.

Within 500 kpc, the baryon fraction $f_b \equiv M_b/M_{\text{total}}$ is therefore 0.024-0.029 (or 0.051 with a flattened halo profile) - far less than the cosmological mean value $f_b = 0.171 \pm 0.006$ (Dunkley et al. 2009). This corresponds to a baryon fraction within R_{500} of 0.023-0.033 for the single-component fit, and a maximum baryon fraction within R_{500} of 0.043 for the two-component flattened β -model. Thus, within the virial radius NGC 1961 is missing over 75% of its baryons, which is surprising since the missing baryon fraction is nearly always smaller in structures of this size (McGaugh et al. 2010).

The baryonic Tully-Fisher relation (Stark, McGaugh, and Swaters 2009) predicts this galaxy should have a baryonic mass of $1.1 \times 10^{12} M_\odot$ (corresponding to a baryon fraction $f_b = 0.052$), so NGC 1961 is slightly below the BTF, but still potentially within the typical scatter about the relation. Including a flattened profile would definitely place this galaxy on the BTF, as would changing the gas metallicity to $0.2Z_\odot$. Additionally, using the higher inclination angle of Combes et al. (2009) (discussed in section 3.3) would also bring NGC 1961 onto the BTF due to the steep dependence of this relation on v_{circ} .

3.8.3 Halo Cooling Rates and Implications for Galaxy Formation

We can estimate the cooling radius of this hot halo and the implied accretion rate onto the galaxy, which has implications for setting and regulating the star formation rate in the galaxy. We define the cooling radius as the radius for which the cooling time is 10 Gyr, using the expression for cooling time from Fukugita and Peebles (2006):

$$\tau(r) = \frac{1.5nkT}{\Lambda n_e (n - n_e)} \approx \frac{1.5kT \times 1.92}{\Lambda n_e \times 0.92} \quad (3.1)$$

where the latter expression assumes primeval Helium abundance so the total particle density $n = 1.92n_e$. For $T = 10^{6.85}$ K and $Z = 0.5Z_\odot$, $\Lambda = 10^{-22.85}$ erg cm³ s⁻¹ (Sutherland and Dopita 1993). Thus the cooling radius occurs at

$n_e = 6.8 \times 10^{-4} \text{ cm}^{-3}$. For the range of best-fit β -model profiles listed above, this corresponds to a cooling radius between 17.8 and 18.2 kpc, and an interior hot halo mass of $8.9 - 10.2 \times 10^8 M_\odot$. It is difficult to estimate the accretion rate onto the disk from this hot halo, since the heating rate is unconstrained, but we can make an order-of-magnitude estimate by dividing the hot gas thermal energy within the 10 Gyr cooling radius by the luminosity within that radius; this yields a cooling time of 2.0-2.4 Gyr for material within the cooling radius, or an effective cooling rate of $0.4 M_\odot \text{ year}^{-1}$. In contrast, we can estimate the star formation rate in NGC 1961 from the total H α luminosity ($7.6 \pm 0.9 \times 10^{41} \text{ erg s}^{-1}$) using the relation in Kennicutt (1998): $\text{SFR} = 7.9 \times 10^{-42} L(\text{H}\alpha) = 6.0 \pm 0.7 M_\odot \text{ year}^{-1}$. The halo accretion rate is therefore insufficient to produce the star formation rate of the galaxy. More relevant for galaxy formation, the halo accretion rate is two orders of magnitude too low to assemble the stellar mass of this galaxy within a Hubble time. If we preserve β and r_0 for the halo, but increase S_0 to add the present-day stellar mass of $3.1 \times 10^{11} M_\odot$ to the halo, the cooling rate becomes $1.2 - 1.8 M_\odot \text{ year}^{-1}$, which is still insufficient to assemble the stellar mass by a factor of 20.

These results are also evidence against the emission around this galaxy being dominated by a galactic fountain or other internal processes related to star formation. One of the brightest known galactic fountains, in NGC 891, has a bolometric luminosity of $4 \times 10^{39} \text{ erg s}^{-1}$ (Bregman and Houck 1997) and a star formation rate of about $4 M_\odot \text{ year}^{-1}$ (Strickland et al. 2004b). For the NGC 1961 SFR of $6 M_\odot \text{ year}^{-1}$, the highest expected luminosity of a fountain is therefore $\sim 6 \times 10^{39} \text{ erg s}^{-1}$, whereas we measure an (absorbed) luminosity within 50 kpc of $2.6 - 3.0 \times 10^{40} \text{ erg s}^{-1}$. Moreover, the cooling time for the gas we observe beyond 18 kpc is greater than 10 Gyr, so even if the material has an internal origin, it should be treated as quasi-static instead of as a fountain.

Within the context of galaxy formation models, the physical explanation for the missing baryons from NGC 1961 is unclear. The amount of missing matter in this galaxy and the depth of the potential well make it difficult for supernova-driven winds to expel the missing baryons. The escape velocity for material originating in the

galactic disk (at an assumed radius of 10 kpc) in an NFW halo of mass $2.1 \times 10^{13} M_{\odot}$ is about 1300 km s^{-1} . The missing baryonic mass is $3.2 \times 10^{12} M_{\odot}$, so the energy needed to unbind the missing baryons from the gravitational potential of the galaxy is $\sim 5 \times 10^{61}$ erg. A supernova with perfect coupling to the interstellar medium can inject up to $\sim 10^{51}$ erg of kinetic energy, so expelling the missing baryons from this galaxy requires $\sim 5 \times 10^{10}$ supernovae. Based on the present-day stellar mass of the galaxy and the average supernova rates per stellar mass per century from Mannucci et al. (2005), the expected total amount of supernovae (Type I + II) over 13 Gyr is 4×10^9 , so NGC 1961 would need to have 13 times more stellar mass to supply enough energy from supernovae to unbind its missing baryons. Performing this calculation in more detail (e.g. accounting for smaller stellar mass at earlier times, allowing for cooling of supernova-heated gas) in general exacerbates the discrepancy.

We also considered the possibility of AGN feedback for ejecting the missing baryons. However, NGC 1961 is a late-type spiral galaxy, so it has a relatively small bulge and presumably a correspondingly small central black hole, which makes an AGN-driven superwind less plausible. For example, if we assume the bulge contains a sixth of the stellar mass for this galaxy, and use the bulge mass-black hole mass relation (Marconi and Hunt 2003), we estimate a central black hole of mass $1 \times 10^8 M_{\odot}$. Assuming accretion converts 10% of the infalling mass into energy, over its lifetime this black hole could have produced 2×10^{61} erg, of which it is generally assumed only about 2% can couple to the interstellar medium (Shankar et al. 2006, Bower et al. 2008) for a total available AGN feedback energy of only 4×10^{59} erg. The most likely explanation for the missing baryons from this galaxy is probably some form of very early pre-heating which prevents the baryons from falling deeply into the potential well in the first place.

We also note that other giant spirals show similar baryon deficits. For example, UGC 12591, at a distance of 134 Mpc (NED average) has a peak (inclination-corrected) HI rotation velocity of $476 \pm 23 \text{ km s}^{-1}$ (Giovanelli & Haynes 1985), which yields a virial mass of $3.5 \times 10^{13} M_{\odot}$ according to the NFW relation cited above. The HI mass is $4.3 \times 10^9 M_{\odot}$ (Ho 2007), and, using the 2MASS K_S -band total magnitude

(8.9) as above we estimate a total stellar mass of $5.9 \times 10^{11} M_{\odot}$, for a baryon fraction of 0.017 before including the hot halo. We also obtained XMM-Newton observations of the hot halo around this galaxy (Dai 2011, in preparation; see Chapter 4) which point to a similarly underluminous halo, to be discussed further in an upcoming paper.

In future work we hope to increase the sample of hot halos detected around giant spirals and further quantify the halo properties such as metallicity. If the trend of underluminous halos is common for the giant spirals, it would signal a break in the BTF somewhere between Milky-Way mass galaxies and the giant spirals.

3.9 Acknowledgements

The authors would like to thank the anonymous referee for a very thoughtful report which greatly improved the paper. We would also like to thank E. Bell, N. Calvet, X. Dai, A. Muratov, M. Ruszkowski, and C. Slater for helpful comments and discussions during the research and writing of this paper. We would also like to thank X. Dai for sharing with us his preliminary results and methods, S. Haan for helpfully sending his HI observations of this galaxy, and N. Bonaventura at the CXC help desk for assistance in reducing the Chandra data. This research has made use of NASA's Astrophysics Data System. This research has made use of the NASA/IPAC Extragalactic Database (NED) which is operated by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. This research has made use of software provided by the High Energy Astrophysics Science Archive Research Center (HEASARC), which is a service of the Astrophysics Science Division at NASA/GSFC and the High Energy Astrophysics Division of the Smithsonian Astrophysical Observatory. We gratefully acknowledge financial support from NASA through Chandra grant GO9-0089X. M.E.A. also acknowledges support from the NSF in the form of a Graduate Research Fellowship.

3.10 References

- Anderson, M. E. and Bregman, J. N. 2010, *ApJ* 714, 320
Bell, E. F. and de Jong, R. S. 2001, *ApJ* 550, 212

Benson, A. J. et al. 2000, MNRAS 314, 557
Benson, A. J. 2010, PhR 495, 33
Bower, R. G. et al. 2006, MNRAS 370, 645
Bower, R. G. et al. 2008, MNRAS 390, 1399
Bregman, J. N. and Houck, J. C. 1997, ApJ 485, 159
Combes, F. et al. 2009, A&A 503, 73
Crain, R. A. et al. 2010a, MNRAS 407, 1403
Crain, R. A. et al. 2010b, MNRAS submitted, astro-ph/1011.1906
Dai, X. et al. 2011, in preparation
De Lucia, G. et al. 2006, MNRAS 366, 499
Dekel, A. and Birnboim, Y. 2006, MNRAS 368, 2
Dickey, J. M. and Lockman, F. J. 1990, ARA&A 28, 215
Dressler, A. 1980, ApJ 236, 351
Dunkley, J. et al. 2009, ApJS 180, 306
Forman, W., Jones, C., and Tucker, W. 1985, ApJ 293, 102
Fukugita, M. and Peebles, P. J. E. 2006, ApJ 639, 590
Giovanelli, R and Haynes, M. P. 1985, AJ 90, 2445
Gottesman, S. T. et al. 2002, MNRAS 337, 34
Grimm, H. J., Gilfanov, M., and Sunyaev, R. 2002, A&A 391, 923
Guedes, J., Callegari, S., Madau, P., & Mayer, L. 2011, arXiv:1103.6030
Haan, S. et al. 2008, AJ 135, 232
Ho, L. C. 2007, ApJ 668, 94
Humphrey, P. J. et al. 2011, ApJ 729, 53
Irwin, J. A., Athey, A. E., and Bregman, J. N. 2003, ApJ 587, 356
Jarrett, T. H. et al. 2003, AJ 125, 525
Kaufmann, T. et al. 2009, MNRAS 396, 191
Kennicutt, R. 1998. ARA&A 36, 189
Kim, M. et al. 2007. ApJ 659, 29
Klypin, A., Zhao, H. S., and Somerville, R. S. 2002, ApJ 573, 597
Li, Z. et al. 2006, MNRAS 371, 147

Li, Z., Wang, Q. D., and Hameed, S. 2007, MNRAS 376, 960
O'Sullivan, E., Forbes, D. A., and Ponman, T. J. 2001, MNRAS 328, 461
Owen, R. A. and Warwick, R. S. 2009, MNRAS 394, 1741
Mannucci, F. et al. 2005, A&A 433, 807
Marconi, A. and Hunt, L. K. 2003, ApJ 589, 21L
McGaugh, S. S. et al. 2010, ApJ 708, 14
Mulchaey, J. S., and Jeltema, T. E. 2010, ApJ 715, L1
Navarro, J. F. 1998, preprint (astro-ph/9807084)
Pedersen, K. et al. 2006, New Astron. 11, 465
Rasmussen, J. et al. 2009, ApJ 697, 79
Read, A. M. and Ponman, T. J. 1998, MNRAS 297, 143
Revnivtsev, M. et al. 2008. A&A 490, 37
Roberts, T. P. and Warwick, R. S. 2000, MNRAS 315, 98
Robertson, B. et al. 2006, ApJ 645, 896
Rubin, V. C. et al. 1979, ApJ 230, 35
Sarazin, C. L. 1986, RevModPhys 58, 1
Shankar, F. et al. 2006, ApJ 643, 14
Shattow, G. and Loeb, A. 2009, MNRAS 392, L21
Sommer-Larsen, J. 2006, ApJ 644, L1
Stark, D. V., McGaugh, S. S., and Swaters, R. A. 2009, AJ 138, 392
Strickland, D. K. et al. 2004a, ApJS 151, 193
Strickland, D. K. et al. 2004b, ApJ 606, 829
Sutherland, R. S. and Dopita, M. A. 1993, ApJS 88, 253
Tüllmann, R. et al. 2006, A&A 448, 43
Voss, R. et al. 2009, ApJ 701, 471
White, S. D. M. and Frenk, C. S. 1991, ApJ 379, 52
White, S. D. M. and Rees, M. J. 1978, MNRAS 183, 341
Yamasaki, N. Y. et al. 2009, PASJ 61, 291




Figure 3.1 An ESO DSS image of NGC 1961, with the layout of our four ACIS-I observations overlaid. North is up, East is left, and each box is about 17 arcmin (280 kpc) on a side.




Figure 3.2 Our four Chandra ACIS-I observations of NGC 1961. Observation 10528 is top left, observation 10529 is top right, observation 10530 is bottom left, observation 10531 is bottom right. Each image uses 5-pixel Gaussian smoothing and square root scaling, and North is up and East is left. The annuli show our in-field conjugate background subtraction technique for measuring a background-subtracted surface brightness profile around NGC1961. We use annuli around the source and subtract background from the annulus of equivalent size on the opposite side of the aimpoint. The solid annulus denotes an example source region, and the dashed annulus an example conjugate region. The red shaded region represents the area on each detector for which we were unable to subtract a conjugate background from the 180° annulus; for data within these regions in observations 10528 and 10529 we used conjugate regions at 120° and 240° instead. For observation 10531, at radii larger than 4.5 arcseconds the annuli cross each other, so we cannot use points at larger radius.




Figure 3.3 Background-subtracted radial surface brightness profile of NGC 1961, including binned data from all four observations in the same figure. The black line is the best-fit β -model to the data. The model has parameters $S_0 = 9.77 \times 10^{-8}$ count s^{-1} cm^{-2} $arcsec^{-2}$, $r_0 = 1.00$ kpc, $\beta = 0.47$. The χ^2 is 66.9 for 50 degrees of freedom, so the fit is inconsistent with the data at less than 95% confidence.




Figure 3.4 Background-subtracted radial surface brightness profiles for all four observations after background smoothing (see text). The smoothed fits to the background are represented by the colored dashed lines (the dotted lines begin where the background data are extrapolated beyond the edge of the detector). The black line is the best-fit β -model to the data. The model has parameters $S_0 = 8.90 \times 10^{-8}$ count s $^{-1}$ cm $^{-2}$ arcsec $^{-2}$, $r_0 = 1.00$ kpc, $\beta = 0.46$. The χ^2 is 62.1 for 50 degrees of freedom, so the fit is inconsistent with the data at less than 95% confidence. The gray shaded region denotes all the acceptable fits to the data at this confidence level; we use this region as the fiducial range throughout the paper.




Figure 3.5 Log-log plot of radial surface brightness profiles for all four observations. The black dashed line is the estimated contribution of stars, and the colored dashed lines are the estimated contributions of X-ray binaries. The colored data points are the surface brightness profile with resolved and unresolved point sources subtracted. Unlike Figures 3.3 and 3.4, we have not subtracted the sky X-ray background from the surface brightness profile. The smoothed sky X-ray background is indicated by the four dotted colored lines. We detect emission above the background out to 40-50 kpc which is more spatially extended than the other galactic components.

CHAPTER 4

XMM-Newton Detects a Hot Gaseous Halo in the Fastest Rotating Spiral Galaxy UGC 12591

Note: This chapter is adapted from a paper of the same title, published in the *Astrophysical Journal*, on which I was the second author. Because I am only the second author, I have rewritten the paper for this thesis. This chapter highlights my contributions to the paper, and the paper's contributions to our understanding of hot halos around galaxies. Exceptions are the section on spatial analysis and much of the discussion (Sections 4.4 and 4.5), which were written by me in the original paper, and have therefore been reproduced here closer to their original form. The reference is Dai, X., Anderson, M. E., Bregman, J. N., and Miller, J. M. 2012, *ApJ*, 755, 107-117. The paper is copyright 2012, the American Astronomical Society, reproduced here under the non-exclusive right of republication granted by the AAS to the author(s) of the paper.

4.1 Introduction

Here we describe our observation and analysis of one of the most rapidly rotating galaxies known, the early-type spiral UGC 12591. This galaxy has a redshift of $z = 0.023$, corresponding to a distance of about 100 Mpc. Its HI rotational velocity is $466 - 500 \text{ km s}^{-1}$ (Giovanelli et al. 1986, Paturel et al. 2003), and its K-band luminosity is about nine times the K-band luminosity of M31. Thus this galaxy is even more massive than NGC 1961; it rotates nearly 100 km s^{-1} faster, and is 50% more luminous in the near-IR. It also significantly more bulge-dominated than NGC 1961, with an S0/a morphology (de Vaucouleurs et al. 1993). For this analysis, we

adopt $H_0 = 72 \text{ km s}^{-1} \text{ Mpc}^{-1}$.

4.2 Observation and Data Reduction

In contrast to our strategy for NGC 1961 (featuring four observations around the galaxy with the galaxy itself placed near the edge of the detectors) we placed UGC 12591 at the aimpoint of the observation. We observed this galaxy for 80 ks with XMM-Newton. The data reduction was simple and standard. We quote Dai et al. (2012) for the details of the data reduction: "We reprocessed the PN and MOS data using the SAS9.0.0 software tools `epchain` and `emchain`, and filtered the events with the patterns ≤ 4 and ≤ 12 for the PN and MOS chips, respectively. We filtered background flares by excluding the intervals with background count rates of $\text{CR} > 0.4 \text{ counts s}^{-1}$ and $\text{CR} > 0.2 \text{ counts s}^{-1}$ in the 10-12 keV band in PN and MOS observations, respectively, following the standard suggestion from SAS. We also applied a low energy flare filter to exclude flares in the 0.61.4 keV band. We obtained net exposure times of 31.2, 50.0, and 46.0 ks for PN, MOS1 and MOS2 CCDs, respectively. We detected serendipitous sources in the field using the SAS tools, and masked the bright serendipitous source regions in the subsequent analysis with a flux limit of $1.5 \times 10^{-14} \text{ erg cm}^{-2} \text{ s}^{-1}$ ".

We estimate the background in-field, as with NGC 1961. The details of this process are somewhat different from that observation, however, as might be expected from the different observing configuration and the difference of telescope. This procedure is based on the discussions in Kuntz & Snowden (2008) and Leccardi & Molendi (2008). The idea is to decompose the background into vignettted and unvignettted components by fitting the spectrum of the background at large off-axis angles, where very little source contamination is expected. The background is modeled as cosmic rays and a soft proton background (unvignettted), and Galactic and extragalactic photons (vignettted). The flare filtering described above removes most of the time-variable soft proton background as well, since the flares are largely composed of soft protons. The remaining components are then estimated from a spectral fit to the outer regions of the image; once we fit for the various components, we can then extrapolate them

inwards.

The quiescent soft photons are assumed to follow an unvignetted $\Gamma = 0.4$ powerlaw on the PN chip and a broken powerlaw ($\Gamma_1 = 0.4$, $\Gamma_2 = 0.8$, $E_b = 5$ keV) on the MOS chips, with free normalizations. The extragalactic background was parameterized with a $\Gamma = 1.4$ powerlaw, and the Galactic background with an APEC (Smith et al. 2001) model with Solar metallicity and $kT = 0.197$ keV. Both these components also were given free normalizations across all three chips.

The fits to the off-axis spectra on each chip give fairly similar normalizations in each case, and the models fit the data with $\chi^2/\text{d.o.f.}$ of 327.5/261, 117.4/74, and 96.4/66 on the PN, MOS1, and MOS2 respectively. The cosmic X-ray backgrounds in the model also match well with observations of the X-ray background in this region of the sky (i.e. Cowie et al. 2002, Snowden et al. 1997).

4.3 Spectral Analysis

Having measured the background, we can also now analyze the source spectrum. We extracted the spectrum within a circle centered on the galaxy with radius 50'' (25 kpc), since the hot gas has a very high S/N within this radius. We binned the photons into bins of 20 photons, and fit the spectrum from 0.4-2.0 keV. We used a similar spectral model to that of NGC 1961: an APEC model for the hot gas and an absorbed powerlaw model ($\Gamma = 1.56$) for X-ray binaries in the galaxy. We fixed the normalization of the powerlaw model to the expected value based on extrapolation from the 2-8 keV spectrum, in similar fashion to our extrapolation of the X-ray binary (XRB) contribution to the surface brightness profile for NGC 1961. We also fixed the Galactic absorption and we fixed the hot halo metallicity to $0.5Z_{\odot}$.

The best-fit model fits the data with $\chi^2/\text{d.o.f.} = 129.2/115$, which is an acceptable fit. We measure the hot gas temperature to be 0.64 ± 0.03 keV, which is identical to within the uncertainties to the temperature of the hot gas measured around NGC 1961. We also infer an internal absorption of $9 \pm 2 \times 10^{20}$ cm⁻² in this galaxy. The point source component has a luminosity of 1.1×10^{41} erg s⁻¹, which is quite large, so we conclude that this component must also contain AGN emission from the galaxy.

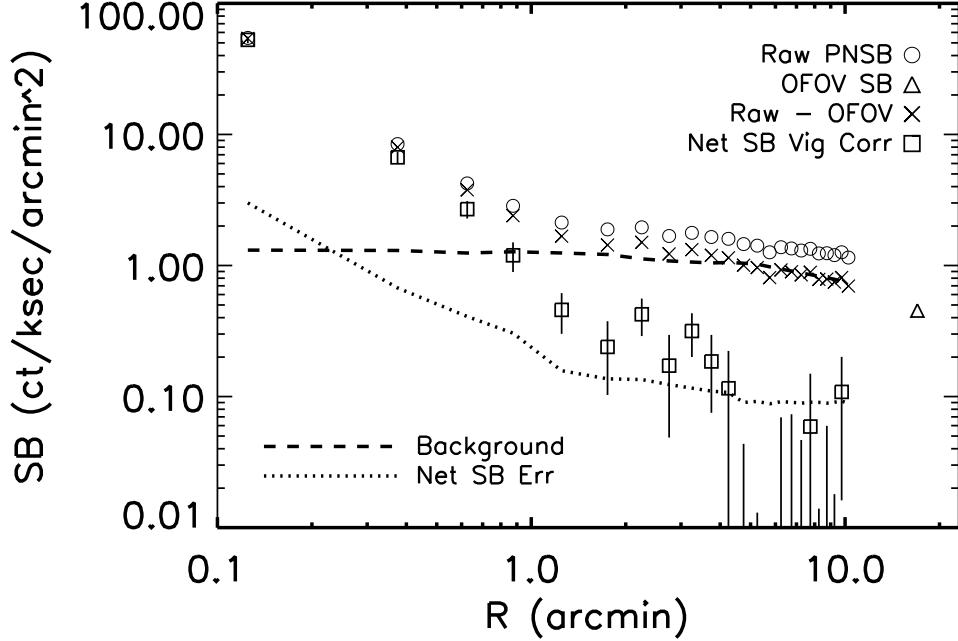


Figure 4.1 X-ray surface brightness profile of UGC 12591 in the XMM-Newton PN observation (0.6-1.4 keV). The circles represent the raw surface brightness profile, and the triangle represents the OFOV surface brightness. The cross symbols are results of subtracting the the OFOV surface brightness from the raw surface brightness profile. The dashed line is our model for the background, including vignettted and unvignettted components, based on the outer regions. The squares and their associated error-bars denote the net surface brightness profile and its uncertainty (crosses minus the dashed line) corrected for vignetting. We also plot the uncertainties as dotted lines.

This is consistent with a radio detection of this galaxy with the VLA (Condon et al. 1991).

4.4 Spatial Analysis

Now we examine the surface brightness profile of the X-ray emission in order to infer the density profile of the hot gas. The surface brightness profiles for the PN, MOS1, and MOS2 chips are shown in Figures 4.1, 4.2, and 4.3.

We fit the surface brightness profile in the 0.61.4 keV band using three components, the hot gas, AGN/accretion onto X-ray binaries, and the contribution from the stellar population (including cataclysmic variables). The AGN/XRB contribu-




Figure 4.2 X-ray surface brightness profile of UGC 12591 in the XMM-Newton M1 observation (0.6-1.4 keV). The circles represent the raw surface brightness profile, and the triangle represents the OFOV surface brightness. The cross symbols are results of subtracting the the OFOV surface brightness from the raw surface brightness profile. The dashed line is our model for the background, including vignettted and unvignettted components, based on the outer regions. The squares and their associated error-bars denote the net surface brightness profile and its uncertainty (crosses minus the dashed line) corrected for vignetting. We also plot the uncertainties as dotted lines.

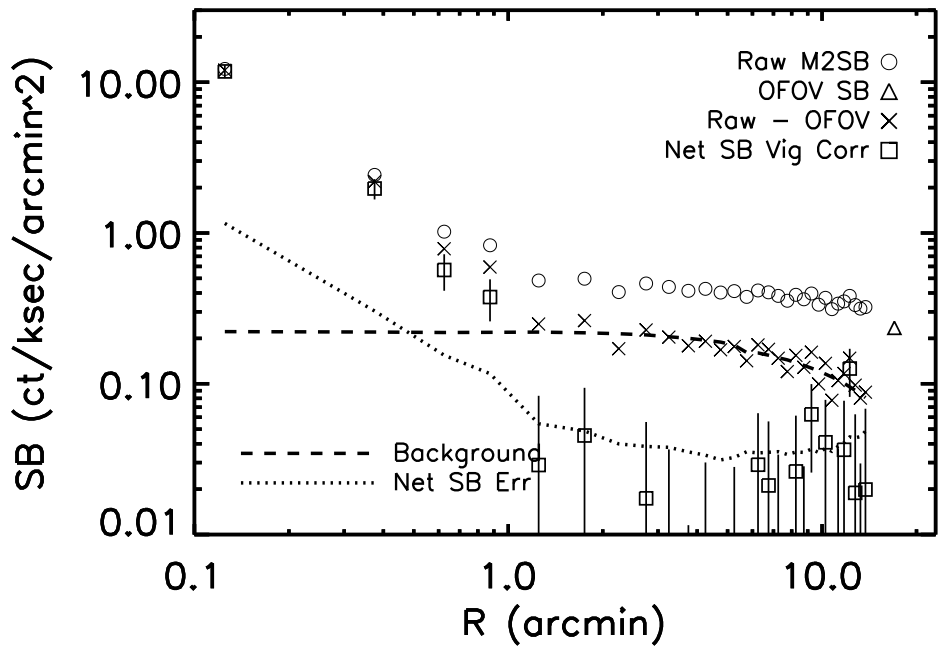


Figure 4.3 X-ray surface brightness profile of UGC 12591 in the XMM-Newton M2 observation (0.6-1.4 keV). The circles represent the raw surface brightness profile, and the triangle represents the OFOV surface brightness. The cross symbols are results of subtracting the the OFOV surface brightness from the raw surface brightness profile. The dashed line is our model for the background, including vignettted and unvignettted components, based on the outer regions. The squares and their associated error-bars denote the net surface brightness profile and its uncertainty (crosses minus the dashed line) corrected for vignetting. We also plot the uncertainties as dotted lines.

tion can be measured using the surface brightness profiles from the 28 keV band, as described above. We estimated the x-ray emission from the stellar population (e.g., M stars, but excluding the AGN/binaries) using the stellar mass to x-ray luminosity conversion factor of Revnivtsev et al. (2008). This relation is calibrated for old stellar populations, and since UGC 12591 is fairly bulge-dominated, the relation should still work approximately. To compute the radial surface brightness profile from the stellar population in the X-ray band, we derived a K-band radial surface brightness profile from the K-band magnitudes within circles of different angular sizes for this galaxy listed in the 2MASS Extended Source Catalog. We used a distance modulus of 35.0 and a K-band mass-to-light ratio of 0.6 (Bell & de Jong 2001) for this galaxy, and then applied the Revnivtsev et al. (2008) conversion to derive an X-ray surface brightness profile. We also added a systematic uncertainty of $0.08 \text{ counts ks}^{-1} \text{ arcmin}^{-2}$ in each data point to account for the uncertainties in modeling the background. The remaining emission is attributed to the hot gas in the halo. We fit the surface brightness profile of this emission using a standard β -model.

Figure 4.4 shows the XMM-Newton PN and MOS data, after vignetting correction and background subtraction, as well as the estimates of the various components we fit to the surface brightness profile. It is clear that out to 30 or 40 kpc the X-ray binary and stellar components are insufficient to account for all the emission, and therefore that a hot halo component is necessary. We only allowed β -models for the hot gas if they cannot be excluded at greater than 95% confidence. This results in a narrow range of acceptable fits for the hot halo profile. The true uncertainties in the surface brightness profile are probably somewhat larger than the statistical errors, however, due to inevitable systematic errors in the flat fielding and background subtraction. This means that the range of formally acceptable fits for the hot halo might be a little wider than we have indicated (Figure 4.4). Deviations from a simple β -model for the hot gas distribution are also possible, and we tested by adding a second, flatter β -model component into our fits as well.

The acceptable fit with the highest enclosed mass has $\beta = 0.40$, $r_0 = 0.53 \text{ kpc}$, and $S_0 = 87.0 \text{ counts ks}^{-1} \text{ arcmin}^{-2}$. This corresponds to a count rate within a




Figure 4.4 Background-subtracted, vignetting-corrected XMM-Newton PN and MOS radial surface brightness profiles of the 0.6-1.4 keV emission around UGC 12591 (open circles denote negative values). The PN emission has been multiplied by a factor of 0.26 to convert it into effective counts for the MOS detectors. We fit the surface brightness profile as the sum of three components, the hot gas halo, AGN/binaries, and the stellar contribution. The black dashed line is the emission from coronally active stars and cataclysmic variables in the galaxy, estimated by scaling the 2MASS K-band luminosity. The colored dashed lines are the emission from X-ray binaries/AGN, estimated by scaling from the 2-8 keV surface brightness profile. Finally, the shaded region denotes the emission from hot gas, which makes up the rest of the observed surface brightness profile. The black solid line is the best-fit β -model for the hot halo component with a χ^2 of 61.1 for 45 degrees of freedom, and the yellow solid line shows our best fit by adding an additional flattered β -profile.

projected radius of 50 kpc of 0.0039 counts s⁻¹. Assuming an APEC model with $kT = 0.64$, $Z = 0.5Z_{\odot}$, and $N_H = 9 \times 10^{20}$ cm⁻², we infer the mass within 50 kpc to be $5.1 \times 10^9 M_{\odot}$ and the unabsorbed luminosity (0.61.4 keV) to be 2.8×10^{40} erg s⁻¹ given a distance of 100 Mpc. The acceptable fit with the lowest enclosed mass has $\beta = 0.83$, $r_0 = 10.6$ kpc, and $S_0 = 6.1$ counts ks⁻¹ arcmin⁻². This corresponds to a count rate within a projected radius of 50 kpc of 0.0025 counts s⁻¹, yielding a mass of $3.3 \times 10^9 M_{\odot}$ and an unabsorbed luminosity of 1.8×10^{40} erg s⁻¹. The best-fit profile has $\beta = 0.48$, $r_0 = 3.04$ kpc, and $S_0 = 14.2$ counts ks⁻¹ arcmin⁻². The χ^2 is 2.7 for 6 degrees of freedom. This corresponds to a count rate within a projected radius of 50 kpc of 0.0032 counts s⁻¹, yielding a mass of $4.4 \times 10^9 M_{\odot}$ and an unabsorbed luminosity of 2.3×10^{40} erg s⁻¹. If we integrate these profiles out to 500 kpc, then the fit with the highest enclosed mass contains $3.1 \times 10^{11} M_{\odot}$ with an unabsorbed luminosity of 1.1×10^{41} erg s⁻¹. The fit with the lowest enclosed mass contains $1.9 \times 10^{10} M_{\odot}$ with an unabsorbed luminosity of 1.9×10^{40} erg s⁻¹. The best fit contains $1.5 \times 10^{11} M_{\odot}$ with an unabsorbed luminosity of 4.7×10^{40} erg s⁻¹.

