GAS DYNAMICS AND STAR FORMATION NEAR SUPERMASSIVE BLACK HOLES

Christopher Charles Frazer

A dissertation submitted to the faculty at the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Physics and Astronomy.

Chapel Hill
2018

Approved by:

Fabian Heitsch

Gerald Cecil

Joaquín E. Drut

Charles R. Evans

Sheila Kannappan

ABSTRACT

Christopher Charles Frazer: Gas Dynamics and Star Formation Near Supermassive Black Holes
(Under the direction of Fabian Heitsch)

The existence of young stars in the central parsec of the Milky Way's Galactic Center (GC) is challenging to explain due to the extreme tidal field of the supermassive black hole (SMBH), Sgr A*, which is expected to disrupt any potentially star forming gas in its vicinity. Yet, numerical models indicate that star formation may occur in this region through two processes. First, the formation and fragmentation of a dense gas disk, which results from rapid gas inflow towards the SMBH, is the most likely origin for the observed stellar disk in the GC. Second, the extreme gravitational compression of low mass gas "clumps" on eccentric orbits around the SMBH provides an alternative explanation for isolated star clusters, as well as newly forming stars, that are not associated with the stellar disk. I build upon these previous theories of nuclear star formation by using three-dimensional grid-based radiative-hydrodynamic simulations to consider the role of ionizing radiation. I begin with an implementation of an adaptive radiative transfer routine for the ATHENA code. I then explore how radiation arising from accretion onto the SMBH may affect the process of gas inflow and stellar disk formation. I then consider the effect of radiation from existing stars on continued star formation occurring via the infall of low-mass gas clumps.

ACKNOWLEDGMENTS

In between the lines, equations, and figures of this dissertation are entire volumes filled with unwritten tales of frustration and failure. The few successful endeavors that I showcase here are built upon a boneyard of mistakes and projects left behind. It is my hope that this document stands as a monument to these untold stories and as a testament to the value of perseverance and patience.

Words cannot repay the innumerable sacrifices made by my fiancèe, Elaine Snyder, during my pursuit of this degree. Her unwavering patience and steadfast love have been the engine behind everything I have done here, and her blood, sweat, and tears have been poured into these projects as much as my own. I cannot imagine what the past few years, or the rest of my life, would be like without her. Elaine, this is for you.

I am most grateful to my research advisor, Fabian Heitsch, for always believing in my abilities and for pushing me to continuously challenge myself. I am immensely proud to have been his student and am thankful for his guidance and mentorship which have made me both a stronger scientist and a better person.

I am fortunate enough to have many friends and a large wonderful family, all of whom have contributed to this work through their support and kindness. At the heart of this research is a lifetime's worth of encouragement and love gifted to me by my parents Donalyn and Brian. It is because of their many sacrifices throughout my life that I have come this far. I am especially thankful to my mother for helping me pack a small black hatchback with all of my earthly belongings and for driving with me from California to North Carolina to begin this venture almost six years ago. I also thank my siblings Shelly, Michael, Kyle, and Matthew for being the best friends that anyone could wish to call their own.

Over the years, I have been lucky to have Brandon Bartell, Christina Haig, Matthew Goodson, Thomas Dombrowski, and John Dupuy as office mates. I want to particularly thank Matthew for being willing to listen as I aired my frustrations, for being an incredible friend, and for always setting an example that was worth following. I also want to singularly thank John, who certainly got the worst of my time and put up with more complaining and whining from me than anyone should ever have to. Lastly, I want to thank Gibson Bennett for choosing to invest his undergraduate research time in my project. He has always been an unexpected source of wisdom and reason throughout the highs and lows of this experience.

The simulations presented in this dissertation were performed on the Killdevil and Dogwood clusters at UNC-Chapel Hill. I am indebted to the staff at UNC Research Computing for their continuous effort to improve the resources available to students and for sparing the time to sort out even my smallest computing issues.

The work done here would not be possible without the generous contributions made to the Department of Physics and Astronomy in the name and honor of Edgar T. Cato and of the Southern Astrophysical Research (SOAR) Telescope. The foundation of this research was built through support from the Cato-SOAR fellowship, and I am deeply grateful for the opportunity to represent this investment in both Astronomy and graduate research.

I am grateful to the North Carolina Space Grant for two summers of research funding that allowed me to explore a variety of different projects, the results of which manifest in this work. I also thank NASA ATP (grant #NNX10AC84G) for funding through the initial stages of this research.

I would lastly like acknowledge the financial support entrusted to me through the Paul Hardin Dissertation Completion Fellowship of the Royster Society at UNC-Chapel Hill. Truly, this fellowship has served as a last, but crucial, piece of the puzzle that has been my graduate experience.

TABLE OF CONTENTS

LIST OF FIGURES . . . . **xi**

LIST OF TABLES . . . . **xiii**

LIST OF ABBREVIATIONS AND SYMBOLS . . . . **xiv**

1 INTRODUCTION **1**

1.1 A Census of Gas in the Galactic Center 1

1.2 The Nuclear Stellar Cluster 3

1.2.1 The Nuclear Stellar Disk 3

1.2.2 The S-Star Cluster 4

1.2.3 On-going Star Formation 5

1.3 Star Formation Theories 5

1.3.1 High Mass Inflow 6

1.3.2 Clump Compression 6

1.4 Overview 7

2 NUMERICAL METHODS **8**

2.1 Grid-Based Hydrodynamics 8

2.2 Source Terms 9

2.2.1 Static Gravitational Potential 10

2.2.2 Energy Gains and Losses 11

2.2.3 Radiative-Hydrodynamics 11

2.3 Static Mesh Refinement 12

2.4 Modifications to Athena 12
2.4.1 Internal Energy Advection 13
2.4.2 Refinement Interpolation 14
2.4.3 Primitive Grid Restriction 15
3 ADAPTIVE RADIATIVE TRANSFER 16
3.1 Motivation 16
3.2 The Equation of Radiative Transfer 16
3.3 Adaptive Ray Tracing 18
3.3.1 Ray-Tree Initialization 19
3.3.2 Parallel Integration 21
3.4 Included Physics 26
3.4.1 Photo-absorption 26
3.4.2 Ionization and Recombination 27
3.4.3 Compton Heating 29
3.4.4 Secondary Ionizations 30
3.4.5 Radiation Pressure 31
3.4.6 Thermal Physics 31
3.4.7 Collisionally Excited Line Radiation 33
3.5 Radiation Timestep 34
3.5.1 Sub-cycling Thermal Physics 36
3.6 Test Cases 37
3.6.1 R-type Ionization Front without Recombination 37
3.6.2 R-type Ionization Front with Recombination 38
3.6.3 D-type Ionization Front 41
3.6.4 Radiation Pressure 42

3.6.5 Secondary Ionization . . . 43
3.7 Additional Features . . . 45
3.7.1 Active Ray Tracing . . . 45
3.7.2 Periodic Rotation . . . 48
3.7.3 Selective Integration . . . 50
3.7.4 Off-Mesh Ray Tracing . . . 52
3.8 Performance . . . 55
3.8.1 Strong Scaling . . . 55
3.8.2 Weak Scaling . . . 57

4 GAS INFLOW AND STAR FORMATION AROUND SUPERMASSIVE BLACK HOLES: THE ROLE OF NUCLEAR ACTIVITY . . . 59
4.1 Introduction . . . 59
4.1.1 Star Formation in the Galactic Center . . . 59
4.1.2 Nuclear Activity in the Galactic Center . . . 60
4.1.3 Motivation and Outline . . . 62
4.2 Methods . . . 62
4.2.1 Radiation . . . 63
4.3 Simulation Set-up . . . 63
4.3.1 Accretion Boundary . . . 64
4.3.2 Inflow Conditions . . . 65
4.3.3 Radiation Field . . . 67
4.3.4 Models . . . 69
4.3.5 Disk Finding and Stability Measure . . . 70
4.3.6 Disk Finding Algorithm . . . 72
4.4 Results . . . 73

4.4.1 Dynamics 73
4.4.2 Thermal Evolution 77
4.5 Discussion 82
4.5.1 The Effects of X-Rays 82
4.5.2 The Effects of UV Photons 85
4.5.3 The Role of Radiation Pressure 87
4.5.4 Disk Formation 89
4.5.5 Mass Accretion 91
4.5.6 Disk Substructure 94
4.6 Conclusions 97

5 ECCENTRIC GAS CLUMP ORBITS AND STAR FORMATION IN THE GALACTIC CENTER: THE ROLE OF IONIZING RADIATION 101

5.1 Introduction 101
5.1.1 Star Formation in the Galactic Center 101
5.1.2 Motivation 102
5.1.3 Rationale for Including Radiation 102
5.1.4 Outline 104
5.2 Methods 104
5.3 Simulation Set-up 104
5.3.1 Orbits 106
5.3.2 Radiation Fields 106
5.3.3 Co-moving Mesh 108
5.4 Results 108
5.4.1 Orbital Evolution 108
5.4.2 Gas Compression 111

5.4.3 Gas Heating and Cooling 115
5.5 Discussion 117
5.5.1 Tidal Densities 117
5.5.2 The Effect of Clump Temperature 119
5.5.3 Resolution Limits and Convergence 121
5.5.4 Other Considerations 123
5.6 Conclusions 125
6 CONCLUSION 127
BIBLIOGRAPHY 133

LIST OF FIGURES

2.1 Static Mesh Refinement demonstration 13

3.1 Restricted radius calculation flow chart 22

3.2 2D representation of ray passing geometries 23

3.3 Parallel ray walk flow chart 25

3.4 Parallel ray integration flow chart 26

3.5 CLE cooling grid and piecewise approximation 35

3.6 Ionization fraction profile for R-type ionization front without recombinations and analytic comparison 39

3.7 Ionization fraction profile for R-type ionization front with recombinations and analytic comparison. 40

3.8 Number density profile for D-type ionization front and analytic comparison . . . 42

3.9 Analytical comparison for radiation pressure 44

3.10 Temperature and density profiles for secondary ionization test case 45

3.11 Parallel active ray integration flow chart 48

3.12 D-type ionization front with and without ray-tree rotation 49

3.13 Off-Mesh Ray Tracing comparison 54

3.14 Strong scaling for fixed and active ray tracing 56

3.15 Weak scaling for fixed and active ray-tree integration 57

4.1 Initial conditions and computational mesh 66

4.2 Approximate AGN spectral energy distribution 68

4.3 Dynamics for lower mass inflow 74

4.4 Dynamics for higher mass inflow 76

4.5 Temperature-density evolution for lower mass models 79

4.6 Temperature-density evolution for higher mass models 80

4.7 Column density of monochromatic models 83

4.8 Temperature-density histogram of monochromatic models 84

4.9 Midplane gas density and ionized gas density for central disk 87

4.10 Midplane gas density for models with and without radiation pressure 88

4.11 Disk parameters and Toomre Q evolution 90

4.12 Mass accretion and mass accretion rate 92

4.13 Peak density spatial distribution over time 96

4.14 Spatial distribution of peak density relative to the tidal limit over time 98

4.15 Spatial distribution of peak density relative to the tidal limit for high-resolution models 99

5.1 Mesh and initial conditions 106

5.2 Orbital evolution for $\epsilon = 0.5$ 110

5.3 Orbital evolution for $\epsilon = 0.9$ 112

5.4 Density distribution of clump gas over time 113

5.5 Temperature distribution of clump gas over time 114

5.6 Peak density vs. time for multiple radiation strengths 118

5.7 Peak density vs. time for multiple minimum temperatures 120

5.8 Peak density vs. time for multiple grid resolutions 122

LIST OF TABLES

3.1 Ray structure variables 20

3.2 Ionization test refinement geometry 37

4.1 Mesh and refinement geometry for inflow models 64

4.2 Inflow parameters 69

4.3 Radiation parameters for inflow models 70

5.1 Mesh and refinement geometry for clump models 105

5.2 Orbital parameters 107

5.3 Radiation parameters for clump models 107

LIST OF ABBREVIATIONS

AMR	Adaptive Mesh Refinement
AGN	Active Galactic Nuclei
CMZ	Central Molecular Zone
CND	Circumnuclear Disk
GC	Galactic Center
HEALPix	Hierarchical Equal Area isoLatitude Pixelation
IMBH	Intermediate Mass Black Hole
IMF	Initial Mass Function
ISM	Interstellar medium
I-front	Ionization Front
MPI	Message Passing Interface
NSC	Nuclear Stellar Cluster
RHD	Radiative Hydrodynamics
SMBH	Supermassive Black Hole
SMR	Static Mesh Refinement
UV	Ultraviolet
VL	van Leer (method)
WR	Wolf-Rayet (stars)
YSO	Young Stellar Object

CHAPTER 1: INTRODUCTION

The Milky Way's Galactic Center (GC) sits roughly 8 kpc from our Sun (Reid, 1993; Ghez et al., 2008; Eisenhauer et al., 2005; Gillessen et al., 2009b) and hosts the only nuclear stellar cluster (NSC) observed at sufficiently high resolution to identify individual stars. Over the past 20 years, several observational programs following the orbits of individual stars in the NSC have unveiled a central supermassive black hole (SMBH), Sgr A*, (Schödel et al., 2002; Ghez et al., 2003; Gillessen et al., 2009a) and ≈ 100 young, massive stars (Genzel et al., 2003; Paumard et al., 2006; Bartko et al., 2009, 2010). For this reason, the GC has been regarded as an "exquisite laboratory" for understanding gas dynamics and star formation both in extreme conditions characteristic of galactic nuclei and in the immediate vicinity of a SMBH where, until recently, star formation was believed to be impossible.

1.1 A Census of Gas in the Galactic Center

Roughly 10% of the Milky Way's molecular gas content is found within the few$\times$100 pc of Sgr A* (Morris & Serabyn, 1996) in a region referred to as the Central Molecular Zone (CMZ). Many of the most massive gas complexes in this region orbit along a large-scale "warped ring" (Kruijssen et al., 2015; Molinari et al., 2011), possibly due to an accumulation of mass resulting from orbital resonances with the Galactic bar (Krumholz & Kruijssen, 2015). Despite high gas densities in the CMZ ($n \gtrsim 10^4$ cm^{-3}), star formation is suppressed by a factor of 10 with respect to the solar neighborhood (Kauffmann et al., 2017; Longmore et al., 2013; Guesten & Downes, 1983), though the reasons for this remain a matter of debate (see discussion in Kruijssen et al. (2014)).

Two massive ($> 10^5\ M_\odot$) molecular gas complexes in the CMZ, M-0.13-0.08 (the "50 km s^{-1}

cloud") and M-0.02-0.07 (the "20 km s^{-1} cloud"), are currently positioned within 10–50 pc of Sgr A* (Lee et al., 2008; Coil & Ho, 1999; Mezger et al., 1989; Herrnstein & Ho, 2005; McGary et al., 2001; Coil & Ho, 2000). It is unclear how these gas structures interact with the inner few parsecs of the GC as the three-dimensional structure of the gas is uncertain (Ferrière, 2012). There is evidence, though, that both M-0.13-0.08 and M-0.02-0.07 are evolving in response to the environment surrounding the SMBH. First, a "molecular ridge" ($M = 10^4\ M_\odot$) extending from M-0.13-0.08 appears to follow the outer edge of the non-thermal shell, Sgr A East, which surrounds the central parsecs of the GC (Coil & Ho, 2000). Furthermore, Sgr A East may be expanding into M-0.13-0.08 itself (Mezger et al., 1989), possibly inciting nuclear gas inflow (Ho et al., 1991). Second, the so-called "southern streamer" flows from M-0.02-0.07 towards the SMBH, channelling cold gas towards the central parsecs of the Galaxy.

At the center of the CMZ, where the gravitational potential is dominated by the SMBH ($r \lesssim 2$ pc; (Schödel et al., 2007; Genzel et al., 2003)), a coherently orbiting group of clumpy molecular gas structures extends from $\approx 2 - 5$ pc. This so-called Circumnuclear Disk (CND[1]) is possibly the remnant of past star formation activity (Sanders, 1998; Alig et al., 2013; Mapelli & Trani, 2016) and is likely on the brink of star formation in the near future itself (Oka et al., 2011; Christopher et al., 2005; Yusef-Zadeh et al., 2008). Early mass estimates from HCN/HCO+ observations ($M_{\mathrm{CND}} = 10^6\ M_\odot$; Christopher et al. (2005); Montero-Castaño et al. (2009)) suggest that the CND is a semi-permanent feature of the GC ($\tau \approx 10^7$ yr). Yet, more recent measurements have determined that individual gas clumps in the CND are not gravitationally bound, thus leading to a considerably lower total mass estimate ($M_{\mathrm{CND}} \approx 10^4\ M_\odot$; Requena-Torres et al. (2012); Lau et al. (2013); Mills et al. (2013)) and projected lifetime ($\approx 10^5$ yr). The transient nature of the CND remains an on-going discussion. It is clear, though, that the CND is actively accreting material from the

[1] The CND is interchangeably referred to as the Circumnuclear Molecular Ring (CMR) or the Circumnuclear Ring (CNR).

CMZ (Oka et al., 2011; Coil & Ho, 1999), and that gas from its inner edge is kinematically linked to the ionized mini-spiral, Sgr A West, that surrounds Sgr A* (Lo & Claussen, 1983; Christopher et al., 2005; Zhao et al., 2010). Interior to the CND, intense ultraviolet (UV) radiation from a few $\times$ 100 OB/WR stars (Paumard et al., 2006; Gillessen et al., 2009b; Lu et al., 2009, 2013; Yelda et al., 2014) heats dense gas to temperatures of a few $\times 10^3$ K (Zhao et al., 2010; Moser et al., 2017). The manner by which cold gas can penetrate into the central parsecs of the GC, as is required for a nuclear star formation episode, is unclear.

1.2 The Nuclear Stellar Cluster

NSCs are extremely crowded regions with total stellar masses of $\approx 10^5 - 10^7\ M_\odot$ (Walcher et al., 2005) within few $\times$ pc of the galactic nucleus (Böker et al., 2004) and are found in the centers of most spiral galaxies (Côté et al., 2006). Approximately $\approx$ 10000 orbital solutions for individual stars in the Milky Way's NSC constrain the measurements of the SMBH mass ($\approx 4 \times 10^6\ M_\odot$) and the distance to the GC ($\approx$ 8.5 kpc) (Genzel et al., 2003; Trippe et al., 2008; Schödel et al., 2007; Chatzopoulos et al., 2015). Genzel et al. (2003) and Schödel et al. (2007) both show that the mass distribution of the NSC follows an isothermal profile ($\rho \propto r^{-2}$) for $r \gtrsim 0.25$ pc, indicating that the stellar orbits are dynamically relaxed. Interior to this region, the profile is cusped. The NSC is predominantly an aged population, but the presence of young, massive (OB/WR) stars in the central parsec of the GC has been known for over two decades (Morris & Serabyn (1996) and references therein). Their origin is a topic of current research.

1.2.1 The Nuclear Stellar Disk

Many of the young stars in the NSC orbit within a clockwise (on the plane of the sky) eccentric stellar disk at 0.05 pc $< r <$ 0.5 pc (Paumard et al., 2006; Lu et al., 2009, 2013; Bartko et al., 2009; Yelda et al., 2014). The initial mass function (IMF) of the disk stars follows a slope of $\alpha \approx 1.7$ (as compared to the Salpeter value of $\alpha = 2.35$), showing a clear preference for massive star formation (Lu et al., 2013). The disk is also characterized by a mild eccentricity ($\epsilon = 0.2 - 0.4$; Yelda et al. (2014)). The cluster age of 2.5-5.8 Myr (Paumard et al., 2006; Lu et al., 2013) is compatible with a local episode of star formation in the the last 3-7 Myr (Krabbe et al., 1995). Yet, in order for stars to have formed in the region inhabited by the stellar disk, the density of the progenitor gas structure must have exceeded the tidal limit ($n \gtrsim 10^8$ cm^{-3}; Yusef-Zadeh et al. (2013)) to avoid disruption in the steep potential of the SMBH. Gas densities in the CMZ ($n \lesssim 10^5$ cm^{-3}) fall well below this limit, therefore a compression mechanism is required to support in situ formation of the stellar disk. One possible avenue for sufficient gas compression suggests a massive ($\approx 10^5$ $M_\odot$) inflow from the CMZ can result in a sufficiently dense gas disk to allow for gravitational fragmentation (Wardle & Yusef-Zadeh, 2008). Numerical models of this process result in stellar populations that are consistent with the orbital structure and IMF of the nuclear stellar disk (see § 1.3.1).

1.2.2 The S-Star Cluster

The inner-most stars in the GC follow randomly distributed orbits with pericenter distances within 0.04 pc of the SMBH (i.e. the "S-star cluster; Ghez et al. (2003); Schödel et al. (2003); Ghez et al. (2005); Eisenhauer et al. (2005)). Orbital solutions obtained for a majority of these stars show a preference towards higher eccentricities ($n(e) \approx e^{2.6\pm0.9}$) than the stellar disk, indicating they are dynamically distinct from other stellar populations in the vicinity (Gillessen et al., 2009a). The minimum age of the S-star cluster members ($\tau \approx 6$ Myr; Eisenhauer et al. (2005); Gillessen et al. (2009a)) also appears to require an in situ formation avenue. Models exploring the formation of disk stars are unable to replicate the observed stellar orbits for the S-stars, thus relaxation processes are typically invoked to explain their presence (see (Mapelli & Gualandris, 2016) and references therein). Currently there is no theoretical compression mechanism that allows for gas to achieve the densities required for in-situ star formation in the region occupied by the S-stars ($n > 10^{12}$ cm^{-3}).

1.2.3 On-going Star Formation

The S-star cluster and the nuclear stellar disk suggest that in situ star formation may have occurred in the past; however, there are several indications that star formation is actively occurring throughout the central 2 pc of the GC. As previously mentioned, the stability of the CND is highly contentious, though continued accretion (Oka et al., 2011) and maser activity within the disk (Yusef-Zadeh et al., 2008) indicate that star formation in its vicinity is possible. Observations of molecular gas clumps in the central parsec (which may originate from within the CND (Jalali et al., 2014)) show that cold gas survives both the tidal shear of the SMBH and the intense radiation field of massive stars (Yusef-Zadeh et al., 2013; Moser et al., 2017). Several protostellar disk candidates have also been identified in this region (Yusef-Zadeh et al., 2013, 2015). Recently, the discovery of 11 bipolar outflows appears to concretely confirm that star formation continues to occur near the SMBH (Yusef-Zadeh et al., 2017), revitalizing the search for mechanisms that drive isolated bursts of nuclear star formation.

1.3 Star Formation Theories

It is difficult to reconcile the apparent need for in-situ star formation with the extreme tidal field in the immediate presence of the SMBH. For the S-stars in particular, this issue has been coined "The Paradox of Youth" (Ghez et al., 2003). Alternative explanations for the existence on young stars hinge on the presence of external forces (i.e. an over-abundance of massive stars in the GC (Gürkan & Rasio, 2005) or an intermediate mass black hole (IMBH; Hansen & Milosavljević (2003))) to expedite the transit time of star clusters migrating to the inner parsecs of the GC, though observational and timescale constraints do not support such a picture (Stolte et al., 2008; Genzel et al., 2010). There are currently two mechanisms that allow for star formation near a supermassive black hole. I describe these below.

1.3.1 High Mass Inflow

Stellar disk formation may result from the gravitational fragmentation of an extended accretion disk (Nayakshin et al., 2007) following the collision between the SMBH and a $\approx 10^5\ M_\odot$ gas stream infalling from the CMZ (Wardle & Yusef-Zadeh, 2008). Several numerical studies have considered this process over the past decade, collectively demonstrating that it is a viable candidate for the origin of the observed stellar disk in the GC (Sanders, 1998; Lucas et al., 2013; Bonnell & Rice, 2008; Mapelli et al., 2012; Alig et al., 2011). What remains uncertain is the source of such massive gas inflow required to drive disk formation. As noted by Hobbs & Nayakshin (2009), the orbits of both the Arches cluster (Stolte et al., 2008) and gas clouds along the warped ring in the CMZ (Kruijssen et al., 2015) follow an elliptical trajectory, therefore it is reasonable to believe that such a collision between a large gas cloud and the SMBH is possible. Alternative mechanisms for driving sufficient inflow rely on gas stream collisions with the CND (Alig et al., 2013) or inner-CMZ clouds (Hobbs & Nayakshin, 2009), though these models typically result in multiple, distinct stellar disks. Previous work on this subject also shows that this star formation process would spark an accretion episode for the SMBH, possibly giving rise to AGN activity (Bonnell & Rice, 2008; Hobbs & Nayakshin, 2009). The role of radiative feedback resulting from accretion has yet to be explored, and is the topic considered in Chapter 4.

1.3.2 Clump Compression

An alternative mode of star formation that has been proposed relies on the extreme compression of infalling gas clumps (Eckart et al., 2004). Jalali et al. (2014) explore this effect through hydrodynamic simulations of 100 $M_\odot$ gas clumps on eccentric orbits with pericenter distances of $\gtrsim 0.1$ pc. These models produce individual stellar clusters similar to those observed in the GC (i.e. IRS-13N (Eckart et al., 2004; Mužić et al., 2008)). Jalali et al. (2014) also show that clump-clump collisions within the CND can give rise to periodic star formation episodes through this process on a timescale of 10 Myr^{-1}. Given the conclusive evidence of protostellar outflows in the study of Yusef-Zadeh et al. (2017), a continuation of the models of Jalali et al. (2014) is warranted. I build on this previous work by including the effect of ultraviolet radiation from existing stars on infalling gas clumps. I detail this study and my findings in Chapter 5.

1.4 Overview

I begin in Chapter 2 with an overview of the ATHENA code, and discuss some minor changes necessary for this work. The implementation of an adaptive radiative-transfer code for ATHENA is detailed in Chapter 3 along with several benchmark tests to demonstrate its accuracy. Chapter 4 revisits the "cloud inflow" scenario for in situ stellar disk formation in the GC, including radiative feedback from accretion onto the central SMBH. In Chapter 5, I consider the "clump infall" scenario which is well known to result in isolated star formation episodes, and include radiative feedback from existing stars to determine the effect of ionizing radiation in this process. I conclude with a brief summary in Chapter 6.

CHAPTER 2: NUMERICAL METHODS

2.1 Grid-Based Hydrodynamics

The evolution of inviscid compressible fluid systems is governed by the Euler equations which enforce conservation of mass, momentum, and energy:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0 \tag{2.1.1}$$

$$\frac{\partial \rho \mathbf{v}}{\partial t} + \nabla \cdot (\rho \mathbf{v}\mathbf{v} + \mathbf{I}P) = 0 \tag{2.1.2}$$

$$\frac{\partial E}{\partial t} + \nabla \cdot (\mathbf{v}(E + P)) = 0 \tag{2.1.3}$$

with the total mass density ρ, fluid velocity vector $\mathbf{v}$, gas pressure P, unit dyad $\mathbf{I}$, and energy density

$$E = \frac{1}{2}\rho \mathbf{v} \cdot \mathbf{v} + \frac{P}{\gamma - 1} \ , \tag{2.1.4}$$

which depends on the adiabatic index γ such that the internal energy density $e = P(\gamma - 1)^{-1}$. To trace individual parcels of gas in a fluid system, an additional scalar advection equation can be employed:

$$\frac{\partial C\rho}{\partial t} + \nabla \cdot (C\rho \mathbf{v}) = 0 \ , \tag{2.1.5}$$

where C is the passive scalar quantity. C is interchangeably referred to as a color field.

For this work, I use the ATHENA hydrodynamics code (Stone et al., 2008) which numerically integrates the fluid equations (Eqs. 2.1.1–2.1.3) iteratively on a three-dimensional grid via a higher-order Godunov scheme. In ATHENA, the conserved variables (mass, momentum, and energy) are computed as volume averaged quantities that are represented as cell center values on the grid. The solution to the conservation equations at each computational timestep follows from these steps: (i) spatial interpolation (typically linear, Colella & Woodward (1984)) is used to approximate the fluid quantities (conserved or primitive[1]) at cell interfaces. Interpolations are used on each side of the cell interface to determine the initial condition for the so-called Riemann problem. The Riemann problem describes how an initially discontinuous fluid system evolves in time. The fluid responds to the discontinuity causing waves to propagate from the cell interface at speeds depending on the initial condition. (ii) Time averaged fluxes are calculated by imposing conservation across the regions separated by the wave characteristics in the Riemann problem. I use the HLLC Riemann solver (Toro, 2009) which builds upon the HLL method via a three-wave approximation that restores the contact wave. (iii) Intercell fluxes are applied to each cell to update the conserved quantities in time. I use the Van Leer integrator (VL; Stone & Gardiner (2009)), a predictor-corrector scheme, to achieve second order time accuracy. For a fixed grid resolution, Δx, ATHENA solves these equations in timesteps that are restricted by the Courant-Friedrichs-Lewy condition (i.e. $\Delta t = \alpha_{\rm CFL}\, \Delta x\, \max(|\mathbf{v}| + c_s)^{-1}$, where c_s is the gas sound speed; Courant et al. (1928)). The value of the Courant number, $\alpha_{\rm CFL}$, depends on the problem, but I find that $\alpha_{\rm CFL} \approx 0.1$ provides for sufficient integration stability in most problems.

2.2 Source Terms

ATHENA solves the conservation equations as detailed above. Yet, for the applications considered in this work, non-conservative processes must be introduced via source terms to account for additional physics. Here, I detail the required source terms for my models.

[1]Primitive variables are the fluid velocity vector $\mathbf{v}$ and the pressure, P. Density is also included with the primitive variables.

2.2.1 Static Gravitational Potential

Self-gravity of the gas is not included in my models (I discuss the implications of this simplification in Chapters 4 and 5). I include the effect of the point sources of gravity (i.e. the SMBH) via a static gravitation potential. Because the conservative formulation in ATHENA does not include gravitational energy in the total energy (Eq. 2.1.4), the potential must be implemented as a source term. Work done on the gas by the gravitational field reflects in changes to the kinetic energy so that Eqs. 2.1.2 and 2.1.3 are rewritten as:

$$\frac{\partial \rho \mathbf{v}}{\partial t} + \nabla \cdot (\rho \mathbf{v}\mathbf{v} + \mathbf{I}P) = -\rho \nabla \phi \tag{2.2.1}$$

$$\frac{\partial E}{\partial t} + \nabla \cdot (\mathbf{v}(E + P)) = -\rho \mathbf{v} \cdot \nabla \phi \ , \tag{2.2.2}$$

where ϕ is the static gravitational potential, which, for a point mass (M) takes the form:

$$\phi = -\frac{GM}{r} \tag{2.2.3}$$

In the original implementation for static gravitational potentials in ATHENA, the momentum density is updated via an Euler step, but the energy density update depends on the density flux. In rare cases, the differences in the kinetic energy density implied in these two independent processes was sufficient to cause the internal energy of the cell ($e = E - \frac{1}{2}\rho \mathbf{v} \cdot \mathbf{v}$) to be negative. As an alternative to this approach, a forward Euler step is applied to both the momentum and the energy, so that, in one dimension:

$$\rho u' = \rho u + \rho \Delta t \frac{\phi_{i+1/2} - \phi_{i-1/2}}{\Delta x} \tag{2.2.4}$$

$$E' = \frac{\rho (u')^2}{2} + \frac{P}{\gamma - 1} \tag{2.2.5}$$

where primed values indicate updated quantities, and u is the gas velocity. This change guarantees that variations in kinetic energy are reflected uniformly in both the momentum density and total energy density, guaranteeing temperature positivity.

2.2.2 Energy Gains and Losses

Approximate thermal physics is incorporated to account for the changes of gas temperature due to optically thin radiation and ambient gas heating. In Chapter 3, I discuss a wide variety of heating and cooling mechanisms in interstellar gases, particularly in the presence of an ionizing radiation source. In the absence of additional source terms, energy gains (G) and radiative losses (L) are included to the total energy equation (Eq. 2.1.3) so that:

$$\frac{\partial E}{\partial t} + \nabla \cdot (\mathbf{v}(E+P)) = G - L \ , \tag{2.2.6}$$

ATHENA has been modified to allow for second order time accuracy for these thermal source terms within the VL integrator. For simulations with radiation, these terms are calculated in an operator split fashion in conjunction with the radiation update which is first order accurate in time.

2.2.3 Radiative-Hydrodynamics

In Chapter 3, I detail a method for incorporating ionization and heating from point sources of radiation. To include radiative effects, I modify Eqs. 2.1.2 and 2.1.3 so that:

$$\frac{\partial \rho \mathbf{v}}{\partial t} + \nabla \cdot (\rho \mathbf{v}\mathbf{v} + \mathbf{I}P) = \rho \mathbf{a}_\gamma \tag{2.2.7}$$

$$\frac{\partial E}{\partial t} + \nabla \cdot (\mathbf{v}(E+P)) = \rho \mathbf{v} \cdot \mathbf{a}_\gamma + G - L \tag{2.2.8}$$

where I have introduced the an acceleration term $\mathbf{a}_\gamma$ to account for radiation pressure. To track the ionization of the gas, I introduce an additional equation for the conservation of ionized hydrogen mass density (ρ_{HII}):

$$\frac{\partial \rho_{\mathrm{HII}}}{\partial t} + \nabla \cdot (\rho_{\mathrm{HII}} \mathbf{v}) = m_{\mathrm{H}}(I - R) \ . \tag{2.2.9}$$

The ionized hydrogen mass density depends on the ionization rate per volume, I, and the recombination rate per volume R (see § 3.4.2). Photo-heating is incorporated into the energy gains term G, and a variety of cooling processes for both ionized and neutral gasses is built into L (see § 3.4.6).

2.3 Static Mesh Refinement

Athena is outfitted with Static Mesh Refinement (SMR) which allows for the targeted placement of additional refinement in regions where higher resolution is required. In contrast to Adaptive Mesh Refinement (AMR), SMR is not dynamic to the flow of the gas, therefore the geometry of refinement zones must be known a priori. Although this is somewhat limiting, SMR has the added benefit of increased computational efficiency and ease in load balancing of computational resources. In athena, SMR zones are restricted to resolution increases in factors of two, though many refinement zones can be used to provide for a large range of spatial scales. In Figure 2.1, I show a sample application of SMR in which a gas clump is centered on the computational domain. In one frame, low resolution results in a highly discretized gradient structure on the edge of the clump where individual grid cells are clearly visible. Adding nested levels of refinement around the clumps allows for a smooth gas distribution without increasing the resolution (and computational cost) globally. The placement of SMR zones is highly dependent on the problem being considered. I discuss applications of SMR in Chapters 3, 4, and 5.

2.4 Modifications to Athena

Several tools were added to athena to improve integration stability and prevent occasional failures in the hydrodynamic update. Here, I detail a fall-back internal energy integration scheme (§ 2.4.1) and two additions to the refinement synchronization for prolongation (§ 2.4.2) and coarse grid restriction (§ 2.4.3).

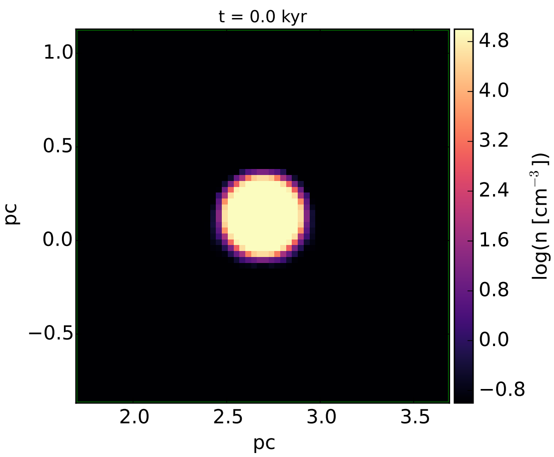

(a) No refinement

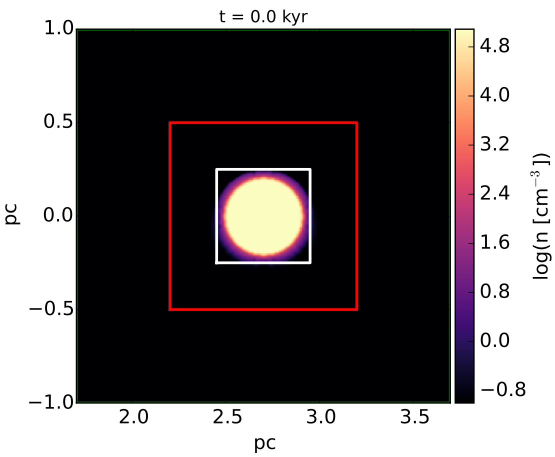

(b) Multiple tiers of refinement

Figure 2.1: Demonstration of an application of SMR. (a) Gas number density of a singular gas clump which sits at the center of a uniform resolution grid. Grid cells are clearly visible leading to sharp transitions between cells. (b) Same as in (a), but with two levels of SMR (square boxes). The resolution increases in factors of two within each nested refinement domain.

2.4.1 Internal Energy Advection

In grid cells where the kinetic energy is a significant fraction of the total energy, it is possible that the resulting internal energy density ($e = (E - \frac{1}{2}\rho\mathbf{v} \cdot \mathbf{v})$) will be negative. For this reason I have implemented a procedure that is similar to that of Bryan et al. (2014) by additionally solving the internal energy equation:

$$\frac{\partial e}{\partial t} + \nabla \cdot (\mathbf{v}e) = -P\nabla \cdot \mathbf{v} \; . \tag{2.4.1}$$

The solution to this equation is split into two stages. First, thermal energy is transported proportionally to the density flux by solving the left hand side of Eq. 2.4.1 in conservative form:

$$\frac{\partial (C_e \rho)}{\partial t} + \nabla \cdot (C_e \rho \mathbf{v}) = 0 \; , \tag{2.4.2}$$

where

$$C_e = \frac{e}{\rho} \tag{2.4.3}$$

is the specific internal energy. Interface states are calculated independently for internal energy, and fluxes are determined in the same manner as a color field. The resulting transport update in one dimension is then,

$$e^{n+1} = e_i^n + \frac{\Delta t}{\Delta x}(F_{e,i-1/2} - F_{e,i+1/2}) \ , \tag{2.4.4}$$

where $F_{e,i\pm1/2}$ are the internal energy fluxes at the boundaries of the cell. The right hand side of Eq. 2.4.1 is included as a source term which is solved in tandem with the source terms for the conserved variables within the VL integrator. The resulting change in the internal energy is calculated via central difference:

$$\Delta e = -P\Delta t(\frac{v_{x,i+1} - v_{x,i-1}}{2\Delta x} + \frac{v_{y,j+1} - v_{y,j-1}}{2\Delta y} + \frac{v_{z,k+1} - v_{z,k-1}}{2\Delta z}) \tag{2.4.5}$$

Source terms are applied at the half timestep using the pressure calculated from the initial condition. In the corrector step of the VL integrator, the pressure is calculated from the half timestep values. My approach differs from Bryan et al. (2014) as I do not continually track the evolution of the internal energy independently of the total energy. Instead, the internal energy is calculated from the total energy at the beginning of every integration step. Pressure positivity is checked both at the predictor and corrector stages of the VL integrator. If the pressure calculated from the conserved variables is negative, the internal energy solution is used as a replacement. The change in the total energy imposed by this process is typically $\delta E/E < 10^{-5}$.

