

UC Santa Barbara

UC Santa Barbara Electronic Theses and Dissertations

Title

Cool Things That Matter: Multiphase Dynamics of Galactic Atmospheres

Permalink

<https://escholarship.org/uc/item/4qw919v9>

Author

Tan, Zun Yi Brent

Publication Date

2023

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

Cool Things That Matter: Multiphase Dynamics of Galactic Atmospheres

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Physics

by

Brent Tan Zun Yi

Committee in charge:

Professor Siang Peng Oh, Chair
Professor Crystal L. Martin
Professor Joseph F. Hennawi

September 2023

The dissertation of Brent Tan Zun Yi is approved.

Professor Crystal L. Martin

Professor Joseph F. Hennawi

Professor Siang Peng Oh, Committee Chair

May 2023

Copyright © 2023

by Brent Tan Zun Yi

To my parents
Who in the length of eternity
And the breadth of the cosmos
Gifted me the depths of
Life, love, and hope

Acknowledgements

*O snail
Climb Mount Fuji
But slowly, slowly*

Kobayashi Issa

By and large, the course of a PhD is a winding uphill task, the path sometimes arduous and the route indirect. Some problems lead down a rabbit hole and become a test of tenacity, and some wrong turns and dead ends become reminiscent of Sisyphus' perpetual toil. Progress in uncharted territory is usually incremental, and slow but persistent perseverance is the name of the game. At the end of the day, it is the people we meet along the way who carry us to the finish line. It has been my great fortune in this endeavour to have been guided, and often inspired, by a number of people. Without them, none of this would be possible.

First and foremost, I would like to thank my advisor and mentor Siang Peng Oh, to whom I owe a deep debt of gratitude. Not just for the scientific development and endless source of critical insights he provided, but also for really shaping new perspectives and ways of tackling difficult and complex problems. He has been a tremendous source of inspiration, encouragement, and motivation, and a role model to aspire towards.

I would also like to thank the other members of my committee, Crystal Martin and Joseph Hennawi, for all the nuggets of wisdom they have shared and their support and guidance over the years.

I thank my mentors and collaborators, Max Gronke and Drummond Fielding, for taking the time, patience and effort to impart lessons that have made me a better scientist, and for enabling me to work on interesting problems.

I thank Shirley Ho, Rachel Mandelbaum, and Tom Quinn for giving me early opportunities to explore research as an undergrad.

I thank all my peers and friends for being a home away from home.

I thank all the people I have met along the way but have not explicitly acknowledged, whether our interactions were big or small, for the difference they have made.

And last but certainly not least, I thank my family—my parents and my sister—for their unconditional and unwavering love and support throughout this journey.

Finally, I would like to acknowledge the immense privilege to have had the opportunity to be a part of human exploration in a field where the universe is your oyster. All science is done standing on the shoulders of giants, upon all that have come before, and with the timeless endeavouring spirit that is human curiosity.

Curriculum Vitæ

Brent Tan Zun Yi

Education

2023 Ph.D. in Physics, University of California, Santa Barbara.
2020 M.A. in Physics, University of California, Santa Barbara.
2017 B.Sc. in Physics, Carnegie Mellon University.

Publications

Tan, B., & Fielding, D. B. (2023). Cloud Atlas: Navigating the Multiphase Landscape of Tempestuous Galactic Winds. *Monthly Notices of the Royal Astronomical Society*, submitted.

Tan, B., Oh, S. P., & Gronke, M. (2023). Cloudy with a chance of rain: accretion braking of cold clouds. *Monthly Notices of the Royal Astronomical Society*, 520(2), 2571-2592.

Tan, B., & Oh, S. P. (2021). A model for line absorption and emission from turbulent mixing layers. *Monthly Notices of the Royal Astronomical Society: Letters*, 508(1), L37-L42.

Tan, B., Oh, S. P., & Gronke, M. (2021). Radiative mixing layers: insights from turbulent combustion. *Monthly Notices of the Royal Astronomical Society*, 502(3), 3179-3199.

Lin, C. H., **Tan, B.**, Mandelbaum, R., & Hirata, C. M. (2020). The impact of light polarization effects on weak lensing systematics. *Monthly Notices of the Royal Astronomical Society*, 496(1), 532-539.

