

BRIDGING THE GAP BETWEEN BLACK HOLE
ACCRETION DISK THEORY AND X-RAY SPECTRAL
OBSERVATIONS

Brianna Sue Wu Mills

Louisville (“lull-vul”), Kentucky

B.S. Physics with Honors, Minor in Philosophy, University of Louisville, 2018

M.S. Astronomy, University of Virginia, 2020

A Dissertation Presented to the
Graduate Faculty of the

University of Virginia

in Candidacy for the Degree of

Doctor of Philosophy

Department of Astronomy

University of Virginia

May 2023

Committee Members:

Shane W. Davis

Craig Sarazin

Phil Arras

Kent Yagi

© Copyright by
Brianna Sue Wu Mills
All rights reserved
May 20, 2023

ABSTRACT

Accretion of material onto black holes and other compact objects is a fundamental process that shapes the growth and evolution of galaxies, star forming regions, and powers some of the most energetic phenomena in our universe (e.g. X-ray binaries, active galactic nuclei, gamma ray bursts, etc.). Black holes are notorious for their immense gravitational pull, stripping gas from stars which venture too close. Due to the angular momentum of the in-falling gas, a disk is formed around the black hole releasing incredible amounts of radiation, often at X-ray wavelengths. These accretion disks are extremely hot and luminous, which not only provide a way for us to detect and observe black holes, but also to study their only astrophysical properties: mass and spin. Modeling the complex accretion flows around black holes serves as one of the most direct ways we can obtain the spin of a black hole and test Einstein's theory of general relativity (GR) where the curvature of spacetime is at its most extreme.

The foundational theory of accretion disks was developed almost 50 years ago, and it assumes that the gaseous material falling onto a black hole will take the form of a geometrically thin, optically thick accretion disk. The thin disk model is widely applicable to other areas of astronomy, such as proto-planetary disk systems and disks around accreting neutron stars and white dwarfs. It is commonly used as a first-order fit to the observed X-ray spectrum of accreting black holes. However, our theoretical understanding of black hole accretion has been challenged by observations of X-ray sources which are not well described by the thin disk model. An example of such sources are the well-known class of objects called ultraluminous X-ray sources (ULXs), which appear to accrete well above their Eddington limit, or the point at which the outward radiation pressure near the black hole exceeds the inward pull of gravity. Super-Eddington accretion onto black holes implies a much thicker accretion disk accompanied by radiatively driven outflows. This raises fundamental questions

about the nature and physical mechanisms governing super-Eddington accretion, motivating the development of X-ray spectral models that can more accurately describe these flows. Due to the complex three-dimensional nature of super-Eddington accretion, numerical radiation magnetohydrodynamic (RMHD) simulations are required to replicate these environments.

In Chapter 2 of this dissertation, I explore commonly used X-ray spectral fitting models to re-examine the black hole spin in the black hole X-ray binary (BHXB) system GRS 1915+105, in light of new constraints on its mass, distance, and inclination. The thin disk assumptions underlying the spectral fitting models tend to work well for the spectra of BHXBs, whose accretion rates remain below their Eddington limit. The thin disk model begins to break down at super-Eddington accretion rates, in which case we turn to using a state-of-the-art RMHD simulation code, **Athena++**, in order to simulate these flows. In Chapter 3, we describe a Monte Carlo radiation transfer module that is designed to produce spectra from the **Athena++** simulations. We use this module in Chapter 4 in conjunction with simulations of super-Eddington accretion onto a stellar mass black hole in order to produce simulated spectral models. We then compare these simulated models with the phenomenological models in Chapter 2, and I perform direct fits to the observed X-ray spectrum of the ULX NGC 1313 X-1. We find that the simulated spectral models from the RMHD simulations are qualitatively able to reproduce the observed spectrum of NGC 1313 X-1, providing a promising direction for bridging the gap between black hole accretion theory and observations of super-Eddington accretors. A summary of each chapter is provided in Chapter 5, along with future directions and code developments on the (event) horizon.

ACKNOWLEDGEMENTS

I don't remember which elementary school teacher told me that I had two choices in life: get a real job after high school, or continue with school and write a 100 page paper to get a PhD. Apparently I did the exact thing that I did *not* want to do at that time. So, here is my 100+ page paper which wouldn't have been possible without the unknown forces which have swayed me here, but also to the following people who encouraged such a thing to keep progressing.

To my committee members: Kent, thank you for hosting me at one of the gravity group meetings and for agreeing to be on my committee – and congrats again on being a dad!! Craig, thanks for all of your support and the interesting tidbits you've shared about X-ray astronomy. I've enjoyed your fun commentary at various meetings and clubs over the years. Lastly, thanks to Phil for also being part of my committee, and for giving one of the best Kovalenko speeches that ended promptly at 6 minutes past the time the bar closed. I had planned to go back to the department after the dinner to continue writing my dissertation, however I did not plan for a few unused drink tickets to sit longingly next to my computer. It was a long (but painfully productive) night. Thank you for unfortunately helping me write a coherent dissertation.

To my advisor, Shane: Most of the people around me are probably tired of hearing how great of an advisor you are, since I've discussed it plenty of times. I'm so thankful for your support, encouragement, advice, and how truly devoted you are to the progress and well-being of your graduate students. Even if I didn't have live footage of you "singing" the Duke Fight Song at the Charleston conference as leverage, I think you would still be as supportive as you have been. And don't worry, I'm absolutely, completely, totally over that time you said to the entire High Energy class, "Bri and Xiaoshan are working on research with me and they'll likely never get

a Nobel Prize.”¹ Truly, I’ve enjoyed the research projects I’ve done, and I have gained so much from working with you.

I’m incredibly thankful to have such great collaborators: Yan-Fei, Matt, Jim, Lizhong, Sergei, and the entire **Athena++** collaboration. I’ve experienced so much growth since my first **A++** meeting, when I was but a young grad student still trying to understand the world of computational astrophysics. Everyone was so kind and made an effort to tell me how well my talks went. It gave me motivation and courage to achieve the bulk of the work I did in Chapter 4. I’m excited to see everyone again and to present these results at the conference in May.

To the true MVP’s of the department: Whitney & Gabby. Thank you so much for all of the help and support over the years. It was also really nice to know someone from back home in good ol’ Kentucky, so thanks Whitney. Also thanks to the Jefferson Scholars Foundation for all of the wonderful events, opportunities, and spaces to grow as an academic and community member. I am so sad to be leaving this year, but I will always cherish the wonderful memories and friends I’ve made at the Foundation.

To my grad school cohort: Anne, Andy, Richard, Hannah, Daniel, Nick, (Renato + Samer = 1 full 5th year?). There’s no dynamic quite like our group (the best group). All the board game nights, Halloween parties, potlucks, sand volleyball, Pokemon Go walks, starting the \hbar tradition after quals, and so many engaging conversations over the years. You all made grad school so much more enjoyable and welcoming. I was also fortunate to spend my first few years as a grad student naively looking up to the elder grads. Thanks to Robby W. and Chris H. for frequently coming into the big grad office “for the candy drawer” when really it was to avoid work or maybe just to distract us. I really missed those distractions after you all left. Also thanks to Paul and the old D&D group. Even though we never finished that campaign, it was hands down the best cancelled show I’ve ever experienced. So many laughs – mostly because Robby just had to be extra, but also we did try to get away with anything and everything. From finishing pitchers at Nelly’s (version 1.0), to the fire pit beer talks at Matt/Sean’s and Matt/Molly’s, to drinking and dancing at Thankful and Riley’s, to cat sitting, Murder in the Dark, making dumplings for Lunar New Year with Yiqing, Xiaoshan, et al... I will remember most of my grad years fondly. Also thank you to Connor, Xinlun, Mark, Prakanya, Deryl, Jordan, Mihika, Simran,

¹We know you meant the research would never get a Prize, not the graduate students, but the results are still the same.

Lianis, Andres, Deven, and all of the other grads my extremely caffeinated brain cannot name in this moment who've also made the department feel fun, interesting, and welcoming. The grad students at UVA were one of the biggest reasons I chose to come here, and I'm happy that it turned out to be the right choice for me. Finally, thanks to Trey Wenger, you L^AT_EX wizard. Thank you for all of the help and resources you left for future, struggling grad students – this dissertation template included.

To my best friend, Anne. We bonded over Gilmore Girls our first year and we've grown together since ("better together!!"). Sometimes it's really felt like a Film By Kirk, so thank you for all of the incredible support, friend. Hope you're staying warm in Buffalo, and I can't wait to go on more of our accomplished walk n' talks (and the Falls!) together.

To Xiaoshan, thank you so much for all of your help and chats throughout grad school. I've thoroughly enjoyed our talks coming up with so many great cat puns and acronyms. From adventuring in Vegas, to our nice walks, to our Tofu Exchange Law, thank you for being such a great friend and research sister. I'm excited to see where our lives take us!

To Allison M., thank you for being such a sweet soul and introducing me to Patches, a.k.a. Gutter Cat. Some of the best times of my grad school career was doing work with him purring SO loud next to me, just living his best life, always sleeping on my face. I'll never forget how thoughtful your welcome gifts were to our first year class. It really improved my transition to grad school.

To Allison T., although UHeights was hardly a great place, it was so nice living near you and Sibelius. Thank you both for all of the long, heartfelt talks when I really needed someone to talk with, and for reaching out to grab lunch when you were in town for ALMA things. It was also a treasure spending time with Mr. Big Paws, especially when I needed a good break.

To Sara, Sam, and Skye. I'm so glad that I met you in my first year of grad school. It's truly been wonderful to witness your family grow, and it's something I look forward to in my own future. Having wonderful human beings to connect with during the waves of grad school and a global pandemic has been so comforting. I will remember our lunches at Revolutionary Soup and spending time with you all in your new home with your new little one. Still to this day, Skye is my favorite little person. I can't wait to see what doors she kicks down when she's older!

To my undergrad physics friends, Olivia, David, Nick, Samir, and Austin. You

are all some of the most enriching people in my life, and it was only fair since we all met through an interest in physics. I'm happy that we've kept in touch since leaving Natty Sci and I'm looking forward to many more warm conversations. Olivia, I'm so glad we can be our complete unabridged selves with each other. Our podcast is going to be so left field and existential that it'll end up as a collection of tasty but forgotten memes, a true feat. Can't wait to see you, friend.

To my childhood best friends, Priscilla and Chloe. You've both cheered me on since elementary/middle school. Great friendships are hard to come by, and I feel really lucky to have found you both so early in my life. Thank you for your unwavering support while I trudged through grad school, and for dealing with me not responding to messages for weeks/months at a time. I can't wait for our girls' trip to Europe now that I'm going to be done with grad school and somewhat more alive.

To my wonderful partner. We never intended to find each other especially during one of the most chaotic times in both of our lives, but I'm so glad we did. They say that some of the best relationships are ones where you feel safe enough to begin healing parts of yourself and grow into a more mature human being. (Paws for dramatic effect). Thank you for holding that space with me and for learning, growing, and communicating together in this partnership. I love you, JW.

To Spirit and Meat(Leo)f, my babies, who will never read this but know I love them enough to keep buying their favorite wet foods. I'm thankful for all of the dumb things you did that made me laugh when grad school became heavy, and for all of the big loves and purrs. I love you both so much.

To my sister goblins who I almost had to murder in order to finish this (ahem, JJ, why did you miss your bus 437978 times the week my dissertation was due??). To the rest of my family: Hannah, Scott, Boots, Beth, Oma, Coufue, Kevin, Pua Pua, Gong Gong, Jason, Brandon, Stepfan, Nana, Poppy, Michelle, Maddie, Steven, Brayden, Jaiden, and Lileigh. Thank you for being part of my life. For shaping who I am and who I choose to care for. For teaching me that unconditional love is not a bargaining chip, but a mutual love.

Most importantly, to my mother. Words can't describe how grateful I am for you and everything you've done to protect us from the chaotic and unforgiving circumstances of our childhood. Thank you for everything you have sacrificed so that we could end up here, together, with me writing this and you being my biggest supporter. The first PhD in the family is happening because of you and your incredible

love, support, resilience, and goofiness. I'm so, so glad you're my mom. I love you, ma.

Lastly, I would also like to thank the author for persevering through an extraordinary period of loss, grief, challenge, and growth. The work wasn't possible without you at the (overw)helm, navigating incredulous storms especially when it was clear that another storm was fast approaching. I'm indebted to Dr. DB – thank you for always answering my questions and ensuring that I'm physically cared for. And to my therapist DKF, who is still helping me work through trauma and other complexities during the pinnacle of my educational career, thank you for reminding me that I am human, not some superhero, and that my resilience will still be here when I'm ready. I give myself permission to rest and just *be*. Thank you all for believing in me so that I could, too.

For my mom and sisters.
Without you, I probably could have had more cats.

TABLE OF CONTENTS

Abstract	iii
Acknowledgements	v
List of Figures	xiv
List of Tables	xv
1 An Introduction to Black Holes, X-ray Astrophysics, and Radiation Magnetohydrodynamics	1
1.1 Properties of black holes	7
1.2 Accretion Theory	9
1.2.1 Thin accretion disk model	10
1.2.2 Super-Eddington accretion	12
1.3 X-ray Observations: Methods, Modeling, and Introductions	14
1.3.1 Continuum-fitting method	15
1.3.2 BHXB: GRS 1915+105	17
1.3.3 ULX: NGC 1313 X-1	20
1.4 Radiation Magnetohydrodynamics and Monte Carlo Radiation Transfer	21
1.4.1 Athena++ RMHD code	21
1.4.2 Monte Carlo radiation transfer	24
1.5 Outline of this thesis	25
2 The Black Hole Spin in GRS 1915+105, Revisited	27
2.1 Abstract	27
2.2 Introduction	28
2.3 RXTE Data Reduction	30
2.4 Results	32
2.4.1 Non-relativistic accretion disk model	32
2.4.2 Relativistic accretion disk models	35
2.4.3 Color correction factor	38
2.4.4 Exploring System Uncertainties	41

2.5	Discussion	42
2.5.1	Implications for the GRS 1915+105 system	42
2.5.2	Impact and Uncertainties in Interstellar Absorption	46
2.5.3	Uncertainties in Models and System Parameters	47
2.6	Summary	50
3	The Athena++ Monte Carlo radiation transfer module	52
3.1	Adopting the Athena++ Code Framework	52
3.2	Photon samples	54
3.3	Emission and Statistical Weights	54
3.4	Photon Movement	58
3.5	Scattering and Absorption	61
3.6	Outputs	62
3.7	Tests of the Method	62
3.7.1	Absorption and Thomson scattering	65
3.7.2	Boost velocities	67
3.7.3	Compton scattering	70
3.7.4	Transport in Curvilinear Coordinates	70
3.8	Using the MC module	72
4	Spectral Calculations of 3D RMHD Simulations of Super-Eddington Accretion Onto A Stellar-Mass Black Hole	75
4.1	Abstract	75
4.2	Introduction	76
4.3	Methods	78
4.3.1	MC radiation transfer code	78
4.3.2	Athena++ RMHD simulation snapshots	81
4.3.3	Spectral post-processing	88
4.4	Results	89
4.4.1	Comparing Athena++ with Monte Carlo	89
4.4.2	Spectra & Compton Cooling of Gray RMHD simulations	94
4.4.3	A multi-group RMHD approach	98
4.4.4	Simulated spectra in comparison with observational models	102
4.4.5	Example analysis: NGC 1313 X-1	108
4.5	Discussion	113
4.5.1	Caveats and Sensitivity to Model Assumptions	113
4.5.2	Comparison with ULX modeling	114
4.5.3	Comparison with previous work	115
4.6	Summary	116
5	Summary and Future Directions	119
	References	123

LIST OF FIGURES

1.1	Artist's illustration of Cygnus X-1	2
1.2	EHT Images of Sgr A* and M87*	4
1.3	Optical, infrared, and X-ray images of GRS 1915+105	18
1.4	Fig. 4 from Bachetti et al. (2013) of spectral fits to <i>XMM-Newton</i> and <i>NuSTAR</i> data of NGC 1313 X-1.	19
2.1	Spectral fits to three RXTE observations from Middleton et al. (2006)	36
2.2	Color correction values as a function of black hole spin for three RXTE observations from McClintock et al. (2006)	39
2.3	Spectral fits to three RXTE observations from McClintock et al. (2006)	40
2.4	Range of parameter space (D, M, i) as a function of black hole spin for the three RXTE observations from Middleton et al. (2006)	43
2.5	Range of parameter space (D, M, i) as a function of black hole spin for the three RXTE observations from McClintock et al. (2006)	44
3.1	Test spectrum of an isothermal atmosphere with pure absorption . . .	63
3.2	Test convergence of a simple isothermal atmosphere with density gradient	64
3.3	Test spectrum of an isothermal atmosphere with polarized Thomson scattering and absorption.	66
3.4	Test convergence of E_r with boosted velocities in the Eulerian and Comoving frames.	68
3.5	L1 Norm convergence of boosted velocities in the Eulerian and Comoving frames.	69
3.6	Test spectra for uniform spherical case with Compton scattering . . .	71
4.1	$\dot{M}/\dot{M}_{\text{Edd}}$ vs. r_g of the RMHD snapshots.	82
4.2	Gas density of grey RMHD snapshot ULX2.5	85
4.3	Gas temperature of the grey RMHD snapshot ULX2.5	87
4.4	Comparison of the radiation energy density between the <i>Athena++</i> grey RMHD simulation and the MC calculation	90
4.5	Ratio of the MC mean photon energy $\langle h\nu \rangle$ to the grey RMHD $4kT_r$ for ULX2.5	92

4.6	Cooling comparison between the grey RMHD simulation and the MC calculation for ULX2.5	93
4.7	Monte Carlo post-processed X-ray spectra for grey RMHD snapshots	95
4.8	Flux fraction as a function of inclination angle for grey RMHD snapshots	97
4.9	Images of emergent radiation for ULX2.5-MG and ULX2.5 for different inclination angles	99
4.10	Radiation energy density comparison between the multi-group RMHD snapshot and the MC	101
4.11	Comparison between grey RMHD spectra and the corresponding multi-group RMHD spectra	103
4.12	Spectral model fits to simulated multi-group RMHD spectrum	107
4.13	Best-fit to <i>XMM-Newton</i> and <i>NuSTAR</i> data of NGC 1313 X-1 with the ULX4a-MG spectral model	111
4.14	Best-fit to <i>XMM-Newton</i> and <i>NuSTAR</i> data of NGC 1313 X-1 with the ULX2.5-MG spectral model	112

LIST OF TABLES

2.1	List of XSPEC models and notation	33
2.2	Best-fit values for the MID06 observations	37
4.1	Athena++ RMHD simulation snapshot parameters	83
4.2	Model comparison fits to simulated RMHD spectra	105
4.3	Best-fit parameters for fits to NGC 1313 X-1 data using simulated spectral models	109

CHAPTER 1

AN INTRODUCTION TO BLACK HOLES, X-RAY ASTROPHYSICS, AND RADIATION MAGNETOHYDRODYNAMICS

A black hole is a region of strongly curved spacetime where nothing – not even light – can escape its immense gravitational pull. The event horizon of a black hole, or the “point of no return,” depends on the properties of the black hole. Stellar mass black holes ($M \sim 3 - 100 M_{\odot}$) are formed when a massive star ($M > 10 M_{\odot}$) dies, collapsing and expelling material into the surrounding interstellar medium (ISM) and leaving behind a dense, compact core that becomes a black hole. Although massive stars tend to be short-lived and are thus rarer than less massive stars in the universe, they are the biggest and brightest stars whose influence and radiation significantly impact the growth and evolution of their host galaxies. Stellar mass black holes are therefore important for understanding massive star evolution, their consequent supernovae explosions, and pollution of the ISM with heavy elements. The latter is particularly important for understanding the elemental building blocks of nature (e.g. approximately the first half of the period table) from which new stars, planets, and possible life form.

The vast majority of massive stars also exist as part of a binary or higher order multi-body system ([Sana et al., 2012](#)). This can be two massive stars, one massive



Figure 1.1: Artist's illustration of the BHXB Cygnus X-1, discovered in 1971. The companion star shown on the left is a 33 solar mass (M_{\odot}) star, and the smaller object to the right is the black hole which is stripping gas from the star via Roche lobe overflow, forming an accretion disk. Optical spectroscopic observations determined this star has a large radial velocity amplitude of ~ 50 km/s. The orbital period of the binary system is about 5.6 days, inferring that the mass of the heavier object must be a black hole of about $15M_{\odot}$. X-ray radiation is produced from the innermost part of the accretion disk where the temperatures are hottest. (Image credit: NASA's HEASARC Education and Public Information)

star and one smaller star, or other such combinations. When a massive companion star undergoes collapse and forms a black hole, the result can be an energetic black hole X-ray binary (BHXB) system, assuming the other star survives the explosive collapse and is sufficiently close to the black hole.

Figure 1.1 shows an artist’s illustration of the BHXB system Cygnus X-1. The strong gravitational field of the black hole begins to pull or accrete gaseous material from the surviving companion star, forming an accretion disk from which the gas slowly spirals into the black hole. The accretion disk can be thought of as a pancake, donut, or “torus” with the black hole at the center¹. When material is accreted onto a black hole, an immense amount of energy is released in the form of X-ray radiation, hot winds, and possibly relativistic jets. The detection of X-ray radiation has historically been a reliable method to infer the presence of an accreting black hole, or other compact objects such as neutron stars and (in special cases) white dwarfs.

On the other hand, supermassive black holes (SMBHs; $M = 10^6\text{--}10^9 M_\odot$), such as Sgr A^* at the center of our own Galaxy, are thought to form either from the gravitational collapse of giant gas clouds (e.g. during galaxy formation) or massive stars. They then grow by accreting more gas or via the repeated merger of massive black holes in binary systems. The latter process is one focus of an emergent field of astrophysics studying gravitational waves emitted from the in-spiral and coalescence of merging black holes (Abbott et al., 2016; also see Bailes et al., 2021 for a review). Almost every galaxy harbors a supermassive black hole at its center. Supermassive black holes that accrete and release energetic radiation become Active Galactic Nuclei (AGN; see Morganti, 2017; Blandford et al., 2019 for a review). A subset of AGN called quasars exhibit bright, relativistic, supersonic radio jets that can outshine the entire host galaxy and heat their surroundings. AGN/quasars may even trigger or stifle star formation in their host galaxies (Morganti, 2017). Understanding the population of black holes and their impact on host galaxies is an important field of study that will be advanced by current and future surveys. For example, the Vera C. Rubin Observatory’s Large Synoptic Survey Telescope (*LSST*) will compile the largest sample of astrophysical phenomena including AGN sources (Ivezić et al., 2019). Studies of supermassive black holes, AGN, and quasars are thus important for understanding galaxy formation and evolution, as well as the birth places of stars.

¹“A cosmic bakery!” - Xiaoshan Huang




Figure 1.2: The first two direct images ever taken of black holes. The left panel shows M87*, the supermassive black hole at the centre of the Messier 87 galaxy (M87), which is about 55 million light-years away from Earth. The right panel shows Sagittarius A* (Sgr A*), the black hole at the centre of our Milky Way galaxy. These two images show the black holes as they would appear in the sky, with their bright rings appearing to be roughly the same size, even though M87* is about a thousand times larger than Sgr A* (but it appears to be the same size because it's about a thousand times farther away!). These images were captured by the Event Horizon Telescope (EHT), a global network of radio telescopes including the Atacama Large Millimeter/submillimeter Array (ALMA) and the Atacama Pathfinder EXperiment (APEX), in which the European Space Agency (ESO) is a co-owner. (Image credit: EHT Collaboration and the ESO)

There also exists a middle population of black holes with masses between that of a stellar mass black hole and a supermassive black hole called intermediate mass black holes (IMBHs, $M \sim 10^2\text{--}10^5 M_\odot$), but this population is more elusive than its counterparts. Presently there are few solid detections of IMBHs, but there are a growing number of IMBH candidates (for a review, see [Greene et al., 2020](#)). How IMBHs are formed is not yet known, but some theories suggest that IMBHs must have at least been formed in the early universe ([Bañados et al., 2018](#)) in order to give rise to some of the supermassive black holes we see today. Other theories have proposed that IMBHs may have initially been stellar mass black holes which grew by capturing and consuming stars over its lifetime ([Rizzuto et al., 2023](#); [Arca Sedda et al., 2023](#)). Future missions such as the Laser Interferometer Space Antenna (*LISA*, the space-based successor of LIGO, the Laser Interferometer Gravitational-Wave Observatory), will be able to probe this mass range and study the properties of candidate IMBHs. In this dissertation, we do not investigate IMBH properties, but we note that they have been posited as a possible explanation for some ultraluminous X-ray sources, which we discuss below.

Nearly all black holes² can serve as laboratories for probing Einstein’s theory of General Relativity (GR) and modeling the behavior of light and matter in extreme gravity environments. Physicist John Wheeler is credited with coining the term “black hole” and also eloquently explained Einstein’s theory of GR as:

*Matter tells spacetime how to curve, and curved spacetime tells matter
how to move.*

The curvature of spacetime becomes more curved around more massive objects, so the effects of GR are strongest in regions closest to a black hole (e.g. near the event horizon). The Newtonian picture of gravity as a force as we are classically trained to understand is represented in GR as the degree of curvature of spacetime.

Studying the regions nearest the black hole event horizon is extremely difficult due to the combination of the small observable size of the black hole and the large distances between us. For this reason, we have only recently been able to directly image two black holes with the Event Horizon Telescope, seen in Figure 1.2 (M87*: [Event Horizon Telescope Collaboration et al., 2019](#); Sgr A*: [Event Horizon Tele-](#)

²The observable ones, at least. Although, even isolated black holes apparently can’t hide forever (see [Sahu et al., 2022](#)).

scope Collaboration et al., 2022). Nevertheless, since the discovery of black holes and their associated X-ray emission, X-ray telescopes have been able to observe and store a multitude of data for black holes. Some of these missions include cutting-edge NASA/ESA observatories like the *Chandra* X-ray Observatory, *XMM-Newton*, *RXTE*, and *NuSTAR*. Black holes can also be observed with ground-based observatories covering the radio to optical wavelengths, in addition to gravitational wave signals with *LIGO* (and its longer wavelength counterpart, *LISA*).

In this dissertation, we make use of archival X-ray observations of accreting black holes and focus on analyzing their X-ray spectral properties. X-ray observations in particular illuminate the accretion disk, as well as other regions near the black hole, such as the “corona,” which is a region filled with hot electrons named in analogy to the hot corona or atmosphere around our Sun. The black hole corona is generally thought to reside somewhere very near the black hole and consists of hot electrons which produce high energy (hard) X-rays via inverse-Compton scattering. This process happens when the lower energy (soft X-ray) photons coming from the accretion disk scatter off of the hot electrons in the corona, taking some of the electron’s energy and thus becoming higher energy photons that are observed at hard X-ray wavelengths (Gierliński et al., 1999; McClintock & Remillard, 2006; García et al., 2020). The origin, geometry, and location of this hot corona is still under investigation (You et al., 2021; Zhong & Wang, 2021; Krawczynski & Beheshtipour, 2022; Peirano et al., 2023; Serafinelli et al., 2023).

Distinguishing between the different physical processes and their resulting spectral features (e.g. the soft X-ray contribution from a hot, optically thick accretion disk versus the hard X-ray contribution from a Comptonizing corona) can give key insights into the properties of the black hole and its surroundings. The origin and nature of the black hole’s accretion mechanisms, the disk and coronal geometries, and interaction with other companion stars and gas not only allow us to study stars and the host galaxy, but it can also help determine the properties of the black hole itself (e.g. mass, spin, gravity and GR, etc.). A conservative list of some of the long standing questions are:

- Why do some X-ray sources appear to be more luminous than expected? Are these sources really stellar-mass objects or potentially intermediate mass black holes (in which case the observed luminosity would be reasonable given the

larger mass)?.

- What is the nature of the corona around black holes?
- How do light and matter behave under the influence of strong gravity?
- Why do we see more soft X-rays from ULXs and AGN compared to BHXBs?
- Are black holes born with high or low spins?
- Can we improve our theoretical models of black hole accretion?
- Can we reproduce observed X-ray spectra using numerical simulations?

We explore the last three questions in Chapter 2 by using X-ray spectral models to investigate the spin of the black hole in the BHXB GRS 1915+105. Chapters 3 and 4 explore the last question using the **Athena++** numerical radiation magnetohydrodynamics (RMHD) simulations along with a Monte Carlo (MC) radiation transfer module to generate simulated X-ray spectra to compare with observed spectra of the ultraluminous X-ray source (ULX) NGC 1313 X-1. At the heart of this dissertation is bridging the gap between observations of accreting black holes and the theoretical models and simulations of such systems. In the rest of this introductory chapter, we discuss the observable properties of astrophysical black holes, methods for analyzing black hole X-ray spectra, and the **Athena++** RMHD simulations and MC radiation transfer codes. In the last subsection, we give a brief outline of the subsequent chapters of this dissertation.

1.1 PROPERTIES OF BLACK HOLES

Although black holes are some of the most mysterious and extreme astrophysical objects in our universe, they are also the simplest objects to physically describe. Only three quantities are required to completely construct a picture of the black hole³: its mass, spin, and charge. The charge of an astrophysical black hole is expected to be extremely small and thus has a negligible effect on the spacetime curvature. For the purposes of this dissertation, we do not investigate black hole charge and assume it is effectively zero.

The mass of a black hole can range from a few solar masses to billions of solar masses. Measuring the black hole mass is typically accomplished by observing the relative motion of stellar companions in BHXBs. In the case of supermassive black holes, their mass is typically inferred by the relative motions of stars or gas in the

³Assuming general relativity is correct.

black hole’s vicinity. Most spectacularly, this was the subject of the joint Nobel Prize for detecting the supermassive black hole at our Galactic center (Baganoff et al., 2001; Genzel et al., 2003; Ghez et al., 2004). This dissertation will primarily investigate stellar mass black holes, particularly BHXBs. However, the methods and analyses described in Chapters 3 and 4 can be used for future studies of AGN. Chapter 4 focuses on ULXs, which are thought to either be stellar mass black holes or neutron stars in accreting binary systems whose luminosities surpass the Eddington limit (discussed in Section 1.2.2). An alternative explanation for ULXs is an IMBH whose accretion rate does not surpass its Eddington limit, although there is stronger evidence for the former explanation (Skinner et al., 1982; Poutanen et al., 2007; Bachetti et al., 2014; Middleton et al., 2015a,b; Pinto et al., 2016a, 2020; Walton et al., 2016; Kosec et al., 2021; King et al., 2023). In particular, we know that many ULXs must be accreting neutron stars from their observed pulsations (Pintore et al., 2017; Walton et al., 2018; Middleton et al., 2023).

The black hole spin is related to the angular momentum imparted onto the black hole during its formation (e.g. from the angular momentum of the collapsing massive star or gas clouds) and subsequent accretion. The dimensionless spin parameter a_* is defined in terms of the angular momentum J , the speed of light c , and the black hole mass M :

$$a_* = \frac{Jc}{GM^2} \quad (1.1)$$

The spin of the black hole has a direct effect on the innermost stable circular orbit (ISCO) of the accretion disk, defined as the last stable radius at which a particle can orbit the black hole. For a non-spinning black hole ($a_* = 0$), the location of the ISCO is at $r_{ISCO} = 6r_g$ where r_g is the gravitational radius defined by

$$r_g = \frac{GM}{c^2} \quad (1.2)$$

where G is the gravitational constant. The gravitational radius is related to the event horizon or the Schwarzschild radius of a black hole, $r_s = 2r_g$. This is the radius at which light becomes trapped. Inside the event horizon $r < 2r_g$ is the singularity of the black hole, where the spacetime curvature becomes infinitely large. Here GR is thought to break down and quantum gravitational effects may become important,

but consideration of such properties are beyond the scope of this dissertation.

Within the scope of this dissertation, the theoretical spin of a black hole ranges from $a_* = 0$ (for a non-spinning Schwarzschild black hole) to $a_* = 1$ (for a maximally spinning Kerr black hole). It is not thought that $a_* > 1$ as this would result in a "naked" singularity, which is a singularity that is not hidden behind an event horizon (see Roger Penrose's Cosmic Censorship hypothesis: [Penrose, 1978, 1999](#)). In turn, the locations of the ISCO ranges from $r_{\text{ISCO}} = 6r_g$ for a non-spinning black hole to $r_{\text{ISCO}} \sim r_g$ for a maximally spinning black hole. Astrophysically, a "maximally spinning" Kerr black hole would likely have a canonical spin limit $a_* \leq 0.998$ (and thus $r_{\text{ISCO}} = 1.23r_g$) due to a counter torque produced when some of the radiation from the disk falls into the black hole ([Thorne, 1974](#)). An even lower maximal spin limit of $a_* \lesssim 0.94$ was proposed by [Gammie et al. \(2004\)](#) taking into account the effects of MHD torques in the in-falling gas. Presumably during the collapse of a massive star or gas cloud, angular momentum is carried away by radiation and material expelled during the supernova explosion. Simulations of stellar collapse suggest that natal black hole spins are likely slow ($a_* \sim 0.01 - 0.72$; see e.g., [Lousto et al., 2010](#); [Gerosa & Berti, 2017](#); [Rodríguez et al., 2018](#); [Fuller & Ma, 2019](#)). However, spectral modeling of observed black holes often estimate much higher spinning black holes ([Miller et al., 2015](#); [Gallo et al., 2015](#); [Walton et al., 2017](#); [Mills et al., 2021](#); [Coughenour et al., 2023](#); [Draghis et al., 2023](#); [Feng et al., 2023](#); see also [Reynolds, 2021](#) for a review).