2.4.2 Refinement Interpolation

The directionally split boundary value interpolation method used for refinement synchronization in ATHENA can result in negative pressures in regions with steep gradients in momentum or pressure. To avoid this issue, I have included the "second-order A" interpolation scheme from the ENZO code (Bryan et al., 2014). In short, this method uses a tri-linear interpolation to calculate the conserved variables at the corners of the coarse cell. Monotonized slopes are then calculated along each diagonal of the cell, and an additional constraint is applied to the slopes to remove outliers. The resulting slopes are then applied to calculate the refined zone boundary condition. I make two changes to this routine. First, I impose the condition that, for each conserved variable, the corner values must satisfy the condition $0.1q_c < q < 10q_c$, where q is the corner value and q_c is the cell centered value. Additionally, rather than computing the interpolation on the total energy, I use the pressure as is done in the directionally split approach native to athena.

2.4.3 Primitive Grid Restriction

The restriction scheme of the Athena stock version uses conserved variables (density, momenta, and total energy) to synchronize the solutions of coarser grid levels with averaged values from finer levels. Under the physical conditions met in my simulations, this process led to unphysical feedback at the refinement boundaries. Unresolved motion from converging flows on the fine grid can be translated into thermal energy on the coarse grid. This results in an overpressure on the coarse grid that imprints the mesh structure onto the gas and dampens fluid flow onto higher refinement domains. I resolved this issue by restricting on the pressure rather than the total energy. By doing this, mass and momentum are perfectly conserved, but I sacrifice total energy conservation for improved continuity of fluid flow and increased stability. For the conditions in my models, this process leads to an average error of $\delta E/E \lesssim 10^{-3}$.

CHAPTER 3: ADAPTIVE RADIATIVE TRANSFER

3.1 Motivation

The problem of radiative transfer is ubiquitous in astronomy and can be incorporated into numerical simulations in a variety of ways (see Iliev et al. (2009) and references therein). A previous implementation of radiative transfer via ray-tracing in ATHENA is shown in Krumholz et al. (2007), but this feature is not included in the public release of the code. Furthermore, it was not designed to be used with the SMR framework native to ATHENA and does not include all the physics components required for this work. I roughly follow the structure of this previous implementation, but include features and simplifications from Wise & Abel (2011) that were necessary for making my models computationally feasible. In this chapter I detail the radiative transfer routine which I have written for Athena (§ 3.3), discuss the included physics in this routine (§ 3.4), briefly cover the time-stepping restrictions imposed (§ 3.5), demonstrate the accuracy of my code through benchmark tests (§ 3.6), provide an overview of additional features and their applications (§ 3.7), and conclude with a discussion of parallel performance (§ 3.8).

3.2 The Equation of Radiative Transfer

The general problem that must be solved is the time evolution of a gaseous medium with an embedded radiation source. Therefore, the treatment for the absorption of radiation by intervening columns of material is explained here. The equation of radiative transfer that governs this process is (Abel et al., 1999):

$$\frac{1}{c}\frac{\partial I_\nu}{\partial t} + \hat{n}\frac{\partial I_\nu}{\partial x} = -\kappa_\nu I_\nu + j_\nu \tag{3.2.1}$$

where c is the speed of light, I_ν is the specific intensity, κ_ν is the attenuation coefficient, j_ν is the emissivity, and $\hat{n}$ is the direction of propagation. To simplify this equation, the radiation field is represented using monochromatic bins which eliminates the frequency dependence (Wise & Abel, 2011; Rosen et al., 2017). I also ignore the emissivity of the gas ($j_\nu = 0$) and take the attenuation coefficient to be constant within each timestep of the routine, thus eliminating the time dependence and reducing Eq. 3.2.1 to:

$$\hat{n}\frac{dI}{dx} = -\kappa I \ . \tag{3.2.2}$$

The solution to this equation is a simple attenuation law,

$$I_f = I_i \exp(-\kappa \Delta s) \ , \tag{3.2.3}$$

where Δs is the path length along which the light travels. I_i and I_f are the initial and final specific intensities, respectively. The attenuation coefficient, κ, can be expressed in terms of the photo-ionization cross section of the intervening gas, so that $\kappa = \sigma_{\rm pi} n_{\rm abs}(1-x)$, where $\mathrm{n_{abs}}$ is the number density of the absorbing gas and x is the ionization fraction. A pure atomic hydrogen gas is assumed for simplicity, thus $n_{\rm abs} = n_{\rm H}$. With this assumption, the photo-ionization cross section depends only on the frequency of the absorbed radiation (Draine, 2011):

$$\sigma_{\rm pi} = 6.304 \times 10^{-18} \left(\frac{E_H}{h\nu}\right)^4 \frac{\exp(4(1-\arctan(\eta)/\eta))}{1-\exp(-2\pi/\eta)} \ \mathrm{cm}^2 \tag{3.2.4}$$

with

$$\eta = \sqrt{\frac{h\nu}{E_H} - 1} \ . \tag{3.2.5}$$

E_H=13.6 eV is the binding energy of a hydrogen atom, and ν is the frequency of the photon species. I choose to solve the equation of radiative transfer in my simulations directly through ray tracing. With this approach, the attenuation law (Eq. 3.2.3) can be rewritten explicitly in terms of photon number as (Abel et al., 1999):

$$N_{\gamma,f} = N_{\gamma,i} \exp(-\sigma_{\mathrm{pi}} n_H (1-x)\Delta s), \tag{3.2.6}$$

where $N_{\gamma,i}$ and $N_{\gamma,f}$ are the initial and final photon numbers along a line of sight, respectively.

3.3 Adaptive Ray Tracing

Radiation in these models is treated using long ray characteristics. To prevent oversampling near the source, and to minimize computational cost, I use an adaptive ray tree based on the "nested" geometry of the Hierarchical Equal Area isoLatitude Pixelation (HEALPix, Górski et al. (2005)). The use of the HEALPix geometric library is useful for two reasons. First, this discretization provides equally spaced cell centers (from which rays extend) with equal areas on a spherical surface. This is particularly important for ensuring spherical symmetry of the radiation field. Second, the HEALPix geometry is hierarchical. Using the built-in "nested" scheme, a quad-tree of rays can easily be created and applied to the ray tracing problem as first demonstrated by Abel & Wandelt (2002).

To build the ray tree for a computational mesh, 12 rays emanate from the radiation source using the angles supplied by the coarsest discretization of HEALPix. The rays extend radially outward until the ray density is too low to uniformly sample the grid. The radius at which this occurs is determined by imposing a restriction on the ratio of the number of grid cells on the surface of a sphere of radius R_{max} to the number of rays at a HEALPix resolution level, l, so that

$$R_{\mathrm{max}} = \sqrt{\frac{N_{\mathrm{rays}} \Delta x^2}{f 4\pi}}\ , \tag{3.3.1}$$

where $N_{\mathrm{rays}} = 12 \times 4^l$ is the number of rays, Δx is the cell spacing, and f is the covering factor that approximately sets the minimum number of rays that must trace each cell. When rays reach the termination condition (Eq. 3.3.1), they split into four child rays. In the HEALPix nested scheme, the child address is attained by adding 1 to the parent level and using a bit shift operation on the ray index to get the child indices (i.e. $i_{child} = i_{parent} << 2 + j$; j=0, 1, 2, 3). The direction of propagation is determined using the PIX2ANG_NEST function of HEALPix. Each child ray continues outwards along the directions prescribed by HEALPix until the termination condition (Eq. 3.3.1) is met. This process continues until the grid is completely sampled by the ray tree. The only free parameter involved with the ray trace is the covering factor, f. An increase in the value of f corresponds to a more uniformly sampled radiation field; however, the number of rays required to trace the mesh and the computational cost also increase. Previous implementations of radiative transfer into grid-based hydrodynamic codes recommend either a minimum value of f=2 (Abel & Wandelt, 2002) or $f \geq 3$ (Krumholz et al., 2007; Wise & Abel, 2011), though I find that the appropriate value of f is problem specific.

3.3.1 Ray-Tree Initialization

I have implemented two ray-tracing techniques into ATHENA. The first uses a "fixed" ray tree that is predetermined prior to the radiation cycle. The second uses an "active" ray tree which is built during integration. I detail the fixed ray tree approach here and discuss active ray tree integration in § 3.7.1.

At the beginning of every simulation, the geometry of the ray-tree follows from the restrictions imposed by the computational mesh. For each ray, the code stores the initial and final position, the initial and final photon numbers, the propagation angles prescribed by HEALPix, the unit vector in the propagation direction, the HEALPix resolution level, and the HEALPix ray index. Similar to Wise & Abel (2011), rays are organize into a doubly-linked list such that each ray stores the address for both the parent ray and all four child rays. Parallel jobs only account for rays which are either completely or partially contained within the local boundary assigned to each processor. Because of this, processors are outfitted with a look-up table for all included rays, for which a local ray index (i_{local}) is assigned to each ray. Processor addresses for any required communication are determined and stored for each

Table 3.1: Components of the ray structure

Variable	Description	Type
l_r	HEALPix resolution level	long
i_r	nested HEALPix ray index	long
$i_{\rm local}$	ray ID on local processor	int
θ	polar propagation angle	double
ϕ	azimuthal propagation angle	double
$\hat{u}$	unit vector (propagation direction)	double
$\vec{x}_i$	initial position	double
$\vec{x}_f$	final position	double
$N_{\gamma,i}$	incident photon flux (per species)	double
$N_{\gamma,f}$	attenuated photon flux (per species)	double
PARENT	address of parent ray	ray structure
CHILD	address of child rays (four elements)	ray structure
PASSIN_ID	ID to receive from	int
PASSOUT_ID	ID to send to	int
PASSIN_PID	ID to receive from (for off grid parent rays)	int
PASSOUT_CID	ID to send to (for off grid child rays)	int
N_{cc}	number of cells crossed	int
Δs_{cc}	path lengths of cells crossed	double
$(i,j,k)_{cc}$	grid indices of cells crossed	int

ray (§ 3.3.2). To accelerate the ray tracing step, a list of the cells crossed by each ray and the path lengths through those cells are also stored (Krumholz et al., 2007). A list of the elements in each ray structure is listed in Table 3.1.

A major difference between the implementation of radiation into ATHENA in previous work (Krumholz et al., 2007) and that which is detailed here is that I require the use of grid refinement, thus I have designed this routine to be compatible with the SMR framework native to the ATHENA code. In this case, the ray tree structure is complicated by non-uniform cell spacing on the computational mesh. Yet, the mesh geometry is completely determined and known to all processors, thus initial and final positions of every ray structure are calculated independently on each processor. I follow the procedure shown in Figure 3.1 to compute the maximum radius for each ray. In short, the routine determines if a ray crosses into higher refinement zones. If so, the radius given by Eq. 3.3.1 is imposed if it causes the ray to terminate, or spawn a child ray, on the refinement level. If not, the algorithm determines the minimum requirement for $R_{\rm max}$ to satisfy the termination condition across all lower refinement zones, if any, that it crosses. To improve the efficiency of this process, the algorithm uses an optimized ray-box intersection method (Williams et al., 2005) to determine if rays intersect SMR domain blocks. The minimum radius of each ray is computed by applying this algorithm on the ray's parent. The process outlined in Figure 3.1 functions irrespective of source position on the mesh; however, I have assumed that only one refinement hierarchy exists. This can be easily extended to included multiple refinement zones, though it is unnecessary for the work presented here.

Once the starting and terminating positions of each ray have been determined, the ray is cropped to the boundary of the local grid. I do this by computing a grid walk (Abel et al., 1999) along each ray until one of the following conditions is met: (i) The ray exits the local grid boundary within the mesh (ii) the ray is overlapped by a refinement zone (iii) the ray passes off of the computational mesh. In cases (i) and (ii), parallel communication is required.

3.3.2 Parallel Integration

To improve computational efficiency, Athena is parallelized using distributed memory. Because of this, individual processors are assigned to a unique portion of the computational mesh, only communicating boundary value information to neighboring processors in a predetermined number of synchronized message passing interface (MPI) messages. Such a parallelization framework cannot be applied to the ray tracing problem because of its sequential nature. Furthermore, the geometry of the ray-tree often prohibits singular send-receive sequences between processors. In the simplest case, rays pass directly from one processor to the next, thus requiring a single MPI communication. Yet, there are instances in which rays terminate on a processor but spawn an "abandoned child" on a neighboring processor. In this case, sequenced communications across the processors are required. These ray passing geometries are illustrated in Figure 3.2.

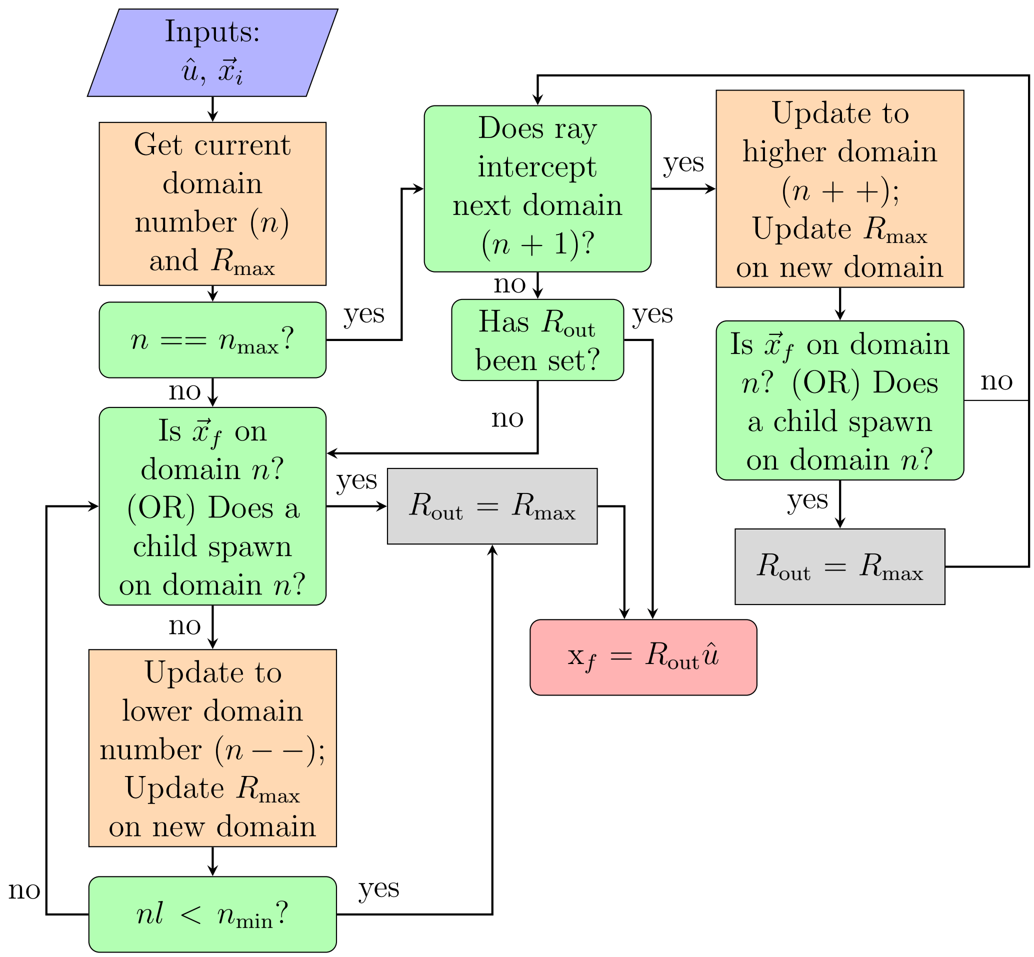

Figure 3.1: Flowchart for determining the ray termination radius with arbitrarily placed refinement. The initial ray position and the unit vector in the direction of propagation are supplied. If the ray intercepts refinement domains, the algorithm determines if the termination condition (Eq. 3.3.1) must be imposed. If not, the process is repeated to determine the maximum radius that meets the termination condition for lower refinement domains.

.

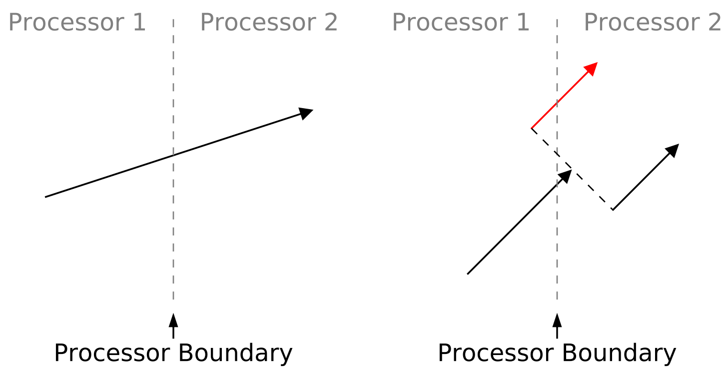

Figure 3.2: Left: Simple ray passing geometry in which a single ray passes a processor boundary, thus requiring a single parallel communication. Right: The "abandoned child" ray passing geometry in which a parent terminates on processor 2, but spawns a child (red) on processor 1, resulting in additional parallel communication between processors.

Because of the sequential nature of the ray-tracing problem, it is imperative to use a parallelization strategy that efficiently dispatches MPI messages as they become available. To this end, the parallelization should be maximally non-blocking. Previous (Wise & Abel, 2011) and recent (Rosen et al., 2017) work on grid-based radiative hydrodynamic models use book-keeping methods that aggressively drain messages between processors sequentially but asynchronously, thus requiring many messages for each processor. This general approach, which this routine roughly follows, provides the best possible performance for ray tracing with distributed memory.

For rays which require simple message passing between processors, the code determines the processor ID, if any, from which the ray will be sent during integration (PASSIN_ID). Similarly, the code determines and stores the processor ID to which any rays will continue after the local grid walk (PASSOUT_ID). To handle "abandoned child" rays, each ray stores a 4 element list of processor IDs to which the parent ray will send in these cases (PASSOUT-CID). Each abandoned child is provided with the processor ID from which their parent ray will post a message (PASSIN-PARID). Although every ray is initialized with these variables, they are only activated in the event that MPI communication is needed. A simple ray-by-ray parallelization check is completed after each ray tree build to guarantee that all MPI messages will be matched across processors. ID's for incoming messages are only used for this validation test.

Ray-tree integration begins on the host processor for the radiation source. The photon number (emission rate $\times$ radiation timestep) is split between twelve rays corresponding to the coarsest discretization of HEALPix. These rays are integrated following the procedure outlined in Figure 3.3. Each ray is walked through a list of pre-determined cells on the grid through which photon deposition occurs (Eq. 3.2.6). Terminating rays pass a quarter of the remaining photon number to child rays. Rays requiring parallel communication must post non-blocking messages containing the photon flux and ray address to the corresponding processor ID before continuing the integration locally. This process continues along a given branch of the ray tree until all rays extending from the root ray have been integrated and all necessary MPI messages have been posted.

Simultaneous integration of branches across multiple processors is accomplished using the procedure seen in Figure 3.4 in conjunction with the ray walk described above. If a ray reaches a processor boundary, a non-blocking send, or MPI_ISEND, is posted for the processor listed in PASSOUT-ID. Once the message is posted, the linked-list sequence can be continued on the local processor. Each processor that will post send requests is initialized with an array equal to the total size of the messages that will be sent. All messages are sent as elements of this list. After a message is sent, the list is incremented to the next position to avoid overwriting the buffer before it is read by the receiving processor. For each processor, the number of rays that will be received is completely predetermined. For these processors, a while loop is initiated so that the processor posts sequential blocking MPI_RECV messages until the number of expected receives occurs. The flag MPI_ANY_SOURCE is used so that messages are received as they become available. Once a ray is received, it behaves as a new root ray from which an entire branch of rays stem. The branch is integrated in the linked-list sequence, posting MPI message sends as ray segments terminate at the grid or SMR boundaries. This process continues until all expected sends and receives are completed across the entire mesh. To free the memory used for MPI communication and check for

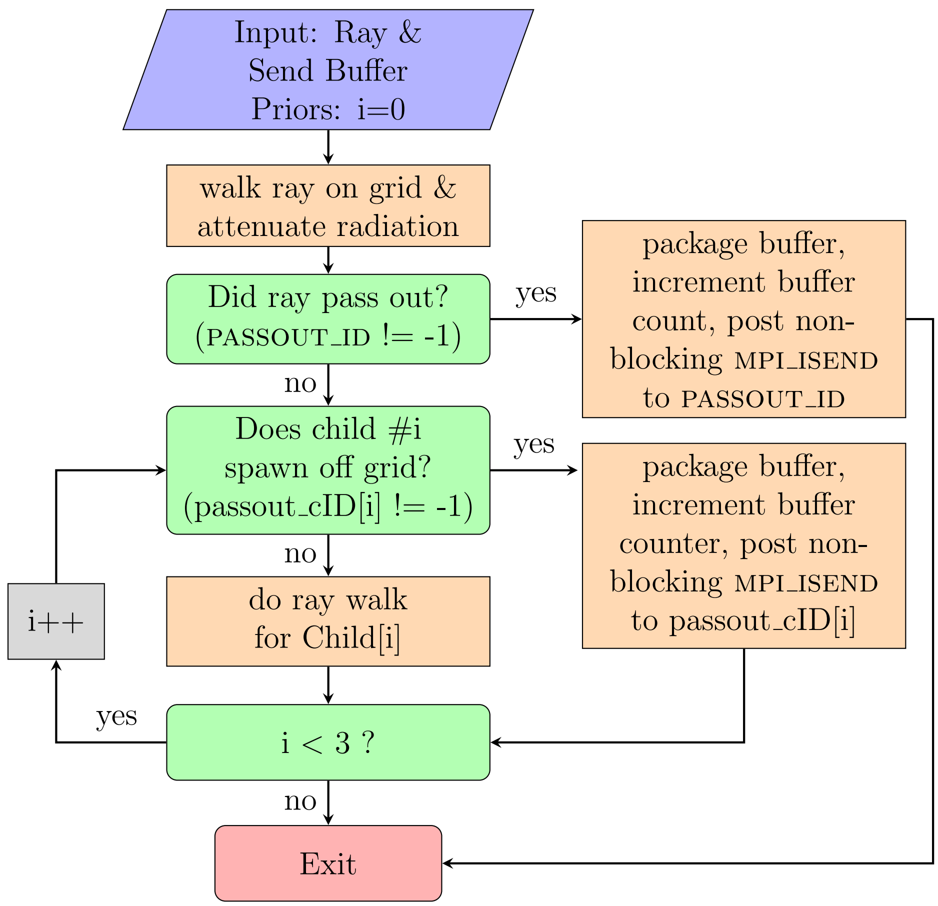

Figure 3.3: Flowchart for an individual ray walk. Rays traverse grid cells until they reach a final position on the local processor. The ray may (i) split into four child rays, (ii) pass off of the local grid by posting ray information using a non-blocking MPI send, (iii) spawn a child off of the grid using a non-blocking MPI send, or (iv) terminate off of the mesh.

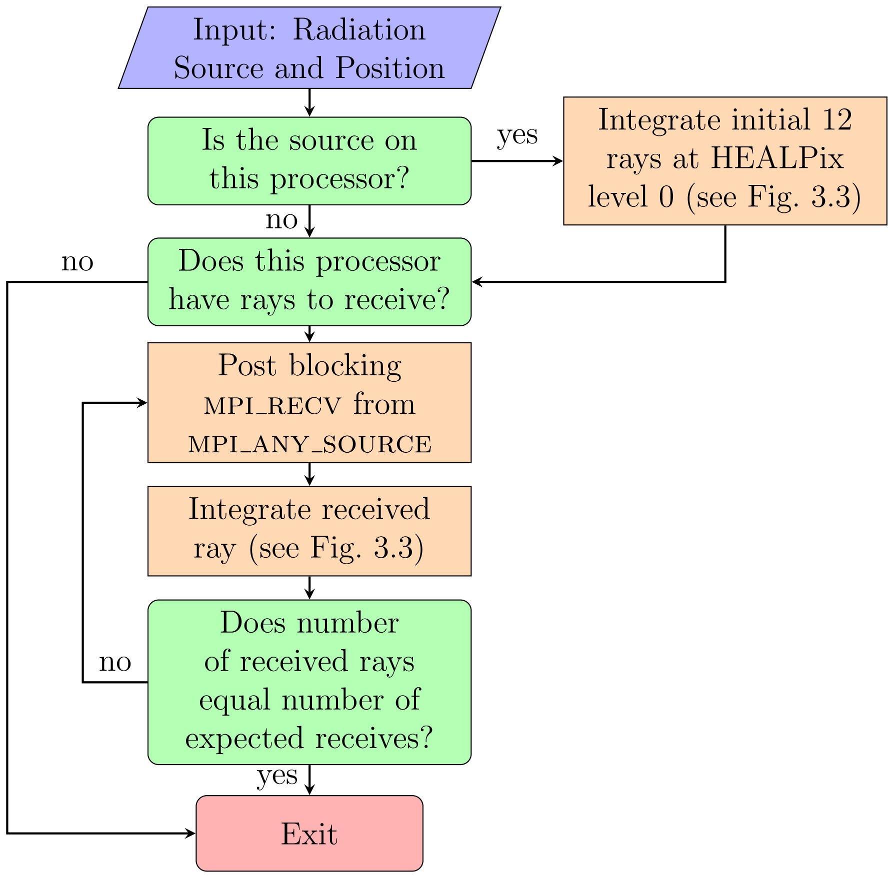

Figure 3.4: Flowchart for parallelized ray integration with a fixed geometry ray tree. The host processors initializes integration from the source. All processors use blocking MPI receives to accept incoming messages until all messages are received.

failed messages, the code loops over all messages using MPI_WAIT.

3.4 Included Physics

3.4.1 Photo-absorption

As rays pass through grid cells, ionization and heating takes place as a consequence of photon deposition. Assuming a zero emissivity gas and a constant absorption cross section, the radiative transfer equation can be solved through each ray segment using Eq. 3.2.6. The photon deposition is calculated by taking the difference between the initial and final photon numbers:

$$\delta N_\gamma = N_{\gamma,i}(1 - e^{-\sigma_{\mathrm{pi}} n_H (1-x)\Delta s}) \ . \tag{3.4.1}$$

Multiple rays will pass through each cell, thus, for every cell, I also assign a sub-volume to each ray for two reasons. First, the radiation update is calculated using a forward Euler method, thus this helps to avoid overshoots in ionization and heating. Second, this helps to improve spherical symmetry by partially absorbing photons from rays that only cross a small fraction of a cell's volume. For each ray segment, I calculate the weighting factor:

$$W_{\rm r} = \frac{\Delta s}{4^l} \tag{3.4.2}$$

where Δs is the path length and l is the resolution level on the ray tree. The sub-volume associated with each ray is then given by the fraction

$$V_{\rm sub} = \frac{W_{\rm r}}{\sum_i W_{\rm r,i}} \tag{3.4.3}$$

where the denominator is the sum of weights for all rays passing through the cell. I limit the photo-absorption in each sub-volume so that the number of photons does not exceed the number of neutral atoms after both collisional ionization and recombination have been accounted for.

3.4.2 Ionization and Recombination

To conserve photon number, each ionizing photon ($E_\gamma > 13.6$ eV) must be exchanged exactly for one ionization of a hydrogen atom, thus the ionization rate is

$$I_i = \frac{\delta N_{\gamma,i}}{\Delta t_\gamma \Delta x^3} \ , \tag{3.4.4}$$

where $\delta N_{\gamma,i}$ is the number of photons deposited per photon species, i. Δt_γ is the radiation timestep, and Δx^3 is the cell volume assuming a uniform aspect ratio for grid cells. In addition to photo-ionization, collisional ionization occurs such that:

$$I_{\mathrm{coll}} = k_{\mathrm{coll}} n_{\mathrm{HI}} n_{\mathrm{HII}} \ , \tag{3.4.5}$$

where the collisional ionization rate coefficient is (Tenorio-Tagle et al., 1986):

$$k_{\mathrm{coll}} = 5.84 \times 10^{-11} \sqrt{\frac{T}{K}} e^{-E_H/k_B T} \mathrm{cm}^3 \mathrm{s}^{-1} \ . \tag{3.4.6}$$

n_{HI} is the neutral hydrogen number density, n_{HII} is the ionized hydrogen number density, T is the temperature of the gas, and k_B is the Boltzmann constant. The net ionization which enters into Eq. 2.2.9 is the sum of collisional and radiative ionization terms:

$$I = \sum_i I_i + I_{\mathrm{coll}} \ , \tag{3.4.7}$$

where the sum is over contributions from each photon species. For photon energies in excess of the the binding energy of hydrogen, photo-heating also occurs so that

$$\Gamma_{\mathrm{ion},i} = \frac{\delta N_{\gamma,i} (E_{\gamma,i} - E_H)}{n_H \Delta x^3 \Delta t_\gamma} \tag{3.4.8}$$

is the volumetric heating rate due to photo-absorption, where $E_{\gamma,i}$ is the energy of the photon species i.

For recombination of ionized hydrogen, I assume the so-called "on-the-spot" approximation in which recombinations of hydrogen atoms to the ground state emit ionizing photons that are immediately re-absorbed by the intervening medium (Osterbrock, 1989). Recombinations to excited states of hydrogen are assumed to emit photons to which the surrounding gas is optically thin. Under this assumption, the recombination rate is then

$$R = \alpha_B(T) n_e n_p \ , \tag{3.4.9}$$

where $\alpha_B(T)$ is the recombination coefficient for case B recombination,

$$\alpha_B(T) = 2.59 \times 10^{-13} (\frac{T}{10^4 K})^{-0.7} \text{cm}^3 \ \text{s}^{-1} \ , \tag{3.4.10}$$

and n_e and n_p are the electron number density and free proton number density of the gas. Under the assumption of a pure hydrogen gas, the number density of electrons is equal to the number density of free protons. Therefore, the recombination rate in this case simplifies to $R = n_{\text{HII}}^2 \alpha_B(T)$.

3.4.3 Compton Heating

In the X-ray regime, Compton heating from the scattering with free electrons leads to photon energy loss rather than absorption. Because the radiative transfer scheme uses monochromatic photon bins, a change in photon energy is not possible. Instead, I follow the approach of Kim et al. (2011) by proportionally decreasing the photon number flux to account for energy loss from Compton scattering such that,

$$\delta N_{\gamma,\text{comp}} = N_\gamma (1 - e^{-\tau_e}) \Delta E(T_e)/E_\gamma \ . \tag{3.4.11}$$

where N_γ is the incident number of photons. $\tau_e \approx n_e \sigma_{KN} \Delta s$ is the optical depth where n_e is the electron number density ($n_e = n_{\text{HII}}$), and σ_{KN} the Klein-Nishina cross section (Rybicki & Lightman, 1979). For the non-relativistic energies considered in this work, $\sigma_{KN} \approx \sigma_T$, where σ_T is the Thomson scattering cross section. The energy lost in a Compton scattering event is

$$\Delta E(T_e) = 4 k_B T_e \frac{E_\gamma}{m_e c^2} \ , \tag{3.4.12}$$

where m_e is the electron mass, and T_e is the electron temperature. The Compton heating term is written in the same manner as heating from photo-absorption:

$$\Gamma_{\rm comp} = \frac{\delta N_{\gamma,\rm comp} E_\gamma}{n_e \Delta x^3 \Delta t_\gamma} \ . \tag{3.4.13}$$

3.4.4 Secondary Ionizations

Photons with energies much greater than the ionization potential of hydrogen ($E_\gamma >$ 100 eV) may result in the ionization of multiple hydrogen atoms. The fractional amount of energy allotted to the heating and ionization of hydrogen for these photons has been calculated as (Shull & van Steenberg, 1985):

$$Y_\Gamma = 0.9971 * (1 - x^{0.2663})^{1.3163} \tag{3.4.14}$$

$$Y_{\rm H,ion} = 0.3908 * (1 - x^{0.4092})^{1.7592} \ , \tag{3.4.15}$$

where $x = n_{\rm HII}/n_H$ is the ionization fraction of the gas. It should be noted that the models presented here do not explicitly include helium, though the energy fraction which contributes to ionization of helium is:

$$Y_{\rm He,ion} = 0.0554 * (1 - x^{0.4614})^{1.6660} \ . \tag{3.4.16}$$

I include ionization of helium atoms implicitly such that the total fractional energy input for ionization is $Y_{\rm ion} = Y_{\rm H,ion} + Y_{\rm He,ion}$. The net ionization and heating rates are then given as:

$$I_{\rm i,secondary} = \frac{\delta N_{\gamma,i} E_{\gamma,i}}{\Delta t_\gamma \Delta x^3} \frac{Y_{\rm ion}}{I_H} \tag{3.4.17}$$

$$\Gamma_{\rm i,secondary} = \frac{\delta N_{\gamma,i} E_{\gamma,i}}{\Delta t_\gamma \Delta x^3} Y_\Gamma \tag{3.4.18}$$

For X-ray photons, these ionization and heating terms are substituted in for the standard absorption into Eqs 3.4.7 and 3.4.8.

3.4.5 Radiation Pressure

As photons are absorbed by the intervening medium, they transfer momentum ($p_\gamma = E_\gamma/c$) to the gas. The force due to this process is

$$F_{\gamma,i} = \frac{\delta N_{\gamma,i} E_{\gamma,i}}{c\Delta t_\gamma}\ \hat{r}\ , \tag{3.4.19}$$

where the index i is again used to distinguish monochromatic photon bins. Momentum injection is aligned with the photon propagation direction, denoted by $\hat{r}$. The resulting acceleration that enters into Eqs. 2.2.7 and 2.2.8 is determined by dividing by the cell mass:

$$\mathbf{a}_\gamma = \frac{\delta N_{\gamma,i} E_{\gamma,i}}{c\rho\Delta t_\gamma \Delta x^3}\ \hat{r} \tag{3.4.20}$$

I apply the source terms for radiation pressure (Eqs. 2.2.7 and 2.2.8) in tandem with the ionization and thermal updates in the radiation routine. An alternative approach would be to include this term in the hydrodynamic integration step with other force terms (Wise & Abel, 2011). Despite this simplification, the implementation here shows good agreement with theory as demonstrated in § 3.6.4.

3.4.6 Thermal Physics

The heating term, G, in Eq. 2.2.8 is the sum of heating from ionization, heating of neutral gas from a constant background radiation field, and heating from Compton scattering:

$$G = n_{\rm HI}\Gamma_{\rm amb} + \sum n_H \Gamma_{ion,i} + n_e \Gamma_{\rm comp} \tag{3.4.21}$$

I take the background heating term to be (Koyama & Inutsuka, 2002)

$$\Gamma_{\rm amb} = G_0\ (2\times 10^{-26})\ {\rm erg\ s^{-1}}\ . \tag{3.4.22}$$

In simulations accounting for the strong interstellar radiation field of the Galactic Center (GC), $G_0 = 1000$ (Clark et al., 2013). Otherwise, $G_0 = 1$ for radiation fields typical of the interstellar medium (ISM). To avoid overheating of low density gas, a hyperbolic tangent function smoothly drives the ambient heating term to 0 for temperatures above 10^4 K. For neutral gas, I use a modified version of the cooling function from Koyama & Inutsuka (2002):

$$\begin{aligned}\Lambda_{\mathrm{n}} = 2 \times 10^{-26} \Big(& 10^7 \exp \frac{-118400}{T + 1000} \\ & + 0.014 \sqrt{T} \exp \frac{-92\ \beta_{\mathrm{T}}}{T} \Big) \text{ erg s}^{-1} \text{ cm}^3 \ .\end{aligned} \tag{3.4.23}$$

For standard ISM conditions $\beta_T = 1$, but the value of β_{T} can be modified to approximate the effect of cosmic rays which penetrate deep into dense gas structures, effectively setting a baseline temperature (Goldsmith & Langer, 1978). The cosmic ray ionization rate in the GC is roughly a thousand times greater than the solar neighborhood (Clark et al., 2013) which results in a minimum gas temperature of ≈ 100 K (Wolfire et al., 1995; Papadopoulos et al., 2011). To approach this minimum temperature smoothly in simulations with GC conditions, I set $\beta_{\mathrm{T}} = 10$.

For ionized gas, recombination cooling and free-free cooling are given by (Osterbrock, 1989):

$$\Lambda_{\mathrm{rec}} = 8.418 \times 10^{-26} T^{0.11} \text{ erg s}^{-1} \text{ cm}^3 \tag{3.4.24}$$

$$\Lambda_{\mathrm{ff}} = 1.427 \times 10^{-27} 1.3 \sqrt{T} \text{ erg s}^{-1} \text{ cm}^3 \tag{3.4.25}$$

I also follow the treatment for collisionally excited radiation in Osterbrock (1989), but reduce the resulting cooling rate to a piecewise approximation which incorporates trace amounts of

NII, NIII, OII, OIII, NeII, and NeIII (see § 3.4.7):

$$\Lambda_{\rm CLE}(T) = \begin{cases} 3.47 \times 10^{-29}\ T^{1.915} & 0 < T < 10^2 \\ 2.34 \times 10^{-26}\ T^{0.500} & 10^2 < T < 10^{2.8} \\ 1.11 \times 10^{-24}\ T^{-0.099} & 10^{2.8} < T < 10^{3.6} \\ 1.08 \times 10^{-32}\ T^{2.127} & 10^{3.6} < T < 10^4 \\ 2.67 \times 10^{-30}\ T^{1.529} & 10^4 < T < 10^{4.5} \\ 1.74 \times 10^{-24}\ T^{0.237} & 10^4.5 < T < 10^5 \\ 1.10 \times 10^{-21}\ T^{-0.323} & 10^5 < T < 10^6 \\ 7.49 \times 10^{-21}\ T^{-0.462} & 10^6 < T \qquad\qquad \text{ergs s}^{-1}\text{ cm}^3 \end{cases} \tag{3.4.26}$$

The net cooling rate is a combination of all cooling terms:

$$L = n_{\rm HI}^2 \Lambda_{\rm n}(T) + n_{\rm HII}^2(\Lambda_{\rm ff}(T) + \Lambda_{\rm rec}(T) + \Lambda_{\rm CLE}(T)) \tag{3.4.27}$$

Unlike Krumholz et al. (2007), I do not restrict cooling in mixed-ionization cells, as this led to incorrect propagation speeds in the HII region expansion test (see § 3.6.3).

3.4.7 Collisionally Excited Line Radiation

I include the effect of optically thin line radiation resulting from collisional excitation of OII, OIII, NII, NeII, and NeIII. Following the treatment of Osterbrock (1989) for conditions typical of HII regions, I assume abundances of $X_{\rm Ne} = 7 \times 10^{-4}$, $X_{\rm N} = 9 \times 10^{-5}$, and $X_{\rm O} = 7 \times 10^{-4}$ with 80% of each species in the singly ionized state and 20% in the doubly ionized state. For each species, I solve the collisional equilibrium equation:

$$\sum_{j \neq i} n_j n_e q_{ji} + \sum_{j>i} n_j A_{ji} = \sum_{j \neq i} n_i n_e q_{ij} + \sum_{j<i} n_i A_{ij} \tag{3.4.28}$$

where A_{ij} is the spontaneous transition rate from level i to level j, and q_{ij} is the collisional excitation rate from level i to level j. The spontaneous transition rates and collision strengths, which are used in the calculation for the collision rate, are taken from Osterbrock (1989). I assume that the hydrogen gas is fully ionized so that $n_e = n_H$. This proves to be a fair assumption given that the resulting cooling rate peaks in a temperature regime where the gas is expected to be fully ionized.