We also examined the possibility of a higher-entropy halo as predicted by many simulations (Maller & Bullock 2004; Kauffmann et al. 2009; Crain et al. 2010; Guedes et al. 2011). Such a halo would have a flatter density profile than the $\beta \sim 0.5$ models we find above, and it would therefore be both more difficult to detect in emission and also more massive than a lower-entropy halo. As in Anderson & Bregman (2011), we chose to model a flattened profile with a two-component fit to the data. The flattened component is a β -model with fixed $\beta = 0.35$ and $r_0 = 50$ kpc, but with free normalization, and the other component is a more concentrated β -model with all three parameters free, used to model the emission at smaller radii. For this galaxy, there is very little statistical space left to add a flattened profile, so the most mass that can be included by adding a flattened component is $4.5 \times 10^{11} M_{\odot}$. As before, however, we caution that this statistical constraint depends on understanding all the systematic uncertainties perfectly, especially the background in XMM-Newton PN/MOS CCDs. These profiles can be independently tested in absorption profiles instead of emission due to the linear dependence on density in absorption. In addition, we note that the

metallicity assumed can be another uncertainty in our analysis since the total mass will depend on the metallicity.

4.5 Discussion

4.5.1 Baryon Mass Components in UGC 12591

We have constrained the hot gas mass of UGC 12591 using the XMM-Newton observation as $(4.1 \pm 0.3) \times 10^9 M_\odot$ within 50 kpc with an average temperature of $T = 0.64 \pm 0.03$ keV. We have also constrained the hot gas mass of $1.5 \times 10^{11} M_\odot$ within 500 kpc regions using our best-fit β -model, and the hot gas mass is below $4.5 \times 10^{11} M_\odot$ within 500 kpc even if we add another flatter β -model component in our fits. Beside the hot gas mass, there are other baryon mass components in the galaxy including the stellar mass and atomic and molecular gas mass components. For the cold gas mass component, Giovanelli et al. (1986) measure the *HI* mass to be $5.3 \times 10^9 h_{72}^{-2} M_\odot$. With a correction of 4/3 to account for Helium, this yields a total atomic gas mass of $7.1 \times 10^9 M_\odot$. Compared to these quantities, we assume the molecular gas mass to be negligible. And we estimate the stellar mass to be $(4.5 \pm 1.0) \times 10^{11} M_\odot$ within a 29 kpc radius, using its K-band total magnitude ($K = 8.89$ mag) from the 2MASS Extended Source Catalog and a range of mass-to-light ratio from 0.6 to 0.95 (e.g., Bell et al. 2003).

We estimate the dark matter mass in two different ways. Extrapolating the rotation curve out to 500 kpc (near the virial radius) gives us a virial mass of $2.7 \times 10^{13} M_\odot$. Using the equation of hydrostatic equilibrium and the X-ray gas temperature, on the other hand, gives us a mass of $1.9 \times 10^{13} M_\odot$. We adopt the lower estimate, both because it is more conservative and because the rotation curve is only measured out to 28 kpc and so the extrapolation for that method is large.

The baryon mass within 500 kpc is $6.1 \times 10^{11} M_\odot$, and we measure a baryon fraction of $f_b \approx 0.03$. Considering a second flattened gas component, the baryon fraction within 500 kpc can reach to $f_b \lesssim 0.05$. Since we use the smaller total mass estimate in the calculation, the baryon fraction quoted should be treated as a conservative upper

limit. To summarize, combining our XMM-Newton observation and the 2MASS and radio data in the literature, we have constrained that UGC 12591 has lost at least 70% of the baryons compared to the cosmological value. The missing baryons do not reside in the hot halos for spiral galaxies. Our result confirms the recent measurements in another giant spiral NGC 1961 using Chandra by Anderson & Bregman (2011), who find that NGC 1961 has also lost 75% of its baryon content.

4.5.2 Cooling of the hot gas

The hot halo is predicted to play an important role in galaxy formation and evolution. With our detection of the hot halo around UGC 12591, we can estimate the cooling time of this hot gas and the implied accretion rate onto the galaxy, which can provide constraints on the gas available for new star formation. We define the cooling radius as the radius where the cooling time is 10 Gyr, using the expression of the cooling time (Fukugita & Peebles 2006):

$$\tau(r) = \frac{1.5nkT}{\Lambda n_e(n - n_e)} = \frac{1.5kT \times 1.92}{\Lambda n_e \times 0.92}$$

where the latter expression assumes a primeval He abundance resulting in a total particle density of $n = 1.92n_e$. For $T = 10^{6.85}$ K, $Z/Z_\odot = 0.5$, and $\Lambda = 10^{-22.85}$ erg $\text{cm}^3 \text{s}^{-1}$ (Sutherland & Dopita 1993), the cooling radius is at $n_e = 6.8 \times 10^{-4} \text{ cm}^{-3}$. For the range of best-fit β -model profiles constrained in this paper, this corresponds to a cooling radius between 15.6 and 18.0 kpc, and a hot halo mass of $6.2 - 9.2 \times 10^8 M_\odot$ within that radius. We can roughly estimate the cooling time and rate by dividing the thermal energy in the hot gas within the cooling radius by the luminosity within that radius, and this yields a wide range in cooling time of 2.8-6.3 Gyr for material within the cooling radius, but a fairly narrow range in the effective cooling rate of $0.15 - 0.21 M_\odot \text{ yr}^{-1}$. This halo accretion rate is two orders of magnitude too low to assemble the stellar mass of this galaxy within a Hubble time. Therefore, significant accretion must have occurred via some other mode, such as cold flows or mergers, to produce the stellar mass seen in this galaxy today, supporting the conclusion drawn for NGC 1961 (Anderson & Bregman 2011).

4.6 References

- Anderson, M. E. & Bregman, J. N. 2011, *ApJ*, 737, 22
- Bell, E. F. & de Jong R. S. 2001, *ApJ*, 550, 212
- Bell, E. F., McIntosh, D. H., Katz, N., & Weinberg, M. D. 2003, *ApJS*, 149, 289
- Condon, J. J., Frayer, D. T., & Broderick, J. J. 1991, *AJ*, 101, 362
- Cowie, L. L., Garmire, G. P., Bautz, M. W., et al. 2002, *ApJ*, 566, L5
- Crain, R. A., McCarthy, I. G., Frenk, C. S., Theuns, T., & Schaye, J. 2010, *MNRAS*, 407, 1403
- Dai, X., Anderson, M. E., and Bregman, J. N. 2012, *ApJ*, 755, 107
- de Vaucouleurs, G., de Vaucouleurs, A., Corwin, H. G., Jr., et al. 1991, *Third Reference Catalogue of Bright Galaxies*, Vol. 13 (Berlin: Springer)
- Fukugita, M., & Peebles, P. J. E. 2006, *ApJ*, 639, 590
- Giovanelli, R., Haynes, M. P., Rubin, V. C., & Ford, W. K., Jr. 1986, *ApJ*, 301, L7
- Guedes, J., Callegari, S., Madau, P., & Mayer, L. 2011, *ApJ*, 742, 76
- Kaufmann, T., Bullock, J. S., Maller, A. H., Fang, T., & Wadsley, J. 2009, *MNRAS*, 396, 191
- Kuntz, K. D., & Snowden, S. L. 2008, *A&A*, 478, 575
- Leccardi, A., & Molendi, S. 2008, *A&A*, 486, 359
- Maller, A. H., & Bullock, J. S. 2004, *MNRAS*, 355, 694
- Paturel, G., Theureau, G., Bottinelli, L., et al. 2003, *A&A*, 412, 57
- Revnivtsev, M., Churazov, E., Sazonov, S., Forman, W., & Jones, C. 2008, *A&A*, 490, 37
- Smith, R. K., Brickhouse, N. S., Liedahl, D. A., and Raymond, J. C. 2001, *ApJ*, 556, 91L
- Snowden, S. L., Egger, R., Freyberg, M. J., et al. 1997, *ApJ*, 485, 125
- Sutherland, R. S., & Dopita, M. A. 1993, *ApJS*, 88, 253

CHAPTER 5

Extended Hot Halos Around Isolated Galaxies Observed in the *ROSAT* All-Sky Survey

Note: This chapter is a reproduction of a paper of the same title, published in the *Astrophysical Journal*. The reference is Anderson, M. E., Bregman, J. N., and Dai, X. 2013, *ApJ*, 762, 106-127. The paper is copyright 2013, the American Astronomical Society, reproduced here under the non-exclusive right of republication granted by the AAS to the author(s) of the paper.

5.1 Abstract

We place general constraints on the luminosity and mass of hot X-ray emitting gas residing in extended “hot halos” around nearby massive galaxies. We examine stacked images of 2165 galaxies from the 2MASS Isolated Galaxy Catalog (2MIG), as well as subsets of this sample based on galaxy morphology and K-band luminosity. We detect X-ray emission at high confidence (ranging up to nearly 10σ) for each subsample of galaxies. The average L_X within 50 kpc is 1.0 ± 0.1 (statistical) ± 0.2 (systematic) $\times 10^{40}$ erg s⁻¹, although the early-type galaxies are more than twice as luminous as the late-type galaxies. Using a spatial analysis, we also find evidence for extended emission around five out of seven subsamples (the full sample, the luminous galaxies, early-type galaxies, luminous late-type galaxies, and luminous early-type galaxies) at 92.7%, 99.3%, 89.3%, 98.7%, and 92.1% confidence, respectively. Several additional lines of evidence also support this conclusion and suggest that about 1/2 of the total emission is extended, and about 1/3 of the extended emission comes from hot gas. For the sample of luminous galaxies, which has the strongest evidence for extended

emission, the average hot gas mass is $4 \times 10^9 M_\odot$ within 50 kpc and the implied accretion rate is $0.4 M_\odot \text{ yr}^{-1}$.

5.2 Introduction

In addition to the central black hole, the stellar component (bulge and/or disk) and stellar halo, the dark matter halo, and interstellar gas, a fundamental component of an isolated massive galaxy is the circumgalactic medium. The CGM contains all the material outflowing from or accreting onto the galaxy, and may contain a quasi-static hot gaseous halo as well. As the medium through which baryons flow into and out of a galaxy, the CGM therefore encodes a wealth of information about the galaxy’s past and future. Understanding the CGM is critical for understanding the baryonic processes responsible for galaxy formation. A portion of the CGM gas has been observed extensively by studying UV absorption lines, at redshifts from 0 to ~ 2 , probing temperatures from 10^4 K to $10^{5.8}$ K (e.g. Tripp et al. 2000, Adelberger et al. 2003, Tumlinson et al. 2011, Rudie et al. 2012) . But a significant fraction is expected to lie at higher temperatures, in the X-ray regime (see next paragraph). This hot component of the CGM is the focus of this paper.

The hot gas is generally assumed to be associated with a quasi-static hot gaseous halo (White and Rees 1978), although outflows can also reach the required temperatures (Strickland et al. 2000), and hot filamentary accretion may be possible as well Sijacki et al. (2012). Present-day simulations predict that hot halos form as gas accretes onto a galaxy; if the dark matter halo is sufficiently massive ($\sim 10^{12} M_\odot$), some of the inflowing gas is heated by an accretion shock to approximately the virial temperature of the system (around a quarter of a keV for an L^* galaxy; Birnboim and Dekel 2003). Below the critical galaxy mass, the majority of the gas is thought to accrete in a “cold flow”.

The hot gaseous halo is thought to be responsible for a number of observed effects. It makes accretion onto a galaxy much slower, strangling star formation and potentially explaining the color-magnitude bimodality (Dekel and Birnboim 2006) as an artifact of the transition between the cold and hot modes of accretion. The

increasing prevalence of hot mode accretion at lower redshift also can help to explain the observed “down-sizing” in the star formation history (Bower et al. 2006, De Lucia et al. 2006). For the very largest galaxies, with inferred dark matter halos as massive as small groups, there is some conceptual overlap between the hot halo and the intergroup medium (as in the case of the so-called “fossil groups”, e.g. Mulchaey and Zabludoff 1999).

Unfortunately, there is scant observational evidence even for the existence of extended gaseous halos caused by hot mode accretion onto large galaxies. This is especially true for late-type galaxies. Searches for hot halo emission around late-type galaxies were first conducted with the Einstein Observatory (e.g. Bregman and Glassgold 1982, McCammon and Sanders 1984) and then later with *ROSAT* (e.g. Benson et al. 2000); these studies were only able to place upper limits on the extended diffuse emission. With the newer generation of X-ray telescopes, diffuse emission has been detected around late-type galaxies (e.g. Strickland et al. 2004, Li et al. 2006, Tüllmann et al. 2006, Owen and Warwick 2009, Yamasaki et al. 2009). This emission is generally confined to the innermost 10-20 kpc and is seen perpendicular to the disk, and widely assumed to stem from internal processes such as galactic fountains. Hot gas is also observed around the Milky Way (e.g. Snowden et al. 1998), but its radial extent is unclear (Anderson and Bregman 2010, Gupta et al. 2012, Miller and Bregman 2012). So far, the only spiral galaxies with strong evidence for genuinely extended hot gaseous halos are NGC 1961 and UGC 12591, two of the largest spiral galaxies in the Universe (Anderson and Bregman 2011; Dai et al. 2012).

The situation around early-type galaxies is rather different, since the Einstein Observatory was able to detect hot halos around massive ellipticals (Foreman et al. 1979; Foreman, Jones, and Tucker 1985). These hot halos vary by more than an order of magnitude in their X-ray luminosity for a fixed stellar luminosity, and studying this scaling relation continues to be an important area of research today (O’Sullivan et al. 2001, Mulchaey and Jeltama 2010, Bogdán et al. 2012). However, the hot halo metallicities are generally super-solar (Mathews and Brighenti 2003) and so the gas is generally assumed to have an internal origin (i.e., AGB stellar winds shock-heated to

the virial temperature) instead of revealing hot mode accretion, although it is possible that the high-metallicity gas is screening a lower-metallicity component (Crain et al. 2010).

There are also significant theoretical uncertainties about the amount of mass contained in hot gaseous halos. Analytic calculations (White and Frenk 1991; Fukugita and Peebles 2006) typically assume that the hot halo contains enough baryonic mass to bring the baryon fraction within the virial radius up to the cosmic mean of 0.17 (Dunkley et al. 2009). In simulations, the baryon fraction of a galaxy can vary substantially, depending on the initial conditions and the feedback prescription. To illustrate this, we can take the example of the Milky Way. The stars and ISM gas of the Milky Way only contribute about 20-30% of the baryons required to bring our Galaxy’s baryon fraction to the cosmic mean, and these quantities can be reproduced in simulations. But different simulations predict very different values for the total baryon fraction of the Galaxy, and these differences are nearly entirely due to the mass of the hot halo. As examples, Guedes et al. (2011) ran a self-consistent simulation of a Milky Way-like galaxy and found a total baryon fraction of 70% of the cosmic value, while Moster et al. (2011) ran simulations of Milky Way-like galaxies initialized to 50% and 100% of the cosmic value. Finally, Faucher-Giguère et al. (2011) compared various feedback prescriptions and were able to produce L^* galaxies with baryon fractions ranging from 30% (essentially no hot halo) to 100% of the cosmic value. Modern simulations of galaxy formation use baryon fractions all over this range, such that the baryon fraction of galaxies is essentially a free parameter.

Unfortunately, observations have not been sufficient to help constrain the mass of hot halos around L^* galaxies. We collected various observational constraints in Anderson and Bregman (2010) which collectively suggest that hot halos do not contain more than 10–20% of the “missing” baryons from galaxies. We also detected extended X-ray emission around two giant spiral galaxies, NGC 1961 (Anderson and Bregman 2011) and UGC 12591 (Dai et al. 2012), and interpreted the emission as hot halo emission; these galaxies only seem to contain about 10-20% of the missing baryons in their hot halos as well. On the other hand, Humphrey et al. (2011) detected a hot

halo around the isolated L^* elliptical NGC 720, and that hot halo seems to bring the baryon fraction of the galaxy up to the cosmic mean. The recent observation of NGC 1521 Humphrey et al. (2012) seems to find a similar hot halo around this isolated elliptical as well. The discrepancy between the inferred (low) baryon fractions in our giant spirals and the (high) baryon fractions in the L^* ellipticals of Humphrey et al. merits further study, which we will undertake in subsequent work.

With only a few individual cases, and the aforementioned discrepancy in the results, it has not yet been possible to draw broader conclusions about the mass contained within hot halos. In this paper, we examine the aggregate properties of thousands of galaxies, by stacking images from the *ROSAT* All-Sky Survey (RASS). We are therefore able to provide much more general constraints on hot halo masses around L^* galaxies.

5.3 Sample

In choosing a sample of galaxies to stack, our highest priority was selecting very isolated galaxies. Galaxies in group or cluster environments are surrounded by hot, X-ray emitting gas in the form of the intragroup/intracluster medium, and it is far too difficult to distinguish a galactic hot halo from an intragroup medium using *ROSAT*. Moreover, there could be complex interactions between this medium and the galactic hot halo which would complicate interpretation of the results. We therefore chose the 2MASS Isolated Galaxy Catalog (2MIG, Karachentseva et al. 2010), which contains 3227 nearby galaxies selected using the following criteria: 1. detected in the 2MASS Extended Sources Catalog (XSC). 2. K_S magnitude < 12 . 3. angular diameter $a > 30''$. 4. the galaxy must be “isolated” from any other 2MASS XSC companions with $K_S < 14.5$.

We took the 3227 2MIG galaxies and narrowed the sample down to galaxies with measured radial velocities, and required the radial velocity to be at least 500 km s^{-1} . This left 2496 galaxies in our total sample. Using their de Vaucouleurs classifications, we divided these galaxies into early-type (Elliptical and S0) and late-type (Sa-Sd, and irregulars), with 756 in the former and 1740 in the latter (including 7 irregulars).




Figure 5.1 Histograms of (left) K_S absolute magnitudes and (right) distances for our sample of galaxies in the 2MASS Isolated Galaxies catalog (2MIG, Karachentseva et al. 2010). We differentiate between the subsample of galaxies with $M_{K_S} < -24.1$ (indicated in the bolder colors) and the fainter subsample (indicated in the more muted colors). This characteristic magnitude is also denoted on the plot on the left with a black dashed line. We also place vertical lines at $d = 20$ Mpc and $d = 250$ Mpc; the 67 galaxies that fall outside this region are excluded from our stacking analysis. In total, we examine 2496 galaxies, divided into subsamples based on luminosity and morphological type, and after rejecting galaxies based on distance or point source contamination we eventually stack 2165 galaxies.

We also divided the galaxies at L^* , which we took to be at $M^*_{K_S} = -24.1$. This is the characteristic magnitude of the K_S galaxy luminosity function (Bell et al. 2003; Kochanek et al. 2001), and roughly corresponds to the magnitude at which hot mode accretion is expected to dominate over cold mode accretion. To be precise, samples of nearby red and blue galaxies can have slightly different values for M^* , but we ignored that small difference for this analysis. The luminous subsample has 1016 galaxies total, with 393 ellipticals and 623 spirals. The faint subsample has 1480 galaxies total, with 363 ellipticals and 1117 spirals. In Figure 5.1, we present histograms of our sample as functions of distance and absolute magnitude.

5.4 Constructing Stacked Images

The RASS offers several useful features for this sort of analysis. The most important advantage is that the images have been carefully flattened (Voges et al. 1999), which

reduces the effect of spatial variations in the background which could otherwise be confused with diffuse emission. The effect of the flat fielding is visible in the flatness of the background in the surface brightness profiles which we will present in this paper. Additionally, the RASS has nearly full sky coverage, so we have a large number of objects available for our sample, and can therefore measure down to extremely low surface brightness. The major disadvantage of using the RASS is its poor angular resolution (for which we compensate by constructing an empirical PSF, see section 5.4).

To stack the galaxies in our sample, we used a modified version of the stacking code written by Dai, Kochanek, and Morgan (2007) to stack RASS images of nearby galaxy clusters. We add all the photons in the 0.5-2.0 keV range within 500 projected kpc of each galaxy (each galaxy has a known redshift listed in Karachentseva et al. 2010). Since each galaxy is at a different distance, each image therefore covers a different angular size on the sky, and we re-scale each image to the same angular size while preserving the photon count within 500 kpc.

Since 500 kpc covers a larger angular size for nearby galaxies, this naturally leads to the nearby galaxies contributing more background photons per image than the more distant galaxies. Most of the galaxies in our sample lie around 100 Mpc, but the handful of images corresponding to much nearer galaxies can dominate the total stacked image. We therefore imposed a distance cut, excluding any galaxies with an angular diameter distance of less than 20 Mpc. This removes 67 galaxies from the sample. We also imposed a cut on the distant end ($d > 250$ Mpc) because the 20 galaxies at larger distances were contributing so few photons to the stack that we were adding noise.

We also construct a stacked image of the RASS map of exposure time for each observation, again extending out to 500 projected kpc for each galaxy. We weight each individual exposure map by the quantity $(100 \text{ Mpc} / d)^2$ in order to account for the different angular sizes. This exposure map is necessary because some galaxies lie near the edge of a RASS $6.4^\circ \times 6.4^\circ$ frame, and so we are missing some counts within 500 projected kpc on one side of these galaxies. This produces a sort of

“effective vignetting”, which we can model with knowledge of the exposure map. The total weighted exposure time is also useful in later analysis for converting the total measured counts into flux and luminosity units (section 5.8.1).

We also cross-match each frame with the *ROSAT* bright source catalog (BSC, Voges et al. 1999). We exclude any point sources in our frames that are listed in the BSC. After masking regions with known point sources, we randomly populate the masked regions with artificial photons matching the photon density of nearby regions. We also visually inspected each individual frame, and removed frames with extremely bright point sources that were not masked by our software. In total, there were 81 such frames. There were also 183 galaxies that were more than $190'$ from the center of the nearest RASS frame, and therefore rejected from the stacking algorithm (the frames are squares with sides of length $384'$). We create a stacked image of all remaining 2165 galaxies, as well as images of four subsamples. Two subsamples are divided by morphology, with 659 early-type galaxies and 1506 late-type galaxies. The other two subsamples are divided by luminosity, with a cutoff at $M_{K_s} = -24.1$, such that there are 911 luminous galaxies and 1254 faint galaxies. The resulting stacked images are shown in Figure 5.2. We binned the images into “pixels” of 5 projected kpc ($10.3''$) on a side, and computed the average surface brightness in 1-pixel bins, displayed in Figure 5.3.

These images generally show fuzzy emission in the central region, but we need to consider the point spread function for the *ROSAT* PSPC to determine to what extent the emission is more diffuse than a point source. The off-axis PSPC PSF is fairly well-understood (with an analytic approximation available¹), but our image is a composite of a variety of images at different off-axis angles. Moreover, the all-sky survey was conducted in scanning mode, so each source has been observed at a variety of (unknown) off-axis angles. This makes it impractical to construct an analytic composite PSF, so we chose instead to construct an empirical composite PSF.

We generated a sample of point sources from the *ROSAT* Faint Source Catalog

¹http://heasarc.gsfc.nasa.gov/docs/journal/ROSAT_off-axis_psf4.html




Figure 5.2 Stacked images of our samples of 2MIG galaxies. Going clockwise from upper left, the samples are: early-type galaxies, luminous galaxies, faint galaxies, random positions on the sky (i.e., a null sample), all 2MIG galaxies, and late-type galaxies. All images have been smoothed with a 3-pixel Gaussian kernel. Note the different colorbars used in the three rows of images. Emission is clearly visible in the center of all the samples of 2MIG galaxies, and no emission is visible in the stack of random positions on the sky.




Figure 5.3 These plots contain the same information as the images in Figure 5.2, expressed as surface brightness profiles instead of images. The shading shows the number of counts per 5-kpc pixel, as a function of radius. The shading increases with radius because there are more pixels at large radii. The yellow line shows the mean surface brightness. The red line shows the mean surface brightness after dividing the mean by the average effective exposure at that radius. The difference between the red and yellow lines represents the effective vignetting described in section 5.4.

(FSC) and the BSC designed to match the observed flux distribution of the galaxies in our sample. We assumed the galaxies in a given sample all have similar X-ray luminosities, and so the flux distribution is determined by the distribution of distances to the stacked galaxies. We created point source distributions for three sets of galaxies: the full sample, the luminous sample, and the faint sample. For each sample, we also examined the galaxies between 30 and 150 Mpc (a range which contains most of the galaxies in each case, but only produces a factor of 25 range in flux instead of 156 from the full 20-250 Mpc sample). For each galaxy in each sample, we picked the ten point sources in the FSC and/or the BSC with the fluxes that were closest to this galaxy's position in the distance-squared distribution. We assigned these point sources the redshift of the corresponding galaxy (in order to convert the stacked PSF into the same physical units we use for the real analysis), and stacked all the point sources. After rejecting a handful of outliers with other bright point sources in the frame, we have empirical PSFs corresponding to the stacks of the full 2MIG sample, the faint galaxies, and the luminous galaxies, and containing 15440, 10564, and 4833 stacked point sources, respectively. In Figure 5.4, we present the radial surface brightness profiles for our fiducial PSFs. Note that our empirical PSF has a similar angular size to the PSF measured by Dai et al. (2010).

This method is subject to a Malmquist bias, because within each subsample the more distant galaxies will be on average slightly more luminous than the nearer galaxies. This is a higher-order correction which we do not apply to our data because we do not know the intrinsic X-ray luminosities of our target galaxies (only their redshifts and K-band luminosities). Instead, we construct alternate empirical PSFs by stacking 25,000 point sources at random from the FSC, and assigning them redshifts of 2MIG galaxies at random. This is essentially an overcorrection for the Malmquist bias because we are effectively assuming here that every galaxy has the same X-ray flux and therefore we are over-weighting the more distant galaxies. Still, the resulting empirical PSFs are fairly similar to the fiducial PSFs; they have slightly more power in the 200-300 kpc range, and less power at larger radii, but the shapes of the cores are very similar. We ran the entire subsequent analysis with these alternate PSFs

and the results were nearly identical.

We also constructed another set of empirical PSFs by stacking 1946 K-type stars from the Michigan Spectral Catalogue (Houk et al. 1975). We used the same technique for stacking these images as for the FSC sources. The MSC is potentially less likely to contain extended sources, however, compared to the FSC, so it provides a valuable cross-check of our estimated PSF. On the other hand, the MSC sources are generally fainter than the FSC, and we have fewer than 10% as many sources, so the PSF constructed from the MSC is not sampled nearly as fully. Still, these PSGs were very similar to the FSC PSFs above, further demonstrating the robustness of our determination of the PSF.

We also examined the effect of varying the distribution of redshifts assigned to the FSC sources. We will discuss this in section 5.4, but it turns out that using the wrong redshift distribution introduces an error of up to 30% to the total flux. To avoid this error, we use the PSFs constructed with the appropriate redshift distributions in the subsequent analysis for these two samples.

5.5 Analysis

5.5.1 Parameterizing the surface brightness profile

The primary analysis in this paper is a Markov-Chain Monte Carlo fit to the various surface brightness profiles we constructed in the above section. An “MCMC” analysis efficiently samples likelihoods over parameter spaces of moderate dimensionality, yielding posterior probability distribution functions (pdfs) for each parameter of interest. It also easily allows us to marginalize over uninteresting parameters, and to compute pdfs for quantities derived from the surface brightness profiles. As it is a Bayesian technique, it requires specification of prior probability distributions for each parameter. In this analysis, we use the Python package `pymc` (Patil and Fonnesbeck 2010) to perform the MCMC fit, which uses a standard Metropolis-Hastings sampling algorithm.

We begin by motivating the parameterization of the surface brightness profile. As




Figure 5.4 Mean surface brightness profiles of the primary stacked PSFs, as described in section 5.4, without the “effective vignetting” correction. The thick solid line represents the surface brightness profile computed by assigning each point source a redshift at random from the full distribution of 2MIG galaxies. The thin solid line is the PSF that results if redshifts are assigned at random from only the luminous 2MIG galaxies, and the dashed line is the PSF that results if redshifts are assigned at random from only the faint 2MIG galaxies. Because the luminous galaxies are on average more distant than the full sample (and the faint galaxies on average are less distant), the aperture corresponding to 500 projected kpc includes more (fewer) photons, which affects the normalization of the PSF. The shaded regions around each surface brightness profile show the full range of MCMC fits to the data, using the model defined by eq. 5.4.

Figure 5.3 shows, at large projected radii from the stacked galaxies (if we account for the “effective vignetting” discussed in section 5.4), the profile looks flat; we therefore assume the background is uniform across the detector, and define the parameter A_c as the constant background in units of counts/pixel. In fact, the remarkable flatness of the background is one of the primary advantages of using the RASS for this analysis (in particular, note the flatness of the random profile in Figure 5.3). We assume that the extended emission from the galaxies can be parameterized with a β -model, $b(r)$, defined as

$$b(r) = \left[1 + \left(\frac{r}{r_0} \right)^2 \right]^{0.5-3\beta} \quad (5.1)$$

In this model, r_0 is the core radius, β is the slope, and we pull the normalization A_β out in front of the convolution. There are many advantages of using a β -model for the hot gas component. If we assume the gas is isothermal and has a constant metallicity (which we do assume for this analysis), then its density profile is given by

$$n(r) = n_0 \left[1 + \left(\frac{r}{r_0} \right)^2 \right]^{-1.5\beta} \quad (5.2)$$

which we can integrate to find the hot gas mass within a given radius. Empirically, these models fit fairly well to the hot gas around elliptical galaxies (Forman, Jones, and Tucker 1985) as well as the hot gas in galaxy groups and clusters (Sarazin 1986), and we used them to fit the surface brightness profile around the isolated giant spiral galaxies NGC 1961 (Anderson and Bregman 2011) and UGC 12591 (Dai et al. 2012) as well.

We also include an additional component intended to mimic the signal from the center of the galaxy and any active nucleus it may contain. We parametrize this with a δ -function, $\delta(r)$, which is approximated as A_δ in the first pixel and 0 in the other pixels. Any emission which is restricted to the nucleus, the bulge, or in the inner 5 kpc of the stellar component is intended to be captured in this component. Then we convolve $b(r)$ and $\delta(r)$ with a PSF (using the symbol $*$ to denote a discrete convolution) to produce the full signal from the source.

Finally, as mentioned above (and as can be seen in Figure 5.3), the surface brightness profile of the data declines slightly at large radius. This is because some of the images we stack do not extend to 500 projected kpc on each side of the galaxy, meaning that the final image is missing some photons at large radii. We account for this by measuring the surface brightness profile of the weighted, stacked exposure map for each sample, and normalizing by the central (maximum) exposure time to produce an effective exposure profile for each sample, $f_{\text{vig}}(r)$. At 500 projected kpc, f_{vig} is still over 90%, but this 10% correction can still be important for measuring the background accurately. The final model is therefore:

$$\begin{aligned}
S(r) &= A_\delta \times [\delta(r) * \text{psf}(r)] \times f_{\text{vig}}(r) \\
&+ A_\beta \times [b(r) * \text{psf}(r)] \times f_{\text{vig}}(r) \\
&+ A_c \times f_{\text{vig}}(r)
\end{aligned} \tag{5.3}$$

5.5.2 Computing the PSF

To estimate PSFs from the surface brightness profiles in Figure 5.4, we need to subtract the uniform background component. We estimate the level of this background by parameterizing the PSF analytically and running an MCMC fit to the surface brightness profile to compute the posterior probability distribution for the level of the background. Motivated by the analytic expression for the off-axis PSF in targeted PSPC observations: (Hasinger et al. 1993)

$$\begin{aligned}
\frac{\text{psf}(r)}{f_{\text{vig}}(r)} &= A_g \times \exp\left(\frac{-r^2}{2\sigma^2}\right) \\
&+ (r > r_c) \times A_p \left(\frac{r}{r_0}\right)^{-\gamma} \\
&+ (r \leq r_c) \times \frac{A_L}{1 + (2r/r_L)^2} \\
&+ A_c
\end{aligned} \tag{5.4}$$

Table 5.1. PSF fitting parameters

PSF name	A_p	A_g	A_c	γ	σ	r_c	r_L
all	0.46 ± 0.03	4366^{+78}_{-73}	239.8 ± 0.1	2.57 ± 0.03	0.80 ± 0.01	11.2 ± 0.2	5.5 ± 0.1
faint	0.30 ± 0.02	3658^{+76}_{-74}	193.7 ± 0.1	2.63 ± 0.04	0.76 ± 0.01	10.7 ± 0.3	5.2 ± 0.1
luminous	$0.20^{+0.02}_{-0.03}$	877^{+31}_{-27}	44.8 ± 0.1	$2.32^{+0.08}_{-0.06}$	0.93 ± 0.02	$12.3^{+1.5}_{-0.8}$	6.8 ± 0.2

Note. — Median values and 68% confidence regions for the MCMC fits to the three psfs examined in section 5.5.2. Note that all of these parameters are computed in units of 5-kpc pixels, so to convert r_c into kpc, for example, multiply it by 5.