The question of how black holes born with low to moderate spins would be able to achieve such high or near maximal spins is not well understood. The effect on the spin via mass accretion from the binary companion star is not significant enough to explain the discrepancy ([King & Kolb, 1999](#); [Bavera et al., 2021](#)) implying one of the following scenarios: 1) that black holes may be born with much higher spins than the stellar evolution models predict, or 2) the spectral modeling methods may be overestimating the actual spins of black holes, or 3) black holes born with low spin could have undergone a short period of rapid or extreme accretion in order to significantly spin them up. In Chapter 2, we present an analysis of the black hole spin in the BHX system GRS 1915+105.

1.2 ACCRETION THEORY

Chapters 2 – 4 rely on the theoretical foundations of black hole accretion, which we briefly introduce here. [Shakura & Sunyaev \(1973\)](#) is one of the seminal papers which

set the foundation of accretion disk theory that is still used today. The following subsections describe the theoretical framework for the thin accretion disk model proposed by [Shakura & Sunyaev \(1973\)](#) as this is commonly used as a first approximation to describe observed black hole accretion disks. The last subsection describes when this model begins to break down in the cases of super-Eddington accretion, a process that has been suggested is present in ULXs and is important for understanding the project discussed in Chapter 4.

1.2.1 Thin accretion disk model

The thin disk model, often called the standard disk model or α -disk prescription, describes a geometrically thin, optically thick accretion disk. This model is widely used as a first-order approximation to black hole accretion disks around BHXBs and (less appropriately) ULXs. It also has applications in proto-planetary (early solar system) disks and the disks around neutrons stars and white dwarfs (cataclysmic variable stars). Many simplifying assumptions are made in order to describe a thin disk. We summarize some of the key points:

- Geometrically thin ($H \ll r$) and optically thick
- Conservation of mass, energy, and angular momentum
- The disk is not self-gravitating ($M_{\text{disk}} \ll M_{\text{BH}}$)
- The disk is axisymmetric and in steady-state (no outflows)
- Keplerian orbital velocity is much larger than the radial velocity
- Stress is proportional to the total pressure (P) and viscosity parameter (α)

For the disk to be geometrically thin, the vertical scale height of the disk H must be much smaller than the disk's radial extent r . The optically thick accretion disk has a viscosity parameter α which describes the transport of angular momentum outward via stress between adjacent layers in the disk,

$$\tau_{r\phi} = \alpha \tag{1.3}$$

Consider the scenario of two adjacent layers of gas in the disk. The inner layer does work on the outer layer during their viscous interaction, transporting angular momentum out and allowing material to spiral into the black hole. Molecular viscosity, however, is not large enough to account for the inferred accretion rates ([Pringle, 1981](#)), leading to a nearly two decade puzzle of what process could account for this

anomalously high viscosity. It was realized that turbulent magnetic fields driven by the magnetorotational instability (MRI) can drive black hole accretion (Balbus & Hawley, 1991). In this dissertation we do not discuss the impact of magnetic fields at length, however they are a necessary component underlying the RMHD simulations in Chapter 4.

In the thin disk model, the flux radiated from the surface of the disk per unit time is given by

$$F(r) = \frac{3}{8\pi} \frac{GM\dot{M}}{r^3} \left[1 - \left(\frac{r_0}{r} \right)^{1/2} \right] \quad (1.4)$$

where r_0 is defined as the inner edge of the disk where the stress is assumed to vanish, \dot{M} is the mass accretion rate onto the black hole, and r is the disk radius. For black holes, r_0 is chosen to be the innermost stable circular orbit (ISCO). This expression is for a Newtonian potential, but can be generalized to relativistic spacetimes (Novikov & Thorne, 1973).

If we assume that the accretion disk is a blackbody, we can start from the Stefan-Boltzmann law and approximate the flux emergent from an isotropically emitting blackbody as a function of temperature:

$$F(r) = \sigma T_{\text{eff}}^4(r) \quad (1.5)$$

where σ is a constant defined as $\sigma = 5.67 \times 10^{-5} \text{ erg cm}^{-2} \text{ deg}^{-4} \text{ s}^{-1}$, and T_{eff} represents the effective temperature of the disk. We can then get the effective temperature of the disk as a function of mass M , accretion rate \dot{M} , and radius using Equation 1.4,

$$T_{\text{eff}}(r) = \left(\frac{F(r)}{\sigma} \right)^{1/4} \approx \left(\frac{3GM\dot{M}}{8\pi\sigma r^3} \right)^{1/4} \quad (1.6)$$

We see that the effective temperature scales with the mass of the black hole, the mass accretion rate, and the radius. At the innermost radius, the effective temperature is then

$$T_{\text{eff}} \propto \left(\frac{M\dot{M}}{r_0^3} \right)^{1/4} \quad (1.7)$$

and since r_0 represents the ISCO which scales with the gravitational radius, it is

proportional to M . So, the effective temperature for the inner disk then becomes

$$T_{\text{eff}} \propto \left(\frac{M\dot{M}}{M^3} \right)^{1/4} \propto \left(\frac{\dot{M}}{M^2} \right)^{1/4} \propto \left(\frac{\dot{m}}{M} \right)^{1/4} \quad (1.8)$$

where the following dimensionless parameters are defined (Shakura & Sunyaev, 1973):

$$m = \frac{M}{M_{\odot}} \quad (1.9)$$

$$\dot{m} = \frac{\dot{M}}{\dot{M}_{\text{Edd}}} \quad (1.10)$$

\dot{M}_{Edd} is the critical Eddington accretion rate at which the total energy released in the disk L is equal to the Eddington critical luminosity L_{Edd} . The effective temperature at the ISCO thus scales with just the mass of the black hole and its mass accretion rate. The outer radii of the disk have subsequently lower temperatures (assuming a fixed mass accretion rate). Toward the inner radii of the disk, the temperature is higher. In the context of X-ray observations, the spectrum of black hole accretion disks typically show a modified blackbody spectrum with lower (softer) X-ray energies likely coming from further out in the accretion disk, while higher (harder) X-ray energies likely originate from the innermost disk regions. This also explains why X-ray binaries (on the order of tens of solar masses) peak in the X-ray wavelengths, and why AGN (on the order of millions to billions of solar masses) peak in the optical/UV wavelengths.

Astrophysical accretion disks are not expected to be perfect blackbody emitters, however it serves as a good first-order approximation. The standard thin disk is applicable for accretion rates that remain relatively low such as those of BHXBs, but what happens when the mass accretion rate is much higher? In the case of super-Eddington accretion flows, the disk is no longer geometrically thin and is thought to have radiatively driven winds/outflows (Shakura & Sunyaev, 1973).

1.2.2 Super-Eddington accretion

ULXs are point-like, off-nuclear (not AGN) extragalactic objects observed to have X-ray luminosities comparable to, or in excess of, the critical Eddington luminosity $L_X \gtrsim 10^{39}$ erg/s (assuming isotropic emission for a $10 M_{\odot}$ black hole; see Pinto &

Walton, 2023; King et al., 2023 for a review of ULXs). The Eddington limit is the point at which the outward radiation pressure force near the black hole exceeds the inward force due to gravity, thus driving outflows and winds. This limit is defined for a gas of pure ionized hydrogen (Rybicki & Lightman, 1979):

$$L_{\text{Edd}} = \frac{4\pi GMcm_{\text{p}}}{\sigma_{\text{T}}} = 1.25 \times 10^{38} \left(\frac{M}{M_{\odot}} \right) \text{ erg s}^{-1} \quad (1.11)$$

where G is the gravitational constant, c is the speed of light, m_{p} is the mass of a proton, σ_{T} is the Thomson scattering cross-section, and M is the mass of the black hole. For a $10 M_{\odot}$ black hole, the Eddington luminosity is $L_{\text{Edd}} = 1.25 \times 10^{39} \text{ erg s}^{-1}$.

We can relate the Eddington luminosity to its associated Eddington mass accretion rate \dot{M}_{Edd} and dimensionless radiative efficiency parameter η :

$$L = \eta \dot{M} c^2 = L_{\text{Edd}} = 10^{38} \frac{M}{M_{\odot}} \text{ erg s}^{-1} \quad (1.12)$$

The radiative efficiency $0 \leq \eta \leq 1$ signifies how efficient the accretion flow is at being converted into escaping radiation via outflows. If all of the accreted matter went into the black hole without radiating, then $\eta = 0$. The efficiency is inversely proportional to the mass accretion rate, so a larger \dot{M} would imply a lower η (for a fixed bolometric luminosity). Processes like advection (Abramowicz et al., 1988) and radiatively driven outflows (Shakura & Sunyaev, 1973; Ohsuga & Mineshige, 2011) may reduce radiative efficiency and be associated with geometrically thicker flows in the super-Eddington regime. Strong optically thick winds are also expected to be launched in ULXs (Middleton et al., 2014, 2015b; Pinto et al., 2016b, 2020; Walton et al., 2016; Kosec et al., 2021), which likely shroud the outer accretion disk and can contribute additional low energy flux for preferential sight lines. We discuss the radiative efficiencies calculated from our RMHD simulations in Chapter 4.

The majority of ULXs are now commonly thought to be X-ray binary systems with super-Eddington accretion rates onto a compact object, namely a stellar mass black hole (Poutanen et al., 2007; Middleton et al., 2015a) or neutron star (Skinner et al., 1982; Bachetti et al., 2014). Some small fraction of ULXs may yet harbor sub-Eddington accretion rates onto IMBHs (Farrell et al., 2009; Mezcua et al., 2013; Earnshaw, 2016; Brightman et al., 2016; Webb et al., 2017; Oskinova et al., 2019). The physical mechanisms which drive super-Eddington accretion are still under in-

investigation and require numerical simulations in order to evaluate existing models of black hole accretion. If some ULXs are indeed IMHBs, the spectra are expected to resemble scaled up versions of BHXB spectra, showing cooler accretion disks as the black hole mass increases (e.g. [Miller et al., 2004](#)).

1.3 X-RAY OBSERVATIONS: METHODS, MODELING, AND INTRODUCTIONS

Spectral observations of BHXBs typically show a soft, thermal emission component along with a second, harder X-ray component. The soft component is widely believed to be emission from an optically-thick, geometrically thin accretion disk ([Shakura & Sunyaev, 1973](#)), while the hard component is thought to be Comptonized emission from hot electrons near the disk (the corona). Most BHXBs display variability between spectral states where the relative strengths of these components vary, with the high/soft state referring to cases where the disk component dominates (so called disk-dominated states; [Remillard & McClintock, 2006](#); [Done et al., 2007](#)).

Observations of ULXs, however, typically show two thermal components consisting of a soft, thermal X-ray component and a hard thermal component with a rollover beginning below ~ 10 keV ([Gladstone et al., 2009](#); [Bachetti et al., 2014](#)). The latter supports the interpretation of super-Eddington accretion onto a compact object. Early models debated whether this hard X-ray emission originates from coronal emission from IMBHs ([Miller et al., 2004](#)) or Comptonized emission from super-Eddington accretion ([Gladstone et al., 2009](#); [Socrates & Davis, 2006](#)). However, classically, one would expect the innermost regions to have a different spectral shape due to optical depth effects and anisotropy ([Poutanen et al., 2007](#)). In Chapter 4 we use a combination of commonly used phenomenological spectral models that are typically used to fit BHXB spectra. Although the disk structure of ULXs is expected to be different from that of BHXBs, we find that fitting a combination of these spectral fitting models can phenomenologically fit the spectra of the ULX NGC 1313 X-1. We discuss the relevant implications and caveats of this in Chapter 4.

In the next section we discuss one observational method for analyzing BHXB spectral observations to determine the spin parameter of black holes. This method is used to determine the spin of the black hole in GRS 1915+105, the results of which are presented in Chapter 2.

1.3.1 Continuum-fitting method

When BHXBs enter strongly disk-dominated high/soft states, one might expect the emission to be well-represented by a bare accretion disk model, which accounts for the relativistic effects on photon emission and the possible change in flow properties at or near the black hole’s ISCO. This has motivated a number of relativistic accretion disk spectral models (Hanawa, 1989; Gierliński et al., 1999; Li et al., 2005; Davis & Hubeny, 2006), which have had success in fitting the spectrum and its variation with accretion rate in the high/soft state of many BHXBs (Gierliński & Done, 2004; Davis et al., 2005; Shafee et al., 2006; McClintock et al., 2011). Spectral fits are typically performed using X-ray spectral fitting software such as XSPEC (Arnaud, 1996), which provides a number of spectral models used to fit the data.

Placing constraints on the black hole angular momentum, or spin, is a key motivation of many such studies (for a review, see Middleton, 2016). The spectrum of the disk is sensitive to the spin through location of the ISCO as well as the relativistic effects on the photon propagation through the black hole spacetime. This technique of spin measurement is generally referred to as the continuum-fitting method (Zhang et al., 1997; Done et al., 2007; McClintock et al., 2011), which distinguishes it from other spectral-fitting spin measurements such as those that fit the reflected emission features, including the prominent Fe $K\alpha$ line (Fabian et al., 1989), high frequency quasi-periodic oscillations (Motta et al., 2022), as well as polarization constraints (Mikusinova et al., 2023). For the purposes of this dissertation, we only discuss the continuum-fitting method.

Typically when fitting for black hole spin, the disk-dominated continuum part of the spectrum is usually fit with some accretion disk model such as a multi-temperature disk blackbody model `diskbb` (Mitsuda et al., 1984), a similar but more self-consistent accretion disk model `bhspec` (Davis et al., 2005; Davis & Hubeny, 2006), or an accretion disk model which includes a color correction (or sometimes called the spectral hardening) parameter `kerrbb` (Li et al., 2005). These models all incorporate the assumptions of a thin accretion disk (Shakura & Sunyaev, 1973), and are discussed more in Chapter 2.

Determining black hole spin thus relies on current observational capabilities for measuring the emission coming from near r_{ISCO} . Observing the inner accretion disk is also subject to determining the accretion disk luminosity and temperature, which for

spinning black holes is larger and hotter than that of non-spinning black holes at the same mass accretion rate. Observations of the disk luminosity and temperature as well as the radius of the ISCO would allow the spin of the black hole to be determined, except that black hole spectra are not as simple as we could hope (given that black holes are relatively simple objects themselves). One might assume that fitting a `diskbb` or `bhspec` model would be enough to acquire the disk continuum emission. However there are several components which contaminate the accretion disk spectrum, each of which need to be carefully modeled such that the disk emission can be isolated.

The disk continuum can be contaminated by the presence of absorption or emission lines coming from the intervening ISM, the presence of any local gas clouds, as well as the Fe-K α line seen at approximately 6.4 – 7 keV (Basko, 1978; Fabian et al., 1989). The neutral hydrogen column along the line of sight is usually modeled with an absorption component such as `TBabs` or `varabs`, whereas atomic transition lines are usually modeled with a few Gaussian absorption/emission components such as `gabs`.

In addition to emission and absorption contaminants, there is also the hard X-ray coronal component which modifies the tail end of spectrum out to hard X-ray energies, causing a flatter profile than the standard blackbody Wien tail. In terms of fitting the X-ray continuum, the hard X-ray tail can contribute a significant fraction to the bolometric luminosity, thus contaminating the continuum (Kubota et al., 2004). This high energy tail is usually fit with one of a number of models including power law models `pow` or a cutoff power law model `cutoffpl`, or Compton scattering models such as `nthcomp` (Zdziarski et al., 1996; Życki et al., 1999) or `simpl` (Steiner et al., 2009).

For some context in Chapter 2, we discuss the continuum-fitting methods used in two different papers for the BHXB GRS 1915+105: Middleton et al. (2006) and McClintock et al. (2006). The continuum-fitting model combination used by Middleton et al. (2006) included `varabs`, a variable absorption model to account for the column density along the line of sight to GRS 1915+105, two different models for the accretion disk (`diskbb` and `bhspec`), a Comptonization model (`nthcomp`), a Gaussian model centered at 6 – 7 keV to account for the location of the broad Fe K α line, and a smeared edge model `smedge` (Ebisawa, 1991) to account for the smeared edge of the broad Fe K α line at 6.9 – 9 keV. McClintock et al. (2006) used a similar model combination as Middleton et al. (2006) except they utilized `phabs`, a photoelectric

absorption model instead of `varabs`, and a modified disk model which includes both the relativistic `kerrbb` disk model and `bhspec`. They also used a combination of `pow` and the usual `Gaussian` and `smedge` to model the Fe $K\alpha$ line and the high energy tail. [McClintock et al. \(2006\)](#) claimed that their results are “robust to the details of the analysis,” that is, they depend weakly on the choice of the model for the high energy tail component, and weakly depend on the `Gaussian` and `smedge` models for getting a good fit. We discuss the details of these continuum-fitting models in the thesis Chapter.

Due to the sensitivity of this method on isolating the disk continuum emission, BHXBs are largely better candidates for employing this method. In theory this could be used for ULXs and AGN, but more accurate models of their accretion disk structures are needed if this method is to be useful for determining black hole spin in those sources.

1.3.2 BHXB: GRS 1915+105

In Chapter 2, we discuss the black hole spin results for the BHXB GRS 1915+105. Figure 1.3 shows a combined optical and infrared image of the sky containing the X-ray source, with the smaller image in the upper left showing GRS 1915+105 in X-rays. GRS 1915+105 is a Galactic microquasar discovered in 1992, exhibiting extraordinary observational features including superluminal radio jets and extremely variable X-ray emission. This is a very unusual system compared to most other Galactic BHXBs with low mass companions because of its longevity in outburst (although it has recently entered some low luminosity “obscured” state, see [Imazato et al., 2021](#)). The lower right image in Figure 1.3 shows the “heartbeat” of GRS 1915+105 in X-rays observed over time, showing the periodic variability this system exhibits. The GRS 1915+105 system includes a low-mass companion star whose Roche lobe overflow fuels the accretion disk around the black hole.

Since its discovery, GRS 1915+105 has continuously switched between a “hot and bright” spectral state (implying a smaller ISCO) to a “cooler and dimmer” state (implying a larger ISCO). This state switching has been attributed to an instability in the inner disk where the disk is constantly emptying and refilling. There is particular interest in the spin of GRS 1915+105, partially because it has been observed more than any other Galactic microquasar and has evidence of relativistic jets. However, the variability in GRS 1915+105 presents quite a challenge for selecting a purely disk-




Figure 1.3: An optical and infrared image from the Digitized Sky Survey which shows a crowded field of view around the black hole X-ray binary system GRS 1915+105, located near the plane of our Milky Way Galaxy. The top-left image shows a close-up of GRS 1915+105, taken by the *Chandra* X-ray telescope. The bottom right image shows the “heartbeats” of GRS 1915+105 in X-rays over time. GRS 1915+105 contains a black hole that is about $12 - 14M_{\odot}$ and is accreting material from its companion star. Credits: X-ray: NASA/CXC/Harvard/J. Neilsen et al.; Optical: Palomar DSS2.

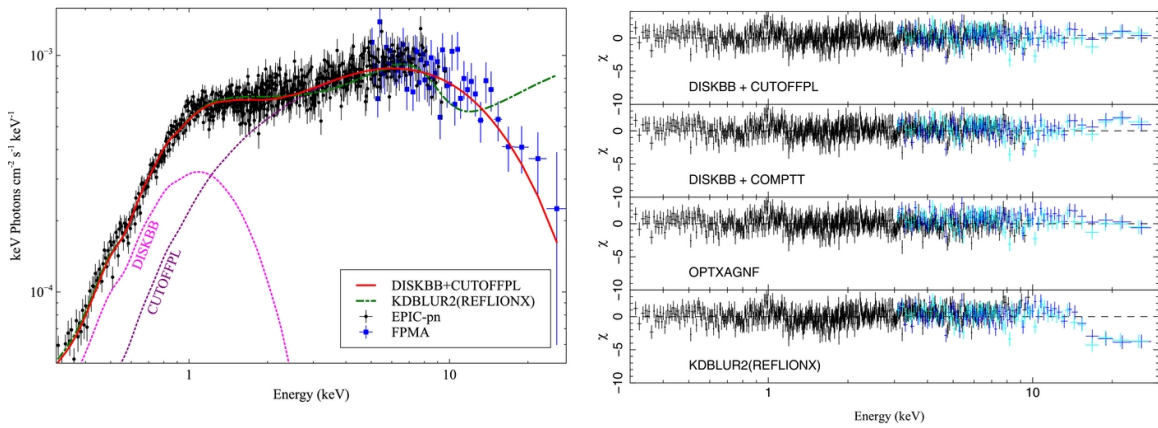


Figure 1.4: Figure 4 from [Bachetti et al. \(2013\)](#) showing the combined XMM-Newton and NuSTAR data for NGC 1313 X-1 plus several model fits of the data. The left plot shows the XMM-Newton EPIC-pn data represented by the black circles, and the NuSTAR Focal Plane Module A data is represented by the blue squares. The solid red line is the combined `diskbb+cutoffpl` model, with the individual model components shown by the dashed pink and purple lines, respectively. The best-fit reflection model is shown by the dashed green line. The right plot shows the fit residuals from several models chosen by the authors: XMM-Newton EPIC-pn (black), NuSTAR Focal Plane Module A (blue) and Focal Plane Module B (light blue).

dominated spectral state, which is crucial for determining the black hole spin using the continuum-fitting method, which we describe in more detail in Chapter 2.

The production of relativistic jets is a topic of ongoing research, but the jets are believed to be related to the spin of the black hole via the Blandford-Znajek effect (Blandford & Znajek, 1977), as well as magnetic flux carried by the accretion of material from the companion star. Many black hole binary systems in the Milky Way exhibit signs of relativistic jets, so measuring the spin of these systems may provide insights into the jet-spin relation. Middleton et al. (2006) claim that jet-powered systems likely do not require maximal spin, as there have been moderately spinning BHXBs (determined via the continuum-fitting method) with the presence of jets (Blaes et al., 2006; Shafee et al., 2006). Observations seem to suggest that jet power scales with accretion rate, favoring the idea that relativistic jets are probably produced by gravitational means. This has yet to be conclusively tested, due to the difficulty of measuring black hole spin and other properties of the BHXB system (e.g. inclination of the disk viewed from a perspective on Earth).

1.3.3 ULX: NGC 1313 X-1

The other black hole system of interest discussed in Chapter 4 is the ULX source, NGC 1313 X-1. The spiral galaxy NGC 1313 hosts two of the brightest, closest, and hardest ULXs: NGC 1313 X-1 and NGC 1313 X-2. Before the hard X-ray telescope *NuSTAR* was launched, there was some speculation as to what the spectral shape above 10 keV was for this class of BHXBs. A spectral break or cutoff around 10 keV was proposed by Gladstone et al. (2009), evidenced by the *XMM-Newton* archival data for several sources. Below 10 keV, at least two models were thought to describe ULX spectra: (1) a reflection-dominated regime where the spectral break is caused by a relativistically broadened iron line and edge, and (2) a low-temperature Comptonized emission that breaks right before 10 keV. The addition of *NuSTAR* data broke this degeneracy. Bachetti et al. (2013) show a distinct cutoff above 10 keV for NGC 1313 X-1 concluding that the ULX is most likely a stellar mass black hole ($70\text{-}100M_{\odot}$) accreting close to its Eddington limit with a cold corona obscuring the inner part of the disk. This is consistent with the low-temperature Comptonized emission model posed by Gladstone et al. (2009) and not with the reflection-dominated regime model as shown in the left plot in Figure 1.4. The data for NGC 1313 X-1 are shown in black circles for *XMM-Newton* EPIC-pn, and blue circles for *NuSTAR* Focal Plane

Module A (FPMA). The dashed green line represents the reflection model, and the red line represents the combined `disk+cutoffpl` model (its separate components are shown in pink and purple). Above ~ 10 keV, the `diskbb+cutoffpl` model fits the data well, whereas the reflection model does not accurately capture the downturn in the spectrum.

1.4 RADIATION MAGNETOHYDRODYNAMICS AND MONTE CARLO RADIATION TRANSFER

Understanding the X-ray spectra of astrophysical black holes requires using numerical simulations in order to replicate the complex multi-dimensional environments surrounding black holes. These simulations aim to include all the necessary physics and theoretical frameworks, including GR and radiation transfer. Although the thin accretion disk theory is typically a good first approximation to use, there are many astrophysical ingredients (such as magnetic fields or outflows) that can complicate or completely violate this theoretical framework. Our goal is to use RMHD simulations to replicate super-Eddington accretion flows onto a stellar mass black hole and study its spectral properties in comparison with observations. The complex three-dimensional nature of the physical accretion processes in BHXBs, ULXs, and AGN motivates using numerical simulations to study these systems. Specifically, we use a radiation magnetohydrodynamics (RMHD) code, `Athena++` (Stone et al., 2020), which we summarize briefly in the next subsection.

`Athena++` incorporates the MHD equations and a treatment of the radiation transfer equation in order to simulate the turbulent, radiation mediated flows around a black hole. These simulations are computationally expensive, thus some simplifying assumptions are made in order to run a full calculation. For example, the simulations are “gray” or frequency integrated in order to simplify the RMHD equations and reduce computation time. However, in order to look at the X-ray spectra of these simulations, this requires some post-processing in order to extract the frequency information. The primary way we do this is by using a Monte Carlo radiation transfer code, which we discuss in subsection 1.4.2.

1.4.1 `Athena++` RMHD code

`Athena++` is an RMHD code written in C++ and contains a number of improvements over its predecessor, `Athena` (Stone et al., 2008). It now includes flexible

coordinates and grids including adaptive mesh refinement (AMR), general relativity, and improved performance and modularity. It adopts a dynamical task-list design in which overlapping communication between code cells can occur in any order (as opposed to a strict hard-coded order where some cells must wait until other cells have finished computing), thus improving the parallel performance and lowering the computational expense of the code. It also incorporates a block-based mesh refinement structure which exists independently such that it allows a variety of discretizations (e.g. cell-center volume averages, face-center area averages, etc.) to be calculated and stored on the mesh. This also allows the user to specify the desired refinement level in chosen domain blocks, utilizing a new AMR functionality (Tomida & Stone, 2023) where the computational grid originally divided into structured mesh blocks can be refined into smaller mesh blocks. In the case of AMR, the number and location of the mesh blocks can change as the simulation runs. This is useful for simulations that wish to fully resolve only certain parts of the computational domain (e.g. near the black hole) while keeping the rest of the domain at coarser refinement level. This way, the computational time can be reduced while the relevant simulation properties can be analyzed with high resolution. For the purposes of this work, we only consider static (and uniform) mesh refinement, as the version of the MC transfer module at the time of this work could easily be used with RMHD simulations that had a static mesh refinement.

Athena++ adopts a finite-volume method and a time-dependent, angular resolving radiation transfer module (Jiang, 2021) where the following equations are solved:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0, \quad (1.13)$$

$$\frac{\partial(\rho \mathbf{v})}{\partial t} + \nabla \cdot (\rho \mathbf{v} \mathbf{v} - \mathbf{P}) = -\mathbf{G} + \rho \mathbf{a}_{\text{grav}} \quad (1.14)$$

$$\frac{\partial E}{\partial t} + \nabla \cdot [(E + P)\mathbf{v}] = -cG^0 + \rho \mathbf{a}_{\text{grav}} \cdot \mathbf{v}, \quad (1.15)$$

$$\frac{\partial I}{\partial t} + c\mathbf{n} \cdot \nabla I = cS_I \quad (1.16)$$

$$S_I \equiv \Gamma^{-3}[\rho(\kappa_s + \kappa_a)(J_0 - I_0) + \rho(\kappa_a + \kappa_{\delta P})\left(\frac{a_{\text{R}}T^4}{4\pi} - J_0\right)] \quad (1.17)$$

$$cG^0 \equiv 4\pi c \int S_I d\Omega \quad (1.18)$$

$$\mathbf{G} \equiv 4\pi \int \mathbf{n} S_I d\Omega \quad (1.19)$$

Equations 1.13 – 1.15 are the hydrodynamic fluid equations where ρ , \mathbf{v} , $E = E_{\text{g}} +$

$(1/2)\rho v^2$ are the fluid density, velocity and total energy density. Assuming an ideal gas, the internal gas energy E_g can be related to the pressure P by $E_g = P/(\gamma - 1)$. The pressure tensor is represented as \mathbf{P} . The right-hand sides of Equations 1.14 and 1.15 represent the coupling between the radiation and the gas through the components of the radiation four force $-\mathbf{G}$ and $-cG^0$, which are the momentum and energy that the radiation imparts to the gas. The generalized Newtonian gravitational force and gravitational potential are given by $\rho \mathbf{a}_{\text{grav}}$ and $\rho \mathbf{a}_{\text{grav}} \cdot \mathbf{v}$, respectively.

Equations 1.16 – 1.19 represent the frequency-integrated “gray” radiation transfer equation. The source terms are coupled to the momentum and energy of the gas. In the *Athena++* implementation, the specific intensities I are first transformed to the comoving frame of the fluid and the radiation source terms S_I , \mathbf{G} , and cG^0 are evaluated and updated implicitly along with Equation 1.15 in the comoving frame. The intensities are then transformed back into the lab frame, and the resulting source terms are evaluated by integrating the intensities over angle as described in Jiang (2021).

I_0 and J_0 are the comoving frame intensity and mean intensity, and c is the speed of light. \mathbf{n} is a unit vector which corresponds to rays in discretized angular grids, with the same angular discretization in Davis et al. (2012). Equations 1.18 and 1.19 represent the impact of the radiation on the momentum and energy of the gas. In the source term, a_R is the radiation constant, κ_s and κ_a are the scattering opacity and Rosseland mean absorption opacity, and $\kappa_{\delta P}$ is the difference between the Planck mean and Rosseland mean opacity. This utilizes free-free opacity for κ_a based on the local gas density and temperature, assuming $\kappa_s = 0.34 \text{ cm}^2 \text{ g}^{-1}$. *Athena++* solves the unit-less equations with the scaling of density unit $\rho_0 = 10^{-10} \text{ g cm}^{-3}$, velocity unit $v_0 = 0.01c$ and length unit $l_0 = r_s = 2.95 \times 10^{12} \text{ cm}$.

The *Athena++* simulation outputs HDF5 files with the relevant quantities (e.g. flux, radiation energy densities, pressures, velocities, etc.) which we then use in our post-processing analysis. A more detailed description of black hole simulation setup is described in Chapter 4. We discuss the use of simulation snapshots, which are a snapshot of the simulation at a moment in time. These 3D snapshots are then used in our post-processing analysis which we describe more in Chapters 3 and 4.

1.4.2 Monte Carlo radiation transfer

Here we briefly introduce MC methods generally. In Chapter 3, we give an overview of the `Athena++` Monte Carlo (MC) radiation transfer module (Davis et al., 2009, Davis et al. in prep) and present code tests to show its robustness. In Chapter 4, we use the MC module to post-process `Athena++` RMHD simulation snapshots to generate X-ray spectra which we compare with X-ray observations.

Monte Carlo methods have a wide range of applications in non-astrophysical fields that are interested in the emission and transmission of light particles. For example, MC methods are used in biomedical research to understand how focused light interacts with human tissue for applications in oncology and radiation therapy (Doronin & Meglinski, 2011; Sarrut et al., 2014; Watson et al., 2020) as well as modeling brain pathways in mice (Tang & Yao, 2021)⁴. MC methods are generally useful for gathering statistics and representing the possible outcomes of complex multi-variable systems. In the context of this dissertation, we use MC to simulate photons of light near a black hole.

The MC method utilizes a large number of test particles. These are sometimes called photon packets, or super photons, but we prefer the term photon samples, since they are sampling the statistical properties of the radiation field. The samples are initialized and transported in some simulated domain of a specified size. Wood et al. (2013) framed the basic concept of MC radiation transfer methods quite nicely:

A photon is emitted, it travels a distance, and something happens to it.

One can imagine a box that is sectioned into smaller block regions, with particles distributed across all of the blocks. The particles are then allowed to move freely in any direction, depending on the specific conditions set by the researcher (e.g. applying specific laws of physics such as gravity). The MC method then facilitates any movement of the particles across different regions within the box. In the context of our MC setup, we could allow particles to undergo scattering physics which can change the direction, energy, and statistical weights of the scattered particles involved. When a particle reaches say the top of the computational box, we can set boundary conditions such as designating that the particle should now be removed (escaped or absorbed), destroyed completely, or if the boundary is periodic then the particle will

⁴The authors use MC on a digital mouse brain atlas called “Digimouse”

re-enter the reverse side (bottom) of the box and continue. This process continues until all of the particles either escape the simulation domain, are absorbed, or are terminated by some other user-specified condition.