To calculate an approximate cooling function, I solve Eq. 3.4.28 for densities in the range $10^{-2} - 10^{6}\mathrm{cm}^{-3}$ and in the temperature range of $1 - 10^{8}K$. For each species, cooling rates are calculated at collisional excitation equilibrium. The net cooling is the sum of these terms:

$$\Lambda_{\mathrm{CLE}} = \sum_s \sum_i n_{i,s} \sum_{j<i} A_{ij} h\nu_{ij} \tag{3.4.29}$$

where the first sum is over species, $n_{i,s}$ is the number density of an excited state for a given species, and ν_{ij} is the frequency of the photon emitted by a transition from level i to level j. The resulting cooling rates are temperature and density dependent and can be placed into a look-up table as seen in Figure 3.5a. To use this table in simulations, a bi-linear interpolation must be computed for every cell at every iteration in the thermal physics routine. Thus, in the interest of computational efficiency, I average the two dimensional table of cooling rates along the density axis to get an approximate cooling rate solely as a function of temperature. The average cooling function is show in Figure 3.5b alongside the piecewise fit used in my models.

3.5 Radiation Timestep

Radiation is included into my simulations in an operator split fashion. Because the radiation timescale is typically shorter than the dynamical timescales, the radiation module sub-cycles on the hydrodynamic timestep but is limited to ten sub-cycles. In the event that the hydrodynamic timestep is greater than the radiation timestep, the radiation is

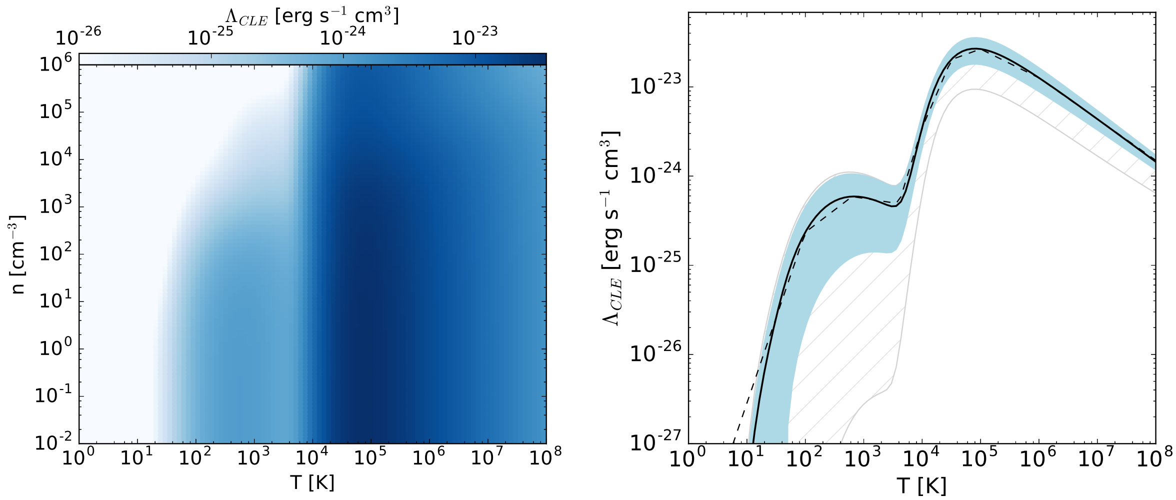

(a) (b)

Figure 3.5: (a) Cooling rate for collisionally excited line radiation. (b) Average cooling rate (solid) and piecewise approximation (dashed) for collisionally excited line radiation. The shaded region indicates the 1σ error region, and the hatched region highlights the range of values averaged over to obtain the approximate cooling function.

calculated for the hydrodynamic timestep without sub-cycling. For each radiation cycle, the radiation routine follows the structure seen in figure 1 of Krumholz et al. (2007) with the exception that the timestep is calculated at the beginning of the cycle based on the conditions from the previous timestep. As a consequence, this scheme is more akin to Wise & Abel (2011) in which photon numbers are traced rather than rates. This helps to alleviate stringent restrictions imposed on the radiation time step, helping to mitigate computational cost. It should be noted, however, that this formalism does not preserve the time derivative of the radiative transfer equation. This assumption is fair given that the light transit time in my simulations is much shorter than the dynamical timescale of the gas ($\tau_\gamma/\tau_{\rm gas} \propto |\mathbf{v}_{\rm max}|/c < 10^{-2}$). By eliminating time dependence of the radiation, shielding occurs instantaneously because optical depths are calculated to all grid cells at every radiation time step. Similarly, the radiation field ionizes newly exposed gasses without delay.

For each radiation cycle, the radiation timestep at each cell is

$$\Delta t_\gamma = \alpha \frac{x}{|dx/dt|} = \alpha \frac{n_{\rm HII}}{|I - R|} \tag{3.5.1}$$

where x is the ionization fraction, and α is the maximum fractional change allowed in the ionization fraction. Unless otherwise noted, I set $\alpha = 0.1$ for my models. The radiation timestep is taken to be the minimum value across all cells in the simulation. To avoid prohibitively small timesteps, grid cells with $0.1 < x$ are excluded from this timestep calculation. I enforce a minimum timestep of $\Delta t_\gamma = \frac{\Delta x}{c}$, which is the light crossing time of a cell (Wise & Abel, 2011).

3.5.1 Sub-cycling Thermal Physics

To avoid over-cooling, ionization and thermal effects are included in a sub-cycle within each radiation timestep. After the heating rates are calculated, the sub-cycle timestep is determined via:

$$\Delta t_{\rm sub} = \min\left(\beta \frac{e}{|\frac{de}{dt}|_{\rm therm} + |\frac{de}{dt}|_\gamma}, \Delta t_\gamma\right) \tag{3.5.2}$$

where $\beta = 0.1$ is the maximum fractional change allowed in the internal energy. Changes in internal energy due to the thermal physics prescription are calculated at each step in the sub-cycle; however, changes due to radiation are assumed to be linear throughout this process. The internal energy is updated using a forward Euler step,

$$e_f = e_i + \left(\left(\frac{de}{dt}\right)_{\rm therm} + \left(\frac{de}{dt}\right)_\gamma\right)\Delta t_{\rm sub}\ . \tag{3.5.3}$$

The ionization fraction, which affects the cooling rate, is updated as

$$x_f = x_i + \frac{dx}{dt}\Delta t_{\rm sub} \tag{3.5.4}$$

Table 3.2: Ionization Tests Refinement Geometry

Level	Dimensions pc	x_0 pc	y_0 pc	z_0 pc	Resolution
1	20×20×20	-10	-10	-10	64×64×64
2	11.25×10×10	-1.25	-5	-5	72×64×64
3	10×5×5	0	-2.5	-2.5	128×64×64

The sub-cycle continues until the total time equals the radiation timestep.

3.6 Test Cases

To demonstrate the fidelity of this radiative transfer method, I begin with three tests originally used for the first implementation of ray-tracing in the Athena code (Krumholz et al., 2007). These test cases differ for two reasons: (1) The radiation module used for this work is compatible with the refinement framework of Athena, therefore it is important to demonstrate agreement across grids of varying resolution. (2) The time-step restrictions imposed by the radiation routine have been designed to maximize the radiation timestep, thus lowering the computational cost. In addition, I present a new test for my implementation of radiation pressure, and demonstrate the effect of including secondary ionizations from high energy X-rays.

For all tests, I use the mesh and refinement geometry shown in Table 3.2. The radiation source in all cases is placed at the origin of a box which extends 10 pc in each direction. Two elongated SMR regions enclose the radiation source, and extend in the $+\hat{x}$ direction to the edge of the computational mesh. I use a covering factor of $f = 3$ and use a single rotation of the HEALPix sphere to minimize strong grid effects (see § 3.7.2).

3.6.1 R-type Ionization Front without Recombination

As a first test, I follow the evolution of an R-type ionization front (I-front) as detailed in Krumholz et al. (2007). This test imitates the initial evolution of an I-front in which rapid ionization occurs prior to the onset of either recombination or hydrodynamic response to photo-heating. To disable photo-heating, I set $E_\gamma = 13.6$ eV, equal to the binding energy of hydrogen. I also set the recombination coefficient to $\alpha_B = 0$ and disable radiation pressure. This test serves as an excellent probe of photon conservation as the number of ionizations that occur must exactly match the number of photons emitted. By equating these two quantities, an analytic expression for the radius of the expanding I-front is found to be

$$R_{nr}(t) = \Big(\frac{3Q_\star t}{4\pi n_H}\Big)^{1/3} , \tag{3.6.1}$$

where $Q_\star$ is the rate of ionizing photon emission, n_H is the density of the ambient hydrogen gas, and t is the elapsed time. I adopt parameters identical to those in § 4.1 of Krumholz et al. (2007) so that $n_H = 100$ cm^{-3} and $Q_\star = 4.0 \times 10^{49}$ s^{-1}. The gas within the box is initially neutral with an arbitrarily set temperature of 10 K. In Figure 3.6a, I show a midplane slice of the ionization fraction for the model of the R-type I-front without recombinations at t= 3 kyr. The extent of the ionzation front is uniform across refinement levels, demonstrating that the SMR treatment is sound. To calculate the numerical value of R_{nr}, I compute the average radius to cells with ionization fractions between 1% and 99% which marks the edge of the I-front. In Figure 3.6b, I show a comparison of the numerically calculated I-front radius with the analytic solution. When the radius of the I-front is sufficiently large ($R_{nr} >$ few $\times\, \Delta x$), the error between the numerical result and analytic result is on the order of 0.1%, which corresponds to an error less than that of a cell size on the highest refinement level (i.e. $\Delta x/R_{nr} \approx 1\%$). I conclude that my radiative transfer code is sufficiently photon conservative.

3.6.2 R-type Ionization Front with Recombination

Following Krumholz et al. (2007), I repeat the previous test, but enable recombinations by setting $\alpha_B = 2.59 \times 10^{13}$ cm^3 s^{-1}. Physically, this test simulates an early stage of I-front evolution prior to the onset of hydrodynamic response to the radiation field. An equilibrium

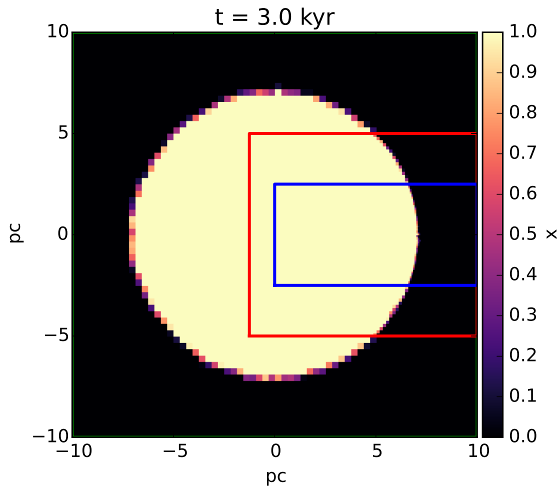

(a)

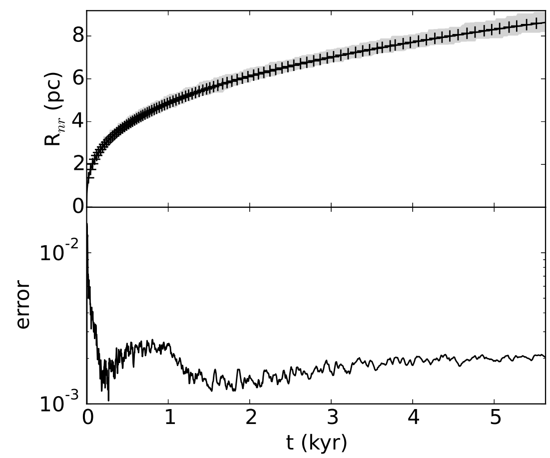

(b)

Figure 3.6: (a) Midplane slice of the ionization fraction in a simulation of an R-type I-front without recombination. The rectangular boxes outline the refinement domains used in this model. (b) Radius of the R-type ionization front without recombinations R_{nr} vs time t (top), and the error relative to the analytic solution (bottom). The top figure shows the analytical solution (solid line) and average radius of the I-front in the simulation (plus signs). The gray shaded region shows the range of all radii included in this average. The bottom panel shows the relative error of the radius in the simulation with respect to the analytic solution.

radius exists at which the photon emission rate perfectly balances the net recombination rate of gas interior to the I-front. By equating these two terms, the I-front radius has the analytic solution,

$$R_s = \left(\frac{3Q_\star}{4\pi\alpha_B n_H^2}\right)^{1/3}, \tag{3.6.2}$$

which is the well-known Strömgren radius (Strömgren, 1939). By assuming a constant recombination coefficient, the time dependence of the ionization front radius as it approaches the Strömgren radius also has the analytic solution

$$R_r(t) = R_s(1 - e^{-t/\tau_r})^{1/3} \tag{3.6.3}$$

Where τ_r is the recombination timescale:

$$\tau_r = \frac{1}{n_H \alpha_B} \tag{3.6.4}$$

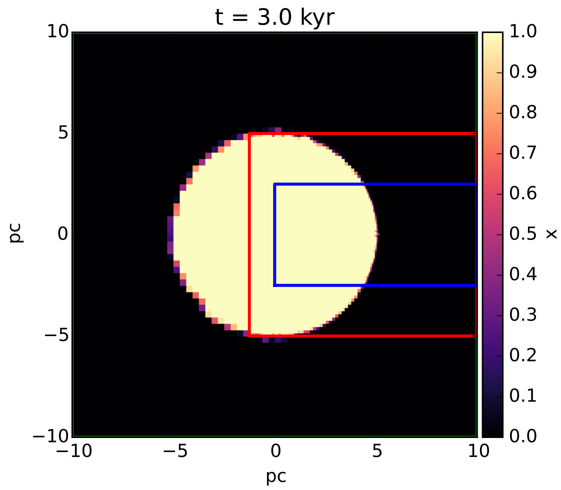

(a)

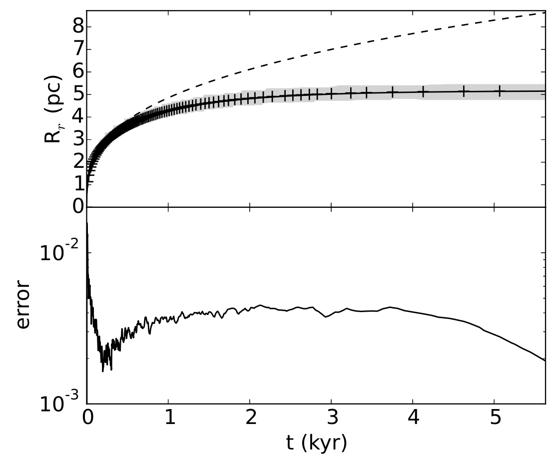

(b)

Figure 3.7: (a) Midplane slice of the ionization fraction in a simulation of an R-type I-front with recombinations. The rectangular boxes outline the refinement domains used in this model. (b) Radius of the R-type ionization front with recombinations R_r vs time t (top), and the error relative to the analytic solution (bottom). The top figure shows the analytical solution (solid line) and the numerical approximation to the radius of the I-front (plus signs). For reference, the dashed line shows the radius of the I-front without recombination. The gray shaded region shows the range of all radii included in this average. The bottom panel shows the relative error of the radius in the simulation with respect to the analytic solution.

For these initial conditions, the Strömgren radius is expected to be roughly 5pc. In Figure 3.7a I show the ionization fraction in a midplane slice through the computational domain at 3 kyr, which again demonstrates good agreement across refinement levels. The I-front radius is computed numerically in the same manner as the previous test. The result and the comparison to the analytic solution can be seen in Figure 3.7b. The error in the result obtained from the radiative transfer method is again on the order of 0.1%, demonstrating that my treatment of radiation in the absence of hydrodynamic response is theoretically sound.

3.6.3 D-type Ionization Front

To demonstrate that my code correctly models radiative-hydrodynamic coupling, I present an idealized model of a D-type I-front as seen in Krumholz et al. (2007). Here, gas surrounding the radiation source will be ionized and rapidly heated. Shortly after, the over-pressured gas will expand into the surrounding medium, sweeping up a dense gas shell whose radius can be approximated with the analytic solution (Draine, 2011):

$$R_i(t) = R_D\Big(1 + \frac{7}{4}\frac{c_s(t - t_D)}{R_D}\Big)^{-4/7} \tag{3.6.5}$$

where R_D and t_D are the radius and time at which the I-front transitions from R-type to D-type, and c_s is the isothermal sound speed of the ionized gas interior to the dense shell. On the timescales considered in this test problem, $t_D \approx 0$ and $R_D \approx R_s$. I use an identical simulation set-up as in the previous test cases; however, the photon emission rate is lowered to 4×10^{46} s^{-1}, corresponding to a Strömgren radius of $R_s = 0.5$ pc. I set the photon energy to be 16 eV to include photo-heating and enable both the temperature dependent recombination rate in Eq. 3.4.10 and the full thermal physics prescription (§ 3.4.6). The equilibrium temperature of ionized gas in this model is 5803K.

In Figure 3.8a, I show a midplane slice of the density in the simulation of the D-type I-front at 1.4 Myr, and in Figure 3.8b, I show the comparison between I-front radius computed in the simulation with the analytical value. To calculate the numerical value of R_i, I find the radius to all cells that meet two criteria. First, the cell must have $\rho < \rho_{\text{ambient}}$. Second, the nearest neighbor in the $+\hat{r}$ direction must have $\rho > 1.1 * \rho_{\text{ambient}}$. These two conditions mark the inner edge of the dense gas shell and the extent of the I-front. I take an average of the radii of cells in this sample. The error with respect to the analytical solution is $\approx 1\%$, corresponding roughly to a cell size on the highest refinement level.

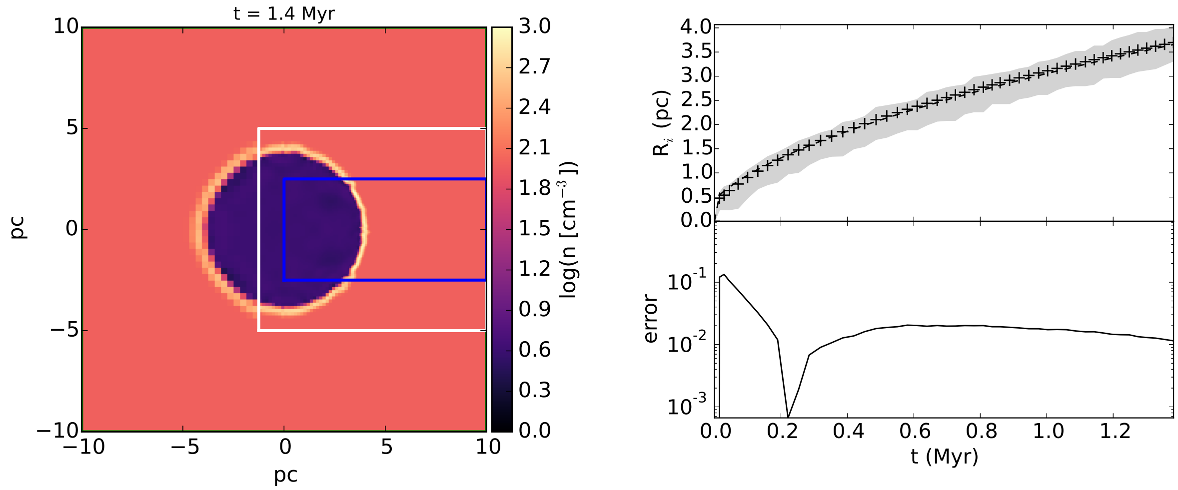

(a) (b)

Figure 3.8: (a) Midplane slice of the density in the simulation of an D-type I-front. The rectangular boxes outline the refinement domains used in this model. (b) Radius of the D-type ionization front R_i vs time t (top), and the error relative to the analytic solution (bottom). The top figure shows the analytical solution (solid line) and the numerical approximation to the radius of the I-front (plus signs). The gray shaded region shows the range of all radii included in this average. The bottom panel shows the relative error of the radius in the simulation with respect to the analytic solution.

3.6.4 Radiation Pressure

Here, I demonstrate the effect of radiation pressure through a simple test of a radiation source embedded in a uniform medium. I choose a monochromatic spectrum with $E_\gamma = 13.6$ eV to disable photo-heating and choose an emission rate of 4×10^{48} s^{-1}. The ambient medium is initialized with a density of 5 cm^{-3}, a temperature of T=100 K, and is assumed to be fully ionized. I disable additional thermal physics for both the neutral and ionized gas and use the constant recombination coefficient of $\alpha_B = 2.59 \times 10^{-13}$ cm^3 s^{-1}. The Strömgren radius for this configuration is approximately 30 pc, which extends beyond the computational domain, thus the gas will remain completely ionized for the duration of the simulation. As such, the photon absorption rate depends only on the recombination rate of the gas.

The analytic expression for the acceleration due to radiation pressure is given by:

$$a_{\rm HI}(r) = n_{\rm HI}\frac{1}{c\rho}\int\limits_{\nu_L}^{\infty}\frac{L_\nu e^{-\tau_\nu}}{4\pi r^2}\sigma_{pi}(\nu)d\nu \tag{3.6.6}$$

In the case of a monochromatic spectrum, I write the luminosity in terms of the ionizing photon emission rate:

$$L_\nu = QE_\gamma\delta(\nu_\gamma)\ , \tag{3.6.7}$$

which simplifies the expression to

$$a_{\rm HI}(r) = \frac{n_{HI}Q_\star e^{-\tau}\sigma_{\rm pi}}{4\pi r^2}\frac{E_\gamma}{c\rho}\ . \tag{3.6.8}$$

The rate of photo absorption equals the recombination rate of the ambient gas. Therefore, the expected acceleration within the Strömgren radius is:

$$a_{\rm HI}(r) = n_H^2\alpha_B\frac{E_\gamma}{c\rho}, \tag{3.6.9}$$

which is independent of distance and attenuation terms. I compare the average velocity of all cells with $\rho > 0.95\rho_{\rm amb}$ with the expected constant acceleration in Figure 3.9. Prior to the onset of hydrodynamic effects, the gas accelerates as expected with a relative error of $\approx 1\%$. This figure also shows that the range of velocities calculated for all cells is exceptionally narrow which is in agreement with acceleration that depends exclusively on the local recombination rate.

3.6.5 Secondary Ionization

Lastly, I replicate a test from Wise & Abel (2011), albeit on a smaller scale, to demonstrate the effect of including secondary ionizations. A monochromatic radiation source of 1 keV photons and an emission rate of Q=10^{49} $\rm s^{-1}$ is placed at the center of the box. The

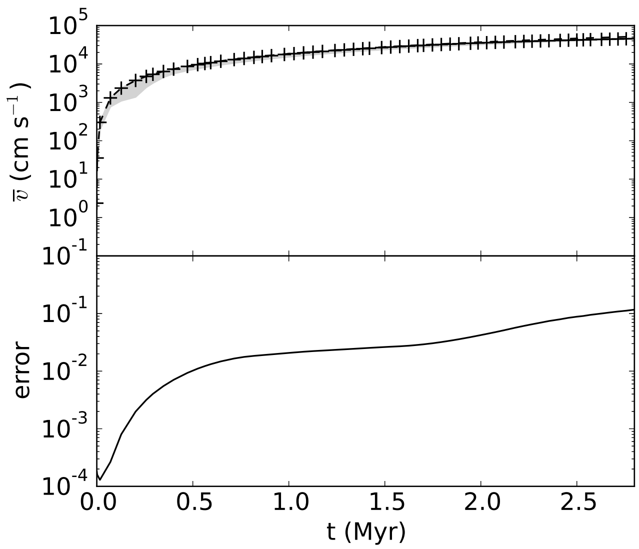

Figure 3.9: Average velocity of an ambient ionized gas exposed to the radiation pressure of a central source vs. time (top), and relative error with the analytic velocity vs. time (bottom). The top panel shows the average velocity of all cells with $\rho > 0.95\rho_{\rm amb}$ in my model (plus signs) as well as the analytic velocity expected due to a constant acceleration. The gray shaded region shows the range of values for all cells that were used in the average. The bottom panel shows the relative error of the numerical result with respect to the analytic result.

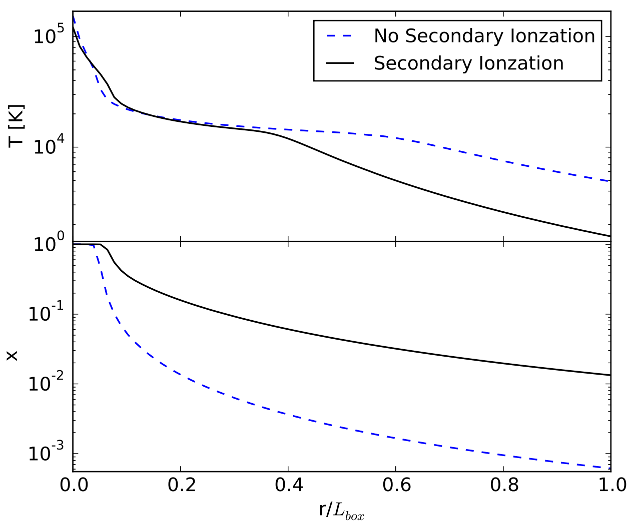

Figure 3.10: Average radial profiles of temperature (top) and ionization fraction (bottom) vs. time for a monochromatic 1 keV source embedded in an initially uniform gas with (black) and without (blue, dashed) secondary ionization.

computational domain is initialized with a uniform gas of $n_{\rm H} = 1$ cm^{-3} and temperature T=10 K. The full thermal physics prescription, including temperature dependent recombinations is included. Average ionization and temperature profiles are shown in Figure 3.10. In general, the inclusion of secondary ionizations extends the ionization profile while simultaneously reducing photo-heating. This qualitative result is in agreement with that of Wise & Abel (2011).

3.7 Additional Features

3.7.1 Active Ray Tracing

There are two cases that I have encountered for which fixed ray-tree integration is impractical. First, in simulations where the spherical symmetry of the radiation field is paramount, periodic rotations of the ray-tree (see § 3.7.2) which would otherwise be a trivial improvement to the accuracy of the routine (Krumholz et al., 2007) are computationally inefficient with the current implementation of fixed ray-tree integration due to the SMR treatment. Second, simulations which include a moving radiation source require a rebuild of the ray-tree at every hydrodynamic timestep due to relative changes in geometry, which is equally inefficient. As a result, I have implemented a method that uses an "active" ray tracing parallelization similar to Wise & Abel (2011).

The ray-tree integration for this method roughly follows Figure 3.3. However, the number of MPI messaging sequences must be determined in real time. The integration follows the fixed ray-tree case such that the twelve rays corresponding to the coarsest resolution of HEALPix are cast outwards until they meet one of four conditions: (i) the ray reaches the termination radius (Eq. 3.3.1) (ii) the ray passes off of the local grid (iii) the ray crosses into a cell which is overlapped by a refinement zone. (iv) $> 99.99\%$ of the photon flux is absorbed in a single cell (Wise & Abel, 2011). In case (i), the ray splits into four child rays. If any of these child rays spawn on a neighboring processor (i.e. the "abandoned child" problem), parallel communication is required. In cases (ii) and (iii) an MPI message is sent to the appropriate processor. Case (iv) occurs in optically thick regions. As the photon flux is rapidly reduced, the ray is terminated, and does not continue integration along the ray-tree branch. Termination of rays in this fashion also accelerates computation due to the decreased number of grid walks and MPI communications required.

For each processor, the number of MPI messages, the processor IDs used in communication, and the number of rays that will enter the local grid are not known a priori. Thus, each array is outfitted with a static buffer on the local processor that is populated and incremented as rays become available. Each ray is integrated using the grid walk method outlined in Abel et al. (1999). If a ray meets termination conditions (ii) or (iii), the final position of the ray on the local processor is used to calculate the processor to which a send must be made. Messages between processors are stored in a local buffer and posted using a non-blocking MPI_ISEND. The buffer is incremented after each send as multiple messages may be posted before prior messages are completed. Because of the relatively low cost of periodic rotations of the ray tree in this case, I choose to not store the grid cells through which rays cross and apply a rotation at every radiation cycle instead to improve spherical symmetry

I adopt the following procedure to track the send and receive sequences globally: (i) All processors are initialized with arrays called MYRECVCOUNT and MYSENDCOUNT which are both equal in length to the number of processors assigned to the simulation. MYRECVCOUNT is used to track the number of receives made by each processor, and MYSENDCOUNT tracks the number of messages sent to other processors. When a message is sent, the sending processor updates the value of MYSENDCOUNT at the index corresponding to the destination processor (i.e. MYSENDCOUNT [PASSOUT_ID]++). As a processor receives a message, it updates MYRECVCOUNT; however each processor can only update the value at the index corresponding to its own processor ID (i.e. MYRECVCOUNT[MYID]++). (ii) Each processor is also outfitted with two arrays, ALLSENDCOUNT and ALLRECVCOUNT which are calculated using an MPI_ALLREDUCE to sum across all processors. At the beginning of the integration cycle, after the initial MPI messages have been posted as a consequence of integrating the 12 root rays on the processor which hosts the radiation source, a number of sub-cycles is calculated as

$$N_{\rm cycle} = \frac{\sum\limits_{i=0}^{N_{\rm procs}} \text{ALLSENDCOUNT}[i] - \text{ALLRECVCOUNT}[i]}{\sum\limits_{i=0}^{N_{\rm procs}} [\text{ALLSENDCOUNT}[i] \; != \text{ALLRECVCOUNT}[i]]} \,, \tag{3.7.1}$$

where $N_{\rm procs}$ is the number of processors. The numerator of this sum calculates the difference between the number of messages sent to each processor and the number of messages which have been processed. The denominator calculates the number of processors which have pending messages. (iii) Each processor sequentially posts blocking MPI receive requests (MPI_RECV) with the MPI_ANY_SOURCE flag until either the number of sub cycles is complete or the processor has completed all pending requests. (iv) The previous process continues until $\sum$ ALLSENDCOUNT $= \sum$ ALLRECVCOUNT, at which point the integration terminates.

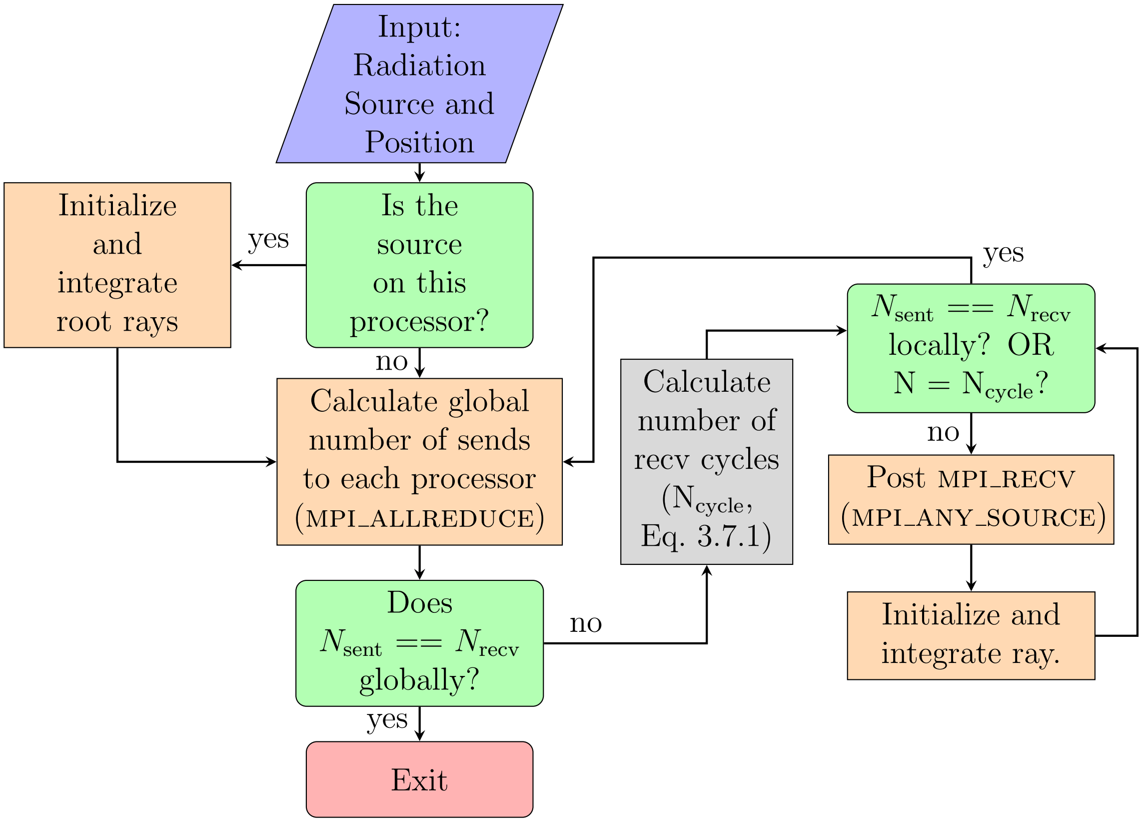

Figure 3.11: Flowchart for parallelized ray integration with active ray tracing.

A schematic of this process is shown in Figure 3.11.

3.7.2 Periodic Rotation

The HEALPix discretization provides the best possible representation of a radiation field through the use of ray-tracing; however, for spherically symmetric problems, such as the HII region expansion test, the ray-tree structure can be imprinted on the gas if only one orientation is used. This can be remedied by the use of geometric corrections to photo-absorption for rays which only partially intercept grid cells (Wise & Abel, 2011). Alternatively, the ray-tree can be periodically rotated (Krumholz et al., 2007). I have implemented the latter option for both fixed and active ray tracing.

At each radiation cycle requiring a ray rotation, I randomly select three angles α, β, and γ in the interval $[0,2\pi]$. These angles are then applied to each ray using the classical Euler

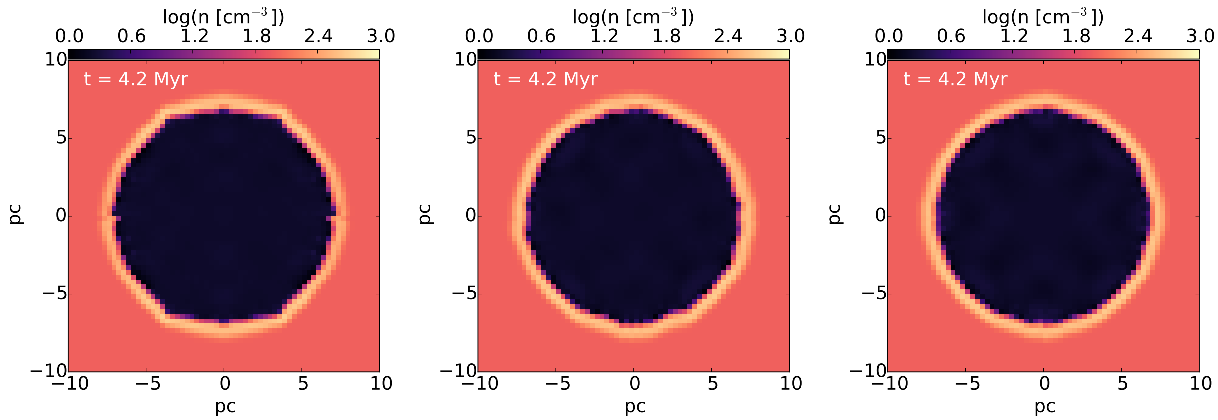

(a) no rotation (b) single rotation (c) periodic rotation (δ N = 10)

Figure 3.12: (a) Midplane slice of the density in a D-type I-front simulation with no ray-tree rotation (b) Same as (a), but is a single ray tree rotation. (c) Same as (a), but with periodic rotations every 10 hydrodynamic timesteps.

angle method:

$$\vec{x}' = \begin{pmatrix} \cos\gamma & \sin\gamma & 0 \\ -\sin\gamma & \cos\gamma & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos\beta & \sin\beta \\ 0 & -\sin\beta & \cos\beta \end{pmatrix} \begin{pmatrix} \cos\alpha & \sin\alpha & 0 \\ -\sin\alpha & \cos\alpha & 0 \\ 0 & 0 & 1 \end{pmatrix} \vec{x} \tag{3.7.2}$$

where primed coordinates are used to indicate rotated values. After rotation, the propagation angles of all rays are updated with respect to grid coordinates. To calculate HEALPix indices using grid coordinates through the ANG2PIX_NEST an inverse rotation must be applied to the cell position.

In Figure 3.12, I show a comparison of simulations with and without rotation. For these models, I use initial conditions identical to those used in § 3.6.3; however, I do not include the additional refinement levels. In Figure 3.12a, in which rotations are not included, grid effects are clearly visible along the principle axes. An improvement on this can be made by including a singular rotation, as is seen in Figure 3.12b. In Figure 3.12c, I show the effect of periodic rotations applied at every 10 hydrodynamic timesteps. An increase in the number of ray-tree rotations reduces the prominence of grid effects and improves spherical symmetry.

The use of periodic rotations is costly for the fixed ray implementation in simulations with SMR, thus I do not make use of this feature generally. I have found that a singular rotation is sufficient to remove strong grid-effects that occur when the HEALPix geometry is aligned perfectly with the principle axes of the mesh. In simulations with active ray-tracing, rotations can be implemented at every radiation cycle with essentially no additional computational cost.

3.7.3 Selective Integration

In this section, I briefly detail an acceleration technique, which I refer to as "selective integration." This process allows for the partial-integration of ray-trees in systems where the region of interest only occupies a small fraction of the computational volume. One particular application of this method considers an isolated gas "clump" in orbit about a central radiation source (see Chapter 5), but this process is generally applicable to systems where the volume filling fraction of "active" gas is low. For these systems, I make the following assumptions: (i) The solution to the radiative transfer equation along lines of sight that do not intercept regions of interest is trivial, therefore the corresponding ionization rates can be determined without requiring a ray trace. (ii) The dynamical influence of heating, ionization, and radiation pressure on gas excluded from the ray trace is not significant to the dynamical evolution of the system.

I have implemented selective ray-tree integration for both active and fixed ray trees. For both cases, I specify a condition to determine which grid cells qualify as "active." For example, in a simulation of a gas cloud with an average density of $n = 10^4$ cm^{-3}, I might set the threshold density for active cells to $n = 10^{-1}$ cm^{-3}. Cells which do not meet the criterion are not considered when determining which rays must be traced. For this particular application the assumptions of selective integration are justified as the recombination timescales ($\tau_{\rm rec} = (n_H \alpha_B)^{-1}$) are sufficiently long so that the number of incident ionizing photons greatly exceeds the number of recombinations in low density, unobscured gas.