So the PSF is defined by the sum of a Gaussian component, a Lorentzian component that becomes a powerlaw at large radii (A_L is fixed so that the Lorentzian and the powerlaw have the same value at $r = r_c$), and a uniform background component. We use uniform priors on all parameters except A_c , which uses Gaussian priors centered around the mean value of the surface brightness profile at large radii (and standard deviation of 0.1). The powerlaw slope γ is also required to be less than -2, in order to guarantee that the PSF converges to zero at infinitely large radii. The variable r_0 is just a normalization constant; we take $r_0 = 100$ pixels = 500 projected kpc. Note that we have moved $f_{\text{vig}}(r)$ to the left just to make the equation easier to read; $f_{\text{vig}}(r)$ is always applied to the model, not to the data, for the MCMC analysis.

We multiply both the data and the model by the number of pixels in each 1-pixel radial bin, since we are actually measuring the total number of photons and not the average per radial bin. We then compute uncertainties in each bin according to Poisson statistics. We ran 3×10^6 iterations, discarded the first 5×10^4 , and selected only every 1000th remaining iteration in order to ensure each element is independent of each subsequent element. The results of the MCMC fitting are shown above, graphically in Figure 5.4, and the probability distribution functions (pdfs) of the fit parameters are listed in Table 5.1.

Because the stacked FSC and MSC data (the latter in particular) are so noisy, we use the best-fit profiles from the MCMC analysis to generate the PSFs. We divide this best-fit profile by $f_{\text{vig}}(r)$ to correct for the artificial vignetting, and we subtract

the median value of A_c to yield an empirical PSF. The mean value of the PSF surface brightness at large radii turns out to be slightly above A_c . The discrepancy is about 4%, suggesting that this fraction of the signal at large radii comes from the wings of the PSF and not from the uniform background.

In Figure 5.5a, we show the resulting PSFs, after subtracting off the background component and normalizing the result so the integrated sum is unity. We also show average power at each radius, multiplying the PSF at each radius by the number of pixels it subtends, in Figure 5.5b (the discrete pixel size causes the apparent noise). The two sets of point spread functions, although computed from different objects and using very different sample sizes, have similar shapes. The cores seem weaker in the PSFs constructed by assigning redshifts from the sample of luminous galaxies, compared to the other PSFs (Fig. 5.5a, inset), and the PSFs are less well constrained at large radii where there is very little signal, but there are not significant systematic differences otherwise.

Considering how much smoother the underlying data are, we adopt the FSC PSFs for subsequent analysis. The fiducial FSC PSF is used to analyze the full sample and the early-type and late-type subsamples. The FSC PSF constructed with the redshift distributions of the faint and the luminous galaxies are used to analyze these two subsamples, respectively.

5.6 Simulations

Now that we have a functional form for the PSF, we are able to begin analyzing the stacked 2MIG images. However, to verify the reliability of our stacking technique and to aid in the interpretation of the results, we first apply our analysis to simulated stacked images. We constructed the simulated images based on the model in equation 5.3. For each image, we specified all of the free parameters in equation 5.3 (i.e. A_δ , A_β , A_c , r_0 , β , the functional form of the PSF, and the exposure map from the stack of real 2MIG galaxies corresponding to the simulated image). For each of the three components of the model (the point source, the extended component, and the uniform background - note that the first two components are each convolved with the specified




Figure 5.5 Profiles showing the normalized strength of the PSF as a function of projected radius. In (a), we show the PSF value per 5-kpc pixel, normalized so the total across the image adds to unity. In (b) we show the PSF value per 5-kpc radial bin (i.e., the PSF weighted by the azimuthal area it subtends), again normalized so the total out to 500 kpc adds to unity. The inset in each plot shows the same data on a linear scale. In each plot, the thick solid line corresponds to the PSF constructed from the redshift distribution of all 2MIG galaxies, the thin solid line is the PSF constructed from redshifts of only the luminous galaxies, and the dashed line is the PSF constructed from redshifts of only the faint galaxies.

PSF as well), we computed the expected number of counts in each $5 \text{ kpc} \times 5 \text{ kpc}$ “pixel”, and then placed in that pixel a number of counts equal to a random Poisson deviate with the given expected value. We added the three components together and multiplied by the normalized exposure map to produce a full simulated image with the same “effective vignetting” as the real image.

In the rest of this section, we present and discuss a number of simulations. Note that instead of A_δ and A_β , we found it more useful in analyzing the simulations to specify N_δ and N_β - the total number of counts in the point source and extended components, respectively. Recovering N_δ and N_β from a simulated image is our first priority in this analysis. We define the total number of counts within a circle of radius 500 projected kpc, or 100 pixels. This means that, in theory, we could also estimate the number of counts N_c , which should be within integrated pixel-by-pixel Poisson errors of $10^4 \pi A_c$. However, in every simulation we run, N_c is much larger than either N_δ or N_β , and we find it conceptually easier to work with A_c . Table 5.2 lists the basic properties of our simulated images. The only differences between the various sets of simulations are the quantities listed in Table 5.2.

We analyzed these simulated images using the same pipeline as the real stacked images. This means for each simulation we specify the appropriate FSC PSF and the stacked exposure map corresponding to the image (see Table 5.2). The choices of priors for the model parameters are described below, in section 5.8. The only prior that uses information about the image is A_c , to which we assign a Gaussian prior centered at the the average number of counts per pixel at large radii, after attempting to correct for the artificial vignetting. We tested the choice of this prior, and it converges to the same value if we move the mean upwards or downwards slightly, or if we widen the width of the Gaussian. After running the MCMC analysis, we examined the resulting posterior probability distributions for β and r_0 , A_β , and A_δ . From the first three of these distributions, we computed the posterior probability distribution for N_β , and from the fourth distribution we computed N_δ .

Table 5.2. Properties of simulations

name	$N_\delta + N_\beta$	A_c	PSF	exposure map
All galaxies	1600	22.6	all	all
Late-type galaxies	800	16.6	all	late
Early-type galaxies	800	6.0	all	early
Faint galaxies	600	18.4	faint	faint
Luminous galaxies	1000	4.2	luminous	luminous
Extra counts	10000	22.6	all	all

Note. — Basic properties of the sets of simulated images we examine. The simulated images are designed to mimic the properties of the real stacked images of each subsample of 2MIG galaxies. The $N_\delta + N_\beta$ column shows the number of source counts spread across the image, and the numbers are chosen to approximate the estimated number of source counts in the corresponding real image. The A_c column shows the average number of background counts per pixel, and again these numbers are chosen to approximate the estimated background in the corresponding real images. The PSF column shows the redshift distribution of 2MIG galaxies used to produce the PSF for the image; these are the same PSFs used to simulate the real images. The exposure map column shows which real stacked exposure map we use to simulate the “effective vignetting” in the simulated images. Finally, note that “Extra Counts” simulations have more photons than any real image and are designed to reduce the Poisson uncertainties and therefore help diagnose which errors in the simulations stem from photon noise and which errors are systematic.

5.6.1 Recovering total counts and fit parameters

We first examine the simplest case, where we set either N_β or N_δ to zero. In Figure 5.6 we compare the pdfs (we show the medians and 68% confidence intervals) for $N \equiv N_\beta + N_\delta$ to the true value for the simulations of each galaxy sample. We integrate the models over the entire image for this plot (out to 500 kpc); as we will show, it is easier to recover the emission just within 50 kpc or so, where it is significantly above the background. Still, across the whole image we are able to recover $N_\beta + N_\delta$ adequately for entirely point source emission and for $\beta > 0.5$ (i.e. extended emission with a fairly steeply declining slope). For flatter extended emission, the uncertainty grows, but in general the MCMC analysis fails to recover all of the emission, and the discrepancy increases as the slope becomes flatter. What seems to be happening is that the flatter profiles are becoming indistinguishable from Poisson noise at larger radii, and so the MCMC analysis is attributing the missing counts to a slightly higher background instead of to the extended component.

To help verify this, we checked if the MCMC analysis was able to recover the value of β correctly (Figure 5.7). For the simulation with extra counts, β is recovered approximately, but the other stacks almost always recover a value of β around 0.5 – 0.6, regardless of the true value. Unfortunately, this means we cannot use the recovered value of β as a reliable guide to the true value of β . This also suggests that measurements which rely on extrapolating with β out to large radii are unreliable if we suspect $\beta \lesssim 0.5$.

5.6.2 Recovering counts within 50 kpc

While it is difficult to distinguish a flattened profile from the background at very large radii, this emission is stronger at smaller radii, and therefore should be easier to recover. In Figure 5.8 we show the pdf for N if we only integrate the fit parameters out to 50 kpc instead of out to 500 kpc. We compare the recovered counts within this radius to the expected number of counts within this radius, which can be computed easily since we know β for our simulations (Table 5.2). Note that the number of counts within 50 kpc varies significantly with β , ranging from 11% of the counts




Figure 5.6 Probability distribution functions (pdfs) of the ratio of the measured value of $N \equiv N_\beta + N_\delta$ to the true value of N for simulations of each galaxy sample. We measure N using the MCMC analysis to fit the simulated image, and we integrate the resulting models out to the edge of the image (500 kpc). The true values of N are listed in Table 5.2. 68% confidence intervals around the medians of the pdfs for N are shown as well. The points labeled 'psf' are simulations where all the source counts are in the point source component. All the other points are simulations where all the source counts are in the extended component, the slope β of the extended component is indicated for each set of simulations. The points correspond to various simulations, as indicated in the legend. The points have been scattered by up to $\beta = 0.01$ for readability; each set of simulations actually uses the same values of β at the indicated values of β . For discussion of this plot, see the text (section 5.6.1).




Figure 5.7 Probability distribution functions (pdfs) for the measured value of the slope β of the extended component of the simulated images, compared to the true value of β for those images. The 68% confidence intervals are shown for each simulation as well. The dashed line denotes equality, i.e. where the measured value equals the true value. Points are colored/shaped as indicated in the legend. The points have been scattered horizontally by up to $\beta = 0.01$ for readability; each set of simulations actually uses the same values of β at the indicated values of β on the horizontal axis. This plot shows that, while with more counts it would be possible to recover β approximately, we are not able to recover β in simulations corresponding to our actual data.

within 500 kpc for $\beta = 0.35$ to 94% of the total counts for $\beta = 0.70$. As Figure 5.8 shows, we are able to recover the number of counts within 50 kpc for most cases. The uncertainties appear much larger for the flatter profiles, but this is mostly because the total number of counts within 50 kpc is so much smaller.

Additionally, there appears to be a slight trend of overestimating N within 50 kpc, for $\beta \approx 0.5$, in the images with all the emission in an extended component. This happens because some extended emission is mistakenly attributed to the point source component, which increases the inferred number of counts originating from radii within 50 kpc.

5.6.3 Identifying extended emission

We now examine how well the MCMC analysis can distinguish point source emission from extended emission. In other words, we have examined how well we can recover $N \equiv N_\beta + N_\delta$, and now we examine how well we can recover N_β and N_δ individually. To do this, we focus on the simulations of the stack of all 2MIG galaxies. In Figure 5.9 we show the recovered counts within 50 and 500 kpc, as well as the pdfs for N_β and N_δ . There are two separate but related failure modes. First, there is a constant – but fairly small – level of confusion in each simulation; the emission should be entirely in one component, but some emission is always attributed to the other component as well. Because all the point source emission is by definition within 5 kpc, the point source comprises a larger fraction of the true signal within 50 kpc compared to 500 kpc, so this failure mode hurts our ability to measure the correct fraction of extended emission within 50 kpc, and it has less of an effect at 500 kpc. Second, at $\beta \gtrsim 0.6$, the MCMC analysis is unable to distinguish extended emission from point source emission, so it mistakenly attributes most of the emission to the point source. This is because the PSF is wider than the profile at $\beta \gtrsim 0.5$, so the convolution of the two looks very similar to the PSF. This failure mode is still quite significant at 500 kpc, because the steeper profiles have very little emission at larger radii.

The failure modes identified above - failure to recover counts at large radii if β is too low, inability to distinguish point source emission from extended emission if




Figure 5.8 Probability distribution functions (pdfs) of the ratio of the measured value of $N \equiv N_\beta + N_\delta$ to the true value of N for simulations of each galaxy sample. We measure N using the MCMC analysis to fit the simulated image, and we integrate the resulting models out to a radius of 50 kpc. The true values of N range from 11% of the total listed in Table 5.2 to 94% of the total listed in Table 5.2, depending on the slope of the profile. 68% confidence intervals of the pdfs for N are shown as well. The points labeled 'psf' are simulations where all the source counts are in the point source component. All the other points are simulations where all the source counts are in the extended component, the slope β of the extended component is indicated for each set of simulations. The points correspond to various types of simulations, as indicated in the legend. The points have been scattered by up to $\beta = 0.01$ for readability; each set of simulations actually uses the same values of β at the indicated values of β . We are able to recover the number of counts within 50 kpc for most cases. The uncertainties appear much larger for the flatter profiles, but this is primarily because the total number of counts within 50 kpc is so much smaller.

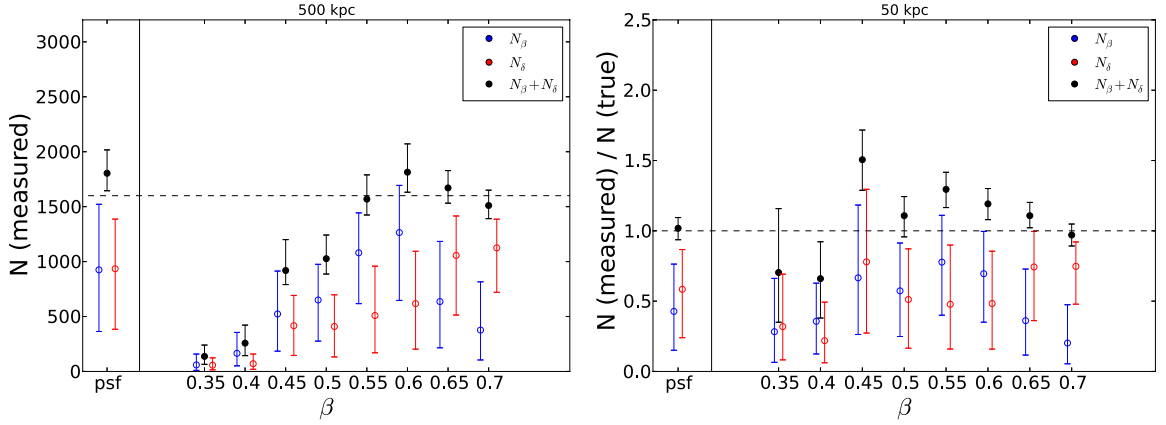


Figure 5.9 Probability distribution functions (pdfs) for (a) the values of N_β and N_δ integrated out to 500 kpc and (b) the ratios of the values of N_β and N_δ integrated out to 50 kpc compared to the true values (dashed lines). 68% confidence intervals on each measured value are shown as well. The blue points denote the pdf for the extended component N_β , the red points the pdf for the point source component N_δ , and the black points the pdf for the total source counts $N = N_\beta + N_\delta$. These pdfs are all measured for the “all” simulations. In the leftmost simulation, the true source is entirely composed of point source emission, and the other simulations have sources entirely composed of extended emission. Unfortunately, the MCMC analysis does not distinguish between the two types of emission reliably: for steeper profiles ($\beta \gtrsim 0.5$) the extended emission is mistaken for point source emission, and all simulated images have some emission mistakenly attributed to the wrong component as well. Later in the paper, we combine this spatial analysis with a likelihood ratio test and a spectral analysis to improve our constraints on the extended component.

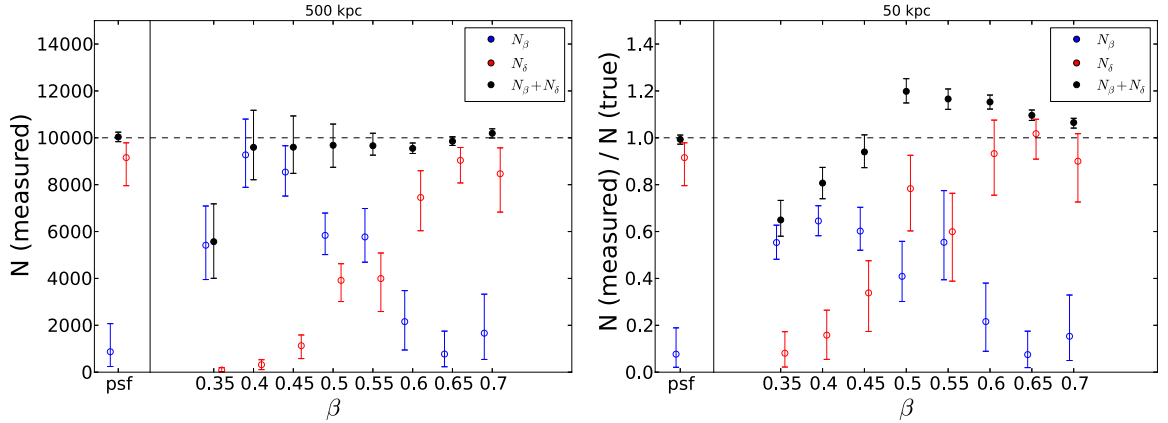


Figure 5.10 Same as Figure 5.9, but for the “extra counts” simulations. The MCMC analysis is still not able to distinguish between the two types of emission reliably for $\beta \gtrsim 0.5$, but the amount of misclassified emission is much smaller in the other profiles, compared to the “all” simulations (Figure 5.9). Later in the paper, we combine this spatial analysis with a likelihood ratio test and a spectral hardness ratio analysis to improve our constraints on the extended component.

β is too high, and some confusion between point source and extended emission at all values of β - could be due to insufficient photons or to systematic failures of the method. We attribute the first failure mode to insufficient photons (and note that this conclusion is supported by the fact that the “all” simulation, which has over twice as many source photons as the other simulations, better recovers the number of photons in the cases with low β), and the second failure mode to a fundamental systematic uncertainty: the poor angular resolution afforded by the PSF. To help support this claim, we now examine the “extra counts” simulations. These images have the same background and exposure map as the stack of all 2MIG galaxies, but we have 10000 source counts instead of 1600. The results are shown below, in Figure 5.10.

This Figure shows that the total number of counts ($N_\beta + N_\delta$) can be recovered to within 10% for $\beta \geq 0.4$ within 500 kpc. So failures to recover this parameter are due to insufficient photons to distinguish the surface brightness profile from Poisson variations in the background. However, these simulations still confuse extended emission and point source emission for $\beta \gtrsim 0.5$, so this seems to be a systematic failure of

the analysis: the PSF is broad enough that the steeply declining extended profiles do not look very different from a point source. This also helps to explain why N is only recovered to 30% accuracy within 50 kpc: the misattribution of extended emission to a point source component increases the inferred number of counts at relatively small radii, although the total number of counts across the full 500 kpc image is still recovered correctly. The same issue applies in Figure 5.9b, although the statistical uncertainties are larger so this systematic effect is less important.

Finally, the amount of confusion between point source and extended emission at lower values of β and in the case of pure point source emission is much lower than it was in Figure 5.9. This implies that this confusion is due to insufficient photons - the MCMC analysis does work in the cases where the extended emission looks different from a point source. Thus, these failure modes should not affect the stacking analysis of high S/N data such as galaxy clusters (e.g. Dai et al. 2010).

5.6.4 Effects of changing the PSF

While we believe the FSC PSF is the most realistic empirical PSF we have been able to produce, we examined the effect of varying the PSF as well. We explore effects of varying the PSF redshift distribution in Figure 5.11. Here, we use the PSF computed with three different redshift distributions to analyze the “all” simulated images (which were created using the PSF with the redshift distribution of all 2MIG sources). While the MCMC analysis slightly systematically overestimates the number of counts within 50 kpc for this simulation (see section 5.6.2), changing to the PSF with the redshift distribution of one of the subsamples of 2MIG galaxies (faint or luminous) also systematically changes the amount of recovered counts. Using the PSF constructed with the redshift distribution of faint galaxies reduces the total inferred number of counts by $\sim 10\%$, while the PSF constructed with the redshift distribution of luminous galaxies increases the total inferred number of counts by $\sim 30\%$. The luminous galaxy PSF is noticeably flatter than the other PSFs in Figure 5.5, so this makes sense. In our analysis of the 2MIG data, we match the redshift distributions of the point sources to the distributions of the 2MIG sources, in order




Figure 5.11 Comparison of pdfs of recovered counts within 50 kpc after variations in the redshift distribution/aperture size of point sources we stack to produce the empirical PSFs. We fit the “all” simulated images (created using the redshift distribution of all 2MIG galaxies) with empirical PSFs created using three different redshift distributions as noted in the legend. Using the faint redshift distribution introduces a bias of $\sim 10\%$, while the luminous redshift distribution introduces a bias of $\sim 30\%$ around the fits using the correct relation (the black points), which are themselves systematically high by $\sim 20\% - 30\%$.

to minimize these errors.

5.6.5 Including a combination of point source and extended components

Finally, we examine the effect of mixing an extended component and a point source component in the same image. We created a series of simulated images with half the source counts in the extended component and half in the point source component and analyzed them. Figure 5.12a shows the total fraction of recovered counts for each set of simulated images, and Figure 5.12b shows the allocation of the recovered counts into extended and point source components for the set of simulations designed to mimic the full stack of all galaxies. The analysis works very well; the results actually




Figure 5.12 Comparison of pdfs of recovered counts within 50 kpc to the true number of counts within 50 kpc for simulations where $N_\beta = N_\delta$, i.e. half the source counts are in the extended component and half the source counts are in the point source component. In (a) we present the pdfs of the recovered counts for all five sets of simulations; this Figure shows that the recovered counts are closer to the true value than in the simulations with all the source counts in one component (see Figure 5.8). In (b), we present the pdfs of the individual components of the “all” simulation, and again the recovered values for N_β and N_δ (both of which should be half the total) are closer to the true value than they are when all the source counts are in just one component (see Figure 5.9b).

look better than Figures 5.9 and 5.10 because the failure to recover very flat extended emission only applies to half of the total source counts, and the point source counts are recovered fully. For the luminous galaxies, the MCMC analysis systematically underestimates the source counts within 50 kpc by $\sim 30\%$; this seems to be due to misattribution of some counts to the extended component, since the number of counts within 500 kpc is recovered correctly. More generally, the division between the two components leads to significant uncertainties on the total counts within each, but in Figure 5.12b the total number of counts is recovered more accurately than in Figure 5.9.

5.6.6 Summary of simulations

The major findings from our simulations are:

1. The MCMC analysis should be able to measure the number of source counts within 50 kpc to an accuracy of 50% or better.

2. The MCMC analysis might be able to measure the number of source counts all the way out to 500 kpc, although if $\beta \lesssim 0.5$ it can underestimate the total number of source counts significantly. We will principally be concerned with the emission within 50 kpc for the subsequent analysis.

3. The MCMC analysis cannot reliably recover β or r_0 with the quality of data we are using.

4. The MCMC analysis has trouble distinguishing purely extended and purely point source emission, especially when focusing on the emission within 50 kpc. Most notably, for $\beta \gtrsim 0.5$, it cannot distinguish one component from the other, and attributes most of the emission to the point source component.

5. If the source is composed of a mixture of an extended component and a point source component, the MCMC analysis can recover the total number of source counts within 50 kpc to an accuracy of 30% or better.

6. Reasonable variations in the PSF due to the size of the aperture produce an error in the measured counts within 50 kpc up to 50%, which we attempt to avoid by choosing the correct redshift distributions when making the empirical PSFs. Reasonable variations in the sample of point sources used to create the empirical PSFs produce a negligible error in the total measured counts.

5.7 Statistical significance of detections

After running the simulations we have confidence in being able to recover the total amount of source emission, especially within a 50 kpc radius of the galaxy, but less confidence in identifying the fraction of emission in the extended component. An important question to consider before running the MCMC analysis is whether we have good reason to believe there *is* an extended component in our stacked images, and here we quantify the statistical significance of the detection of extended emission, as well as the statistical significance of the detection of any emission at all.

To quantify the statistical significance of our detections of *any* emission, we apply Poisson statistics on the counts within a radius of 50 projected kpc. We model the background as a uniform component multiplied by the normalized exposure map to

account for the “effective vignetting” (which is only a few percent within 50 projected kpc). We estimate the level of the uniform background using a bootstrap analysis of various circular annuli around the central region. For each background annulus, we compute the pixel-by-pixel likelihood of obtaining the observed counts within 50 projected kpc given the background annulus and vignetting model. We compare this likelihood to 10000 randomly selected regions of radius 50 projected kpc chosen across the stacked image (but outside the central region). We repeat this analysis for 100 variations in background annuli, yielding a total of 10^6 null likelihoods. The statistical significance of our source detection is the fraction of null likelihoods which are lower than the observed likelihood of the central emission. These values are listed in Table 5.3.

We can also approximate the significance of the detection as the conjugate of the probability of the measured number of counts within 50 projected kpc, given our estimate of the background count rate from the MCMC fitting. This method is simpler, but does not account for non-uniformity in the background, and therefore may overestimate the significance of the detections if the background has significant variations above Poisson noise. Therefore we only use this technique to approximate the significance of detections at confidences of greater than 0.999999 (the highest confidence we can measure from the bootstrap technique above). These values are listed in Table 5.3 for the stacks of all galaxies, early-type galaxies, luminous galaxies, and luminous early-type galaxies.

To quantify the statistical significance of our detections of extended emission, we employ a likelihood ratio test. We compare the likelihood of the data for a model with an extended component (eq. 5.3) to the likelihood of the data for a model with the extended component set equal to zero (eq 3 with $A_\beta = 0$). The latter model therefore only contains a uniform background and a point source convolved with the empirical PSF (both terms subject to effective vignetting as well). We use a truncated Newton algorithm² to find the maximum likelihood best-fit to the data for each of our two models. In the model with the nonzero extended component, we require r_0

²the tnc algorithm included in the Python package SciPy.optimize, v. 0.11

Table 5.3. Significance of detections

sample	$1 - p(\text{any emission } r < 50 \text{ kpc})$	$1 - p(\text{extended emission})$
All galaxies	4.4×10^{-16} (8.1σ)	0.073
Late-type galaxies	0.000446	0.347
Early-type galaxies	3.0×10^{-9} (5.9σ)	0.107
Faint galaxies	0.005312	0.528
Luminous galaxies	2.6×10^{-23} (9.9σ)	0.007
Luminous late-type galaxies	0.000025	0.013
Luminous early-type galaxies	3.2×10^{-15} (7.9σ)	0.079
Random locations	0.412603	-

Note. — Estimated statistical significance of our detections of: any emission above the background (column 2), and extended emission (column 3). Significances expressed in scientific notation are higher than 99.9999% and therefore estimated using a different, less precise, technique. See section 6 for details of how these significances were computed. The image of “random locations” can be seen in Figure 5.2; we stacked images of 2500 random locations in the sky to serve as a null hypothesis. Because emission is not detected, we did not compute the significance of extended emission for this image.

to be between 0.5 and 5.0 kpc and we require β to be between 0.35 and 0.70. We then compute the likelihood ratio for the best fits, defined as the difference between the log-likelihoods of the best fits for the two models. Since the model without an extended component is a subset of the model with the extended component, the likelihood never decreases when we allow for an extended component, so the likelihood ratio is always non-negative.

To interpret the likelihood ratios, we compare the likelihood ratios for the data to likelihood ratios computed for simulated data. For each stacked image, we generate 1000 simulated images using the parameters of the simulations we ran in section 5.6. These simulated images have all the source counts in the point source component, however. We find the maximum-likelihood fits to each image for both models, and compute the likelihood ratio for each simulated image. This gives us a set of 1000 likelihood ratios for images without an extended component, and we therefore can estimate the probability of measuring a given likelihood ratio for an image without any extended emission. The statistical significance of our detection of an extended component is then just the conjugate of this probability, i.e. the fraction of point source simulations with likelihood ratios lower than the value we measure for the data. These values are listed in Table 5.3.

We computed the significance of emission in the full stack and the four subsamples used throughout the analysis, but we also examine here a few other interesting subsamples. We stacked the 556 luminous late-type galaxies and the 355 luminous early-type galaxies separately. There are not enough photons in these images to get as reliable a measurement of the shape of the extended component for these stacks, so we do not emphasize them in much of the subsequent analysis, but the background is so low in these images that we can detect the emission at very high confidence, so we include these figures in Table 5.3. We also include the stack of random pointings in the sky shown in Figure 5.2, as a null hypothesis, which indeed is not detected at any significance.

Except for the faint galaxies, all of the samples are detected at $> 3\sigma$ confidence (and in some cases, the signal is much stronger). None of the extended emission is detected at $> 3\sigma$ confidence, though the luminous galaxies and the luminous late-type galaxies are close to 3σ . This latter result appears to be a fundamental limitation of the *ROSAT* data. After all, we only have approximately one source photon per luminous galaxy in our stacked images, and less than one source photon per faint galaxy. Most of our samples have extended emission detected above or near 90% confidence (the full stack, the early-type galaxies, the luminous galaxies, the luminous late-type galaxies, and the luminous early-type galaxies), while two of our samples (the late-type galaxies and the faint galaxies) have extended emission at much lower confidence. As we will show in section 5.8.3, these samples also have stronger evidence for hot gas in their spectral hardness ratios than the other two samples, and the preponderance of evidence will suggest (but not definitively) that these galaxies do indeed have extended gaseous emission. However, with *ROSAT*, the result is still only a strong suggestion of extended emission; a definitive detection will probably have to wait for the next generation X-ray all sky survey (i.e E-Rosita).

5.8 Results

Based on our simulations, we have confidence in being able to use MCMC analysis to extract the total number of source counts from our stacked 2MIG data. The simulations suggest that the MCMC analysis is less reliable in recovering the fraction of source counts in the extended component, although it is fairly reliable in the (likely) case of an equal mixture of point source emission and extended emission. We also have very strong (up to nearly 10σ) detections of soft X-ray emission from all of our stacked samples, and moderate confidence that there is extended emission around the stacks of all galaxies (92.7% significance), early-type galaxies (89.3% significance), luminous galaxies (99.3% significance), luminous late-type galaxies (98.7% significance), and luminous early-type galaxies (92.1% significance).

We now present the MCMC analysis on the real stacked images. Just as with the simulations, we specify the appropriate PSF and the stacked exposure map for each image, as well as prior distributions for the five model parameters in eq. 5.3. The prior for A_δ was uniformly distributed between 0 and 2×10^4 , and the prior for A_β was uniformly distributed between 0 and 2×10^{53} . A_c was normally distributed around the average pixel value at large radii, after attempting to correct for the “effective vignetting”, with a standard deviation of 0.1. The core radius r_0 was uniformly distributed between 0.5 and 5.0 kpc, and β was normally distributed around $\beta = 0.5$, with a standard deviation of 0.1. Note that our priors for β and r_0 reflect our prior knowledge of the shapes of hot gaseous halos around other galaxies, including the two giant spirals NGC 1961 and UGC 12591.