An important characteristic of any MC method is that it employs pseudorandom⁵ number sampling based on a given probability distribution. Statistically, a sufficiently large number of particles or samples should be used in order to effectively represent true physical processes. After the MC simulation is completed (e.g. all particles are finished moving) we can then analyze how many particles underwent the specific processes relevant to our scientific motivations such as creating a spectrum of all particle energies after undergoing scattering and absorption. We discuss the details of our MC methods in Chapter 3.

One downside of utilizing MC radiation transfer methods is stochastic error (counting statistics) that are always prevalent. The signal-to-noise in the MC calculation scales as $\sqrt{N_{\text{ph}}}$. To decrease the stochastic noise in the calculation, the number of photon samples must be increased. However, increasing the number of photon samples also increases the computation time, especially if scattering is enabled in optically thick regimes. Then the number of scatterings scales as τ^2 where τ is the optical depth. Some of the limitations we discuss in Chapter 4 are related to the computational expense brought on by the number of photons that we could feasibly simulate.

1.5 OUTLINE OF THIS THESIS

Chapter 2 is adapted from an article published in the *Astrophysical Journal* (Mills et al., 2021). This work re-examined the black hole spin estimates for the BHXB GRS 1915+105 from two previous papers (McClintock et al., 2006; Middleton et al., 2006). Following those publications, new constraints on the mass, distance, and disk inclination of GRS 1915+105 (Reid et al., 2014) motivates analyzing the change in the previous spin estimates. In this Chapter we discuss the impact of these new constraints on the spin estimate of the black hole in this BHXB source.

Chapter 3 is adapted from a co-authored article in preparation (S. W. Davis, B. S. Mills, et al., in preparation). We present the *Athena++* Monte Carlo radiation transfer module, its capabilities and methodology, relevant tests of the code, and its

⁵Pseudorandom because computers are, of course, deterministic.

usage and application in Chapter 4.

Chapter 4 is based on an article submitted to the *Astrophysical Journal* (Mills et al. 2023, submitted to ApJ). In this Chapter we use the Monte Carlo module from Chapter 3 to post-process spectra from a set of **Athena++** radiation magnetohydrodynamic simulations of super-Eddington accretion onto a stellar mass black hole. We discuss the comparison of these simulated spectral models with commonly used phenomenological models (see Chapter 2) and direct fits to the observed spectrum of the ultraluminous X-ray source, NGC 1313 X-1.

Finally, Chapter 5 summarizes the main points of each chapter in this dissertation and conveys the impact of this work on bridging the gap between black hole theory and observation. We also discuss future extrapolations of this work in addition to future developments of the Monte Carlo module.

CHAPTER 2

THE BLACK HOLE SPIN IN GRS 1915+105, REVISITED

2.1 ABSTRACT

We estimate the black hole spin parameter in GRS 1915+105 using the continuum-fitting method with revised mass and inclination constraints based on the very long baseline interferometric parallax measurement of the distance to this source. We fit Rossi X-ray Timing Explorer observations selected to be accretion disk-dominated spectral states as described in McClintock et al. (2006) and Middleton et al. (2006), which previously gave discrepant spin estimates with this method. We find that, using the new system parameters, the spin in both datasets increased, providing a best-fit spin of $a_* = 0.86$ for the Middleton et al. data and a poor fit for the McClintock et al. dataset, which becomes pegged at the BHSPEC model limit of $a_* = 0.99$. We explore the impact of the uncertainties in the system parameters, showing that the best-fit spin ranges from $a_* = 0.4$ to 0.99 for the Middleton et al. dataset and allows reasonable fits to the McClintock et al. dataset with near maximal spin for system distances greater than ~ 10 kpc. We discuss the uncertainties and implications of these estimates.

This chapter is adapted from an article published in the *Astrophysical Journal*. The original citation is as follow: B. S. Mills, S. W. Davis, M. J. Middleton. “The Black Hole Spin in GRS 1915+105, Revisited.” *Astrophysical Journal*, 914:1, June 2021.

2.2 INTRODUCTION

A soft, apparently thermal emission component is frequently observed in the spectra of black hole candidate X-ray binaries (hereafter BHXRBs) along with a second, harder X-ray component. The soft component is widely believed to be emission from an optically-thick, geometrically thin accretion disk (Shakura & Sunyaev, 1973), while the hard component is thought to be Comptonized emission from hot electrons near the disk (the corona). Most BHXRBs display variability between spectral states where the relative strengths of these components vary, with the high/soft state referring to cases where the disk component dominates (Remillard & McClintock, 2006; Done et al., 2007).

When BHXRBs enter strongly disk-dominated high/soft states, one might expect the emission to be well-represented by a bare accretion disk model, which accounts for the relativistic effects on photon emission and the possible change in flow properties at or near the black hole’s innermost stable circular orbit (ISCO). This has motivated a number of relativistic accretion disk spectral models (Hanawa, 1989; Gierliński et al., 1999; Li et al., 2005; Davis & Hubeny, 2006), which have had success in fitting the spectrum and its variation with accretion rate in the high/soft state of many BHXRBs (Gierliński & Done, 2004; Davis et al., 2005; Shafee et al., 2006; McClintock et al., 2011). Placing constraints on the black hole angular momentum, or spin, is a key motivation of many such studies (for a review, see Middleton, 2016). The spectrum of the disk is sensitive to the spin through location of the ISCO as well as the relativistic effects on the photon propagation through the black hole spacetime. This technique of spin measurement is generally referred to as the continuum-fitting method (Zhang et al., 1997; McClintock et al., 2011), which distinguishes it from other spectral-fitting spin measurements such as those that fit the reflected emission features, including the prominent Fe K α line (Fabian et al., 1989).

The continuum-fitting method has previously been applied to the BHXRB GRS 1915+105, yielding inconsistent estimates for the black hole spin (McClintock et al., 2006; Middleton et al., 2006). Although these studies used similar spectral models and fitting methods, Middleton et al. (2006) (hereafter MID06) found a more moderate spin parameter ($a_* \sim 0.7$) while McClintock et al. (2006) (hereafter MCC06) favored high spin ($a_* \gtrsim 0.98$), where we define the dimensionless spin parameter $a_* = Jc/(GM^2)$ and J is the angular momentum of the black hole. This discrepancy

in spin can be primarily attributed to differences in the selection of spectra used in their analyses. These differences arise from the difficulties in unambiguously identifying a disk-dominated state in this source, which is famous for its complex variability, with a diversity not generally seen in other low mass X-ray binaries (Belloni et al., 2000) except for the black hole source IGR J17091-3624 (Altamirano et al., 2011).

Both studies focused on analysis of spectral observations using data taken by the PCA detector on board the Rossi X-Ray Timing Explorer (hereafter *RXTE*). MCC06 identified a selection of apparently disk-dominated spectra in the 3-25 keV range based on a number of screening criteria, including RMS variability and hardness ratio, ultimately arriving at 20 candidates. MID06 argued against identifying the MCC06 sample as disk-dominated, instead arguing these observations are more like very high/steep power law states (Remillard & McClintock, 2006; Done et al., 2007), in which a low temperature Comptonization component is present (Zdziarski et al., 2005, MID06). Instead, MID06 generated a large library of spectra in the 3-20 keV range with 16 second exposures from observations within the β and κ variability classes of Belloni et al. (2000), ultimately identifying 34 disk-dominated candidates. MCC06, in turn, criticized this selection process, raising concerns about potential systematic errors arising from the short 16 second exposures while also worrying that the implied luminosity of observations were above or sufficiently close to the Eddington limit as to invalidate the assumptions of the underlying disk model. In contrast, MCC06 had focused their spectral analysis on relatively low Eddington observations, where the assumption of a thin accretion disk is more self-consistent. In fact, the highest luminosity observations among the MCC06 sample showed a trend toward decreasing best-fit spin, in better agreement with the MID06 results. The result is that the community has been left to decide for themselves which selection criteria seems preferable, or whether either is robust. Nevertheless, some support for the higher spin estimate of MCC06 is provided by efforts to model the reflection component (Miller et al., 2013), which also favors nearly maximal spins.

It is notable, however, that both of these previous papers (MCC06, MID06) made assumptions about the distance and black hole mass that are not well-supported by more recent efforts to constrain these system parameters based on very long baseline interferometric (VLBI) parallax distance measurements (Reid et al., 2014, hereafter R14), which place GRS 1915+105 at a smaller distance from us than originally assumed, with a lower black hole mass and a slightly lower inclination. R14 report

a preliminary analysis of black hole spin $a_* = 0.98 \pm 0.01$ (statistical error only), which would be consistent with the previous estimate from MCC06. Unfortunately, the continuum-fitting analysis is not a primary focus of R14 so the results are not described in extensive detail. Nor does it cover the MID06 selected data. Therefore, the goal of this study is to reanalyze both datasets, using these updated system parameters and their associated uncertainties to explore the uncertainty on the best-fitting spin.

The plan of this work is as follows: In Section 2.3 we summarize our methods of data selection and data reduction. In Section 2.4, we describe our spectral fitting models and best-fit results. We discuss our results in Section 2.5 and summarize our conclusions in Section 2.6.

2.3 RXTE DATA REDUCTION

The spectral states of GRS 1915+105 are known to vary quite rapidly on timescales of seconds to days, making it difficult to obtain disk-dominated spectra for continuum-fitting analyses (Greiner et al., 1996). MID06 and MCC06 sifted through archival *RXTE* data and generated large libraries of spectra determined to be disk-dominated. We briefly summarize the key differences in the two data selections, but refer the reader to the respective papers for further details.

MID06 generated a large library of spectra in the 3-20 keV range. They then selected out intervals of 16 seconds in which the disk contribution was more than 85% of the total bolometric luminosity. The 16 second exposures were set by the shortest timing resolution of *RXTE* and were chosen to avoid the variability seen in the longer exposures. However, MID06 note that the variability of GRS 1915+105 can be seen on timescales shorter than 16 seconds and thus require that the 16 second intervals have a rms variability less than 5%. Ultimately, 34 disk-dominated spectra across 6 *RXTE* observations in 16 second intervals were identified. These observations are within the β and κ variability classes of Belloni et al. (2000), in which the transition between spectral states is slow. Three of the 34 spectra were chosen by MID06 for their continuum-fitting analysis, which we also adopt in this paper: *RXTE* observation IDs 20402-01-45-03, 10408-01-10-00, and 10408-01-38-00, hereafter referred to as MID06a, MID06b, and MID06c, respectively. The following start and stop times for each observation's 16 second interval used in our data reduction are: MID06a (116417059 - 116417075), MID06b (75756947 - 75756963), and MID06c (87295987

- 87296003) (see MID06 their Figure 2). These start and stop times are in *RXTE* mission elapsed time (seconds). In their analysis, MID06 did not include the conventional 1% systematic error that is often added while performing spectral fitting to account for residuals that can be as large as 1% in the power law fits to the Crab Nebula. Since these observations are very short, the systematics are not expected to dominate the observations. We found that our results did not change significantly when we included a 1% systematic error, but nevertheless decided to retain it in the rest of our analysis.

MCC06 generated a large library of observations determined to be in a disk-dominated state using the following criteria: the disk contribution was more than 75% of the bolometric luminosity (in the 2-20 keV range) with QPOs either absent or weak and only allowing a small rms variability (<7.5%; McClintock & Remillard, 2006). This resulted in 20 candidate disk-dominated observations. In contrast to the three 16 second interval spectra used in the MID06 analysis, the spectra used for the MCC06 analysis remained thousands of seconds long. MCC06 identified five of the 20 observations as “key low-luminosity” spectra critical to their continuum-fitting analysis, where $L/L_{\text{Edd}} < 0.3$, and L_{Edd} is the Eddington luminosity. We note that this inferred luminosity and mass depends on the distance to GRS 1915+105, which given the closer VLBI distance and smaller mass should push the inferred L/L_{Edd} up. This implies that the MCC06 Eddington ratio criterion was more strict than was required. Contrary to the MCC06 observations, MID06 did not impose an Eddington ratio cut-off for their spectra. From the five, key, low-luminosity MCC06 observations, we chose three for our re-analysis: *RXTE* ObsIDs 10408-01-20-00, 10408-01-20-01, and 30703-01-13-00, hereafter referred to as MCC06a, MCC06b, and MCC06c, respectively. Note that the spectral energy range of interest for these spectra is 3-25 keV, which is slightly larger than the 3-20 keV range used for the MID06 spectra since the MCC06 observations are much longer and afford more signal-to-noise in the highest energy bins.

We emphasize that neither of the above selection criteria rely on any assumptions about the GRS 1915+105 system parameters. Therefore, we do believe it is necessary to repeat the selection analysis in response to the new VLBI distance constraints. The only exception is that MCC06 chose to focus on a low-luminosity subset of their selected data (with $L/L_{\text{Edd}} < 0.3$) for their discussion and we retain that focus here.

We used data reduction software tools from HEASOFT version 6.26.1. Follow-

ing the same reduction steps in both MID06 and MCC06, Standard-2 PCA spectra were extracted using FTOOLS `saextract` and corrected for background using `runpcabackest`, where all individual xenon gas layers were included. PCU gain variations were not corrected for and xenon layer spectra were not expanded to 256 channels. All spectra were corrected for dead-time. A 1% systematic error was added to all spectra using `grppha`. During the data reduction, a Good Time Interval (GTI) is usually specified to screen out undesirable data from events such as earth occultations, passage through the South Atlantic Anomaly, the target being at the edge of the field of view, etc. For the MID06 observations, we did not use any GTI criteria as these were only 16 second exposures. For the MCC06 observations, the GTI criterion specified was only data intervals in which all five PCUs were active during the observation.

2.4 RESULTS

The primary focus of the continuum-fitting method is to apply relativistic accretion disk models such as `kerrbb` (Li et al., 2005) and `bhspec` (Davis et al., 2005; Davis & Hubeny, 2006) to disk-dominated X-ray spectra and fit for the spin of the black hole. These models can fit for all parameters but degeneracies in how the model parameters affect the spectrum mean that prior knowledge of the distance to the source, D , the mass of the black hole, M , and the inclination of the accretion disk, i , are required for robust constraints. The most recent constraints on these values for the GRS 1915+105 system come from R14: $D = 8.6$ kpc, $M = 12.4M_{\odot}$, and $i = 60^{\circ}$, hereafter referred to as the R14 preferred values. We utilize XSPEC (Arnaud, 1996) for all of our spectral fitting, and the models used in this paper are collected in Table 2.1 with their corresponding XSPEC notations.

2.4.1 Non-relativistic accretion disk model

We first confirmed that our data are consistent with those reported in MCC06 and MID06 by comparing our fits with the non-relativistic accretion disk model, `diskbb` (Mitsuda et al., 1984), with corresponding fits in the two papers. Following the same fit procedure as MID06, we fit MID06a, MID06b, and MID06c tied together with the model `diskbb+nthcomp` (see Table 2.1). This model includes the variable abundance photoelectric absorption model `varabs`, the `diskbb` model, the thermal Comptonization model `nthcomp` (Zdziarski et al., 1996; Życki et al., 1999),

Table 2.1: List of models used in this paper and their full corresponding XSPEC notation. The prefix “mid” refers to models used to fit the selected MID06 observations following the same parameter and abundance prescriptions in MID06, and the prefix “mcc” is similarly used for the selected MCC06 observations for their parameter and abundance prescriptions (see Section 2.4.1).

Model Name	XSPEC Notation
diskbb+nthcomp	varabs*smedge(diskbb+nthcomp+gaussian)
diskbb+powerlaw	phabs*smedge(diskbb+powerlaw+gaussian)
mid:bhspec+simpl	varabs*simpl(bhspec)
mid:bhspec+nthcomp	varabs*smedge(bhspec+nthcomp+gaussian)
mcc:bhspec+simpl	phabs*edge*smedge*simpl(bhspec+gaussian)
mcc:bhspec+comptt	phabs*edge*smedge(bhspec+comptt+gaussian)
mcc:kerrbb+simpl	phabs*edge*smedge*simpl(kerrbb+gaussian)

the smeared edge component `smedge` (Ebisawa, 1991), and Gaussian line component `gaussian`. MID06 used abundances from Anders & Ebihara (1982), fixing all column densities in `varabs` to $4.7 \times 10^{22} \text{ cm}^{-2}$, except for Si and Fe which were fixed to $16.4 \times 10^{22} \text{ cm}^{-2}$ and $10.9 \times 10^{22} \text{ cm}^{-2}$, respectively (Lee et al., 2002). The smeared edge energy was fixed to lie between 6.9 – 9.0 keV, following MCC06 (as MID06 did not specify any restriction for this parameter), and width fixed at 7.0 keV Shafee et al. (2006). The `gaussian` line energy was fixed to lie between 6 – 7 keV, and the width was fixed at 0.5 keV. We obtained a fit with χ^2 per degree of freedom = 109.66/113 for all three observations tied together, with `diskbb` seed photon temperatures of $1.38_{-0.06}^{+0.06}$ keV, $1.68_{-0.11}^{+0.06}$ keV, and $1.93_{-0.17}^{+0.14}$ keV for MID06a, MID06b, and MID06c, respectively, which are within 10% of the values reported for the same model fit in MID06.

We also fit the MID06 observations with `simpl` (Steiner et al., 2009) in place of `nthcomp`. The `simpl` model relies on an approximate treatment of inverse Compton scattering, but it assumes the observed soft model component (in this case `diskbb`) provides the seed photon distribution that is Comptonized to give the hard X-ray emission. With `simpl`, the additional `smedge` and `gaussian` components that are necessary for fitting with `nthcomp` no longer significantly improve the `diskbb+simpl` fits. Our best fit $\chi^2_{\nu} = 173.86/113$, is notably worse with `simpl` than `nthcomp`.

Following the fit procedure in MCC06, the observations MCC06a, MCC06b, and MCC06c were all fit separately using the model `diskbb+powerlaw`. While MID06 used `varabs` for the absorption component, MCC06 used the photoelectric absorption model `phabs` with relative abundances from Anders & Grevesse (1989) and a lower fixed column density of $4.0 \times 10^{22} \text{ cm}^{-2}$ (see Section 2.5.2 for a discussion on the impact that chosen absorption models and column densities have on our analysis). The smeared edge energy was again fixed to lie between 6.9 – 9.0 keV and the width fixed at 7.0 keV. The Gaussian energy was fixed to lie between 6.3 – 7.5 keV and the normalization was allowed to go to negative values to allow for absorption, following MCC06. The `diskbb` temperatures and best-fit χ^2 per degree of freedom we obtained for each observation are $2.05_{-0.03}^{+0.03}$ keV with $\chi^2_{\nu} = 48.84/44$ for MCC06a, $2.06_{-0.03}^{+0.03}$ keV with $\chi^2_{\nu} = 45.76/44$ for MCC06b, and $2.11_{-0.02}^{+0.02}$ keV with $\chi^2_{\nu} = 58.08/44$ for MCC06c. Note that the 44 degrees of freedom reflect that each observation was fit independently of the others, in order to compare with the results from MCC06. We find that the results are consistent with the fits reported in MCC06.

2.4.2 Relativistic accretion disk models

We fit both the MCC06 and MID06 datasets using the relativistic accretion disk model `bhspec` (Davis & Hubeny, 2006) for $\alpha = 0.01$. In their analysis, MID06 used the model `mid:bhspec+nthcomp` with the same `varabs` prescription and parameter ranges discussed previously in Section 2.4.1. Instead, we chose to fit MID06a, MID06b, and MID06c simultaneously with the model `mid:bhspec+simpl`. This model differs from the model MID06 used in that we chose to use `simpl` (Steiner et al., 2009) to fit the hard X-ray emission rather than `nthcomp`. We performed fits with `bhspec` and `nthcomp`, but only report the `simpl` results here since we believe that tying the seed photon distribution to the soft model component is more self-consistent with a physically motivated accretion disk model. Furthermore, when used with `bhspec`, we do not consistently find best-fit results that are disk-dominated since `nthcomp` can account for a fraction of the softer emission when the temperature of the Comptonizing gas T_e is only slightly larger than seed photon temperature. `simpl` only has two free parameters (photon power-law index and photon scattered fraction) and we also drop the additional `smedge` and `gaussian` components, which do not significantly improve the fit when `simpl` is used. We constrained the `simpl` photon index to $\Gamma > 2$. If $\Gamma < 2$ is allowed, fits with high `simpl` scattering fractions are favored for observation MID06c, resulting in almost all of the Comptonized emission being present outside the limit of our data $E > 20$ keV. Hence, instead of fitting a power law in the hardest observed X-ray channels, `simpl` simply depresses the `bhspec` model flux fitting the softer photons, effectively renormalizing it. Fixing the values for the mass, distance, and inclination in `bhspec` to the values MID06 assumed ($M = 14 M_\odot$, $D = 12.5$ kpc, and $i = 66^\circ$), we find that MID06c became pegged at the luminosity limit of `bhspec` ($L/L_{\text{Edd}} = 1.77$), causing the spin to unrealistically drop to 0. Fixing the mass, distance, and inclination to the new R14 preferred values, we obtained a reasonable fit with $\chi^2_\nu = 136.3/125$ and a moderately high spin of $a_* = 0.863^{+0.014}_{-0.015}$. The best-fit parameters are reported in Table 2.2, and the three spectra fit simultaneously with the model `mid:bhspec+simpl` assuming the R14 preferred values are shown in Figure 2.1.

When we fit only observation MID06a, the closest in luminosity to the MCC06 observations, the best-fit preferred a high `simpl` scattering fraction and a significantly lower spin. We found that when separately fitting MID06a, MID06b, and MID06c




Figure 2.1: The top panel shows the three MID06 *RXTE* observations, MID06a (blue circles), MID06b (red diamonds), and MID06c (black squares), tied together and fit with the model `mid:bhspec+simpl`. The mass, distance, and inclination in `bhspec` were fixed to the R14 preferred values for the GRS 1915+105 system ($M = 12.4 M_{\odot}$, $D = 8.6$ kpc, and $i = 60^{\circ}$). The solid lines show the total model fit for each observation, and the dashed lines show the disk contribution for each spectral fit by setting the `simpl` scattering fraction to zero. The bottom panel shows the respective fit residuals.

Table 2.2: Best-fit values for the three MID06 observations, MID06a, MID06b, and MID06c tied together and fit with the model `mid:bhspec+simpl` (see Table 2.1). From top to bottom, the parameter values are the `simpl` model photon index, fraction of scattered photons, the `bhspec` luminosity, dimensionless spin parameter, and the reported χ^2 per degree of freedom for the entire fit. The mass, distance, and inclination were fixed at the R14 preferred values ($12.4M_{\odot}$, 8.6 kpc, and 60°).

Model Component	Parameter	MID06a	MID06b	MID06c
<code>simpl</code>	Γ	$3.52^{+0.34}_{-0.33}$	$4.15^{+0.47}_{-0.48}$	5.00^a
	f_{sc}	$0.22^{+0.05}_{-0.07}$	$0.33^{+0.12}_{-0.10}$	$0.24^{+0.06}_{-0.01}$
<code>bhspec</code>	L/L_{Edd}	$0.30^{+0.01}_{-0.01}$	$0.47^{+0.01}_{-0.01}$	$0.87^{+0.01}_{-0.01}$
	a_*	$0.863^{+0.014}_{-0.015}$	(tied)	(tied)
	χ^2_{ν}	136.3/125		

^a Parameter was completely unconstrained.

assuming the R14 preferred values and the model `mid:bhspec+simpl`, the two lower luminosity datasets (MID06a and MID06b) preferred higher scattering fractions with lower spins, whereas the higher luminosity dataset (MID06c) preferred a lower scattering fraction with a slightly higher spin compared to the best-fit values when all three datasets were tied together. This preference for high scattering fraction in the lower luminosity observations is partially attributable to `simpl` depressing the flux of the `bhspec` model at soft X-ray energies, an effect that is absent when additive models like `nthcomp` or `comptt` are fit for the harder X-ray component.

In contrast to `bhspec`, the relativistic accretion disk model `kerrbb` allows the color correction factor f_{col} to vary as a free parameter. However, fitting for a_* while allowing f_{col} to be free did not give reliable spin estimates, as the two parameters share a strong degeneracy (Salvesen & Miller, 2020). We discuss fitting `kerrbb` to the selected MCC06 observations fixing f_{col} at different values in the next section.

2.4.3 Color correction factor

We fit MCC06a, MCC06b, and MCC06c tied together and fit with the model `mcc:kerrbb+simpl`. The same parameter restrictions for the `edge`, `smedge`, and `gaussian` components discussed in Section 2.4.2 were again used in this model. For the `kerrbb` parameters, we assumed zero torque at the inner boundary, limb-darkening, and self-irradiation. The mass, distance, and inclination were fixed at the R14 preferred values. We then fixed f_{col} at different values ranging between 1.4 – 3.1 and show each resulting best-fit a_* and χ^2_ν (126 degrees of freedom) plotted as black dots in Figure 2.2. Over-plotted in the figure are two estimates of f_{col} obtained by fitting `kerrbb` to `bhspec`: one for fixed $L/L_{\text{Edd}} = 0.1$ (shown as blue diamonds), and one for fixed $L/L_{\text{Edd}} = 1$ (shown as red X's). We obtain the f_{col} estimates by running the XSPEC `fakeit` command to generate artificial datasets with a `phabs*bhspec` model, the response from the MCC06b observation, and assuming an exposure of 10^4 seconds. We then fit these datasets with `phabs*kerrbb`, fixing a_* , M , D , i , n_{H} and the accretion rate to match the values assumed in the faked spectrum, but allowing f_{col} to be a free parameter.

A reasonable fit can be obtained for color correction factors $f_{\text{col}} \gtrsim 1.7$ if $a_* = 0.999$. Lower spins only provide acceptable fits with higher values of f_{col} , but sufficiently high values of f_{col} only occur for L/L_{Edd} greater than inferred for the MCC06 data. A representative best-fit to the three spectra by arbitrarily fixing $f_{\text{col}} = 2.0$ is shown in

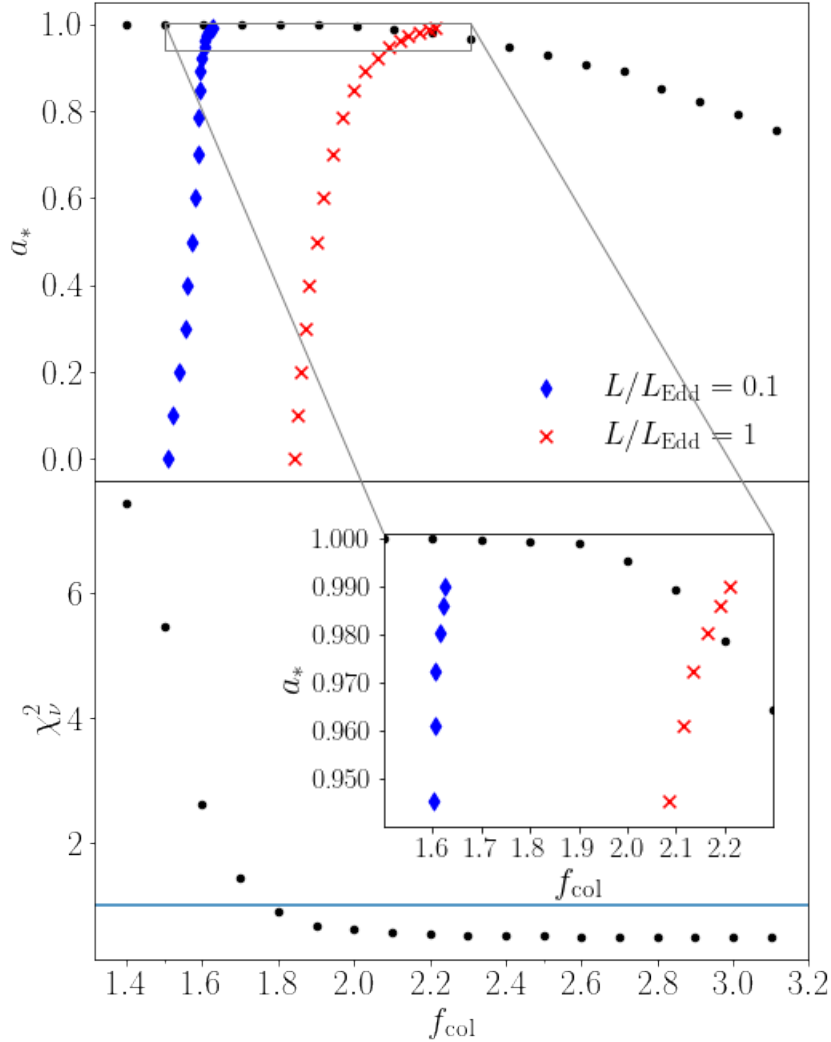


Figure 2.2: Plot showing the range of color correction factor values f_{col} for the three chosen MCC06 observations MCC06a, MCC06b, MCC06c tied together and fit with the `mcc:kerrbb+simpl` model, shown as black dots. The corresponding best-fit dimensionless spin parameter a_* is shown in the top panel, and χ^2_ν (126 degrees of freedom) is shown in the bottom panel. These fits were calculated assuming a mass, distance, and inclination fixed at the R14 preferred values ($M = 12.4 M_\odot$, $D = 8.6$ kpc, and $i = 60^\circ$). Estimates for f_{col} , found by fitting `kerrbb` to `bhspec` are overplotted for a fixed $L/L_{\text{Edd}} = 0.1$ (blue diamonds) and a fixed $L/L_{\text{Edd}} = 1$ (red X's). The horizontal line in the bottom panel shows $\chi^2_\nu = 1$ for reference.




Figure 2.3: The top panel shows the three MCC06 *RXTE* observations MCC06a, MCC06b, MCC06c tied together and fit with the model `mcc:kerrbb+simpl`. The mass, distance, and inclination were fixed at the R14 preferred values for GRS 1915+105 ($M = 12.4 M_{\odot}$, $D = 8.6$ kpc, and $i = 60^{\circ}$). The `kerrbb` color correction factor, or spectral hardening factor, was arbitrarily fixed at $f_{\text{col}} = 2.0$. The dashed lines show the contribution of the disk emission for each spectral fit by setting the `simpl` scattering fraction to zero. The bottom panel shows the fit residuals for each spectrum.

Figure 2.3 where the best-fit spin is $a_* = 0.995_{-0.003}^{+0.002}$ and $\chi_\nu^2 = 77.6/126$.

2.4.4 Exploring System Uncertainties

The uncertainty on the best-fit spin depends directly on the uncertainties in the distance, mass, and inclination. We explored this uncertainty in parameter space by fixing the mass, distance, and inclination at a range of different values above and below the R14 preferred values. A distance was randomly sampled from a Gaussian distribution centered on 8.6 kpc, with a 2.0 kpc width chosen to approximately match their uncertainty. From R14, the dependence of the inclination on a given distance is constrained from VLBI proper motion constraints, assuming ballistic trajectories for the plasma emitting in the jet. This gives

$$\tan i = \left(\frac{2D}{c} \right) \left(\frac{\mu_a \mu_r}{\mu_a - \mu_r} \right), \quad (2.1)$$

where i is the inclination of the accretion disk with respect to the line of sight ($i = 0$ is a face-on disk), D is the distance to the black hole, μ_a is the apparent speed of the approaching radio jet, and μ_r is the apparent speed of the receding radio jet. The values for μ_a and μ_r were also sampled from Gaussian distributions centered on their reported values of 23.6 ± 0.5 milliarcseconds/yr and 10.0 ± 0.5 milliarcseconds/yr, respectively (R14). The mass of the black hole is then determined by using the inclination from equation (2.1) in the following expression:

$$M = \frac{\mathcal{M}}{\sin^3 i}, \quad (2.2)$$

where M is the black hole mass, and \mathcal{M} is a constant adopted from the values for the mass function and binary mass ratio in Steeghs et al. (2013). For each fit, the randomly selected distance and subsequent inclination and mass were held fixed while the three MID06 observations were simultaneously fit with the model `mid:bhspec+simp1`. The results from each fit are plotted in Figure 2.4 which shows the spread in parameter space for mass, distance, and inclination, as well as a histogram of all best-fit a_* obtained. The blue dots are fits which are within 99% confidence, $\chi_\nu^2 \leq 164.7/125$, and the red dots are fits with $\chi_\nu^2 > 164.7/125$ which highlight regions where fits either became pegged at the maximum spin or the maximum luminosity limit of `bhspec`. The pile-up of fits at high spin have pegged at

the maximum spin limit of `bhspec` ($a_* = 0.99$), and fits below $a_* \sim 0.5$ signify observation MID06c has pegged at the luminosity limit of `bhspec` ($L/L_{\text{Edd}} = 1.77$). The R14 preferred values are marked as black X's at our best-fit spin, $a_* = 0.86$, for the MID06 observations. The previously assumed values for the mass, distance, and inclination from MID06 are marked as green X's at the best-fit spin reported in MID06, $a_* \sim 0.72$. The bottom x-axis is the logarithm of $(1 - a_*)$ to better show the portion of moderate to maximal spins, while the top x-axis is just a_* . Calculating a 1-sigma confidence interval on either side of our best-fit spin for only the acceptable fits (in blue) gives a spread of $a_* \sim 0.60 - 0.97$.