To apply selective integration to fixed-ray trees, I calculate and store a list of rays that pass through each cell on the mesh. In addition, I store a grid array, GRIDACTIVE, which indicates whether or not grid cells meet the specified criteria. On each processor, I also store a one dimensional array, RAYSWITCHCOUNT, with length equal to the number of rays in the simulation. A ray's index in RAYSWITCHCOUNT can be determined by summing the total number of rays on previous levels and adding the ray's HEALPix index ($i_{\mathrm{uniq}} = i + \sum_{i=0}^{i=l-1} 12 \times 4^{l}$). GRIDACTIVE and RAYSWITCHCOUNT are both initialized to zero. At the beginning of each radiation step, I determine, based on the preset condition, if the value of GRIDACTIVE has changed. If a cell is newly active, I loop over all rays in the cell and add 1 the value of RAYSWITCHCOUNT count at both the ray's index and all the indices corresponding to the ray's ancestors. For newly inactive cells, one is subtracted from the same group of indices. An MPI_ALLREDUCE is used to total the values of RAYSWITCHCOUNT across all processors at the beginning of every radiation cycle. In this way rays, which are not contained in the local grid but are necessary for the ray trace, are activated. Rays are only integrated if their corresponding value of RAYWITSCHCOUNT is non-zero.

For active ray-tracing, the process of selective integration is much simpler, though not as robust as in the fixed ray-tree case. At the beginning of each radiation cycle, I determine the bounding box of active cells on the computation mesh. Rays are only integrated if they meet one of two conditions: (i) The ray directly intercepts the bounding box of the active gas. To determine this, I assume an infinite radius for the ray and use an optimized ray-box interception algorithm (Williams et al., 2005). (ii) Any of the eight corners of the bounding box are contained within the solid angle subtended by the ray. This is determined using the ANG2PIX_NEST function from the HEALPix library. This is necessary because rays which do not physically intercept the bounding box of active gas may be ancestors of rays that do. Rays which extend beyond the bounding box of the active gas are also terminated. This method is inefficient in the event that multiple bounding boxes are needed; however, for my models this approach was sufficient. In more complicated geometries, a process similar to that outlined for fixed ray-trees would likely be the best option.

3.7.4 Off-Mesh Ray Tracing

As an extension to selective integration (§ 3.7.3), I have also developed a method that allows for the the inclusion of point radiation sources which exist outside of the computational domain. I refer to this method as "off-mesh ray tracing." This is particularly useful in simulations which require a point radiation source, but extending the mesh to include the source is computationally prohibitive (see Chapter 5). To use off-mesh ray tracing, I assume that gas beyond the computational domain is characterized by a constant density and temperature. These are specified in the routine as AMBIENTN and AMBIENTT. Given these two parameters, the radiation field is computed following three steps. (i) Using a ray-box intersection method (Williams et al., 2005), the code determines whether rays extending from the source will intercept the computational domain. This calculation can be expedited by determining the maximum and minimum radius from the source to the mesh walls which entering rays will intercept. This sets the maximum and minimum allowable HEALPix resolution levels required to satisfy the coverage condition in Eq. 3.3.1. I compute this for each computational grid which has a boundary wall that is directly exposed to the radiation source. (ii) For each ray entering the grid, I calculate the radius, r, at which the ray enters the mesh. The photon flux is then calculated via a simple attenuation law:

$$N_\gamma = \frac{Q_\star - n_{\rm ambient}^2 \alpha_B(T_{\rm ambient})\frac{4}{3}\pi r^3}{12\times 4^l}\Delta t_\gamma \ , \tag{3.7.3}$$

where N_γ is the incident photon number flux, $Q_\star$ is the photon emission rate, and Δt_γ is the photon timestep. $n_{\rm ambient}$ and $T_{\rm ambient}$ are the assumed ambient gas number density and temperature. The numerator of this equation accounts for recombinations that occur along the line of sight prior to mesh entry. The denominator accounts for geometric dilution of the radiation field. (iii) Once rays enter the mesh, ray-tracing continues following the process outlined in Figure 3.11. It should be noted that off-mesh ray tracing has only been included for a monochromatic spectrum. For multi-frequency emission, the attenuation in Eq. 3.7.3 would need to be split amongst the photon species.

To demonstrate the use of off-mesh ray-tracing, I provide the following test case. I place a radiation source with $Q_{\rm UV} = 10^{48}{\rm s}^{-1}$ at $\vec{x}_{\rm source} = (5, 0, 0)$ pc on a computational domain extending from -10 pc to 10 pc in the x direction and -5 pc to 5 pc in both the y and z directions. The grid resolution is set to $128 \times 64 \times 64$ with no additional refinement zones. A gas cloud of radius 2.5pc and number density of 10^2 cm^{-3} is embedded in an ambient gas with number density 1 cm^{-3}. The gas cloud is initialized with a temperature of 10 K, and the ambient gas is set to a temperature of 1000 K so that the gas is in pressure equilibrium. The gas cloud is centered at $\vec{x}_{\rm cloud} = (-5, 0, 0)$ pc. I do not include gravity in this test. The results of this model are shown in the left column of Figure 3.13. The exposed face of the cloud is photo-compressed over time. Photo-heating at the leading edge of the cloud results in the expansion of a low density shell. I repeat this simple model with off-mesh ray tracing by using a computational domain of 64^3 cells that extends from -10 pc to 0 pc in the x direction and -5 pc to 5 pc in both the y and z directions. The radiation source is situated off of the mesh. The results of this model are shown in the second column of Figure 3.13. As in the previous case, the radiation compresses the face of the cloud and photo-heating drives the expansion of a shell extending from the irradiated face. Similar features are seen between the two simulations; however, as the low density shell expands off of the computational domain the assumption of a constant density and temperature ambient gas is no longer accurate. Because of this, compression in the model using off-mesh ray tracing is diminished, and the error with respect to the on-mesh simulation rapidly increases (right column of Figure 3.13). Yet, the general structure is consistent between the two cases, suggesting that the radiation field calculated via off-mesh ray-tracing is a viable approximation.

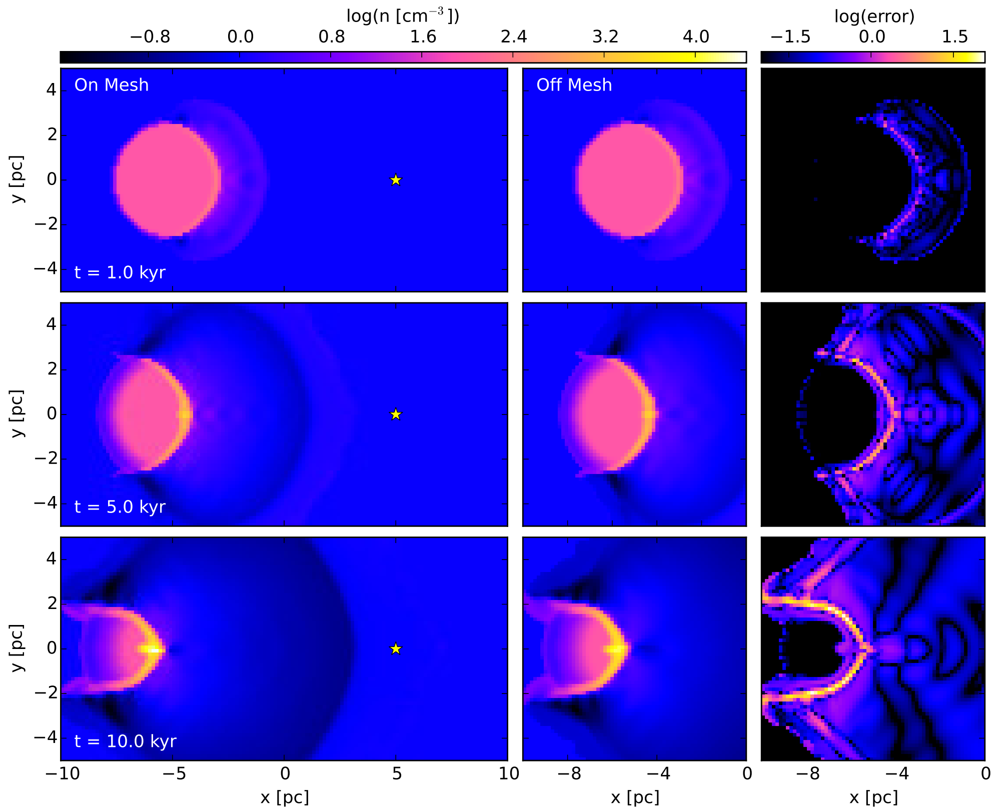

Figure 3.13: Comparison of on-mesh ray-tracing and off-mesh ray-tracing for a model of an irradiated gas cloud. Time increases from top to bottom. The left panel shows the results of the simulation with the radiation source located on the mesh (indicated with a yellow star). The central panel shows the results for a simulation with the approximate radiation treatment used in the off-mesh ray-tracing treatment. The right panel shows the error between the two simulations, and suggests that off-mesh ray-tracing only provides an approximate solution to the radiative transfer problem.

3.8 Performance

I lastly present a brief discussion on the parallel performance of the radiative transfer module. Assessing the scalability of radiative transfer routines is difficult by virtue of the fact that the parallel communication is highly sequential. This is particularly challenging for a singular source as processors which are assigned to grids that are separated from the radiation source must remain idle until rays traverse through many other interior grids. Optimization via asynchronous messaging helps to mitigate this issue, and scaling tests can be done with a large number of sources distributed across the computational volume to diminish processor latency (Wise & Abel, 2011). The implementation of radiative transfer into ATHENA detailed here has yet to be fully completed for multiple radiation sources, thus the tests presented are limited to a singular source and do not necessarily represent the general behavior of the radiative transfer routine for all problems.

3.8.1 Strong Scaling

Strong scaling tests are designed to assess how a code performs with an increase in the processor load for a fixed job size. To demonstrate the strong scaling in my radiative transfer routine, I repeat the test outlined in § 3.6.3. For this test, however, I use four perfectly nested domains centered on the origin, with a resolution of 64^3 at each level. I run this test for both the fixed and active ray tracing techniques used in this work, and use 8 – 2048 processors, increasing the number of processors in factors of 2.

I show the results of the strong scaling test for the fixed ray tree integration in Figure 3.14a, which includes the total time per simulation, total time spent in the radiation cycle, and the average time spent on ray-tracing and parallel communication for the radiative transfer. In general, the code shows good scaling (perfect strong scaling would show $t \propto N^{-1}$). For this particular application, parallel communication in the radiative transfer routine is the most costly component. The bottleneck in this test is the idle time of processors which must wait for rays to traverse many interior grids. As the MPI communication

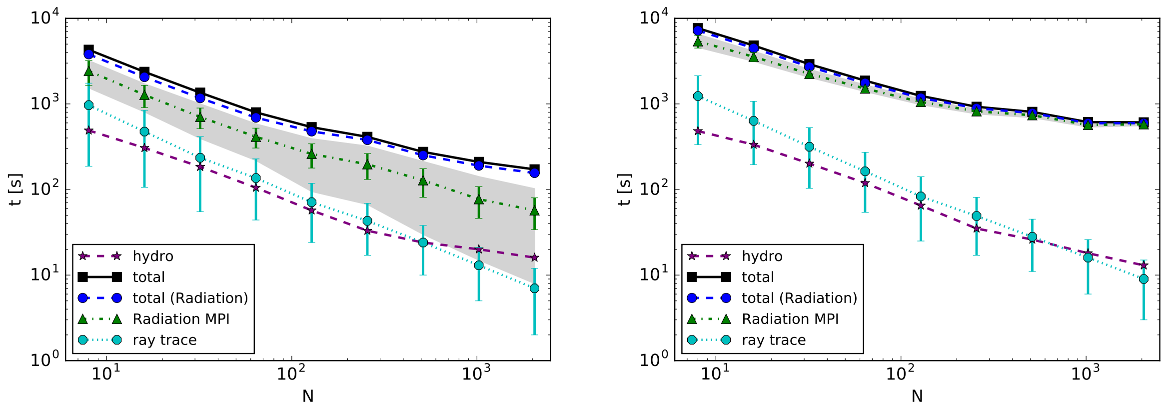

(a) (b)

Figure 3.14: Strong scaling test for fixed ray (left) and active ray (right) integration using the D-type ionization front test with four centered SMR levels of resolution 64^3 each (effective 512^3). Total times are shown for the cumulative integration (black), radiation module (blue), MPI communication (green), ray-tracing (cyan), and hydrodynamic update (purple). The error bars correspond to the standard deviation of times. The grey shaded region shows the range of values used in the MPI communication time average. $t \propto N^{-1}$ is ideal.

is completely non-blocking, the range of times spent on MPI communication widely varies because of this processor latency (grey shaded region in Figure 3.14a). Despite good scaling in the average MPI communication time, the total cost of the radiation routine depends completely on the poorest performance across all processors. A turn over is seen in the total time for the hydrodynamic solver at N_{proc}=512. This occurs because the domain decomposition becomes such that the ratio of the active cells to ghost cells on a given grid approaches 1. At this point, the hydrodynamic update is dominated by MPI communication required for calculating boundary conditions on each grid. The average cost of local ray tracing shows nearly perfect scaling for the full range of processor number.

I show the results of the strong scaling test for "active" integration in Figure 3.14b. Both ray tracing and the hydrodynamic update show good strong scaling. The overall timing shows good scaling up to $N_{\mathrm{procs}} \lesssim 128$, beyond which the MPI communication becomes inefficient. This occurs because the process which is detailed in Figure 3.11 makes use of the blocking MPI_ALLREDUCE function, thus rays are processed in batches, increasing the idle time of each processor. In addition, all processors complete the final cycle of communication

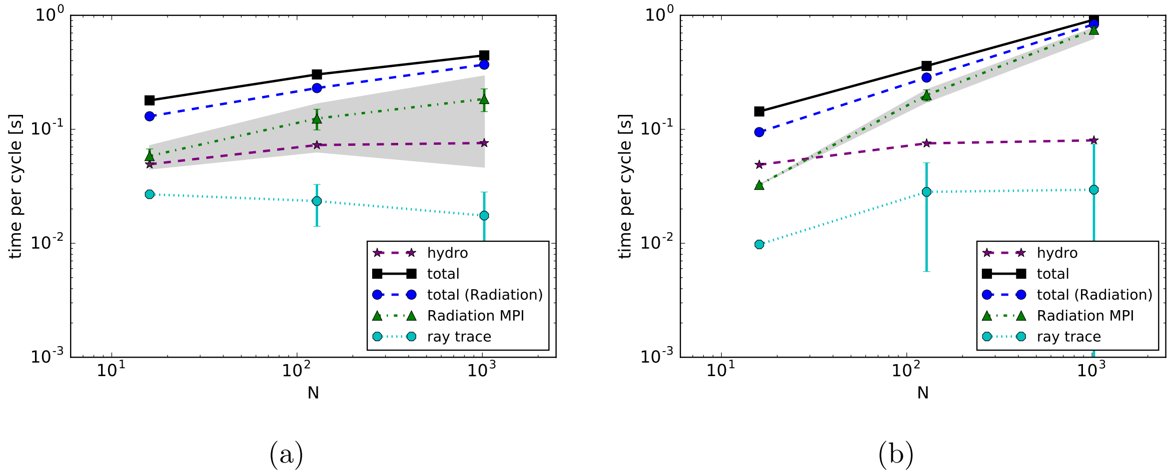

Figure 3.15: Weak scaling test for fixed ray (left) and active ray (right) integration using the D-type ionization front test. Total times are shown for the cumulative integration (black), radiation module (blue), MPI communication (green), ray-tracing (cyan), and the hydrodynamic update (purple). The error bars correspond to the standard deviation of times. The grey shaded region shows the range of values in the average of MPI communication time. A flat line, indicating a constant time per cycle, would be ideal.

simultaneously, thus the time spread in MPI communication time seen in Figure 3.14a is not seen. The parallel communication scaling can be improved by implementing a completely non-blocking algorithm (Rosen et al., 2017).

3.8.2 Weak Scaling

Weak scaling tests are used to determine how performance varies with a constant job size per processor. For radiative-hydrodynamic codes, this test typically uses a number of computational blocks equal to the number of processors, with a radiation source at the center of each block (Rosen et al., 2017; Wise & Abel, 2011). As stated before, the radiation routine developed for this work does not yet include multiple radiation sources, thus I cannot replicate this test exactly. As an alternative, I repeat the D-type ionization test (§ 3.6.3) without the additional SMR domains. I run this test with 16, 128, and 1024 processors at resolutions of 64^3, 128^3, and 256^3, respectively. This sets the number of grid cells per processor to 16384. As the resolution varies between these simulations, the number of hydrodynamic cycles required is not identical, thus I compare the average time per cycle.

In Figure 3.15a, I show the weak scaling for fixed ray integration. As the parallelization in this scheme is completely non-blocking, the weak scaling is acceptable (perfect scaling results in a flat line). Both the hydrodynamic update and the ray tracing routines show good scaling up to 1024 processors. The MPI communication for the radiative transfer increases with the job size and essentially sets the total simulation time. The resulting time per cycle increases by a factor ≈ 2 from 16 to 1024 processors.

I show the weak scaling for active ray tracing in Figure 3.15b. As in the previous case, the hydrodynamic update and ray tracing show excellent weak scaling. MPI communication in the radiation routine scales poorly, with an order of magnitude drop in processor efficiency over this processor range. This effect is driven by the fact that the parallelization used in Figure 3.11 requires several synchronizations per radiation cycle. With an increase in the processor load, an increased number of MPI_ALLREDUCE calls decreases the overall processor efficiency. Again, using a more complex asynchronous book keeping strategy (i.e. Rosen et al. (2017)) would remedy this issue. I plan on implementing an approach similar to this in future work.

CHAPTER 4: GAS INFLOW AND STAR FORMATION AROUND SUPERMASSIVE BLACK HOLES: THE ROLE OF NUCLEAR ACTIVITY

4.1 Introduction

4.1.1 Star Formation in the Galactic Center

Stellar orbits in the Milky Way's Galactic Center (GC) serve as direct evidence for the existence of a $4 \times 10^6\ M_\odot$ supermassive black hole (SMBH), coinciding with the radio source, Sgr A* (Schödel et al., 2002; Ghez et al., 2003; Gillessen et al., 2009a). Several hundred massive stars orbit within the central parsec of Sgr A*, many of which belong to a clockwise orbiting disk extending from 0.05 to 0.5 pc (Genzel et al., 2003; Paumard et al., 2006; Bartko et al., 2009, 2010; Lu et al., 2009, 2013; Yelda et al., 2014). The stellar disk age of < 6 Myr (Paumard et al., 2006; Lu et al., 2013) strongly suggests that these stars formed in situ. However, the immense tidal field of the SMBH is expected to inhibit star formation within the central few parsecs of the GC (see discussion in Mapelli & Gualandris (2016)).

Possible explanations for the observed stellar populations include newly formed massive star clusters migrating inwards via dynamical friction under the gravitational influence of an intermediate mass black hole (Hansen & Milosavljević, 2003), or an overabundance of massive stars (Gürkan & Rasio, 2005), though observational and timescale constraints do not strongly support such scenarios (Stolte et al., 2008; Genzel et al., 2010). The alternative, in-situ star formation, remains favored.

Theory suggests that star formation in the immediate vicinity of a SMBH can occur via rapid cooling and fragmentation of an accretion disk (Levin & Beloborodov, 2003; Nayakshin et al., 2007). The formation of a sufficiently dense accretion disk is a natural consequence of the tidal disruption of a $\approx 10^5\ M_\odot$ molecular gas stream on a low angular momentum orbit about the GC (Wardle & Yusef-Zadeh, 2008). Hydrodynamic models of this process reproduce stellar disks in rough agreement with observed stellar orbits (Sanders, 1998; Lucas et al., 2013; Bonnell & Rice, 2008; Mapelli et al., 2012; Alig et al., 2011). The origin of such gas inflow remains uncertain, though models suggest that gas clump collisions at ≈ 1 pc could supply sufficient inflow to incite a star formation episode (Hobbs & Nayakshin, 2009; Alig et al., 2013). Furthermore, observational estimates for an inflow rate of 0.1-1 $M_\odot$ yr^{-1} in the GC (Morris & Serabyn, 1996) are consistent with the infall of $\approx 10^5$ $M_\odot$ gas streams on a timescale of a few $\times$ Myr.

4.1.2 Nuclear Activity in the Galactic Center

Models which explore the process of stellar disk formation resulting from gas stream capture also show evidence of accretion rates onto the SMBH at considerable fractions of the Eddington limit (Bonnell & Rice, 2008; Hobbs & Nayakshin, 2009; Alig et al., 2011):

$$\dot{M}_{\rm edd} = 2 \times 10^{-8} \left(\frac{M_{\rm BH}}{M_\odot} \right) M_\odot \text{ yr}^{-1} \ . \tag{4.1.1}$$

During such an accretion episode, an active galactic nucleus (AGN) can radiate at large fractions of the Eddington luminosity:

$$L_{\rm edd} = \dot{M}_{\rm edd} c^2 \epsilon_r = 3 \times 10^4 \left(\frac{M_{\rm BH}}{M_\odot} \right) L_\odot \ , \tag{4.1.2}$$

where $\epsilon_r = 0.1$ is the radiative efficiency, and c is the speed of light. Despite estimates for large scale mass inflow in the GC, Sgr A* shows no evidence of current accretion activity (see Morris & Serabyn (1996) and references therein). Furthermore, observations of the GC limit the bolometric luminosity of Sgr A* to $\approx 10^{-10} - 10^{-9} L_{\rm edd}$ over the past few hundred years (Sunyaev et al., 1993; Baganoff et al., 2003). Yet, there are two pieces of evidence that point to past AGN activity in the GC. First, the existence of two extended gamma-ray sources referred to as the Fermi bubbles (Su et al., 2010) may be the result of either a

Galactic outflow triggered by AGN activity 6 Myr ago (Zubovas & Nayakshin, 2012) or an AGN jet that existed 1-3 Myr ago (Guo & Mathews, 2012). Second, as previously noted, the population of several hundred massive stars within the central parsec of Sgr A* is difficult to explain in the absence of rapid gas inflow towards Sgr A*.

So far, no models exploring tidal disruption of inflowing gas and central star formation have considered the effect of radiative feedback from an accretion episode onto the SMBH. Yet, radiative-hydrodynamic (RHD) simulations of gas clouds subject to AGN radiation have been presented in several works. Using two-dimensional models of infalling dusty gas clouds, Schartmann et al. (2011) demonstrated that the fate of such objects largely depends on the column density of the gas. Only in cases with sufficiently strong shielding is gas able to withstand the impinging radiation and complete its approach to the SMBH. Hocuk & Spaans (2011, 2010) considered the evolution of clouds at distances of ≈ 10 pc to explore the effect of X-ray feedback on the initial mass function (IMF) of stars forming in this region. Assuming that UV radiation is obscured by interior gas and dust, these models showed that X-ray radiation alone leads to significant gas compression and heating, the latter of which promotes the formation of higher mass protostars. Most recently, Namekata et al. (2014) explored the effect of both UV and X-ray radiation on infalling gas clouds with galactocentric distances of 5 pc and 50 pc using three-dimensional hydrodynamic simulations, including a detailed chemical network. They parameterize the radiation field by the ionization parameter $U_{\rm ion}$, which measures the ratio of the photon density to the number density of the irradiated gas. These models show that photo-evaporation dominates when $U_{\rm ion}$ is low, whereas radiation pressure becomes more important for large values of $U_{\rm ion}$. Collectively, the models demonstrate the impact of AGN radiation on the surrounding medium. Yet, models which follow the evolution of such clouds through a direct collision with the central SMBH, as is required for the birth of a nuclear stellar disk, have yet to be considered.

4.1.3 Motivation and Outline

We explore the role of radiative feedback from accretion onto the central SMBH on the formation and evolution of a circum-nuclear gas disk. Specifically, we consider the gas inflow scenario which is known to result in both the formation of a stellar disk as well as accretion rates at large fractions of the Eddington limit.

Numerical methods are outlined in § 4.2. In § 4.3, we describe the initial conditions for our stream inflow models. We present our results in § 4.4 and discuss the implications in the context of nuclear star formation in § 4.5. We provide a summary of key results from this study in § 4.6.

4.2 Methods

We use a modified version of ATHENA 4.2 (Stone et al., 2008) to solve the system of equations:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0 \tag{4.2.1}$$

$$\frac{\partial \rho \mathbf{v}}{\partial t} + \nabla \cdot (\rho \mathbf{v}\mathbf{v} + \mathbf{I}P) = -\rho \nabla \phi + \rho \mathbf{a}_\gamma \tag{4.2.2}$$

$$\frac{\partial E}{\partial t} + \nabla \cdot (\mathbf{v}(E + P)) = -\rho \mathbf{v} \cdot \nabla \phi + \rho \mathbf{v} \cdot \mathbf{a}_\gamma + \mathbf{G} - \mathbf{L} \tag{4.2.3}$$

$$\frac{\partial C\rho}{\partial t} + \nabla \cdot (C\rho \mathbf{v}) = 0 \tag{4.2.4}$$

$$\frac{\partial \rho_{\mathrm{HII}}}{\partial t} + \nabla \cdot (\rho_{\mathrm{HII}} \mathbf{v}) = m_H (I - R); \tag{4.2.5}$$

with the gas density ρ, the fluid velocity vector $\mathbf{v}$, the gas pressure P, the unit dyad $\mathbf{I}$, the energy density

$$E = \frac{1}{2}\rho \mathbf{v} \cdot \mathbf{v} + \frac{P}{\gamma - 1}\ , \tag{4.2.6}$$

and a static gravitational potential

$$\phi = -\frac{G M_{\mathrm{BH}}}{r} \tag{4.2.7}$$

where M_{BH} is the black hole mass which is situated at the origin. We include two source terms, G and L, which account for energy gains, and for losses due to radiative processes (see § 3.4.6). A color field, C, is advected to trace inflowing gas.

For models that include radiation, we track the ionization of hydrogen gas via Eq. 4.2.5. The mass density of ionized hydrogen, ρ_{HII}, depends on two source terms, I and R, which are the ionization and recombination rates per volume (see Chapter 3). In addition, we include an acceleration, $\mathbf{a}_\gamma$, to account for radiation pressure.

For all models in this work we use the directionally un-split Van-Leer (VL) integrator (Stone & Gardiner, 2009) with second order reconstruction in the primitive variables (Colella & Woodward, 1984) and the HLLC Riemann solver (Toro, 2009). We assume a pure hydrogen gas ($\mu = 1$) with an adiabatic ($\gamma = C_P/C_V = 5/3$) equation of state. We assume Cartesian geometry for our computational mesh.

4.2.1 Radiation

Radiation is included via an adaptive ray-tracing routine that roughly follows both a previous implementation into the code (Krumholz et al., 2007) and the radiation module from ENZO-MORAY (Wise & Abel, 2011). At each radiation cycle, an adaptive ray tree is traced sequentially from the radiation source outwards throughout the computational mesh. As rays traverse through cells, attenuation of incident radiation leads to photon deposition, gas heating, ionization, and radiation pressure. A full description of our radiative transfer module as well as standard tests of its accuracy are provided in Chapter 3.

4.3 Simulation Set-up

For all models in this work, we use a $(4\ \mathrm{pc})^3$ computational box centered on the origin. Boundary conditions on the box are set to allow outflow but prohibit inflow. A SMBH of $4 \times 10^6\ M_\odot$ (Gillessen et al., 2009a) is placed at the origin, implemented via a static gravitational potential. To resolve fluid flow around the SMBH while also minimizing computational

Table 4.1: Mesh and Refinement Geometry

Level	Dimensions pc	x_0 pc	y_0 pc	z_0 pc	Res.	Δx mpc
1	4×4×4	-2	-2	-2	64^3	62.5
2	2×2×2	-1	-1	-1	64^3	31.3
3	1×1×1	-0.5	-0.5	-0.5	64^3	15.6
4	0.5×0.5×0.5	-0.25	-0.25	-0.25	64^3	7.8

cost we use three levels of static mesh refinement (SMR), resulting in an effective resolution of 512^3 or 7.8 mpc at the finest level. The geometry of the computational mesh, including refinement zones, is shown in Table 4.1.

We initialize a uniform ambient medium with a number density of $n_0 = 1$ cm^{-3} and a gas temperature of $T_0 = 5803$ K, corresponding to the equilibrium temperature of gas ionized by UV radiation in our thermal model. The ambient gas is assumed to be fully ionized and is initialized with zero velocity. To ensure that the ambient medium does not collapse under the influence of the central SMBH, we set the color field of ambient gas to zero. At every hydrodynamic timestep, we reset cells with color fields of $C < 10^{-10}$ to the ambient initial condition. This approach is similar to Burkert et al. (2012), though we have lowered the threshold color field as the ambient gas profile is not convectively unstable.

The densities and temperatures in our models range over several orders of magnitude, requiring density and temperature floors to avoid occasional failures in the integration scheme. We choose a minimum number density of $n = 1$ cm^{-3}, consistent with the ambient background. Similarly, we assume a temperature minimum of $T = 100$ K, which is consistent with the minimum equilibrium temperature in our thermal model. We enforce these floors at the end of each hydrodynamic update. Imposing a density floor is equivalent to adding mass. Yet, over the duration of the simulation, the mass accumulated in this way is negligible.

4.3.1 Accretion Boundary

An accretion boundary with radius $R_{\rm acc} \approx 40$ mpc, or 5 cells on the highest refinement level, is centered on the SMBH at the origin. Outflow is not explicitly imposed on the accretion boundary. Instead, we smoothly remove momentum and mass from gas which enters this region using the radial smoothing profile,

$$s = \left(1 - \frac{r}{R_{\mathrm{acc}}}\right)^2 \min(\frac{\Delta t}{t_{\mathrm{s}}}, 1) \ , \tag{4.3.1}$$

with the hydrodynamic timestep Δt, and the smoothing timescale $t_{\mathrm{s}} = 0.1$ yr. The density, velocity, and temperature within the accretion boundary are then rescaled as

$$n' = n(1 - s) + n_0 s \tag{4.3.2}$$

$$T' = T(1 - s) + T_0 s \tag{4.3.3}$$

$$\mathbf{v}' = \mathbf{v}(1 - s) \ , \tag{4.3.4}$$

where primed variables are the updated values. We scale the color field proportionally with the density, and we calculate the change of mass in this region ($\rho C \Delta x^3$) throughout this process to track accretion.

4.3.2 Inflow Conditions

We model infalling gas streams via an inflow condition on the $+\hat{x}$ face of the computational box. We show a two-dimensional representation of this set-up in Figure 4.1. Density perturbations are imposed onto the inflowing gas for two reasons. First, inhomogeneity in the inflow mimics the substructure observed in interstellar gas. Second, in the process of gravitational focusing during the cloud's infall, streams of gas passing the SMBH in opposite directions collide, leading to specific angular momentum cancellation. Uniform inflow results in the well known Bondi-Hoyle-Lyttleton accretion process (Bondi & Hoyle, 1944), whereas inhomogeneity in the gas leads to the retention of angular momentum that is essential for the formation of a dense gas disk and stars (Yusef-Zadeh et al., 2008).

Perturbations are calculated as a sum of incoherent sine waves. We first determine the

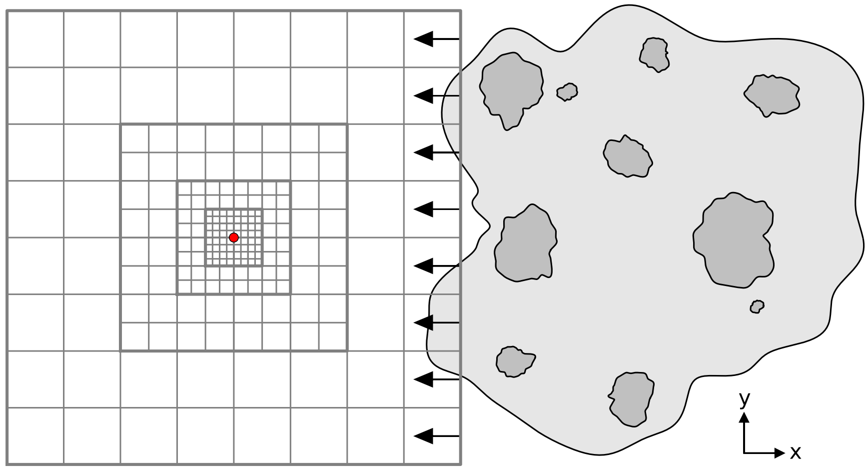

Figure 4.1: Two-dimensional representation of the computational mesh and inflow conditions used for our models. Three SMR domains are nested and centered on an SMBH which rests at the origin (red circle). A gas stream with density perturbations is injected into the mesh via a time-dependent inflow condition denoted by black arrows.

perturbation exponent as

$$\delta(\vec{x}) = \sum_{k_k,k_j,k_i=1}^{k_{\max}} |k|^{-\alpha} \sin\left(\frac{2\pi k_i}{L_{\max}}(x - v_{\rm in}t) + \phi_x(\vec{k})\right) \times \sin\left(\frac{2\pi k_j y}{L_{\max}} + \phi_y(\vec{k})\right) \sin\left(\frac{2\pi k_k z}{L_{\max}} + \phi_z(\vec{k})\right) , \tag{4.3.5}$$

where ϕ_x, ϕ_y, and ϕ_z are randomly chosen phases between 0 and 2π for all values of $\vec{k} = (k_i, k_j, k_k)$. $v_{\rm in}$ is the inflow velocity of the gas which is assumed to be constant. $L_{\max}$ is the longest length scale included, which is set to 5 pc. It should be noted that, in our case, the inflow length ($L_{\rm in} = v_{\rm in}t_{\rm in}$) is shorter than the box dimensions. In cases where the inflow is extended, $L_{\max}$ must be set accordingly to avoid repetitious inflow structure. The maximum wavenumber is set by the ratio of the inflow scale and the "clump scale", L_c, which represents the size of the smallest structures in the inflow. We set $L_c = 0.1$ pc, corresponding to a maximum normalized wavenumber $k_{\max} = \text{ceil}(L_{\max}/(2L_{\rm c})) = 26$. This is consistent with lower limits of observed clump sizes in both the circumnuclear disk (CND) ($r \approx 0.125$ pc) (Christopher et al., 2005) and the so-called 50 km s^{-1} cloud ($r > 0.15$ pc) (Tsuboi & Miyazaki, 2012). Somewhat motivated by turbulent cloud structure, we choose a power law index of $\alpha = 3$. This results in a density distribution steep enough to prevent strong angular momentum cancellation, while it generates a sufficient amount of "cloud fragments". We then scale the perturbation exponent to the inflow density as

$$\delta_{\rm norm} = \frac{\delta - \min(\delta)}{\max(\delta) - \min(\delta)} \ \log_{10}\left(\frac{\rho_{\rm in,max}}{\rho_{\rm in,min}}\right) , \tag{4.3.6}$$

where $\rho_{\rm in,min}$ and $\rho_{\rm in,max}$ are the minimum and maximum mass densities of the inflow. The inflow mass density at each position is then calculated as:

$$\rho_{in}(\vec{x}) = 10^{\ \delta_{\rm norm}(\vec{x})} \tag{4.3.7}$$

Scaling the density perturbations logarithmically is necessary to generate a gas stream with multiple isolated clumps similar to interstellar gas residing in the central 100 pc of the GC (Zylka et al., 1990). Lastly, to give the cloud finite extent in the direction perpendicular to the inflow, we use a hyperbolic tangent function to smoothly bring the inflow density to the ambient gas density for $y^2 + z^2 > 2$ pc. This also effectively removes high angular momentum clumps from the inflow.

To save computational time, the inflow structure is calculated and stored in an array with the cell spacing of the coarsest grid. Inflowing densities are then linearly interpolated from this grid, and fed into the boundary cells of the computational box. We assume the inflowing gas is neutral, thus we set the gas temperature to the density-dependent equilibrium temperature for neutral gas ($\Gamma_{\rm n} = n_{\rm HI}\Lambda_{\rm n}$). We also set the color field $C = 1$ for all inflowing mass in order to distinguish it from the ambient background.

4.3.3 Radiation Field

In our treatment of radiation we use a binned monochromatic spectrum with photon energies of 16 eV for UV radiation and 1 keV for X-ray radiation. The selected UV photon energy is consistent with Krumholz et al. (2007) and is intentionally in excess of the hydrogen

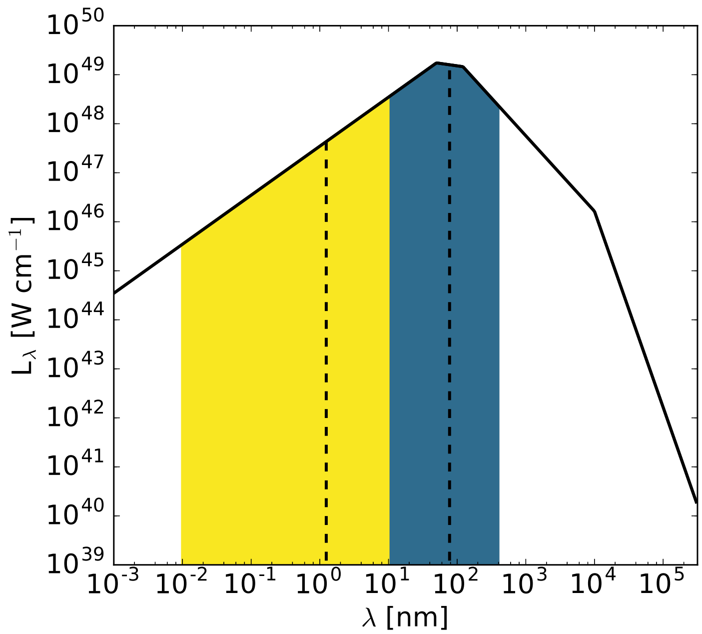

Figure 4.2: Model spectral energy distribution for a $4 \times 10^6\ M_\odot$ SMBH radiating at the Eddington luminosity. The colored regions represent the wavelength ranges for X-ray (left, yellow) and UV (right, blue) photons. The vertical lines at 1.24 nm and 77.5 nm correspond to the photon energies of 1 keV and 16 eV used in this work.

binding energy (13.6 eV) to allow for photo-heating from this species. To determine the photon emission rate for each species, we assume a piecewise spectral energy distribution for an AGN (Schartmann et al., 2005):

$$L_\lambda \propto \begin{cases} \lambda^{-1} & \lambda < 500\text{Å} \\ \lambda^{-0.2} & 500\text{Å} < \lambda < 121.5\text{nm} \\ \lambda^{-1.54} & 121.5\text{nm} < \lambda < 10\mu\text{m} \\ \lambda^{-4} & 10\mu\text{m} < \lambda \end{cases} \tag{4.3.8}$$

In Figure 4.2 we show the normalized spectral energy distribution for a $4 \times 10^6\ M_\odot$ SMBH radiating at the Eddington luminosity (Eq. 4.1.2). Assuming an input bolometric luminosity, we integrate the spectral energy distribution and calculate the normalization factor $L_{\rm norm} = L_{\rm bol} / \int_0^\infty L_\lambda d\lambda$. For each photon species, we calculate the net luminosity

Table 4.2: Inflow Parameters

prefix	n_{min} cm^{-3}	n_{max} cm^{-3}	$\bar{n}$ cm^{-3}	M_{in} $M_\odot$	seed
C1	10^4	5×10^6	9.7×10^4	1.0×10^5	1
C2	10^4	10^7	1.6×10^5	1.6×10^5	2

within the appropriate wavelength range and multiply by this normalization factor. For UV photons, we integrate in the range of 10–400 nm, and for X-rays we set this range to 0.01–10 nm. The photon emission rate is then determined by dividing by the photon energy. Using this model, the photon emission rates at the Eddington luminosity of a 4×10^6 $M_\odot$ SMBH are calculated to be approximately $Q_{\rm UV} = 10^{55}$ s^{-1} and $Q_{\rm X} = 10^{51}$ s^{-1}.