We carefully explored the effects of various choices for the prior distributions of the free parameters listed above. Varying the width of the priors for A_δ and A_β does not affect the result significantly, since the total amount of emission above the background is fairly well constrained by the data. Similarly, small variations in the width or the shape of the prior on A_c do not have much effect, since the background

³Note that A_δ translates directly into the number of point source counts in the image, while the number of extended counts depends on A_β , r_0 , and β . The adopted priors on A_β and A_δ correspond to roughly equal ranges in counts for each component




Figure 5.13 Probability distribution functions (pdfs) for the total number of source counts in the stacked 2MIG images, within a radius of (a) 50 kpc and (b) 500 kpc from the central galaxies. The four samples are shown in four different colors as indicated, and the median of each pdf is marked with a vertical dash at the top of each plot. The longer tails to the right in (b) reflect the possibility of extremely flattened profiles.

is also well constrained by the data. Varying r_0 has almost no effect on the luminosity and only a small effect on the mass, but the simulations showed that the data are not of sufficient quality to constrain r_0 at all so we leave the priors wide. Finally, varying β does have a significant effect on the posterior mass distribution; however, we think we understand β well, and nearly all existing observations of β around other galaxies place it between 0.4 and 0.6, with a median right around 0.5. Priors on β that conform to these constraints all yield similar results.

Figure 5.13 shows the pdf of the total number of source counts within physical radii of 50 kpc and 500 kpc from the central galaxies, for each of the four subsamples of the 2MIG galaxies. The number of counts within 50 kpc is fairly well constrained, but the number of counts within 500 kpc has a long tail extending to much larger values. This tail reflects the possibility of flattened profiles, as examined in our simulated images. $1-\sigma$ confidence intervals miss this long tail, but it appears plainly in the full marginalized probability distribution.

5.8.1 Average Luminosity

The next step is to convert the number of source counts in each image into an average luminosity per galaxy. This requires three steps. First, we must assume a spectral model for the soft X-ray emission, and use the spectral model to convert the number of counts into units of energy (erg). At the same time we can fold in the instrumental effective area as a function of energy to get energy per unit area (erg cm^{-2}). Then we divide by the exposure time, which we estimate from the stacked and weighted exposure map, to get a flux ($\text{erg cm}^{-2} \text{s}^{-1}$). Finally, we multiply by $4\pi d_L^2$ to produce a luminosity (erg s^{-1}), which is simple because the exposure map was weighted such that each galaxy's exposure time corresponds to flux from at an effective distance of 100 Mpc (see section 5.4).

There is also a small additional dependence on redshift, which changes the effective energy band of our observations by a factor of $(1+z)$ and introduces cosmological surface brightness dimming which goes as $(1+z)^4$. The maximum redshift in our sample is 0.06 and the mean is 0.02. The galaxies which dominate the observed flux have even lower redshifts than that, so these redshift effects are small and we do not attempt to correct for them.

For the assumed spectral model, we examine various linear combinations of an APEC model (Smith et al. 2001; used for the hot gas) and a powerlaw (for the X-ray binaries / possible AGN), as well as various amounts of absorption for both components. For the APEC model, we vary the metallicity between $0.2Z_\odot$ and $1.0Z_\odot$, and we explore three different gas temperatures. The $kT = 0.2$ keV model is approximately the virial temperature for hot gas with a truncated singular isothermal sphere density profile inhabiting an L^* galaxy's halo, and represents an approximately quasi-static hot gas halo. We also explored the effect of hotter gas (0.3 keV) and cooler gas (0.1 keV). For the X-ray binaries, we assumed a powerlaw with spectral slope of 1.56 (Irwin, Athey, and Bregman 2003), and allowed the slope to vary by ± 0.2 . For the AGN, we assumed a spectral slope of 2 ± 0.5 . Finally, for each model we vary the absorbing Hydrogen column between 10^{20} cm^{-2} and 10^{21} cm^{-2} , a range which encompasses 80% of our 2MIG galaxies based on the Dickey and Lockman (1990)

Galactic values.

We compute the conversion factors for each model using the HEASARC PIMMS tool⁴ (version 4.6). In addition to calculating the spectral models, this tool automatically folds in the effective area for the instrument, so the output is in units of erg cm^{-2} . We will estimate the conversion factor from counts to energy per area by computing the linear combination of these models. We attribute the AGN-like spectrum to the point source emission, and the combination of hot gas and X-ray binaries to the extended component. The point source being an AGN is plausible both based on the subsequent point source luminosities and a recent optical emission line analysis of isolated galaxies that found that nearly all of these galaxies harbor a weak AGN (Hernandez-Ibarra et al. 2012). Since the X-ray binary spectrum ($\Gamma = 1.56$) is within the range of possible AGN spectra we consider ($\Gamma = 1.5 - 2.5$), this also allows for the possibility of X-ray binaries in the central region of the galaxy. We can then estimate the appropriate linear combination of these components based on the fraction of emission in each spatial component, as measured by our MCMC analysis. This fraction is shown for each sample in Figure 5.14.

The median extended fractions increase with the statistical significance of the extended component, up to a maximum of $\sim 70\%$ for the luminous galaxies, though the uncertainties are very large and these images are all also consistent with an equal mixture of extended and point source emission (i.e. an extended fraction of 50%). In order to estimate the counts-to-energy ratio, we use the fraction of extended emission for the full sample instead of the subsamples, since they are all statistically indistinguishable from one another, but the full sample has the most photons and theoretically the most accurate constraints. This fraction is $60 \pm 30\%$. We assign the remaining fraction of the emission to the point source component, and assume it has the AGN-like spectrum described above. For the extended component, we explore models with 100% hot gas, 100% XRB, and 50% of each. We present the computed counts-to-energy ratios in Figures 5.15a and 5.15b.

We think the most plausible model is the leftmost model in each plot, so we adopt

⁴<http://heasarc.nasa.gov/Tools/w3pimms.html>

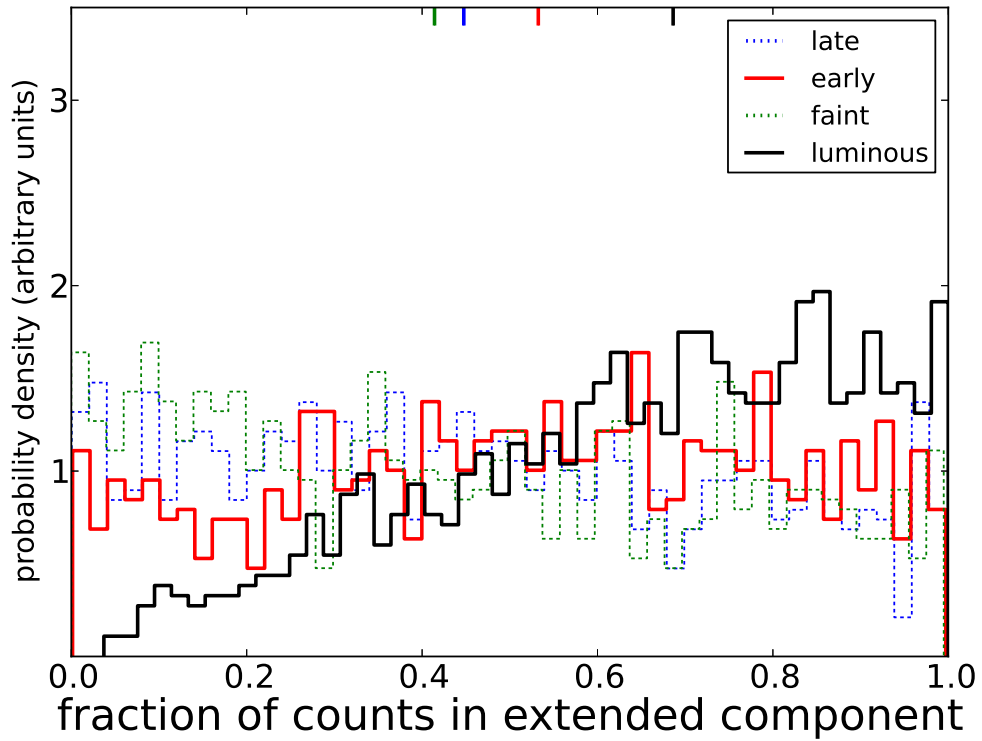


Figure 5.14 The pdfs for the fraction of source counts within 50 kpc of the galaxy, attributed to the extended component for all four subsamples. The median of each distribution is marked with a vertical dash at the top of the plot. The pdfs for late-type and faint subsamples are shown as dashed lines because the statistical significance of the extended component in these samples is much lower than for the others. For all four samples the uncertainties are very large, but the median extended fractions increase with the statistical significance of the extended component, up to a maximum of $\sim 70\%$ for the luminous galaxies.

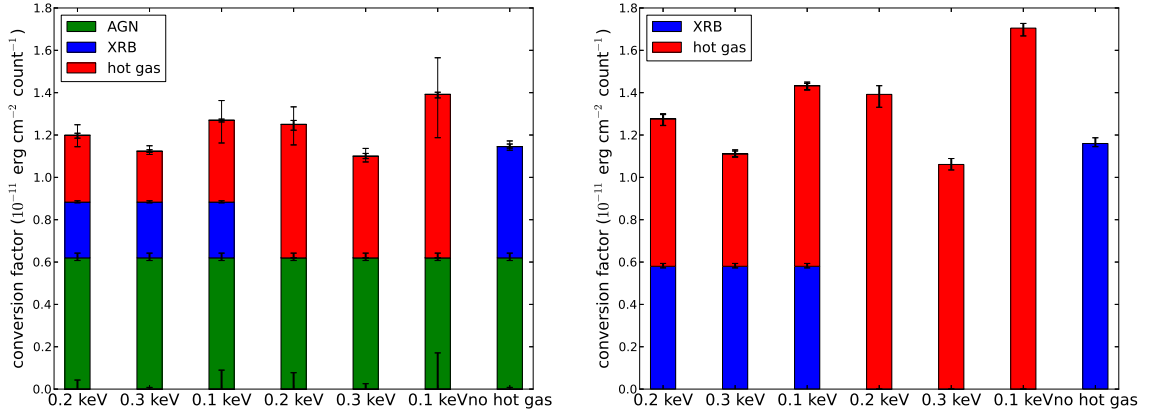


Figure 5.15 Estimates of the conversion factor from counts to energy per unit area for our data, assuming various spectral models for the emission. In (a) we consider models with an AGN component, an X-ray binary component, and a hot gas (APEC) component - i.e. a point source and extended emission. In (b) we only consider the latter two components - i.e. extended emission only. The assumed temperature of the hot gas is noted at the bottom of each model. The error bars at the top of each component show the uncertainties contributed to the total conversion factor by uncertainties in the shape of the spectral model for that component. The error bars at the very bottom of each model show the uncertainty in the total conversion factor from the uncertainties in the fraction of AGN emission. The largest error bar at the top of each model shows the sum of all the uncertainties, added in quadrature. We adopt the leftmost conversion factors from each plot for the rest of the analysis - $1.21 \times 10^{-11} \text{ erg cm}^{-2} \text{ count}^{-1}$ and $1.28 \times 10^{-11} \text{ erg cm}^{-2} \text{ count}^{-1}$ for the total emission and the extended emission, respectively.

conversion factors of 1.21×10^{-11} erg cm⁻² count⁻¹ and 1.28×10^{-11} erg cm⁻² count⁻¹ for the total emission and the extended emission, respectively. The uncertainty in these conversion factors can be estimated from the range of possible values visible in Figures 5.15a and 5.15b, and seems to be no more than 30% upwards and no more than 15% downwards. Most of this upwards scatter comes from the models with the extended component consisting entirely of hot gas at 0.1 keV, and we will show later (section 5.8.3) that these models are unlikely to describe the real data, since X-ray binaries almost certainly contribute significantly to the observed emission.

For the total effective exposure, as discussed above in section 5.4, we already have a stacked exposure map for each sample, where the exposure map of each image is weighted by the quantity $(100 \text{ Mpc} / d)^2$ so that the effective distance of the sample is 100 Mpc. These exposure maps peak at the center of the image and decline slowly at larger radii (by about 10% or so) because not every image in the stack extends out to 500 projected kpc in every direction. This “effective vignetting” is included in our model for the data, so the appropriate exposure time to use for converting energy/area to flux is the value at the center of the image. Thus the stack of late-type galaxies has an effective exposure time of 1.51×10^6 s, the early-type galaxies have an effective exposure of 5.63×10^5 s, the faint galaxies have an effective exposure of 1.71×10^6 s, and the luminous galaxies have an effective exposure of 3.65×10^5 s.

We can now combine the counts-to-energy factor, the distance, and the effective exposure time to convert the counts probability distribution function into a luminosity probability distribution function (Figure 5.16). The luminosity within a radius of 500 kpc is not a physically interesting quantity, so we show the luminosity within 300 kpc instead, as this is closer to the virial radius of an L^* galaxy. There is very little emission beyond 50 kpc, however, so the choice for the outer radius is not very important.

Finally, from the MCMC analysis, we can separate out just the extended component of the emission by integrating only the emission parameterized by the β -model. The average luminosity from this extended emission is shown in Figure 5.17.




Figure 5.16 Probability distribution functions (pdfs) for the average 0.5–2.0 keV X-ray luminosity within radii of (a) 50 kpc and (b) 300 kpc of the galaxy, for each of our four subsamples of 2MIG galaxies. The median of each distribution is marked with a vertical dash at the top of the plot. Note that the X-axis can be slightly shifted left or right due to uncertainties in the conversion from source counts to luminosity, as illustrated in Figure 5.15.

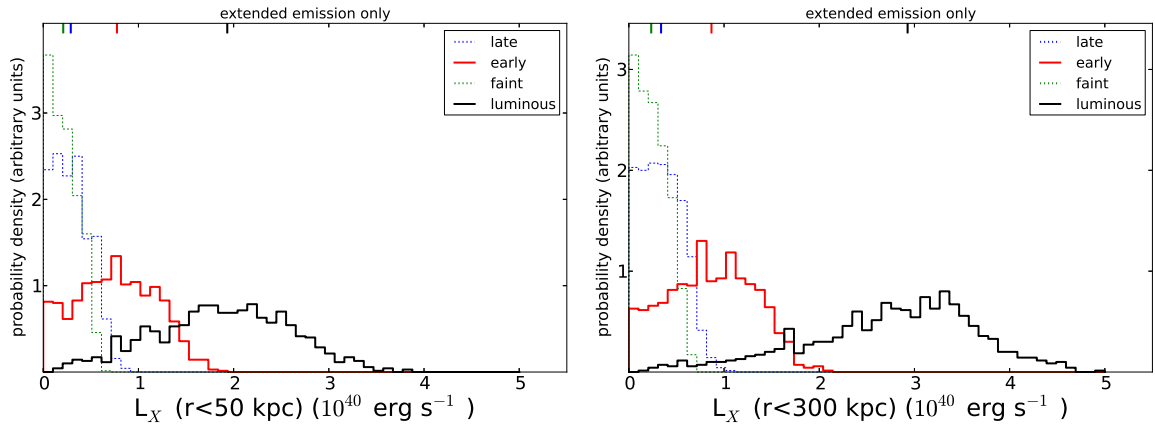


Figure 5.17 Probability distribution functions (pdfs) for the average 0.5–2.0 keV X-ray luminosity of the *extended component* of the emission within radii of (a) 50 kpc and (b) 300 kpc of the galaxy for each of our four subsamples of 2MIG galaxies. The median of each distribution is marked with a vertical dash at the top of the plot. The pdfs for the late-type and faint subsamples are shown as dashed lines because the statistical significance of the extended component is well below 90%. Note that the X-axis can be slightly shifted left or right due to uncertainties in the conversion from source counts to luminosity, as illustrated in Figure 5.15.

5.8.2 AGN emission

While the inclusion of an AGN component does not change the estimation of the total luminosity significantly, it is worth reflecting on the expected luminosity of accretion onto supermassive black holes in these galaxies. Based on the spatial analysis, we are placing about 50% of the emission into the point source component (which includes both AGN emission and X-ray binaries in the central 5 kpc of the galaxy). In detail, the average soft X-ray luminosities for the point source component are 3×10^{39} erg s^{-1} , 7×10^{39} erg s^{-1} , 3×10^{39} erg s^{-1} , and 9×10^{39} erg s^{-1} , for the late, early, faint, and luminous samples, respectively. We know there are no AGN accreting at luminosities of $10^{41.5}$ erg s^{-1} or more in our galaxies, because we examined each image and removed any images with extremely bright point sources, as well as automatically removing any objects in the *ROSAT* BSC. But we also expect most of these galaxies to harbor at least some SMBH accretion, based on high excitation lines visible in optical spectra of isolated nearby galaxies (Hernandez-Ibarra et al. 2012).

To study the effect of AGN emission more quantitatively, we examined the AMUSE-field sample (Miller et al. 2012), a Chandra survey of nearby early-type field galaxies designed to characterize the intensity of low-luminosity AGN activity in these galaxies. We divided the 57 AMUSE-field galaxies in the lowest-density environments into two stellar mass bins, $\log M_* = 9.5-10.5$ (23 galaxies, roughly corresponding to the range covered by our “faint” sample) and $\log M_* = 10.5-11.5$ (34 galaxies, roughly corresponding to the range covered by our “luminous” sample). The faint galaxies all had AGNs fainter than $10^{39.5}$ erg s^{-1} , and only 17% of them even had emission at the 10^{39} erg s^{-1} level. The luminous galaxies all had AGNs fainter than $10^{41.1}$ erg s^{-1} , 41% of them had emission below 10^{39} erg s^{-1} , and the remaining 59% had a median luminosity just under $10^{39.5}$ erg s^{-1} . Both of these samples are fully consistent with our average point source luminosities, and in fact our point source luminosities are about twice the average measured SMBH luminosities for the AMUSE-field galaxies. The other half of the point source component can be attributed to X-ray binaries in the central 5 kpc of the 2MIG galaxies.

5.8.3 X-ray binaries

It is therefore also worth examining the global effect of X-ray binary contamination. They are the most important source of soft X-rays in galaxies other than hot gas and AGN (e.g. Boroson et al. 2011), and they can extend out to tens of kpc. We constrain X-ray binary emission using two independent methods, detailed below.

Hardness Ratio

To attempt to distinguish between hot gas emission and X-ray binary emission, we stacked each sample of 2MIG galaxies in a “soft” 0.5-1.0 keV band and a “hard” 1.0-2.0 keV band, and fit for the total amount of emission within 50 kpc above the background. Ideally we would examine just the extended component, but we do not have enough counts in these narrower bands to constrain the extended component significantly. We computed the hardness ratio (HR), defined as

$$HR \equiv \frac{H - S}{H + S} \quad (5.5)$$

where H and S are the number of counts in the hard and soft components, respectively. The 68% confidence intervals for HR for the four subsamples are shown in Figure 5.18. In this Figure we also plot the expected HR for various emission models. The green lines are powerlaw emission models with a slope of -2 and absorbing columns of 1,5, and $10 \times 10^{20} \text{ cm}^{-2}$. These are intended to mimic AGN emission, although as discussed above it is difficult to distinguish AGN and XRB emission based on their spectral hardness in the soft X-rays. The blue lines are powerlaw emission models with a slope of -1.56 and absorbing columns of 1,5, and $10 \times 10^{20} \text{ cm}^{-2}$, designed to mimic XRB emission. And the three sets of red lines are APEC emission models with $kT = 0.1, 0.2,$ and 0.3 keV , respectively. For the APEC models, the metallicity ranges from $0.2Z_{\odot}$ to $1.0Z_{\odot}$ and the absorbing columns range from 10^{20} to 10^{21} cm^{-2} . Higher absorption increases the hardness ratio.

We can see that all five samples have statistically indistinguishable hardness ratios. In general no linear combination of AGN and XRB can reproduce the observed hardness ratios, except pure AGN absorption, and that is disfavored by both the




Figure 5.18 Hardness ratio of source counts within 50 kpc of the galaxy, for the full sample and all four samples of 2MIG galaxies, along with hardness ratios predicted by various spectral models. The shaded red region corresponds to APEC models for hot gas emission at the indicated temperatures; the region encompasses different choices for metallicity (ranging from $0.2Z_{\odot}$ to $1.0Z_{\odot}$ and absorption (ranging from $1 \times 10^{20} \text{ cm}^{-2}$ to $1 \times 10^{21} \text{ cm}^{-2}$), with the darkest shading at the central values for these choices. The shaded green region corresponds to a model for AGN emission (an absorbed $\Gamma = 2$ powerlaw), and the blue shaded region corresponds to the union of two different models for XRB emission; note that the green and blue regions overlap significantly. For full details of these models, see the text. In all models, increasing absorption increases the hardness ratio. The data primarily fall between the XRB and APEC models, and overlap with the AGN models. The fact that no linear combination of XRB + AGN model can reproduce the observed hardness ratios suggests that an APEC (hot gas) component is required as well. The three black lines correspond to three linear combinations of the three components. The dotted black line is an equal combination of all three components, the dashed black line is the leftmost model from Figure 5.15a with approximately a 2:1:1 ratio of AGN : APEC : XRB emission. The solid black line has 1/4 AGN emission and finds the best fit for the APEC and XRB fractions in the full sample of 2MIG galaxies: 60% and 15%.

spatial analysis and the expected average AGN luminosity. If we include any XRB emission, we need to include APEC emission as well to keep the total hardness ratio soft enough to fit the data. Thus, since we expect XRB emission in each sample, the measured hardness ratios imply the presence of hot gas in each of our samples as well.

We can make a simple estimate of the XRB and APEC contributions. We adopt what we consider the likeliest values for each spectral model - the XRB and AGN models with $N_H = 5 \times 10^{20} \text{ cm}^{-2}$, and the APEC model with $kT = 0.2 \text{ keV}$, $Z = 0.5Z_\odot$, and $N_H = 5 \times 10^{20} \text{ cm}^{-2}$. We attribute 1/4 of the emission to SMBH accretion (with an absorbed $\Gamma = 2$ powerlaw), and the rest to XRBs and the hot gas. Since the spatial analysis suggests $\approx 1/2$ of the emission belongs in the point source component, this implies the rest of the point source emission comes from X-ray binaries. Solving for the resulting fraction of the total emission attributed to the other components, we find $f_{\text{XRB}} = 60\%$ and $f_{\text{APEC}} = 15\%$.

Scaling Relations

We reach a similar conclusion if we estimate the X-ray binary luminosity from established scaling relations. There is a correlation between L_K for a galaxy and its low-mass XRB (LMXB) luminosity. The most recent estimate (Boroson et al. 2011) finds $L_{\text{LMXB}} \approx 1 \times 10^{29} (L_K/L_\odot) \text{ erg s}^{-1}$. This is a broad-band measurement of L_X (0.3-8 keV); we use their spectral model (thermal bremsstrahlung with $kT = 7 \text{ keV}$, although they note a powerlaw with $\Gamma = 1.4 - 1.8$ works about as well) to convert this L_X into our 0.5-2.0 keV band. We assume $N_H = 5 \times 10^{20} \text{ cm}^{-2}$. The resulting scaling relation is

$$L_{\text{LMXB}} \approx 3 \times 10^{28} \frac{L_K}{L_\odot} \text{ erg s}^{-1} \quad (5.6)$$

For high-mass X-ray binaries, the luminosity scales with the galactic star formation rate (SFR). The most recent estimate (Mineo, Gilfanov, and Sunyaev 2012) is $L_{\text{HMXB}} \approx 2.6 \times 10^{39} \text{ SFR erg s}^{-1}$ for a SFR in $M_\odot \text{ yr}^{-1}$. Again, this is a broad-band measurement of L_X (0.5-8 keV), and we use their spectral model (a $\Gamma = 2$ powerlaw)

with $N_H = 5 \times 10^{20} \text{ cm}^{-2}$ to convert into our 0.5-2.0 keV band. The resulting scaling relation is

$$L_{\text{HMXB}} \approx 1.4 \times 10^{39} \frac{\text{SFR}}{M_{\odot} \text{ yr}^{-1}} \text{ erg s}^{-1} \quad (5.7)$$

We can use these scaling relations to estimate the expected XRB luminosity for each of our subsamples. We compute the mean exposure-weighted L_K for the late-type, early-type, faint, and luminous subsamples: $4.4 \times 10^{10} L_{\odot}$, $7.4 \times 10^{10} L_{\odot}$, $3.3 \times 10^{10} L_{\odot}$, and $1.4 \times 10^{11} L_{\odot}$, respectively. This yields L_{LMXB} of $1.3 \times 10^{39} \text{ erg s}^{-1}$, $2.2 \times 10^{39} \text{ erg s}^{-1}$, $1.0 \times 10^{39} \text{ erg s}^{-1}$, $4.3 \times 10^{39} \text{ erg s}^{-1}$, respectively. We assume mean SFRs of $0.1 M_{\odot} \text{ yr}^{-1}$ for the early-type galaxies (including S0s) and $1.0 M_{\odot} \text{ yr}^{-1}$ for the late-type galaxies, and since about 2/3 of the galaxies are late-type, we assume mean SFRs of $0.7 M_{\odot} \text{ yr}^{-1}$ for the faint and the luminous samples. This yields L_{HMXB} of $1.4 \times 10^{39} \text{ erg s}^{-1}$, $1.4 \times 10^{38} \text{ erg s}^{-1}$, $9.8 \times 10^{38} \text{ erg s}^{-1}$, and $9.8 \times 10^{38} \text{ erg s}^{-1}$, respectively. Therefore the average total XRB luminosity for the late-type, early-type, faint, and luminous galaxies, respectively, is predicted to be $2.7 \times 10^{39} \text{ erg s}^{-1}$, $2.4 \times 10^{39} \text{ erg s}^{-1}$, $2.0 \times 10^{39} \text{ erg s}^{-1}$, and $5.2 \times 10^{39} \text{ erg s}^{-1}$.

Comparing these predictions to the median total luminosities within 50 kpc (Figure 5.16), we see that XRBs are predicted to comprise 43%, 17%, 41%, and 18% of the total emission for the late-type, early-type, faint, and luminous galaxies, respectively. These fractions are lower than the 59% fraction inferred from the spectral hardness ratios, but given the huge uncertainties in each method, the rough agreement is still encouraging. Our (conservative) conclusion is that the data seem to suggest about half of the emission belongs in the point source component, and about 2/3 of the remaining (extended) emission is due to X-ray binaries, but there is a lot of room for variations between different samples and for complications in the spectral modeling. In particular, the luminous and the early-type galaxies seem to have a higher fraction of hot gas emission than the other samples.

5.8.4 The L_X - L_K relation

Now that we have X-ray luminosities for the hot gas around our samples of isolated galaxies, we present the L_X - L_K relation for our samples. Figure 5.19 shows an L_X - L_K plot with our three samples with the strongest suggestion of extended emission - the luminous galaxies and the early-type and late-type luminous subsamples. We also include data from other recent measurements (in the same 0.5–2.0 keV energy band) of isolated elliptical galaxies (Mulchaey and Jeltema 2010, Memola et al. 2009). For our samples, we plot the mean exposure-weighted L_K (with errorbars enclosing 68% of the galaxies in each sample) and the median of the pdf of the X-ray luminosity of the extended component out to 50 kpc (along with the 68% confidence interval on this luminosity). We also indicate with an X symbol the median luminosity after subtracting the X-ray binary luminosity as estimated from the scaling relations (section 5.8.3). This is probably an over-estimate of the effect of XRBs, however, since $\sim 50\%$ of the X-ray binary emission may be in the point source component instead of the extended component. Our 2MIG luminous galaxies seem to fall within the scatter of the L_X - L_K relation, especially after subtracting the estimated XRB emission.

5.8.5 Hot Gas Mass

Finally, after analyzing the X-ray luminosities of the hot gas, we turn to the hot gas mass. We estimate the hot gas mass simply by integrating the β -model and marginalizing over β , r_0 , and A_β . Unfortunately, the measured pdf for β does not provide a reliable guide to the true value of β (see Figure 5.7). Moreover, we do not know the outer extent of the hot gas since the detected emission drops off beyond about 50 kpc. Therefore we restrict the analysis of the gas mass to within a radius of 50 kpc, where the gas luminosity is fairly well constrained and variations in β are not yet too significant (the slope becomes increasingly important for the mass at larger radii).

We also need a way to relate β , r_0 , and A_β to the central density n_0 . We can connect these parameters using the spectral model (section 5.8.1). For an APEC model with $kT = 0.2$ keV, $Z = 0.3Z_\odot$, and $N_H = 5 \times 10^{20}$ cm $^{-2}$, we simulated obser-




Figure 5.19 L_X - L_K relation for all the luminous, the luminous late-type, and the luminous early-type galaxies. The value of L_K is the mean exposure time-weighted value for each sample, and the error bars enclose the central 68% of the galaxies in each sample. The value of L_X is the median X-ray luminosity of the extended component of the emission, integrated out to 50 kpc, and the error bars enclose the 68% confidence interval on this luminosity. The X symbols show the estimated L_X after subtracting the maximum estimated emission from X-ray binaries. For comparison we also include data from other studies of isolated elliptical galaxies.

vations with Webspec⁵ and found that the observed 0.5-2.0 keV count rate is equal to 19.32 times the normalization of the APEC model. Unfortunately, at this temperature the proportionality constant is nearly linearly dependent on the metallicity of the gas, and this metallicity could plausibly be anywhere from $0.1Z_{\odot}$ to $1.0Z_{\odot}$. This metallicity uncertainty introduces a factor of three uncertainty in the estimate of the normalization. Even worse, the gas could even be incompletely mixed, or host a metallicity gradient, and we have no way to account for these possibilities, so the normalization is potentially even more uncertain.

After estimating the normalization of the APEC model from the count rate, we can relate the normalization to the central density using the relation

$$\text{norm} \equiv \frac{10^{-14}}{4\pi d_L} \int n_e n_H dV \quad (5.8)$$

where d_L is the luminosity distance to the galaxy, and n_e and n_H and the electron and Hydrogen column densities, respectively. Substituting the β -profile into the equations for n_e and n_H , converting n_H into n_e , and setting $d_L = 100$ Mpc, we can use this equation to compute the pdf for n_0 for each sample, and then integrate the mass profile out to 50 kpc to get a gas mass pdf.

To account for contamination from X-ray binaries, we divide the values of A_{β} from the MCMC fit by 3, since we estimate above (section 5.8.3) that 2/3 of the extended emission comes from X-ray binaries. This implicitly assumes the X-ray binaries have the same density profile as the hot gas, which is probably not correct, but with the data we have this is the best we can do. The resulting pdfs for the gas mass are shown in Figure 5.20.

Accretion Rates

Finally, we can combine our knowledge of the mass profiles and the luminosity profiles to estimate the cooling rate of the hot gas. The total amount of cooling is set by the X-ray luminosity of the gas at small radii, and we have some confidence in our ability to measure this quantity. In the absence of nongravitational heating, the cooling rate

⁵<http://heasarc.gsfc.nasa.gov/webspec/webspec.html>




Figure 5.20 Probability distribution functions (pdfs) for the hot gas mass within 50 kpc of the galaxy. The median of each distribution is marked with a vertical dash at the top of the plot. The masses were corrected downwards to account for contamination from X-ray binaries. The pdfs for the late-type and faint subsamples are shown as dashed lines because the statistical significance of the extended component well below 90%. The gas was assumed to have $kT = 0.2$ keV and $Z = 0.3Z_{\odot}$; uncertainties in the metallicity of the gas lead to additional uncertainties in the inferred mass on the order of 75%.

is equal to the hot-mode accretion rate onto the central galaxy. The accreted gas can then serve as fuel for star formation in the galaxy.

However, at least for early-type galaxies, the cooling rates measured in pointed observations ($\sim 1M_{\odot} \text{ yr}^{-1}$, e.g. Nulsen, Stewart, and Fabian 1984, Fabbiano 1989) are typically much larger than the (very low) star formation rates typically measured in these galaxies (e.g. O’Connell 1999). This is the so-called “cooling-flow” problem in elliptical galaxies. It is thought that the X-ray cooling is counteracted by heating from AGN feedback or other heating processes (Rosner and Tucker 1989, Tucker and David 1997, Mathews and Brighenti 2003), so that the actual hot-mode accretion rate is closer to the star formation rate than the cooling rate.