The same random sampling continuum-fitting analysis was done for the three MCC06 observations, the results of which are shown in Figure 2.5. This plot shows the range of distances, masses, and inclinations which produce fits to this dataset within 99% confidence ($\chi^2_\nu < 172.66/132$) using the model `mcc:bhspec+simpl` and the same parameter prescription in Section 2.4.2. Each dot is one realization for a fixed mass, distance, and inclination. When compared to the MID06 results from Figure 2.4, the MCC06 results show poorer fits for distances below ~ 10 kpc, with most of the best-fit spins pegging at the maximum spin limit of `bhspec` ($a_* \sim 0.99$). There are also a number of poor fits that prefer more moderate spins ($a_* \gtrsim 0.90$), but the `simpl` scattering fractions become high in these cases and are no longer consistent with disk-dominated results. Note that the range of the x-axes differ from Figure 2.4 in that the MCC06 fits do not reach spins below $a_* \sim 0.9$.

2.5 DISCUSSION

2.5.1 Implications for the GRS 1915+105 system

The relative merits of the MCC06 and MID06 data selection were debated in those papers and are summarized in Section 2.2. We will not discuss this at further length here but we note that one of the objections to the MID06 datasets is that their larger Eddington ratios imply thicker accretion disks, potentially invalidating the assumptions of the thin disk model underlying the `bhspec` model. With the revised system parameters, the implied Eddington ratios are now slightly lower, with the lowest Eddington observation being in a range where the disk model remains relatively thin. More generally, the relatively high Eddington ratio of GRS 1915+105 is sometimes hypothesized to account for its relatively unique variability, but our




Figure 2.4: Range of parameter space for the mass, distance, inclination, and resulting best-fit spin, a_* , for fits to the three MID06 observations, MID06a, MID06b, MID06c tied together and fit with the model `mid:bhspec+simp1` (see Section 2.4.2). These fits account for the parameter dependence of inclination and mass on the distance to GRS 1915+105 via equations (2.1) and (2.2) (R14). Each distance was randomly sampled from a Gaussian distribution centered on the R14 preferred value $D = 8.6$ kpc. The blue dots indicate fits with $\chi_\nu^2 \leq 164.7/125$ which are within 99% confidence. The red dots in each panel indicate fits with $\chi_\nu^2 > 164.7/125$ which are outside the 99% confidence. The pile-up of red dots at high spins is due to fits which have pegged at the maximum spin limit of `bhspec` ($a_* = 0.99$). The red fits at lower spins $a_* \lesssim 0.5$ indicate fits in which observation MID06c has pegged at the luminosity limit of `bhspec` ($L/L_{\text{Edd}} = 1.77$, where L_{Edd} is the Eddington luminosity). The R14 preferred values ($D = 8.6$ kpc, $M = 12.4 M_\odot$, $i = 60^\circ$) are marked with a black “X” at the best-fit spin for these assumed values, $a_* = 0.863^{+0.014}_{-0.015}$. The best-fit parameter values are listed in Table 2.2 for the R14 preferred values). The distance, mass, and inclination previously assumed by MID06 (12.5 kpc, $14 M_\odot$, 66° , respectively) are marked with a green “X” for comparison. The bottom right panel shows a histogram of all the resulting best-fit spins obtained for the MID06 observations. Note that the histogram is stacked rather than superimposed, and the area under the histogram integrates to 1.




Figure 2.5: The same analysis in Figure 2.4 is performed for the three MCC06 observations MCC06a, MCC06b, MCC06c. Note that when comparing this figure to Figure 2.4, the spin axis here is truncated at $a_* = 0.900$ as there were no best-fit, moderate spin values below this for the MCC06 fits. These observations were similarly fit together with spins tied in the model `mcc:bhspec+simpl`. Each dot represents one realization for a fixed mass, distance, and inclination. All blue fits shown here are within 99% confidence ($\chi^2_\nu \leq 172.7/132$). All red fits have a $\chi^2_\nu > 172.7/132$ where the `simpl` scattering fractions have become high and are no longer consistent with disk-dominated results. A pile-up of spins is shown at $a_* = 0.99$ where fits have pegged at the maximum spin limit of `bhspec`, similar to Figure 2.4. The lower right panel shows a stacked histogram of all best-fit spins where the sum over all bins equates to 1.

best-fit constraints imply the source is generally sub-Eddington or at most slightly super-Eddington.

Although more moderate spins are allowed by the MID06 datasets, relatively high spin is implied for the R14 best-fit system parameters. Black holes in X-ray binaries are expected to be born with low to moderate spins ($a_* \lesssim 0.7$), although this is subject to uncertainties in the core-collapse process (Gammie et al., 2004). It is also not clear that they can be significantly spun up by accretion under standard assumptions about mass transfer (King & Kolb, 1999). The high spins inferred here would then imply that either black holes are born with higher natal spin than expected or experience phases of high mass transfer to spin them up.

It is perhaps notable that the best-fit values from MID06 are in the ballpark where general relativistic magnetohydrodynamic simulations suggest magnetic torques would balance the spin-up due to accretion (Gammie et al., 2004). This limit ($a_* \lesssim 0.94$) is more stringent than the commonly cited limit of $a_* = 0.998$ from Thorne (1974), which only accounts for the angular momentum carried by the radiated photons. The MID06 results are thus consistent with GRS 1915+105 being spun up by accretion and reaching an equilibrium with magnetohydrodynamic torques provided by field lines connected to the black hole, while MCC06 results exceed this nominal limit. Note, however, that the presence of such torques may have an effect on the accretion disk emission (Gammie, 1999; Agol & Krolik, 2000; Kulkarni et al., 2011; Schnittman et al., 2016) and are not accounted for in the present analysis.

Our results for the MID06 data are inconsistent with those of R14, who report best-fit $a_* \simeq 0.98$. Accounting for systematic uncertainties, they report $a_* > 0.92$, which is consistent with spin estimates in Figure 2.5. Our results for the MCC06 data are in better agreement in that both analyses favor near maximal spin but differ in that R14 managed to find suitable fits for the revised distance of 8.6 kpc. This may owe in part to R14 reanalyzing a large sample of *RXTE* observations, selecting all observations that obey a criterion $L/L_{\text{Edd}} \leq 0.3$, $\chi^2_\nu \leq 2$, and $f_{\text{sc}} < 0.25$. It is possible that the MCC06 selected datasets may have been selected out in the process using the revised system parameters. Sreehari et al. (2020) also find a near maximal best-fit spin for *kerrbb* fits to *AstroSat* observations of GRS 1915+105. These results are notable in that, like MID06, they are for observations that would nominally place the emission above the Eddington limit. Their analysis differs in allowing the mass to be a free parameter, although their best-fit mass is consistent with the R14 constraints.

In addition to the continuum-fitting method, the spin of GRS 1915+105 has also been estimated by fitting the relativistically broadened reflection spectrum due to irradiation of the accretion disk by a corona (Blum et al., 2009; Miller et al., 2013) or via modeling of quasiperiodic oscillations (QPOs; Török et al., 2011; Šrámková et al., 2015). Although a range of results have been reported with both methods, the reflection fitting efforts are both consistent with high spins $a_* \simeq 0.98$ while the QPO model favors somewhat lower spins $a_* \sim 0.7-0.9$. Based on previous results, this puts the reflection fitting results in good agreement with MCC06 and the QPO estimates in better agreement with MID06. For our results, the MCC06 fits are still broadly in agreement with near maximal spin as long as GRS 1915+105 lies at the far end of the distance distribution allowed by VLBI parallax. The MID06 data now provide a larger best-fit spin in better agreement with reflection fitting. More moderate spins are still favored albeit with large uncertainty.

2.5.2 Impact and Uncertainties in Interstellar Absorption

Typically when fitting the continuum of a BHXRb, a model for the photoelectric absorption along the line of sight (e.g. `phabs`, `varabs`) is needed. Using the model `varabs` and abundances from Anders & Ebihara (1982), MID06 assumed a column value of $n_{\text{H}} = 4.7 \times 10^{22} \text{ cm}^{-2}$ for all elements except Si and Fe which were fixed at $16.4 \times 10^{22} \text{ cm}^{-2}$ and $10.9 \times 10^{22} \text{ cm}^{-2}$, respectively. These values were reported by Lee et al. (2002) for *Chandra* X-ray observations of GRS 1915+105 assuming ISM abundances. Relativistic disk reflection studies constraining the spin in GRS 1915+105 report best-fit values which also favor a high absorption column using the `phabs` model ($n_{\text{H}} = 4.15 - 5.64 \times 10^{22} \text{ cm}^{-2}$, Blum et al. 2009; $n_{\text{H}} = 6.1 \times 10^{22} \text{ cm}^{-2}$, Miller et al. 2013). Lee et al. (2002) also reported a S- and Mg-derived hydrogen column assuming solar abundances, giving a more moderate column value of $n_{\text{H}} \sim 3.1 \times 10^{22} \text{ cm}^{-2}$. This is in better agreement with other modest column estimates from ASCA X-ray observations ($n_{\text{H}} = 3.8 \times 10^{22} \text{ cm}^{-2}$, Ebisawa et al., 1994) along with millimeter and radio observations ($n_{\text{H}} = 3.6 \times 10^{22} \text{ cm}^{-2}$, Chapuis & Corbel, 2004). Following these modest estimates, MCC06 adopted a value of $n_{\text{H}} = 4.0 \times 10^{22} \text{ cm}^{-2}$ (assuming abundances from Anders & Grevesse 1989).

Not only did MCC06 and MID06 assume different values for the column, they also chose different XSPEC absorption models. MCC06 selected `phabs` for their analysis, while MID06 chose `varabs`. We found that when using the same n_{H} value and all

other variables kept the same, **varabs** tended toward a lower a_* than **phabs** did. Fits to the MCC06 data with either **phabs** or **varabs** tended to fit better with low n_H values, while fits to the MID06 data tended to fit better with higher n_H values. An overall trend between both **varabs** and **phabs** is that the value for a_* decreased as the n_H column increased for both datasets. To maintain consistency in our re-analysis we kept the same XSPEC absorption models and values for n_H chosen by each group, but note the sensitivity and dependence of the spin on the assumed n_H column value as a source of uncertainty in the a_* estimate in this work.

Aside from the line-of-sight hydrogen column estimates, kinematic studies near GRS 1915+105 have also located a molecular cloud at a distance of 9.4 ± 0.2 kpc (Chaty et al., 1996; Chapuis & Corbel, 2004). The new VLBI constraints may have implications for the history and conception of the GRS 1915+105 system given the distance of 8.6 kpc, which could place it with the observed interstellar structure (R14).

2.5.3 Uncertainties in Models and System Parameters

The R14 VLBI parallax measurements provide much stronger constraints on the system parameters than were previously available, but our analysis shows that remaining uncertainties still allow for a rather large range of spins. If we treat the models as robust to systematic uncertainties, then our fitting constraints nominally imply strict limits on the system parameters. Figure 2.4 implies that for relatively low distances and inclinations, the implied spin from the MID06 data would be higher than $a_* = 0.99$, challenging the theoretical understanding on black hole spin limits. The constraints are even stronger for the MCC06 data, which would limit the distances to $D \gtrsim 10$ kpc, near the outer limits of what is allowed by the R14 constraints. In fact, this result is consistent with predictions made in Figure 18 of MCC06, which predicted GRS 1915+105 to lie within an error triangle whose minimum distance was just under ~ 10 kpc.

The model for system parameters implied by equations (2.1) and (2.2) is also subject to systematic uncertainties. First, equation (2.1) assumes the observed superluminal motion can be interpreted as emission from plasma following ballistic trajectories and that these trajectories lie along the spin axis of the black hole. Furthermore, it is conceivable that the plane of the binary is not perpendicular to the black hole spin axis (Fragile et al., 2001; Maccarone, 2002), although there are theoretical arguments that such misalignments should typically be modest (Fragos &

McClintock, 2015). If such misalignment is present, it seems likely that the inner accretion disk would align with the spin axis due to the action of Lense-Thirring precession (Bardeen & Petterson, 1975). In this case, our use of the jet to fix the disk inclination would still be reasonable as long as the jet is aligned with the spin axis, although GRMHD simulations of misaligned disks indicate this may not be guaranteed (Liska et al., 2018) and observations of V404 Cygni show the jet angle to precess (Miller-Jones et al., 2019) perhaps due to Lense-Thirring precession associated with the high mass accretion rates inflating the disk (Middleton et al., 2018, 2019). The inclination implied by observations of the jet would then not correspond to the binary inclination in equation (2.2), and the inferred black hole mass would be incorrect.

An independent constraint on the inclination comes from the reflection spectral modeling, where the relativistic line profiles are sensitive to the viewing inclination. The best-fit inclinations from Blum et al. (2009) are $i = 55^\circ$ or $i = 69^\circ$ depending on the reflection model used. Miller et al. (2013) found inclinations ranging from 65° to 74° depending on the model. Blum et al. (2009) constrained the inclination to lie between 55° and 75° , while Miller et al. (2013) constrained it to be between 65° and 80° , both based on interpretations of constraints from the superluminal jet model (Fender et al., 1999). The allowed inclination ranges are generally higher than those used in our analysis because these papers predate the R14 measurements. The higher inclinations ($i \gtrsim 70^\circ$) would pose a challenge to the super-luminal motion interpretation for the new parallax distance, but if we assume they imply that the inclination should be towards the high end of the allowed range, they would push the MID06 results to low spins that would be at odds with the best-fit spins from these reflection models. The MCC06 results could remain consistent with near maximal spin, but again this requires GRS 1915+105 to be located at the more distant end of the range allowed by VLBI parallax measurements.

We emphasize that these constraints are subject to unquantified systematic uncertainties in the underlying accretion disk model. This could arise from inaccuracies in the underlying thin disk model (Shakura & Sunyaev, 1973; Novikov & Thorne, 1973) or because of errors in the TLUSTY atmosphere models (Davis & Hubeny, 2006). The former is perhaps most worrisome for the highest Eddington ratio observations from MID06, where the thin disk assumption would begin to break down. It is also possible that some aspect of the model is inaccurate, such as the stress prescription (Done & Davis, 2008), assumption of a torque-free inner boundary (Gammie, 1999;

Agol & Krolik, 2000) and truncation of emission at the ISCO (Krolik & Hawley, 2002; Abramowicz et al., 2010). The relativistic thin disk models employed here are broadly consistent with GRMHD models, but emission interior to the ISCO and magnetic torques would provide a modest bias towards higher spins when using standard models (Kulkarni et al., 2011; Schnittman et al., 2016). In that case, the near maximal best-fit spins might still be indicative of high spin, but not necessarily maximal spin. These models also assume a razor thin disk even though they can be at accretion rates where the thin disk assumption breaks down. Zhou et al. (2020) found that considering a model with a finite disk thickness led to a modestly higher spin for fits to *RXTE* observations of GRS 1915+105, but their thin disk fits were already near maximal.

The spectra derived from atmosphere modelling with TLUSTY are another potential source of error. Errors in the atmosphere models and spectra could arise from inaccurate assumptions about the vertical distribution of dissipation, contributions from magnetic pressure support, inhomogeneities in the turbulent disk, and lack of irradiation of the surface (Davis et al., 2005; Blaes et al., 2006; Davis et al., 2009; Tao & Blaes, 2013; Narayan et al., 2016). Figure 2.2 provides a sense of the degree to which spectral hardening errors would impact our spin results. Further discussion of the spectral hardening implied by TLUSTY calculations can be found in Davis & El-Abd (2019) while Salvesen & Miller (2021) provide a thorough review of uncertainties and quantitative estimates of their impact on spin measurements.

We note that the selection criteria is also a clear source of uncertainty, since two different methods provide nominally disk-dominated spectra that yield different results. This concern for GRS 1915+105 contrasts with other sources (mainly soft X-ray transient low mass X-ray binaries) that tend to approximately follow a luminosity proportional to temperature to the fourth power relation (Gierliński & Done, 2004; Dunn et al., 2011). Since color corrections tend to vary relatively weakly (Davis & El-Abd, 2019), this means different observations of the same source likely yield consistent inner disk radii and consistent spins. GRS 1915+105 tends to be highly variable and less consistent, which is partly why Dunn et al. (2011) exclude it from their analysis. Nor is it clear that its variability properties are consistent with other sources in its nominally disk-dominated states, possibly indicating contamination from a warm Comptonizing component (Zdziarski et al., 2005; Ueda et al., 2010). The quality of fit for our spectral modelling is sensitive to the chosen hard X-ray model. Fits

with the `nthcomp` model can provide a lower χ^2_ν than with the `simpl` model, and did not provide consistently disk-dominated fits when paired with `bhspec`. Although the `simpl` model provides a poorer fit, it provides relatively disk-dominated results, consistent with Comptonization models where the seed photons are provided by the accretion disk and scattered in a hot corona. Further work is needed to robustly characterize the disk-dominated states of GRS 1915+105 and ascertain how this may affect the best-fit spin constraints.

2.6 SUMMARY

We re-examine the continuum-fitting based spin estimates for GRS 1915+105 in light of new constraints on the mass, distance, and inclination from VLBI parallax (R14). We find that the discrepancies between data selected by MID06 and MCC06 persist, implying that the selection criteria of one (or both) is inconsistent with the assumptions of the thin disk model. MCC06 showed a trend towards lower spin as the luminosity of the observations increased, indicating the discrepancy may be driven primarily by different Eddington ratio ranges of the two datasets. The revised system parameters lower the mass, but lead to relatively smaller implied luminosities, leading to lower overall Eddington ratios for fits to both datasets. This somewhat mitigates concerns that the Eddington ratios in the MID06 models were too high, but the highest Eddington ratio is still close to unity ($L/L_{\text{Edd}} = 0.87$), where the scale height of the disk is unlikely to remain small compared to the radius, as assumed in the model.

The new system parameters drive both datasets to higher spins. Since MCC06 were already fitting for near maximal spin, this presents a challenge. For the `bhspec` model (or `kerrbb` model with color corrections set to match `bhspec`), we cannot obtain a good fit for the preferred (R14) system parameters. Good fits to these data can only be found if the color correction is allowed to vary to values significantly higher ($f \gtrsim 2.2$) than implied by `bhspec` or the distance to GRS 1915+105 is near or greater than 10 kpc, consistent with the prediction of MCC06 (their Figure 18). The spin would remain near maximal ($a_* \simeq 0.99$) for a distance of 10 kpc, consistent with constraints from modeling of the reflection spectrum.

For the MID06 data, the best-fit spin is moderately high ($a_* \simeq 0.86$) for the best-fit R14 system parameters. We find, however, that a fairly broad range of spins are allowed when the uncertainty in the parallax distance and jet model inclination are

accounted for, as indicated by Figure 2.4. In principle, this allows the spin to match constraints from either the near maximal spins from reflection modeling or the more moderate spins from QPO models. Near maximal spin, however, would result for fairly low inclinations in our model, which would be inconsistent with the best-fit inclinations from reflection modelling. The low end of the allowed spin distribution is sensitive to the maximum Eddington ratio permitted by `bhspec` ($L/L_{\text{Edd}} = 1.77$). Therefore, the lower limit on a_* is tied to the Eddington ratio beyond which one says the thin disk model is no longer valid.

Although the VLBI parallax measurements are an impressive achievement, our results indicate that even stronger constraints are necessary to provide a tight constraint on the spin with the continuum-fitting method, and to help resolve the discrepancies driven by data selection. We note that such constraints also have implications for the reflection spectrum modeling through the dependence of relativistic Doppler shift and beaming on the observer inclination. This analysis would also benefit from better constraints on the interstellar absorption toward GRS 1915+105. The datasets modeled here prefer different models for the absorption and are particularly sensitive to the hydrogen column assumed. The range of columns used in the literature vary by more than a factor of two, which is enough to modify the best-fit spin, with larger assumed columns generally providing lower spins.

CHAPTER 3

THE Athena++ MONTE CARLO RADIATION TRANSFER MODULE

In this Chapter we present the Monte Carlo radiation transfer module (Davis et al., 2005, Davis et al. in preparation) methods and tests, followed by its application for ULX spectra in Chapter 4. An overview of the MC module code structure and implementation into the Athena++ code framework is presented in Section 3.1. We discuss the MC photon samples in Section 3.2, their emission and weighting properties in Section 3.3, the movement of photon samples in Section 3.4, scattering and absorption processes in Section 3.5, and the MC outputs in 3.6. In Section 3.7, we present several code tests of the MC module and compare (when possible) to known analytic solutions. Finally in Section 3.8, we briefly describe how the MC module can be used.

3.1 ADOPTING THE Athena++ CODE FRAMEWORK

The Athena++ framework is described in detail in Stone et al. (2020) and we briefly review the most salient points relevant to MC calculations here. The primary components of the Athena++ code are a mesh and a task list. The mesh provides a structure on which physical variables (e.g. fluid properties, magnetic fields, etc.) are defined. The task list provides a general framework for time integration of the physical variables. A typical Athena++ calculation begins with creation of the mesh

This chapter is adapted from an article in preparation. Note that the complete author list and general order after the first few authors has not yet been finalized. The current working citation is as follows: S. W. Davis, B. S. Mills, J. Stone, B. C. McClellan, et al. in preparation.

and initialization of all physical variables. The code is parallelized with the message passing interface (MPI) framework so the meshblocks may be on separate processes corresponding to different CPUs or CPU cores. The variables defined on the mesh are then integrated with each time step following a dynamic task list implementation, which reduces latency due to communication between processes. The calculation is finished once it meets some termination condition such as a time or cell cycle limit.

The Athena++ mesh is structured as a collection of meshblocks, with each meshblock having an identical layout of $Nx1 \times Nx2 \times Nx3$ cells. The organization of the mesh into meshblocks facilitates parallelization by subdomain decomposition and mesh refinement. Subdomain decomposition allows one to spread the computational cost of integration across a number of different MPI processes, with each process integrating its own meshblocks. Information required for integration can then be passed to, or received from, neighboring meshblocks on other processes. Athena++ implements a block-based mesh refinement scheme. For each additional refinement level requested, the mesh is divided by two along each coordinate direction. For example, in a three-dimensional mesh if refinement by one level is requested, an $Nx1 \times Nx2 \times Nx3$ meshblock at the coarse level will be replaced by eight $Nx1 \times Nx2 \times Nx3$ meshblocks at the finer level.

Each cell on the meshblock contains a number of physical variables to be integrated. The variables of primary interest to the MC calculations described in Chapter 4 are the fluid variables contained in the Hydro class: the fluid density ρ , the fluid pressure P , and the fluid velocity \mathbf{v} . Athena++ also supports the advection of passive scalars, which can be used to follow radiation transfer properties, such as ionization fractions or level populations.

Athena++ can be used to run specific problems by setting configuration flags and creating two user-defined files. The first is an input file which defines the mesh properties, specifies desired outputs, and sets physics-specific or problem-specific parameters. The second is a problem generator file. This file specifies a user-defined C++ function to initialize physical variables on the simulation mesh. It also provides hooks into the code for a variety of user-defined functions, such as boundary conditions, explicit source terms, output variables, and a variety of other physics components.

3.2 PHOTON SAMPLES

The MC module is fundamentally an integration scheme for a vector of photon samples (often referred to as superphotons or photon packets). Each photon sample represents an ensemble of a large number of photons with common properties such as energies, positions, directions, and statistical weights. All Athena++ photons have the following attributes: a statistical weight w , a direction four vector \mathbf{k} , and a position four vector \mathbf{x} . These are the primary properties that are evolved as the MC photon proceeds. Additional physical variables are introduced as needed. Examples include Stokes vectors (q, u, v) and polarization tensors to model polarization properties, as well as acceleration vectors to handle movement by geodesic integration (Section 3.4). Additional problem-specific user variables can also be specified.

A number of auxiliary variables are also used for algorithmic convenience. Each photon retains a vector of three integers \mathbf{i}_x , which identifies the location of its current cell within mesh. It is also useful to retain the extinction coefficients due to absorption $\alpha_\nu = \kappa_\nu^{\text{abs}} \rho$ and scattering $\sigma_\nu = \kappa_\nu^{\text{scat}} \rho$, where ρ is the density, κ_ν^{abs} is the absorption opacity, and κ_ν^{scat} is the scattering opacity in the mesh cell currently containing the photon. A status flag is used to indicate whether the photon is evolving, escaped, absorbed, or in some other state.

Photons are represented by a C++ class, with each instance containing multiple photon samples. Each photon property listed above is implemented as a vector of integers, real numbers, or complex numbers corresponding to the number of samples within the instance. Each property of a given photon is labeled by the same integer index across these vectors. There is one photon instance per meshblock, which is used to integrate all the photons on that meshblock. This vector implementation has a number of advantages from a performance standpoint, particularly when photons need to be moved between meshblocks on different processes. Integrating a large number of photons in parallel allows us to collect the photons crossing meshblock boundaries, reducing the number of MPI communication calls.

3.3 EMISSION AND STATISTICAL WEIGHTS

The mechanism for emitting photons can vary greatly from one type of calculation to the next. For example, a user may want the emission to represent discrete point sources within the simulation domain, irradiation by an external source outside the

domain, or continuum processes that are continuous throughout the domain. Given the variety and varying degrees of complexity involved in different emission scenarios, we largely leave emission to a problem-specific user supplied function within the problem generator file (Davis et al. in preparation).

This user-defined function must set w , E , \mathbf{x} , \mathbf{k} and any other photon or auxiliary variables required by the configuration. The module provides functions for initializing the direction and photon energy for some common cases, such as isotropic emission, semi-isotropic emission from a boundary, emission energies drawn to correspond to the Planck distribution, and setting the weights for emission from point sources of prescribed luminosity. We now consider the case of continuum emission to provide a concrete example.

We consider any continuum emission process that can be written as a function of fluid variables specified on the mesh block and lab frame photon frequency ν . Usually this is the fluid density ρ and temperature T . Since the fluid generally provides pressure P , the module provides a default function for computing the fluid temperature,

$$T = \frac{\mu m_p P}{\rho k_B}, \quad (3.1)$$

where m_p is the proton mass, μ is the mean molecular weight ($\mu = 1$ by default), and k_B is Boltzmann's constant. This function can be replaced by a user-defined function, if desired. The comoving frame emissivity is then assumed to be $j_0(\nu_0, T, \rho)$ although more general functions of additional fluid variables are possible. The subscript "0" denotes that the variables are evaluated in the comoving frame.

In the following discussion, we take care to differentiate photon samples and the physical photons that they model. The initial weight w of the photons emitted in a cell is then determined by the emissivity, the cell volume, and the integration time (which is arbitrary for post-processing calculations). We label cells by index i and photons samples with index j . We assume a total number of photon samples N_s and N_{cell} cells in the mesh. The total number of physical photons emitted in cell i can be written as

$$N_i = \int \frac{j_i(\nu)}{h\nu} d\Omega d\nu \mathcal{V}_i \Delta t_{\text{int}}, \quad (3.2)$$

where all variables are evaluated in the lab frame. Here, Ω is the solid angle, \mathcal{V}_i is the volume of cell i , and Δt_{int} is the integration time interval. Note that N_i is a relativistic invariant. This can be seen by recognizing that the product of the volume and time interval is relativistically invariant, that $j_0(\nu_0) = (\nu_0/\nu)^2 j(\nu)$, and that $d\Omega_0 d\nu_0 = (\nu/\nu_0) d\Omega d\nu$.

The statistical weights are defined such that

$$\sum_{j=1}^{N_s} w_j = \sum_{i=1}^{N_{\text{cell}}} N_i = N_{\text{ph}}, \quad (3.3)$$

where N_{ph} is the total number of physical photons emitted within the entire mesh. We can define the probability P_i for a photon to be emitted in zone i as $P_i = N_i/N_{\text{ph}}$. Then, the average number of photon samples emitted in cell i is $P_i N_s$ and we have

$$w_i = \frac{N_i}{P_i N_s}, \quad (3.4)$$

where w_i is the default weight for samples emitted in cell i . In principle, we are free to set P_i as we like to determine which cells will get more or less samples, as long as we adjust w_i accordingly. Two obvious limits are $P_i = 1/N_{\text{cell}}$ and $P_i = N_i/N_{\text{ph}}$. We refer to the former as the “equal probability” case since all cells have equal probability of emitting a photon sample. In this case, the initial weights $w_i = N_i N_{\text{cell}}/N_s$ can span a large range, since N_i can vary greatly from cell to cell. We call the latter limit the “equal weight” case, as each photon is emitted with the same initial weight $w_i = N_{\text{ph}}/N_s$.

There are pros and cons to each limit. The equal probability is easier to implement when the mesh is broken up into multiple meshblocks, as each block simply receives an equal fraction of the total number of photon samples. It can also be advantageous when one wants to compute quantities like the radiation energy density, which are evaluated in each cell as it ensures that all cells are likely to receive at least some minimum photons samples. In problems with large scattering optical depths, it is also frequently the case that photons with large N_i are regions deep within the photosphere so the average photon sample experiences more overall scatterings, increasing the run time.

The equal weight case can be harder to implement on a distributed mesh. One has

to compute the emissivity of each meshblock and then do a global calculation over all meshblocks on all MPI processes to assess the relative number of photon samples that each meshblock receives. Then photon samples are distributed to each cell within the meshblock accordingly. Such a scheme is implemented in **Athena++**. The main advantage for starting all photons with equal weights is that this is generally variance reducing, which should allow one to obtain better statistical averages of output quantities of interest with fewer photon samples, nominally decreasing the run time. Tests show, however, that this advantage is often offset by the large number of scatterings when scattering optical depths are large. The optimal choice will therefore depend on the problem being solved.

Once the statistical weight has been assigned, one must set the initial direction and energy of the sample. These may be further supplemented by choices for the polarization (Stokes parameters) or any user-defined sample properties, if needed for the problem under consideration. Again, the code is designed to allow significant freedom in these choices. In the case of free-free emission, the starting photon direction is drawn as a unit vector. For example, in Cartesian coordinates:

$$\begin{aligned} k_x &= \sqrt{1 - k_z^2} \cos(2\pi\chi_2), \\ k_y &= \sqrt{1 - k_z^2} \sin(2\pi\chi_2), \\ k_z &= 2\chi_1 - 1, \end{aligned} \tag{3.5}$$

where χ_1 and χ_2 are pseudorandom numbers distributed uniformly on the interval $(0, 1)$. One could set the initial photon energy (frequency) by directly sampling the free-free emissivity function (e.g. using a rejection method), but this has the disadvantage of producing relatively few samples in the high energy tail. Since we frequently desire spectra with reasonable signal-to-noise at higher energies, we often find it better to sample uniformly in the logarithm of the photon energy over some specified energy range and then adjust the statistical weight accordingly. The energy limits can be set either by a global (E_{\min}, E_{\max}) or a local (x_{\min}, x_{\max}) , where $x = E/(k_B T_i)$ and T_i is the temperature of the zone in which the sample is initialized. If tracked, polarization is initialized as $q = 0$, $u = 0$, and $v = 0$, corresponding to unpolarized radiation.

Finally, if frame transformations are requested, these photon properties are Lorentz boosted from the comoving frame to the lab frame. In the case where general relativis-

tic affects are included, the photons are transformed to the coordinate frame. Photon samples are then transported (moved) in the lab or coordinate frames, respectively.

3.4 PHOTON MOVEMENT

The transport, or movement, of photon samples is another property of MC radiation transfer that can vary depending on the problem and method of solution. In many MC schemes, integration of photons generally proceeds in a series of steps where one draws the distance a photon travels between matter interaction events (scattering or absorption) as a rational number of photon mean-free-paths. This can also be thought of as the optical depth τ a photon travels before the next interaction. We implement this method in *Athena++*. In the simplest, unbiased distribution we draw

$$\tau = -\log(1 - \chi) = -\log(\chi), \quad (3.6)$$

where χ is a pseudorandom number distributed uniformly on the interval $(0, 1)$. This corresponds to the probability decreasing exponentially as the path length in units of the current mean-free-path increases.

In some cases, this path length will exceed the length within the current cell and the photon mean-free-path will change as the photon moves into a new cell with different opacities. Generally, the photon is moved in a series of steps (indexed by m) such that

$$\sum_{m=1}^{m_{\max}} \lambda_m [\alpha(\nu) + \sigma(\nu)] = \tau. \quad (3.7)$$

At the end of this sequence, the scattering and absorption processes are evaluated. If the photon continues evolving, a new τ value is drawn. In *Athena++* we implement two primary schemes for movement. In both methods, we treat cells as uniform in fluid properties with discontinuous jumps at cell faces. Therefore each cell has a uniform $\alpha_0(\nu_0)$ and $\sigma_0(\nu_0)$, which are then boosted to obtain $\alpha(\nu)$ and $\sigma(\nu)$ in the lab frame where the faces are defined.