4.3.4 Models

We present two sets of inflow models. The naming convention of our models uses a prefix to denote the inflow structure of the gas and a suffix to indicate the radiation field. The inflow is initialized with a minimum number density of 10^4 cm^{-3} for both cases. The peak number density for the C1 inflow condition is set to 5×10^6 cm^{-3} which yields an average number density of $n = 9.7 \times 10^4$ cm^{-3} and total inflow mass of 10^5 $M_\odot$. For the C2 model inflow, the peak number density is set to 10^7 cm^{-3} which results in an average number density of $n = 1.6 \times 10^5$cm^{-3} and a total mass of 1.6×10^5 $M_\odot$. We use unique random seeds to calculate perturbations in C1 and C2. The inflow velocity is set to 100 km s^{-1} in all models, but velocity perturbations are not included as they are expected to be unimportant due to the highly supersonic bulk flow (Yusef-Zadeh et al., 2008). We follow the evolution of this system for 100 kyr, during which gas is smoothly injected into the $+\hat{x}$ boundary on the coarsest level for the first kyr and continues until $t = 25$ kyr. Parameters for the two inflow conditions are shown in Table 4.2.

For each inflow condition, we run simulations without radiation as a control case. These simulations are denoted with the suffix C. We include our full radiative transfer scheme at

Table 4.3: Radiation Parameters

suffix	$Q_{\rm UV}$ $\rm s^{-1}$	$Q_{\rm X}$ $\rm s^{-1}$	notes
C	0	0	no radiation
C128*	0	0	resolution increase ×2
C256*	0	0	resolution increase ×4
RL	1e54	1e50	
RH	1e55	1e51	$\rm L = L_{edd}$
X†	0	1e51	X-rays only
UV†	1e55	0	UV photons only
NORP†	1e55	1e51	no radiation pressure

∗ used for testing convergence, only for C2 (see § 4.5.6)

† used for testing radiation field components (see § 4.5.2 and § 4.5.3)

10% and 100% of the Eddington luminosity. These simulations are labelled with the suffixes RL ("low") and RH ("high"), respectively. We repeat runs C1-RH and C2-RH three times each to explore the influence of various components of the radiation field. In these models, we scale the radiation field to the Eddington limit and (i) only include X-rays (ii) only include UV photons (iii) exclude radiation pressure. These simulations are labelled X, UV, and NORP respectively. A list of the radiation parameters used in our models is given in Table 4.3. We also repeat model C2-C twice at two and four times the resolution listed in Table 4.1. These models are named C2-C128 and C2-C256, respectively.

4.3.5 Disk Finding and Stability Measure

We do not include self-gravity in our models, thus we are unable to follow the evolution of formed gas disks to the point of star formation. This is partly due to the fact that our simulations are resolution limited, and we cannot sufficiently resolve the Jeans length ($\lambda_J = (\pi c_s^2/G\rho)^{1/2}$, Jeans, 1902) of dense gas. Peak densities in our models reach $n = 10^9$ cm^{-3}, thus, assuming an equilibrium temperature of 100 K, the Jeans length of 6 mpc is roughly equal to the cell size on the highest refinement level. When considering additional constraints on the spatial resolution required for monitoring gravitational collapse in grid-based simulations (i.e. Truelove et al. (1997)) or threshold densities motivated by the tidal limit (Bonnell & Rice, 2008; Lucas et al., 2013; Hobbs & Nayakshin, 2009) the required resolution is not computationally feasible with our radiative transfer module.

As an alternative to directly monitoring disk fragmentation, we test for gravitational instability of formed disks via an approximate measurement of the Toomre Q parameter (Toomre, 1964):

$$Q_{\mathrm{T}} \approx \frac{c_s \Omega}{\pi G \Sigma} \tag{4.3.9}$$

where c_s is the sound speed of the gas, Ω is the orbital frequency, G is the gravitational constant, and Σ is the surface density of the disk. A Keplerian disk is subject to gravitational instability for values of $Q_{\mathrm{T}} < 1$.

Because of the random perturbation structure of the inflow used in our models, we cannot perfectly predict the orientation or extent of gas disks that form in our simulations. Therefore, we have implemented a parallelized disk finding routine that is detailed in § 4.3.6. In short, the routine uses search volumes with radii of 0.25 pc, 0.5 pc, 1 pc, and 2 pc centered on the SMBH to calculate the net angular momentum of enclosed gas. A disk is found if a significant fraction of the mass within the search volume has an angular momentum vector roughly parallel to the net angular momentum vector. The routine then interpolates the computational mesh to a frame aligned with the net angular momentum vector, preserving the mesh refinement when possible. We exclude gas that does not have a parallel angular momentum vector from this interpolation step in order to isolate the disk. To extract only the dense central component of the disk, we also exclude gas with a mass density $\rho < 10^{-3}\rho_{\mathrm{peak}}$, where ρ_{peak} is the peak mass density within the disk. We run this routine at 1 kyr intervals throughout the simulation. We then use the interpolated disk frame mesh to calculate the mass-weighted average sound speed ($\bar{c_s} = \sum(c_s \rho \Delta x^3)/\sum(\rho \Delta x^3)$), average surface density, average orbital velocity, and the total mass of each disk. These values are then inserted into Eq. 4.3.9. To avoid redundant measurements of the disk at each time interval, we only retain values for the smallest search volume which contains at least 99% of the total disk mass. Our values of Q_{T} should be treated as conservative estimates as they do not account for the radial dependence of any of the quantities which enter into the Toomre Q calculation.

4.3.6 Disk Finding Algorithm

Here we detail an efficient parallelized disk finding routine that allows us to track the formation and evolution of formed disks in our models irrespective of scale or orientation. We implement this routine directly into ATHENA for two reasons. First, this tool provides the advantage of "on-the-fly" analysis which yields much greater time accuracy than is reasonably managed through post-processing. Second, the method is designed to take advantage of the parallelized structure of ATHENA, which dramatically reduces the amount of time required to execute the necessary operations.

The algorithm takes the following steps: (1) On each processor, or local grid, the total mass and angular momentum with respect to the origin are calculated within a search radius, $R_{\rm search}$. For simulations with refinement, overlapped cells are excluded from this total. The net angular momentum ($\mathbf{L}_{\rm net}$) and mass ($M_{\rm net}$) are then calculated across all processors. (2) For each cell, the deviation angle is calculated as

$$\theta_{\rm dev} = \left| \cos^{-1}\left(\frac{\mathbf{L}_{\rm net}\cdot\mathbf{L}_{\rm cell}}{|\mathbf{L}_{\rm net}||\mathbf{L}_{\rm cell}|}\right)\right| \tag{4.3.10}$$

where $\mathbf{L}_{\rm cell}$ is the angular momentum vector of the cell. (3) The total mass for cells with $\theta_{\rm dev} < \theta_{\rm thresh}$ is then calculated. For our simulations, we found that setting $\theta_{\rm thresh} = 30°$ was a sufficiently strict condition for capturing the formation of a disk. (4) A mass fraction is then calculated as

$$f_M = \frac{M(\theta_{\rm dev} < \theta_{\rm thresh})}{M_{\rm net}} \ . \tag{4.3.11}$$

For our models, we consider $f_M > 0.75$ to be indicative of a potential disk. (6) For disk candidates, we continue to a three-dimensional rotation of the mesh from the simulation reference frame into a reference frame in which the $\hat{z}'$ axis is parallel to $\mathbf{L_{net}}$, which we call the disk frame. We construct a new mesh in the disk frame with identical dimensions and hierarchical structure to the simulation frame of reference. We loop over all cells on the disk frame mesh and calculate the corresponding position in the simulation frame:

$$\mathbf{x} = \begin{pmatrix} \cos\phi_L & -\sin\phi_L & 0 \\ \sin\phi_L & \cos\phi_L & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \cos\theta_L & 0 & \sin\theta_L \\ 0 & 1 & 0 \\ -sin\theta_L & 0 & \cos\theta_L \end{pmatrix} \mathbf{x}' \tag{4.3.12}$$

With $\mathbf{x}$ and $\mathbf{x}'$ representing simulation frame and disk frame coordinates, respectively. The angles used in the rotation are computed from the net angular momentum vector:

$$\theta_L = \tan^{-1}\left(\frac{\sqrt{L^2_{\rm net,x} + L^2_{\rm net,y}}}{L_{\rm net,z}}\right) \tag{4.3.13}$$

$$\phi_L = \tan^{-1}\left(\frac{L_{\rm net,y}}{L_{\rm net,x}}\right) \tag{4.3.14}$$

(7) The eight cells surrounding $\mathbf{x}$ in the simulation frame are used to tri-linearly interpolate conserved variables in the disk frame. The interpolation is only performed if at least one of the surrounding cells in the simulation frame has $\theta_{\rm dev} < \theta_{\rm thresh}$. (8) We use the disk frame mesh to calculate the average orbital velocity, sound speed, and column density of the disk.

4.4 Results

4.4.1 Dynamics

Figure 4.3 shows the time evolution for models C1-C, C1-RL, and C1-RH. In the control model, C1-C, inflowing gas is gravitationally focused by the SMBH, leading to stream collisions and partial angular momentum loss. From $\approx 10 - 30$ kyr, residual angular momentum from collisions leads to the onset of disk formation. From $\approx$ 30–70 kyr the disk continues to accumulate mass from a post-collision stream. Angular momentum accretion only mildly alters the orientation of the disk over time. An eccentric disk forms within ≈ 1 pc of the origin. High surface density streams form in the central 0.5 pc, but these features are transient

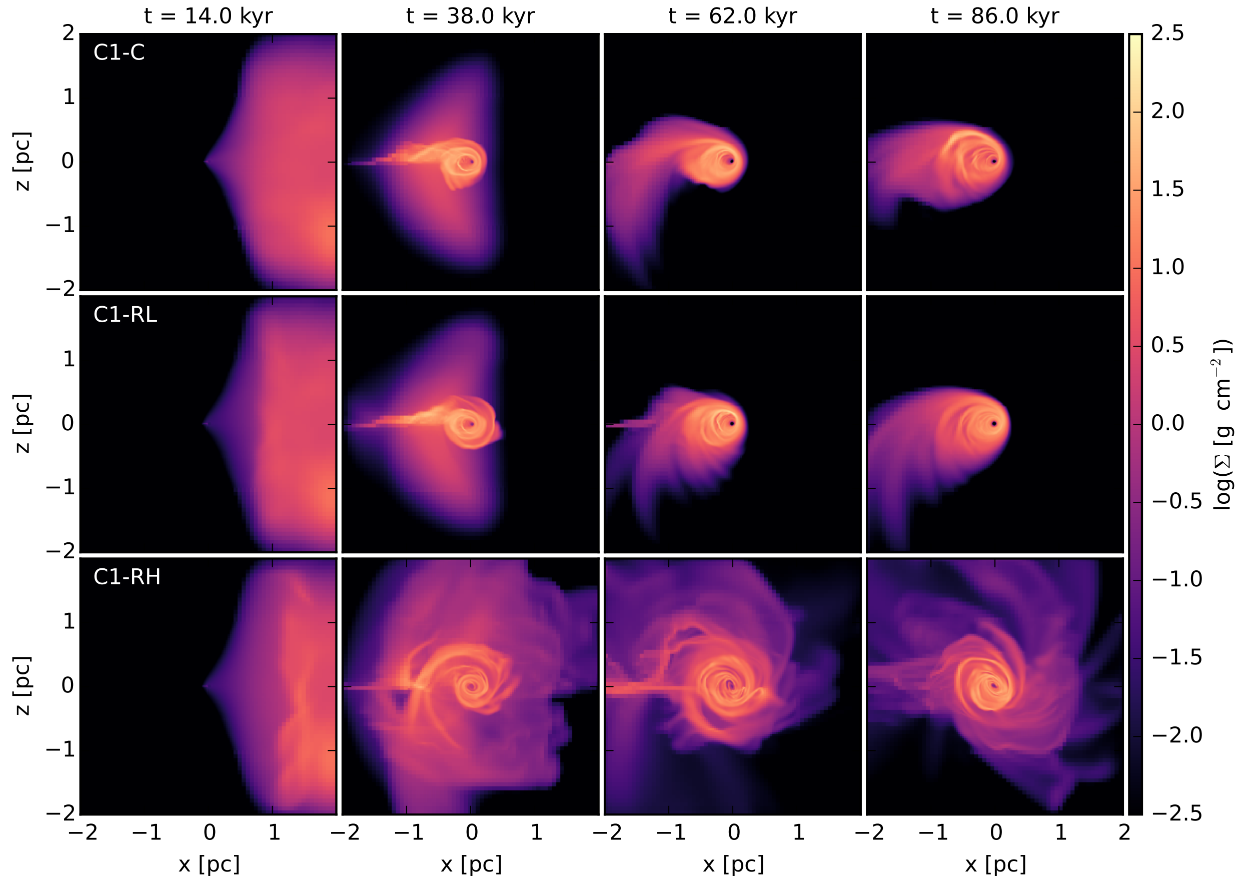

Figure 4.3: Gas column density in a time sequence of models C1-C (top), C1-RL (middle) and C1-RH (bottom). Time increases from left to right, and radiative strength increases from top to bottom.

due to strong shearing.

The time evolution of C1-RL ($L = 0.1L_{\rm edd}$) is nearly identical to C1-C with only minor morphological differences. During inflow, competing forces from radiation and the gravitational pull of the SMBH compress the irradiated face of the inflowing gas, leading to a build-up of mass at x $\approx$ 1 pc. Stream collisions provide sufficient angular momentum loss for disk formation from $\approx$ 10–30 kyr. The sub-structure of the disk differs from the control model for two reasons. First, compression of the inflowing gas leads to partial angular momentum loss that is only supplied through collisions in C1-C. Second, continued photo-heating and photo-compression of unobscured streams of gas cause the disk to be slightly more extended. At $t = 62$ kyr, high surface density streams are seen within the central 1 pc. At $t = 86$ kyr, a dense central component of the disk is surrounded by a low density stream that extends to the boundary of the computational domain. In comparison to C1-C, a more uniform surface density in seen within the central 0.5 pc.

For C1-RH ($L = L_{\rm edd}$), the time evolution follows the general sequence noted above, though the influence of radiation is more pronounced. Evidence of photo-compression during initial inflow at $t = 14$ kyr extends beyond $x = 1$ pc. The build-up of inflowing gas competing with the impinging radiation leads to higher density sub-structure. A period of disk building occurs from $t \approx 30$–70 kyr, though radiation partially inhibits inflow. As a result, the forming disk is surrounded by a low density envelope of gas. Stream collisions continue to supply mass to the inner 1 pc, supporting a more gradual period of disk growth. This process continues through ≈ 62 kyr. The central component of the disk eventually forms with similar extent to those seen in C1-C and C1-RL, albeit with a different orientation. Peak densities seen in C1-RH are lower than in the previous cases, likely due to the prolonged period of disk building.

In Figure 4.4, we show the time evolution for the C2 (higher mass) inflow models. For C2-C, inflowing streams of gas collide after passing the SMBH, leading to angular momentum loss, and disk formation. A high surface density disk is present at $t = 38$ kyr and continues to accumulate mass for the remainder of the simulation. Throughout this process the disk grows in extent, and peak surface densities rapidly depreciate. Yet, the resulting disk is still both larger and more dense than C1-C. This is consistent with the expectation that higher mass inflow should produce denser structures. High surface density streams are seen within $r \approx 0.5$ pc after the period of disk building. As before, these structures are transient due to strong shear and the absence of a self-gravity.

The effect of radiation in the C2 inflow models also follows the previous trend seen in the lower mass case. As gas inflows towards the SMBH, the irradiated face is compressed. In contrast to the lower mass cases, compression of the inflowing gas is not as uniform.

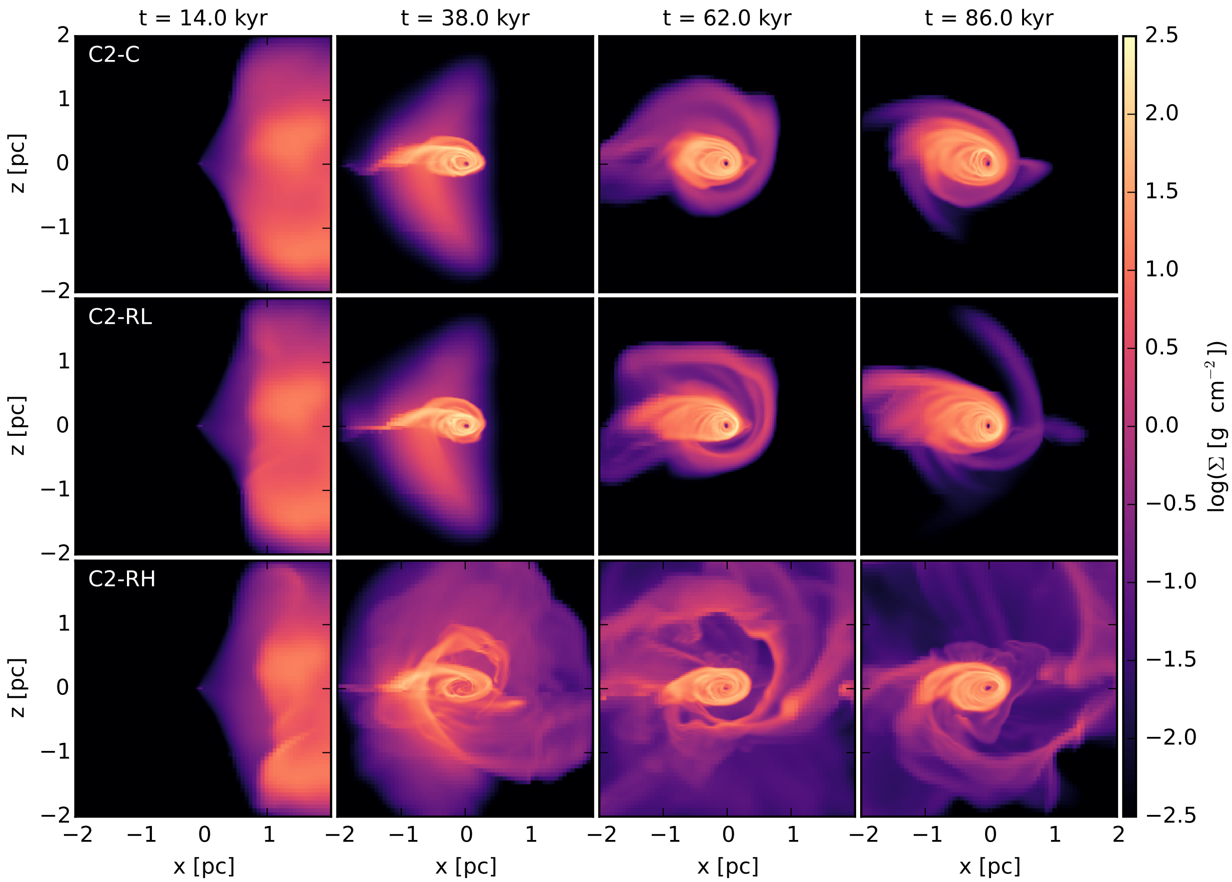

Figure 4.4: Gas column density in a time sequence of models C2-C (top), C2-RL (middle) and C2-RH (bottom). Time increases from left to right, and radiative strength increases from top to bottom.

This is likely due to higher average and peak densities which increase the optical depth. In C2-RL, disk formation occurs in a manner similar to C2-C. The surface density of the disk within $r \approx 0.5$ pc at $t = 38$ kyr is slightly elevated. Evidence of photo-compression is seen at $t = 62$ kyr in a stream at $r \approx 1$ pc surrounding the central disk, but the radiation field is not sufficient to disrupt continued accretion. The resulting disk in C2-RL and C2-C are similar both in extent and structure. For C2-RH, a central disk forms even though stream collisions are less efficient in delivering mass to the central disk. At $t = 38$ kyr, two streams orbit the origin, one which directly feeds the central disk, and another which is radiatively disrupted. At $t = 62$ kyr the later of these two streams extends to larger radii, reducing the rate of accretion onto the outer portion of the disk. However, the central component of the disk is of comparable structure to the lower radiation case, though with slightly different orientation.

4.4.2 Thermal Evolution

We show the mass distribution in density-temperature space over time for C1-C, C1-RL and C1-RH in Figure 4.5. The times correspond to those used for the snapshots shown Figure 4.3. The top row of the figure shows the thermal evolution of C1-C. At $t = 14$ kyr, inflowing gas has reached the central SMBH. A majority of the mass congregates at density-dependent equilibrium temperatures for neutral gas (i.e. $\Lambda_n = \Gamma_n / n_{\rm HI}$) so that the equilibrium curve appears as a thin line which asymptotes to 5×10^3 K at low density and to 100 K for high densities. A transition between temperature extremes occurs between $n = 10^3$ cm^{-3} and $n = 10^8$ cm^{-3}. Adiabatic pressure response ($P \propto \rho^\gamma$) to gravitational compression or shocks drives a fraction of the mass away from thermal equilibrium. One-dimensional histograms along the density and temperature axes show singular peaks. As inflowing mass forms into a disk from $t \approx 10 - 70$ kyr, a second peak emerges along the density axis, consistent with the high density central disk shown in Figure 4.3. At $t = 86$ kyr, the mass distribution peaks at 10^8 cm^{-3} roughly marking the average disk density. Throughout the simulation a transition from $\approx$ 5800K to 100 K begins at $n = 10^6$ cm^{-3}. Above this density, cooling rates are sufficiently high to maintain the neutral equilibrium temperature. As time increases, a high temperature tail that is seen at $t = 38$ kyr is suppressed as more gas cools to thermal equilibrium.

The left column of Figure 4.5 demonstrates how radiation affects the initial inflow in two ways. First, the irradiated face of the inflowing gas is ionized and photo-heated, thus both C1-RL and C1-RH show an increased amount of mass in the region between the equilibrium curve for neutral gas and the equilibrium temperature of 5803 K for a primarily UV-ionized gas. This change is reflected in the one dimensional temperature histograms which show both high temperature tails, and a peak at 5803 K. The prominence of this peak increases in proportion to the radiation field strength. Second, the effect of radiation pressure on the infalling gas is seen in the one dimensional density histograms. For both C1-C and C1-RL, the density distribution peaks below 10^6cm^{-3}. However, in C1-RH, where photo-compression is strong (Figure 4.3), the density peak migrates towards 10^6 cm^{-3}. Low density ($n < 10^2$ cm^{-3}) gas in both C1-RL and C1-RH reaches to temperatures in excess of 10^6 K. The discrepancy seen at low densities between C1-C and models with radiation is likely caused by the assumption of neutral gas in C1-C which leads to overly-efficient cooling in low density gas. This effect is not likely to be dynamically significant as the pressure of this gas is orders of magnitude lower than the pressure of both the disk and the surrounding gas streams.

In C1-RL, both the temperature and density distributions of the gas roughly mirror C1-C. As the disk builds, the mass fraction of ionized gas drops due to shielding of gas interior to the disk. The density distribution extends to $\approx 10^9$ cm^{-3} at $t = 86$ kyr, showing a preference to higher densities than in the control case. This is likely a consequence of photo-compression occurring during the initial inflow which lowers the angular momentum of the gas and increases the density of inflowing clumps. In contrast, the formation of a high density central disk is delayed in C1-RH. A tail in the temperature distribution extends to 1000 K for

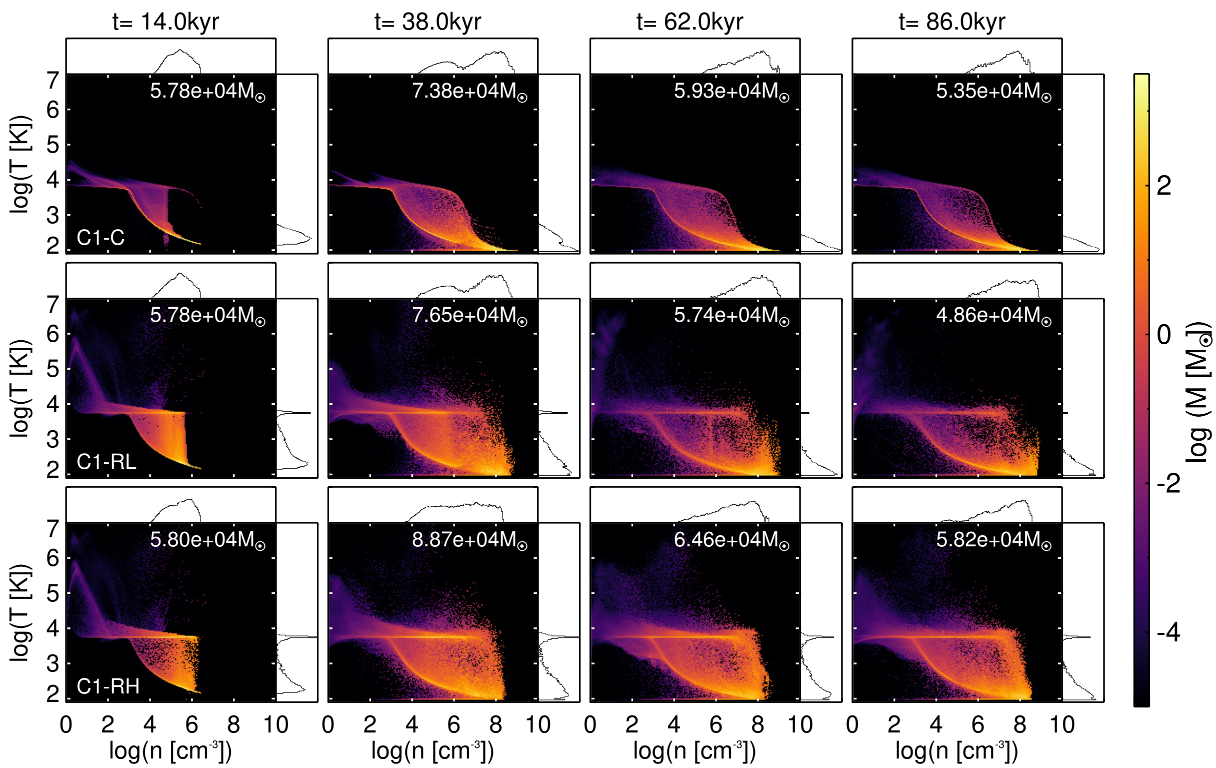

Figure 4.5: Time evolution of the temperature vs. number density mass distribution for C1-C (top), C1-RL (middle), and C1-RH (bottom). Time increases from left to right with times identical to those shown in Figure 4.3. The one dimensional histograms along each axis show the total mass at each temperature or density, and are plotted logarithmically with a range extending from 10^2 $M_\odot$ to $10^{4.2}$ $M_\odot$. The total mass in each histogram is plotted in the top right of each image.

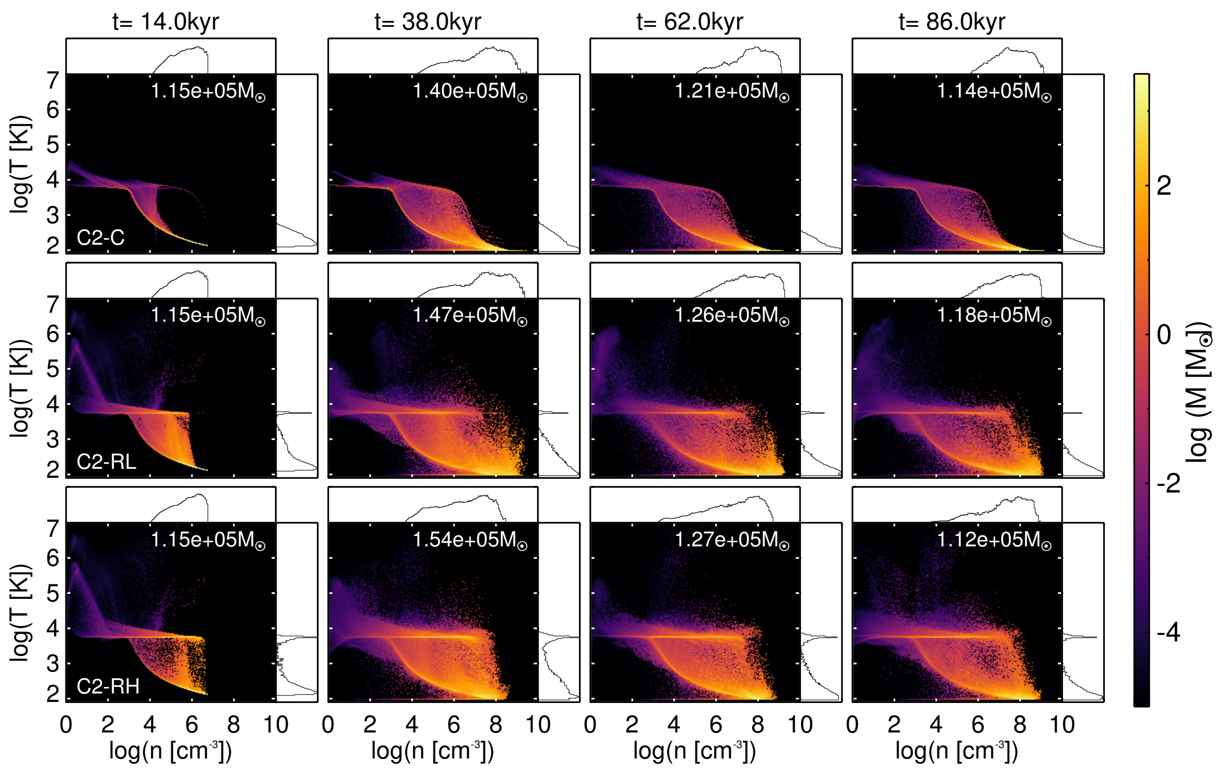

Figure 4.6: Time evolution of the temperature vs. number density mass distribution for C2-C (top), C2-RL (middle), and C2-RH (bottom). Time increases from left to right with times identical to those shown in Figure 4.4. The one dimensional histograms along each axis show the total mass at each temperature or density, and are plotted logarithmically with a range extending from $10^2\ M_\odot$ to $10^{4.2}\ M_\odot$. The total mass in each histogram is plotted in the top right of each image.

the duration of the simulation. A peak is also seen at 5803 K, thus a fraction of ionized gas survives even after disk formation. The density distribution in C1-RH at $t = 38$ kyr shows a suppression of the high density peak. Densities do reach to $> 10^8$ cm^{-3}, but this only occurs at late times. In both radiation models, a turn off is seen above 10^8 cm^{-3}, marking the density above which recombination rates are too high, and shielding too effective, for gas to remain fully ionized.

We show the mass distribution in density-temperature space over time for the C2-C, C2-RL, and C2-RH in Figure 4.6. The results are similar to that seen for the C1 inflow condition. For C2-C, a thermally neutral gas extends to high densities during the disk building phase seen in the top panels of Figure 4.4. The density turn off for dominant cooling at $n = 10^6$ cm^{-3} is also present throughout the time evolution of C2-C. A vast majority of the gas cools to the density dependent equilibrium temperature for neutral gas. In contrast to the lower mass inflow, the density distribution extends to values $\approx 10^9$ cm^{-3} at $t = 38$ kyr indicating the presence of a more compact central disk.

The effect of radiation for the C2 inflow condition is similar to that discussed previously. During the initial inflow ($t = 14$ kyr), photo-heating increases in proportion to the radiation field. Along the temperature axis, mass congregates at both the equilibrium temperatures for UV-ionized and neutral gases. Increases in the density caused by photo-compression are not evident in the density histograms, likely due to the increase of the peak density, and thus shielding, of the gas. For C1-RL, the density distribution follows C2-C, showing a preference towards higher densities. The temperature distribution at $t = 86$ kyr shows a tail extending to 1000 K and a small mass fraction of fully ionized gas. Again, the formation of dense structures is delayed at the Eddington luminosity. C2-RH shows an increase in mass in the parameter space between the neutral thermal equilibrium curve and T=5803 K. This is evident in the temperature histograms which maintain a large fraction of photo-heated and ionized gas for the duration of the simulation. As compared to both C2-C and C2-RL, C2-RH also shows a more extended low density tail which is consistent with the envelope of low density gas in which the central gas disk is embedded (Figure 4.4). Photo-compression during inflow and continued radiative feedback occurring during disk formation inhibit the rapid formation of high density gas disks seen in C2-C and C2-RL.

4.5 Discussion

In the previous section, we qualitatively demonstrated that gas inflow during AGN activity (here implemented as a static UV and X-ray radiation field) may still result in the formation of a central gas disk. The dynamical and thermal evolution of these disks suggests that radiation has two noteworthy effects: (i) photo-heating and ionization drives a small fraction of the disk mass away from neutral thermal equilibrium; (ii) radiation inhibits the initial inflow of material resulting in both increased disk density for weak radiation fields and delayed disk formation for strong radiation fields. Here, we consider the relative impact of UV photons (§ 4.5.2) and X-rays (§ 4.5.1), and the dynamical influence of radiation pressure in disk formation (§ 4.5.3). We discuss the gravitational stability of formed gas disks in our models via conservative measurements of the Toomre Q parameter (§ 4.5.4), and consider the formation of gravitationally bound structure within the disk (§ 4.5.6). We discuss the mass accretion rates measured in our models and estimate the expected radiative feedback for such gas inflow assuming approximate viscous transport timescales for the unresolved accretion flow (§ 4.5.5).

4.5.1 The Effects of X-Rays

In the top row of Figure 4.7, we show the column density of the disk in models C1-X and C2-X for which the radiation field is set to the Eddington limit and UV photons are removed. These snapshots are taken at t = 86 kyr and can be compared to right panel in Figures 4.3 and 4.4. The disks are morphologically similar to models without radiation, suggesting that X-rays do not dramatically affect the formation of the central gas disk. The top row of Figure 4.8 shows the temperature-density distribution for these models at the

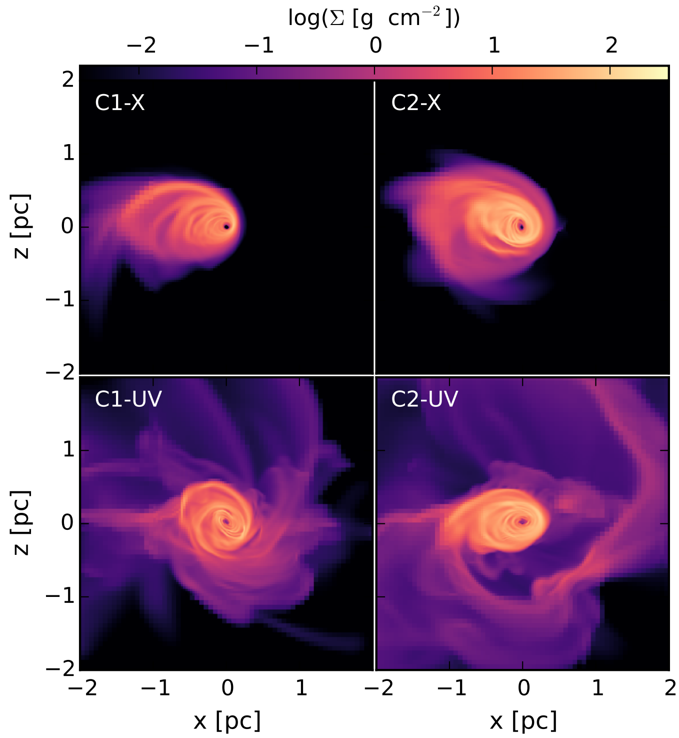

Figure 4.7: Column density in the x-z plane for models C1-X, C2-X, C1-UV, and C2-UV. All snapshots are taken at t = 86 kyr.

same time. The mass distribution along the density axis is most similar to the control models. A collection of mass is seen in between $n = 10^4$ cm^{-3} and $n = 10^7$ cm^{-3} around $T = 10^4$ K. This mass distribution extends to low densities at temperatures in excess of 10^6 K. In the absence of UV photons, mass does not congregate at 5803 K. Furthermore, the equilibrium temperature of an X-ray ionized gas is $\approx 1.267 \times 10^6$K, thus the lack of mass in this temperature regime indicates that X-rays only partially ionize high density gas.

To understand the lack of dynamical influence by high energy X-rays, we can characterize the radiation field in terms of the ionization parameter, $U = Q_{\text{ion}}/(4\pi r^2 c n_H)$, which serves as a measure of the strength of the radiation field. First, we can consider the effect of X-rays during the initial inflow. Taking the photon emission rate to be $Q_{\text{ion}} = Q_{\text{X}} = 10^{51}$ s^{-1}, and assuming a distance of ≈ 1 pc and an average gas density of 10^5 cm^{-3} for the inflow, the ionization parameter is 3×10^{-3}. As stated in Namekata et al. (2014), an ionization parameter

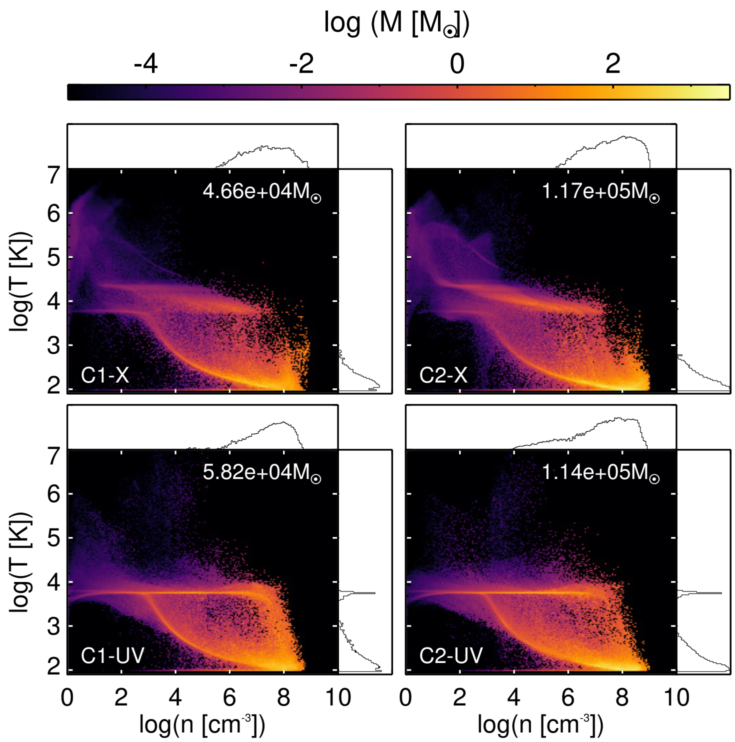

Figure 4.8: Temperature vs. number density mass distribution for C1-UV, C2-UV, C1-X, C2-X at t = 86 kyr. The one dimensional histograms along each axis show the total mass at each temperature or density, and are plotted logarithmically with a range extending from 10^2 $M_\odot$ to $10^{4.2}$ $M_\odot$. The total mass included is plotted in the top right of each figure.

of 10^{-2} is considered a "low" radiation field in which the evolution of the irradiated gas is dominated by photo-evaporation. Therefore, the X-ray flux is not sufficient to significantly influence the inflowing gas via radiation driven outflow. Second, we can consider the impact X-rays have on the formed disk. For the average disk density of 10^8 cm^{-3}, and a minimum distance of $r = 40$ mpc, which marks the disk's inner edge, the ionization parameter is equally low at $U = 2 \times 10^{-3}$. Yet, it is unclear to what extent the conclusions of Namekata et al. (2014) generalizes to nuclear disk structures. It is clear though, that the X-ray photon density is too low to strongly influence the evolution and structure of the disk.