On the other hand, it obviously should not be possible for the hot mode accretion rate to be higher than the X-ray cooling rate, so the cooling rate sets an upper bound on the hot mode accretion rate (although other modes of accretion can still be extremely important). Moreover, we are unlikely to underestimate significantly the average cooling rate, since our simulations showed that we can recover the luminosity to 30% or better, and that if the true luminosity were higher we could recover it even more accurately. Since the cooling rate is set by the behavior of the gas at small radii, this quantity is also not very sensitive to the slope of the density profile.

We compute the cooling rate by defining a cooling radius, estimating the thermal energy of the hot gas within the cooling radius, and dividing by the X-ray luminosity to infer a cooling time. The cooling rate is then the gas mass within the cooling radius, divided by the cooling time. Following Fukugita and Peebles (2006), we start with the cooling time:

$$\tau(r) = \frac{1.5nkT}{\Lambda n_e(n - n_e)} \approx \frac{1.5kT \times 1.92}{\Lambda n_e \times 0.92} \quad (5.9)$$

where the latter expression assumes the total particle density $n = 1.92n_e$. For $T = 10^{6.35}$ K and $Z = 0.3Z_{\odot}$, $\Lambda = 10^{-22.54} \text{ erg cm}^3 \text{ s}^{-1}$ (Sutherland and Dopita 1993). Thus for $\tau = 10$ Gyr, the cooling radius occurs at $n_e = 1.1 \times 10^{-4} \text{ cm}^{-3}$. Marginalizing over our density profiles, the median cooling radii for the luminous

galaxies, the luminous late-type galaxies, and the luminous early-type galaxies are 54 kpc, 48 kpc, and 46 kpc, respectively. The median implied cooling rates are $0.4M_{\odot} \text{ yr}^{-1}$, $0.3M_{\odot} \text{ yr}^{-1}$, and $0.4M_{\odot} \text{ yr}^{-1}$, respectively.

The inferred cooling rate around our luminous early-type galaxies is consistent with cooling rates measured around ellipticals in pointed observations, and illustrates the same cooling-flow problem, as noted above. On the other hand, for the luminous late-type galaxies, our inferred accretion rate is a few times lower than the expected star formation rate. A similar situation entails for the Milky Way, which also has a hot gas accretion rate several times smaller than the observed star formation rate (Anderson and Bregman 2010, Putman et al. 2012), and the implied accretion rates from the hot coronae around the Tüllmann et al. (2006) sample of star-forming spirals are also lower than the observed star formation rates. Therefore, the conclusion must be that hot mode accretion on its own is insufficient to fuel star formation in isolated massive spirals. Either additional modes of accretion must supplement the hot gas cooling, or mass lost from stellar winds (Leitner and Kravtsov 2011) must make up the difference.

Baryon budgets

As mentioned above, in our model the total hot gas is primarily determined by the value of β and by the outer radius of the hot halo. We do not have strong constraints on either of these parameters, other than that the hot halo seems to extend out to at least 50 kpc, so we do not integrate the mass profiles beyond 50 kpc. In this section, however, we speculate briefly on the possible behavior at larger radii.

Based on previous observations described in section 5.8, we believe the likeliest value for β is $\beta \approx 0.5$, but the flattest density profiles in simulations have $\beta \approx 0.35$ (e.g. Kauffmann et al. 2009) and the steepest observed profiles approach $\beta \approx 0.7$ (numerical predictions by Feldmann et al. (2012) also predict a similarly large value of β), thus motivating the range of values we explored in our simulated images. If we integrate the β -model from 50 kpc outwards, we can get a sense of the potential hot gas mass at larger radii. Figure 5.21 shows the ratio of the hot gas mass out to




Figure 5.21 Ratio of the hot gas mass within 200 or 300 kpc to the hot gas mass within 50 kpc, assuming the hot gas follows a β -model. The range on each distribution reflects the effect of varying the core radius r_0 between 0.5 kpc and 5.0 kpc.

200 or 300 kpc to the hot gas mass within 50 kpc, as a function of β . The range on each distribution reflects the effect of varying the core radius r_0 between 0.5 kpc and 5.0 kpc – the effect of varying r_0 is not very significant. Within 200 kpc, the hot gas mass can increase from a factor of a few up to a factor of 15, and if we extend the profile to 300 kpc the total mass can increase very significantly for the flatter profiles.

From section 5.8.5, the median masses of the hot gas within 50 kpc around the early-type galaxies and the luminous galaxies are $1.9 \times 10^9 M_\odot$ and $4.1 \times 10^9 M_\odot$. For $\beta = 0.5$, integrating out to 200 kpc would increase these masses to $1.5 \times 10^{10} M_\odot$ and $3.3 \times 10^{10} M_\odot$, respectively. These inferred gas masses are much larger than the observed masses, and are now comparable to the average exposure-weighted stellar masses of the galaxies in their respective samples. Since the stars in the average massive galaxy contain less than half of the baryons expected in the halo based on the cosmological baryon fraction (e.g. McGaugh et al. 2010), the hot halo appears to be a

tempting possible reservoir for these missing baryons. The slope of the density profile would have to be somewhat flatter than 0.5 to account for all the missing baryons, but the *ROSAT* observations do not constrain this extrapolation, and indeed it is very difficult to place any constraints on X-ray emission at these radii. However, as we pointed out in Anderson and Bregman (2010), massive hot halos extending to such large radii violate X-ray absorption-line constraints (and decreasing β exacerbates the problem). Therefore we think the hot halo probably does not follow a β -model out to hundreds of kpc.

5.8.6 Comparison with previous work

We briefly compare our results to results from X-ray stacking analyses of other samples of galaxies. Often, this sort of stacking has been performed for high-redshift galaxies in the deep fields (e.g. $z \sim 3$ Lyman-break galaxies; Nandra et al. 2002, Lehmer et al. 2005, and Laird et al. 2006); $z \sim 1$ UV-luminous galaxies; Laird et al. 2005) which have very different properties from the nearby isolated galaxies in our sample. A few papers do examine “normal” galaxies at lower redshift, however. Lehmer et al. (2007) stacked 539 early-type galaxies from $0.1 \lesssim z \lesssim 0.7$, and found mean $z \approx 0.25$ X-ray luminosities (0.5-2.0 keV) for their luminous and faint samples of $10^{40.2 \pm 0.2}$ erg s⁻¹ and $10^{39.6 \pm 0.4}$ erg s⁻¹, respectively, with no evidence for redshift evolution. While they use a different optical classification for luminous and faint galaxies, and a smaller aperture than our analysis (based on maximizing the observed signal-to-noise), their results are consistent with ours. Watson et al. (2009) perform a similar analysis on 6146 galaxies in the AGN and Galaxy Evolution Survey (AGES) over $0.1 \lesssim z \lesssim 0.5$. They parameterize their results in terms of L_K and $(1+z)$; plugging in our mean values for these quantities, we infer a prediction of $L_X = 2.4 \pm 0.7 \times 10^{40}$ erg s⁻¹, which is not too different from our value of $L_X = 1.0 \pm 0.3 \times 10^{40}$ erg s⁻¹. Their values could also be slightly higher because their sample is not restricted to isolated galaxies; galaxies in denser environments might be expected to have higher rates of SMBH accretion, and also they have additional emission from a hot intracluster/intragroup medium.

5.9 Conclusions

We examine 2496 nearby isolated galaxies from the 2MASS Isolated Galaxies survey, and stack *ROSAT* All-Sky Survey images of 2165 of these galaxies. We also divide the full 2165 galaxies by morphology (into subsamples of late-type and early-type galaxies) and luminosity (into subsamples of luminous and faint galaxies), and examine each subsample in turn. We also examine further subsamples of the luminous galaxies, divided into luminous early-type and luminous late-type galaxies, and we produce a stacked image of random positions on the sky as well as a null sample. For each sample, we stacked all the 0.5-2.0 keV photons within 500 projected kpc of the galaxy, and also stacked the RASS exposure map out to the same radius, weighing the exposure map by the inverse square of the distance to the galaxy.

In the stacked images, we model the surface brightness profiles as the sum of a uniform background, a point source convolved with an empirically determined PSF, and an extended component (a β -model) convolved with the same PSF. We also include in the model an “effective vignetting” caused by incomplete coverage at larger radii in some images, which we measured from the stacked effective exposure map. We employ an MCMC analysis to compare the model to the observed surface brightness profile.

For each of our samples, except for the null sample, we also create a variety of simulated images which appear similar to the real images, and test how well the MCMC analysis can recover the various image parameters. In general the MCMC analysis can recover the number of source counts within 50 kpc of the central galaxy, but it cannot recover the slope and core radius of the extended component reliably. It also has mixed success in distinguishing emission in the point source component and the extended component, but the MCMC analysis, a likelihood ratio test, and the hardness ratio of the emission all suggest the emission is composed of a mixture of the two components in fairly comparable amounts.

We examine the Poisson statistics of the counts within 50 projected kpc of the center of each image in order to estimate the statistical significance of detecting

emission in each stacked image. We detect emission in all stacked images, except the null image, at $> 99\%$ confidence. We then employ likelihood ratio tests and simulate images in order to estimate the statistical significance of detecting *extended* (beyond 5 kpc) emission in each stacked image. For five out of seven samples (all galaxies, early-type galaxies, luminous galaxies, luminous late-type galaxies, and luminous early-type galaxies) an extended component has a statistical significance of $\gtrsim 90\%$ (and the luminous galaxies have an extended component at $> 99\%$); the other two samples have extended emission at much lower significance. We compute the average X-ray luminosity within 50 kpc of each sample of galaxies, and we compute the average X-ray luminosity of the extended component within 50 kpc as well, with the caveat that the extended component may not be statistically necessary in some cases.

From the observed hardness ratios of the emission, and from empirical scaling relations, we estimate the contribution of X-ray binaries to the luminosity within 50 kpc of the galaxy in each of our samples. While the uncertainties are very large, both methods generally agree, and suggest that about 2/3 of the extended emission is due to X-ray binaries, with the remaining 1/3 of the extended emission attributed to hot gas. We also estimate that the point source component is roughly evenly distributed between a weakly accreting SMBH and X-ray binaries in the central 5 kpc of the galaxy. We can then infer the mass of the hot gas within 50 kpc, again with the caveat that the extended component may not be statistically required for some of the stacked samples of galaxies.

We find that the average X-ray luminosity of the early-type galaxies is about twice the X-ray luminosity of the late-type galaxies, and that the statistical significance of the extended emission around the early-type galaxies is also much higher. The early-type galaxies therefore seem to have larger hot gas halos than the late-type galaxies. It is unclear, however, if this is related to the morphology of the galaxies or simply a consequence of the early-type galaxies having larger average stellar masses and luminosities than the late-type galaxies. For the luminous subsamples, the late-type galaxies and early-type galaxies have the same X-ray luminosities and hot gas masses, within the (large) uncertainties, although the luminous late-type galaxies

have extended emission with higher statistical confidence.

Finally, for the samples where we measure a hot gas mass, we also infer a hot gas accretion rate onto the central galaxy. These accretion rates are a few times smaller than the expected star formation rates in the luminous late-type galaxies, and a few times larger than the expected star formation rates in the luminous early-type galaxies. We comment briefly on the expected hot gas masses if the β -model is extrapolated out to the virial radius. If the hot gas has a typical value of $\beta \approx 0.5$, and extends to the virial radius, the total hot gas mass would be expected to be 10-20 times larger than the value we measure. Flatter values of β would increase this mass significantly, as would abundance gradients in the hot halo, but at present we have no observational constraint on either of these possibilities.

All of the measured and derived quantities discussed above are presented in Table 5.4. Median values of the pdf are listed, along with 68% confidence intervals (identified as the statistical errors). We also estimate the size of the systematic uncertainties on each quantity and list those uncertainties as well; for detailed discussion of the systematic uncertainties, see the appropriate sections of the paper.

In the next few years, the E-Rosita mission is expected to launch, and to produce another all-sky survey with coverage in the soft X-rays. This survey could be up to an order of magnitude deeper than the RASS. It will be very interesting to examine the properties of stacked galaxies in this survey, as we should be able to characterize the extended emission with much more precision.

5.10 Acknowledgements

The authors would like to thank E. Bell, E. Gallo, O. Gnedin, B. Miller, C. Miller, J. Miller, M. Ruszkowski, C. Slater, R. Smith, and A. Vikhlinin for very helpful conversations and suggestions which greatly improved this manuscript. We would also like to thank the referee for suggestions and references that greatly improved the result. This research has made use of data and/or software provided by the High Energy Astrophysics Science Archive Research Center (HEASARC), which is a service of the Astrophysics Science Division at NASA/GSFC and the High Energy Astrophysics

Division of the Smithsonian Astrophysical Observatory. This research has made use of NASA's Astrophysics Data System. M.E.A. also gratefully acknowledges support from the NSF in the form of a Graduate Research Fellowship.

Table 5.4 Summary of measurements and derived quantities

sample	L_K ($10^{10} L_\odot$)	$1 - p(\text{any})$	$1 - p(\text{ext})$	$L_{X,<50}$ ($10^{39} \text{ erg s}^{-1}$)	$L_{X,\text{ext},<50}$ ($10^{39} \text{ erg s}^{-1}$)	$L_{X,<300}$ ($10^{39} \text{ erg s}^{-1}$)	$M_{\text{gas},<50}$ ($10^9 M_\odot$)	\dot{M} ($M_\odot \text{ yr}^{-1}$)
All galaxies	$5.1^{+10.5}_{-2.2}$	4.4×10^{-16} (8.1 σ)	0.073	$10.0 \pm 0.8(\text{stat}) \pm 2.0(\text{sys})$	$5.5^{+3.1}_{-3.4}(\text{stat}) \pm 1.1(\text{sys})$	$10.6^{+0.9}_{-0.8}(\text{stat}) \pm 2.1(\text{sys})$	$1.6^{+0.5}_{-0.6}(\text{stat}) \pm 1.2(\text{sys})$	0.09 ± 0.06
Late-type galaxies	$4.4^{+9.7}_{-1.9}$	0.000446	0.347	$6.3 \pm 0.9(\text{stat}) \pm 0.9(\text{sys})$	$2.9^{+2.5}_{-2.0}(\text{stat}) \pm 0.4(\text{sys})^*$	$6.7^{+1.1}_{-0.9}(\text{stat}) \pm 1.0(\text{sys})$	$1.2^{+0.5}_{-0.6}(\text{stat}) \pm 0.8(\text{sys})^*$	$0.05 \pm 0.04^*$
Early-type galaxies	$7.4^{+12.0}_{-3.1}$	3.0×10^{-9} (5.9 σ)	0.107	$14.2^{+1.6}_{-1.8}(\text{stat}) \pm 4.3(\text{sys})$	$7.7^{+4.6}_{-4.8}(\text{stat}) \pm 2.3(\text{sys})$	$15.0^{+1.8}_{-1.7}(\text{stat}) \pm 4.5(\text{sys})$	$1.9 \pm 0.7(\text{stat}) \pm 1.3(\text{sys})$	0.1 ± 0.1
Faint galaxies	$3.3^{+4.0}_{-1.2}$	0.005312	0.528	$4.9^{+0.7}_{-0.8}(\text{stat}) \pm 1.5(\text{sys})$	$2.1^{+2.0}_{-1.6}(\text{stat}) \pm 0.6(\text{sys})^*$	$5.2 \pm 0.8(\text{stat}) \pm 1.6(\text{sys})$	$0.9^{+0.5}_{-0.4}(\text{stat}) \pm 0.6(\text{sys})^*$	$0.04 \pm 0.03^*$
Luminous galaxies	$14.2^{+8.0}_{-4.1}$	2.6×10^{-23} (9.9 σ)	0.007	$28.3^{+3.9}_{-4.1}(\text{stat}) \pm 4.2(\text{sys})$	$19.3^{+7.2}_{-8.3}(\text{stat}) \pm 2.9(\text{sys})$	$36.2^{+4.5}_{-4.0}(\text{stat}) \pm 5.4(\text{sys})$	$4.1^{+0.6}_{-1.0}(\text{stat}) \pm 2.9(\text{sys})$	$0.4^{+0.1}_{-0.2}$
Luminous late-type galaxies	$13.5^{+7.2}_{-3.6}$	0.000025	0.013	$25.9^{+4.6}_{-4.3}(\text{stat}) \pm 3.9(\text{sys})$	$16.7^{+7.6}_{-8.0}(\text{stat}) \pm 2.5(\text{sys})$	$31.5^{+5.0}_{-4.4}(\text{stat}) \pm 4.7(\text{sys})$	$3.6^{+0.8}_{-1.1}(\text{stat}) \pm 2.5(\text{sys})$	$0.3^{+0.1}_{-0.2}$
Luminous early-type galaxies	$15.2^{+9.0}_{-4.6}$	3.2×10^{-15} (7.9 σ)	0.079	$35.6^{+5.6}_{-5.2}(\text{stat}) \pm 5.3(\text{sys})$	$20.0^{+11.0}_{-11.5}(\text{stat}) \pm 3.0(\text{sys})$	$41.0^{+6.0}_{-5.1}(\text{stat}) \pm 6.1(\text{sys})$	$3.5^{+1.2}_{-1.4}(\text{stat}) \pm 2.5(\text{sys})$	0.4 ± 0.2

Measured and derived quantities for stacks of various samples of 2MIG galaxies. L_K is the mean exposure-weighted K-band luminosity of a galaxy in each sample, and the range enclosing the central 68% of luminosities is indicated. The next two columns show the statistical significance of detections of any emission and of extended emission within the central 50 projected kpc of the stacked images. $L_{X,<50}$ is the median of the pdf of the 0.5-2.0 keV X-ray luminosity within 50 physical kpc of the galaxy in each stack. The 68% confidence interval is indicated with (stat) and the estimated impact of systematic errors is indicated with (sys). $L_{X,\text{ext},<50}$ shows the luminosity of extended emission within 50 kpc, in the same format. $L_{X,<300}$ is the same as $L_{X,<50}$ except the emission is integrated out to 300 kpc from the galaxy instead of 50 kpc. $M_{\text{gas},<50}$ is the estimated hot gas mass within 50 kpc of the galaxy, after correcting for X-ray binary contamination. The systematic uncertainty accounts for metallicities ranging from $0.1Z_\odot$ to $1.0Z_\odot$. Finally, \dot{M} is the estimated accretion rate of hot gas onto the central galaxy (and 68% confidence intervals on the median value). Asterisks denote the quantities based on extended emission detected at $\lesssim 90\%$ significance.

5.11 References

- Adelberger, K. L., Steidel, C. C., Shapley, A. E., & Pettini, M. 2003, *ApJ*, 584, 45
- Anderson, M. E., & Bregman, J. N. 2010, *ApJ*, 714, 320
- Anderson, M. E., & Bregman, J. N. 2011, *ApJ*, 737, 22
- Bell, E. F., McIntosh, D. H., Katz, N., & Weinberg, M. D. 2003, *ApJS*, 149, 289
- Benson, A. J., Bower, R. G., Frenk, C. S., & White, S. D. M. 2000, *MNRAS*, 314, 557
- Birnboim, Y., & Dekel, A. 2003, *MNRAS*, 345, 349
- Bogdán, A., David, L. P., Jones, C., Forman, W. R., & Kraft, R. P. 2012, *ApJ*, 758, 65
- Boroson, B., Kim, D.-W., & Fabbiano, G. 2011, *ApJ*, 729, 12
- Bower, R. G., Benson, A. J., Malbon, R., et al. 2006, *MNRAS*, 370, 645
- Bregman, J. N., & Glassgold, A. E. 1982, *ApJ*, 263, 564
- Crain, R. A., McCarthy, I. G., Frenk, C. S., Theuns, T., & Schaye, J. 2010, *MNRAS*, 407, 1403
- Dai, X., Anderson, M. E., Bregman, J. N., & Miller, J. M. 2012, *ApJ*, 755, 107
- Dai, X., Bregman, J. N., Kochanek, C. S., & Rasia, E. 2010, *ApJ*, 719, 119
- Dai, X., Kochanek, C. S., & Morgan, N. D. 2007, *ApJ*, 658, 917
- De Lucia, G., Springel, V., White, S. D. M., Croton, D., & Kauffmann, G. 2006, *MNRAS*, 366, 499
- Dekel, A., & Birnboim, Y. 2006, *MNRAS*, 368, 2
- Dickey, J. M., & Lockman, F. J. 1990, *ARA&A*, 28, 215
- Dunkley, J., Komatsu, E., Nolte, M. R., et al. 2009, *ApJS*, 180, 306
- Fabbiano, G. 1989, *ARA&A*, 27, 87
- Faucher-Giguère, C.-A., Kereš, D., & Ma, C.-P. 2011, *MNRAS*, 417, 2982
- Feldmann, R., Hooper, D., & Gnedin, N. Y. 2012, arXiv:1205.0249
- Forman, W., Jones, C., & Tucker, W. 1985, *ApJ*, 293, 102
- Forman, W., Schwarz, J., Jones, C., Liller, W., & Fabian, A. C. 1979, *ApJL*, 234, 27
- Fukugita, M., & Peebles, P. J. E. 2006, *ApJ*, 639, 590

Guedes, J., Callegari, S., Madau, P., & Mayer, L. 2011, *ApJ*, 742, 76

Gupta, A., Mathur, S., Krongold, Y., Nicastro, F., & Galeazzi, M. 2012, *ApJL*, 756, 8

Hasinger, G., Boese, G., Predehl, P., et al. 1993, MPE/OGIP Calibration Memo, 93, 015, <http://heasarc.gsfc.nasa.gov/docs/journal/rosat-off-axis-psf4.html>

Hernandez-Ibarra, F., Dultzin, D., Krongold, Y., et al. 2012, arXiv:1206.6777

Houk, N., & Cowley, A. P. 1975, University of Michigan Catalogue of two-dimensional spectral types for the HD stars (Vol. 1; Ann Arbor, MI: Univ. of Michigan)

Humphrey, P. J., Buote, D. A., Canizares, C. R., Fabian, A. C., & Miller, J. M. 2011, *ApJ*, 729, 53

Humphrey, P. J., Buote, D. A., OSullivan, E., & Ponman, T. J. 2012, *ApJ*, 755, 166

Irwin, J. A., Athey, A. E., & Bregman, J. N. 2003, *ApJ*, 587, 356

Karachentseva, V. E., Mitronova, S. N., Melnyk, O. V., & Karachentsev, I. D. 2010, *AstBu*, 65, 1

Kaufmann, T., Bullock, J. S., Maller, A. H., Fang, T., & Wadsley, J. 2009, *MNRAS*, 396, 191

Kochanek, C. S., Pahre, M. A., Falco, E. E., et al. 2001, *ApJ*, 560, 566

Laird, E. S., Nandra, K., Adelberger, K. L., Steidel, C. C., & Reddy, N. A. 2005, *MNRAS*, 359, 47

Laird, E. S., Nandra, K., Hobbs, A., & Steidel, C. C. 2006, *MNRAS*, 373, 217

Lehmer, B. D., Brandt, W. N., Alexander, D. M., et al. 2005, *AJ*, 129, 1

Lehmer, B. D., Brandt, W. N., Alexander, D. M., et al. 2007, *ApJ*, 657, 681

Leitner, S. N., & Kravtsov, A. V. 2011, *ApJ*, 734, 48

Li, Z., Wang, Q. D., Irwin, J. A., & Chaves, T. 2006, *MNRAS*, 371, 147

Mathews, W. G., & Brighenti, F. 2003, *ARA&A*, 41, 191

McCammon, D., & Sanders, W. T. 1984, *ApJ*, 287, 167

McGaugh, S. S., Schombert, J. M., de Blok, W. J. G., & Zagursky, M. J. 2010, *ApJL*, 708, 14

Memola, E., Trinchieri, G., Wolter, A., Focardi, P., & Kelm, B. 2009, *A&A*, 497, 359

Miller, B., Gallo, E., Treu, T., & Woo, J.-H. 2012, *ApJ*, 747, 57

Miller, M. J., & Bregman, J. N. 2012, ApJ, submitted

Mineo, S., Gilfanov, M., & Sunyaev, R. 2012, MNRAS, 419, 2095

Moster, B. P., Maccio, A. V., Somerville, R. S., Naab, T., & Cox, T. J. 2011, MNRAS, 415, 3750

Mulchaey, J. S., & Jeltema, T. E. 2010, ApJL, 715, 1

Mulchaey, J. S., & Zabludoff, A. I. 1999, ApJ, 514, 133

Nandra, K., Mushotzky, R. F., Arnaud, K., et al. 2002, ApJ, 576, 625

Nulsen, P. E. J., Stewart, G. C., & Fabian, A. C. 1984, MNRAS, 208, 185

O'Connell, R. W. 1999, ARA&A, 37, 603

O'Sullivan, E., Forbes, D. A., & Ponman, T. J. 2001, MNRAS, 328, 461

Owen, R. A., & Warwick, R. S. 2009, MNRAS, 394, 1741

Patil, A., & Fonnesbeck, C. J. 2010, J. Stat. Softw., 35, 1

Putman, M. E., Peek, J. E. G., & Joung, M. R. 2012, ARA&A, 50, 491

Rosner, R., & Tucker, W. H. 1989, ApJ, 338, 761

Rudie, G. C., Steidel, C. C., Trainor, R. F., et al. 2012, ApJ, 750, 67

Sarazin, C. L. 1986, RvMP, 58, 1

Sijacki, D., Vogelsberger, M., Kereš, D., Springel, V., & Hernquist, L. 2012, MNRAS, 424, 2999

Smith, R. K., Brickhouse, N. S., Liedahl, D. A., & Raymond, J. C. 2001, ApJL, 556, 91

Snowden, S. L., Egger, R., Finkbeiner, D. P., Freyberg, M. J., & Plucinsky, P. P. 1998, ApJ, 493, 715

Strickland, D. K., Heckman, T. M., Colbert, E. J. M., Hoopes, C. G., & Weaver, K. A. 2004, ApJS, 151, 193

Strickland, D. K., & Stevens, I. R. 2000, MNRAS, 314, 511

Sutherland, R. S., & Dopita, M. A. 1993, ApJS, 88, 253

Tripp, T. M., Savage, B. D., & Jenkins, E. B. 2000, ApJL, 534, 1

Tucker, W., & David, L. P. 1997, ApJ, 484, 602

Tüllmann, R., Pietsch, W., Rossa, J., Breitschwerdt, D., & Dettmar, R.-J. 2006, A&A, 448, 43

- Tumlinson, J., Thom, C., Werk, J. K., et al. 2011, *Sci*, 334, 948
- Voges, W., Aschenbach, B., Boller, T., et al. 1999, *A&A*, 349, 389
- Watson, C. R., Kochanek, C. S., Forman, W. R., et al. 2009, *ApJ*, 696, 2206
- White, S. D. M., & Frenk, C. S. 1991, *ApJ*, 379, 52
- White, S. D. M., & Rees, M. J. 1978, *MNRAS*, 183, 341
- Yamasaki, N. Y., Sato, K., Mitsuishi, I., & Ohashi, T. 2009, *PASJ*, 61, 291

CHAPTER 6

Modeling X-ray Emission Around Galaxies

Note: This chapter is adapted from a paper of the same title, which is currently in preparation for submission to the *Astrophysical Journal*. The reference is Anderson, M. E. and Bregman, J. N. 2013, *ApJ* in preparation.

6.1 Abstract

Extended X-ray emission can be studied either spatially (through its surface brightness profile) or spectrally (by analyzing the spectrum at various locations in the field). Both techniques have advantages and disadvantages, and when the emission becomes particularly faint and/or extended, the two methods can disagree. Here we describe a flexible new method to perform spatial analysis, which can be extended to leverage spectral information simultaneously. We construct a model for the entire X-ray image in a given energy band, and generate a likelihood function to compare the model to the data. Employing either maximum likelihood or Markov Chain Monte Carlo, we can derive probability distribution functions for the source and background parameters together. We calibrate and demonstrate this method against a variety of simulated images, and then apply it to a Chandra observation of the hot gaseous halo around the elliptical galaxy NGC 720. We are able to follow the X-ray emission to a tenth of the background, and to derive much stronger constraints on the surface brightness profile than previous studies, which employed the spectral method. Extrapolating to the virial radius, we estimate that this galaxy has a total hot halo mass of less than $1 \times 10^{11} M_{\odot}$. This gas mass is slightly less than the stellar mass of the galaxy, and combined with the stars still leaves over half of this galaxy's baryons

missing.

6.2 Introduction

The study of very extended emission comprises a large portion of the work of extragalactic X-ray astronomy. All galaxy clusters (Gursky and Schwartz 1977, Forman and Jones 1982, Rosati et al. 2002) and most galaxy groups (Mulchaey et al. 1993, Ebeling et al. 1994, Ponman et al. 1996, Mulchaey 2000) are suffused with a hot ($kT > 10^6$ K), X-ray emitting gaseous medium. In all clusters and many groups, this medium contains the majority of the baryons in the system (Ettori 2003, Gonzalez et al. 2007, Giodini et al. 2009, Andreon 2010, Dai et al. 2010, Sanderson et al. 2013), and extends to hundreds of kpc. In recent years, X-ray observations have even been able to push outwards to the virial radius of some nearby clusters (George et al. 2009, Reiprich et al. 2009, Bautz et al. 2009, Kawaharada et al. 2010, Hoshino et al. 2010, Simionescu et al. 2011, Akamatsu et al. 2011, Walker et al. 2012ab, Morandi et al. 2012, Sato et al. 2012, Bonamente et al. 2012, Ichikawa et al. 2013, Walker et al. 2013).

Individual galaxies are also surrounded by extended X-ray emitting halos. Around elliptical galaxies, these hot gaseous halos have been studied for decades (e.g. Forman et al. 1979, Forman et al. 1985, Mathews and Brighenti 2003, O’Sullivan et al. 2001, Mulchaey & Jeltema 2010). Starbursting spirals have extended coronae above and below the disk extending to a few tens of kpc (Strickland et al. 2004, Li et al. 2006, Tüllmann et al. 2006, Owen & Warwick 2009, Yamasaki et al. 2009, Li et al. 2013). We also recently reported the detection of hot gaseous halos around more quiescent massive spiral galaxies, extending out to ~ 50 kpc (Anderson and Bregman 2011, Dai et al. 2012, Anderson et al. 2013); Bogdán et al. (2013) have confirmed one of these detections and have discovered another hot halo as well.

As both of these fields continue to detect emission at larger radii and lower X-ray surface brightness, it is becoming increasingly important to have effective observational techniques for studying faint, extended X-ray emission. At present, there are two major approaches to this analysis: spectral fitting and spatial binning. Spatial

binning is conceptually simple: one measures the X-ray radial surface brightness profile in a given band (optionally deprojecting the emission, if desired), and infers a gas density profile from the surface brightness profile. The major uncertainties with this method are flat-fielding the image and estimating the background. For bright sources, blank-sky backgrounds are sufficient, but for faint emission the background should be estimated in-field, which requires accurate flat-fielding. A secondary concern is separating various components of emission, if they exist; examples of this separation are illustrated in Anderson et al. (2011), Dai et al. (2012), and Bogdán et al. (2012).

For the spectral method, one studies the X-ray spectrum in radial bins instead of the surface brightness profile. Analyzing a spectrum requires more photons than measuring a broad-band surface brightness, in part because most realistic spectra have many more free parameters than a surface brightness profile. Thus, the image is usually broken into large bins, which sacrifices some location information. On the other hand, the various instrumental and background components are included in the spectral model, so in theory this method does not require separate flat-fielding or background subtraction. Also additional source components can be included, obviating concerns about confusion between hot gas emission and other X-ray sources such as X-ray binaries or background point sources. The primary downsides to this method are model specification and the need for more photons.

The aforementioned issues with these methods can be quite important, and can lead to conflicting results. For the spatial method, one major failure mode in flat-fielding is incorrect estimation of the “vignetting” - the decrease in sensitivity of the detectors as off-axis angle increases. The vignetting profile varies as a function of time (typically getting worse as the telescope degrades), and must be computed separately for each observation. It also varies as a function of energy, so a given exposure map (which contains the information about the vignetting) should only be computed for and applied to a narrow energy band over which the vignetting effects do not vary significantly. An example of this difficulty occurs in Pedersen et al. (2006), where the authors found evidence for a hot halo around the giant spiral NGC 5746 at 4σ significance. Later, after applying an updated calibration file which accounted more

correctly for the time-dependent degradation of the instruments, the vignetting profile changed and the signal disappeared (Rasmussen et al. 2009).