Cell-by-cell Movement

The first method is a cell-by-cell approach that is only used for flat spacetimes. At the start of the move, one computes the distance $l_\tau = \tau/(\alpha(\nu) + \sigma(\nu))$ that the photon would need to travel in the current zone to move the drawn optical depth. This is compared to the distance l_f along \mathbf{k} to the nearest cell face. This method is simplest to implement in Cartesian coordinates, in which case

$$\begin{aligned} l_{f,x} &= (x_f - x)/k_x, \\ l_{f,y} &= (y_f - y)/k_y, \\ l_{f,z} &= (z_f - z)/k_z, \end{aligned} \tag{3.8}$$

and $l_f = \min(l_{f,x}, l_{f,y}, l_{f,z})$. Here x_f is the x location of the nearest face in the photon's direction, x is the photon's current value, and k_x is the current direction. Note that in the cell-by-cell approach, we normalize \mathbf{k} so that $|\mathbf{k}| = 1$.

We then choose the step size $\lambda_m = \min(l_\tau, l_f)$ and the new photon position is updated with

$$\begin{aligned} x_{m+1} &= x_m + k_x \lambda_m, \\ y_{m+1} &= y_m + k_y \lambda_m, \\ z_{m+1} &= z_m + k_z \lambda_m. \end{aligned} \tag{3.9}$$

If the photon moves to a new cell, we update the integers \mathbf{i}_x designating the photon's current cell and compute the values of $\alpha(\nu)$ and $\sigma(\nu)$ in the new cell.

In flat space time, this method can be generalized to curvilinear coordinates with slight modifications that introduce greater complexity and increased computational cost. Noting that k_x , k_y , and k_z are unchanged during the photon's displacement, we consider the spherical polar case. Displacement of the photon is similar to the above algorithm, except that one first transforms k_r , k_θ , and k_ϕ to k_x , k_y , and k_z , then uses equation (3.9) to update x , y , and z , and finally transforms both \mathbf{x} and \mathbf{k} back to spherical polar at the final position. This requires computation of the sine and cosine of the θ and ϕ coordinates before and after the displacement.

The primary difficulty is that equation (3.8) no longer applies. A similar expression can be written for motion in the ϕ direction, but a photon moving in a constant

direction can have turning points in both the r and θ coordinates. For example, a photon with $k_r < 1$ at radius r may intersect a cell face with either $r_f < r$ or $r_f > r$ depending on the cell geometry. Similarly, a photon with $k_\theta > 0$ at polar angle θ may intersect the face with $\theta_f < \theta$ and vice versa. One must check for turning points before evaluating distances to the nearest face. This added complexity makes moving photons through a spherical polar grid more expensive than Cartesian in the cell-by-cell approach.

Geodesic Integration

In curved spacetime, photons follow geodesics so we implement a general photon mover that integrates these geodesics. This method can also be adapted to solve curvilinear coordinates in flat spacetime. The definition and normalization of \mathbf{k} are slightly different. The contravariant components k^α are related to the coordinate position of the photon via

$$\frac{dx^\alpha}{d\lambda} = k^\alpha, \quad (3.10)$$

where λ being an affine parameter. This equation is supplemented with geodesic equation

$$\frac{dk^\alpha}{d\lambda} = \Gamma_{\beta\gamma}^\alpha k^\beta k^\gamma, \quad (3.11)$$

where the standard definition for the connection coefficients applies in a coordinate frame

$$\Gamma_{\beta\gamma}^\alpha = \frac{g^{\alpha\delta}}{2} \left(\frac{\partial g_{\beta\delta}}{\partial x^\gamma} + \frac{\partial g_{\delta\gamma}}{\partial x^\beta} - \frac{\partial g_{\beta\gamma}}{\partial x^\delta} \right). \quad (3.12)$$

Currently, the integration uses the Verlet algorithm (Dolence et al., 2009), but other integration will be implemented in the near future. The Verlet algorithm has the advantage that the connection coefficients only need to be computed once per integration step.

In this method, the samples are moved with either a fixed step size or a step size that is chosen to be a constant ratio (e.g. 0.1, 0.01, etc.) of the current minimum of the cell size in all directions. Thus a typical integration involves multiple steps per cell. After each step, a check is made to see if the photon sample has changed cells

or progressed the required optical depth. When the photon leaves the cell it is either added to the output (if it has left the domain) or is Lorentz transformed into the new frame (if boosts are calculated) where integration is continued until it scatters or moves to the next neighboring cell and the update is repeated.

The applications presented in Chapter 4 use the cell-by-cell approach, but geodesic integration will be used in future general relativistic calculations.

3.5 SCATTERING AND ABSORPTION

The MC module has several functions for scattering processes including isotropic, Thomson, resonance line, and Compton scattering, as well as user-defined options. Free-free emission and absorption are also included in addition to user-defined options.

Photon samples are moved between scattering/absorption events as described above. If the movement terminates within the simulation domain, scattering and/or absorption will occur. The outcome will depend on the quantity

$$\epsilon = \frac{\alpha_\nu}{\alpha_\nu + \sigma_\nu}, \quad (3.13)$$

which is the ratio of absorption opacity to total opacity. This quantity is evaluated in the comoving frame when boosts between the comoving and lab frames are included.

There are two common approaches and Athena++ implements both methods, with the choice set by flags in the input file. In the first case, a pseudorandom number χ is drawn on the interval $(0, 1)$. If $\chi < \epsilon$, the photon is absorbed and further evolution is discontinued. If $\chi \geq \epsilon$, the photon is scattered using the requested scattering method, and evolution continues. In the second case, scattering always occurs, but the photon sample's weight is reduced with $w' = w\epsilon$, where w' is the new weight after scattering. In this latter case, the photon may never be explicitly absorbed and evolution will continue until it escapes the domain or some other user-defined termination condition is met. If the statistical weight falls below a pre-specified small threshold value (usually based on the initial emissivity), the photon is considered absorbed and further evolution is terminated.

When the photon scatters, a new direction, energy, and polarization may be determined for the outgoing photon depending on the scattering method requested. For this thesis, we focus on the methods of electron scattering, which can be evaluated in the Compton or Thomson limits, with or without polarization included. The out-

going photon energy and direction after Compton scattering follow from procedures described in [Pozdnyakov et al. \(1983\)](#), except that we tabulate the scattering cross-section using a method similar to that described in [Dolence et al. \(2009\)](#). Polarized Thomson scattering is implemented following procedures described in [Wood et al. \(2013\)](#). Polarized Compton scattering utilizes the algorithm described in [Davis et al. \(2009\)](#).

3.6 OUTPUTS

When photons escape through the domain boundary, their energies, locations, and angles are tabulated in a variety of output formats. The MC module provides a few different MC specific outputs including a spectrum, a photon list, images, and photon trajectories. The most versatile output is the photon list, which simply tabulates the properties of each escaping photon. Since no binning is involved, the photon list output tends to produce the largest file sizes. It can, however, be used to generate spectra and images. This is particularly versatile when one does not know a priori what the optimal binning procedure should be for the spectra, or if one wants to generate a non-standard output, such as an escape time distribution. Spectra are computed by binning the photons in angle and/or photon energy, weighted by their statistical weights. Uncertainties are computed by evaluating the corresponding error on the mean within the bin.

The MC calculation can also tabulate cell-averaged radiation moments such as the energy density, radiation flux vector, and pressure tensor, as well as user-defined quantities (e.g. radiative cooling, average photon energy, and average energy mean opacity, etc.). These are output in standard `Athena++` HDF5 and VTK formats and can take advantage of a wide variety of visualization tools.

3.7 TESTS OF THE METHOD

In this section we present some simple, but constraining test cases to show the robustness of the MC module. Here we focus on tests of the module that are relevant to its application in Chapter 4. Since we do not consider time dependence in our MC calculations in Chapter 4, we also do not address time dependence in following test cases.

In the following, we will find it useful to compare our MC results to either a known analytic formula or solutions that are computed using an alternative numerical




Figure 3.1: Spectrum of the isothermal atmosphere test case with just pure absorption shown as the black dots. The solid blue line shows a comparison blackbody spectrum with the same temperature and area as the test case.

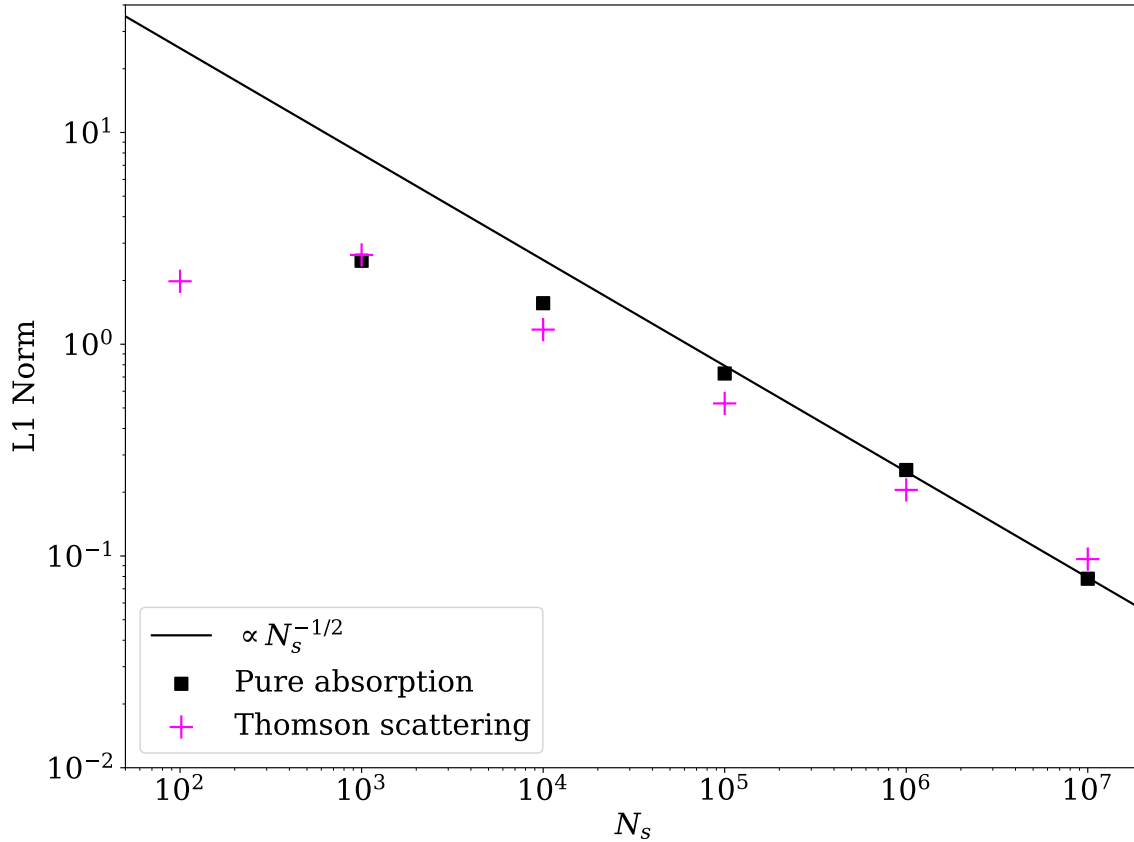


Figure 3.2: Plot of the convergence of two test cases: 1) isothermal atmosphere with only absorption (black squares), and 2) isothermal atmosphere with both absorption and Thomson scattering (pink crosses), as the number of samples increases. The solid black line is proportional to $N^{-1/2}$ and serves as a guide to judge the convergence. The y-axis is the logarithm of L1 Norm.

method. We evaluate the convergence of the MC calculations to these comparison solutions by calculating the L1 norm:

$$\text{L1 Norm} = \frac{1}{N} \sum_i^N \frac{|X_{\text{MC},i} - X_{\text{sol},i}|}{X_{\text{sol},i}}, \quad (3.14)$$

where the sum runs over the resolution elements i of interest (e.g. spectral bins, simulation cells, etc.). Here $X_{\text{MC},i}$ is the MC result and $X_{\text{sol},i}$ is the comparison solution, which is positive for all cases considered here. The L1 Norm is thus the average over all of the the deviations computed in each cell, and we report the mesh sizes for each test case below.

3.7.1 Absorption and Thomson scattering

We first consider the case of a simple isothermal ($T = 10^5$ K) atmosphere with an exponentially declining ρ such that the optical depth ranges from $10^{-3} < \tau < 10^4$ along the $Nx3$ direction. In the case of only free-free absorption (no scattering, polarization, or boosted velocities) photons are initialized with energies on a logarithmic grid from 1 eV to 100 eV.

The mesh setup is $64 \times 64 \times 128$ in $(Nx1, Nx2, Nx3)$, which is distributed over eight meshblocks of size $32 \times 32 \times 64$ using subdomain decomposition and MPI. Periodic boundary conditions are specified along $Nx1$ and $Nx2$, and an absorbing and escaping boundary condition is specified for the inner and outer $Nx3$ boundaries, respectively.

Figure 3.1 shows the spectrum for the case of pure absorption. The solid blue line is a blackbody spectrum with the same temperature and area as the test case. The MC spectrum shown as black dots matches the blackbody spectrum, with a luminosity of $L = 5.7 \times 10^{37}$ erg s^{-1} . We also show the convergence of this case in Figure 3.2. The y-axis is the L1 Norm as a function of N_s . The pure absorption case is shown as the black squares. The solid black line is proportional to $N_s^{-1/2}$, and serves as a guide for judging the convergence. The pure absorption case appears to converge toward $N_s^{-1/2}$ after about $N_s = 10^4$.

We next consider the same isothermal atmosphere setup, but include polarized Thomson scattering. For comparison, we also compute a Feautrier solution (Mihalas & Mihalas, 1984) with which to compare to the MC results. This is a close approximation to the true solution, but we note that the slight differences in the way the two codes are discretized in optical depth and angle: finite differences at specific angles and optical depth for Feautrier, compared to uniform cells and averaged bins for MC. Therefore, we do not expect the methods to converge to arbitrarily high precision. Figure 3.3 shows the spectra (top panel) and polarization (bottom panel) for this test case. Specifically, we plot the energy dependence of the normalized Stokes parameter Q_ν/I_ν , where positive $Q_\nu > 0$ corresponds to polarization in the plane of the atmosphere.

The convergence is shown as the pink crosses in Figure 3.2. In contrast to the pure absorption case, the case with Thomson scattering does not asymptote exactly to the expected $N_s^{-1/2}$ relation and begins to flatten slightly. We attribute this to the slight




Figure 3.3: Spectrum and polarization of the isothermal atmosphere test case including absorption and polarized Thomson scattering. These MC spectra shown as dots are compared to a Feautrier solution (Mihalas & Mihalas, 1984) shown as the solid lines. The top panel shows the spectra for two of the eight angle bins. The bottom panel shows the normalized Stokes parameter Q_ν/I_ν .

difference in discretization between the two calculations as mentioned above. We generally find that convergence continues to higher N_s when we use more optical depth points (Feautrier) and more simulation cells (MC) to resolve the density variation in the vertical direction.

3.7.2 Boost velocities

We consider a uniform box with constant density $\rho = 3 \times 10^{-8} \text{ g cm}^{-3}$, unpolarized Thomson scattering, and free-free absorption. The mesh setup is $8 \times 8 \times 8$ in $(Nx1, Nx2, Nx3)$ distributed over eight meshblocks of size $4 \times 4 \times 4$, and the boundary conditions are periodic boundary in all directions. We test the boosts functionality in which a background velocity field is applied to the box along the $Nx3$ direction. The radiation energy density and flux vector component is then calculated in either the Eulerian or comoving frame. The velocities are in terms of $\beta = v/c = 0.1, 0.2, 0.5, 0.9$ and 0.99 .

Figure 3.4 shows the agreement of the radiation energy density with the boosted analytic expression for radiation in thermal equilibrium with the isothermal temperature. We consider three cases for a range of boost velocities: unboosted radiation, boosted radiation evaluated in the Eulerian frame, and boosted radiation evaluated in the comoving frame. The dotted black line shows the analytic blackbody E_r , which the colored markers indicate the MC tests when the E_r is computed in the Eulerian frame (blue cross mark and purple X mark) and the comoving frame (green diamond mark). The MC module E_r converges exactly in the boosted Eulerian case. In the non-boosted Eulerian and the boosted comoving cases, the E_r is expected to be equivalent to the analytic thermal equilibrium $E_r = aT^4$, and the MC module shows this convergence. Similar results apply for the radiation flux in the boosted direction, although we do not show this result here for the sake of brevity.

We also show the L1 Norm convergence for these three cases in Figure 3.5 for a fixed velocity $v/c = 0.5$. As the number of photons increases, the E_r for each case converges toward the expected $N_s^{-1/2}$ proportionately shown as the solid black line. Note that the number of photons is distributed over a total of 64 cells per meshblock, which effectively raises the variance computed especially at lower numbers of photons.




Figure 3.4: Convergence of the radiation energy density in case of a uniform periodic box with constant density. The analytic blackbody radiation energy density E_r is shown as the dotted black line, and the MC tests are as follows: boosted velocity with E_r computed in the Eulerian frame (blue cross marks), same as mentioned except no boosted velocity (purple X marks), and boosted velocity with E_r computed in the comoving frame (green diamond marks). The boosted velocities are $v/c = 0.1, 0.2, 0.5, 0.9, 0.99$.




Figure 3.5: Convergence of the L1 Norm for each uniform box test case with fixed velocity $v/c = 0.5$ for three cases: boosted radiation evaluated in the Eulerian frame (blue cross marks), unboosted radiation (black X marks), and boosted radiation evaluated in the comoving frame (purple star marks). The solid black line is proportional to $N_s^{-1/2}$ which shows the convergence of the MC tests.

3.7.3 Compton scattering

Next we consider the important test case of a uniform isothermal ($T = 3 \times 10^6$ K) sphere with only unpolarized Compton scattering. The effects of Compton scattering are extremely important for the results in Chapter 4. This case was done for several fixed optical depths, $\tau = 15, 30, 60, 120$. The domain consists of a single zone containing a sphere of radius of 1×10^{10} cm. We specify absorbing boundary conditions, but provide a user-defined escape condition which is triggered when the photon leaves the sphere. The density is adjusted to obtain the target optical depth and no absorption is included. Photons are all initialized at the center of the sphere with an isotropic distribution of directions and an initial energy corresponding to $E = 0.1$ kT. The weights are arbitrarily initialized to 1.

Figure 3.6 shows the resulting spectra computed for the range of optical depths: $\tau = 15$ (black dots), $\tau = 30$ (red dots), $\tau = 60$ (blue dots), and $\tau = 120$ (green dots). The corresponding solid lines show the analytic Green’s function solution to the Kompaneets equation (see equation 33 in Becker, 2003), to which the MC module converges fairly well. Note that the slight excess of the MC runs above the Green’s function solution at the lower optical depths is likely due to the slow convergence of the Green’s function solution (an integral over rapidly oscillatory functions) in this limit, rather than a problem with the MC calculation.

3.7.4 Transport in Curvilinear Coordinates

For a final test, we initialize a spherical polar grid with $N_r \times N_\theta \times N_\phi = 128 \times 64 \times 64$. We then specify that both absorption and scattering opacity are zero. The initial cells and positions of the photon samples are randomly selected along with a random isotropic direction. The photon sample is then propagated without scattering across the grid until it escapes the domain. Upon escaping, the final position \mathbf{x}_f is converted to Cartesian coordinates and recorded for each sample.

Since the photons travel along straight lines, the intersection of the photon with the spherical boundary \mathbf{x}_{sol} can easily be computed analytically for each photon given its initial position and direction. The difference for each photon from its recorded position and its analytically determined position are then tabulated and added to a running total. We compute the average relative error δx_f (similar to equation 3.14)




Figure 3.6: Test spectra for a uniform sphere of fixed optical depth including unpolarized Compton scattering. The optical depths are: $\tau = 15$ (black dots), $\tau = 30$ (red dots), $\tau = 60$ (blue dots), and $\tau = 120$ (green dots). The solid lines of corresponding color show the analytic Green's function solution to the Kompaneets equation (see equation 33 in [Becker, 2003](#)).

defined as

$$\langle \delta x_f \rangle = \frac{1}{N} \sum_i^N \frac{|\mathbf{x}_{f,i} - \mathbf{x}_{\text{sol},i}|}{|\mathbf{x}_{\text{sol},i}|}. \quad (3.15)$$

We evaluate this relative error using 10^4 photon samples. We first evaluate this quantity for photons using the cell-by-cell integration approach and find $\langle \delta x_f \rangle = 1.6 \times 10^{-11}$. Hence, this scheme is almost exact for the position accuracy required by most MC applications. For the geodesic approach, we set the step size parameters to 0.1, 0.01, and 0.001 times the minimum cell width and found $\langle \delta x_f \rangle = 3.6 \times 10^{-3}$, 3.7×10^{-4} , and 3.6×10^{-5} , respectively. Hence, the error in the final position converges roughly linearly with the step size per cell. Although geodesic integration is not as precise as the cell-by-cell transport, these are likely good enough for many MC calculations, with a step size of 0.1 times the cell width generally being computationally faster than the cell-by-cell approach. Although we do not report the results here, we find similar convergence and comparable error for general relativistic integrations.

Finally, we check that neither method produces rare large deviations. For the cell-by-cell approach no values of $\delta x_f > 10^{-9}$ were observed in test with $N = 10^9$ samples. For the geodesic integration, less than 0.01% of samples had deviations that exceeded $10 \langle \delta x_f \rangle$ of the corresponding step size.

3.8 USING THE MC MODULE

We briefly outline the basic steps of utilizing the Athena++ MC module. The module can be run in a time-dependent mode as the Athena++ simulation evolves, with options to couple the MC radiation field back to the fluid. Alternatively, it can be run in a post-processing mode where the fluid properties are fixed and the photons continue evolving until all samples either leave the domain or are absorbed. Here we focus on the post-processing mode, which is most relevant for our calculations presented in Chapter 4.

1. **Configuring and Compiling.** The MC module can be configured and compiled using the standard Athena++ interface. Unlike MHD integrations, most of the properties of MC calculations are set in the input file at the start of the run. In addition to specifying an MC calculation, only the name of the problem

Documentation for configuring and compiling the Athena++ code is available on their GitHub Wiki page: <https://github.com/PrincetonUniversity/athena/wiki>

generator, the coordinate system, and the choice of whether to utilize MPI need to be set prior to compilation.

2. **Initializing the Simulation Grid.** The grid may be initialized in two ways: 1) by reading in a MHD simulation snapshot (in VTK or HDF5 format) with the relevant fluid variables, or 2) by directly initializing the fluid variables in the problem generator file. For the post-processing calculation described in Chapter 4, we use Athena++ RMHD snapshots converted to a uniform static mesh refinement level grid as the input for the MC module. In general, choosing a specific static mesh refinement level is no longer required, but at the start of this project non-uniform meshes were not supported for the MC calculations.
3. **Input File.** Since the MC module makes use of the same Athena++ mesh structures, the standard Athena++ input file simply requires a few additional input blocks to specify the MC mesh domain size, boundary conditions (e.g. reflecting, absorbing, escaping, periodic, user-defined, etc.), and problem-specific input variables. Contained within the MC input blocks are variables such as the desired number of photons, the pseudorandom number generator seed value, and methods of scattering, absorption, emission, polarization, etc. The problem-specific variables relevant to the user-defined problem generator are also defined in the problem block of the input file (e.g. the temperature and density in the case of an isothermal atmosphere calculation).
4. **Running the module.** The MC module can be executed simultaneously with the Athena++ simulation (coupled or uncoupled to the RMHD), or it can be run as a post-processing tool after the RMHD simulation. The total computation time is primarily dominated by the optical depth in the specific problem setup and by the total number of photons used in the simulation. If the user is utilizing the parallel computing capabilities of the module, then the computation time is also affected by the number of cores. For example, the highest signal-to-noise post-processed spectra in Chapter 4 were run with 10^8 photons on 320 parallel cores using the Rivanna HPC at the University of Virginia, which took about 42 hours to complete.

For these computations, we were also limited to no more than 3 days computation time, as set by the Rivanna HPC for parallel computing.

5. **Outputs.** As discussed above, the MC module supports several output formats, with the HDF5 dumps of cell-averaged quantities, spectra files, and photon lists being the outputs most commonly utilized. The HDF5 file contains the radiation moments such as the radiation energy density, flux components, etc. as well as any user-defined quantities. The photon list output collects the properties of the escaping photons which can then be used to make a customized spectrum (e.g. different bin numbers, energy ranges, etc. as opposed to fixed values for spectra requested directly in the input file). To analyze these outputs, we have a collection of Python scripts designed to handle the files and create visuals of the data. Users may also modify these scripts or create their own.

CHAPTER 4

SPECTRAL CALCULATIONS OF 3D RMHD SIMULATIONS OF SUPER-EDDINGTON ACCRETION ONTO A STELLAR-MASS BLACK HOLE

4.1 ABSTRACT

We use the Athena++ Monte Carlo (MC) radiation transfer module to post-process simulation snapshots from non-relativistic Athena++ radiation magnetohydrodynamic (RMHD) simulations. These simulations were run using a gray (frequency integrated) approach but were also restarted and ran with a multi-group approach that accounts for Compton scattering with a Kompaneets operator. These simulations produced moderately super-Eddington accretion rates onto a $6.62 M_{\odot}$ black hole. Since we only achieve inflow equilibrium out to 20-25 gravitational radii, we focus on the hard X-ray emission. We provide a comparison between the MC and RMHD simulations showing that the treatment of Compton scattering in the gray RMHD simulations underestimates the gas temperature in the funnel regions above and below the accretion disk. In contrast, the restarted multi-group snapshots provide a treatment for the radiation field that is more consistent with the MC calculations,

This chapter is adapted from a paper submitted to the *Astrophysical Journal*. The citation is as follow: B. S. Mills, S. W. Davis, Y. F. Jiang, M. J. Middleton, 2023, *submitted to the Astrophysical Journal*.

and result in post-processed spectra with harder X-ray emission compared to their gray snapshot counterparts. We characterize these MC post-processed spectra using commonly employed phenomenological spectral fitting models. We also attempt to fit our MC spectra directly to observations of the ultraluminous X-ray source (ULX) NGC 1313 X-1, finding best fit values that are competitive to phenomenological model fits, indicating that first principle models of super-Eddington accretion may adequately explain the observed hard X-ray spectra in some ULX sources.

4.2 INTRODUCTION

Ultraluminous X-ray sources (ULXs) are point-like, off-nuclear extragalactic objects observed to have X-ray luminosities comparable to or in excess of the critical Eddington luminosity $L_X \gtrsim 10^{39}$ erg/s (assuming isotropic emission for a $10 M_\odot$ black hole; see [Pinto & Walton, 2023](#); [King et al., 2023](#) for a review of ULXs). The majority of ULXs are now accepted to be X-ray binary systems with super-Eddington rates of accretion onto a compact object, namely a stellar mass $M < 100 M_\odot$ black hole ([Poutanen et al., 2007](#); [Middleton et al., 2015a](#)) or neutron star ([Skinner et al., 1982](#); [Bachetti et al., 2014](#)). Some small fraction of ULXs may yet harbor sub-Eddington accretion rates onto intermediate-mass black holes $M \gtrsim 100 M_\odot$ (IMBHs; [Farrell et al., 2009](#); [Mezcua et al., 2013](#); [Earnshaw, 2016](#); [Brightman et al., 2016](#); [Webb et al., 2017](#); [Oskinova et al., 2019](#)). The physical mechanisms which drive super-Eddington accretion are still under investigation and require numerical simulations in order to evaluate existing models of black hole accretion.

The classical picture of an optically thick, geometrically thin accretion disk ([Shakura & Sunyaev, 1973](#)) is used to model black hole X-ray binaries (BHXBs) and is a generally applicable when the accretion rate is sub-Eddington ($L/L_{\text{Edd}} < 0.3$) where the disk geometry (defined by the disc semi-thickness, H and radius, R) remains thin $H \ll R$. If ULXs are indeed IMBHs, the spectra are expected to resemble scaled up versions of BHXB spectra, showing cooler accretion disks as the black hole mass increases (e.g. [Miller et al., 2004](#)).

Observations of ULXs typically show a soft, thermal X-ray component and a hard thermal component with a rollover below ~ 10 keV ([Gladstone et al., 2009](#); [Bachetti et al., 2014](#)), the latter supporting the interpretation of super-Eddington accretion. Early models debated whether this hard X-ray emission originated from coronal emission from IMBHs ([Miller et al., 2004](#)) or Comptonized emission from

super-Eddington accretion (Gladstone et al., 2009; Socrates & Davis, 2006). However, classically, one would expect the innermost regions to have a different spectral shape due to optical depth effects and anisotropy (Poutanen et al., 2007).

Super-Eddington accretion is expected to deviate from the classical Shakura & Sunyaev (1973) thin disk approximation, as the radiation pressure exceeds gravity. Processes like advection (Abramowicz et al., 1988) and radiatively driven outflows (Shakura & Sunyaev, 1973; Ohsuga & Mineshige, 2011) may reduce radiative efficiency and result in geometrically thicker flows in the super-Eddington regime. Strong optically thick winds are also expected to be launched in these systems (and widely detected in ULXs: Middleton et al., 2014, 2015b; Pinto et al., 2016b, 2020; Walton et al., 2016; Kosec et al., 2021), which likely shroud the outer accretion disk and can contribute additional low energy flux for preferential sight lines.

Due to the complex nature of describing three-dimensional super-Eddington accretion flows, numerical simulations are required to replicate these environments. Several radiation hydrodynamic (RHD; Ohsuga et al., 2005), radiation magnetohydrodynamic (RMHD; Ohsuga & Mineshige, 2011; Jiang et al., 2014), and general relativistic RMHD (GRRMHD; McKinney et al., 2014; Fragile et al., 2014; Sądowski et al., 2015; Sądowski & Narayan, 2016) simulations have been performed to understand the physical mechanisms involved in super-Eddington accretion. In these simulations, the radiation transfer equation is often integrated over frequency and angle (so called “grey” radiation) to reduce the computational expense, although a multi-group approach has been shown to efficiently allow for frequency dependence (Jiang, 2022) which we discuss in Section 4.4.3. A consequence of grey radiation simulations is that the radiation moments are not explicitly solved for, which requires either a closure relation (e.g. flux limited diffusion, Turner & Stone, 2001; Howell & Greenough, 2003; Krumholz et al., 2007; Moens, N. et al., 2022; M1 closure, Levermore, 1984; González, M. et al., 2007; Skinner & Ostriker, 2013; Wibking & Krumholz, 2022; or variable Eddington tensor method, Stone et al., 1992; Davis et al., 2012; Jiang et al., 2014; Stone et al., 2020; Asahina et al., 2020; Menon et al., 2022) to solve for the radiation moments.

Spectral post-processing is then required to obtain frequency information from grey simulations. Radiation transfer codes have been utilized (Dolence et al., 2009; Schnittman & Krolik, 2013; Narayan et al., 2016, 2017, Davis et al. in prep) to post-process grey radiation simulation snapshots and generate spectra. This is typically

done with a Monte Carlo method (Dolence et al., 2009) or ray tracing method (Davis et al., 2012). These post-processed spectra have also reproduced some observed spectral features such as the Fe $K\alpha$ line profiles (Schnittman et al., 2013; Kinch et al., 2019) commonly seen in most XRBs, AGN, and currently one ULX (Mondal et al., 2021). These post-processing codes often include a treatment for Compton scattering which plays an important role in describing the radiation field of astrophysical systems and their resulting spectra.

In this work, we use the MC radiation transfer module in **Athena++** to post-process **Athena++** RMHD simulation snapshots and aim to describe these results with current black hole accretion models, as well as compare the simulated spectra to data for the ULX NGX 1313 X-1.

The plan of this work is as follows: In Section 4.3 we discuss the MC and **Athena++** methods used in our spectral post-processing analysis. In Section 4.4 we present our spectral modeling and image results. We discuss the caveats and implications of our results in Section 4.5, and finally we summarize the key points of this work in Section 4.6.

4.3 METHODS

We utilize the **Athena++** code (Jiang et al., 2019; Stone et al., 2020; White et al., 2016) in two configurations – using the **Athena++** RMHD simulation snapshots of super-Eddington accretion onto a $6.62 M_{\odot}$ black hole, and using the **Athena++** Monte Carlo radiative transfer module (Davis et al. in prep) to post-process the snapshots. Here we describe both configurations separately, and discuss the methods used for post-processing in the last subsection.

We note that Section 4.3.1 was adapted from the more detailed descriptions of the MC module in Chapter 3.