Our results differ from previous models that highlight the importance of X-ray driven compression in gas clouds at distances of ≈ 10 pc from an AGN (Hocuk & Spaans, 2010, 2011). This discrepancy may be caused by several factors. First, the circumnuclear distances considered in this work are over an order of magnitude lower than in previous models, thus the gas suffers from the effect of strong tides. Second, our inflow includes high density substructure for which the optical depth is sufficiently high to rapidly absorb X-rays. Given the photo-ionization cross-section for 1 keV photons of $\sigma_{\mathrm{pi}} = 10^{-23}$ cm^2, the mean free path of a photon in gas clumps with densities of $n > 10^6$ cm^{-3} is ≈ 30 mpc which is less than a cell size on the coarsest resolution level. Lastly, the density of the disk rapidly reaches values much greater than those considered in previous models. With an average density of 10^8 cm^{-3} the mean free path of X-ray photons drops to 0.3 mpc, or a fraction of the cell size on the highest refinement level. Alternatively, as a consequence of the low ionization parameter for X-rays, the Stömgren length ($l_s \approx \frac{Q}{4\pi r^2} \frac{1}{\alpha_B n_H^2}$) within the disk is $\approx 5 \times 10^{-9}$pc, also indicating that X-ray flux is insufficient to penetrate into the disk. This effect is compounded by the fact that the disk is geometrically thin, limiting photon absorption.

4.5.2 The Effects of UV Photons

In the bottom row of Figure 4.7, we show the column density of the disk in models C1-UV and C2-UV for which the radiation field is set to the Eddington limit and X-rays are removed. These snapshots are taken at $t = 86$ kyr. As in the full radiation cases C1-R and C2-R, both of these models show a dense central disk surrounded by low density streams extending to the edge of the computational domain. In the bottom row of Figure 4.8 we show the temperature-density distribution of theses models at $t = 86$ kyr. For both C1-UV and C2-UV, the mass distribution along the density axis is most similar to the full radiation models. For C1-UV, mass congregates at the 5803 K equilibrium temperature for a UV-ionized gas. Gas temperatures also extend upwards to 1000K due to photo-heating on the surface of the disk. In C2-UV, the density distribution shows a low density tail that is also present in C2-R. The temperature distribution extends upwards to 1000 K, with an isolated peak at 5803K. It should be noted that the high-temperature, low-density gas seen in the C1-X and C2-X is not present, confirming that this parameter space is only accessed via X-ray photo-heating. It is clear that UV photons more strongly influence both the formation and evolution of the central disk and provide a more uniform heating of high density gas.

To quantify limits to the effects of UV photons, we follow the same arguments used for X-rays. First, considering the period of initial inflow, we take the photon emission rate to be $Q = Q_{\rm UV} = 10^{55}{\rm s}^{-1}$, assume a density of $10^5{\rm cm}^{-3}$, and a distance of ≈ 1 pc. For these values, the ionization parameter is ≈ 30. Following Namekata et al. (2014), this constitutes a "high" radiation field where the role of radiation pressure becomes dominant, suppressing photo-evaporation. The structure of the initial inflow of C1-R and C2-R shown in Figures 4.3 and 4.4 show evidence of strong compression driven by radiation pressure. Given the considerably low ionization parameter of X-rays, it is clear that UV photons alone dominate the inflow structure. Second, assuming the disk density of $10^8{\rm cm}^{-3}$ and inner edge of 40 mpc, the ionization parameter for UV photons is ≈ 18, suggesting that ionization and heating still occur on the inner portion of the disk.

The Strömgren length of UV photons within the disk is $\approx 6 \times 10^{-5}$ pc, thus strong shielding allows obscured portions of the disk to remain neutral despite the considerable radiation field. We demonstrate this effect in Figure 4.9 where we show the gas number

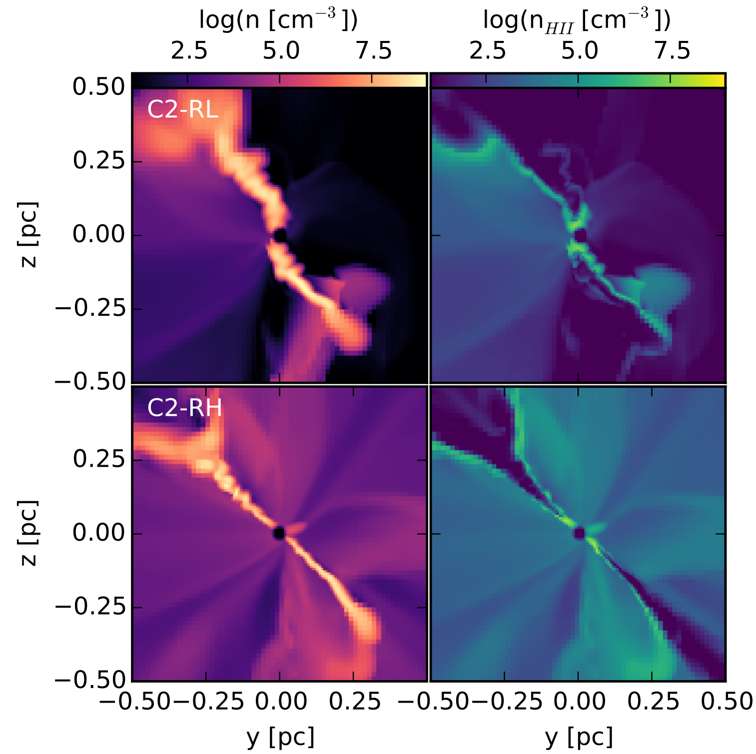

Figure 4.9: Midplane slice in the y-z plane of the total density (left) and ionized hydrogen density (right) in C2-RL and C2-RH. The snapshots are taken at $t = 86$ kyr and are zoomed into the second refinement zone.

density and ionized gas number density through the disk in C2-RL and C2-RH. In both cases, photo-ionization and photo-heating of the disk are restricted to the surface. Because the disks are partially warped, edges of the disk that are not obscured by dense gas are also photo-ionized. Fully ionized, high-density gas is seen at the inner edge of the disk around $r \lesssim 0.1$ pc. Beyond this ionized region, the midplane of the gas disk remains almost completely neutral. Low density gas surrounding the disk is nearly fully ionized with the exception of regions shielded by the central disk.

4.5.3 The Role of Radiation Pressure

In Figure 4.10, we show a midplane slice perpendicular to the inflow direction of the number density for models C1-R, C2-R, C1-NORP, and C2-NORP. In the absence of radiation pressure, a forming disk can be seen in C1-NORP and C2-NORP within the central

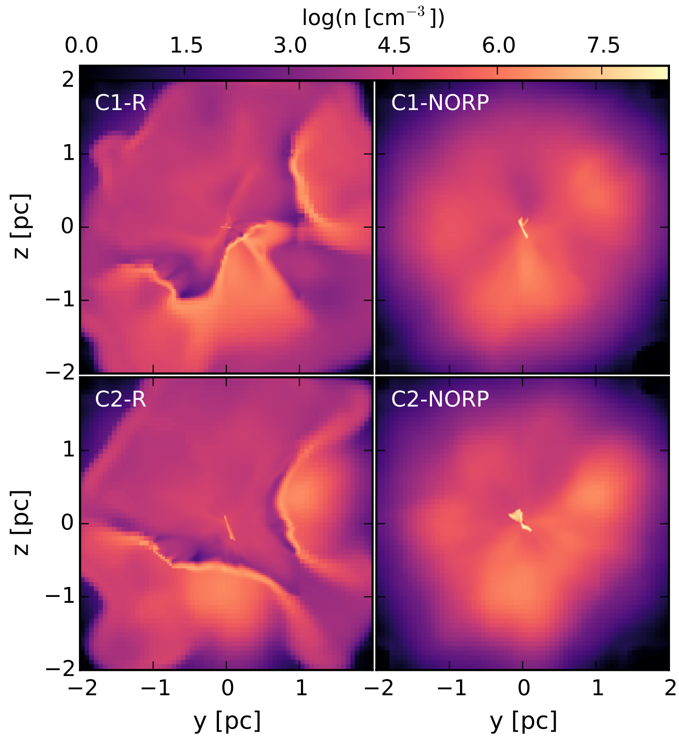

Figure 4.10: Midplane slice of the number density in the y-z plane for C1-R, C2-R, C1-NORP, and C2-NORP. Snapshots are taken at $t = 28$ kyr, during gas inflow.

0.5 pc. In contrast, C1-R and C2-R show that streams of gas which approach the origin are photo-compressed. Only gas with sufficiently low angular momentum and sufficiently high column density with respect to the origin is able to continue the approach towards the SMBH. As shown in the bottom row of both Figure 4.3 and 4.4, streams of gas are forced to larger radii, delaying collisions, angular momentum loss, and the formation of a central disk.

In contrast to both Namekata et al. (2014) and Schartmann et al. (2011), radiation pressure does not completely disrupt the inflowing gas in our models. This is firstly because the inflow models considered here do not begin at rest, thus the gas is exposed to the radiation field for only a fraction of the time. Second, the average gas density in our inflow is roughly an order of magnitude higher than in those previous models, therefore shielding diminishes both the mean free path of photons and the Strömgren length. We note that we do not include self-gravity in our models. As such, we are unable to track the formation of filamentary structures in the photo-compressed gas as is seen in gas cloud models at larger radii. It is clear, though, that radiation pressure dominates the structure of the inflowing gas both inhibiting initial gas disk formation and continued growth.

4.5.4 Disk Formation

In Figure 4.11 we show the disk parameters and Toomre Q parameter for all models. In C1-C, a disk is first found at $t \approx 20$ kyr. Over time, the sound speed of the disk remains roughly constant, indicating that a bulk of the disk material is in thermal equilibrium. The surface density of the disk rapidly increases from 20–40 kyr which is reflected in a increase of total disk mass. Peak surface densities occur between 40–60 kyr, which results in $Q_T \lesssim 2$ during this period. The value of Q_T does not drop below unity during the disk building phase, though our measurements of Q_T are conservative averages and do not account for local density enhancements that are clearly present around the period of minimum Q_T in Figure 4.3.

For C1-RL, a disk is formed at nearly the same time as C1-C. The sound speed rapidly drops to values slightly above the equilibrium value for neutral gas, indicating that the gas disk is only partially ionized. The orbital velocity rises in time, which is reasonable considering the fact that the central disk is less extended at late times than in C1-C. The surface density of the disk increases in a similar fashion to the control model, though it does not drop as steadily. This is likely due to the fact that angular momentum losses also occur during the streams initial approach. The enhanced sound speed of the disk is compensated by the increased surface density, thus the values of Q_T are comparable to the model without radiation. The disk evolution of C1-RH differs from the previous case. The surface density of the disk rises gradually over time, eventually approaching the values seen in C1-C at t=100 kyr. The disk is characterized by $Q_T > 3$ for the duration of the simulation. The sound speed also remains enhanced because of photo-heating.

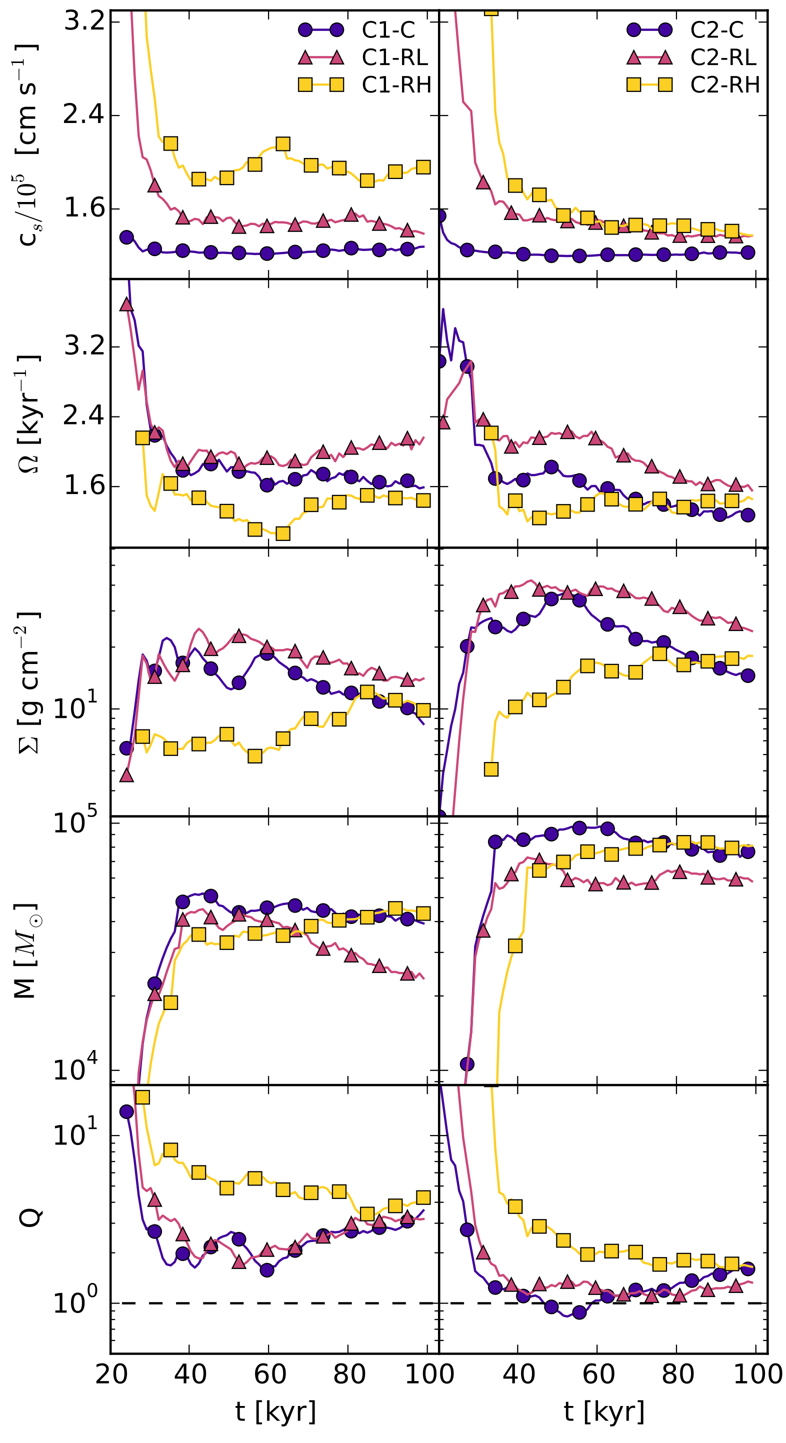

Figure 4.11: Disk parameters and Toomre Q parameter (Q_{T}) for models with the C1 (left) and C2 (right) inflow conditions. The gas sound speed, orbital frequency, surface density, disk mass, and Q_T are shown from top to bottom for models without radiation (purple, circles), with a radiation field 10% of the Eddington luminosity (pink, triangles), and with a radiation at 100% of the Eddington luminosity (yellow, squares).

The disk parameters for the C2 (higher mass) inflow models are shown in the right panel of Figure 4.11. In C2-C, the sound speed of the gas is roughly constant. The surface density increases rapidly from 20–60 kyr, with a maximum occurring at ≈ 55 kyr. For $t > 60$ kyr, the surface density steadily depreciates. At ≈ 55 kyr, $Q_T < 1$, indicating that such conditions may be gravitationally unstable if self-gravity were to be included. The resulting disk mass is roughly twice that seen in C1-C, as expected for the higher mass inflow.

C2-RL shows both higher temperatures and surface densities with respect to the control model. As a result, the values of Q_T follow the trend seen in C2-C. A period of $Q_T < 1$ is not seen for C1-RL, though $Q_T \approx 1$ from $t = 40$ kyr through the end of the simulation. The disk mass in C2-RL is suppressed with respect to C2-C. The sound speed in C2-RH evolves similarly to C2-RL, and is only slightly larger. This is likely due to the fact that the disk which forms in this model is extremely thin (Figure 4.9), therefore very little disk mass is exposed to the radiation source. The surface density of the disk steadily increases over time, giving rise to a central disk with comparable mass to C2-C. The lack of a rapid initial disk building phase results in $Q_T > 1$ for the entirety of the simulation. We do note the downward trend in Q_T, which may indicate a delayed period of $Q_T < 1$ beyond the simulation time.

4.5.5 Mass Accretion

As mass passes through the inner boundary of the computational domain, mass and momentum are slowly removed from the simulation as detailed in § 4.3.1. In Figure 4.12 we show the total mass accreted over time for all models. The figure also shows the mass accretion rate which is averaged in 1 kyr intervals. The general trend is the same for both inflow conditions. The mass accreted in C1-C, C2-C, C1-RL, and C2-RL grows throughout the simulation approaching an accreted mass of $\approx 5 \times 10^4\ M_\odot$, corresponding to 30–50% of the inflow mass. For both C1-RH and C2-RH, the total accreted mass is diminished with respect to the other models, yet the accreted mass still exceeds 10% of the inflow mass in both cases. These values are in agreement with mass accretion reported for high mass

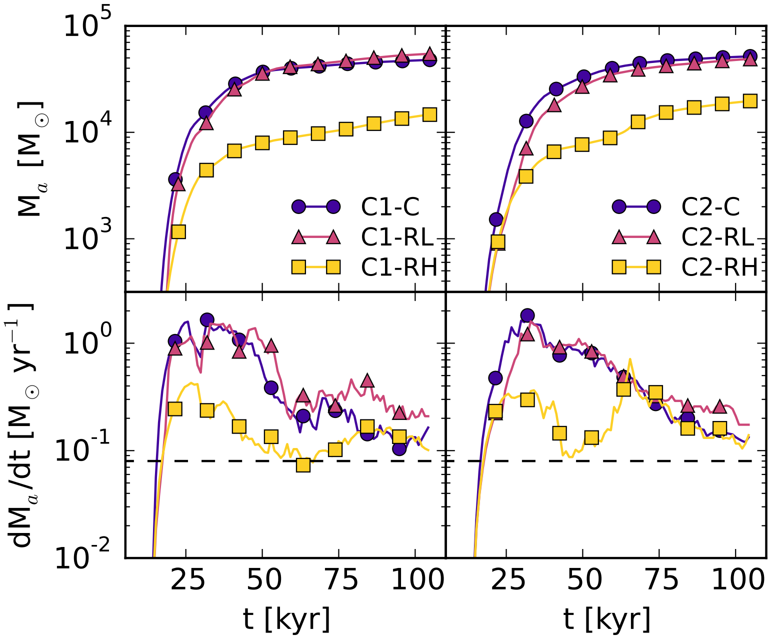

Figure 4.12: Total mass accreted (top) and mass accretion rate (bottom) versus time for models with the C1 (left) and C2 (right) inflow conditions. The horizontal dashed line in the bottom panels is set at 0.08 $M_\odot$ yr^{-1} corresponding to the mass accretion rate required for the Eddington luminosity assuming a 10% radiative efficiency.

misaligned-streamer models shown in Lucas et al. (2013).

For all models, the mass accretion rate exceeds the Eddington rate of 0.08 $M_\odot$ yr^{-1} starting at $t \approx 20$ kyr, well before a dense central disk has formed. Peak accretion rates are over an order of magnitude greater than the Eddington rate, which is consistent with the maximum accretion rates expected for inflow driven by cloud-cloud collision (Hobbs & Nayakshin, 2009). For low radiation fields, high accretion rates are reasonable given that disk formation is uninhibited. In simulations where disk formation is delayed, it follows that mass accretion would also be inhibited as mass is repelled from the origin.

In all of our radiation models, we include a constant radiation field, not accounting for the variability in radiative output expected to be driven by the accretion rate. This is not unlike previous work on the subject (Hocuk & Spaans, 2010, 2011; Namekata et al., 2014; Schartmann et al., 2011) in which the radiation field is assumed to be constant. Yet, in these models, gas clouds are not tracked to the point of the formation of an accretion disk, and in the most extreme case, the radiation source is removed by 50 pc. Our assumption of a constant radiation field may not be crippling, for the following reasons. First, the accretion rates rapidly rise above the Eddington limit at $t = 20$ kyr. At this time, gas is just beginning to collide with the SMBH. Stream collisions provide the necessary angular momentum loss for the onset of disk formation, though a disk is not detected until after the accretion rates have already risen above the Eddington rate. Second, as the disk is rather thin, shielding of the remaining gas inflow is nearly negligible.

Viscous Accretion Timescales

The accretion boundary of our simulations is set at ≈ 40 mpc, roughly twice that used in smoothed particle hydrodynamic (SPH) models (Bonnell & Rice, 2008). Material which enters this boundary should have sufficient angular momentum to form an inner accretion disk which then feeds the central SMBH. The accretion rates measured in our models are therefore instantaneous values and should be considered as generous upper limits for the actual accretion rate onto the SMBH.

To provide a more realistic estimate of the accretion rate, and resulting radiation strength, we consider the viscous timescale on which accretion is expected to occur:

$$t_{\rm visc} = \alpha^{-1}\left(\frac{H}{R}\right)^{-2}\frac{1}{\sqrt{GM_{\rm BH}}}R^{3/2} \ , \tag{4.5.1}$$

where H is the disk height, R is the disk radius, and α is the viscosity parameter (Shakura & Sunyaev, 1973). The value of H/R can be approximated by rearranging the tidal stability criterion (Eq. 4.3.9) to find $H/R \approx M_{\rm disk}/M_{\rm BH}$ (Gammie, 2001). In our models, a disk mass of few $\times\ 10^4\ M_\odot$ yields $H/R \approx 0.01$. The value of α is less certain, but it is often assumed to be in the range of 0.01–1.

Assuming an inner disk radius of a few x 1000 AU, the viscous timescale for $\alpha = 0.01$ is $\approx 10^7$yr. In comparison, the age of the central stellar disk in the GC is only a fraction of this value (Paumard et al., 2006), thus we would expect to see evidence of accretion resulting from this process if accretion occurs on this timescale. For a more liberal estimate of accretion with $\alpha = 0.1$, the accretion timescale is $\approx 10^6$ yr. The accreted mass in our models ranges from $10^4 - 5 \times 10^4$ $M_\odot$. Spreading this accretion over a period of 10^6–10^7yr yields accretion rates of 0.001 – 0.05 $M_\odot$ yr^{-1}, which is approximately 1%–60% of the Eddington limit. This agrees with the findings of both Bonnell & Rice (2008) and Hobbs & Nayakshin (2009). Given these estimates, it is a reasonable assumption that radiative feedback should play a role in the formation of a nuclear stellar disk.

4.5.6 Disk Substructure

For the duration of each simulation, we calculate peak densities in 80 logarithmically spaced radial bins extending from the accretion boundary to the corner of the simulation domain at 10 year intervals. The results of this process are shown in Figure 4.13. Due to the logarithmic spacing, the number of grid cells available to the inner-most bins is low which causes streaks to appear $\log(r/\mathrm{pc}) < -1$ at early times. Furthermore, due to the use of grid refinement, vertical lines can be seen in various locations where gas transitions from low to high grid refinement where compression and cooling of the gas is better resolved.

For C1-C, the inflowing gas stream appears at $\log(r/\mathrm{pc}) > 0$ for $t < 20$ kyr. The presence of the gas disk is clearly seen as a region with peak densities in excess of 10^8 cm^{-3} extending from the accretion boundary to $\log(r/\mathrm{pc}) \approx -0.25$ for $t > 20$kyr. For times $> t = 60$ kyr, peak densities are not contiguous due to the transient high density streams seen in Figure 4.3. The results of C1-RL are similar to C1-C, though peak densities are more continuous at late times, consistent with the more uniform surface density disk that manifests in this model. For C1-RH, peak densities are nearly an order of magnitude lower than in the previous cases, though they are sustained.

The presence of the disk is clearly seen in C2-C, and the densest portion of the disk can be seen between $\log(r/\mathrm{pc}) = -1$ and $\log(r/\mathrm{pc}) = -0.5$ at late times. As the radiation field increases to 10% of the Eddington limit in C2-RL, peak densities are increased in this same region. For C2-RH, the disk density is lower than in the previous case, but peak densities increase over time at $\log(r/\mathrm{pc}) \approx -0.5$, which may be due to ongoing accretion.

Tidal Limit

Though we do not include self-gravity and are unable to monitor the fragmentation of gas to the point of star formation, we can consider the tidal stability of dense gas structures approximately. From the tidal stability condition,

$$\frac{Gm_c}{R_c^2} = \frac{2GM_{\mathrm{BH}}R_c}{r^3} \ , \tag{4.5.2}$$

we can determine the gas density required for gas clump of mass m_c and radius R_c to remain self-gravitationally bound in the presence of a SMBH at a distance of r:

$$\rho_{\mathrm{tidal}} = 1.29 \times 10^{-16}\ \mathrm{g\ cm^{-3}} \Big(\frac{M_{\mathrm{BH}}}{4 \times 10^6\ M_\odot}\Big)\Big(\frac{r}{1\mathrm{pc}}\Big)^{-3} \tag{4.5.3}$$

The corresponding number density, assuming a pure hydrogen gas, is:

$$n_{\mathrm{tidal}} = 7.73 \times 10^7\ \mathrm{cm^{-3}} \Big(\frac{M_{\mathrm{BH}}}{4 \times 10^6\ M_\odot}\Big)\Big(\frac{r}{1\mathrm{pc}}\Big)^{-3} \tag{4.5.4}$$

To compare the peak gas densities in our models to the radially dependent tidal limit, we divide the maximum densities seen in Figure 4.13 by the corresponding tidal density using Eq. 4.5.4. In Figure 4.14 we show the ratio of the peak densities to the tidal density for all models. We exclude bins with $n_{\mathrm{max}}/n_{\mathrm{tidal}} < 0.1$, and use gray-scale for values below unity and colors for super-tidal densities. The inflowing stream again appears in the lower right corner of each plot at $\log(r/\mathrm{pc}) = 0.5$ and $t = 0$ kyr. As gas first collides with the SMBH and stream collisions occur, strong compression and cooling drives the gas density above the tidal limit at $\log(r/\mathrm{pc}) \lesssim 0$ at ≈ 30 kyr in all models. Densities above the tidal limit are not found for $\log(r/\mathrm{pc}) < -0.5$ in any model. Therefore the high densities seen in Figure 4.13

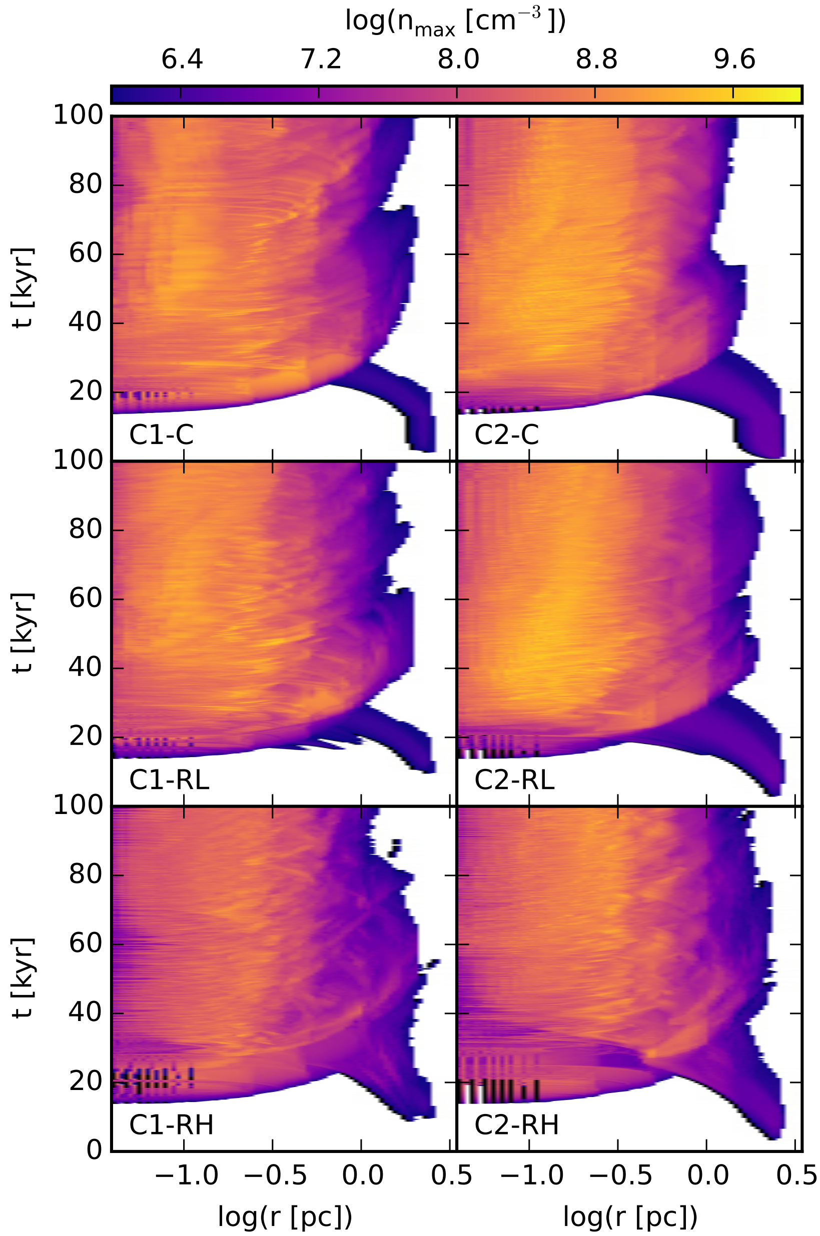

Figure 4.13: Peak density distribution versus time and radius. The left boundary of each figure represents the accretion boundary, and the right edge corresponds to the maximum radius of the computational domain. Vertical lines in the image are a result of gas transition through refinement boundaries. Vertical lines at log(r) < -1 and t < 30 kyr are a result of the sparsely populated radial bins as a consequence of the logarithmic bin spacing.

in this region are well below the tidal limit.

For the C1 inflow condition, an increase in the radiation field decreases the amount of gas near the tidal limit. For C1-C, densities above the tidal limit are found between $\log(r/\mathrm{pc}) = -0.5$ and $\log(r/\mathrm{pc})=0.0$ at $t \approx 80$ kyr. For C1-RL, a stream of super-tidal densities is seen around $\log(r/\mathrm{pc}) = -0.5$ at $t \approx 50$ kyr. As the radiation field increases to C1-RH, the gas disk is completely sub-tidal. Both C2-C and C2-RL are characterized by sub-tidal disks. As the radiation field increases, peak values migrate to larger radii. Periodic super-tidal values are seen at $\log(r/\mathrm{pc}) \approx -0.25$ for $t > 60$ kyr for C2-RH.

The lack of prevalent super-tidal densities in our models is a result of resolution limitations imposed by the exhaustive computational cost of radiative transfer. Because of this we are unable to resolve the cooling length at which dense gas cores are expected to form ($\lambda_{\mathrm{cool}} = c_s \tau_{\mathrm{cool}}$, Iwasaki & Tsuribe (2009)). The volume averaged densities in our models can therefore be treated as lower limits. To demonstrate this effect, we re-ran C2-C at base resolutions of 128^3 and 256^3 (i.e. twice and four times above our fiducial models). We show the peak densities relative to the tidal density for these test models in Figure 4.15. As the resolution increases, the peak densities resulting from both stream collisions and disk formation increase. At a resolution of 256^3, the central disk is characterized by super-tidal structures that are not present at the base resolution of 64^3. However, it should be noted that these results are far from converged. For the average disk density of $10^8 \mathrm{cm}^{-3}$, the cooling length is ≈ 20 μpc, well below our resolution limit of 2.0 mpc. The super-tidal densities seen at a resolution of 256^3 are consistent with $Q_{\mathrm{T}} < 1$ in these models. We cannot perform the same experiment for models with radiation because of current computational limitations.

4.6 Conclusions

Near radial gas flow towards nuclear SMBHs is thought to provide the most viable mechanism for the formation of a nuclear stellar disk on sub-parsec scales in the Galactic Center (Lucas et al., 2013; Bonnell & Rice, 2008; Mapelli et al., 2012; Alig et al., 2011).

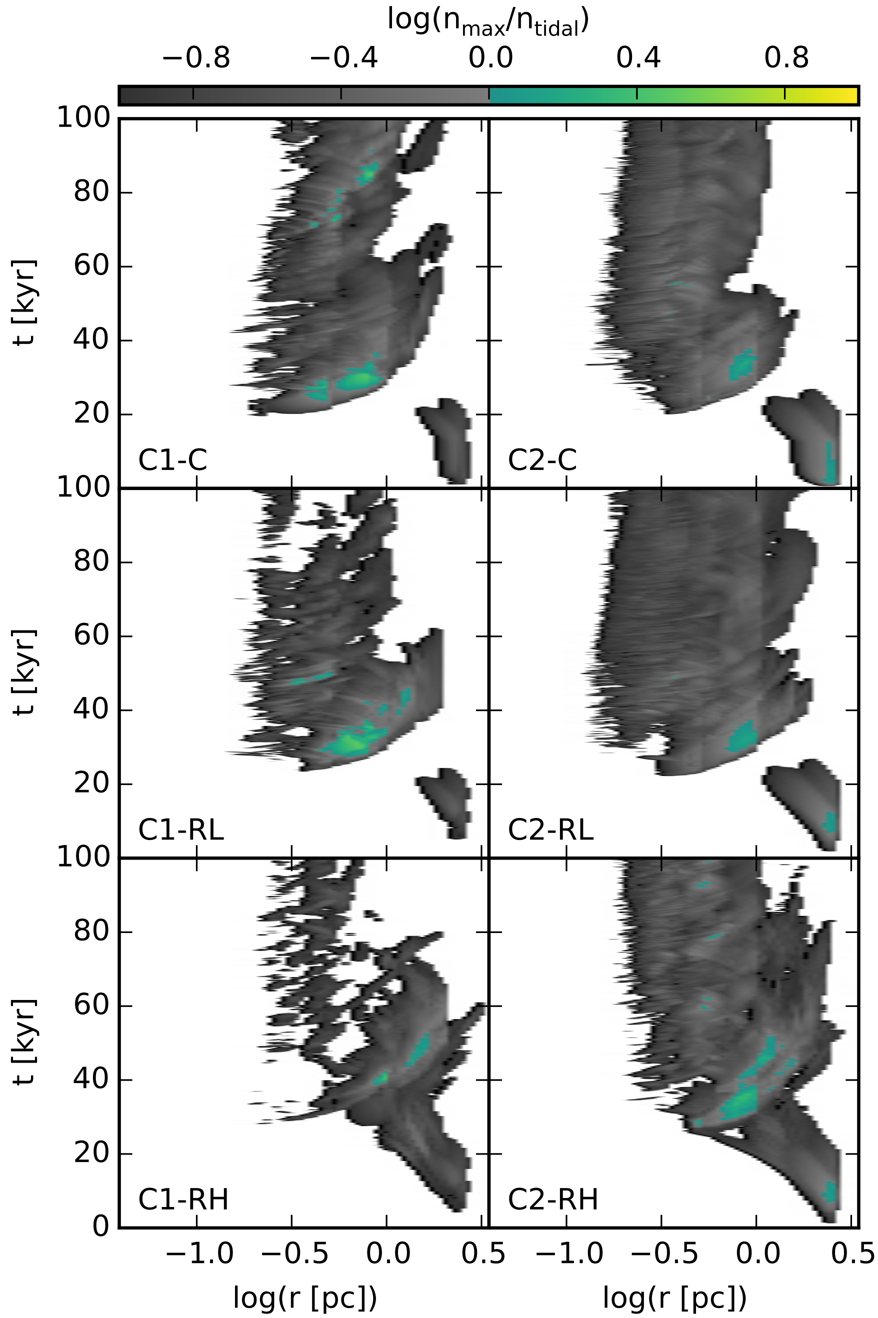

Figure 4.14: Peak density distribution versus time and radius scaled to the radially dependant tidal density for both C1 and C2 models. The left boundary of each figure represents the accretion boundary, and the right boundary marks the maximum radius from the origin on the computational domain. Radiation strength increase from top to bottom, beginning with the control models.

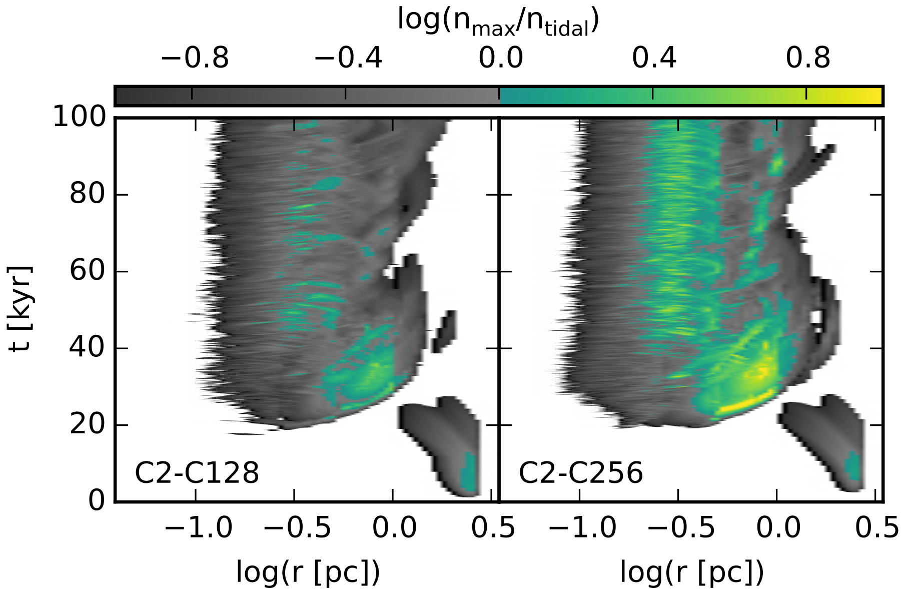

Figure 4.15: Same as Figure 4.14 for models C2C-128 and C2C-256.

Resulting accretion rates onto the SMBH throughout this process can surge to large fractions of the Eddington limit (Hobbs & Nayakshin, 2009; Bonnell & Rice, 2008), suggesting that radiative feedback due to accretion could affect the disk formation and evolution. We present the first 3D RHD simulations following the infall of massive gas streams onto a 4×10^6 $M_\odot$ SMBH. The gas streams are exposed to a constant radiation field at 10% or 100% of the Eddington luminosity, approximating radiative feedback due to accretion. We consider inflow masses of $\approx 10^5$ $M_\odot$, and include the effects of ionization, photo-heating, and radiation pressure in our models. We find the following:

1. A direct collision between a SMBH and a clumpy gas stream can produce gas disks that are characterized by $Q_T \lesssim 1$ (Figure 4.11). This suggests that the disks may be gravitationally unstable, and that an episode of star formation may occur for such a scenario.