For the spectral method, model specification is particularly important. In the observations of interest today, the signal from hot gas is lower than the background, so the model for the background components in the spectrum can significantly influence the inferred properties of the signal. It is not trivial to construct a model for the X-ray background, and most studies use slightly different prescriptions. The components generally include the extragalactic AGN background at hard energies, and Solar-scattered X-rays (time variable), emission from the Local Hot Bubble, and the Galactic hot halo at softer energies. These components are all variable (spatially and sometimes temporally) and so their normalizations are not typically known a priori. Moreover, the emission itself can be quite spectrally complex - especially for the Local Bubble and Solar-scattered X-rays, where charge exchange can affect the signal and the gas is not necessarily in collisional equilibrium. Without sufficient photons to fit all of these spectral components, this method is very susceptible to systematic errors, either from degeneracies between free model parameters, or as a result of overfitting the data with an inappropriate model.

An example of this issue can be seen with the isolated elliptical galaxy NGC 720. This galaxy has a massive hot halo that has been observed once with Suzaku, twice with XMM-Newton, and several times with Chandra. One of the Chandra observations was studied by Humphrey et al. (2006), who used the spectral method in eight annuli to measure the hot gas density profile out to 90 kpc; extrapolating their density profile to 300 kpc yields a hot halo mass of $1 \times 10^{11} M_{\odot}$, which implies that the galaxy is missing about 1/2 of its expected baryons. The rest of the data is analyzed in Humphrey et al. (2011), with slight changes in their model, yielding a hot halo mass of $3 \times 10^{11} M_{\odot}$ which implies that the galaxy is baryon complete. Based on the statistical errors quoted in these papers, this is about a 3σ discrepancy in the mass at the virial radius. About half the discrepancy is caused by the addition of the Suzaku data, and half by the changes in their modeling; however, at a smaller radius like 100 kpc, the discrepancy is still nearly 3σ , and this difference is caused

almost entirely by their modeling. We will discuss this galaxy in much more detail in section 6.7.

Finally, there are also situations where the spectral method and the spatial method yield different conclusions. A notable example is in the estimation of galaxy cluster density profiles near the virial radius - where the cluster emission is much fainter and the systematic uncertainties in these methods become more important. There has been significant debate over the putative detection of a flattening in the radial decrease of the hot gas density profile, and this debate seems largely to fall along the lines between these two methods. The flattening was first observed with Suzaku, using spectral methods (George et al. 2009, Simionescu et al. 2011); spatial methods, using ROSAT, have generally not confirmed this result (Ettori and Balestra 2009, Eckert et al. 2011, Eckert et al. 2012). Thus, while there seems to be general consensus about a number of other features in the gas properties near the virial radius (such as decreasing temperature and flattening entropy), more work needs to be done to understand the behavior of the gas surface brightness and density (and therefore the total gas mass and baryon fraction of the cluster). Discrepancies at the 5%-15% level have also been noted by Rozo et al. (2012) in samples of clusters analyzed with different techniques, for derived parameters such as L_X , T_X , and Y_X .

In this paper we explore new ways to improve the spatial method, incorporating more information about the spatial distribution of the X-ray emission as well as building a foundation to include more spectral information as well. We construct a likelihood function for the event file, and simultaneously fit the background and the source emission across the entire image. While we restrict our analysis in this paper to a single energy band (0.5-2.0 keV), we note that the likelihood function is applied to the events file, and can therefore naturally and straightforwardly be extended to include spectral information, event by event, as well as spatial information. Thus the method holds the promise of combining the best features of both the spatial and the spectral analysis, while minimizing the systematic errors associated with each.

We begin by discussing the model and motivating the various components we include within the likelihood function. We then present simulated images generated

from the model. We compare the simulated images to real images, and detail the reduction of the real images used for the comparison. We then discuss the likelihood function, and compare several different forms for the likelihood function before settling on one function. We show that this likelihood function is able to recover the input parameters in a variety of simulated images. Finally, we apply the method to the case of NGC 720.

6.3 The model

Our model for a typical X-ray image consists of three major components: the background, the source emission, and X-ray point sources. We discuss each component in turn below.

6.3.1 Background

In an infinitely long X-ray observation, most of the hard cosmic X-ray background could be resolved into individual AGNs (Comastri et al. 1995, Ueda et al. 2003, Bauer et al. 2004, Hickox and Markevitch 2007 ab). Below about 2 keV, star-forming galaxies begin to contribute background X-ray point sources as well, at about the 10%-20% level. There is also a diffuse X-ray component to the soft emission, due to the local hot ISM (i.e. the Local Hot Bubble) and the Galactic hot halo (Snowden et al. 1990, Snowden et al. 1998). These various components are all spatially variable across the sky, but in practice this is not as much of a concern as one might expect. The faint point sources, which are not resolved in typical X-ray observations, are so numerous as to wash out most statistical fluctuations in a typical field of view. The bright point sources are treated separately in our model (section 6.2.3). And the diffuse soft emission typically shows features on large (i.e. degree) scales, which are not very important in a single field of view, although it is still unclear to what extent the diffuse emission also contains substructure on smaller scales (Soltan et al. 2005, Galeazzi et al. 2009). Because all of these sources are approximately uniform across a Chandra or XMM field of view, we parameterize these backgrounds as a single uniform component in the image, with free normalization. All of these backgrounds

are focused by the telescope’s mirrors, which introduces a vignetting effect across the image: emission that falls at larger radii from the aimpoint is slightly attenuated compared to emission right on the aimpoint (the magnitude of the attenuation is usually roughly a quarter to a third in the soft band).

There is another type of background which is not focused by the telescope optics. The particle background – mostly soft protons from the Solar Wind – impinges directly on the X-ray detectors, and also triggers X-ray fluorescence from the instruments themselves. This background also registers as events on the detectors. Many of the cosmic rays are automatically flagged by the standard data reduction scripts, and during periods of especially high particle flux the entire event file is generally excluded as well. However, some particles cannot be automatically distinguished from X-rays and remain in the image as false positives; instrumental X-rays are also not automatically removed and will appear in the events file. We model this background as a separate uniform component, but unlike the cosmic X-ray background this component is unvignetted. Therefore the total background is the sum of a uniform vignettted component and a uniform unvignetted component, each of which has one free parameter (its normalization).

6.3.2 Source

The source emission can be parameterized in a number of ways, depending on the details of the source in question. For a hot gaseous halo, we have typically opted for the β model (Cavaliere and Fusco-Femiano 1976), which describes an isothermal, azimuthally symmetric gas distribution. In this model, the gas has an assumed density distribution

$$\rho(r) \equiv \rho_0 \left[1 + \left(\frac{r}{r_0} \right)^2 \right]^{-\frac{3}{2}\beta}$$

and the projected surface brightness profile of the X-ray emission as a function of projected radius r is

$$S_x(r) \equiv S_0 \left[1 + \left(\frac{r}{r_0} \right)^2 \right]^{\frac{1}{2} - 3\beta}$$

This profile is standard for describing the observed X-ray emission from gas in galaxy clusters and in hot halos around individual galaxies. It generally gives a good fit to the data and can be extended if necessary to account for radial temperature gradients, although for hot halos these gradients tend not to be very large. Another possible class of density profiles are the adiabatic profiles (Maller and Bullock 2004, Fang et al. 2013), which assume the gas is adiabatic instead of isothermal. This tends to produce a profile with a much flatter slope, which yields much lower surface brightness and is therefore difficult to constrain observationally. These profiles are employed less frequently than beta models, but for galactic hot halos they are still consistent with observations, and will be considered in future work.

It is straightforward to include other sources in the model as well. For galactic hot halos, emission from X-ray binaries is an important confounding factor, and when observing real hot halos we include a model for this emission, illustrated in section 6.7.

The Chandra psf

The source emission is convolved with the instrumental point spread function (psf), and for extended emission when we are trying to measure a surface brightness profile this can be an important effect. The Chandra psf is much smaller than the XMM-Newton psf or the Suzaku psf, subtending less than an arcsecond at the aimpoint, but it grows to several times this size at larger off-axis angles. In order to convolve our source emission model with the psf, we therefore need a functional form for the psf which is defined at every location on the detector. For this purpose, we specify a two-dimensional, azimuthally symmetric Gaussian at each point, so that the psf can be parametrized as a function of its standard deviation σ at each point. We estimate σ using the `mkpsfmap` function, which estimates the radius corresponding to a given




Figure 6.1 Ratios of the radii corresponding to the (a) 75% ecf, (b) 90% ecf, and (c) 95% ecf, each divided by the radius corresponding to the 39.3% ecf. This entire ratio is then divided by the expected value of the ratio if the psf were perfectly Gaussian; Thus the plots illustrate the amount of nongaussianity in the psf at the given ecf. There is little nongaussianity in the psf until radii corresponding to the 95% ecf, suggesting that a Gaussian is a good approximation to the psf to about the 5% level. There is a region in the center of the ACIS-I array where the psf is more strongly non-Gaussian, but the size of the psf here is small enough that deviations do not effect the image substantially.

ecf (encircled energy fraction) at every location across an image¹. We set the ecf to 0.393 so that the resulting radius is equal to σ .

The true psf is not exactly Gaussian. It is close to Gaussian in the core, but has wings that are more extended than a Gaussian. We can check how important these wings are by running `mkpsfmap` with larger ecfs. We compare the radii corresponding to higher ecfs to the expected radii if the emission were Gaussian with σ equal to the radius enclosing 39.3% of the emission. The results are below, in Figure 6.1, for ecfs of 75%, 90%, and 95%.

Overall the psf does seem very close to Gaussian, with deviations becoming more visible at higher ecfs as the extended wings start to dominate. The deviations are not very significant across most of the detector even at an ecf of 95%, suggesting that a Gaussian is a good approximation to about the 5% level. The psf is more non-Gaussian at the center of the detector, and a roughly 1-arcminute region near the very center is more extended than Gaussian even in the 75% ecf comparison. Fortunately, the size of the psf in this region is still less than an arcsecond, so the

¹The function `mkpsfmap` does neglect aspect variations over the image, which can also broaden the effective psf, but these variations are uniform, random, and on the order of arcseconds, so they are not a major source of concern for this sort of analysis.

absolute magnitude of these nongaussianities is small, and we neglect them here.

We convolve our source emission model with the psf model to generate the total image model². We also multiply the image model by the vignetting profile obtained from the exposure map. We do not convolve the psf model with the background, since the background is assumed to be uniform, and we do not convolve the psf model with the point sources described in the next subsection, since the point sources are computed with a more exact method, and so the approximation of Gaussianity is not needed.

6.3.3 Point sources

It is necessary to include a correction for bright point sources in our model. While we mask out bright point sources (section 6.3.2) before analyzing the images, we want to be sure the standard image masking procedures do not introduce a bias. We might expect a bias towards higher inferred fractions of unvignetted background, because the size of the psf increases with off-axis angle. Therefore a point source with the constant flux will decrease in surface brightness as its off-axis angle increases, which introduces a radial gradient in point source detectability. We therefore add another component to our model to account for undetected point sources. To estimate this component, we first estimate the number of point sources in the image as a function of their count rate. We use the point source number counts of Bauer et al. (2004) for this estimate, parameterizing the number counts from their Figure 4 as

$$N(> S) = \begin{cases} 3039 \left(\frac{S}{10^{-16} \text{ erg s}^{-1}} \right)^{-0.55} & , S < 2 \times 10^{-15} \\ 585 \left(\frac{S}{2 \times 10^{-15} \text{ erg s}^{-1}} \right)^{-1.0} & , 2 \times 10^{-15} \leq S < 1 \times 10^{-14} \\ 117 \left(\frac{S}{10^{-14} \text{ erg s}^{-1}} \right)^{-1.5} & , S \geq 1 \times 10^{-14} \end{cases}$$

where $N(> S)$ is in units of deg^{-2} and S is in $\text{erg s}^{-1} \text{ cm}^{-2}$ over the 0.5-2.0 keV energy range. From this function and the length of a given Chandra observation, we

²Technically, we take the fast Fourier transform of the image and of the model for the spatially varying psf, multiply the two together, and then take the inverse Fourier transform. This is mathematically very similar to a convolution (identical in the limit of infinite spatial resolution), and is much quicker to compute.

estimate the expected number of point sources in the image as a function of counts. We generate a psf map using `mkpsfmap` (using the 90% ecf as the characteristic psf radius). For every integer number of counts in a given point source, we use the psf map to determine where in the image a point source with that number of counts could be detected above a given probability threshold (which we set to the value used by `wavdetect` - 1×10^{-6}). We distribute the expected counts from undetected point sources over the area in which the point source would not be detected, and add together all the point sources expected in the image to get a map of the expected number of counts from undetected point sources at any given location in the field. An example of such a map is shown in Figure 6.2.

However, there are not enough bright point sources in a real observation to distribute their counts smoothly across every pixel, so we also explore the effect of binning the counts from regions with significant undetected point sources together (section 6.5). This offers a modest improvement if performed correctly.

6.4 Simulating images

Now that we have a model for all the X-ray emission in an image, we can generate simulated images from the model. Our fiducial model has five free parameters - two normalizations for the vignettted and unvignettted backgrounds, and the three parameters that define our source emission (β , r_0 , and S_0). We also need to specify shapes for the vignettted and unvignettted backgrounds, although these are completely determined by the details of the observation and contain no free parameters. In practice, we also fix the average flux per pixel to the measured value in a given observation, which removes one of the free parameters from the normalizations of the two backgrounds (since their sum is fixed). We therefore instead parameterize the backgrounds in terms of f_{unvig} , the fraction of the diffuse background counts that we attribute to the unvignettted component.

We assume the vignettted background should be proportional to the exposure map generated with `mkexpmap` (either in units of $\text{cm}^2 \text{ count photon}^{-1}$ or $\text{cm}^2 \text{ s count photon}^{-1}$, although we use the latter). For a given observation, `mkexpmap` computes

the effective exposure time and/or area across the image, so we can use its output as a vignetting profile for the observation. The inputs are an aspect histogram and an instrument map file. The aspect histogram encodes the movement of the telescope during the observation, and the instrument map encodes the quantum efficiency and detector response as a function of energy for the specified source spectrum (see the CXC helpfile or Davis et al. 2001 for more details).

When making the exposure map, care must be taken that both the assumed source spectrum and the instrumental response do not vary significantly across the energy band. This should be checked for each observation, but is generally a safe assumption for fairly narrow wavebands. Note that in the future we can extend this method at this stage to include exposure maps at a variety of energies, instead of an average exposure map over a given energy band. We used the automated `fluximage` script to create an exposure map, with the non-standard setting of `expmapthresh = 67%` instead of the more common value of 2%. This zeros out every pixel with effective exposure time less than 67% of the peak value. The vignetting across the ACIS-I chips is typically fairly modest, so this has little effect except near the edges of the chips. However, it has the advantage of removing from our analysis the chip gaps, where the exposure map is the most uncertain. We can then define the shape of the unvignetted background to be uniform across the image wherever the exposure map is non-zero.

We can now simulate an X-ray image, by multiplying the vignettted background by the vignetting profile specified by the exposure map, multiplying the unvignetted profile by the area of the image specified by the exposure map, and adding a source component to the image as desired. The final step in order to get a realistic-looking image, however, is to include bright background point sources.

6.4.1 Simulating Point Sources

The shape of the Chandra psf varies as a function of location on the detector and of energy, so it is not trivial to generate a simulated point source, but this process has been automated with the Chandra Ray Tracer (ChaRT; Carter et al. 2003) and

the Model of AXAF Response to X-rays (MARX; <http://space.mit.edu/cxc/marx/>) software. We used ChaRT to construct a library of psf shapes across the ACIS field of view. We generated simulated psfs in polar coordinates, recording a psf at every arcminute of off-axis angle (also sampling every half arcminute near the aimpoint) and every 30° of position angle. We used a $\Gamma = 2$ powerlaw for the spectral shape of the point sources and a wide energy band (0.5-8.0 keV). We then passed the simulated psfs to MARX, and traced 10^4 photons through the telescope optics in order to construct an events file for a point source at each location. We repeated this process for both the ACIS-I and ACIS-S configuration, as described in the MARX help files, in order to be able to simulate point sources on the ACIS-S chips in an ACIS-I observation, or vice versa.

Armed with this library of events files, we can simulate a point source at any location in an observation. We find the nearest point source events file from our library, center it on the location we want to simulate a point source, and draw events from the events file at random in order to build up the desired count rate. We use the point source number counts function estimated from Bauer et al. (2004) listed in section 6.2.3 (equation 6.1) to estimate the number and count rate of point sources to generate for a given image.

6.4.2 Detecting Point Sources

To detect bright point sources for masking, we run `wavdetect` on the simulated images. We set the wavelet radii to 1,2,4,8,and 16 pixels (with the image binning at 2, so one pixel is 0.984 arcseconds on a side), and the significance threshold to 10^{-6} , corresponding to approximately one false positive per chip. The one non-standard setting we use with `wavdetect` is to increase the size of the ellipse to mask (*ellsigma*) from 3 to 8. This excludes nearly all of the emission from the point source, as well as a larger region around the point source than is typical. Our goal is to measure the background precisely, not to capture all of the background photons, so it is an acceptable tradeoff to lose some background photons in order to exclude as many of the photons from the bright point sources as possible.

During this process, we noticed that `wavdetect` was failing to detect some point sources at large off-axis angles (especially on the S2 chip). In order to study this effect further, we also tried an alternative method of detecting point sources. For every photon on the detector, we computed the 90% encircled energy radius at the location of that event (using `mkpsfmap` and assuming the same $\Gamma = 2$ powerlaw). We compared the number of counts within this radius to the estimated background (estimated either assuming a 100% unvignetted fraction or an 100% vignetted fraction; both methods give very similar results). If any circle has an excess of photons significant at 10^{-6} or stronger, we masked that circle.

This method is much more computationally intensive than `wavdetect`, since it has to perform aperture photometry tens of thousands of times (once per photon). It also has no ability to find the centroid of a point source; instead it just masks out any photons that could conceivably be associated with the source. While these constraints are not ideal for a robust multi-purpose point source detection algorithm, for our purposes we found this method to be more robust than `wavdetect`. We illustrate the two methods side-by-side in the next section (section 6.4).

The remaining image is composed of a vignetted background, an unvignetted background, a few dozen bright simulated point sources which have been masked out, and dozens of faint simulated point sources which were not detected by `wavdetect`.

6.5 Example: Chandra Deep Field South

In this section we illustrate the process of generating a simulated image, and show how the vignetted and unvignetted backgrounds can be distinguished. We use as our reference image Chandra observation 8595. This is a 115 ks ACIS-I observation of the Chandra Deep Field South (CDF-S) taken in late 2007.

We download the events file and apply a standard reduction in order to process the data for analysis. We use CIAO version 4.5 and CALDB version 4.5.6. Most of the reduction is performed by the new `chandra_repro` and `fluximage` routines in CIAO. `chandra_repro` generates a level 2 event file from the raw data, applying standard filters for bad pixels and for cosmic rays. For this analysis the default settings are

acceptable; in particular we adopt the default value for `check_vf_pha = False` since this routine can remove source events in addition to lowering the background; we prefer to keep all the source events and account for the background in our model.

We examine the lightcurve for any periods of flaring, but there seems to be no significant flaring in this observation.

We then run `fluximage` on the cleaned level 2 event file. As discussed above, this script requires some caution in the choice of energy band and assumed source spectrum. We use the 0.5-2.0 energy band, which covers the region of highest effective area for the Chandra ACIS instrument, with significant declines in sensitivity at both higher and lower energies. In addition, the background increases significantly at lower energies. Finally, for $\lesssim 1$ keV hot gas, most of the emission is expected to fall within this band. For the source spectrum, we specify an absorbed powerlaw, with Galactic absorption and a powerlaw slope of 1.56. We set the image binning scale to 2, so one pixel is 0.984 arcseconds on a side.

Figure 6.2a shows the image generated by `fluximage`. Figure 6.2b shows the exposure map, before applying the threshold at 67% of the maximum, and Figure 6.2c shows the exposure map after applying the threshold to exclude locations on the image with low exposure. As mentioned earlier, we are much more liberal in excluding pixels than the `fluximage` default, because our goal is to measure the background in regions where the exposure map is best-calibrated, not to attempt to recover every single background photon.

We then run `wavdetect` on the image. We mask out the 8σ ellipses around each point source, again eschewing total background counts in favor of having an uncontaminated image. In Figure 6.3, we show the results from `wavdetect` with our parameters, and compare them to the CDF-S 4 MS catalog of point sources (Xue et al. 2011). We do not recover all these point sources in the 110 ks image from observation 8595, which is unsurprising, but we do recover the bright point sources. There is one false positive in the image as well. Note that false positives are not a significant problem for us: they cost us a few background photons that are unnecessarily masked out, and probably bias downwards our estimate of the background flux rate slightly,




Figure 6.2 Illustrations of various steps in our analysis pipeline, using Chandra obsid 8595 (the CDF-S) as our example. In (a) we show the 0.5-2.0 keV image generated with *fluximage*. In (b) we show the exposure map generated with *mkexpmap*, and in (c) we show the exposure map after applying a much more stringent threshold in order to exclude the region with lower exposure.

but this effect is negligible for the low number of false positives generated by our significance threshold..

As a comparison, in Figure 6.4 we also illustrate the point sources detected in this image using our manual point source detection algorithm (section 6.3.2). This method measures the likelihood of obtaining the observed number of counts within a radius of the 90% enclosed counts fraction (ecf) of the psf, at the location of every photon on the image. It is therefore more computationally intensive than *wavdetect*. It also detects slightly fewer point sources overall, but at large off-axis angles it outperforms *wavdetect*; this is especially notable on the S2 chip where *wavdetect* misses a fairly bright and extended source. None of the sources detected using the manual method are false positives. We compare the effects of the two different methods on recovering the correct unvignetted fraction of the background in section 6.5.

After masking the bright point sources, we can compute the average remaining flux per pixel. We also now have models for both the vignetted and the unvignetted backgrounds. The vignetted background is just proportional to the exposure map, and the unvignetted background is uniform across the exposure map, except for




Figure 6.3 Comparison of point sources in observation 8595 detected with *wvdetect* (green; note our unusually large ellipses) to all the known point sources in the field (red; for these sources the size of the circles is arbitrary) from the 4 MS analysis (Xue et al. 2011). We detect most of the bright point sources, and have one false positive (slightly to the left of the center of the ACIS-I array).




Figure 6.4 Comparison of point sources detected with *wavydetect* (green; note our unusually large ellipses) and our manual point source detection algorithm (red). The two methods generally agree with one another, although each catches a few point sources missed by the other, and the manual method performs better on the ACIS-S chip very far off-axis.




Figure 6.5 Simulated images with mock point sources added and then masked with *wavdetect*. Image (a) has 100% of its diffuse background in the vignetted component, and (b) has 100% of its diffuse background in the unvignetted component. The difference is subtle, but the images can be distinguished by eye at large off-axis angles, and as we show they can also be distinguished numerically.

the areas of the exposure map that have been masked out. In Figure 6.5 we show simulated images created using the exposure map generated above, after masking the simulated point sources in each image. Figure 6.5a has 100% of its photons in the vignetted background, and Figure 6.5b has 100% of its photons in the unvignetted background.

Now all we need to do is to differentiate between these images in order to estimate the fraction of the background in the real observation belonging to the vignetted and the unvignetted components. We compare images using a maximum likelihood approach. We compute the likelihood of obtaining the observed image given an arbitrary model image, and find the fraction of unvignetted background that returns the maximum likelihood. Before we can run the maximum likelihood analysis, however, we need a likelihood function.

6.6 The Likelihood Function

We explored a number of different likelihood functions for comparing observed and simulated images. The natural choice, which is of the best, is the pixel-by-pixel function, defined as

$$L \equiv \prod_{\text{all pixels}} p(c_{\text{pix}}|m_{\text{pix}})$$

where $p(c|m)$ is the Poisson probability of obtaining the observed number of counts c in a pixel, given a model prediction m for that pixel. We can also bin the data in various ways

$$L \equiv \prod_{\text{all bins}} p(c_{\text{bin}}|m_{\text{bin}})$$

where now we are adding pixels together within a bin before evaluating the likelihood. We explored binning by off-axis angle and by the size of the psf, in various combinations. We also explored the effect of adding an additional component to the model to account for undetected point sources, which are more likely to be undetected at large radii. We discussed how to generate this component above, in section 6.2.3.

To examine which of these likelihood functions gives the most reliable estimator, we simulated 11 images based on the exposure map for observation 8595 described above. The 11 images had the same average flux per pixel as observation 8595 (i.e. 0.03536 counts pix^{-1} in the diffuse background), although f_{unvig} ranged from 0% to 100%, in steps of 10%. For each image, we computed the likelihood as a function of f_{unvig} , using each of our different likelihood functions.

We estimate uncertainties in the recovered value of the unvignetted fraction using the procedure discussed in Martin et al. (2008). As long as the likelihood is well-behaved near the maximum, then the $k\sigma$ confidence interval around the maximum is “bounded by the values that that correspond to the function $2\ln(L)$ dropping by k^2 .” So the $1 - \sigma$ region around the best-fit value of the unvignetted background fraction is just the region where the log-likelihood is within 0.61 of the maximum

value. We evaluate the effectiveness of recovering f_{unvig} for each likelihood function by calculating the χ^2 , i.e. $(\text{recovered } f_{\text{unvig}} - \text{true } f_{\text{unvig}})^2 / (\sigma f_{\text{unvig}} (\text{recovered}))$. This is not technically correct at the extreme values, since f_{unvig} cannot go below 0 or above 1 and the errors are therefore not Gaussian at the edges. However, we checked the subset of simulations with the extreme points excluded (values of f_{unvig} between 0.1 and 0.9) and obtained essentially the same results, so this effect does not appear to bias our results in any meaningful way.

In Figure 6.6 we present three examples of various likelihood functions. The plots show the inferred value of f_{unvig} , in terms of the value of f_{unvig} used to simulate each image. Figure 6.6a uses the first likelihood function (multiplying the likelihood in each pixel together to generate the total likelihood), without applying the correction for undetected point sources to the model (see section 6.4). The method gives an acceptable χ^2 for the manual point source detection method, a marginally acceptable χ^2 for point sources masked with either method, and unacceptable χ^2 for `wavdetect` and for no point source masking. Figure 6.6b shows the recovered values for f_{unvig} when we include the correction for undetected point sources in the model. Now the χ^2 is marginally acceptable for either point source masking method, and slightly better for the combination of both methods. Also, the average recovered values of f_{unvig} have decreased slightly, owing to the point source correction replacing some of what had been interpreted as unvignetted background at large radii in Figure 6.6a. Finally, Figure 6.6c shows our favored likelihood function. In this method, we bin the 10% of pixels with the largest psf into a single bin, and evaluate the likelihood within that bin. We repeat this for the next 10% of pixels with the next-to-largest psf, and once more for the subsequent 10% of pixels. For the remaining 70% of pixels, we evaluate the likelihood pixel by pixel. We then multiply all these likelihoods together to generate the final likelihood. We also include the correction for undetected point sources in the model for this version.

In Table 1 we list the values of χ^2 for each of the various likelihood functions we examined. The *pixel-by-pixel* likelihood function was defined above; we compute the likelihood in every pixel and multiply them all together to get the final likelihood. For




Figure 6.6 Comparison of a few example likelihood functions in recovering the shape of the background, parametrized by f_{unvig} (see text). 11 simulated images were generated and analyzed; the dashed line shows the true values of f_{unvig} across these 11 images. The black points are the recovered values if the point sources in the image are not masked. The red points are the recovered values if the point sources are masked with *wavdetect*, and the blue points are the recovered values if the point sources are masked with our manual method. Finally, the magenta points are the recovered values if the point sources are masked using the combination of *wavdetect* and our manual method. Plot (a) shows the *pixel-by-pixel* likelihood function, (b) shows the *pixel-by-pixel* likelihood function with the additional component to account for undetected point sources, and (c) shows the *pixel/psf* (*3 psf bins at 70%*) likelihood function, again with the additional component to account for undetected point sources. For reference, with 10 d.o.f., the fits from a likelihood function can be excluded at 95% confidence if $\chi^2 > 18.31$ and at 99% confidence if $\chi^2 > 23.21$.

psf bins likelihood functions, we bin pixels by the size of the psf at the location of that pixel (90% encircled energy fraction, or ecf). The *pixel / psf* likelihood functions are a hybrid of these two functions: we compute the likelihood pixel-by-pixel for the $N\%$ of pixels where the psf is smallest, and we bin the remaining $(100-N)\%$ pixels into M bins determined by the size of the psf within that pixel. For *uniform radial* bins, we binned the pixels in annuli centered around the aimpoint, with constant radius. The $r^{-1/2}$ *radial* bins are the same, except the radius of the annuli decreases as $r^{-1/2}$, such that the number of pixels within each bin is roughly constant. Finally, the *ACIS-I / ACIS-S radial* bins use $r^{-1/2}$ bins across the ACIS-I chips, and an additional 10 uniform bins across the ACIS-S chips, in order to sample the ACIS-S chip in more detail. We also examine each of these likelihood functions both *w/nondetect* and without (default) the correction for undetected point sources included in the model. For reference, with 10 d.o.f., the fits from a likelihood function can be excluded at 95% confidence if $\chi^2 > 18.31$ and at 99% confidence if $\chi^2 > 23.21$.

Table 6.1 Likelihood Function Comparison

Likelihood function	χ^2 (no masking)	χ^2 (manual method)	χ^2 (wavdetect)	χ^2 (manual & wavdetect)
pixel-by-pixel	37.4	16.1	33.1	20.7
pixel-by-pixel w/nondetect	71.4	21.9	21.3	18.9
10 psf bins	149.2	12.8	34.3	15.3
10 psf bins, w/nondetect	186.6	55.4	36.6	31.8
20 psf bins	48.1	18.8	51.3	26.0
20 psf bins, w/nondetect	79.5	33.7	19.9	15.3
40 psf bins	47.7	13.2	37.7	16.8
40 psf bins, w/nondetect	69.4	33.8	18.7	15.1
pixel/psf (1 psf bin at 90%)	> 999	> 999	> 999	> 999
pixel/psf (1 psf bin at 90%), w/nondetect	78.9	25.8	22.4	21.2
pixel/psf (2 psf bins at 80%)	> 999	> 999	> 999	> 999
pixel/psf (2 psf bins at 80%), w/nondetect	75.5	26.4	23.6	21.9
pixel/psf (3 psf bins at 70%)	> 999	> 999	> 999	> 999
pixel/psf (3 psf bins at 70%), w/nondetect	45.4	15.4	26.7	18.3
10 uniform radial bins	120.0	14.5	42.7	21.1
10 uniform radial bins, w/nondetect	89.1	51.4	24.5	28.0
20 uniform radial bins	62.9	16.0	41.5	19.2
20 uniform radial bins, w/nondetect	83.0	53.9	25.3	29.7
30 uniform radial bins	45.5	13.6	40.2	22.0
30 uniform radial bins, w/ps corr	82.7	52.5	23.8	19.9
10 $r^{-1/2}$ radial bins	145.5	13.9	32.9	15.6
10 $r^{-1/2}$ radial bins, w/nondetect	81.9	38.5	21.7	19.6
20 $r^{-1/2}$ radial bins	55.2	15.3	38.9	21.6
20 $r^{-1/2}$ radial bins, w/nondetect	83.0	54.8	34.6	31.2
30 $r^{-1/2}$ radial bins	50.9	13.7	37.4	17.6
30 $r^{-1/2}$ radial bins, w/nondetect	89.8	55.7	25.7	30.9
10 ACIS-I / ACIS-S radial bins	29.5	11.1	28.9	14.9
10 ACIS-I / ACIS-S radial bins w/nondetect	122.1	37.5	20.8	23.1
20 ACIS-I / ACIS-S radial bins	40.5	13.4	41.3	20.0
20 ACIS-I / ACIS-S radial bins, w/nondetect	81.7	53.1	25.2	29.6
30 ACIS-I / ACIS-S radial bins	47.4	12.4	33.5	16.1
30 ACIS-I / ACIS-S radial bins, w/nondetect	82.0	54.8	34.5	29.6

List of χ^2 goodness of fit statistics for recovered values of f_{unvig} using different likelihood functions and point source detection algorithms. Formally, the fits from a likelihood function can be excluded at 95% confidence if $\chi^2 > 18.31$ and at 99% confidence if $\chi^2 > 23.21$.