4.3.1 MC radiation transfer code

The standard **Athena++** RMHD simulations utilize grey opacities (frequency averaged opacities) and thus do not directly provide any spectral information. To extract frequency information needed to produce the spectra, we utilize the **Athena++** Monte Carlo (MC) radiation transfer module (Davis et al., 2009, Davis et al. in prep) to compute the radiation field throughout an **Athena++** simulation snapshot. The MC module utilizes the **Athena++** code structure and mesh, allowing it to be run con-

currently with the simulations. It can also be utilized to read in output simulation snapshot for post-processing, which is how it is used here. Although the module can be used to perform MC transfer on the full three-dimensional refined simulation mesh, we focused here on two dimensional axisymmetric calculations, where finer/coarser levels are prolonged/restricted to an intermediate refinement level uniform mesh. The MC calculation proceeds by creating and then tracking photon samples throughout the mesh. The samples (often referred to as photon packets or superphotons) can be viewed as statistical ensembles of a large number of photons with common properties. These properties of the photons are initialized and evolved using pseudorandom numbers to draw from distributions in positions, photon energies, scattering angles, etc. until they are either absorbed or leave the domain.

In this work we model free-free emission and absorption and unpolarized Compton scattering as the primary radiative processes. Each photon sample has a statistical weight corresponding to the number of photons in the packet. We model emission by randomly sampling each zone and assigning a weight corresponding to the volume integrated free-free emissivity from the sampled cell. We assume a total number of photon samples N_s and N_{cell} cells in the mesh. If we label cells by index i and photon samples with index j , the total number of physical photons emitted in cell i can be written

$$N_i = \int \frac{j(\nu, T_i, \rho_i)}{h\nu} d\Omega d\nu \mathcal{V}_i \Delta t_{\text{int}}, \quad (4.1)$$

Here, Ω is the solid angle, \mathcal{V}_i is the volume of cell i , $j(\nu, T_i, \rho_i)$ is the free-free emissivity as function of temperature and density within the cell (Rybicki & Lightman, 1979), and Δt_{int} is the (arbitrary) integration time interval. The statistical weights are defined so that

$$\sum_{j=1}^{N_s} w_j = \sum_{i=1}^{N_{\text{cell}}} N_i = N_{\text{ph}}, \quad (4.2)$$

where N_{ph} is the total number of physical photons emitted within the entire mesh. We can define the probability P_i for a photon to be emitted in zone i as $P_i = N_i/N_{\text{ph}} = 1/N_{\text{cell}}$. Then, the average number of photon samples emitted in cell i is $P_i N_s$, and we have

$$w_i = \frac{N_i N_{\text{cell}}}{N_s}. \quad (4.3)$$

This procedure yields photon weights that can differ by orders of magnitude. This is

often frowned upon in the MC literature because more uniform weighting is generally variance reducing. We have, however, also implemented an equal weighting scheme where the initial cells of photon samples are chosen proportional to their volume weighted emissivity and found this scheme ultimately results in larger statistical errors in our output spectra per computational second when compared with the scheme used here. This is primarily due to the large scattering optical depths to escape for photons launched in the highest emission cells (Davis et al. in preparation).

Finally, the direction of the photon is randomly sampled from an isotropic distribution, and the energy of the photon is drawn from a log normal distribution in photon energy. We then further adjust the weight so that binned photons match free-free distribution in photon frequency. Photon movement is handled in the Eulerian (coordinate) frame, while emission, scattering, and absorption occur in the comoving fluid frame. Photon sample properties are Lorentz boosted between the coordinate and fluid (comoving) frame for these interactions.

Photon samples are moved between scattering/absorption events by drawing an exponentially distributed dimensionless path length τ to the next absorption/scattering event via $\tau = -\ln \xi$, where ξ is a pseudorandom number uniformly distributed in the interval (0,1). This dimensionless path length can be thought of as the optical depth to the next scattering/absorption event, and is computed as a series of steps l_k (enumerated with subscript k) so that

$$\tau = \sum_k \frac{l_k}{(\alpha_{\nu,k} + \sigma_{\nu,k})^{-1}}, \quad (4.4)$$

where α_{ν} is the absorption extinction coefficient, and σ_{ν} is the scattering extinction coefficient. The scattering and absorption coefficients are the products of the corresponding opacities and density, which are evaluated in the comoving frame and then boosted to the Eulerian frame. In the scheme used here each step k represents a movement of the photon sample to the location of the next scattering/absorption event or the nearest cell face, whichever comes first. This continues until the requisite value of τ is reached or the photon sample escapes the domain. Photon samples are assumed to travel along straight lines, but we use a spherical mesh, so that computing where the photon sample leaves the current cell requires solving quadratic relations and accounting for possible turning points in r and θ (Davis et al., in preparation).

Each interaction of photon sample with matter results in a combination of absorption and scattering, which is handled by reductions in w . We have $w' = w\epsilon$, where w' is the new weight after scattering and

$$\epsilon = \frac{\alpha_\nu}{\alpha_\nu + \sigma_\nu}. \quad (4.5)$$

If the statistical weight falls below a small threshold value (based on the initial emissivity), the photon is considered absorbed and further evolution is terminated. The outgoing photon energy and direction after Compton scattering follow from procedures described in [Pozdnyakov et al. \(1983\)](#), except that we tabulate the scattering cross section using a method similar to that described in [Dolence et al. \(2009\)](#).

When photons escape through the domain boundary, their energies, locations, and angles are tabulated in a photon list output that is then used to generate spectra. The MC calculation also tabulates cell-averaged radiation moments such as the energy density, radiation flux vector, and pressure tensor, as well as user defined quantities such as the net radiative cooling, average photon energy, and average energy mean opacity in each cell. These are output in standard `Athena++` formats, such as HDF5 and VTK.

4.3.2 `Athena++` RMHD simulation snapshots

`Athena++` has been rewritten in C++ compared to its predecessor, `Athena` ([Stone et al., 2008](#)). `Athena++` now includes adaptive mesh refinement ([Stone et al., 2020](#)) and special and general relativistic capabilities ([White et al., 2016, 2023](#)). In the current work, however, a pseudo-Newtonian potential is used to mimic the effects of general relativity around a Schwarzschild black hole ([Paczynski & Wiita, 1980](#)). Results from a GRRMHD implementation of `Athena++` and subsequent spectra will be reported in future work.

We performed a series of global, three-dimensional RMHD simulations for a $6.62 M_\odot$ black hole accreting at several super-Eddington mass accretion rates assuming a 10% radiative efficiency so that $\dot{M}_{\text{Edd}} \equiv 10L_{\text{Edd}}/c^2$. We used the explicit integration RMHD module in `Athena++`, which uses an algorithm similar to [Jiang et al. \(2014\)](#), but with updates that solve a radiation transfer equation of the form presented in [Jiang \(2021\)](#). The simulation setup for these snapshots is similar to the setup described in [Huang et al. \(2023\)](#), where the ideal MHD equations are coupled with




Figure 4.1: The mass accretion rate \dot{M} as a function of gravitational radius r_g from the black hole for each Athena++ RMHD simulation snapshot (see Table 4.1). ULX4a is the dot-dash line, ULX4b is the dashed line, ULX2.5 is the dotted line, and ULX1.3 is the solid line.

Table 4.1: Athena++ RMHD simulation snapshots of a $6.62M_{\odot}$ black hole used in this analysis. All snapshots are azimuthally averaged and are limited to the inner $25r_g$. The first column corresponds to the ratio of the radially averaged mass accretion rate $\langle \dot{M} \rangle$ in terms of the Eddington mass accretion rate \dot{M}_{Edd} . The negative sign indicates accretion towards the black hole. The second column corresponds to the polar funnel angle θ_f , the opening angle relative to the polar axis representing the approximate boundary between the funnel region and the accretion disk. Photons emerging from this polar funnel angle are collected for spectral post-processing and have corresponding funnel luminosities L_f . The last column is the calculated radiative efficiency η_f of the funnel region. Snapshots ULX4a and ULX4b were taken from the same simulation run (at different times), whereas Snapshots ULX2.5 and ULX1.3 are independent simulation runs. The snapshots with the suffix “-MG” correspond to the two grey simulations chosen for the multi-group RMHD implementation (Jiang, 2022).

Snapshot	$\langle \dot{M} \rangle / \dot{M}_{\text{Edd}}$	θ_f	L_f (erg cm ⁻² s ⁻¹)	η_f
ULX4a	-4.15	37°	1.03e+39	2.56%
ULX4b	-3.93	37°	8.77e+38	2.29%
ULX2.5	-2.53	50°	2.78e+38	1.13%
ULX1.3	-1.31	55°	1.51e+38	1.18%
ULX4a-MG	-4.02	37°	1.31e+39	3.34%
ULX2.5-MG	-2.53	50°	4.73e+38	1.92%

the time-dependent radiation transfer equation (see Jiang et al., 2014 equations 1-4, and Jiang 2021 equations 4-6). A rotating gas torus was initialized in hydrostatic equilibrium and threaded with toroidal magnetic fields. Accretion onto the black hole happens via the magnetorotational instability (Balbus & Hawley, 1991) and the mass accretion rate is varied for each simulation based on the initial magnetic field configuration (see e.g., Huang et al., 2023).

The simulations self-consistently form an accretion disk and reach a quasi-steady state for the inner disk. Figure 4.1 shows the mass accretion rate in terms of \dot{M}_{Edd} for a $6.62M_{\odot}$ black hole as a function of radius within the inner $25 r_g$ where $r_g = GM/c^2$ is the gravitational radius. The mass accretion rates are relatively steady-state within $25r_g$ of the black hole, but not at larger radii. The four snapshots and their radially-averaged mass accretion rates (over the $25r_g$) are listed in Table 4.1. Snapshots ULX4a and ULX4b have nearly the same mass accretion rate ($\dot{M} \simeq -4\dot{M}_{\text{Edd}}$) and are both from the same simulation run (at different times), thus we named them ULX4a and ULX4b. Snapshot ULX2.5 and Snapshot ULX1.3 are independent simulation runs with average mass accretion rates $\dot{M} \simeq -2.5$ and $-1.3\dot{M}_{\text{Edd}}$, respectively.

Although the Athena++ RMHD calculations have adaptive mesh refinement capabilities, the MC code works most efficiently on a uniform grid. For efficient parallelization, we chose one uniform refinement level for our analysis. We selected an appropriate refinement level such that all snapshot grids were approximately the same size $256 \times 128 \times 256$ cells in r , θ , and ϕ (respectively). The accretion disk located in the inner $25r_g$ roughly corresponds to the 80 innermost zones in radius at this level, and covers a range of θ from 0 to π , and a range of ϕ from 0 to 2π . Due to the approximately axisymmetric nature of the simulations, we chose to azimuthally average each snapshot for our post-processing analysis. This has little effect on the output spectra, but greatly improves the statistics for cell averaged quantities, examples of which are presented in Figures 4.3 – 4.6.

Figure 4.2 shows the gas density in Snapshot ULX2.5 for the inner $25r_g$ where the accretion disk has roughly reached inflow equilibrium, and the small inset plot shows the the full simulation grid out to $500r_g$. The full simulation grid includes the geometrically thick gas torus extending from $\sim 100r_g$ to $\sim 300r_g$. The densities in the funnel regions are several orders of magnitude lower than the densities in the optically thick accretion disk and gas torus. We discuss the implications of the low density funnel region and the impact of the torus geometry in Section 4.4.1.




Figure 4.2: Azimuthally averaged gas density ρ (g cm^{-3}) of the RMHD simulation snapshot ULX2.5 showing the inner $25 r_g$ with a small inset plot that shows the gas density out to $500 r_g$.

The net cooling in the **Athena++** RMHD simulations is given by

$$\dot{C} = c\rho (\kappa_{\text{P}} a T_{\text{g}}^4 - \kappa_{\text{E}} E_{\text{r}}) + c\rho \kappa_{\text{es}} \frac{4kT_{\text{g}} - \langle h\nu \rangle}{m_e c^2} E_{\text{r}}, \quad (4.6)$$

where c is the speed of light, ρ is the gas density, κ_{P} is the Planck mean opacity, a is the Planck temperature constant, T_{g} is the gas temperature, κ_{E} is the energy mean opacity, E_{r} is the radiation energy density, κ_{es} is the electron scattering opacity, k is the Boltzmann constant, h is the Planck constant, $\langle h\nu \rangle$ is the average photon energy, and m_e is the electron mass. The first term is the frequency and angle integrated free-free emissivity $\eta_{\text{ff}} = c\rho \kappa_{\text{P}} a T_{\text{g}}^4$. The second term is the heating term associated with absorption, and the last term is the net Compton cooling. In the RMHD simulations, the radiation field is assumed to be blackbody so $\langle h\nu \rangle = 4kT_{\text{r}}$, where T_{r} is the radiation temperature $T_{\text{r}} = (E_{\text{r}}/a)^{1/4}$. The simulations also assume that $\kappa_{\text{E}} = \kappa_{\text{P}}$. These assumptions and their impact on the gas temperature distribution and the spectra that result will be discussed further in Section 4.4.

The gas temperature of the same snapshot in Figure 4.2 is shown in Figure 4.3. Note that the apparent asymmetry of the gas temperature in the funnel regions above and below the disk is due to the randomness in the flow at the time this snapshot was taken. Prior to post-processing, we set a lower limit on the gas temperature of 10^6 K and an upper limit of 3×10^8 K (except for the multi-group snapshots, which we set the upper limit to 10^9 K). The gas temperature is hottest in the funnel region where it hits the temperature cap of 3×10^8 K, and the gas in the accretion disk peaks at a few $\times 10^7$ K. Although the temperatures in the funnel regions are large, the corresponding gas densities from Figure 4.2 are small (10^{-8} g cm $^{-3}$) so the contribution to the emission from the hottest simulation cells is relatively weak. The white contour lines roughly define the effective photosphere boundary between the accretion disk and the funnel region, defined by $F_{\text{r}}/cE_{\text{r}} = 0.3$, where F_{r} is the $r\phi$ component of the radiative flux, and E_{r} is the radiation energy density. This flux ratio is consistent with methods that define the photosphere by integrating to an optical depth $\tau = 1$ surface (Chandrasekhar, 1960; Kinch et al., 2019). The polar angle of this boundary is used to approximate the funnel opening angle θ_{f} which is then used to calculate the luminosity, spectra, and images in Section 4.4.2.




Figure 4.3: Gas temperature T_g (K) shown for Snapshot ULX2.5. A white contour line corresponds to $F_r/cE_r = 0.3$, which is roughly equivalent to the effective photosphere boundary and defines the polar funnel angle $\theta_f = 50^\circ$ from the polar axis. The apparent anisotropy of the gas temperature in the funnel regions above and below the disk is a result of this particular simulation taken at this moment in time. The temperatures were capped at a maximum of 3×10^8 K and a minimum of 10^6 K.

4.3.3 Spectral post-processing

Here we describe the methods used for performing our spectral analysis. Spectra were generated for each azimuthally averaged snapshot, truncating the calculation to only model MC transfer within $25r_g$. The properties of all photon samples leaving the domain at $25r_g$ are tabulated in a list, which is then used to generate spectra. This truncation radius was chosen primarily because the outer disk radii are not yet in steady state. In particular, the initial torus is thick, which requires an extremely large radiation pressure. Hence, this torus is not in thermal equilibrium, and is rapidly cooling. We then ran the MC code on a copy of each snapshot grid, initializing 10^7 photons for Snapshots ULX4a and ULX4b, and about 10^8 photons for the other two snapshots. The reason for the difference is due to the larger optical depths and mass accretion rates in ULX4a and ULX4b that result in a factor of ~ 10 difference in number of scatterings per photon sample. Recall that the number of scatterings per photon sample is proportional to the square of the optical depth. Increasing the number of photons in the MC calculations greatly improves the counting statistics, however we found that more than 10^7 photons for those snapshots became too computationally expensive due to the large scattering optical depths.

Photons that escaped the $25r_g$ simulation domain were collected and distributed into 64 photon energy bins ranging from 0.1 keV to 60 keV, and eight direction angles. By direction angle, we mean the angle θ_p that the photon momentum vector makes with the polar axis. We use the subscript p to distinguish angles related to the photon momentum from those related to spherical polar coordinate angles. For example,

$$\theta_p = \arccos\left(\frac{\mathbf{p} \cdot \hat{\mathbf{z}}}{|\mathbf{p}|}\right), \quad (4.7)$$

where \mathbf{p} is the photon momentum vector. These angle bins are distributed uniformly in $\cos\theta_p$ and integrated over azimuthal direction angle ϕ_p . When binning, we do not distinguish between photons leaving above or below the disk. For example, photons with $\theta_p \sim 0$ will be placed in the same bin as photons with $\theta_p \sim \pi$.

We select only photons which escape through a “funnel” like region above and below the disk. Specifically, we only bin photons within a coordinate opening angle of θ_f from the polar axes, retaining photons leaving the domain at $\theta < \theta_f$ or $\theta > \frac{\pi}{2} - \theta_f$. This excludes photon samples that leave domain closer to the midplane. Such photons

would almost certainly be further scattered in the optically thick accretion disk if we extended our domain outwards. Hence, we select our funnel opening angle θ_f to roughly correspond to the location of the disk photosphere at $25r_g$. The approximate values for this funnel opening angle are listed for each snapshot in Table 4.1.

Due to these selections, the resulting spectra are only expected to be useful estimates of the hard X-ray emission as the softer X-ray will have a significant contribution from regions with $r > 25r_g$. We also cannot infer much about the angular distribution of the escaping photons for angles that are more edge-on than θ_f as such photons would likely interact with an optically thick flow beyond $r = 25r_g$.

In the case of snapshot ULX2.5 we also perform a MC calculation using the full simulation domain. In this case we collect all photon samples leaving the domain, but find that the spectrum of the escaping radiation is dominated by contributions from the torus. Due to large optical depths in the outer torus, the calculation is computationally expensive and run with fewer photons, yielding a lower signal-to-noise spectrum. For these reasons, we do not report spectra from these runs, but we do use the cell-averaged radiation outputs for comparison with the truncated runs described above.

4.4 RESULTS

4.4.1 Comparing Athena++ with Monte Carlo

We first compare cell averaged quantities. Figure 4.4 shows a comparison of E_r computed with MC to the azimuthally-averaged E_r from the Athena++ simulation snapshot. This figure is for Snapshot ULX2.5, but the result is representative of all four snapshots in this analysis. The RMHD simulation result is plotted in the left panel, and two MC calculations are plotted in the middle and right panels respectively. The middle panel shows the results of an MC calculation using the full simulation domain out to $\sim 500r_g$, whereas the right panel shows the MC calculation when the grid is truncated at $25r_g$. The two MC calculations show precise agreement in the accretion disk, where the radiation field is nearly in radiative equilibrium with the gas. They also agree reasonably well in the funnel regions, deviating by only a small factor near the outer edge of the truncated domain. This suggests that the E_r in the inner $25r_g$ is dominated by the locally emitted radiation field, since the truncated calculations have no incoming photons on the boundary. Therefore, the radiation from




Figure 4.4: Azimuthally-averaged radiation energy density of Snapshot ULX2.5 for the Athena++ RMHD simulation (left panel), the MC calculation (middle panel), and the MC calculation with the simulation grid truncated at $r = 25r_g$ (right panel). In the left and middle panels, only the inner $25r_g$ are plotted here for comparison, but the full simulation grids extend out to $500r_g$.

the cooling torus, which dominates the overall emission in the full domain calculation, is not providing a significant contribution in the inner disk region. Our comparison suggests that radiation outside the truncated domain is contributing $\lesssim 30\%$ near $25r_g$, and $\lesssim 15\%$ near the photosphere boundary of the disk.

Note that at the very edge of the truncation boundary, E_r is slightly lower compared to the radiation energy density in the full domain calculation, as the truncated calculation assumes no incoming radiation flux. The more noticeable streaks of high E_r noise in the funnel regions in the MC full domain calculation are attributed to the factor 10 fewer photons used to compute the full grid, resulting in a larger statistical variance.

Comparing the gray RMHD module E_r to the MC calculations in Figure 4.4, they also appear to agree within a factor of order unity in the accretion disk midplane, but start to deviate more significantly as one transitions into the funnel region. In the funnel, this deviation is as much as a factor 10. The MC calculations find a significantly lower E_r in the funnel region. This mismatch is even more evident in Figure 4.5, which shows the ratio of the two calculated energy quantities: the mean photon energy $\langle h\nu \rangle$ calculated by the truncated MC calculation, and the mean photon energy $4kT_r$ assumed in the *Athena++* RMHD module. The dark regions where the ratio is of order unity show that the MC and *Athena++* generally agree in the accretion disk, but deviate in the funnel region above and below the disk. In these regions, the MC calculates that average photon energy is at least three times higher than assumed in the RMHD run. The assumption that the radiation field is approximately blackbody works well for the optically thick accretion disk regions, but is inadequate in the optically thin funnel regions.

In Figure 4.6, we compare the resulting cooling computed by the *Athena++* RMHD simulation (left panel) to the same term evaluated by the MC calculation (right panel) for the same snapshot. The net cooling is calculated using Equation 4.6 where positive values indicate cooling and negative values indicate heating. Since the Compton cooling is the dominant term in equation (4.6) within the funnel region, this comparison is strongly dependent on the degree to which $\langle h\nu \rangle$ differs from $4kT_r$ and the ratio of E_r in the MC calculations relative to RMHD (see Figures 4.4 and 4.5). We find that the cooling calculated by the MC code deviates significantly from that of the RMHD simulation, particularly in the funnel regions, where the MC code shows significantly more heating and less cooling. In the accretion disk, the MC code pro-




Figure 4.5: Ratio of the mean photon radiation energy $\langle h\nu \rangle$ calculated in the Monte Carlo code and the radiation energy $4kT_r$ in the RMHD simulation for Snapshot ULX2.5 where T_r is the radiation temperature (assuming the blackbody approximation). The streaks in the funnel region are artifacts of low photon statistics in the Monte Carlo calculation.




Figure 4.6: Comparison of the net cooling of Snapshot ULX2.5 in the **Athena++** RMHD simulation (left panel) and the Monte Carlo $25r_g$ calculation (right panel). The snapshot has been azimuthally averaged in both cases. The net cooling is given by Equation 4.6 where positive values signify net cooling and negative values imply net heating. The black cells extending into the photosphere in the RMHD calculation are artifacts of this particular moment in the simulation.

vides slightly less cooling than the RMHD simulation does. We attribute this to the MC calculation not including stimulated scattering, which would contribute slightly more cooling in the disk.

These results suggest that if the RMHD simulations had a better estimate for $\langle h\nu \rangle$ it would be higher than $4kT_r$. In an approximate steady-state with the Compton cooling term dominating in the funnel region, one expects the $\langle h\nu \rangle \simeq 4kT_g$. By underestimating $\langle h\nu \rangle$, the RMHD simulations tend to underestimate T_g in the optically thin regions above the disk. This underestimate also tends to increase T_r to better balance T_g , causing the RMHD simulations to overestimate E_r , consistent with our findings above. This also means that spectra computed from these snapshots will have lower average photon energies than one might obtain in a simulation with more self-consistent thermodynamics, which would yield higher T_g and harder X-ray spectra. We explore the implications of this in Section 4.4.3.

4.4.2 Spectra & Compton Cooling of Gray RMHD simulations

We present post-processed X-ray spectra for the four gray RMHD snapshots in Figure 4.7. For the lower \dot{M} snapshots (ULX1.3 and ULX2.5), the spectral peaks are roughly around 5 keV, whereas for the higher \dot{M} snapshots (ULX4a and ULX4b) have peaks that are shifted slightly to around 7 keV. The hard X-ray tails appear to follow power laws, which we characterize with `XSPEC` (Arnaud, 1996) model fits in Section 4.4.4. Due to the truncation of the simulation grids at $25r_g$ prior to post-processing, the softer X-ray emission that should be coming from larger radii is largely absent in these spectra so only the hard X-ray are self-consistently modeled.

The frequency-integrated luminosities for each spectrum are tabulated in Table 4.1. We label these funnel luminosities L_f to emphasize that we only tabulate the contributions from photons leaving the domain at coordinate θ within an angle θ_f of the polar axes. Note that these are mostly hard X-ray luminosities due to the missing soft emission from the outer disk. The contour lines in Figure 4.3 approximate the funnel opening angle for ULX2.5 ($\theta_f = 50^\circ$) which we show as a representative snapshot. The θ_f and corresponding funnel luminosities L_f are listed for each snapshot in Table 4.1.

Photons emerging closer to the disk midplane than θ_f are excluded because their escape from the truncated domain at $r = 25r_g$ is largely artificial. If we had instead extended our MC calculation domain outward in radius, these photons would likely




Figure 4.7: Monte Carlo post-processed X-ray spectra from the grey RMHD simulation snapshots. From top to bottom: ULX4a (magenta dash-dot line), ULX4b (green dotted line), ULX2.5 (blue dashed line), and ULX1.3 (black solid line). Snapshots ULX4a and ULX4b were taken from the same simulation run, while ULX2.5 and ULX1.3 are both from independent simulations. Note that these spectra only include the inner $25r_g$ emission escaping out through a polar funnel angle θ_f specified in Table 4.1 for each snapshot.

experience additional scattering and absorption in the optical thick flow before escaping. Slight variations in θ_f can have a modest effect on the funnel luminosity and the resulting spectral shape. For example, in the case of ULX2.5 the funnel luminosity varied by less than 17% when varying θ_f by $\pm 10^\circ$. Choosing a narrower funnel angle ($\theta_f = 40^\circ$) gave a luminosity of $L_f = 2.04 \times 10^{38}$ erg/s, whereas choosing a wider funnel ($\theta_f = 60^\circ$) gave a slightly higher luminosity of $L_f = 3.36 \times 10^{38}$ erg/s. Increasing θ_f , however, results in a slight softening of the spectrum as there is an increase in the flux of photons escaping below what would be the photosphere in a more extended domain. These photons tend to be softer because they are emitted from the cooler regions of the disk. We found this to be true for all snapshot spectra in this analysis.

The radiative efficiency calculated for Snapshot ULX2.5 is $\eta_f = 1.51\%$ for a nominal mass accretion rate of $-2.53\dot{M}_{\text{Edd}}$. We report the calculated η_f for each snapshot in Table 4.1. Generally for super-Eddington accretion, it is expected that the radiative efficiency will be lower than the $\sim 5 - 10\%$ inferred for thin disks, decreasing as accretion rate increases. We do generally find lower efficiencies, but the results are not completely consistent with expectations. Comparing the efficiencies for snapshots ULX2.5 and ULX1.3 which have similar θ_f , we infer a slightly lower efficiency as accretion rate increases. Snapshots ULX4a and ULX4b, however, have higher \dot{M} than the other snapshots, but also show a higher η_f . It is possible that this deviation from the expected trend is a result of our truncation of the calculation at $r = 25r_g$ and merits more consideration in future work exploring a wider range of \dot{M} .

We also examine the angular distribution of the emission, which we model as the flux fraction (ratio of specific intensity I to the flux F) for each spectrum in Figure 4.8. For observations, this should roughly correspond to the inclination viewing angle dependence with respect to the polar axis. A face-on view of the emission corresponds to $\cos \theta_p = 1$, and an edge-on view corresponds to $\cos \theta_p = 0$. Snapshot ULX1.3 is shown as the solid black line, ULX2.5 is the dashed blue line, ULX4b is the dotted green line, and ULX4a is the dash-dot pink line. The flux fraction has been integrated over all frequencies ν for improved statistics, so the resulting distribution is most representative of the angular distribution near the spectral peak.

Although we show the full distribution for $\cos \theta_p$ ranging from 0 to 1, we emphasize that the $\cos \theta_f$ ranges from 0.57 for ULX1.3 to 0.8 for snapshots ULX4a and

We define \dot{M}_{Edd} assuming 10% efficiency, but \dot{M} itself is independent of our assumed efficiency.




Figure 4.8: Limb darkening (flux fraction of the specific intensity I to the isotropic flux F) as a function of inclination angle in terms of $\cos\theta$ for Snapshots ULX4a (pink dash-dotted line), ULX4b (green dotted line), ULX2.5 (blue dashed line), and ULX1.3 (black solid line). Note that each snapshot spectrum was generated using only photons which escape through a funnel opening of polar angle θ_f specified in Table 4.1. Face-on viewing corresponds to $\cos\theta_p = 1$ and edge-on viewing corresponds to $\cos\theta_p = 0$.

ULX4b. Hence, only the bins with $\cos \theta_p$ greater than these values are likely to be well-characterized. Over this limited range, the angular distributions are relatively flat, but notably do not peak at the most face on inclination bin. This is contrary to standard expectations where a face on view provides the largest projected area and, thus, the largest flux. As θ_p approaches the edge-on view the intensity declines by factors of several. However, we emphasize that the intensity distribution at these angles will undoubtedly be impacted by the extension of the optically thick disk outside of the calculations domain. For example, the slight rise in the most edge-on bin is almost certainly a result of our artificial truncation of the simulation domain. Hence, our current results cannot provide reliable predictions about geometric beaming factors.

To better interpret these results, we show a set of reconstructed images in Figure 4.9, which shows the frequency-integrated intensity from the funnel region at different inclination angles $\theta_p \sim 49^\circ$ (left column; funnel edge view) and $\theta_p = 0^\circ$ (right column; face-on view) for two snapshots: the gray RMHD snapshot ULX2.5 (bottom row), and the multi-group RMHD snapshot ULX2.5-MG (top row). We discuss the latter snapshot in detail in the next section. The corresponding opening angle for both snapshots is $\theta_f = 50^\circ$. Photons escaping the funnel were extrapolated out to a distance of $\sim 250,000 r_g$ to form these images. In the face-on case, we see a deficit for photons near the polar axis, along the line of sight to the black hole. In this region the densities are so low that relatively few photons are scattered or emitted toward the observer. Near the edge of the funnel, the intensity of the emission appears to brighten compared to the face-on inclination. This enhancement is consistent with modest amounts of relativistic beaming in the mildly relativistic outflowing gas. This beaming is largest at these moderate inclinations where both the line-of-sight outflow velocities and scattering optical depths are large.

4.4.3 A multi-group RMHD approach

As discussed in Section 4.4.1, the $\langle h\nu \rangle = 4kT_r$ assumption in the gray RMHD simulations likely results in gas temperatures being underestimated in the regions above the optically thick disk. This, in turn, means that the MC spectra we compute are probably softer than they should be if the temperatures were computed with a more self-consistent treatment of Compton scattering. In an effort to recompute the gas temperature, we first tried to use the MC code to calculate the net cooling everywhere in the simulation and balanced this with the dissipation from the RMHD




Figure 4.9: A series of frequency-integrated images showing the emergent radiation intensity for snapshots ULX2.5-MG (top row) and ULX2.5 (bottom row) for two inclination viewing angles. The left column shows a viewpoint from an inclination of $\theta_p \sim 49^\circ$, which is at the edge of the funnel for these snapshots ($\theta_f = 50^\circ$). The right column views the face-on inclination ($\theta_f = 0^\circ$). The photons leaving the simulation domain at $25r_g$ were extrapolated out to a distance of $250,000r_g$ from the black hole.

snapshots. We found, however, that the recomputed temperatures in the funnel had too large of a variance due to the limited photon statistics, and did not consistently converge after several iterations.

Instead, we utilized the multi-group radiation module described in [Jiang \(2022\)](#), which extends the grey radiation scheme in [Jiang \(2021\)](#) to include frequency dependence and treats Compton scattering using a Kompaneets like approximation for the electron scattering source term. We used 20 logarithmically distributed frequency groups to cover the frequency space over three orders of magnitude, which increased the computational cost by a similar factor. Hence, a full three-dimensional simulation with this method would be extremely computationally expensive. Here, we instead start the multi-group simulation by restarting the gray simulation by assuming the initial spectrum to be blackbody, and run for a time just long enough for the gas above the disk to reach a new temperature equilibrium, making the computational expense feasible for this study. This is possible because the thermal timescale in the funnel region is very short.

We performed the multi-group procedure for two of the four snapshots, ULX2.5 and ULX4a, using the restart files from the gray RMHD simulations and running them with 20 frequency groups. We label the new snapshots from these multi-group runs ULX2.5-MG and ULX4a-MG. These snapshots were computed at approximately the same time as their gray counterparts, and thus have the same average mass accretion rates (see [Table 4.1](#)) and the density distributions are quite similar. But, as expected from the MC calculations of cooling rates in the gray snapshots, these new runs find larger gas temperatures in the regions near or above the photosphere of the accretion flow. We post-process these snapshots following the same procedures as we did for the gray snapshots.

[Figure 4.10](#) shows a comparison of the radiation energy density for ULX2.5-MG from the multi-group RMHD simulation (left panel) and the MC calculated radiation energy density (right panel), analogous to the comparison for the grey snapshot ULX2.5 in left and right panels of [Figure 4.4](#). As expected, the E_r in the multi-group approach is lower compared to its gray counterpart, and the comparison with MC for the multi-group approach is in closer agreement, although not exact. Exact agreement is not necessarily expected as the Kompaneets treatment in the RMHD module differs slight from the MC treatment. It also possible that the (computationally expensive) multi-group calculation has not yet reached full equilibrium.