2. At 10% of the Eddington luminosity, the effects of radiation do not strongly influence the process of central disk formation (Figures 4.3 and 4.4). Photo-heating increases the average disk temperature, but an increase in surface density maintains conservative

estimates of $Q_T \approx 1$ (Figure 4.11), consistent with models that do not include radiation.

3. At the Eddington luminosity, radiation pressure from UV photons delays the formation of a dense gas disk (Figures 4.3 and 4.4). Radiative forces are not sufficient to drive gas from the vicinity of the SMBH, thus a disk builds gradually as mass accumulates along shielded lines of sight. Peak surface densities are below those seen for lower radiation fields (Figure 4.11) thus the values of Q_T are larger. Yet, for higher mass inflow, super-tidal densities are seen within the disk even at low resolution thus indicating that star formation may still be possible (Figure 4.14).

4. The instantaneous accretion rates of all models are in excess of the Eddington luminosity (Figure 4.12), although viscous timescale constraints suggest that the resulting radiation strength may fall in the range of 1% – 60% of the Eddington limit. These estimates do not account for on-going mass accretion beyond the simulation time.

5. We lastly note that our results regarding gravitational instability are conservative due to resolution effects. Firstly, we are unable to observe the cooling lengths at which dense structures are expected to form. As a result, the densities observed in the disk are mostly sub-tidal (Figure 4.14), but these values are not converged (Figure 4.15). An increase in resolution is computationally prohibitive for models with radiation.

We conclude that star formation occurring via the collision of inflowing gas streams and the formation of a central gas disk may still occur despite radiative feedback from AGN activity.

CHAPTER 5: ECCENTRIC GAS CLUMP ORBITS AND STAR FORMATION IN THE GALACTIC CENTER: THE ROLE OF IONIZING RADIATION

5.1 Introduction

5.1.1 Star Formation in the Galactic Center

Star formation in the central parsecs of the Milky Way's Galactic Center (GC) should be inhibited by the gravitational influence of the $4 \times 10^6\ M_\odot$ (Gillessen et al., 2009a) supermassive black hole (SMBH) (Morris & Serabyn (1996)). Yet, hundreds of young stars are found in this region, many of which belong to either the so-called "S-star" cluster ($r \lesssim 0.04$ pc) (Ghez et al., 2003; Gillessen et al., 2009a; Eisenhauer et al., 2005; Schödel et al., 2002) or a stellar disk extending from $0.05 - 0.5$ pc (Genzel et al., 2003; Paumard et al., 2006; Lu et al., 2009, 2013; Bartko et al., 2010, 2009; Yelda et al., 2014). For in situ star formation, progenitor gas structures must have densities in excess of the tidal limit (n$\gtrsim 10^8$ cm^{-3}, Eq. 4.5.4) to remain gravitationally bound. In contrast, the gas densities characteristic of this region ($n < 10^6$ cm^{-3}) can be many order of magnitude below this limit (Zylka et al., 1990; Coil & Ho, 2000; Lee et al., 2008).

One explanation for the presence of a young stars proposes that an infalling $\approx\ 10^5\ M_\odot$ gas stream led to the formation of a gas disk that fragmented to form stars (Sanders, 1998; Lucas et al., 2013; Bonnell & Rice, 2008; Mapelli et al., 2012; Alig et al., 2011; Wardle & Yusef-Zadeh, 2008). Yet, models which consider this process are unable to account for the presence of individual star clusters that appear coincident with the disk (i.e IRS-13E (Schödel et al., 2005; Mužić et al., 2008), IRS-16SW (Lu et al., 2005)). Furthermore, the discovery of possible young stellar objects (YSOs) in the IRS-13N complex (Eckart et al., 2004; Mužić et al., 2008) indicate that these objects may be the remnants of an alternative in-situ star formation process (Jalali et al., 2014).

Isolated star formation episodes may occur in the GC through the gravitational compression of high density gas clumps orbiting through the central parsec (Eckart et al., 2004). Jalali et al. (2014) argue that this process naturally occurs on a timescale of 10 Myr^{-1} through clump-clump collisions occurring within the circumnuclear disk (CND, (Christopher et al., 2005; Requena-Torres et al., 2012)). Using smoothed-particle-hydrodynamic (SPH) models they show that orbital compression occurring during pericenter passage drives gas above the tidal limit. The elongated gas stream then fragments to form clusters of stars which are consistent with the YSO cluster in IRS-13.

5.1.2 Motivation

There is considerable evidence for on-going star formation activity in the central parsec of the GC (i.e. SiO emission indicative of protostellar jets (Yusef-Zadeh et al., 2013), several protoplanetary disk candidates (Yusef-Zadeh et al., 2015, 2016), bipolar outflows (Yusef-Zadeh et al., 2017)) that cannot be explained through massive gas disk fragmentation. Currently, the star formation mechanism explored by Jalali et al. (2014) constitutes the only alternative explanation for the presence of YSOs in the central parsec of the GC. Here, we extend their models by including an approximate thermal physics prescription and by incorporating the effect of ionizing radiation from existing massive stars in the GC.

5.1.3 Rationale for Including Radiation

We can qualitatively describe the effect of stellar radiation on in-falling gas clumps by first calculating the Strömgren length of ionizing photons. Assuming the central stellar cluster can aptly be modelled as a point source at a distance r, the Strömgren length to which photons penetrate is:

$$\lambda_{\mathrm{s}} = \frac{Q_\star N_\star}{4\pi r^2} \frac{1}{\alpha_B n_{\mathrm{H}}^2} \ , \tag{5.1.1}$$

where $Q_\star \approx 10^{49}$ s^{-1} is the approximate emission rate of ionizing photons per star (Sternberg et al., 2003), $N_\star \approx 100$ is the number of UV bright stars, and α_B is the case B recombination coefficient. We take the clump distance to be $r \approx 2$ pc, coincident with the inner edge of the Circumnuclear Disk (CND; Christopher et al. (2005)). The gas clump density of $n_\mathrm{H} = 10^5$ cm^{-3} results in $\lambda_s = 10^{-5}$ pc. The gas clumps modelled in Jalali et al. (2014) have radii of $R_\mathrm{c} \approx 0.2$ pc, thus ionizing photons are confined to the surface of the clump. Interior gas is shielded from the radiation field.

For the irradiated gas, the effect of the ionization can be quantified via the ratio of the ionization timescale to the recombination timescale:

$$\frac{\tau_\mathrm{ion}}{\tau_\mathrm{rec}} = \frac{4\pi r^2 n_H \alpha_B}{Q_\star N_\star \sigma_\mathrm{ion}} \tag{5.1.2}$$

For the conditions listed above, the value of this ratio is ≈ 0.05, indicating that ionization dominates over the recombination of the gas. The front of the clump will be ionized and photo-heated, forming a pressure gradient that results in two effects. First, the gas will be over pressured with respect to the ambient gas, leading to a photo-evaporative flow (Namekata et al., 2014). Second, the pressure gradient on the leading edge of the cloud will resist the tidal field (Yusef-Zadeh & Wardle, 2017). Assuming constant density, so that the pressure gradient is solely driven by the temperature gradient, the ratio of these two forces is given by:

$$\frac{F_\mathrm{pres}}{F_\mathrm{tidal}} \approx \frac{k_B \nabla T}{m_H} \left(\frac{2 G M_\mathrm{BH} R_\mathrm{c}}{r^3} \right)^{-1} \tag{5.1.3}$$

For an ionized gas temperature of 5803 K (see Chapter 3), a gas clump temperature of 275 K, a black hole mass of $4 \times 10^6\ M_\odot$, and a length scale equivalent to the Strömgren length, the value of this ratio is $\approx 2.5 \times 10^3$. This indicates that heating from the central stellar population may partially inhibit tidal stretching of gas clumps during their approach to the SMBH.

5.1.4 Outline

We investigate the effect of photo-compression on in-falling gas clumps in the GC via radiative hydrodynamic (RHD) models. Our simulation set-up and parameters are detailed in § 5.3. In § 5.4, we discuss the dynamical and thermal evolution of our models. We quantify the star formation potential of this process in § 5.5 and conclude with an overview in § 5.6.

5.2 Methods

The hydrodynamics and integration scheme used here is identical to that discussed in § 4.2. Again, we do not include self gravity but discuss this simplification further in § 5.5.3. Radiation is included using the radiative transfer routine and thermal physics prescription discussed in Chapter 3. We set the "covering factor" of the ray-tree to $f = 2$ so that at least two rays trace each cell. Following Krumholz et al. (2007), our models assume a monochromatic spectrum of 16 eV (UV) photons for ionizing stellar radiation. We take $\alpha = 0.25$ and $\beta = 0.25$ which set the maximum allowable fractional change in the ionization fraction (per radiation cycle) and internal energy (per thermal sub-cycle), respectively. In addition, we have implemented two changes to the radiation routine for the purposes of this work. The first, "selective integration", allows for the partial integration of the ray-tree in simulations where the region of interest populates only a small fraction of the computational volume. The second, "off-mesh ray tracing", approximates the radiation field for sources which extend beyond the computational mesh . These methods are discussed in § 3.7.3 and § 3.7.4, respectively.

5.3 Simulation Set-up

In all models we use a computational domain with three static mesh refinement (SMR) zones centered on a singular gas clump. Table 5.1 lists the mesh and refinement zone geometries. We imitate the initial conditions of Jalali et al. (2014) such that a gas clump with

Table 5.1: Mesh and Refinement Geometry

Level	Dimensions pc	Resolution	Δx mpc
1	(4, 4, 1)	(64, 64, 16)	62.5
2	(2, 2, 1)	(64, 64, 32)	31.3
3	(1, 1, 0.75)	(64, 64, 48)	15.6
4	(0.5, 0.5, 0.5)	(64, 64, 64)	7.8

a total mass of $M = 100\ M_\odot$ and a radius of 0.2 pc is placed on an eccentric orbit about a $4 \times 10^6\ M_\odot$ (Gillessen et al., 2009b) SMBH. The clump density of $n = 1.2 \times 10^5$ cm^{-3} agrees with conservative upper limits placed on similar gas structures within the Circumnuclear Disk (CND; Requena-Torres et al. (2012); Lau et al. (2013)).

We use a co-moving geometry so that the computational mesh roughly follows the orbit of the gas (see § 5.3.3). Boundary conditions of the mesh allow for outflow but prohibit inflow. We place the SMBH at the origin which is initially situated off of the computational mesh. When the SMBH crosses onto the computational domain, an accretion boundary of 5 cells in radius consumes mass and momentum smoothly (see § 4.3.1). The SMBH may traverse multiple refinement domains, thus the accretion radius is variable and dependent on the SMBH's position. A schematic of our initial condition is shown in Figure 5.1.

We initialize the ambient medium with a constant gas density of $n_0 = 10^{-1}$ cm^{-3}. We assume the ambient gas is completely ionized, thus we set the temperature to $T_0 = 5803$ K, consistent with a UV-ionized gas in our thermal model. The ambient medium is at rest with respect to the SMBH. As the initial conditions are not in hydrostatic equilibrium with the SMBH, the gas will collapse over time. To prevent this, we set the color field of this gas to zero. At the end of every hydrodynamic update, we reset cells with $C < 10^{-10}$ to the static initial condition. This approach is similar to that used by Burkert et al. (2012), though we employ a lower threshold color field because the gas is not convectively unstable.

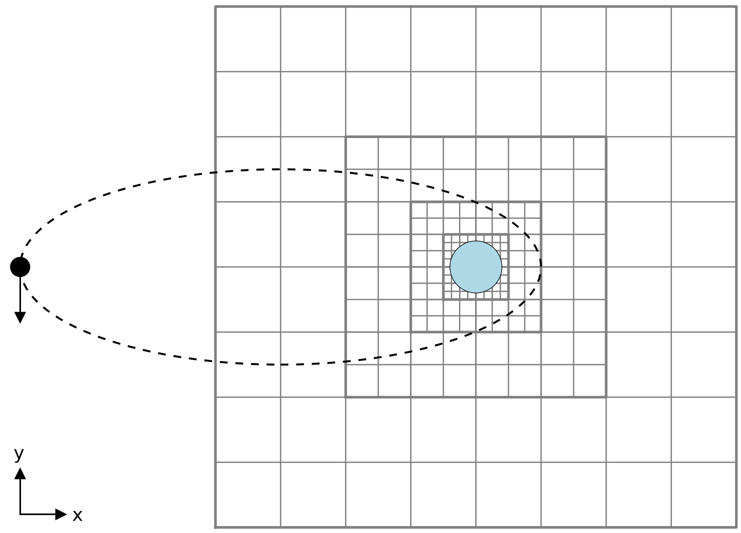

Figure 5.1: Two-dimensional representation of the computational mesh, refinement geometry, and initial conditions used in our models. A gas clump (blue circle) is situated at the center of three nested refinement (four total) zones. A SMBH (black dot) is initially placed off of the computational domain. We use a co-moving frame of reference so that the clump remains centered on the mesh hierarchy during orbit. The orbit of the SMBH in this frame is shown with the black dashed line.

5.3.1 Orbits

The clump follows an eccentric orbit about the SMBH. As in Jalali et al. (2014), we use a semi-major axis of a=1.8 pc for all models, thus the orbital period is uniformly 113 kyr (Jalali et al., 2014). We explore orbital eccentricities (ϵ) of 0.5 and 0.9 for which the pericenter distances are 0.9 pc and 0.2 pc, respectively. The clump is initially situated at apoapsis ($r_{\rm a} = a(1+\epsilon)$), with an initial velocity (i.e. Vis-viva equation, $v = \sqrt{GM_{\rm BH} * \frac{1-\epsilon}{r_a}}$) along the orbital direction. We adopt a naming convention for our models such that the prefix indicates the eccentricity of the orbit, for which the orbital parameters are shown in Table 5.2.

5.3.2 Radiation Fields

We approximate the radiation field of the many massive stars distributed within the central 0.5 pc of the GC (Paumard et al., 2006; Gillessen et al., 2009b; Lu et al., 2013; Yelda et al., 2014) with a singular central source that coincides with the position the SMBH. We

Table 5.2: Orbital Parameters

prefix	ϵ	a	$\mathbf{x}_0$	$\mathbf{v}_0$	$r_p^\dagger$	r_a^*
		pc	pc	km s^{-1}	pc	pc
E5	0.5	1.8	(2.7, 0, 0)	(0, 56, 0)	0.9	2.7
E9	0.9	1.8	(3.4, 0, 0)	(0, 22, 0)	0.2	3.4

†pericenter distance, *apocenter distance

Table 5.3: Radiation Parameters

suffix	$Q_{\rm UV}$ s^{-1}	notes
C	0	no radiation
C128*	0	increase resolution $\times$ 2
C256*	0	increase resolution $\times$ 4
L	1e51	
H	1e52	
J10†	0	$\mathrm{T_{min}} = 10$ (adiabatic)
J50†	0	$\mathrm{T_{min}} = 50$ (adiabatic)
NORP$^+$	1e52	no radiation pressure

† Used to test clump temperature (see § 5.5.2)
* Used for convergence testing (see § 5.5.3)
$^+$ Used to test radiation pressure effects (see § 5.5.4)

run a simulation for each orbit in the absence of a radiation field as a basis for comparison. We denote this simulation with the suffix C. We take the emission rate of 16 eV photons for a single massive star to be $Q_\star \approx 10^{49} - 10^{50}$ s^{-1} (Sternberg et al., 2003), and consider radiation fields equivalent to 100 such stars. The radiation fields used are indicated by the suffixes L (10^{51} s^{-1},"low") and H (10^{52} s^{-1},"high") respectively.

We include two auxiliary runs for each orbit in which the minimum temperature is set to 10 K (J10) and 50 K (J50) for direct comparison with Jalali et al. (2014), though we do not use an isothermal equation of state (see § 5.5.2). To test for convergence, we run two models for each orbit without radiation with a resolution that is two times (C128) and four times (C256) the standard resolution listed in Table 5.1 (see § 5.5.3). We lastly include a model for each orbit (NORP) in which radiation pressure is not included (see § 5.5.4). A list of radiation parameters used in our models can be seen in Table 5.2.

5.3.3 Co-moving Mesh

ATHENA does not include adaptive mesh refinement (AMR), therefore the clump's orbital motion about the SMBH would rapidly remove it from the preset refinement zones. Furthermore, to save on computational time, the mesh is only large enough to allow for the expected tidal stretching and does not contain the entire orbital path of the gas. Therefore, we choose a frame of reference in which the center of the gas clump remains stationary. To do this, we determine the grid cell closest to the center of the clump upon initialization. At the end of every hydrodynamic timestep, we calculate the velocity of this cell in each direction and subtract it globally from the mesh. This is equivalent to performing a Galilean transformation, under which the fluid equations are invariant. To preserve the motion of the clump with respect to the SMBH, we maintain a cumulative mesh velocity. Mesh positions are calculated with respect to the SMBH and are updated at every timestep during this procedure. In the co-moving frame, the "static" ambient gas is assigned to the velocity opposite that of the mesh. This method is similar to that used by Shin et al. (2008), though we do not track the position of the gas clump via the mass averaged velocity. We verify this approach by running comparable simulations that use an extended rest-frame mesh, the dynamical evolution of which shows good agreement with our co-moving mesh models.

5.4 Results

5.4.1 Orbital Evolution

In Figure 5.2 we show the orbital evolution of models E5C, E5L, and E5H. Snapshots are taken at the first refinement level, and the times selected reflect three distinct periods of dynamical evolution for the clump - the approach towards the SMBH ($t = 22.3$ kyr), pericenter passage ($t = 55.7$ kyr), and recession from the SMBH ($t = 89.1$ kyr). In E5C, the clump elongates due to the gravitational pull of the SMBH. Through pericenter passage, the clump forms into a narrow filament at a distance of ≈ 0.9 pc. Peak densities occur during this period as the SMBH compresses the filament vertically and orbital restrictions compress the gas laterally (see discussion in Jalali et al. (2014)). The vertical extent of the filament is limited by the resolution, thus the density increases across refinement zones as compression is better resolved with increasing resolution (see § 5.5.3). This effect imprints the SMR boundary structure onto the gas flow. As the gas returns towards apoapsis, the gravitational pull of the SMBH slows the receding flow, resulting in a pile-up of mass at ≈ 2 pc.

The second and third rows of Figure 5.2 show the effect of UV radiation on the clump's dynamical evolution. During approach towards the SMBH, the irradiated face of the clump compresses in proportion to the strength of the radiation field. A photo-evaporative flow emanates from the clump. The general orbital structure of the gas follows that of E5C. An increase in low density gas surrounding the central narrow filament during pericenter passage reflects the radiative driven mass loss that occurred during infall. This is also seen in the shortening of the stream's leading edge. The peak density of the filament slightly decreases with an increase of radiation strength. Refinement noise appears at the boundaries of each level along the filament as mentioned previously. As the gas recedes from the SMBH, radiation continues to compress the gas, but not considerably for E5L. For E5H, this effect appears to dominate over the flow convergence seen in E5C.

In Figure 5.3, we show the orbital evolution for the higher eccentricity models. Again, the snapshots are taken at the first refinement level during approach, pericenter passage, and retreat from the SMBH. The dynamical evolution follows the general pattern seen in the lower eccentricity case. As the semi-major axes of both orbits are identical, the apocenter distances in the higher eccentricity case are larger. Because of this, tidal effects during the clump's initial approach are not as pronounced, and the gas remains relatively circular. As gas orbits the SMBH during pericenter passage, a narrow filament forms at a distance of ≈ 0.2 pc. Gas on the leading edge of the filament accretes onto the SMBH. In contrast to the previous case, mass does not strongly converge as the gas retreats to apoapsis. The filament

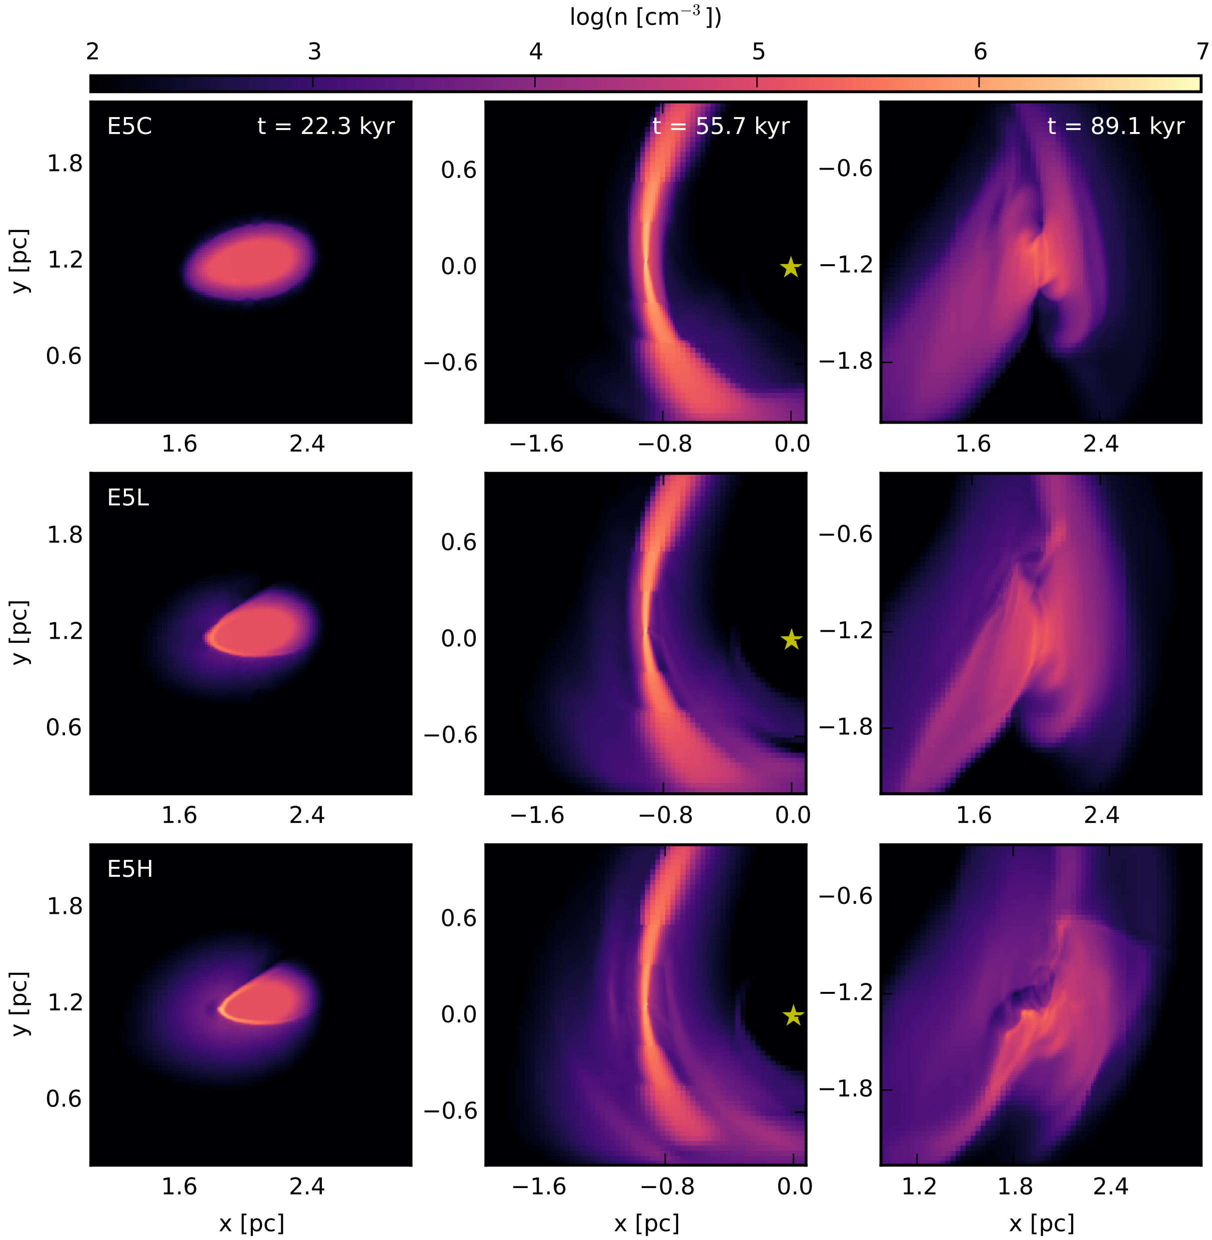

Figure 5.2: Orbital evolution of models E5C, E5L, and E5H. Snapshots are taken at the first refinement level of the computational mesh. Time increases from left to right with individual frames representing the clump approach, pericenter passage, and return to apoapsis. Radiation field strength increases from top to bottom. The coordinates on each figure are given relative to the position of the SMBH (yellow star) which sits at the origin.

extends off the refinement level shown, and reaches beyond the computational mesh. Gas densities rapidly drop as the stream continues to elongate.

Radiation affects the in-falling gas clump as before such that a photo-evaporative flow emanates from the leading edge during approach. The clump is initially situated at a larger distance, thus geometric dilution diminishes of the radiation field and reduces photo-compression. A low density gas envelope surrounds the filament during pericenter passage. The extent of the leading edge of the filament decreases with increasing radiation strength, though the structure of the central stream is unaffected. Tidal stretching of the gas continues after pericenter passage. For E9L photo compression focuses the gas into a narrow stream. This effect becomes destructive for E9H.

5.4.2 Gas Compression

In Figure 5.4 we show the density distribution of clump mass over time. Histograms are shown at the times used in both Figures 5.2 and 5.3. The top left panel shows the evolution of model E5C. During the clump's approach towards the SMBH ($t = 22.3$ kyr), a majority of the mass sits at the initial density of $\approx 10^5$ cm^{-3}. During pericenter passage ($t = 55.7$ kyr) orbital compression drives gas to densities $> 10^7$ cm^{-3}. As gas recedes from the SMBH ($t = 89.1$ kyr), peak densities return to $< 10^5$ cm^{-3}. A larger mass fraction of low density gas is seen at late times due to the tidal shear of the clump. This effect is likely exaggerated as dense structures forming during pericenter passage cannot be retained due to the lack of self-gravity.

The effect of radiation on the low eccentricity clump can be seen by following the first column of Figure 5.4. During the clump's initial approach, photo-compression drives gas to higher densities. For E5H, densities reach values above 10^6 cm^{-3}. Yet, orbital compression occurring during pericenter passage is not as pronounced as in E5C. In particular, the density distribution of E5H does not significantly change, suggesting that photo-compression dominates over orbital compression for this model. As the clump recedes from the SMBH,

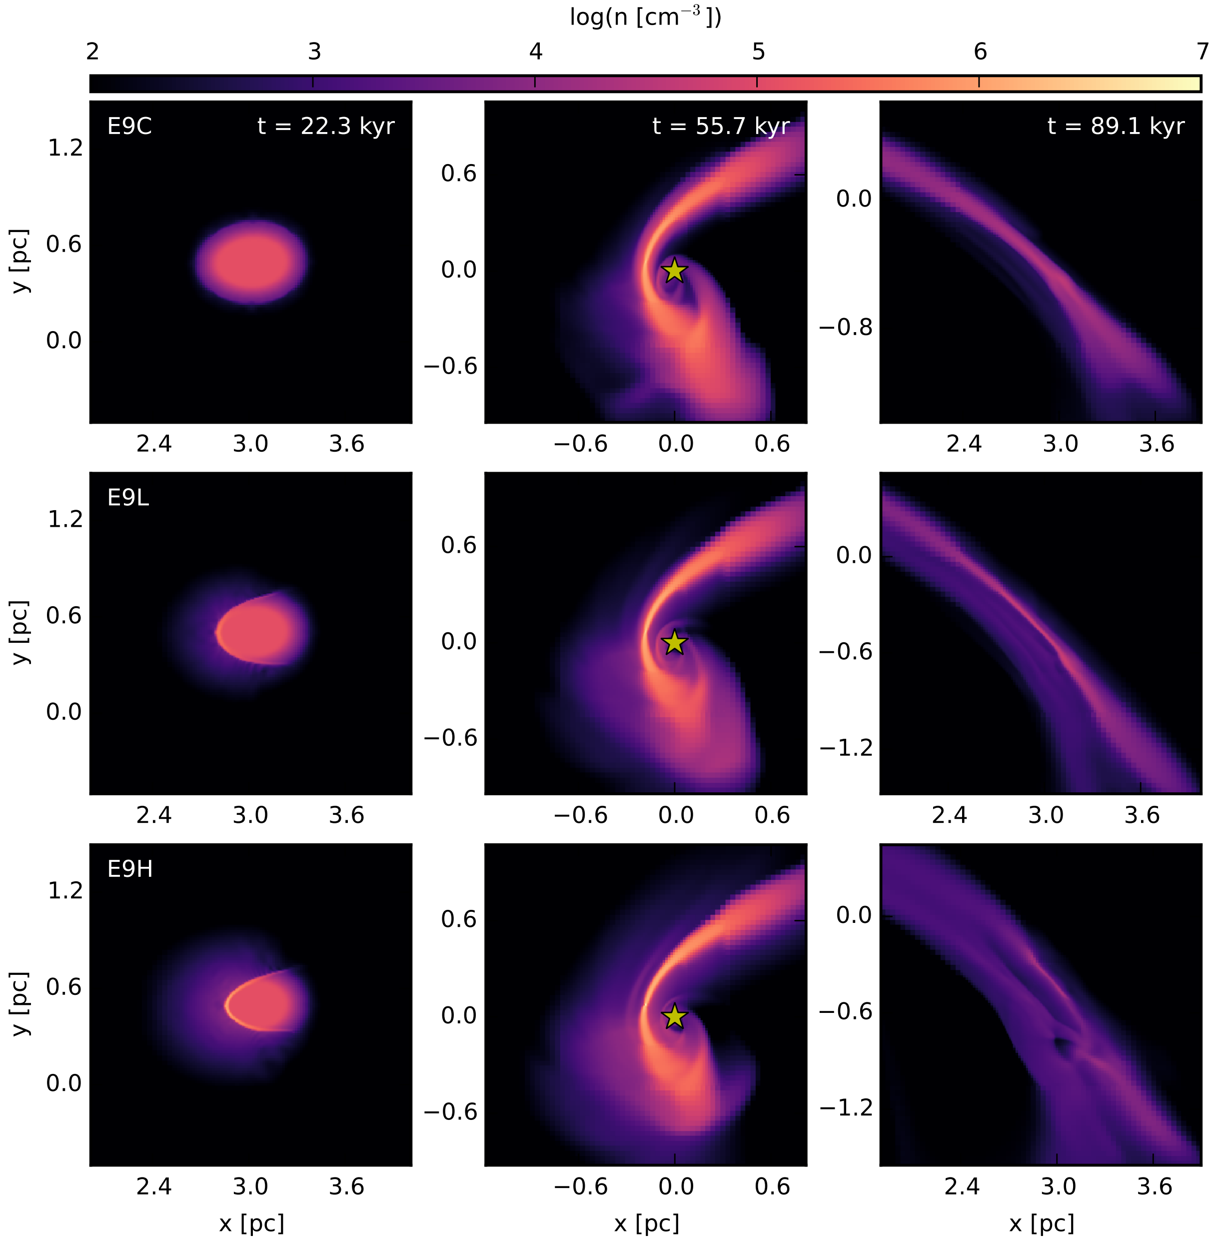

Figure 5.3: Orbital evolution of models E9C, E9L, and E9H. Snapshots are taken at the first refinement level of the computational mesh. Time increases from left to right with individual frames representing the clump approach, pericenter passage, and return to apoapsis. Radiation field strength increases from top to bottom. The coordinates on each figure are given relative to the position of the SMBH (yellow star) which sits at the origin.

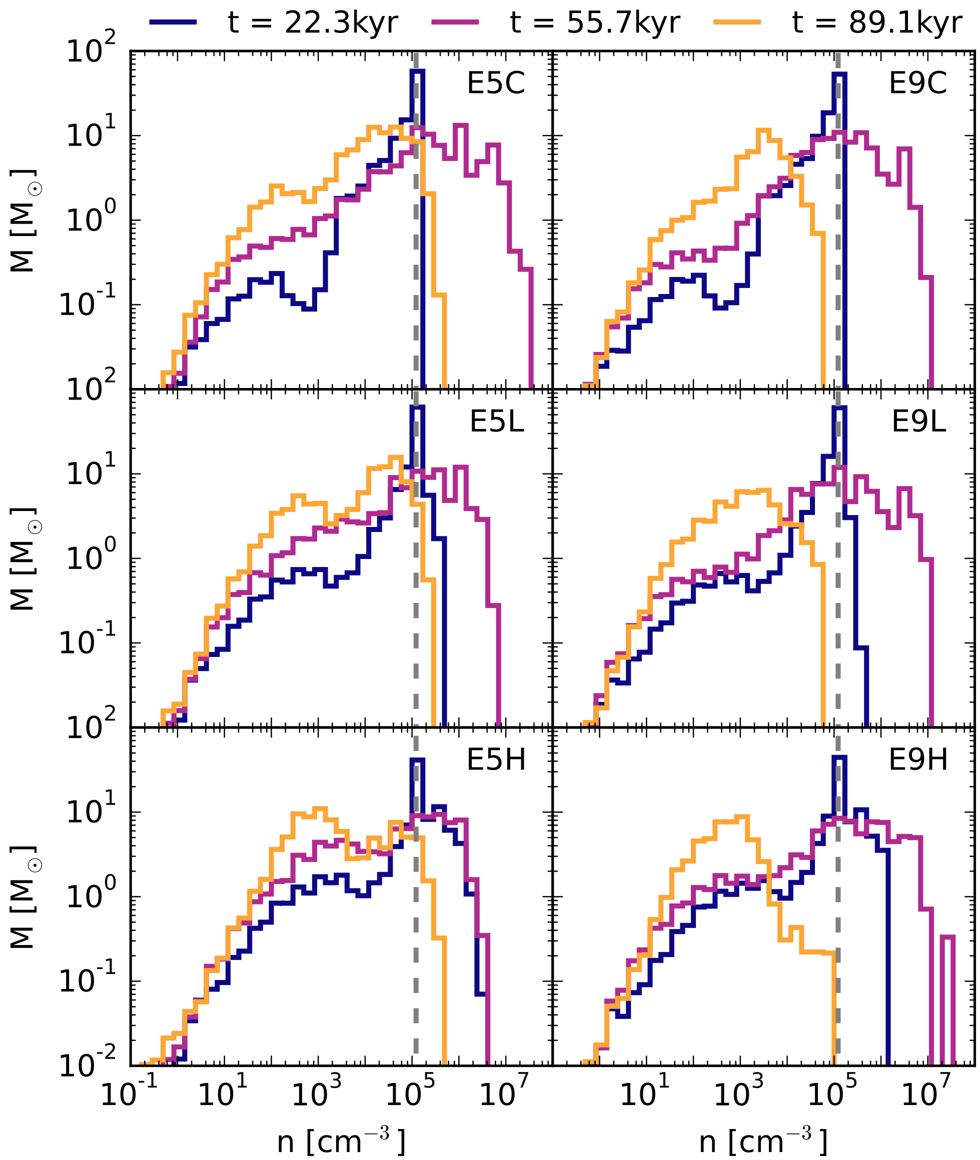

Figure 5.4: Density distribution of clump mass over time. The times selected reflect those shown in Figures 5.2 and 5.3 which include the clump's approach to the SMBH ($t = 22.3$ kyr, blue), pericenter passage ($t = 55.7$ kyr, pink), and retreat from the SMBH ($t = 89.1$ kyr, orange). The radiation strength increases from top to bottom. The left column shows the results for $\epsilon = 0.5$, and the right column corresponds to $\epsilon = 0.9$. The vertical line in each panel shows the initial clump density.

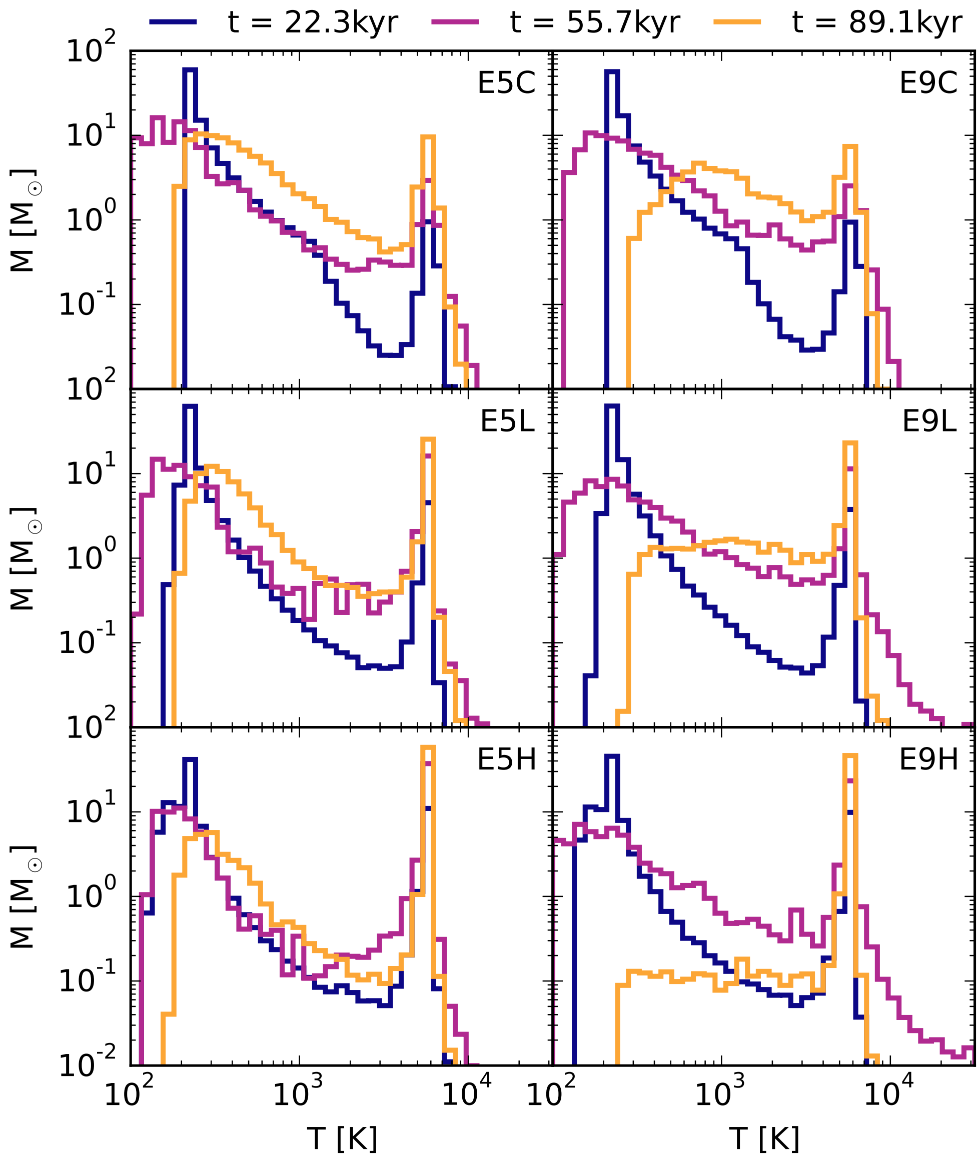

Figure 5.5: Temperature distribution of clump mass over time. The times selected reflect those shown in Figures 5.2 and 5.3 which include the clump's approach to the SMBH (t = 22.3 kyr, blue), pericenter passage (t = 55.7 kyr, pink), and retreat from the SMBH (t = 89.1 kyr, orange). The radiation strength increases from top to bottom. The left column shows the results for $\epsilon = 0.5$, and the right column corresponds to $\epsilon = 0.9$.

the density distribution of both E5L and E5H are similar to E5C, though radiation increases the mass fraction of low density gas.