The wide variation in χ^2 shows that the choice of likelihood function is important. However, the recovered values of f_{unvig} between the different methods are not very different at all; the range in χ^2 values primarily stems from differences in the estimated size of the $1 - \sigma$ uncertainties. Thus, while the method with the absolute minimum χ^2 is the *10 ACIS-I / ACIS-S radial bins* likelihood function applied to images that have been cleaned with the manual point source detection method, this is not our favored likelihood function. The $1 - \sigma$ uncertainties in its fits to f_{unvig} are about twice as large with this method as they are with the pixel-by-pixel methods (i.e 0.1 instead of 0.05), which is why the χ^2 is so low, but it also limits the statistical constraining power for this technique.

The one difference in likelihood functions that does affect the recovered values of f_{unvig} instead of the uncertainties on these values is the *w/nondetect* component in the model. This is the component that corrects for undetected point sources. It is a small component (comprising about 3% of the total counts in the model), but it increases with off-axis angle, and therefore works against the vignettted component in the model, which decreases with off-axis angle. Thus, its inclusion in the model tends to also increase the recovered vignettted fraction of the emission, leading to lower values of f_{unvig} . This generally improves the χ^2 , but not in every case. However, without including this component, we find a systematic bias such that the recovered value of f_{unvig} is too high. This is visible in Figure 6.6, for example. Thus, while the *pixel-by-pixel* method applied to the manually cleaned images actually has a lower χ^2 when the *w/nondetect* component is excluded from the model, we prefer to keep this component in our model to prevent our results from being biased towards high values of f_{unvig} .

In general methods with increased binning have larger uncertainties (because information is lost during the binning process), but methods with less binning are more sensitive to failures caused by failing to mask a point source (since there is less binning to average out the effect of a point source). The method that offers the best compromise between these two tradeoffs seems to be the *pixel/psf (3 psf bins at 70%)*, *w/nondetect* likelihood function applied to manually filtered images. We also

consider this likelihood function applied to images filtered with `wavdetect`, as well as the *pixel-by-pixel* likelihood function applied to images filtered with both point source detection methods, in section 6.5.1.

6.6.1 Measuring f_{unvig} for observation 8595

Finally, we present the results for observation 8595 using these various likelihood functions. In addition to running `wavdetect` or the manual point source detection algorithm, we also masked out about half of the ACIS-S chip due to an extended source. In Table 2, we present the value of f_{unvig} , along with 1σ uncertainties, for the best few likelihood functions as determined above. The inferred values of f_{unvig} are systematically higher when the point source correction is not applied, and the results are also less precise. After applying the point source correction, the recovered values are identical within uncertainties for all four methods.

Table 6.2 Measurements of f_{unvig} with various likelihood functions

Likelihood Function	without nondetect	with nondetect
<i>pixel-by-pixel</i> , manual point source detection	0.56 ± 0.05	$0.53^{+0.06}_{-0.05}$
<i>pixel-by-pixel</i> , <code>wavdetect</code>	0.60 ± 0.05	0.58 ± 0.05
<i>pixel-by-pixel</i> , both methods	0.56 ± 0.05	0.54 ± 0.05
<i>pixel/psf</i> (3 psf bins at 70%), manual point source detection	0.65 ± 0.01	$0.55^{+0.06}_{-0.05}$
<i>pixel/psf</i> (3 psf bins at 70%), <code>wavdetect</code>	0.86 ± 0.01	$0.60^{+0.05}_{-0.06}$
<i>pixel/psf</i> (3 psf bins at 70%), both methods	0.80 ± 0.01	$0.56^{+0.05}_{-0.06}$

6.7 Simulated Images with an Extended Source

Now we demonstrate the effectiveness of this technique for measuring extended emission. In this section we generate simulated images, just like in section 6.4, but now we also add an extended source to the model (see section 6.2.2). We then mask out the bright point sources using `wavdetect`, and attempt to recover the input parameters. For extended emission, we are fitting simultaneously for the three components of the β -model (i.e. β , r_0 , and S_0), as well as the shape of the background (f_{unvig}). This means that the normalization of the background is no longer fixed. Rather, we fix the total number of counts in any model image to be the same as the number of counts in the simulated image, and so the normalization of the background is determined by

the difference between the total counts in the image and the counts attributed to the extended source.

Because we have increased the dimensionality of the problem from one dimension to four dimensions, running the maximum likelihood analysis has become much more cumbersome. We found that a fixed grid search was too inefficient for exploring this space, so we opted for Markov Chain Monte Carlo (MCMC) instead. An MCMC search is a Bayesian technique designed to efficiently sample likelihoods in multidimensional space. For long enough chains, the distribution of samples in multidimensional space is supposed to approximate the posterior probability distribution for the fit parameters. Obviously, since it is a Bayesian technique, MCMC requires specification of priors for each fit parameter as well.

In this analysis, we use the open-source Python package `emcee` (Foreman-Mackey et al. 2012), which implements an affine-invariant MCMC method (Goodman and Weare 2010). This implementation of MCMC uses a user-specified number of independent “walkers” which separately explore the parameter space. The prior sets a uniform region, from within which each walker is randomly assigned an initial position. But as long as the chain is run for a sufficiently long time, and the likelihood function is well-behaved, the walkers can explore the entire space and the result is nearly completely insensitive to the choice of prior.

6.7.1 Source at Large Off-Axis Angle

We first simulate an image of an extended source that approximates a fairly distant ($d \sim 50$ Mpc) hot halo. We start by placing the source fairly far off-axis, and we specify a β -model with $\beta = 0.5$ and $r_0 = 10$ pixels (9.8 arcseconds). We use the same chip layout, vignetting profile, and background rate as observation 8595 studied above, and we set $f_{\text{unvig}}=0.4$, which seems to be a fairly typical value. We simulate an observation of the same length as observation 8595 (110 ks), so we use the same average background as that observation (0.03536 counts pixel⁻¹). The extended source is normalized to have 2000 source counts within an arbitrary large radius of 400 pixels (6.5 arcminutes), which corresponds to a source flux of about 0.05 count




Figure 6.7 Simulated images with mock point sources added at random, and an extended source placed near the bottom of the ACIS-I array that contains 4000 source counts. Image (a) is the raw image, (b) shows the same image after running *wavydetect* to detect the background point sources.

s^{-1} – fairly faint for extended emission. There are about 38400 background counts as well. We also add simulated point sources to the image according to the same procedure described above in section 6.2. We run *wavydetect* in order to mask out the bright point sources. We delete any detections of point sources within $40''$ of the extended source, which serves to prevent the extended source itself from being masked. The simulated image, before and after masking out bright point sources, is displayed in Figure 6.7.

We analyze the simulated image with *emcee*, as discussed above. We use 100 independent walkers, and for each walker we discard the first 100 elements of the chain, before keeping the next 100 for analysis. The autocorrelation length is less than 5 for each parameter, so this gives $100 \times 100/5 = 2000$ independent samples. There




Figure 6.8 Posterior probability distribution functions (pdfs) for the image parameters corresponding to the images in Figure 6.7. The dashed lines are the true values of the parameter. The blue lines are the pdfs if we include f_{unvig} in our fit; the red lines illustrate the effect of neglecting f_{unvig} and setting it to zero, as in a standard spatial analysis. Neglecting this effect biases the result in a number of problematic ways.

are four free parameters in our fit to the image: β , r_0 , S_0 (parameterized in terms of the total source counts within $6.5'$, N_{src}), and f_{unvig} . The posterior probability distribution functions (pdfs) for these four variables, as computed by `emcee`, are displayed below in Figure 6.8.

All four parameters are recovered extremely well. In particular, the total number of source counts is recovered to nearly 10% accuracy, which is fairly impressive considering the faintness of the source. The slope β of the surface brightness profile is also recovered to 10% accuracy. This parameter governs the behavior at large radii, and is therefore particularly important to recover accurately. The core radius r_0 is recovered to 20%, but this parameter is not particularly important outside of the core. Finally, the background parameter f_{unvig} is constrained within about 25%, which is consistent with our expectations from the analysis in section 6.5.

To compare this to traditional fitting procedures, we repeat the same analysis with f_{unvig} set to 0. This mimics the typical assumption that the background is entirely vignetted which is implicit in the standard step of dividing the image by the exposure map. With this implicit assumption, there are now only three free parameters to fit. The resulting pdfs for these three parameters are also shown in Figure 6.8. In the latter case, inaccurate assumptions about the shape of the background cause all the source parameters to be incorrectly determined.




Figure 6.9 Simulated images with mock point sources added at random, and an extended source placed near aimpoint of the ACIS-I array that contains 4000 source counts. Image (a) is the raw image, (b) shows the same image after running *wavydetect* to detect the background point sources.

6.7.2 Source at Aimpoint

We also examine a simulated image with the source placed near the aimpoint. We specify for the extended source the same parameters as before ($\beta = 0.5$, $r_0 = 10$ pixels, 2000 source counts), and we use the same background as the previous image (i.e. $f_{\text{unvig}} = 0.4$). The simulated images are shown in Figure 6.9, and the recovered parameters in Figure 6.10.

There is a systematic error in recovering f_{unvig} in this image. The overestimate of f_{unvig} also causes the total number of counts in the extended source to be overestimated. We hypothesize that this is due to the difficulty in distinguishing the intrinsic radial decrease in surface brightness from the effects of vignetting. However, even with this systematic issue, the analysis with $f_{\text{unvig}} = 0$ does a worse job of recovering




Figure 6.10 Posterior probability distribution functions (pdfs) for the image parameters corresponding to the images in Figure 6.9. The dashed lines are the true values of the parameter. The blue lines are the pdfs if we include f_{unvig} in our fit; the red lines illustrate the effect of neglecting f_{unvig} and setting it to zero, as in a standard spatial analysis. Here the discrepancy between the two methods is even more severe, and our approach is strongly favored.

each of the parameters.

These simulated images suggest some benefits with our approach. The benefits are accurate determinations of the shape of the surface brightness profile: in both cases, β and r_0 are recovered very precisely. These are the most important parameters to recover, since they govern the shape of the density profile at large radii, and therefore the inferred mass in the extended source. The normalization of the density profile is also recovered fairly well: nearly perfectly for the off-axis source, and with a 10% systematic error for the on-axis source. This contrasts favorably with the normalization determined when f_{unvig} is assumed to be zero, which has a 20% systematic offset in both cases. This also suggests, surprisingly, that our approach is better-suited for observations with the source at an appreciable off-axis angle, so that the source’s surface brightness profile can be more easily disentangled from the instrumental vignetting profile.

6.8 Application to Real Observation: NGC 720

We have seen that this method can recover an extended surface brightness profile robustly in realistic simulations, especially if the source is located at an appreciable off-axis angle. In this section we apply our analysis to real data.

We examine the L^* elliptical galaxy NGC 720. This galaxy was discussed briefly in the introduction as an example of an object for which spectral fitting has produced

inconclusive results. Humphrey et al. (2006) (hereafter H06) inferred a hot halo mass of $\sim 1 \times 10^{11} M_{\odot}$ around this galaxy, while Humphrey et al. (2011) (hereafter H11) inferred a hot halo mass about three times larger. It is difficult to understand how such different results were obtained for the same galaxy. The earlier paper is based on Chandra obsid 492, a 40-ks ACIS-S observation of the galaxy, of which only 17 ks were used due to flaring. The later paper uses different observations: four other Chandra pointings with ACIS-S (obs id 7062, 7372, 8448, 8449) adding to 100.5 ks, and a 177 ks Suzaku pointing (obs id 80009010). According to H11, if they exclude the Suzaku data, the inferred hot halo mass decreased by 50%, which suggests that much of the discrepancy is driven by the Suzaku data (they emphasize XIS1 in particular). However, at smaller radii the discrepancy between Chandra and Suzaku essentially disappears, while the discrepancy between the two papers remains.

We investigate this discrepancy in more detail by examining the radial density profiles for the hot gas, as derived in both H06 and H11. In Figure 3 of H06, they present their inferred deprojected gas density profile, extending out to 90 kpc. We read in this profile, including their uncertainties, using the free WebPlotDigitizer tool³ and used `emcee` to fit β models to the data. The best-fit profile had $\beta = 0.45$, $n_0 = 0.04 \text{ cm}^{-3}$, $r_0 = 1.1 \text{ kpc}$ (yielding a hot gas mass within 300 kpc of $1.5 \times 10^{11} M_{\odot}$), but we retained the full range of fits from the MCMC chain. H11 provides a projected density profile, although their definition of “projected density” is fairly difficult to deproject and analyze. We instead are able to obtain a density profile from their Figure 6, which shows the cumulative enclosed gas mass as a function of radius, extending from about 30 kpc to 400 kpc. The cumulative enclosed mass profile in this range is very well fit with a power-law with slope 1.76 ± 0.02 , corresponding to a β -model with $\beta = 0.41 \pm 0.01$. The core is not resolved, so we leave this as a free parameter. The normalization can be obtained from the enclosed mass in H11 Figure 6, for any given value of r_0 . Unfortunately, the profile in their Figure 6 is just their best-fit profile, and the uncertainties around it need to be estimated as well. We estimate the uncertainties by scaling the uncertainties from H06 by the size of the

³<http://arohatgi.info/WebPlotDigitizer>

error bars on the points in the projected density profile in H11. As H11 point out, the errors are much smaller for the density profile in their newer analysis, partially because they do not deproject the profile, but they are also subject to considerably larger systematics.

We also analyze this galaxy with our method. We use a different observation than either H06 or H11, a Chandra observation (obs id 11868) with ACIS-I taken in the same configuration as observation 8595 which we have been studying in sections 6.4-6.6. The observation was for 40.5 ks, with no significant flaring in the lightcurve. We reduced the data following the same procedure as in section 6.4: running `chandra_repro`, then `fluximage` with a 66.7% exposure threshold to generate an exposure map, then `wavdetect` to find and mask out bright point sources. This observation is shorter than the CDF-S observation we studied perviously, so we expect the background not to be as well constrained in this observation. On the other hand, the source (NGC 720) is brighter than the simulated sources we studied in section 6.6, with about twice as many source counts in a third of the exposure time, so we should be able to recover the surface brightness profile straightforwardly. Also, the background flux rate seems to be higher in this observation than in observation 8595.

We analyze this observation in the same way as in the previous sections, except that we include an additional term in our model to account for emission from X-ray binaries (XRBs). X-ray binaries have a different, harder, spectrum than hot gas, and in theory the two could be distinguished even in a spatial analysis by their hardness ratios. But for NGC 720, H11 have already carefully estimated the total XRB luminosity using two different methods. From the spectrum of the source, they infer an XRB luminosity of 2.95×10^{40} erg s⁻¹ in the 0.5-7.0 keV band. From the XRB luminosity function, they infer an unresolved XRB luminosity of 2.8×10^{40} erg s⁻¹ in the 0.5-7.0 keV band. We take the average of these two estimates, which for an assumed XRB spectrum of a $\Gamma = 1.56$ powerlaw (Irwin and Bregman 2003) yields a luminosity of 1.1×10^{40} erg s⁻¹ in our 0.5-2.0 keV band. We assume the XRBs are distributed like the K-band light is distributed, and we resample the K-band image of

the galaxy from 2MASS into the same pixel coordinates as the Chandra image. Our model for the XRBs is just this K-band image, normalized so the total luminosity is 1.1×10^{40} erg s⁻¹. This adds no additional degrees of freedom to the model, but their inclusion is necessary because they represent a non-negligible contribution to the overall soft X-ray flux.

We convert the surface brightness profile into a mass profile assuming $Z = 0.6Z_{\odot}$ and $kT = 0.5$ keV, based on the spectral fits of H11. We neglect the small temperature gradient which they infer, since most of their annuli are within $1 - \sigma$ of 0.5 keV. We also ignore the metallicity gradient, which is not detected at very high confidence.

The results of this analysis are displayed in Figure 6.11. Here we present the hot gas mass enclosed as a function of radius, for the density profiles from H06, H11, and for our analysis. We can see that, as Humphrey et al. (2011) claimed, their mass model does indeed predict the hot halo contains most, if not all, of the missing baryons from the galaxy within 300 kpc (which is close to the virial radius of the galaxy). This is also about a factor of two greater than the prediction of H06, and the two fits are inconsistent by a little more than $1 - \sigma$ at radii larger than ~ 80 kpc. Our model for the hot gas contains significantly less mass than either H11 or H06 at large radius; extrapolated to 300 kpc we only predict the hot halo contains about half as much mass as H06.

We can also compare the predicted hot gas profiles to the images directly. For our method, we have a full model of the image, including the hot gas, the XRBs, the undetected point sources, and the vignetted and unvignetted backgrounds. For the H06 and H11 results, they used a different technique, so they did not generate an estimate of the unvignetted and vignetted components of the background. We therefore take their estimate of the hot gas profile, and add it to our XRB model and our background model in order to construct a full model image. In Figure 6.12, we plot the radial surface brightness profile of the image, and compare it to the radial surface brightness profiles predicted by these three models.

Here it is clear that the surface brightness profiles inferred from H06 and H11 are not acceptable fits to the data. The uncertainties are much larger than the data




Figure 6.11 1σ regions corresponding to the estimated mass in the hot halo around NGC 720, from three different analyses of the galaxy. Blue corresponds to H06, green to H11, and red to our analysis. Colors transition from dark to light at the location of the outermost spectral bin (for green and blue), or at the outermost radius where the surface brightness is $1 - \sigma$ above the background (red). All mass estimates assume a halo with approximately $Z = 0.6Z_{\odot}$ and $kT = 0.5$ keV. At the right we illustrate the total implied baryon fraction at 300 kpc for this galaxy, including the stars and the hot halo, for various hot halo masses.




Figure 6.12 Comparison of predicted surface brightness profiles around NGC 720 to the observed surface brightness profile, after masking point sources in the image. The black points are the observed values, and the shaded regions are the $1 - \sigma$ ranges of predicted surface brightness profiles, according to (green) H11, (blue) H06, and (red) our analysis. Figure (a) is on a log-log scale, and figure (b) is on a linear scale designed to show the behavior at large radii in more detail. In (b) we also illustrate the best-fit model background with a dashed line and the best-fit hot gas model with a dotted line; we do not show the XRB and undetected point source components separately, but they are in the model as well. H06 and H11 clearly overestimate their uncertainties and are systematically biased upwards; for more discussion see the text.

actually justify, and the profiles are systematically biased above the data. These flaws are the systematic failings of relying solely on spectra to measure faint extended X-ray emission, and they stem from overbinning the data spatially (thereby losing information and increasing the uncertainties in the surface brightness profile) and degeneracies or incorrectly specified components in the spectral model (leading to a systematic overestimate of the gas density at all radii).

Our fit to the image matches the surface brightness profile much more precisely. The uncertainties are very small, and our profiles match the shape of the profile very accurately, which is remarkable considering our model only uses four free parameters (three for the β -model and one for the shape of the background). Figure 6.12b shows the surface brightness profile with a linear scale, which emphasizes the behavior at large radii where most of the hot gas mass is contained. Note in particular that the background can be distinguished from the source emission out to nearly 400 pixels (50 kpc), at which radius the source is less than a tenth of the background. This is a significant improvement over previous analyses. In the analyses of Anderson and Bregman (2011) and Dai et al (2012), for example, our source became lost in the background at 1/5 of the background. This method therefore seems to allow us to reduce by a factor of two the systematic error caused by uncertainty in the shape of the background, which will hopefully open up new frontiers for future analysis.

6.9 Conclusion

In this paper, we have accomplished several goals. We argued that neither spectral fitting nor spatial fitting, in their current levels of sophistication, are adequate for the analysis of extended faint X-ray emission below about 1/5 of the background. We then presented an improvement to spatial fitting, wherein we model the entire image and compare the models to the full dataset available in the events file. In this work, we restrict our model to a single energy band (0.5-2.0 keV), discarding more detailed energy information. We construct a model of Chandra images with three broad components: diffuse background, point sources, and the extended source we are analyzing. The diffuse background can be decomposed into a vignetted and an

unvignetted background. The point sources can be decomposed into detectable and undetectable point sources. The former are masked either with `wavdetect` or with our manual point source detection algorithm, and the contribution from the latter is estimated using the point source number counts of Bauer et al. (2004).

We then explored a number of possible likelihood functions for comparing the models to data. We decided upon a hybrid pixel-by-pixel likelihood function, with pixels binned where the psf is largest (and the correspond point source detectability the poorest). We showed that this likelihood function can reliably recover the shape of the background across the full spectrum from 100% unvignetted to 100% vignetted. This method works with either `wavdetect` or with our manual point source detection algorithm, although we find the best results by masking point sources if either method detects them.

We then examined simulated extended sources, and showed that we can recover the source emission, including the slope of the surface brightness profile, to about 10%. The method works better if the source is placed somewhat off-axis, so that its surface brightness profile can more easily be distinguished from the instrumental vignetting profile.

Finally we applied our method to the isolated elliptical galaxy NGC 720. We showed that the previous measurements of the density profile of its hot halo (Humphrey et al. 2006, 2011) systematically over-predicted the X-ray surface brightness, and we argued that this failure is an inherent feature of the spectral method when the emission is very diffuse. The surface brightness profile predicted by our models offers a much better fit to the data. With this method, we are also able to trace the source emission out to a tenth of the background - a significant improvement over previous methods.

The implied mass of the hot halo extrapolated to 300 kpc is about $8 \times 10^{10} M_{\odot}$. This is insufficient by about a factor of four to bring the galaxy to baryonic closure, in contrast to the conclusions of Humphrey et al. (2011). We instead estimate that, after accounting for the stars and the hot gas, NGC 720 is missing over half of its baryons. This brings the galaxy into closer agreement with our estimates in Anderson,

Bregman, and Dai (2013) of the amount of hot gas around typical L^* galaxies based on stacked images from the ROSAT All-Sky Survey. In particular, for early-type galaxies with 75% the stellar mass of NGC 720, we estimated an average hot gas mass of $3.5 \times 10^9 M_\odot$. Our value of $6.5 \times 10^9 M_\odot$ for the hot gas mass within 50 kpc for this galaxy seems consistent with this estimate, especially given the higher stellar mass for NGC 720 and the steep slope of the $L_X - L_K$ relation.

As mentioned in section 6.2, there are several refinements to this spatial analysis technique which we intend to add in future work. A potentially very powerful addition would be the inclusion of energy in our models, so that we can use the energy of each event to help measure its probability of being in the vignettted background, unvignettted background, or the extended source. We can also extend this modeling to more complicated geometry, such as multiple overlapping images. And we would also like to apply this method to XMM and to Suzaku as well.

6.10 References

- Akamatsu, H., Hoshino, A., Ishisaki, Y., et al. 2011, PASJ, 63, 1019
- Anderson, M. E., & Bregman, J. N. 2011, ApJ, 737, 22
- Anderson, M. E., Bregman, J. N., & Dai, X. 2013, ApJ, 762, 106
- Andreon, S. 2010, MNRAS, 407, 263
- Bauer, F. E., Alexander, D. M., Brandt, W. N., et al. 2004, AJ, 128, 2048
- Bautz, M. W., Miller, E. D., Sanders, J. S., et al. 2009, PASJ, 61,1117
- Bogdán, A., David, L. P., Jones, C., Forman, W. R., & Kraft, R. P. 2012, ArXiv e-prints
- Bonamente, M., Hasler, N., Bulbul, E., et al. 2012, New Journal of Physics, 14, 025010
- Carter, C., Karovska, M., Jerius, D., Glotfelty, K., & Beikman, S. 2003, in Astronomical Society of the Pacific Conference Series, Vol. 295, Astronomical Data Analysis Software and Systems XII, ed. H. E. Payne, R. I. Jedrzejewski, & R. N. Hook, 477
- Cavaliere, A., & Fusco-Femiano, R. 1976, A&A, 49, 137
- Comastri, A., Setti, G., Zamorani, G., & Hasinger, G. 1995, A&A, 296, 1

Dai, X., Anderson, M. E., Bregman, J. N., & Miller, J. M. 2012, *ApJ*, 755, 107

Dai, X., Bregman, J. N., Kochanek, C. S., & Rasia, E. 2010, *ApJ*, 719, 119

Davis, J. E. 2001, *ApJ*, 562, 575

Ebeling, H., Voges, W., & Boehringer, H. 1994, *ApJ*, 436, 44

Eckert, D., Molendi, S., Gastaldello, F., & Rossetti, M. 2011, *A&A*, 529, A133

Eckert, D., Vazza, F., Ettori, S., et al. 2012, *A&A*, 541, A57

Ettori, S. 2003, *MNRAS*, 344, L13

Ettori, S., & Balestra, I. 2009, *A&A*, 496, 343

Fang, T., Bullock, J., & Boylan-Kolchin, M. 2013, *ApJ*, 762, 20

Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, 125, 306

Forman, W., & Jones, C. 1982, *ARA&A*, 20, 547

Forman, W., Jones, C., & Tucker, W. 1985, *ApJ*, 293, 102

Forman, W., Schwarz, J., Jones, C., Liller, W., & Fabian, A. C. 1979, *ApJ*, 234, L27

Galeazzi, M., Gupta, A., & Ursino, E. 2009, *ApJ*, 695, 1127

George, M. R., Fabian, A. C., Sanders, J. S., Young, A. J., & Russell, H. R. 2009, *MNRAS*, 395, 657

Giodini, S., Pierini, D., Finoguenov, A., et al. 2009, *ApJ*, 703, 982

Gonzalez, A. H., Zaritsky, D., & Zabludoff, A. I. 2007, *ApJ*, 666, 147

Gursky, H., & Schwartz, D. A. 1977, *ARA&A*, 15, 541

Hickox, R. C., & Markevitch, M. 2007a, *ApJ*, 671, 1523

Hickox, R. C., & Markevitch, M. 2007b, *ApJ*, 661, L117

Hoshino, A., Henry, J. P., Sato, K., et al. 2010, *PASJ*, 62, 371

Humphrey, P. J., Buote, D. A., Canizares, C. R., Fabian, A. C., & Miller, J. M. 2011, *ApJ*, 729, 53

Humphrey, P. J., Buote, D. A., Gastaldello, F., et al. 2006, *ApJ*, 646, 899

Ichikawa, K., Matsushita, K., Okabe, N., et al. 2013, *ApJ*, 766, 90

Irwin, J. A., Athey, A. E., & Bregman, J. N. 2003, *ApJ*, 587, 356

Kawaharada, M., Okabe, N., Umetsu, K., et al. 2010, *ApJ*, 714, 423

Li, J.-T., & Wang, Q. D. 2013, *MNRAS*, 428, 2085

Li, Z., Wang, Q. D., Irwin, J. A., & Chaves, T. 2006, *MNRAS*, 371, 147

Maller, A. H., & Bullock, J. S. 2004, MNRAS, 355, 694

Martin, N. F., de Jong, J. T. A., & Rix, H.-W. 2008, ApJ, 684, 1075

Mathews, W. G., & Brighenti, F. 2003, ARA&A, 41, 191

Morandi, A., Limousin, M., Sayers, J., et al. 2012, MNRAS, 425, 2069

Mulchaey, J. S. 2000, ARA&A, 38, 289

Mulchaey, J. S., Davis, D. S., Mushotzky, R. F., & Burstein, D. 1993, ApJ, 404, L9

Mulchaey, J. S., & Jeltama, T. E. 2010, ApJ, 715, L1

OSullivan, E., Forbes, D. A., & Ponman, T. J. 2001, MNRAS, 328, 461

Owen, R. A., & Warwick, R. S. 2009, MNRAS, 394, 1741

Pedersen, K., Rasmussen, J., Sommer-Larsen, J., et al. 2006, New A, 11, 465

Ponman, T. J., Bourner, P. D. J., Ebeling, H., & Boehringer, H. 1996, MNRAS, 283, 690

Rasmussen, J., Sommer-Larsen, J., Pedersen, K., et al. 2009, ApJ, 697, 79

Reiprich, T. H., Hudson, D. S., Zhang, Y.-Y., et al. 2009, A&A, 501, 899

Rosati, P., Borgani, S., & Norman, C. 2002, ARA&A, 40, 539

Rozo, E., Rykoff, E. S., Bartlett, J. G., and Evrard, A. 2012, ArXiv e-prints

Sanderson, A. J. R., OSullivan, E., Ponman, T. J., et al. 2013, MNRAS, 429, 3288

Sato, T., Sasaki, T., Matsushita, K., et al. 2012, PASJ, 64, 95

Simionescu, A., Allen, S. W., Mantz, A., et al. 2011, Science, 331, 1576

Snowden, S. L., Cox, D. P., McCammon, D., & Sanders, W. T. 1990, ApJ, 354, 211

Snowden, S. L., Egger, R., Finkbeiner, D. P., Freyberg, M. J., & Plucinsky, P. P. 1998, ApJ, 493, 715

Soltan, A. M., Freyberg, M. J., & Hasinger, G. 2005, A&A, 436, 67

Strickland, D. K., Heckman, T. M., Colbert, E. J. M., Hoopes, C. G., & Weaver, K. A. 2004, ApJS, 151, 193

Tüllmann, R., Pietsch, W., Rossa, J., Breitschwerdt, D., & Dettmar, R.-J. 2006, A&A, 448, 43

Ueda, Y., Akiyama, M., Ohta, K., & Miyaji, T. 2003, ApJ, 598, 886

Walker, S. A., Fabian, A. C., Sanders, J. S., & George, M. R. 2012a, MNRAS, 424, 1826

Walker, S. A., Fabian, A. C., Sanders, J. S., George, M. R., & Tawara, Y. 2012b, MNRAS, 422, 3503

Walker, S. A., Fabian, A. C., Sanders, J. S., Simionescu, A., & Tawara, Y. 2013, MNRAS, 432, 554

Xue, Y. Q., Luo, B., Brandt, W. N., et al. 2011, ApJS, 195, 10

Yamasaki, N. Y., Sato, K., Mitsuishi, I., & Ohashi, T. 2009, PASJ, 61, 291

CHAPTER 7

Conclusion

When this thesis began in 2009, hot halos had not been detected around any spiral galaxies, and their properties were largely unknown. Hot halos had been observed around ellipticals for decades, but in general these were not thought to be massive enough to contain the missing baryons from elliptical galaxies. The lack of detections of hot halos was a persistent problem for the modern paradigm of galaxy formation, since hot gaseous halos around massive galaxies were a generic prediction of the theory.

This thesis contains the first detections of hot gaseous halos around spiral galaxies, which represent a significant triumph for galaxy formation theory, as well as a considerable observational success. In the wake of these detections, activity in this field has increased, and the field has advanced considerably on a number of fronts. To conclude this thesis, we will review the major results in recent years relating to hot gaseous halos around galaxies, both observational and theoretical, including the results presented in this thesis as well as work by other scientists. We will focus on spiral galaxies, but also discuss briefly the connections to elliptical galaxies. Finally, we will discuss directions for future research.

7.1 Late-type Galaxies

In this thesis, we have presented the first detections of hot gaseous halos around late-type galaxies. Whether the gas comes from hot-mode accretion or from feedback, these observations offer powerful new evidence for a fundamental prediction of galaxy formation theory. As noted above, they also suggest that halos do not contain the missing baryons from galaxies. At the moment, however, this latter claim is only

suggestive. This is because the hot halo is only detected out to 40-50 kpc, which is about a tenth of the virial radius of these enormous galaxies. While we are able to fairly robustly constrain the surface brightness profile within 50 kpc, most of the volume (and presumably the mass) lies at larger radii which these observations left unexplored. We extrapolated our surface brightness profiles outwards to estimate the behavior at large radii, and found that the observed hot halos do not seem to contain the missing baryons from galaxies. And we are able to rule out some possible profiles, like a β -model that contains the missing baryons. But an actual measurement of the density at larger radii is sorely needed in order to obtain a direct constraint.