Figure 4.10: Comparison of the radiation energy density E_r for ULX2.5-MG from the multi-group RMHD snapshot (left panel) and the same quantity calculated by the Monte Carlo module (right panel). In both cases, the simulation has been azimuthally averaged.

We do not show the angular distributions of the emitted spectra for these multi-group snapshot calculations because they are rather similar to their counterparts shown in Figure 4.8. We show, however, a second set of reconstructed images in the top panel of Figure 4.9 for ULX2.5-MG. Compared to its gray counterpart in the top panel, the overall intensities in the multi-group approach are larger due to the larger temperatures, but the mild relativistic beaming again enhances the intensities for off-axis viewing angles relative to those of the most face-on image.

Figure 4.11 shows the MC post-processed spectra from the multi-group snapshots compared to their corresponding grey snapshot counterparts. The multi-group approach leads to harder spectra due to the larger gas temperatures. This is seen as both a shift in the spectral peak and a somewhat flatter power law dependence at higher energies. The effect is larger for ULX2.5-MG than ULX4a-MG. The overall luminosity of the funnel for the multi-group spectra are also larger, with $L_f = 4.73 \times 10^{38}$ erg cm⁻² s⁻¹ for ULX2.5-MG, and $L_f = 1.31 \times 10^{39}$ erg cm⁻² s⁻¹ for ULX4a-MG. There is also a commensurate increase in the radiation efficiencies since the accretion rates were essentially unchanged.

4.4.4 Simulated spectra in comparison with observational models

Here we quantitatively characterize the post-processed spectra by utilizing X-ray spectral fitting models commonly used to describe observations of black hole sources. The motivation here is to get a sense of how the combinations of phenomenological models describe the hard X-ray emission and quantitatively compare between the post-processed grey RMHD spectra and the multi-group RMHD spectra. Although we utilize spectral fitting methodology as a tool to compare our simulations to other models, we emphasize that these are not fits to data, and we make choices in accordance with these considerations.

We use the X-ray spectral fitting package XSPEC (Arnaud, 1996) version 6.26.1 to explore a few different model combinations. We note that the two-component (soft+hard) phenomenological model combinations we use here are typically used to fit BHXB spectra whilst ULX spectra can also be described by two component models (as shown using variability studies: Middleton et al. 2015a), where the components refer to regions in the super-Eddington disk, modified by opacity in the wind and anisotropy (Poutanen et al. 2007). In addition, ULX spectra sometimes require a




Figure 4.11: MC post-processed spectra from the grey RMHD simulation snapshots (ULX2.5 and ULX4a) shown as the lighter blue and purple colored dotted lines, respectively, and the MC spectra from the multi-group RMHD implementation shown as the corresponding darker solid lines. The spectra for ULX2.5 and ULX2.5-MG were computed for a funnel region of $\theta_f = 50^\circ$, while ULX4a and ULX4a-MG were computed for $\theta_f = 37^\circ$.

third component at higher energies from a pulsing component (an accretion column: [Brightman et al. 2016](#); [Walton et al. 2018](#)) which has led to speculation that a generic hard excess compared to thermal models could indicate the presence of a magnetised neutron star ([Pintore et al. 2017](#); [Walton et al. 2018](#)). By comparison, our spectra only correspond to the innermost regions and so miss a large portion of the soft X-ray emission from the outer radii. Therefore, we only use one soft X-ray model and one hard X-ray component to describe our simulated spectra, and we primarily focus on the hard X-ray part of the spectra. Since we only want characterize our simulation spectra, we do not include any absorption model components that are typically used to account for the interstellar medium along the line of sight.

The simulated spectra were transcribed into table models containing the energies and flux counts that could be loaded into XSPEC. We used the `fakeit none` command to generate artificial “datasets” based on each simulated table model. For the required response file during this process, we input a *NuSTAR* FPMA detector response file provided by one of the [Gúrpide et al. \(2021\)](#) observations of NGC 1313 X-1 (see following subsection). All artificial datasets were generated assuming a 100,000 second exposure time. The systematic error was set to 5%, a large fraction compared to the error from the counting statistics. This is, of course, very different from realistic observations, but a process that relied on counting statistics for the uncertainties would have more heavily weighted softer energies near the spectral peak. We prefer a method that weights energy bins more equally to obtain reasonable matches to the hard X-ray tails. The fitted energy ranged from 3 keV to 50 keV. After the artificial datasets were generated, we chose a sample of model combinations listed in [Table 4.2](#) along with the corresponding fit parameters. The model `tbody` fits a blackbody spectrum with two parameters: a temperature kT (keV) and a normalization parameter that accounts for the emitting area and distance to a source $N_{\text{BB}} = L_{39}/D_{10}^2$, where L_{39} is the source luminosity in units of 10^{39} and D_{10} is the distance to the source in units of 10 kpc. Similarly, the `diskbb` model fits a multi-temperature blackbody accretion disk spectrum (without a colour temperature correction factor) and has two free parameters: the inner disk temperature T_{in} (K) and a normalization parameter defined as $N_{\text{DBB}} = (R_{\text{in}}/D_{10})^2 \cos \theta$ where R_{in} is the apparent inner disk radius in km, D_{10} is the distance to the source in units of 10 kpc, and θ is the disk inclination angle at which $\theta = 0$ is face-on ([Mitsuda et al., 1984](#)).

In addition to the accretion disk models, a hard X-ray component was added to

Table 4.2: Comparison of commonly used spectral fitting models when fit to the MC post-processed grey RMHD snapshot spectra (ULX2.5 and ULX4a) and the MC post-processed multi-group RMHD snapshot spectra (ULX2.5-MG and ULX4a-MG). We only report rough values for each model parameter without errors in an effort to get a sense of the relative spectral shape of the simulated spectra between different model combinations. The accretion disk models used to fit the softer part of the spectrum (`diskbb` and `bbody`) were combined with either the Comptonization model `simpl` or power law model `pow` to fit the hard X-ray spectrum.

Model Component	Parameter	ULX2.5	ULX2.5-MG	ULX4a	ULX4a-MG
<code>bbody</code>	kT	1.13	1.42	1.46	1.39
	N_{BB}	0.23	0.32	0.85	0.99
<code>simpl</code>	Γ_{S}	4.38	2.82	4.34	3.52
	f_{sc}	60%	60%	60%	60%
<code>diskbb</code>	T_{in}	1.71	2.27	2.47	2.04
	N_{DBB}	149.92	56.42	116.21	288.64
<code>simpl</code>	Γ_{S}	4.34	2.88	4.14	3.57
	f_{sc}	41.5%	60%	52.2%	60%
<code>diskbb</code>	T_{in}	2.35	4.30	3.06	3.34
	N_{DBB}	25.39	2.07	37.65	23.36
<code>pow</code>	Γ_{P}	3.36	2.40	3.03	2.80

better characterize the hard X-ray energies. We chose a photon power-law model `pow` which has two free parameters: the photon power law index Γ (such that flux goes as $E^{-\Gamma}$), and a normalization parameter. The other model we chose to characterize the hard X-ray spectrum is the X-ray Comptonization model `simpl` (Steiner et al., 2009). `simpl` is a convolution model that approximately Compton up-scatters a fraction f_{sc} of seed photons from the `body` or `diskbb` models. These up-scattered photons comprise a hard X-ray power law with photon power law index Γ . We assume that photons will only be up-scattered, leaving two free fit parameters, similar to `pow`.

We report these fit results in Table 4.2. We only report a few significant digits without the errors as these are essentially model fits to simulated model data, in contrast to model fits to observed data. Any errors computed here would strongly depend on the chosen systematic and stochastic errors, and are not physically meaningful. We do not report goodness-of-fit for similar reasons.

For the gray snapshots, we find a rather steep power law index for all fits, $\Gamma > 4$ when `simpl` is used. The index is still steep, but somewhat flatter for a pure power law. In the combination of `diskbb+pow`, the `pow` model component dominated the total model fit and could not adequately describe the softer part of the spectrum. This is partly because we are missing a large portion of the soft X-rays from the outer disk in the simulations, but also related to the lack of an absorption model to attenuate the power law at lower energies. In contrast, the multi-group snapshots are much flatter, particularly for ULX2.5-MG. These flatter values are in better agreement with most observed spectra. We also generally find high scattering fractions, indicative of the fact that the power law extends from near the spectral peak. For the model combinations that include `simpl`, we set an upper limit on the scattering fraction of $f_{sc} = 60\%$ as this parameter was not well constrained at higher f_{sc} .

Figure 4.12 shows an example of the three model combination fits to ULX2.5-MG. The ULX2.5-MG spectrum is shown as the black solid line. The total `diskbb+pow` model corresponds to the blue dotted line, which shows the deviation of the fit at the softer end of the spectrum due to the `pow` component dominating the fit. Interestingly, the `body+simpl` model more closely fit the simulated spectra compared to the other two model combinations. For most of the spectra, the `diskbb` component in the `diskbb+simpl` model was slightly broader than the simulated spectrum as seen by the pink solid line in Figure 4.12. This, however, may again be impacted by our deficit of soft X-ray from larger radii.

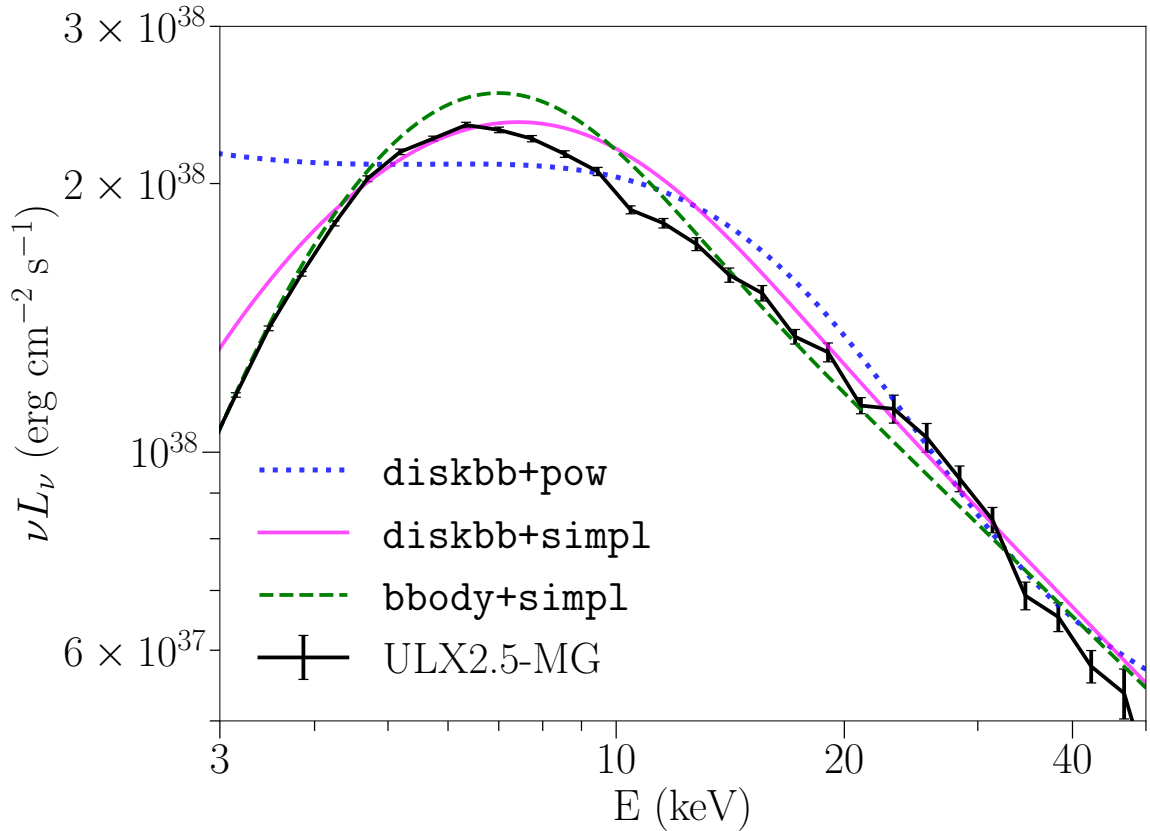


Figure 4.12: Comparison of three X-ray spectral fitting models to the post-processed multi-group RMHD spectrum ULX2.5-MG (shown as the black solid line). The model combinations include two components, one blackbody (bbody) or multi-temperature blackbody accretion disk diskbb model paired with either a power-law (pow) or a hard X-ray Comptonization model (simpl). The total combinations are shown as the blue dotted line for diskbb+pow, pink solid line for diskbb+simpl, and green dashed line for bbody+simpl. Note that the simpl model in XSPEC would be written as simpl×diskbb, since it is a convolution model.

4.4.5 Example analysis: NGC 1313 X-1

NGC 1313 X-1 is a well known ULX ($L_x \sim 10^{40}$ erg s $^{-1}$) located relatively nearby ($D \sim 4.2$ Mpc; Tully et al. 2013). The nature of its compact accretor is not currently known (Walton et al., 2020), although it has been suggested that the large changes observed in its X-ray flux indicate that the system has entered into propeller state consistent with a weakly magnetized neutron star (Middleton et al., 2023) whose relatively weak signal is obscured by being scattered into the wind cone (Mushtukov et al., 2020). Regardless of the nature of its compact object, an interesting feature in the X-ray spectra of this source is the relative stability seen at hard X-ray energies $E \gtrsim 10$ keV, even though variability has been observed at low energies $E \lesssim 10$ keV (Middleton et al., 2015a; Walton et al., 2020; Middleton et al., 2023). In the interest of this analysis, we chose to use NGC 1313 X-1 as an example observation in which to fit our hard X-ray spectral results to.

Table 4.3 shows the best-fit parameters for fits to the combined *XMM-Newton* and *NuSTAR* observations of the ULX NGC 1313 X-1 (Walton et al., 2020). The data for NGC 1313 X-1 were provided by Gúrpide et al. (2021) (see also Middleton et al., 2023). For the *NuSTAR* data we used energies between 3 – 70 keV, and for the *XMM-Newton* data we used energies between 0.3 – 10 keV. In all of the fits, we allow two multiplicative constants to vary freely between the *XMM-Newton* and *NuSTAR* datasets to account for cross-calibration between the different detectors. The FPMA detector constant was set to unity, while the two free constants yielded values $\lesssim 1.38 \pm 0.2$.

We first highlight the “typical model” column which includes two modified disk blackbody components and one hard X-ray model that is commonly used to fit this source (Middleton et al., 2015b; Pinto et al., 2016b; Walton et al., 2020; Gúrpide et al., 2021). To account for the hydrogen column along the line of sight for all model fits in this analysis, we include a neutral absorption component, **TBabs**, adopting the abundances from Wilms et al. (2000) and cross-sections from Verner et al. (1996), and the column N_{H} was left free to vary. See Gúrpide et al. (2021) for some discussion on the variability in the absorption column for NGC 1313 X-1. The two modified disk blackbody components in this model combination are **diskbb** and **diskpbb**, used to model the softer and harder emission components, respectively. The **diskpbb** component includes a free parameter, p , which describes the radial dependence of

Table 4.3: Best-fit parameter values of the simulated spectral model fits to the combined *XMM-Newton* and *NuSTAR* data for NGC 1313 X-1. We show a comparison “Typical model” fit to NGC 1313 X-1, with which to compare with the fits from the four post-processed or simulated snapshot spectral models. The two models ULX4a-MG and ULX2.5-MG include the multi-group implementation, whereas the other snapshots are post-processed from the grey RMHD snapshots. The notation “n/a” indicates that this model parameter was not included in the fit.

Model Component	Parameter	Typical model	ULX2.5	ULX2.5-MG	ULX4a	ULX4a-MG
TBabs	n_H (cm^{-2})	$0.27^{+0.02}_{-0.03}$	$0.26^{+0.02}_{-0.02}$	$0.16^{+0.03}_{-0.02}$	$0.15^{+0.02}_{-0.02}$	$0.19^{+0.02}_{-0.03}$
	T_{in} (keV)	$0.27^{+0.03}_{-0.03}$	$0.27^{+0.01}_{-0.01}$	$0.72^{+0.04}_{-0.03}$	$0.68^{+0.03}_{-0.02}$	$0.72^{+0.03}_{-0.03}$
MC spectrum	N_{DBB}	$11.2^{+7.9}_{-4.3}$	$13.83^{+6.9}_{-6.2}$	$0.30^{+0.07}_{-0.07}$	$0.42^{+0.08}_{-0.7}$	$0.33^{+0.06}_{-0.06}$
	N ($\times 10^{-6}$)	n/a	$100.0^{+1.8}_{-1.8}$	$74.2^{+2.1}_{-2.1}$	$26.5^{+0.6}_{-0.7}$	$22.1^{+0.6}_{-0.5}$
diskpbb	T_{in} (keV)	$3.22^{+0.62}_{-0.46}$	n/a	n/a	n/a	n/a
	p	$0.56^{+0.03}_{-0.02}$	n/a	n/a	n/a	n/a
simpl	N_p ($\times 10^{-4}$)	$4.49^{+5.6}_{-2.5}$	n/a	n/a	n/a	n/a
	Γ	$2.00^{+0.53}_{-2}$	n/a	n/a	n/a	n/a
	f_{sc} (%)	$11.50^{+7.8}_{-6.1}$	n/a	n/a	n/a	n/a
	χ^2/dof	373.60/341	777.71/345	403.37/345	445.72/345	399.05/342

the local disk temperature, $T(r) \propto r^{-p}$. When advection in the disk is considered important, such as in the case of super-Eddington accretion, the p values are typically $p < 0.75$ (Abramowicz et al., 1988). When $p = 0.75$, the model recovers the thin disk `diskbb` solution.

Observations of NGC 1313 X-1 also show emission and absorption lines at energies $E \lesssim 2$ keV that are attributed to the presence of a mildly relativistic disk wind (Middleton et al., 2015a; Pinto et al., 2016b, 2020; Gúrpide et al., 2021). Multiple Gaussian absorption components, `gabs`, are often included to account for some of these atomic features. We limit the `gabs` parameters to $E \leq 2$ keV, line width $\sigma \leq 0.5$ keV, and the line strength was allowed to be positive or negative to represent either emission or absorption. In the typical model, an additional component at moderate to high energies ($\gtrsim 10$ keV) is included to capture the X-ray excess not adequately modeled by the multi-temperature disk components. We apply the same `simpl` convolution model to the `diskpbb` component as in previous works. We set a lower limit on the power law index parameter $\Gamma \geq 2$ as the uncertainties in the data at high energies $E \gtrsim 30$ keV caused `simpl` to fit an unrealistically flat power law in this regime. The typical model is thus written as: `TBabs×gabs×(diskbb+(simpl×diskpbb))`. This provides a reasonably good fit with $\chi^2 = 373.60$ for 341 degrees of freedom, comparable to the best-fits reported in Middleton et al. (2015a); Walton et al. (2020); Gúrpide et al. (2021). One difference in our reproduction of this model is that we only included one `gabs` component with a line energy $E = 1$ keV, line width of $\sigma = 0.01$ keV, and line strength $N_{\text{gabs}} = -0.02$.

In comparison with the typical model, we replaced the `simpl×diskpbb` component with one of the post-processed spectral models denoted in `XSPEC` as:

`TBabs×gabs×(diskbb+MC spectrum)`. The `diskbb` component is included to model the soft X-ray spectrum absent from our spectral models. The spectral models have one fit parameter, $N = (10 \text{ kpc}/D)^2$, where D is the distance to the source. Assuming the distance to NGC 1313 X-1 is $D = 4.25$ Mpc (Tully et al., 2013), this gives a `MC spectrum` model normalization value of $N = 5.67 \times 10^{-6}$. Table 4.3 shows the best-fit `XSPEC` values for four snapshot models: ULX2.5, ULX4a, and their corresponding multi-group runs ULX2.5-MG and ULX4a-MG. We also fit the other two spectral models for ULX1.3 and ULX4b, but for the sake of brevity and lack of a multi-group counterpart for these snapshots we do not include them in Tables 4.2 or 4.3. They provide poor fits to the data. The best-fit for ULX1.3 had a $\chi^2 = 637.87$ for 345




Figure 4.13: Best-fit to the combined *XMM-Newton* (black data points) and *NuSTAR* data (red and blue data points) of the ULX NGC 1313 X-1 using the post-processed spectral model from ULX4a-MG. The top panel shows the spectral fit to the data with individual model components shown for `diskbb` (below 10 keV) and the ULX4a-MG model fitting the rest of the hard X-ray spectrum. The bottom panel shows the fit residuals of the total model (green line) to the data. The best-fit values are collected in Table 4.3.




Figure 4.14: Same as Figure 4.13, but using the ULX2.5-MG spectral model to fit the hard X-ray spectrum.

degrees of freedom, and ULX4b had a $\chi^2 = 432.19$ for 345 degrees of freedom.

All of the spectral model fits included a single **gabs** component except for ULX4a-MG which included two **gabs** components. Most of the fits were insensitive to a second **gabs** component, but the χ^2 for ULX4a improved from $\chi^2 = 433.20$ per 345 degrees of freedom with only one Gaussian absorption component to $\chi^2 = 394.50$ per 342 degrees of freedom with the addition of a second Gaussian component. The two **gabs** components fit lines at 0.34 keV with $\sigma = 0.49$ keV, and a line at 0.67 keV with $\sigma = 0.15$ keV. Generally, modeling the absorption and emission features of this source improves the χ^2 of the fit residuals below 2 keV, but it does not significantly impact the broader continuum fit. Thus we do not attempt model these features in detail as past studies have done (Walton et al., 2016; Pinto et al., 2016b, 2020; Gúrpide et al., 2021; Kosec et al., 2021).

The $\Delta\chi^2$ improves significantly for fits with the multi-group models, as the harder X-ray tails of these spectra better match the observed *NuSTAR* data. We show the two multi-group model combinations in Figures 4.13 and 4.14 for ULX4a-MG and ULX2.5-MG, respectively. The individual model components are shown for the absorbed **diskbb** component below 10 keV, and the spectral model component modeling the higher energies. ULX4a-MG in Figure 4.13 is just slightly steeper than the *NuSTAR* data at $\gtrsim 10$ keV, while ULX2.5-MG in Figure 4.14 are just slightly flatter at $\gtrsim 10$ keV. Both spectral models, however, fit quite well in the 3 – 10 keV range. Even the departures at high energies $\gtrsim 10$ keV are qualitatively good considering the only parameter is the normalization of the simulated spectral component.

4.5 DISCUSSION

4.5.1 Caveats and Sensitivity to Model Assumptions

Our results suggest that RMHD simulations can qualitatively reproduce the observed hard X-ray spectral shape seen in a number of ULX sources as long as the radiative heating/cooling associated with Compton scattering processes are well modeled. Nevertheless, the simulations presented here have only explored a limited range of parameter space and do not yet include all of the relevant physics. Most importantly, the **Athena++** RMHD simulations neglect general relativistic effects such as light bending, relativistic beaming, and relativistic jets (although these simulations do generate radiatively driven disk winds). New GRRMHD simulations are being per-

formed with `Athena++` using direct solution of the radiation transfer solution (White et al., 2023) and we expect that the inclusion of general relativistic effects will have an impact on the accretion flow, disk structure, and associated spectral properties. Post-processed spectra from such GRRMHD simulations will be a focus of future work.

In addition, we only include the inner $25r_g$ of these RMHD simulations in our MC spectral calculations as the simulation is not in steady state at larger radii. Thus, we do not accurately model the soft X-ray spectrum originating beyond $25r_g$. Consequently, we only select photons coming out of a polar funnel angle, θ_f to avoid the impact of photons which would normally interact with the outer disk radii and become trapped in the disk or advected into the black hole. Therefore, we stress that interpretations of these results be limited to the inner regions of the flow.

We also assume that protons and electrons are well coupled and so simulate a single temperature accretion flow $T = 10^7$ K. Some studies have suggested that two temperature accretion flows may become important in areas of low density, such as in the funnel regions (Liska et al., 2022) and may result in softer X-ray spectra (Kinch et al., 2020). We compared the non-relativistic proton-electron relaxation timescale given by Spitzer (1956) and Stepney (1983) assuming a single temperature for both the electrons and protons with the Compton timescale $t_C = (N_e \sigma_T c)^{-1}$ where N_e is the number density of electrons, σ_T is the Thomson cross-section, and c is the speed of light. We found that the relaxation time is much shorter than the Compton timescale and most other dynamical timescales for the temperature and densities in our simulation, except possibly in the very low density, high temperature region near the axis. Hence, our assumption of a single temperature for the protons and electrons seems self-consistent, but the single temperature assumption may need to be revisited in future work, particularly for simulations at lower accretion rates.

4.5.2 Comparison with ULX modeling

When qualitatively fitting X-ray spectral models to the post-processed spectral models, we find that the fits including SIMPL tended to fit the snapshot spectra better than fits with POW. The resulting power law indexes provide a range of values ($\Gamma \sim 2.4 - 4.4$), which are slightly steeper than observed power law tails seen in most observed ULX spectra ($\Gamma \lesssim 2$, Pintore et al., 2017; Dage et al., 2021). The multi-group spectra appear to be in better agreement with the relatively steeper slopes

found for NGC 1313 X-1 ($\Gamma \sim 2.9$) and Holmberg IX X-1 ($\Gamma \sim 2.9 - 3.4$) (Gúrpidé et al., 2021). The power law indexes for the multi-group spectra result in slightly flatter hard X-ray power laws than their grey snapshot counterparts, comparable with the spectral shapes observed in many ULXs.

Fits to the *XMM-Newton* and *NuSTAR* data of NGC 1313 X-1 with the post-processed spectral models qualitatively reproduce the hard X-ray part of the spectrum, although the funnel luminosities for these spectra are at least an order of magnitude less luminous ($L_f = 1.3 \times 10^{39}$ erg s $^{-1}$ for ULX4a-MG, and $L_f = 4.7 \times 10^{38}$ erg s $^{-1}$ for ULX2.5-MG) than that of NGC 1313 X-1 ($L_x \sim 10^{40}$ erg s $^{-1}$). The implied distances from the spectral models also give much closer distances, $D \sim 1 - 2.13$ Mpc, compared to the distance to NGC 1313 X-1, $D = 4.2$ Mpc (Tully et al., 2013). This motivates future work with simulations at higher Eddington ratios and a range of black hole mass, which might better match the observed sources at the known distance. Nevertheless, it is remarkable that a first principles calculation with only the normalization as a free parameter provide best fitting χ^2 values that are quantitatively competitive with commonly used phenomenological models.

4.5.3 Comparison with previous work

Previous simulation have explored a range of accretion rates and masses, using a variety of setups both with and without general relativistic effects included, finding radiative efficiencies that are both relatively large ($\eta \sim 5\%$, e.g. Jiang et al., 2014) or small ($\eta \lesssim 1\%$, e.g. Sądowski et al., 2014). The radiative efficiencies inferred directly from the gray RMHD simulations are typically a few percent, somewhat less than expected for thin accretion disks but not inconsistent with expectations for modestly super-Eddington accretion rates. The simulations emerging in the funnel region computed with MC post-processing are modestly lower than from the RMHD simulations for the gray snapshots, with $\eta_f = 1.13 - 2.56\%$. In contrast, the snapshots produced from the multi-group calculations have slightly larger efficiencies ($\eta_f = 1.92\%$ for ULX2.5-MG, and $\eta_f = 3.34\%$ for ULX4a-MG), in better agreement with the luminosities directly inferred from RMHD simulations.

The Monte Carlo radiation transfer calculations performed in this work are similar to other post-processing codes, particularly those that use MC method such as GRMONTY (Dolence et al., 2009), Pandurata (Schnittman et al., 2013), or RAIKOU (Kawashima et al., 2021) that model Compton scattering and include general rela-

tivistic effects. The `HERIOIC` code (Narayan et al., 2017), which uses a combination of short and long characteristics instead of MC provides similar capabilities. Although the `Athena++` module used here also supports general relativistic transfer, we treat the radiation transfer in Minkowski spacetime to be consistent with the non-relativistic simulations that generate the snapshots. In contrast to Kinch et al. (2019), we do not currently perform any ionization calculations that would investigate atomic transitions. We also do not use any integrated ray tracing algorithms that would integrate back along the photon path in our post-processing. Although to create the images in Figure 4.9, we extrapolate the photons escaping the MC domain out to a distant observer assuming flat spacetime.

Narayan et al. (2017) used `HERIOIC` to post-process simulations from the GRRMHD code `KORAL` (Sądowski et al., 2013; Sądowski et al., 2014; Sądowski & Narayan, 2015; Sądowski & Narayan, 2016), which was used to simulate a broad range of super-Eddington accretion rates onto a $10M_{\odot}$ black hole. They faced the same issue in that their GRMHD simulations only reached inflow equilibrium out to a finite radius. Instead of truncating the disk as we chose to do, they instead extrapolated the flow to larger radii using self-similar approximations. This allowed them to explore the softer X-ray emission and angular dependence, but with the caveat that the outer regions of the calculation were not simulated directly. We find that our spectra are more qualitatively consistent with their results when the gas temperatures in the `HERIOIC` calculation were fixed to the values from `KORAL` (see green curves in Narayan et al., 2017 Fig. 4), but not consistent with the spectra after the radiation field and temperatures were self-consistently solved (red curves in the same figure). The results from `HERIOIC` show that their spectra become much softer after the temperature iteration, whilst our results suggest that a more self-consistent treatment of Compton cooling yields higher temperatures and harder spectra. The origin of the difference is not clear to us, but we note that the `KORAL` simulations use a photon number conservation scheme that is different from what we imply in our gray simulations, so that may explain the differences.

4.6 SUMMARY

We present Monte Carlo post-processed spectral calculations of super-Eddington accretion onto a stellar-mass black hole from the `Athena++` RMHD simulation snapshots. Our calculations suffer from two primary deficiencies. We only achieve inflow

equilibrium out to about $25r_g$, which led us to truncate our spectral calculations at this radius. Hence, the soft X-rays that come from the outer disk are absent. If we instead include the emission from the outer disk, it is significantly overestimated due the cooling of our torus initial conditions. Hence, we focus on the hard X-ray spectrum in this work. These simulations also assume the intensities follow a blackbody spectrum for the purposes of computing Compton cooling and mean opacities. Although this assumption is good for the optically thick disk, we find that using the blackbody assumption to estimate the average photon energy in the Compton cooling term is a poor approximation in the funnel regions where Comptonized electrons dominate the cooling. This leads to an underestimate for the temperatures in the funnel for the gray RMHD simulations. The underestimated temperatures produced spectra that were much softer and led to radiation energy densities above the disk being overestimated. We addressed this underestimate of the temperature by restarting the gray opacities simulations with a multi-group approach (Jiang, 2022) that treats Compton scattering with a Kompaneets like source term. This produced simulation snapshots with higher temperatures in the spectral forming regions above the disk, leading to harder X-ray spectra in better agreement with observed ULX spectra.

We used phenomenological models to fit our Monte Carlo spectra. In most of the two-component (soft X-ray + hard X-ray) models, the hard X-ray component was more accurately described with the SIMPL model compared to the power law POW model, and yielded hard X-ray power law slopes ranging from $\Gamma \sim 2 - 4$ for spectra computed with gray RMHD snapshots. The multi-group snapshot spectra tended to be fit with flatter slopes, with $\Gamma \sim 2 - 3$, comparable to the power laws observed in NGC 1313 X-1 and Holmberg IX X-1.

Finally, we generated an XSPEC table model and directly fit our MC spectra to combined *XMM-Newton* and *NuSTAR* observations of the ULX NGC 1313 X-1. Despite only having one free parameter (the normalization), we find a good fit, which is competitive with the phenomenological models that are commonly used. Close inspection shows that the MC spectra provide a good fit at soft to moderately hard energies $E \lesssim 10$ keV, but are either just slightly too steep in the case of ULX4a-MG or too flat in the case of ULX2.5-MG to exactly describe the hard X-ray power law tail $E \gtrsim 10$ keV. Nor are the best-fit normalizations consistent with the known distance to NGC 1313 X-1, consistent with the lower luminosity inferred for the MC spectra. Given the single black hole mass and the relatively narrow range of accretion rates

explored along with caveats like the neglect of general relativistic effect, this work nonetheless provides a promising direction for super-Eddington ULX accretion simulations, as these post-processed spectral models are close to describing the observed spectrum of NGC 1313 X-1. Simulations with the new GRRMHD implementation of the **Athena++** code (White et al., 2023) are now exploring a range of masses and accretion rates and post-processed spectra from these simulation will be presented in a future work.

This work was supported by NASA Astrophysics Theory Program grant 80NSSC18K1018. BSM also thanks the Jefferson Scholars Foundation Graduate Fellowship for supporting this work. Part of this work was performed using resources provided by the Cambridge Service for Data Driven Discovery (CSD3) operated by the University of Cambridge Research Computing Service (www.csd3.cam.ac.uk), provided by Dell EMC and Intel using Tier-2 funding from the Engineering and Physical Sciences Research Council (capital grant EP/T022159/1), and DiRAC funding from the Science and Technology Facilities Council (www.dirac.ac.uk). Resources supporting this work were also provided by the NASA High-End Computing (HEC) program through the NASA Advanced Supercomputing (NAS) Division at AMES Research Center. The Center for Computational Astrophysics at the Flatiron Institute is supported by the Simons Foundation.