In the right column of Figure 5.4, we show the density distribution for the higher eccentricity models. At early times, the density distribution of E9C is nearly identical to E5C. During pericenter passage, peak densities reach $\approx 10^7$ cm^{-3}. Despite the considerably stronger tidal shearing expected at a pericenter distance of 0.2 pc ($F_{\rm tidal} \propto r^{-3}$), peak densities are only slightly lower than those seen in E5C. This suggests that orbital compression experienced during pericenter passage is balanced by tidal effects occurring prior to and during this period. Yet, at late times, shearing becomes dominant (see Figure 5.3), thus lower density gas with respect to E5C is seen as the clump approaches apoapsis. Gas densities remain $\lesssim 10^5$ cm^{-3} beyond this time.

The effect of radiation follows the same pattern as before. Photo-compression increases in proportion to the strength of the radiation field and is reflected in several solar masses of gas with densities $> 10^5$cm^{-3}. This effect is slightly less pronounced than in the lower eccentricity case, though this is expected as the time averaged orbital distance is larger, effectively lowering the incident photon flux. Orbital compression produces densities $\gtrsim 10^7$ cm^{-3} in both E9L and E9H. Radiation increases the mass fraction of diffuse gas after pericenter passage. The density evolution of these models suggests that for high eccentricity, photo-compression plays a role only in the initial approach of the gas clump. Yet, gravitational compression, which appears uniformly across the radiation fields explored, remains the dominant mechanism for density enhancement.

5.4.3 Gas Heating and Cooling

In Figure 5.5, we show the temperature distribution of clump mass during the initial approach, pericenter passage, and retreat from the SMBH. Again, the times selected reflect those shown in Figures 5.2 and 5.3. Almost all histograms shown in this figure are bimodal. A low temperature ($\approx$ 200 K) peak reflects the presence of cool dense gas, and a high temperature peak (≈ 5000 K) marks the the presence of both low density hot gas and, for models that include radiative transfer, UV-ionized gas.

The temperature evolution of E5C is shown in the top left panel of Figure 5.5. During the initial approach, as the overpressured clump expands to form a low density envelope, a majority of the clump mass sits at a temperature of ≈ 200 K. Compression during pericenter passage drives up the density, allowing the gas to cool to the minimum temperature of 100 K. The fraction of low-density, hot gas grows over time, increasing the mass fraction at ≈ 5000 K. At late times, a few $\times$ 10 $M_{\odot}$ of mass sits near the initial clump temperature of 200 K, consistent with the converging flow seen at late times in Figure 5.2.

The effect of radiation on the temperature distribution of the low eccentricity models can be seen by following the left column of Figure 5.5. During the inflow, radiation has two effects. First, photo-compression increases the gas density of the clump, allowing for efficient cooling and lower temperatures. Second, photo-ionization and photo-heating increase the mass fraction of hot gas. Shielding of clump material allows a bulk of the mass to remain neutral, though ionization of the clump's surface partially contributes to the high temperature peak. Because of these combined effects, E5H shows both the lowest temperatures at early times and the strongest signature of hot gas.

Orbital compression through pericenter passage increases the mass fraction of cold gas for both E5C and E5L. Yet, for E5H, the gas temperature follows the distribution seen during in-fall. This is consistent with the lack of density enhancement seen for E5H in Figure 5.4. The high temperature peak grows in time as low density gas is stripped from the clump. The temperature distribution after pericenter passage remains bimodal in all three models, though the high temperature peak continues to grow in proportion to the radiation field strength.

For E9C, a low temperature peak migrates to lower temperature during pericenter passage, and a high temperature peak grows in time as the clump is disrupted. For E9L and E9H, orbital compression during pericenter passage results in lower temperatures than is seen in E9C, indicating that photo-compression promotes efficient cooling. In all cases, the extreme tidal field completely eliminates high-density, low-temperature gas. A high temperature peak remains in all cases with a mass fractions that increase with the radiation field. Unlike the low eccentricity case, the resulting temperature distribution is only weakly bimodal in E9C, and is no longer doubly-peaked in E9L or E9H.

5.5 Discussion

In the previous section we qualitatively demonstrated the effects of both orbital and radiative compression of gas clumps on eccentric orbits about an SMBH. We have shown that radiative compression dominates the evolution for low eccentricity orbits. Yet for high eccentricity orbits, gravitational compression more effectively provides for high density gas. In this section, we investigate these compression mechanisms further.

In § 5.5.1, we consider the peak densities of our models relative to the tidal density limit required for gravitationally bound structure formation. In § 5.5.2, we explore the effect of minimum clump temperature on these results, and discuss resolution limitations in § 5.5.3. We discuss further simplifications and the effects of radiation pressure on our models in § 5.5.4.

5.5.1 Tidal Densities

In Figure 5.6, we show the peak density and peak density with respect to the tidal limit (Eq. 4.5.4) over time for all models. In the top left panel we show the peak density for models with $\epsilon = 0.5$. In all cases, values increase as the clump approaches periapsis. For E5C, the maximum value occurs just before pericenter passage. As the gas recedes, adiabatic expansion and tidal shear eliminate dense gas structure. For E5L, photo-compression produces higher densities during the clump's approach, though the maximum density is only marginally larger. The maximum value is achieved ≈ 5 kyr before the control model.

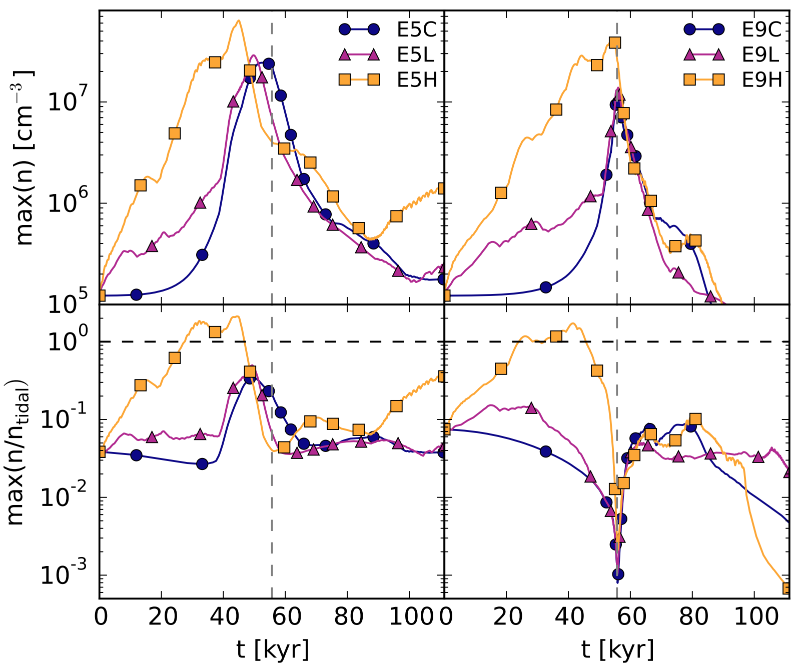

Figure 5.6: Maximum density vs. time (top row) and maximum value of the density divided by the tidal density vs. time (bottom row) for multiple radiation strengths. Results are shown for eccentricities of $\epsilon = 0.5$ (left) and $\epsilon = 0.9$ (right). The horizontal dashed line in the bottom panels at $n/n_{\mathrm{tidal}} = 1$ marks the threshold below which tidal effects are expected to be dominant. The vertical dashed line is placed at the time of pericenter passage ($t = 55.7$ kyr).

Photo-compression dominates in E5H so that peak densities occur well before pericenter passage. Both radiation models show a decrease in density through pericenter passage, though continued photo-compression drives the density to larger values at late times in E9H.

The bottom left panel of Figure 5.6 shows the peak densities with respect to the tidal limit for $\epsilon = 0.5$ (i.e. $\max(\mathrm{n}/\mathrm{n}_{\mathrm{tidal}})$). For both E5C and E5L, peak densities are sub-tidal, indicating that the gas is not gravitationally stable against shearing motion. Yet, this is likely due to a lack of sufficient resolution (see § 5.5.3). Peak values occur slightly before pericenter passage, though this is roughly in agreement with the idea that star formation from orbital compression is spawned during pericenter passage (Jalali et al., 2014). Peak densities in excess of the tidal density occur during infall for E5H, though this occurs well before pericenter passage.

We show the same results for $\epsilon = 0.9$ in the right column of Figure 5.6. Peak densities occur in all three models during pericenter passage. Photo-compression increases the densities of both E9L and E9H at early times. The gas density relative to the tidal limit rapidly drops during pericenter passage where $n_{\rm tidal} \gtrsim 10^{10}$ cm^{-3}. For E9H, peak densities breach the tidal limit during the clump's approach. These results suggest that gravitational compression dominates over photo-compression on eccentric orbits, though the formation of bound structure may begin prior to pericenter passage for sufficiently strong radiation fields.

5.5.2 The Effect of Clump Temperature

We include two models for each orbit in which the minimum temperature of the gas approaches 10K (J10) and 50K (J50) in order to determine the effect of clump temperature. Although these are the temperatures used in Jalali et al. (2014), we do not assume an isothermal gas for our models. We use the thermal prescription with $G_0 = 1000$ for enhanced interstellar heating characteristic of the GC, but set $\beta_T = 1$ and $\beta_T = 5$ in Eq. 3.4.23 to allow for minimum temperatures of 10 K and 50 K, respectively. As a consequence, the initial clump temperature for these two models is 39 K and 150 K, respectively. This modification is equivalent to lowering the cosmic ray heating rate of the gas (Hocuk & Spaans, 2011).

In Figure 5.7 we show the peak densities and peak densities with respect to the tidal limit (Eq. 4.5.4) for these auxiliary models. We include the control case for comparison. Peak densities increase with decreasing temperature for both eccentricities. This is expected for two reasons. First, as the initial temperature of the gas is lower, it is not as strongly overpressured with respect to the ambient gas. Thus, the expansion of the gas seen in Figures 5.2 and 5.3 is not as pronounced. Second, pressure feedback during compression is also lower for the same reason, thus permitting higher gas density to be reached. None of these models show densities above the tidal limit. Yet, for the 10 K models, the ratio of the density to the tidal limit is much higher than in the control cases. This is similar to the density profiles seen in figure 2 of Jalali et al. (2014) in which the densities drop

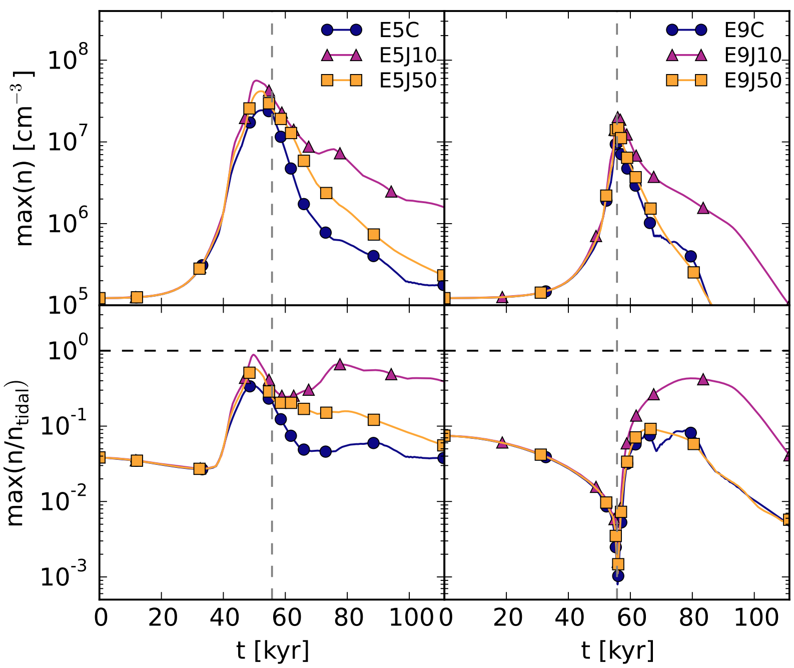

Figure 5.7: Maximum density vs. time (top row) and maximum density divided by the tidal density vs. time (bottom row) for models without radiation at multiple clump temperatures. Results are shown for eccentricities of $\epsilon = 0.5$ (left) and $\epsilon = 0.9$ (right). The horizontal dashed line in the bottom panels at $n/n_{\rm tidal} = 1$ marks the threshold below which tidal effects are expected to be dominant. The vertical dashed line is placed at the time of pericenter passage ($t = 55.7$ kyr).

after pericenter passage (indicating the effect of strong tides), but gravitational collapse still manages to occur as the gas retreats to apoapsis. For such low temperature, the timescale of the pressure-driven expansion is sufficiently low to allow dense gas to recede to larger distances where the tidal limit is less stringent.

5.5.3 Resolution Limits and Convergence

Our models are limited by the computational cost of the radiative transfer routine. Because of this, we are not able to resolve the cooling length ($\lambda_c = c_s \tau_{\rm cool}$, Iwasaki & Tsuribe (2009)) at which dense structure formation is expected to proceed. For the peak densities of 10^8 cm^{-3} seen in our models, the cooling length of ≈ 100 μpc is over an order of magnitude lower than resolution on the highest refinement level. The peak densities seen in our models should therefore be taken as lower limits. For comparison, the peak densities during pericenter passage in Jalali et al. (2014) are on the order of 10^9 cm^{-3}, though we note that their assumption of a lower gas temperature and inclusion of self gravity also increase this value.

To assess the resolution limitations in our models, and to demonstrate the expected lack of convergence, we rerun the control models for each eccentricity with base resolutions of two times (C128) and four times (C256) that of the standard models. In Figure 5.8 we show the results of the peak density and peak density relative to the tidal limit for these models. Because compression is better resolved, peak densities increase with higher resolution. For low eccentricity, the orbital compression is sufficient to drive the density above the tidal limit. Again, due to our lack of self-gravity, the high density gas disperses at late times. For the higher eccentricity models, the density does not breach the tidal limit during pericenter passage, but as the cloud recedes, the density remains sufficiently high for $n/n_{\rm tidal} > 1$. This is similar to the effect noted in § 5.5.2. Our results are far from converged, thus further work is required to diminish the cost of the radiative transfer in order to achieve comparable resolution to previous studies.

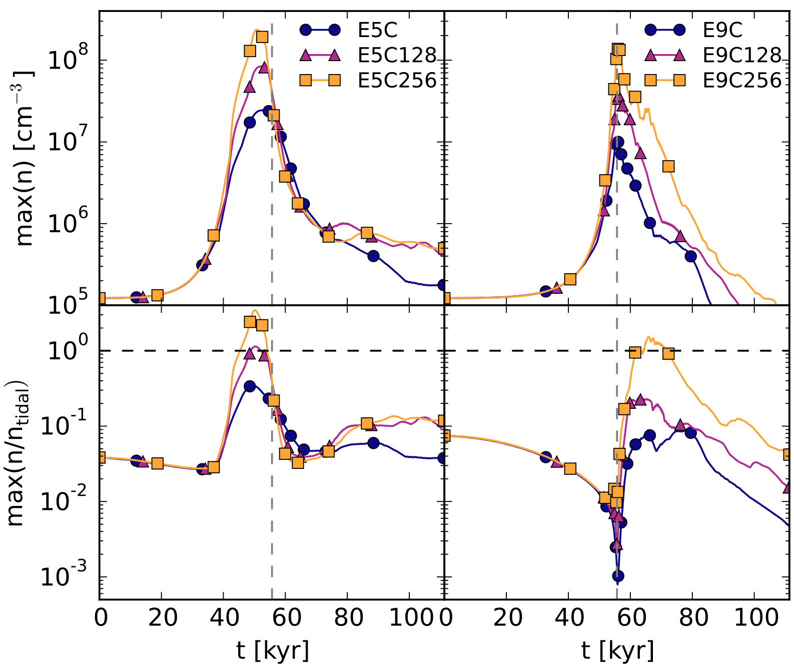

Figure 5.8: Maximum density vs. time (top row) and maximum value of the density divided by the tidal density vs. time (bottom row) for models without radiation at multiple resolutions. Results are shown for eccentricities of $\epsilon = 0.5$ (left) and $\epsilon = 0.9$ (right). The horizontal dashed line in the bottom panels at $n/n_{\rm tidal} = 1$ marks the threshold below which tidal effects are expected to be dominant. The vertical dashed line is placed at the time of pericenter passage ($t = 55.7$ kyr).

Self-Gravity

We do not include self-gravity in our models for two reasons. First, we are interested in exploring the effects of radiation on the clump in-fall scenario, and the inclusion of self-gravity in conjunction with the SMR geometry used would make these models computationally unfeasible. Second, as a further consequence of the radiation routine cost, we are limited to the current resolution. The threshold densities considered in Jalali et al. (2014) ($n \gtrsim 10^{10}$ cm^{-3}) result in Jeans lengths ($\lambda_J = (\pi c_s^2/G\rho)^{1/2} < 1$ pc ; Jeans (1902)) which are an order of magnitude lower than the maximum resolution of our computational mesh. Yet, for high eccentricity orbits, both our models (see Figure 5.8) and the results from Jalali et al. (2014) suggest that peak densities fall well below the tidal limit at pericenter passage. It is only as the gas recedes that the tidal limit decreases and gas of considerably lower density ($n \approx 10^7$ cm^{-3}) becomes gravitationally unstable. This is an effect that should be explored in future work.

5.5.4 Other Considerations

Extended Stellar Potential

In our models, we do not account for the gravitational potential of the nuclear stellar cluster (NSC). Yet, the enclosed mass at 2 pc from stars alone is a few $\times 10^6$ $M_\odot$, comparable to the mass of the SMBH (Genzel et al., 2003; Schödel et al., 2007). This increase in mass, however, does not necessarily indicate that star formation will be further suppressed. On the contrary, the gravitational contribution from this population of stars may aid in ongoing star formation in two ways. First, assuming a spherically symmetric distribution of stars, ambient gas pressures must increase in order to maintain hydrostatic equilibrium in the deeper gravitational potential. The enhanced external pressure resists the outward expansion of gas clumps, resulting in higher gas densities (Yusef-Zadeh & Wardle, 2017). Second, perturbations to the dominant gravitational potential of the SMBH may allow for the tidal force to become compressive in regions where the potential exhibits positive concavity (see case 2 in Figure 1 of Ballesteros-Paredes et al. (2009)).

Nuclear Activity

The larger UV photon emission rate ($Q = 10^{52}$ s^{-1}) likely exceeds that which is produced by the $\approx$ 100 OB/WR stars in the GC. Yet, accretion episodes from Sgr A* may have provided additional UV photons over the past few $\times$ Myr as a result of an extended accretion episode following the formation of the nuclear stellar disk. Conservative lower limits of the accretion rate resulting from massive gas stream collisions with the central SMBH correspond to $\approx$ 1% of the Eddington luminosity (Bonnell & Rice, 2008). For higher mass inflow, the accretion rate can surge to significant fractions of the Eddington limit (Hobbs & Nayakshin (2009), see Chapter 4). Assuming a black hole mass of 4×10^6 $M_\odot$ (Gillessen et al., 2009b) and an approximate spectrum for active galactic nuclei (AGN) emission (Schartmann et al., 2005), 0.1% of the Eddington limit corresponds to a UV (16eV) emission rate of $\approx 10^{52}$ s^{-1}, consistent with the highest central radiation field used in our models. It is therefore likely that past AGN activity in the GC played a role in central star formation arising from a process similar to that detailed in this chapter.

Radiation Pressure

As a final test, we include models for each eccentricity in which the radiation field was set to the maximum value but radiation pressure was excluded. We find that radiation pressure does not significantly alter the dynamics of the gas. Furthermore, radiation pressure does not increase the compression of the clump during infall. This is consistent with the findings of Namekata et al. (2014) that show radiation pressure is sub-dominant when the ionization parameter is low ($U = Q_\star N_\star/(4\pi r^2 c n_{\rm H}) < 10^{-2}$). The distances considered in our models differ from those considered in previous work, though the ionization parameter of $U \approx 10^{-3}$, and the lack of significant influence from radiation pressure, suggest that the claim from Namekata et al. (2014) is also applicable to gas clumps in the central parsec.

5.6 Conclusions

Gas compression of isolated 100 $M_\odot$ gas clumps is thought to provide a non-disk mode of in situ star formation needed to explain both isolated star clusters in the central parsecs of the GC (Eckart et al., 2004) and the presence of YSOs within < 0.1 pc of Sgr A* (Jalali et al., 2014). Evidence of on-going star formation in the GC (Yusef-Zadeh et al. (2017) and references therein) suggests that such a process must contend with the UV radiation from existing massive stars in this region (Gillessen et al., 2009a; Lu et al., 2013; Paumard et al., 2006). Following Jalali et al. (2014), we present 3D RHD simulations of 100 $M_\odot$ gas clumps on eccentric orbits towards SMBH. Our models include an approximate thermal physics prescription using a radiative transfer routine (see Chapter 3) to include a point source with emission rates equivalent to ≈ 100 UV-bright stars. Our results are as follows:

1. The infall of isolated gas clumps on eccentric orbits about a SMBH leads to extreme tidal shearing, though orbital compression occurring through pericenter passage serves as a viable mechanism to rapidly increase the gas density (Figures 5.2 and 5.3).

2. We find that photo-compression of clumps on low eccentricity orbits suppresses orbital compression (Figure 5.4), leading to peak densities prior to pericenter passage (Figure 5.6).

3. For high eccentricity orbits, photo-compression increases the gas density during inflow, though this does not disrupt orbital compression (Figure 5.4). In these models, densities increase through pericenter passage (Figure 5.6), allowing gas to cool to the temperature floor of 100 K (Figure 5.5).

4. For sufficiently high radiation fields, photo-compression results in super-tidal gas densities prior to gas infall (Figure 5.6). This is true for both eccentricities, and may allow for bound structure formation prior to pericenter passage.

5. Adiabatic pressure response for higher clump temperatures rapidly leads to the destruction of gas structures after pericenter passage. Yet, for the low temperatures used in previous work (Jalali et al., 2014), a delay in post-pericenter adiabatic expansion allows for densities to approach the tidal limit at larger radii (Figure 5.7).

6. Peak densities in our models are not yet converged; however, preliminary high resolution simulations suggest that our results are interpreted conservatively (Figure 5.8).

We conclude that star formation resulting from gas clump compression is only strongly affected by the central stellar radiation field for either low eccentricity orbits or for high ionizing photon emission rates ($\gtrsim 10^{52}$ s^{-1}). For high eccentricity orbits, photo-compression provides a fractional increase to orbital compression and may aid in the formation of super-tidal gas structure.

CHAPTER 6: CONCLUSION

The origin of young, massive stars in the GC is best explained by two processes. (i) The formation and fragmentation of a nuclear gas disk via the inflow of a $\approx 10^5\ M_\odot$ gas stream towards the central SMBH remains the most viable origin for the central stellar disk. (ii) The infall of isolated 100 $M_\odot$ gas clumps is currently the only alternative star formation mechanism that can help to explain evidence of recent star formation in the central parsecs of the GC.

In response to evidence that massive gas inflow is likely to incite AGN activity due to accretion onto the SMBH, I have investigated the role of UV and X-ray radiation arising from such a process through the first three-dimensional radiative hydrodynamic simulations of central gas disk formation preceding a nuclear star formation episode. I show that low radiation fields ($L \lesssim 0.1 L_{\rm edd}$) are insufficient to suppress the formation of a nuclear gas disk. Strong shielding allows the disk to sufficiently cool to near-equilibrium temperatures, thus star formation is likely in this radiation limit. For strong radiation fields ($L \approx L_{\rm edd}$), radiation pressure from UV photons inhibits the inflow of gas towards the SMBH. Yet, gas accumulates along lines of sight which are sufficiently shielded from the radiation field. A disk forms gradually, though with considerably lower density than is seen for weaker radiation fields, thus it is unlikely that star formation would occur in this case. Because the accretion of mass onto the SMBH occurs on the viscous timescale ($\approx 10^6 - 10^7$), such a radiation field would be exceptionally strong. As a result, I conclude that nuclear star formation resulting from gas inflow is not strongly affected by nuclear activity.

As an alternative to disk-based star formation, I build upon models exploring star formation sparked by the gravitational compression of infalling gas clumps (Jalali et al., 2014) by considering the effect of UV radiation from existing stars in the GC. I find that photo-compression of the gas partially inhibits the tidal elongation of the clump during its approach towards the SMBH. For sufficiently strong radiation fields, the peak densities can be super-tidal, indicating that gravitational collapse may be possible prior to pericenter passage where maximum gravitational compression occurs. For low eccentricity orbits, the impinging radiation field results in peak densities prior to pericenter passage. For high eccentricity orbits, gravitational compression during pericenter passage remains dominant, though peak densities increase in proportion to the radiation field. I conclude that UV radiation from a nuclear stellar cluster likely aids in on-going star formation in the central parsec through the compression mechanism described above. Future work, including models with self-gravity of the gas, should investigate this process.

To complete these two studies, I have implemented an adaptive radiative transfer routine based on the HEALPix geometry into the ATHENA code, specifically for applications making use of the SMR framework of the code. Great care was taken in balancing computational efficiency and accuracy of this routine. The radiation module shows excellent agreement with benchmark tests of its accuracy, with exquisite consistency across SMR domains. The code also includes two novel techniques, off-mesh ray tracing and selective integration, which extend the applicability of this routine to future studies. However, an improvement to the parallelization scheme for moving radiation sources is necessary to reduce computational cost for higher resolution simulations.

The research presented here considers the influence of ionizing radiation on gas dynamics and star formation near supermassive black holes. Yet, this is only one physical component to the extremely complex environment of the GC. It is clear though, that as computational resources and capabilities increase in the years to come, ionizing radiation will be a vital ingredient to numerical models aiming to improve our understanding of the intricate processes governing star formation in this hostile environment.

BIBLIOGRAPHY

Abel, T., Norman, M. L., & Madau, P. 1999, ApJ, 523, 6671

Abel, T., & Wandelt, B. D. 2002, MNRAS, 330, L53

Alig, C., Burkert, A., Johansson, P. H., & Schartmann, M. 2011, MNRAS, 412, 469

Alig, C., Schartmann, M., Burkert, A., & Dolag, K. 2013, ApJ, 771, 119

Baganoff, F. K., Maeda, Y., Morris, M., et al. 2003, ApJ, 591, 891

Ballesteros-Paredes, J., Gómez, G. C., Pichardo, B., & Vázquez-Semadeni, E. 2009, MNRAS, 393, 1563

Bartko, H., Martins, F., Fritz, T. K., et al. 2009, ApJ, 697, 1741

Bartko, H., Martins, F., Trippe, S., et al. 2010, ApJ, 708, 834

Böker, T., Sarzi, M., McLaughlin, D. E., et al. 2004, AJ, 127, 105

Bondi, H., & Hoyle, F. 1944, MNRAS, 104, 273

Bonnell, I. A., & Rice, W. K. M. 2008, Science, 321, 1060

Bryan, G. L., Norman, M. L., O'Shea, B. W., et al. 2014, ApJS, 211, 19

Burkert, A., Schartmann, M., Alig, C., et al. 2012, ApJ, 750, 58

Chatzopoulos, S., Fritz, T. K., Gerhard, O., et al. 2015, MNRAS, 447, 948

Christopher, M. H., Scoville, N. Z., Stolovy, S. R., & Yun, M. S. 2005, ApJ, 622, 346

Clark, P. C., Glover, S. C. O., Ragan, S. E., Shetty, R., & Klessen, R. S. 2013, ApJL, 768, L34

Coil, A. L., & Ho, P. T. P. 1999, ApJ, 513, 752

—. 2000, ApJ, 533, 245

Colella, P., & Woodward, P. R. 1984, Journal of Computational Physics, 54, 174

Côté, P., Piatek, S., Ferrarese, L., et al. 2006, ApJS, 165, 57

Courant, R., Friedrichs, K., & Lewy, H. 1928, Mathematische Annalen, 100, 32

Draine, B. T. 2011, Physics of the Interstellar and Intergalactic Medium (Princeton University Press)

Eckart, A., Moultaka, J., Viehmann, T., Straubmeier, C., & Mouawad, N. 2004, ApJ, 602, 760

Eisenhauer, F., Genzel, R., Alexander, T., et al. 2005, ApJ, 628, 246

Ferrière, K. 2012, A&A, 540, A50

Gammie, C. F. 2001, ApJ, 553, 174

Genzel, R., Eisenhauer, F., & Gillessen, S. 2010, Reviews of Modern Physics, 82, 3121

Genzel, R., Schödel, R., Ott, T., et al. 2003, ApJ, 594, 812

Ghez, A. M., Salim, S., Hornstein, S. D., et al. 2005, ApJ, 620, 744

Ghez, A. M., Duchêne, G., Matthews, K., et al. 2003, ApJL, 586, L127

Ghez, A. M., Salim, S., Weinberg, N. N., et al. 2008, ApJ, 689, 1044

Gillessen, S., Eisenhauer, F., Fritz, T. K., et al. 2009a, ApJL, 707, L114

Gillessen, S., Eisenhauer, F., Trippe, S., et al. 2009b, ApJ, 692, 1075

Goldsmith, P. F., & Langer, W. D. 1978, ApJ, 222, 881

Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759

Guesten, R., & Downes, D. 1983, A&A, 117, 343

Guo, F., & Mathews, W. G. 2012, ApJ, 756, 181

Gürkan, M. A., & Rasio, F. A. 2005, ApJ, 628, 236

Hansen, B. M. S., & Milosavljević, M. 2003, ApJL, 593, L77

Herrnstein, R. M., & Ho, P. T. P. 2005, ApJ, 620, 287

Ho, P. T. P., Ho, L. C., Szczepanski, J. C., Jackson, J. M., & Armstrong, J. T. 1991, Nature, 350, 309

Hobbs, A., & Nayakshin, S. 2009, MNRAS, 394, 191

Hocuk, S., & Spaans, M. 2010, A&A, 522, A24

—. 2011, A&A, 536, A41

Iliev, I. T., Whalen, D., Mellema, G., et al. 2009, MNRAS, 400, 1283

Iwasaki, K., & Tsuribe, T. 2009, A&A, 508, 725

Jalali, B., Pelupessy, F. I., Eckart, A., et al. 2014, MNRAS, 444, 1205

Jeans, J. H. 1902, Philosophical Transactions of the Royal Society of London Series A, 199, 1

Kauffmann, J., Pillai, T., Zhang, Q., et al. 2017, A&A, 603, A89

Kim, J.-h., Wise, J. H., Alvarez, M. A., & Abel, T. 2011, ApJ, 738, 54

Koyama, H., & Inutsuka, S.-i. 2002, ApJL, 564, L97

Krabbe, A., Genzel, R., Eckart, A., et al. 1995, ApJL, 447, L95

Kruijssen, J. M. D., Dale, J. E., & Longmore, S. N. 2015, MNRAS, 447, 1059

Kruijssen, J. M. D., Longmore, S. N., Elmegreen, B. G., et al. 2014, MNRAS, 440, 3370

Krumholz, M. R., & Kruijssen, J. M. D. 2015, MNRAS, 453, 739

Krumholz, M. R., Stone, J. M., & Gardiner, T. A. 2007, ApJ, 671, 518

Lau, R. M., Herter, T. L., Morris, M. R., Becklin, E. E., & Adams, J. D. 2013, ApJ, 775, 37

Lee, S., Pak, S., Choi, M., et al. 2008, ApJ, 674, 247

Levin, Y., & Beloborodov, A. M. 2003, ApJL, 590, L33

Lo, K. Y., & Claussen, M. J. 1983, Nature, 306, 647

Longmore, S. N., Bally, J., Testi, L., et al. 2013, MNRAS, 429, 987

Lu, J. R., Do, T., Ghez, A. M., et al. 2013, ApJ, 764, 155

Lu, J. R., Ghez, A. M., Hornstein, S. D., Morris, M., & Becklin, E. E. 2005, ApJL, 625, L51

Lu, J. R., Ghez, A. M., Hornstein, S. D., et al. 2009, ApJ, 690, 1463

Lucas, W. E., Bonnell, I. A., Davies, M. B., & Rice, W. K. M. 2013, MNRAS, 433, 353

Mapelli, M., & Gualandris, A. 2016, in Lecture Notes in Physics, Berlin Springer Verlag, Vol. 905, Lecture Notes in Physics, Berlin Springer Verlag, ed. F. Haardt, V. Gorini, U. Moschella, A. Treves, & M. Colpi, 205–272

Mapelli, M., Hayfield, T., Mayer, L., & Wadsley, J. 2012, ApJ, 749, 168

Mapelli, M., & Trani, A. A. 2016, A&A, 585, A161

McGary, R. S., Coil, A. L., & Ho, P. T. P. 2001, ApJ, 559, 326

Mezger, P. G., Zylka, R., Salter, C. J., et al. 1989, A&A, 209, 337

Mills, E. A. C., Güsten, R., Requena-Torres, M. A., & Morris, M. R. 2013, ApJ, 779, 47

Molinari, S., Bally, J., Noriega-Crespo, A., et al. 2011, ApJL, 735, L33

Montero-Castaño, M., Herrnstein, R. M., & Ho, P. T. P. 2009, ApJ, 695, 1477

Morris, M., & Serabyn, E. 1996, ARA&A, 34, 645

Moser, L., Sánchez-Monge, Á., Eckart, A., et al. 2017, A&A, 603, A68

Mužić, K., Schödel, R., Eckart, A., Meyer, L., & Zensus, A. 2008, A&A, 482, 173

Namekata, D., Umemura, M., & Hasegawa, K. 2014, MNRAS, 443, 2018

Nayakshin, S., Cuadra, J., & Springel, V. 2007, MNRAS, 379, 21

Oka, T., Nagai, M., Kamegai, K., & Tanaka, K. 2011, ApJ, 732, 120

Osterbrock, D. E. 1989, Astrophysics of gaseous nebulae and active galactic nuclei (University Science Books)

Papadopoulos, P. P., Thi, W.-F., Miniati, F., & Viti, S. 2011, MNRAS, 414, 1705

Paumard, T., Genzel, R., Martins, F., et al. 2006, ApJ, 643, 1011

Reid, M. J. 1993, ARA&A, 31, 345

Requena-Torres, M. A., Güsten, R., Weiß, A., et al. 2012, A&A, 542, L21

Rosen, A. L., Krumholz, M. R., Oishi, J. S., Lee, A. T., & Klein, R. I. 2017, Journal of Computational Physics, 330, 924

Rybicki, G. B., & Lightman, A. P. 1979, Radiative processes in astrophysics (Wiley-VCH)

Sanders, R. H. 1998, MNRAS, 294, 35

Schartmann, M., Krause, M., & Burkert, A. 2011, MNRAS, 415, 741

Schartmann, M., Meisenheimer, K., Camenzind, M., Wolf, S., & Henning, T. 2005, A&A, 437, 861

Schödel, R., Eckart, A., Iserlohe, C., Genzel, R., & Ott, T. 2005, ApJL, 625, L111

Schödel, R., Ott, T., Genzel, R., et al. 2003, ApJ, 596, 1015

—. 2002, Nature, 419, 694

Schödel, R., Eckart, A., Alexander, T., et al. 2007, A&A, 469, 125

Shakura, N. I., & Sunyaev, R. A. 1973, A&A, 24, 337

Shin, M.-S., Stone, J. M., & Snyder, G. F. 2008, ApJ, 680, 336

Shull, J. M., & van Steenberg, M. E. 1985, ApJ, 298, 268

Sternberg, A., Hoffmann, T. L., & Pauldrach, A. W. A. 2003, ApJ, 599, 1333

Stolte, A., Ghez, A. M., Morris, M., et al. 2008, ApJ, 675, 1278

Stone, J. M., & Gardiner, T. 2009, New Astronomy, 14, 139

Stone, J. M., Gardiner, T. A., Teuben, P., Hawley, J. F., & Simon, J. B. 2008, ApJS, 178, 137

Strömgren, B. 1939, ApJ, 89, 526

Su, M., Slatyer, T. R., & Finkbeiner, D. P. 2010, ApJ, 724, 1044

Sunyaev, R. A., Markevitch, M., & Pavlinsky, M. 1993, ApJ, 407, 606

Tenorio-Tagle, G., Bodenheimer, P., Lin, D. N. C., & Noriega-Crespo, A. 1986, MNRAS, 221, 635

Toomre, A. 1964, ApJ, 139, 1217

Toro, E. 2009, Riemann Solvers and Numerical Methods for Fluid Dynamics: A Practical Introduction (Springer Berlin Heidelberg)

Trippe, S., Gillessen, S., Gerhard, O. E., et al. 2008, A&A, 492, 419

Truelove, J. K., Klein, R. I., McKee, C. F., et al. 1997, ApJL, 489, L179

Tsuboi, M., & Miyazaki, A. 2012, PASJ, 64, 111

Tsuboi, M., Miyazaki, A., & Uehara, K. 2015, PASJ, 67, 109

Walcher, C. J., van der Marel, R. P., McLaughlin, D., et al. 2005, ApJ, 618, 237

Wardle, M., & Yusef-Zadeh, F. 2008, ApJL, 683, L37

Williams, A., Barrus, S., Morley, R. K., & Shirley, P. 2005, in ACM SIGGRAPH 2005 Courses, SIGGRAPH '05 (New York, NY, USA: ACM)

Wise, J. H., & Abel, T. 2011, MNRAS, 414, 3458

Wolfire, M. G., Hollenbach, D., McKee, C. F., Tielens, A. G. G. M., & Bakes, E. L. O. 1995, ApJ, 443, 152

Yelda, S., Ghez, A. M., Lu, J. R., et al. 2014, ApJ, 783, 131

Yusef-Zadeh, F., Braatz, J., Wardle, M., & Roberts, D. 2008, ApJL, 683, L147

Yusef-Zadeh, F., Cotton, B., Wardle, M., et al. 2017, MNRAS, 467, 922

Yusef-Zadeh, F., Roberts, D. A., Wardle, M., et al. 2015, ApJL, 801, L26

Yusef-Zadeh, F., & Wardle, M. 2017, in IAU Symposium, Vol. 324, New Frontiers in Black Hole Astrophysics, ed. A. Gomboc, 111–114

Yusef-Zadeh, F., Wardle, M., Schödel, R., et al. 2016, ApJ, 819, 60

Yusef-Zadeh, F., Royster, M., Wardle, M., et al. 2013, ApJL, 767, L32

Zhao, J.-H., Blundell, R., Moran, J. M., et al. 2010, ApJ, 723, 1097

Zubovas, K., & Nayakshin, S. 2012, MNRAS, 424, 666

Zylka, R., Mezger, P. G., & Wink, J. E. 1990, A&A, 234, 133