Because of the importance of this extrapolation, density profiles of hot halos have received a lot of attention in the last few years. We mostly use β -models in this work, which are a common and flexible choice, but not the only possible parameterization. Gupta et al. (2012) examined the Galactic hot halo and assumed a uniform density profile; with this parameterization they are able to fit all the missing baryons from the galaxy into the hot halo, but this density profile does not seem to be well-motivated physically. Fang et al. (2013) use an adiabatic profile based on a discussion in Maller and Bullock (2004), which is itself based on a simulation (Frenk et al. 1999). This profile is also flatter than most of the β -models we considered. While it is not able to contain as much mass as a uniform density profile, it contains much more than most of the models we considered. Gatto et al. (2013) has also examined isothermal and adiabatic profiles side-by-side, although observations are not yet sufficient to distinguish these various density profiles.

Observations have also pushed slightly deeper, enabling us to measure the density profiles out to larger radii. Bogdán et al. (2013a) observed NGC 1961 and another giant spiral, NGC 6753, with XMM-Newton. They were able to confirm our detection of the hot halo around NGC 1961, and to detect emission out to 60 kpc around both galaxies. Bogdán et al. (2013b) claims another detection of a hot halo around the giant spiral NGC 266 as well, but this detection used ROSAT and does not extend nearly as far. As both our work and Bogdán et al. (2013a) noted, however, at the moment the primary limitation seems to be flat-fielding instead of sensitivity. Our

new method for flat-fielding allows us to push down these systematics by an additional factor of 3-5 or so, which will hopefully allow us to trace emission to larger radii in the future.

Other than the density profile, the other major uncertainty in the hot halo mass comes from the gas metallicity. In the temperature regimes of hot halos around galaxies, this can be a challenging measurement because most of the emission is line emission, so the X-ray thermal bremsstrahlung continuum is difficult to constrain. This introduces a degeneracy between metallicity and gas density. Without sufficient photons to resolve the lines or the continuum, the same observation can be fit with a high-metallicity / low-density plasma, or a low-metallicity / high-density plasma.

These two extremes also roughly correspond to the two different possible origins for the hot halo. High-metallicity ($Z \sim Z_{\odot}$) gas is expected to be the product of stellar feedback (outflows), and low-metallicity ($Z \sim 0.1Z_{\odot}$) gas is expected to come from hot-mode IGM accretion. Measuring the gas metallicity therefore also gives us a clue as to the origin of the hot halo. Of course, it is also possible for a galaxy to have both hot-mode accretion and hot outflows, or metallicity and temperature gradients, or asymmetric distributions (for example, accretion along the major axis and outflow along the minor axis).

Answers to these questions lie far in the future, however, since at present observations do not even contain enough photons to enable a robust spectral measurement of the metallicity of the gas. We generally fixed our metallicities at an intermediate metallicity such as $0.5Z_{\odot}$ to try to minimize the effect of an incorrect guess in our analyses. Bogdán et al. (2012) measure abundances with formal errors of $0.12 \pm 0.03Z_{\odot}$ and $0.13 \pm 0.03Z_{\odot}$ for the hot gas around NGC 1961 and NGC 6573 respectively, which is tentative evidence that these hot halos are formed from hot-mode accretion of the IGM. However, as they note, the spectra are not yet good enough to overcome the degeneracy with density, and higher-metallicity fits are also acceptable. Hodges-Kluck and Bregman (2013) performed an in-depth study of the hot gas at smaller radii (within ~ 20 kpc) around the L^* galaxy NGC 891. They also concluded that the metallicity could yet not be completely constrained, but that lower ($Z \sim 0.1Z_{\odot}$)

metallicities were favored. On the other hand, Miller and Bregman (2013) inferred $Z \gtrsim 0.3Z_{\odot}$ for the hot halo around the Milky Way.

Finally, most of these observations have focused on the most massive spiral galaxies. We noted in Chapters 3 and 4 that the massive spirals seem to lie below the baryonic Tully-Fisher relation for less massive galaxies. The BTF is calibrated with lower-mass galaxies, and it predicts that the baryon fraction of massive galaxies should approach the cosmological value, but when we estimate the hot halo mass for NGC 1961 and UGC 12591 these galaxies are still missing most of their baryons. It is unclear if the BTF needs to be revised to account for mass within the hot halo, or if these giant spirals represent a true “break” in the power-law relation for some reason. Either way, one would want to understand this result by measuring hot halos around less massive galaxies. But the hot gas becomes both cooler and less massive as galaxy size decreases¹, making these hot halos progressively more difficult to observe. And, according to theory, below a certain critical galaxy mass (thought to be $M_{\text{vir}} \sim 10^{12}M_{\odot}$) IGM accretion should transition to the cold mode, leaving feedback as the only way to generate a hot halo. It would be very desirable to observe the run of hot halo properties with galaxy properties, and especially to study this transition regime around L^* .

We took a step forward in this area with our ROSAT stacking project. We were able to detect extended emission at 99.3% significance around a sample of luminous isolated galaxies with average K-band luminosity of $1.4 \times 10^{11}L_{\odot}$, a factor of two lower than NGC 1961 and a factor of three lower than UGC 12591. We were also able to push down to even lower luminosities, albeit at slightly lower significance. For example, our full sample, with weighted mean luminosity of $5 \times 10^{10}L_{\odot}$, has extended X-ray emission at 92.7% confidence, which is not quite 2σ . We do not detect extended emission around the stack of galaxies fainter than L^* , which means that these galaxies have fainter hot halos than more luminous galaxies (unsurprisingly), or perhaps that they have no hot halos at all.

¹If these hot halos follow similar patterns as those around ellipticals, then the central density also decreases with decreasing galaxy mass, but only slowly.

It is also worth pointing out that many of these constraints act as upper limits on the emission at larger radii. If the hot halos genuinely do contain the missing baryons, then they should be detectable out to much larger radii unless the density profiles are extremely flat. And if the density profiles are extremely flat, then they should be detectable in absorption in the spectra of background X-ray quasars. We can use this combination of constraints to further narrow the allowed parameter space, as we discuss more in section 7.4.

7.2 Early-type Galaxies

While the situation around late-type galaxies is the focus of this thesis, some of our results touch upon early-type galaxies as well. As mentioned in Chapter 1, early-type galaxies have much more luminous hot halos than late-type galaxies, and dynamical heating of the ISM is considered a likely explanation for the observed hot halos of these galaxies. However, isolated ellipticals are missing baryons just like spirals (e.g. Jiang and Kochanek 2007), and hot-mode accretion should also occur onto massive ellipticals, motivating a search for fainter, more extended, and possibly more massive hot halos around these galaxies.

Our stacking analysis (Chapter 5) examined samples of isolated ellipticals as well as samples of isolated spirals. We found that the average X-ray luminosity of the early-type galaxies is about twice the X-ray luminosity of the late-type galaxies, and that the statistical significance of the extended emission around the early-type galaxies is also much higher. This accords with the expectation that ellipticals have more luminous hot gas halos than spirals. In our stacking analysis, however, the stellar mass of ellipticals is also on average higher than the spirals, which would also lead to more luminous hot halos around the ellipticals. Moreover, the luminous spirals and the luminous ellipticals seem to have X-ray luminosities consistent with the well-established L_X - L_K relation, so the discrepancy between the spirals and ellipticals in our sample can be explained entirely by the galaxy masses, instead of having morphological significance.

We also examined an isolated elliptical in chapter 6, the L^* galaxy NGC 720. It

had previously been claimed (Humphrey et al. 2011) that this galaxy has essentially all of its missing baryons in its hot halo, but we showed that this claim is unwarranted, and NGC 720 seems to be missing most of its baryons just like a spiral galaxy of the same mass².

In Table 7.1, we summarize our results from sections 7.1 and 7.2. In general more massive galaxies have more mass within 50 kpc and within the virial radius, as one would expect. It is unclear if the hot halo disappears below L^* or if it just becomes too faint to be detectable. There also seems to be a hint of a trend such that elliptical galaxies have more hot gas mass than spiral galaxies, but there are not enough data points to say anything definitive about this yet.

7.3 Systematic Uncertainties

There are a number of uncertainties in our results around both late-type and early-type galaxies, and in this section we review these various uncertainties, including their magnitudes and the regimes in which they apply.

The uncertainties mostly refer to the parametrization of the density profile and the conversion of a surface brightness profile into an inferred density profile. For example, throughout this thesis, we assume the hot halo is isothermal. In some cases (sections 3, 4, and 6) we measure an emission-weighted temperature for at least some of the halo, while in other cases (sections 2 and 5) we assume a gas temperature. How realistic is the assumption of isothermality, and how accurate are our adopted temperatures? For the former question, we can consider two limiting cases. Galaxy clusters, which are also suffused with quasistatic hot gaseous halos, seem to obey a universal temperature profile (e.g. Vikhlinin et al. 2005) outside of the cool core. These temperature profiles are not quite isothermal; the temperature drops by 50% within the cool core, and appears to decline by about the same amount from the peak at $0.1 R_{\text{vir}}$ to the virial radius as well. Outside of the cool core, however, most of the bright emission comes from near the temperature peak, where the scatter (both

²The Milky Way, another L^* galaxy, is also missing a similar fraction of its baryons - roughly 1/2-2/3.

intra-cluster and between different clusters) is less than 20%.

The other limiting case is the interstellar medium, which is multiphase and highly complex. X-ray emission from the ISM is extremely difficult to model spectrally, and nonthermal effects such as charge exchange and shocks can be important. If hot halos around galaxies are similar to hot halos in relaxed galaxy clusters, then the assumption of isothermality is probably acceptable. If the hot gas has a complex structure like the ISM, then isothermality is less reasonable, and the APEC model itself (which assumes thermal bremsstrahlung + line emission, in collision ionization equilibrium) is probably not applicable.

While definitive observations still need to be taken, a number of indirect arguments lead us to believe that galactic hot halos are much more like clusters than like the ISM, and isothermality is a reasonable approximation. First, temperature profiles for hot halos around elliptical galaxies (including NGC 720) have been measured and are typically similar to those around clusters (with negligible cool cores inside the stellar bulge). This is also true in fossil groups (Khosroshahi et al. 2007, Humphrey et al. 2012). The primary counterexamples to these sorts of relaxed X-ray halos are the starbursting galaxies, which do show complex emission (spectrally and spatially) on 10-20 kpc scales that traces the feedback from these galaxies. For this reason, we attempt to exclude starbursts from our studies. We also exclude merging galaxies for similar reasons.

Focusing on isolated, fairly passively star-forming galaxies, we see no evidence that the emission is significantly more complex than the cases of clusters, groups, and massive ellipticals. The emission tends to be well-fit with APEC models (with an additional powerlaw term to account for XRBs), the emission tends to be azimuthally symmetric (in stark contrast to the winds perpendicular to the disks around starbursting galaxies), and there is no sign of a steepening or a truncation in our surface brightness profiles which might suggest a strong temperature gradient. Deeper observations will definitely help to clarify this picture. The emission measure is not strongly sensitive to temperature in this regime, so if the temperature gradients are not severe and the emission is not strongly nonthermal, then temperature uncer-

tainties should not contribute strongly to our errors. As discussed in Section 5.8.1, temperature uncertainties as large as 50% propagate into $\lesssim 30\%$ uncertainties in the soft X-ray luminosity, and similar uncertainties in the inferred mass.

The other dominant parameter in converting X-ray flux into a density profile is the gas metallicity. This has been mentioned in Section 7.1, and in every preceding section, and we consider it the largest single uncertainty in our measurements herein. If the metallicity is bounded by reasonable values of $0.1Z_{\odot}$ and $1.0Z_{\odot}$, and the inferred density is proportional to the inverse square root of the metallicity, then for an assumed metallicity of $1/3Z_{\odot}$ we could err by as much as 70% in our estimate of the gas mass. Fortunately, this is the worst-case scenario, however. In our analysis of NGC 720, we use the metallicity profile measured by Humphrey et al. (2011), which is $0.7Z_{\odot}$ in the ISM and declines to $0.4Z_{\odot}$ around the galaxy. These measurements have less than 50% uncertainties, corresponding to roughly 30% uncertainties in the inferred mass. For the spiral galaxies, however, the metallicity measurements are much less robust, and even the formal measurements give very little information since they allow for either pristine or two-phase (pristine + enriched) solutions, so here the 70% uncertainties do apply to any quoted gas masses.

There is also a potential concern if the halo has a metallicity gradient with radius, since that would affect the measured surface brightness profile. We perform our extrapolations on the inferred density profile instead of the measured surface brightness profile, since the former is not affected by a metallicity gradient. However, there are two second-order effects which may be relevant. First, if the metallicity has a significant gradient within 50 kpc (where we measure the slope of the profiles), the metallicity gradient can mimic the density gradient, causing us to incorrectly infer the slope of the density profile. This is a concern, but in the clusters and in the elliptical galaxies where metallicities can be measured out to large radii, the metallicity gradients are small.

Secondly, enriched gas has a higher emissivity than unenriched gas, so it is possible to hide a large reservoir of unenriched gas, especially if there is also a bright halo of enriched gas to screen out its emission. The details of such a scenario depend on the

metal mixing in the hot halo, which is still debated (Crain et al. 2013), but no such behavior is seen in other systems like elliptical galaxies or galaxy clusters. Still, we consider this possibility in chapters 3 and 4 explicitly, by examining two-component fits to the observed surface brightness profiles, with the second component having a much flatter density profile than the first, and therefore containing more mass. We are still able to put meaningful limits on the mass in this second component, which are well below the total missing baryons from these galaxies. Further constraints may be applied in the future with X-ray absorption limits, which we discuss in the final section (7.4).

There are two other separate, but related, issues which may complicate this sort of analysis: clumping and filling factor. These effects both work in the same direction. If the gas is clumped, then its surface brightness is increased relative to its density, and so we would overestimate the true gas density. Similarly, if the gas has a filling factor below unity, but a covering fraction of unity, then we would typically overestimate the mass by assuming a uniform filling factor. Thus, if present, both of these effects would exacerbate the problem of missing baryons from galaxies, but they could also in theory significantly affect our other conclusions.

However, we do not expect either effect to be significant. In clusters and groups, there is little clumping within r_{500} , although it can become significant at r_{200} and beyond (Simionescu et al. 2011, Vazza et al. 2013) where the cluster is not fully virialized. Our measurements are all well within r_{500} (which is a few hundred kpc for the giant spirals), so if the virialization process operates similarly in galaxies one would not expect significant clumping in these hot halos either. However, as with much of the analysis in this thesis, this might be complicated by the presence of feedback, since the interaction of a wind with a quasistatic medium might produce shocks, and the wind might contain internal shocks as well. This would presumably appear as asymmetry in the X-ray surface brightness, and no such asymmetry is observed in our images.

The filling factor is also expected to be near unity. Around galaxies, gas in the soft X-ray temperatures has enough thermal pressure to fill the virial volume in a

dynamical time or so (less than 1 Gyr), unless it is confined by another medium with higher pressure. In isolated galaxies, no such medium is expected to exist. The only parts of the system which a hot halo is not expected to permeate fully are the ISM, condensed clumps of warm CGM, and possibly a few cold flows. At a significant fraction of the virial radius (far beyond our observational limits) the gas might not be fully virialized as well.

Simple physical arguments support the conclusion that the clumping and the filling factor will both be close to unity. Multiphase gas can only exist in stable form if the cooling time and the heating time are approximately equal (Field, Goldsmith, and Habing 1969; Mckee and Ostriker 1977). These conditions obtain in the ISM, and to some degree for high-velocity clouds around the disk as well (Wolfire et al. 1995). In the hot halo, however, multiphase gas is probably unstable. A cooler cloud will sink in the hot halo and experience Kelvin-Helmholtz instabilities at the interface with the halo. This will destroy the cloud on a timescale comparable to the sound-crossing time (~ 100 Myr for a 10 kpc cloud) unless magnetic fields or more exotic effects stabilize the cloud. Even if the cloud survives, thermal conduction brings it back to the ambient temperature (again on a timescale around 100 Myr; Cowie and McKee 1977) unless the cloud is large and dense enough for radiative cooling to dominate. And if radiative cooling does dominate, then the cloud falls onto the galaxy on a timescale comparable to the free-fall time (~ 500 Myr).

Moreover, even the transition between the hot halo and the ISM does not seem to show up in our galaxies. This is in contrast to starbursting galaxies, where the transition is very complex both spectrally and spatially, and suggests these effects might not be so important in the non-starbursting galaxies we examine.

Finally, there is a possibility of nonthermal pressure support playing a role in these hot gaseous halos. Obviously feedback would be one such source of support, but we attempt to focus on galaxies which do not have strong feedback. Magnetic fields may be important as well if the field strength is significant outside the galaxy (Parrish and Stone 2005, Quataert 2008). Cosmic rays can also drive buoyancy instabilities, at least in galaxy clusters (Sharma et al. 2009). Another effect could

be "stirring" of the hot gas by dwarf galaxies, in a manner analogous to "stirring" of the ICM by L^* galaxies in clusters (Kim 2007, Ruszkowski and Oh 2010). This effect can be produced either by interaction of the galaxies' ISM with the ambient medium, or through resonant interactions of the gravitational wakes of the galaxies. However, unlike in clusters, subhalos around galaxies do not orbit at significantly supersonic velocities, nor do they have massive interstellar media (since the baryon fraction decreases with halo mass). This reduces the strength of both forms of stirring. Moreover, galaxies assemble at higher redshifts than galaxy clusters, so galactic hot halos have additional time to dissipate turbulence.

Still, one or more of these forms of nonthermal pressure support could be important, especially at large radii. This would flatten the density profile at large radii, potentially pushing some hot gas beyond the virial radius, which exacerbates the problem of missing baryons from galaxies. We do not expect there to be significant nonthermal pressure at smaller radii, where our measurements are. In clusters, turbulence is less than 10% of the thermal energy in the cores of relaxed clusters, and only increases to 20%-30% in merging clusters (Vazza et al. 2011). And as mentioned above, in general we expect turbulence to be less significant in the regime of galactic hot halos.

7.4 Missing Baryons From Galaxies?

Now, we return to the question posed in the introduction, of where the missing baryons from galaxies are. This thesis has argued that they do not lie in hot gaseous halos, since such massive hot halos would be easily detectable in our observations, and the hot halos we detect are very faint. Also, this thesis has argued that galactic winds do not expel the gas at low redshift, since at low redshift galactic potential wells are too deep, and since the observed residuals of galaxies around the Baryonic Tully-Fisher relation do not correlate with the energy available for feedback (Chapter 2).

The remaining option is therefore early ejection of the missing baryons from galaxies. This can either occur through pre-heating of the IGM, which prevents baryons

from falling into protogalactic potential wells (at $z \gtrsim 6$), or through ejective feedback occurring at the peak of the star-formation and AGN rates in the Universe (around $z \sim 2 - 3$), or a combination of the two. Both of these possibilities were discussed in Section 1. At present, a number of various mechanisms have been proposed to effect pre-heating at high redshifts, including AGN feedback (Scannapieco and Oh 2004), accretion shocks (Dopita et al. 2011), pop III stars (Madau et al. 2001), although observational evidence is scant. However, at the peak of cosmic star formation and AGN activity ($z \sim 2 - 3$), outflows are observed to be nearly ubiquitous, although it is presently unclear how much mass is entrained in the outflows.

In addition to the proposed mechanisms for pre-heating, there is additional support for this picture from a few other sources. Peirani et al. (2012) showed that early pre-heating can still affect the baryon fractions of Milky Way-mass halos up to the present day, and therefore that this offers a plausible solution to the problem of missing baryons from galaxies. Bregman et al. (2010) argued that an early top-heavy population of stars is necessary to explain the observed metallicities of galaxy clusters. And Zahid et al. (2012) argued that there is a significant problem of missing metals from galaxies (Oxygen in particular), and they also point to early ejection as the preferred solution.

In contrast to this picture, a different option has recently been suggested. A new paradigm has arisen which treats the accreting and outflowing gas as a unified “circumgalactic medium” (CGM). Concomitant with this conception is the so-called “wind mode” or “warm mode” of accretion (Oppenheimer et al. 2010, Joung et al. 2012). In this mode, the hot halo (if it exists) does not contain the missing baryons from the galaxy; instead the missing baryons are at temperatures around $10^5 - 10^{5.5}$ K. At these cooler temperatures, the gas cannot support itself against the gravity of the system, so it accretes onto the galaxy, but is quickly re-heated and re-ejected by powerful forms of stellar feedback. A dynamically recycling CGM offers a tidy solution to the problem of missing baryons, since it does not require the missing baryons to be ejected from the system (which is energetically difficult), to accrete onto the disk (which would produce more star formation than is observed),

or to remain in the hot phase (which violates our observational limits).

However, it does require some fine-tuning at the moment. The stellar feedback must fall within a fairly narrow range of (strong) efficiencies (Stringer et al. 2012, Creasey et al. 2013) in order to match observations. The AGN feedback also requires similar fine-tuning (Puchwein and Springel 2013). Moreover, while there is extensive evidence emerging for the presence of a warm CGM around galaxies (Tumlinson et al. 2011, Werk et al. 2013), it is still far from clear how much mass this component contains. Depending on ionization corrections and filling factors, the possibilities seem to range from a few percent of the missing baryons from galaxies, to all of them (Stocke et al. 2013).

Note that a warm CGM model is inconsistent with a recent result from Planck (Planck Collaboration et al. 2012), which stacks nearby massive galaxies in their Sunyaev-Zel'dovich signal and infers that most of the missing baryons from galaxies are located in the hot phase, and concentrated within a few times the virial radius. The question of whether the missing baryons from galaxies lie in the warm phase or the hot phase is very interesting, and will be explored in much more detail in coming years.

7.5 Future Prospects

The next few years hold the promise of significant new advances. In the next two years, I plan to produce powerful new constraints on the density of hot halos out to the virial radius. This will require a combination of X-ray emission imaging and analysis of X-ray absorption spectroscopy, along with complementary optical multi-object spectroscopy. The X-ray emission imaging will be a continuation of the new flat-fielding techniques discussed in Chapter 6, which can measure gas density out to nearly 100 kpc. At larger radii, the density is too low for the gas to be visible in emission, so I will turn to absorption. Using existing upper limits on O VII and O VIII absorption along sightlines to X-ray bright quasars at modest redshifts ($z \sim 0.2$), along with optical measurements of galaxies that fall near these sightlines, we can place powerful constraints on the amount of hot gas at large radii around these

galaxies. We introduced this sort of analysis in Chapter 2, and it has since been extended by Williams, Mulchaey, and Kollmeier (2013) for another sightline, but a truly systematic study has yet to be performed. Figure 7.1 shows a schematic diagram illustrating the combined plan of attack. As these constraints are developed, I will also compare the observations to predictions from modern simulations of galaxy formation in order to develop constraints and insight into the galaxy formation process.

Finally, in 3-5 years, new instruments will go online with extremely powerful new capabilities. The E-Rosita mission will produce an all-sky survey that goes several times deeper than the ROSAT All-Sky Survey in the soft band, and improves the angular resolution by a factor of 2-3. Stacking these data should produce a much more robust measurement of hot halos around L^* galaxies. In addition, the first X-ray calorimeters are expected to launch in 3-5 years, starting with ASTRO-H. These calorimeters yield a modest improvement in sensitivity, but a huge improvement in spectra resolution, and should allow us to resolve many of the metal lines in the soft X-rays, finally breaking the degeneracy between abundance and normalization and allowing us to trace the propagation of metals into and out of galaxies in the hot mode.

7.6 References

- Bogdán, A. et al. 2013a, ApJ in press
Bogdán, A. et al. 2013b, ApJ in press
Bregman, J. N. et al. 2012, ApJL, 716, 63
Cowie, L. L. and McKee, C. F. 1977, ApJ, 211, 135
Crain, R. et al. 2013, MNRAS, 432, 3005
Creasey, P. et al. 2013, MNRAS, 429, 1922
Dopita, M. A. et al. 2011, Ap&SS, 335, 345
Fang, T., Bullock, J., and Boylan-Kulchin, M. 2013, ApJ, 762, 20
Field, G. B., Goldsmith, D. W., and Habing, H. J. 1969, ApJL, 155, 149
Frenk, C. S. et al. 1999, ApJ, 525, 554
Gatto, A. et al. 2013, MNRAS submitted

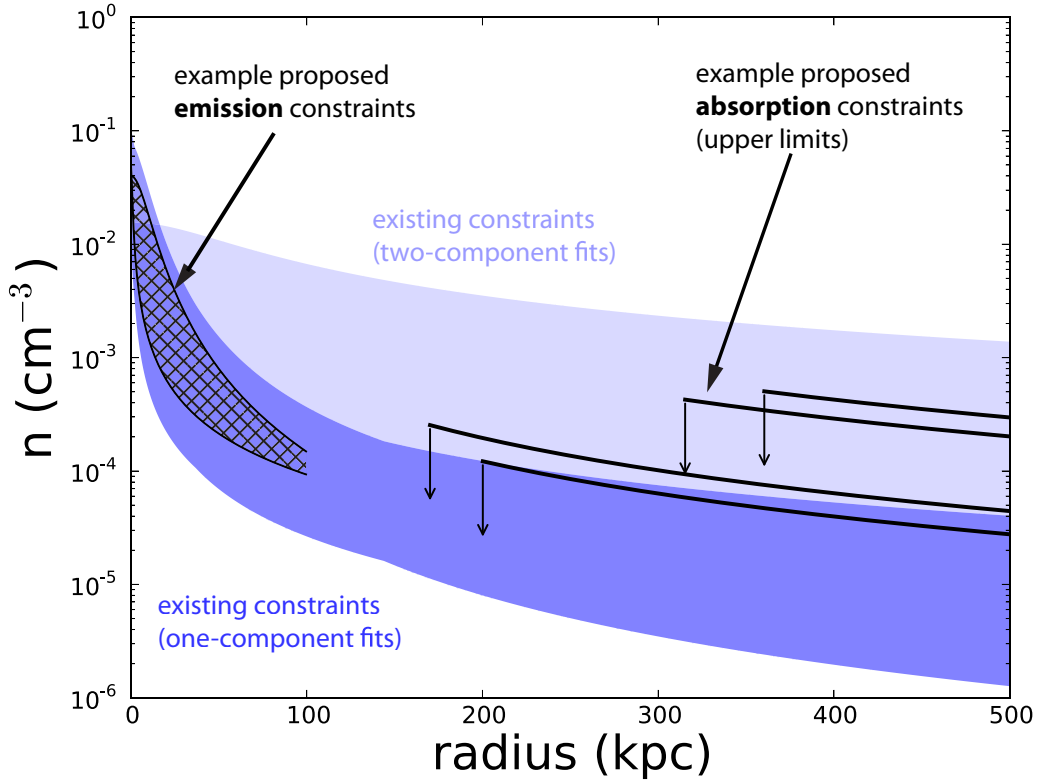


Figure 7.1 A schematic diagram showing example hot halo density profiles and possible constraints. The blue shaded regions show the current constraints based on observations of NGC 1961 and UGC 12591. The dark blue region contains single-component fits based on the surface brightness profiles and the unknown gas metallicity; the lighter shaded region allows a second, flatter, component in the density profile as well. There are orders of magnitude of uncertainty in the gas density at essentially all radii. The black lines show expected constraints at the end of this proposed project; the black hatched region to the left illustrates example emission constraints, and the black lines to the right are example upper limits based on the X-ray absorption analysis. A real analysis will have many more data points, and these quantities will also be studied as functions of stellar mass and star formation rate.

Gupta, A. et al. 2012, ApJL, 756, 8
Hodges-Kluck, E. J. and Bregman, J. N. 2013, ApJ, 762, 12
Humphrey, P. J. et al. 2011, ApJ, 729, 53
Humphrey, P. J. et al. 2012, ApJ, 748, 11
Jiang, G. and Kochanek, C. S. 2007, ApJ, 671, 1568
Joung, M. R. et al. 2012, ApJ, 759, 137
Kim, W.-T. 2007, ApJL, 667, 5
Khosroshahi, H. G. et al. 2007, MNRAS, 377, 595
Madau, P. et al. 2001, ApJ, 555, 92
Maller, A. H. and Bullock, J. S. 2004, MNRAS, 355, 694
McKee, C. F. and Ostriker, J. P. 1977, ApJ, 218, 148
Miller, M. J. and Bregman, J. N. 2013, ApJ, 770, 118
Oppenheimer, B. D. et al. 2010, MNRAS, 406, 2325
Parrish, I. J. and Stone, J. M. 2005, ApJ, 633, 334
Peirani, S. et al. 2012, MNRAS, 427, 2625
Planck Collaboration et al. 2012, A&A submitted
Puchwein, E. and Springel, V. 2013, MNRAS, 428, 2966
Quataert, E. 2008, ApJ, 673, 758
Ruszkowski, M. and Oh, S. P. 2010, ApJ, 713, 1332
Scannapieco, E. and Oh, S. P. 2004, ApJ, 608, 62
Sharma, P. et al. 2009, ApJ, 699, 348
Simionescu, A. et al. 2011, Science, 331, 1567
Stocke, J. T. et al. 2013, ApJ, 763, 148
Stringer, M. J. et al. 2012, MNRAS, 423, 1596
Tumlinson, J. et al. 2011, Science, 334, 948
Vazza, F. et al. 2011, A&A, 529, 17
Vazza, F. et al. 2013, MNRAS, 429, 799
Vikhlinin, A. et al. 2005, ApJ, 628, 655
Werk, J. et al. 2013, ApJS, 204, 17
Williams, R. J., Mulchaey, J. S., and Kollmeier, J. A. 2013, ApJL, 762, 10

Wolfire, M. G. et al. 1995, ApJ, 453, 673

Zahid, H. J. et al. 2012, ApJ, 757, 54

Table 7.1 Summary of Hot Halo Properties

Galaxy	$L_K(10^{11}L_\odot)$	Hubble type	$M_{\text{gas}}(< 50 \text{ kpc}) (10^9 M_\odot)$	$M_{\text{gas}}(< R_{\text{vir}}) (10^9 M_\odot)$	$Z(Z_\odot)$
NGC 1961	5.2	Sb/c	$5.0_{-0.1}^{+0.2}$	$2.3_{-0.9}^{+0.3+5.1}$	0.5
UGC 12591	5.6	S0/a	$4.4_{-0.9}^{+0.7}$	$1.5_{-1.3}^{+1.6+3.0}$	0.5
stacked luminous	$1.4_{-0.4}^{+0.8}$	-	$4.1_{-1.0}^{+0.6}(\text{stat}) \pm 2.9(\text{sys})$	-	0.3
stacked faint	$0.3_{-0.1}^{+0.4}$	-	$0.9_{-0.4}^{+0.5}(\text{stat}) \pm 0.6(\text{sys})^*$	-	0.3
stacked early-type	$0.7_{-0.3}^{+1.2}$	E and S0	$1.9 \pm 0.9(\text{stat}) \pm 1.3(\text{sys})$	-	0.3
stacked late-type	$0.4_{-0.2}^{+1.0}$	Sabc & Irr	$1.2_{-0.6}^{+0.5}(\text{stat}) \pm 0.8(\text{sys})^*$	-	0.3
NGC 720 (this work)	1.6	E5	6.5 ± 0.5	0.8 ± 0.1	0.6

Hot halo masses measured in this thesis. NGC 1961 is discussed in Chapter 3. UGC 12591 is discussed in Chapter 4. The stacked galaxies are discussed in Chapter 5. And NGC 720 is discussed in Chapter 6. For each column, L_K comes from 2MASS and the Hubble type from NED. Our measurements of the hot halo gas mass within 50 kpc and within the virial radius are then listed for each galaxy. In general, the mass within 50 kpc is fairly secure, and the mass within the virial radius depends on extrapolating the density profile out to much larger radii than it is observed. In general the errors quoted are statistical errors based on uncertainties in the surface brightness profile; other sources of error are discussed in the text and in Section 7.3 below (most of these other errors are folded into the systematic errors quoted for the stacked galaxies, however). Note that the asterisks on the stacked faint galaxies and the stacked late-type galaxies denote lower confidence that we are actually detecting and characterizing extended emission around these galaxies. For NGC 1961 and UGC 12591, the second listed uncertainty on the mass within the virial radius accounts for the possibility of a flattened profile, as discussed in the text. For the stacked galaxies, we do not extrapolate the mass to the virial radius since we do not have a strong measurement of the slope of the density profile. The final column lists the assumed (or measured, in the case of NGC 720) metallicity used to convert the surface brightness profile into a gas mass.