CHAPTER 5

SUMMARY AND FUTURE DIRECTIONS

This dissertation has focused on the intersection of X-ray spectral observations, spectral fitting models, and the numerical RMHD simulations of critical accretion onto a black hole. We have examined the spin estimate for the black hole in GRS 1915+105 using several spectral fitting models in Chapter 2. The spectral models used were a combination of phenomenological models (e.g. `nthcomp`, `pow`), and more sophisticated models such as `kerrbb` and `bhspec`. The `bhspec` model is notable in that it uses TLUSTY stellar atmosphere calculations of disk annuli to self-consistently compute the radiative transfer and vertical structure of a [Shakura & Sunyaev \(1973\)](#) thin accretion disk.

However in the case of thicker accretion flows, such as that of ULXs and other sources which exhibit super-Eddington accretors, the models mentioned are not able to accurately describe the disk structure. This dissertation presents spectral calculations from detailed first-principle RMHD simulations and shows how the simulated models compare with the phenomenological models discussed above, as well as direct fits to the ULX NGC 1313 X-1 in Chapter 4. These results represent a strong direction forward for generating self-consistent spectral models using numerical simulations of black hole accretion, as opposed to reliance on phenomenological descriptors of accretion disk spectra. Even with the caveats in the simulations (as discussed in Chapter 4), it is promising that the simulated spectra can already describe most of the observed spectrum for NGC 1313 X-1, with the hard X-ray spectral profiles consistent with the flatter power laws seen in such sources.

We briefly summarize the main conclusions of Chapters 2, 3, and 4, and discuss future directions of this dissertation.

Chapter 2

This chapter was based on a publication in the *Astrophysical Journal* with the following citation: B. S. Mills, S. W. Davis, and M. J. Middleton. In this work, we revisited the spin estimate for the black hole in the BHXB GRS 1915+105, following two prior papers which gave inconsistent spin estimates (Middleton et al., 2006; McClintock et al., 2006). In light of new constraints on the mass, distance, and inclination for the source (Reid et al., 2014), we examined the effect of these new constraints on the spin estimates for the two papers. Below we summarize the main conclusions of this chapter.

- **Higher spins.** We found that the new constraints for the GRS 1915+105 system provided by Reid et al. (2014) resulted in higher spin estimates for both of the previous papers. This pushed the spin from $a_* \simeq 0.72$ to $a_* \simeq 0.86$ for the Middleton et al. (2006) dataset, and the spin from McClintock et al. (2006) was already at maximal spin, $a_* \simeq 0.99$.
- **Explaining the McClintock et al. maximal spin.** We were unable to obtain a good fit for the McClintock et al. (2006) dataset with the new Reid et al. (2014) constraints. However, good fits could be obtained if the distance to GRS 1915+105 is $D \geq 10$ kpc, in which case the spin would remain near maximal ($a_* \simeq 0.99$).
- **A tale of two spins.** The discrepancy between the spin estimates in the two previous papers is largely due to a combination of the differences in their data selection criteria, the extreme variability of GRS 1915+105, and the limitations of the continuum-fitting method.
- **Best-fit spin for GRS 1915+105.** The best-fit spin for GRS 1915+105 using the Middleton et al. (2006) data and the Reid et al. (2014) constraints is moderately high ($a_* \simeq 0.86$), however the uncertainties in parameter space still yield a wide range of spins for GRS 1915+105 ($0.4 \leq a_* \leq 0.99$).
- **Tighter constraints.** Stronger constraints on the mass, distance, and inclination of GRS 1915+105 are needed in order to provide a tighter constraint on the spin of the black hole. Given the current range of uncertainties in the parameter space, the best-fit spins to the Middleton et al. (2006) dataset range between $0.4 \leq a_* \leq 0.99$.

Chapter 3

In this chapter we presented an overview of the MC radiation transfer module (Davis et al., 2005, Davis et al. in preparation) in *Athena++*. This chapter is adapted from an article in preparation with the following (tentative) author list: S. W. Davis, B. S. Mills, et al. We described how the module adopts the mesh structures of the *Athena++* code, how photons are initialized, moved, and output, some relevant code tests of the module, and how the module is used to generate spectra. Some key points of this chapter are as follows:

- **Time-dependent or post-processed spectra.** The MC module is compatible with the *Athena++* mesh structure. It is also able to be run simultaneously with the *Athena++* code to produce time-dependent spectra, or in our case used as a post-processing tool to generate spectra after the RMHD simulations have completed.
- **Statistical weights.** The MC module utilizes a photon weighting scheme where photon weights are allowed to vary. In contrast, many other MC codes implement an equal weighting scheme that is nominally variance reducing, but in practice performs less well when scattering optical depths are large and is more difficult to implement on a distributed mesh such as those used in simulations.
- **Code tests.** We presented several quantitative test of the MC module, including for an isothermal atmosphere with pure absorption and polarized Thomson scattering, a uniform periodic box test with boosted velocities, and a uniform sphere test with unpolarized Compton scattering. The convergence of the MC module shows that the module is robust and converges to the analytic solutions in the test cases presented here.

Chapter 4

This chapter was adapted from a submitted article to the *Astrophysical Journal* with the following author list: B. S. Mills, S. W. Davis, Y. F. Jiang, and M. J. Middleton. We present MC post-processed spectral calculations of *Athena++* RMHD simulation snapshots of mild super-Eddington accretion onto a stellar mass black hole.

- **Multi-group approach.** The frequency-integrated (gray) RMHD simulatons underestimated the temperatures in the funnel region above and below the black

hole. By implementing a multi-group approach (Jiang, 2022), the resulting spectra are harder and more closely resemble observed ULX spectra.

- **Comparison to phenomenological X-ray spectral models.** In comparison to other spectral models, our simulated spectral models are derived from first-principles. Despite the limitations and caveats of the simulated models, they appear to produce competitive fits to the spectrum of NGC 1313 X-1 compared to phenomenological models.
- **Comparison to the spectrum of NGC 1313 X-1.** We fit the simulated MC spectral models directly to the observed spectrum of NGC 1313 X-1, and the models are able to qualitatively reproduce most of the data. At harder energies $E \gtrsim 10$ keV, the spectra are either slightly too steep or too flat compared to the data, but for a first-principles model the fits are nonetheless promising.

Exciting GRRMHD simulations with `Athena++` are underway. The MC module described in Chapter 3 will be used to post-process spectra for the GRRMHD simulations of various black hole masses and accretion rates. Although this dissertation primarily focused on post-processing spectra of super-Eddington accretion onto a stellar mass black hole, spectra could also be made for simulations of sub-Eddington accretion onto a stellar mass black hole (applicable for BHXBs) as well as accretion onto supermassive black holes (AGN). The MC module is particularly useful for `Athena++` users and can be used to calculate spectra for a wide range of astrophysical simulations (e.g. tidal disruption events, cataclysmic variables, neutron star X-ray bursts, etc.) Current work is also being done to expand the capabilities of the MC module. This includes accelerating the MC module via a modified random walk implementation (McClellan et al., 2022, Davis et al. in preparation), as well as a treatment for dust-scattering.

REFERENCES

- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016, *Physical Review X*, 6, 041015
- Abramowicz, M. A., Czerny, B., Lasota, J. P., & Szuszkiewicz, E. 1988, *The Astrophysical Journal*, 332, 646
- Abramowicz, M. A., Jaroszyński, M., Kato, S., et al. 2010, *Astronomy and Astrophysics*, 521, A15
- Agol, E., & Krolik, J. H. 2000, *The Astrophysical Journal*, 528, 161
- Altamirano, D., Belloni, T., Linares, M., et al. 2011, *The Astrophysical Journal*, 742, L17
- Anders, E., & Ebihara, M. 1982, *Geochimica et Cosmochimica Acta*, 46, 2363
- Anders, E., & Grevesse, N. 1989, *Geochimica et Cosmochimica Acta*, 53, 197
- Arca Sedda, M., Mapelli, M., Benacquista, M., & Spera, M. 2023, *Monthly Notices of the Royal Astronomical Society*, 520, 5259
- Arnaud, K. A. 1996, in *Astronomical Society of the Pacific Conference Series*, Vol. 101, *Astronomical Data Analysis Software and Systems V*, ed. G. H. Jacoby & J. Barnes, 17
- Asahina, Y., Takahashi, H. R., & Ohsuga, K. 2020, *The Astrophysical Journal*, 901, 96
- Bachetti, M., Rana, V., Walton, D. J., et al. 2013, *The Astrophysical Journal*, 778, 163
- Bachetti, M., Harrison, F. A., Walton, D. J., et al. 2014, *Nature*, 514, 202
- Baganoff, F. K., Bautz, M. W., Brandt, W. N., et al. 2001, *Nature*, 413, 45
- Bailes, M., Berger, B. K., Brady, P. R., et al. 2021, *Nature Reviews Physics*, 3, 344
- Balbus, S. A., & Hawley, J. F. 1991, *The Astrophysical Journal*, 376, 214
- Bañados, E., Venemans, B. P., Mazzucchelli, C., et al. 2018, *Nature*, 553, 473

- Bardeen, J. M., & Petterson, J. A. 1975, *The Astrophysical Journal Letters*, 195, L65
- Basko, M. M. 1978, *The Astrophysical Journal*, 223, 268
- Bavera, S. S., Fragos, T., Zevin, M., et al. 2021, *Astronomy and Astrophysics*, 647, A153
- Becker, P. A. 2003, *Monthly Notices of the Royal Astronomical Society*, 343, 215
- Belloni, T., Klein-Wolt, M., Méndez, M., van der Klis, M., & van Paradijs, J. 2000, *Astronomy and Astrophysics*, 355, 271
- Blaes, O. M., Davis, S. W., Hirose, S., Krolik, J. H., & Stone, J. M. 2006, *The Astrophysical Journal*, 645, 1402
- Blandford, R., Meier, D., & Readhead, A. 2019, *Annual Review of Astronomy and Astrophysics*, 57, 467
- Blandford, R. D., & Znajek, R. L. 1977, *Monthly Notices of the Royal Astronomical Society*, 179, 433
- Blum, J. L., Miller, J. M., Fabian, A. C., et al. 2009, *The Astrophysical Journal*, 706, 60
- Brightman, M., Harrison, F. A., Barret, D., et al. 2016, *The Astrophysical Journal*, 829, 28
- Chandrasekhar, S. 1960, *Radiative transfer*
- Chapuis, C., & Corbel, S. 2004, *Astronomy and Astrophysics*, 414, 659
- Chaty, S., Mirabel, I. F., Duc, P. A., Wink, J. E., & Rodriguez, L. F. 1996, *Astronomy and Astrophysics*, 310, 825
- Coughenour, B. M., Tomsick, J. A., Mastroserio, G., et al. 2023, arXiv e-prints, arXiv:2303.13718
- Dage, K. C., Vowell, N., Thygesen, E., et al. 2021, *Monthly Notices of the Royal Astronomical Society*, 508, 4008
- Davis, S. W., Blaes, O. M., Hirose, S., & Krolik, J. H. 2009, *The Astrophysical Journal*, 703, 569
- Davis, S. W., Blaes, O. M., Hubeny, I., & Turner, N. J. 2005, *The Astrophysical Journal*, 621, 372
- Davis, S. W., & El-Abd, S. 2019, *The Astrophysical Journal*, 874, 23
- Davis, S. W., & Hubeny, I. 2006, *The Astrophysical Journal Supplement Series*, 164, 530

- Davis, S. W., Stone, J. M., & Jiang, Y.-F. 2012, *The Astrophysical Journal Supplement Series*, 199, 9
- Davis, S. W., Stone, J. M., & Jiang, Y.-F. 2012, *The Astrophysical Journal Supplement Series*, 199, 9
- Dolence, J. C., Gammie, C. F., Mościbrodzka, M., & Leung, P. K. 2009, *The Astrophysical Journal Supplement Series*, 184, 387
- Done, C., & Davis, S. W. 2008, *The Astrophysical Journal*, 683, 389
- Done, C., Gierliński, M., & Kubota, A. 2007, *The Astronomy and Astrophysics Review*, 15, 1
- Doronin, A., & Meglinski, I. 2011, *Biomed. Opt. Express*, 2, 2461
- Draghis, P. A., Balakrishnan, M., Miller, J. M., et al. 2023, arXiv e-prints, arXiv:2303.04164
- Dunn, R. J. H., Fender, R. P., Körding, E. G., Belloni, T., & Merloni, A. 2011, *Monthly Notices of the Royal Astronomical Society*, 411, 337
- Earnshaw, H. M. 2016, *Astronomische Nachrichten*, 337, 448
- Ebisawa, K. 1991, PhD thesis, Institute of Space and Astronautical Science/Japan Aerospace Exploration Agency
- Ebisawa, K., Ogawa, M., Aoki, T., et al. 1994, *Publications of the Astronomical Society of Japan*, 46, 375
- Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019, *The Astrophysical Journal Letters*, 875, L1
- . 2022, *The Astrophysical Journal Letters*, 930, L14
- Fabian, A. C., Rees, M. J., Stella, L., & White, N. E. 1989, *Monthly Notices of the Royal Astronomical Society*, 238, 729
- Farrell, S. A., Webb, N. A., Barret, D., Godet, O., & Rodrigues, J. M. 2009, *Nature*, 460, 73
- Fender, R. P., Garrington, S. T., McKay, D. J., et al. 1999, *Monthly Notices of the Royal Astronomical Society*, 304, 865
- Feng, Y., Steiner, J. F., Ramirez, S. U., & Gou, L. 2023, *Monthly Notices of the Royal Astronomical Society*, 520, 5803
- Fragile, P. C., Mathews, G. J., & Wilson, J. R. 2001, *The Astrophysical Journal*, 553, 955

- Fragile, P. C., Olejar, A., & Anninos, P. 2014, *The Astrophysical Journal*, 796, 22
- Fragos, T., & McClintock, J. E. 2015, *The Astrophysical Journal*, 800, 17
- Fuller, J., & Ma, L. 2019, *The Astrophysical Journal Letters*, 881, L1
- Gallo, L. C., Wilkins, D. R., Bonson, K., et al. 2015, *Monthly Notices of the Royal Astronomical Society*, 446, 633
- Gammie, C. F. 1999, *The Astrophysical Journal Letters*, 522, L57
- Gammie, C. F., Shapiro, S. L., & McKinney, J. C. 2004, *The Astrophysical Journal*, 602, 312
- García, J. A., Sokolova-Lapa, E., Dauser, T., et al. 2020, *The Astrophysical Journal*, 897, 67
- Genzel, R., Schödel, R., Ott, T., et al. 2003, *Nature*, 425, 934
- Gerosa, D., & Berti, E. 2017, *Physical Review D*, 95, 124046
- Ghez, A. M., Wright, S. A., Matthews, K., et al. 2004, *The Astrophysical Journal Letters*, 601, L159
- Gierliński, M., & Done, C. 2004, *Monthly Notices of the Royal Astronomical Society*, 347, 885
- Gierliński, M., Zdziarski, A. A., Poutanen, J., et al. 1999, *Monthly Notices of the Royal Astronomical Society*, 309, 496
- Gladstone, J. C., Roberts, T. P., & Done, C. 2009, *Monthly Notices of the Royal Astronomical Society*, 397, 1836
- González, M., Audit, E., & Huynh, P. 2007, *A&A*, 464, 429
- Greene, J. E., Strader, J., & Ho, L. C. 2020, *Annual Review of Astronomy and Astrophysics*, 58, 257
- Greiner, J., Morgan, E. H., & Remillard, R. A. 1996, *The Astrophysical Journal Letters*, 473, L107
- Gúrpide, A., Godet, O., Koliopanos, F., Webb, N., & Olive, J. F. 2021, *Astronomy and Astrophysics*, 649, A104
- Hanawa, T. 1989, *The Astrophysical Journal*, 341, 948
- Howell, L. H., & Greenough, J. A. 2003, *Journal of Computational Physics*, 184, 53
- Huang, J., Jiang, Y.-F., Feng, H., et al. 2023, arXiv e-prints, arXiv:2301.12679

- Imazato, F., Sasada, M., Uemura, M., et al. 2021, *The Astrophysical Journal*, 916, 114
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *The Astrophysical Journal*, 873, 111
- Jiang, Y.-F. 2021, *The Astrophysical Journal Supplement Series*, 253, 49
- Jiang, Y.-F. 2022, *The Astrophysical Journal Supplement Series*, 263, 4
- Jiang, Y.-F., Stone, J. M., & Davis, S. W. 2014, *The Astrophysical Journal*, 796, 106
- . 2019, *The Astrophysical Journal*, 880, 67
- Kawashima, T., Ohsuga, K., & Takahashi, H. R. 2021, arXiv e-prints, arXiv:2108.05131
- Kinch, B. E., Noble, S. C., Schnittman, J. D., & Krolik, J. H. 2020, *The Astrophysical Journal*, 904, 117
- Kinch, B. E., Schnittman, J. D., Kallman, T. R., & Krolik, J. H. 2019, *The Astrophysical Journal*, 873, 71
- King, A., Lasota, J.-P., & Middleton, M. 2023, *New Astronomy Reviews*, 96, 101672
- King, A. R., & Kolb, U. 1999, *Monthly Notices of the Royal Astronomical Society*, 305, 654
- Kosec, P., Pinto, C., Reynolds, C. S., et al. 2021, *Monthly Notices of the Royal Astronomical Society*, 508, 3569
- Krawczynski, H., & Beheshtipour, B. 2022, *The Astrophysical Journal*, 934, 4
- Krolik, J. H., & Hawley, J. F. 2002, *The Astrophysical Journal*, 573, 754
- Krumholz, M. R., Klein, R. I., McKee, C. F., & Bolstad, J. 2007, *The Astrophysical Journal*, 667, 626
- Kubota, A., Makishima, K., & Done, C. 2004, *Progress of Theoretical Physics Supplement*, 155, 19
- Kulkarni, A. K., Penna, R. F., Shcherbakov, R. V., et al. 2011, *Monthly Notices of the Royal Astronomical Society*, 414, 1183
- Lee, J. C., Reynolds, C. S., Remillard, R., et al. 2002, *The Astrophysical Journal*, 567, 1102
- Levermore, C. D. 1984, *Journal of Quantitative Spectroscopy and Radiative Transfer*, 31, 149
- Li, L.-X., Zimmerman, E. R., Narayan, R., & McClintock, J. E. 2005, *The Astrophysical Journal Supplement Series*, 157, 335

- Liska, M., Hesp, C., Tchekhovskoy, A., et al. 2018, *Monthly Notices of the Royal Astronomical Society*, 474, L81
- Liska, M. T. P., Musoke, G., Tchekhovskoy, A., Porth, O., & Beloborodov, A. M. 2022, *The Astrophysical Journal Letters*, 935, L1
- Lousto, C. O., Nakano, H., Zlochower, Y., & Campanelli, M. 2010, *Phys. Rev. D*, 81, 084023
- Maccarone, T. J. 2002, *Monthly Notices of the Royal Astronomical Society*, 336, 1371
- McClellan, B. C., Davis, S. W., & Arras, P. 2022, *The Astrophysical Journal*, 934, 37
- McClintock, J. E., & Remillard, R. A. 2006, *Black hole binaries*, Vol. 39, 157–213
- McClintock, J. E., Shafee, R., Narayan, R., et al. 2006, *The Astrophysical Journal*, 652, 518
- McClintock, J. E., Narayan, R., Davis, S. W., et al. 2011, *Classical and Quantum Gravity*, 28, 114009
- McKinney, J. C., Tchekhovskoy, Sądowski, A., & Narayan, R. 2014, *Monthly Notices of the Royal Astronomical Society*, 441, 3177
- Menon, S. H., Federrath, C., Krumholz, M. R., et al. 2022, *Monthly Notices of the Royal Astronomical Society*, 512, 401
- Mezcua, M., Roberts, T. P., Sutton, A. D., & Lobanov, A. P. 2013, *Monthly Notices of the Royal Astronomical Society*, 436, 3128
- Middleton, M. 2016, *Astrophysics and Space Science Library*, 99–151
- Middleton, M., Done, C., Gierliński, M., & Davis, S. W. 2006, *Monthly Notices of the Royal Astronomical Society*, 373, 1004
- Middleton, M., Gúrpide, A., & Walton, D. J. 2023, *Monthly Notices of the Royal Astronomical Society*, 519, 2224
- Middleton, M. J., Fragile, P. C., Ingram, A., & Roberts, T. P. 2019, *Monthly Notices of the Royal Astronomical Society*, 489, 282
- Middleton, M. J., Heil, L., Pintore, F., Walton, D. J., & Roberts, T. P. 2015a, *Monthly Notices of the Royal Astronomical Society*, 447, 3243
- Middleton, M. J., Walton, D. J., Fabian, A., et al. 2015b, *Monthly Notices of the Royal Astronomical Society*, 454, 3134
- Middleton, M. J., Walton, D. J., Roberts, T. P., & Heil, L. 2014, *Monthly Notices of the Royal Astronomical Society*, 438, L51

- Middleton, M. J., Fragile, P. C., Bachetti, M., et al. 2018, *Monthly Notices of the Royal Astronomical Society*, 475, 154
- Mihalas, D., & Mihalas, B. W. 1984, *Foundations of radiation hydrodynamics*
- Mikusincova, R., Dovciak, M., Bursa, M., et al. 2023, *Monthly Notices of the Royal Astronomical Society*, 519, 6138
- Miller, J. M., Fabian, A. C., & Miller, M. C. 2004, *The Astrophysical Journal*, 614, L117
- Miller, J. M., Parker, M. L., Fuerst, F., et al. 2013, *The Astrophysical Journal Letters*, 775, L45
- Miller, J. M., Tomsick, J. A., Bachetti, M., et al. 2015, *The Astrophysical Journal Letters*, 799, L6
- Miller-Jones, J. C. A., Tetarenko, A. J., Sivakoff, G. R., et al. 2019, *Nature*, 569, 374
- Mills, B. S., Davis, S. W., & Middleton, M. J. 2021, *The Astrophysical Journal*, 914, 6
- Mitsuda, K., Inoue, H., Koyama, K., et al. 1984, *Publications of the Astronomical Society of Japan*, 36, 741
- Moens, N., Sundqvist, J. O., El Mellah, I., et al. 2022, *A&A*, 657, A81
- Mondal, S., Rózańska, A., De Marco, B., & Markowitz, A. 2021, *Monthly Notices of the Royal Astronomical Society: Letters*, 505, L106
- Morganti, R. 2017, *Frontiers in Astronomy and Space Sciences*, 4, doi:10.3389/fspas.2017.00042
- Motta, S. E., Belloni, T., Stella, L., et al. 2022, *Monthly Notices of the Royal Astronomical Society*, 517, 1469
- Mushtukov, A. A., Portegies Zwart, S., Tsygankov, S. S., Nagirner, D. I., & Poutanen, J. 2020, *Monthly Notices of the Royal Astronomical Society*, 501, 2424
- Narayan, R., Sądowski, A., & Soria, R. 2017, *Monthly Notices of the Royal Astronomical Society*, 469, 2997
- Narayan, R., Zhu, Y., Psaltis, D., & Sądowski, A. 2016, *Monthly Notices of the Royal Astronomical Society*, 457, 608
- Novikov, I. D., & Thorne, K. S. 1973, in *Black Holes (Les Astres Occlus)*, 343–450
- Ohsuga, K., & Mineshige, S. 2011, *The Astrophysical Journal*, 736, 2

- Ohsuga, K., Mori, M., Nakamoto, T., & Mineshige, S. 2005, *The Astrophysical Journal*, 628, 368
- Oskinova, L. M., Bik, A., Mas-Hesse, J. M., et al. 2019, *Astronomy and Astrophysics*, 627, A63
- Paczynski, B., & Wiita, P. J. 1980, *Astronomy and Astrophysics*, 88, 23
- Peirano, V., Méndez, M., García, F., & Belloni, T. 2023, *Monthly Notices of the Royal Astronomical Society*, 519, 1336
- Penrose, R. 1978, in *Theoretical Principles in Astrophysics and Relativity*, ed. N. R. Lebovitz, 217–243
- . 1999, *Journal of Astrophysics and Astronomy*, 20, 233
- Pinto, C., Middleton, M. J., & Fabian, A. C. 2016a, *Nature*, 533, 64
- . 2016b, *Nature*, 533, 64
- Pinto, C., & Walton, D. J. 2023, arXiv e-prints, arXiv:2302.00006
- Pinto, C., Walton, D. J., Kara, E., et al. 2020, *Monthly Notices of the Royal Astronomical Society*, 492, 4646
- Pintore, F., Zampieri, L., Stella, L., et al. 2017, *The Astrophysical Journal*, 836, 113
- Poutanen, J., Lipunova, G., Fabrika, S., Butkevich, A. G., & Abolmasov, P. 2007, *Monthly Notices of the Royal Astronomical Society*, 377, 1187
- Pozdnyakov, L. A., Sobol, I. M., & Syunyaev, R. A. 1983, *Soviet Scientific Reviews, Section E: Astrophysics and Space Physics Reviews*, 2, 189
- Pringle, J. E. 1981, *Annual Review of Astronomy and Astrophysics*, 19, 137
- Reid, M. J., McClintock, J. E., Steiner, J. F., et al. 2014, *The Astrophysical Journal*, 796, 2
- Remillard, R. A., & McClintock, J. E. 2006, *Annual Review of Astronomy and Astrophysics*, 44, 49
- Reynolds, C. S. 2021, *Annual Review of Astronomy and Astrophysics*, 59, 117
- Rizzuto, F. P., Naab, T., Rantala, A., et al. 2023, *Monthly Notices of the Royal Astronomical Society*, 521, 2930
- Rodriguez, C. L., Amaro-Seoane, P., Chatterjee, S., & Rasio, F. A. 2018, *Physical Review Letters*, 120, 151101
- Rybicki, G. B., & Lightman, A. P. 1979, *Radiative processes in astrophysics*

- Sahu, K. C., Anderson, J., Casertano, S., et al. 2022, *The Astrophysical Journal*, 933, 83
- Salvesen, G., & Miller, J. M. 2020, *Monthly Notices of the Royal Astronomical Society*, arXiv:2010.11948
- . 2021, *Monthly Notices of the Royal Astronomical Society*, 500, 3640
- Sana, H., de Mink, S. E., de Koter, A., et al. 2012, *Science*, 337, 444
- Sarrut, D., Bardiès, M., Bousson, N., et al. 2014, *Medical Physics*, 41, 064301
- Schnittman, J. D., & Krolik, J. H. 2013, *The Astrophysical Journal*, 777, 11
- Schnittman, J. D., Krolik, J. H., & Noble, S. C. 2013, *The Astrophysical Journal*, 769, 156
- . 2016, *The Astrophysical Journal*, 819, 48
- Serafinelli, R., Braitto, V., Reeves, J. N., et al. 2023, arXiv e-prints, arXiv:2301.13223
- Shafee, R., McClintock, J. E., Narayan, R., et al. 2006, *The Astrophysical Journal Letters*, 636, L113
- Shakura, N. I., & Sunyaev, R. A. 1973, *Astronomy and Astrophysics*, 500, 33
- Sądowski, A., & Narayan, R. 2016, *Monthly Notices of the Royal Astronomical Society*, 456, 3929
- Sądowski, A., Narayan, R., Tchekhovskoy, A., et al. 2015, *Monthly Notices of the Royal Astronomical Society*, 447, 49
- Sądowski, A., Narayan, R., Tchekhovskoy, A., & Zhu, Y. 2013, *Monthly Notices of the Royal Astronomical Society*, 429, 3533
- Skinner, G. K., Bedford, D. K., Elsner, R. F., et al. 1982, *Nature*, 297, 568
- Skinner, M. A., & Ostriker, E. C. 2013, *The Astrophysical Journal Supplement Series*, 206, 21
- Socrates, A., & Davis, S. W. 2006, *The Astrophysical Journal*, 651, 1049
- Spitzer, L. 1956, *Physics of Fully Ionized Gases*
- Sreehari, H., Nandi, A., Das, S., et al. 2020, *Monthly Notices of the Royal Astronomical Society*, 499, 5891
- Steeghs, D., McClintock, J. E., Parsons, S. G., et al. 2013, *The Astrophysical Journal*, 768, 185

- Steiner, J. F., Narayan, R., McClintock, J. E., & Ebisawa, K. 2009, *Publications of the Astronomical Society of the Pacific*, 121, 1279
- Stepney, S. 1983, *Monthly Notices of the Royal Astronomical Society*, 202, 467
- Stone, J. M., Gardiner, T. A., Teuben, P., Hawley, J. F., & Simon, J. B. 2008, *The Astrophysical Journal Supplement Series*, 178, 137
- Stone, J. M., Mihalas, D., & Norman, M. L. 1992, *The Astrophysical Journal Supplement Series*, 80, 819
- Stone, J. M., Tomida, K., White, C. J., & Felker, K. G. 2020, *The Astrophysical Journal Supplement Series*, 249, 4
- Sądowski, A., & Narayan, R. 2015, *Monthly Notices of the Royal Astronomical Society*, 453, 3213
- Sądowski, A., Narayan, R., McKinney, J. C., & Tchekhovskoy, A. 2014, *Monthly Notices of the Royal Astronomical Society*, 439, 503
- Tang, Y., & Yao, J. 2021, *Quantitative Imaging in Medicine and Surgery*, doi:10.21037/qims-20-815
- Tao, T., & Blaes, O. 2013, *The Astrophysical Journal*, 770, 55
- Thorne, K. S. 1974, *The Astrophysical Journal*, 191, 507
- Tomida, K., & Stone, J. M. 2023, arXiv e-prints, arXiv:2302.13903
- Török, G., Kotrlová, A., Šrámková, E., & Stuchlík, Z. 2011, *Astronomy and Astrophysics*, 531, A59
- Tully, R. B., Courtois, H. M., Dolphin, A. E., et al. 2013, *The Astronomical Journal*, 146, 86
- Turner, N. J., & Stone, J. M. 2001, *The Astrophysical Journal Supplement Series*, 135, 95
- Ueda, Y., Honda, K., Takahashi, H., et al. 2010, *The Astrophysical Journal*, 713, 257
- Verner, D. A., Ferland, G. J., Korista, K. T., & Yakovlev, D. G. 1996, *The Astrophysical Journal*, 465, 487
- Šrámková, E., Török, G., Kotrlová, A., et al. 2015, *Astronomy and Astrophysics*, 578, A90
- Walton, D. J., Middleton, M. J., Pinto, C., et al. 2016, *The Astrophysical Journal Letters*, 826, L26

- Walton, D. J., Mooley, K., King, A. L., et al. 2017, *The Astrophysical Journal*, 839, 110
- Walton, D. J., Fürst, F., Heida, M., et al. 2018, *The Astrophysical Journal*, 856, 128
- Walton, D. J., Pinto, C., Nowak, M., et al. 2020, *Monthly Notices of the Royal Astronomical Society*, 494, 6012
- Watson, C. C., Hu, J., & Zhou, C. 2020, *IEEE Transactions on Radiation and Plasma Medical Sciences*, 4, 570
- Webb, N. A., Guérou, A., Ciambur, B., et al. 2017, *Astronomy and Astrophysics*, 602, A103
- White, C. J., Mullen, P. D., Jiang, Y.-F., et al. 2023, arXiv e-prints, arXiv:2302.04283
- White, C. J., Stone, J. M., & Gammie, C. F. 2016, *The Astrophysical Journal Supplement Series*, 225, 22
- Wibking, B. D., & Krumholz, M. R. 2022, *Monthly Notices of the Royal Astronomical Society*, 512, 1430
- Wilms, J., Allen, A., & McCray, R. 2000, *The Astrophysical Journal*, 542, 914
- Wood, K., Whitney, B., Bjorkman, J., & Wolff, M. 2013, *Introduction to Monte Carlo Radiation Transfer*, ,
- You, B., Tuo, Y., Li, C., et al. 2021, *Nature Communications*, 12, 1025
- Zdziarski, A. A., Gierliński, M., Rao, A. R., Vadawale, S. V., & Mikołajewska, J. 2005, *Monthly Notices of the Royal Astronomical Society*, 360, 825
- Zdziarski, A. A., Johnson, W. N., & Magdziarz, P. 1996, *Monthly Notices of the Royal Astronomical Society*, 283, 193
- Zhang, S. N., Cui, W., & Chen, W. 1997, *The Astrophysical Journal Letters*, 482, L155
- Zhong, X., & Wang, J. 2021, *The Astrophysical Journal*, 912, 113
- Zhou, M., Abdikamalov, A. B., Ayzenberg, D., et al. 2020, *Monthly Notices of the Royal Astronomical Society*, 496, 497
- Życki, P. T., Done, C., & Smith, D. A. 1999, *Monthly Notices of the Royal Astronomical Society*, 309, 561