Abstract

Cool Things That Matter: Multiphase Dynamics of Galactic Atmospheres

by

Brent Tan Zun Yi

Galaxies are extremely complex systems. A multitude of open problems still surround galaxy formation and evolution today. One that sits at the very heart is the challenge to understand the observed multiphase nature of not just the galaxies and their surrounding environment (circumgalactic medium (CGM)), but also that of the Cosmic Baryon Cycle. Galactic outflows driven by feedback mechanisms carry material outwards while inflowing gas provides fuel for new star formation. This cycling connects processes on stellar ($\sim \text{pc}$) scales to galactic ($\sim \text{kpc}$) and cosmological ($\sim \text{Mpc}$) scales. The interactions between phases in multiphase systems leads to coupling across this large range of scales. In short, understanding the small scales is essential for being able to accurately model larger scales. A large focus of the research presented in this dissertation hence lies in studying the physics that govern the dynamic nature of these multiphase systems and their interactions. Despite being ubiquitous, many uncertainties remain due to their surprisingly rich complexity. Combining analytic theory with numerical simulations, we delve into their inner workings so as to be able to understand and model them. We start at the smallest scales in the problem with a deep dive into the interfaces between phases and how they determine the bulk evolution. We then explore the connection between these mixing layers and observables. Applying these results to larger scales, we look at cold clouds moving through hot backgrounds, both infalling under gravity and in turbulent outflowing winds.

Contents

1	Introduction	1
2	Radiative Mixing Layers: Insights from Turbulent Combustion	4
2.1	Introduction	5
2.2	Methods	10
2.2.1	Radiative Cooling	12
2.2.2	Thermal Conduction	12
2.2.3	Parameter studies	14
2.3	1D Simulations: Laminar Fronts	15
2.3.1	Setup	15
2.3.2	Results	19
2.4	Analytic Estimates from Turbulent Combustion	24
2.4.1	Introduction & Terminology	24
2.4.2	Characteristic Regimes	27
2.4.3	Scalings of the turbulent velocity S_T	29
2.4.4	Details of the Fiducial $S_T/u' = Da^{1/4}$ Scaling	31
2.4.5	Implications for the fractal nature of mixing layers	34
2.4.6	Implications for energetics and convergence criteria	35
2.5	3D Simulations: Turbulent Fronts	36
2.5.1	Setup	37
2.5.2	Morphology of Mixing Layers: Transition from Single Phase to Multiphase	38
2.5.3	Scaling Relations	43
2.5.4	Comparing Simulations to 1D Mixing Length Models	55
2.5.5	Thermal Conduction	61
2.5.6	Convergence	63
2.5.7	Comparison with Previous Work	65
2.5.8	Conclusions	67

3	A Model for Line Absorption and Emission from Turbulent Mixing Layers	75
3.1	Introduction	76
3.2	Methods	77
3.3	1D Mixing Layer Models	78
3.3.1	1D Mixing Layers	79
3.3.2	Temperature Distribution	80
3.3.3	Length Scales	83
3.3.4	Non-Constant Conductivity	84
3.3.5	Column Densities and Line Ratios	86
3.4	Results	87
3.4.1	Temperature PDFs	87
3.4.2	Resolution	89
3.4.3	Column Densities and Line Ratios	90
3.5	Discussion	93
4	Cloudy with A Chance of Rain: Accretion Braking of Cold Clouds	96
4.1	Introduction	97
4.2	Dynamics of Infalling Clouds	103
4.2.1	Cloud Evolution and Terminal Velocities	103
4.2.2	Cloud Growth	109
4.2.3	Cloud Survival	115
4.3	Methods	121
4.4	Results : Constant Background	124
4.4.1	Time Evolution	124
4.4.2	Area Growth Rate	128
4.4.3	Scalings	130
4.4.4	Survival	136
4.4.5	Growth and Free-fall Timescales	138
4.5	Results : Stratified Background	141
4.5.1	Time Evolution	141
4.5.2	Terminal Velocity	143
4.5.3	Scaling With Cloud Size and Gravity	147
4.5.4	Resolution Convergence	149
4.5.5	Survival in a stratified background	149
4.6	Discussion	153
4.6.1	High Velocity Clouds	153
4.6.2	Clusters	157
4.6.3	Other Implications	162
4.6.4	Further Considerations	165
4.7	Conclusions	168

5	Cloud Atlas: Navigating the Multiphase Landscape of Tempestuous Galactic Winds	172
5.1	Introduction	173
5.2	Seeding Clouds In Outflows	178
5.2.1	Cloud Size Distribution	179
5.2.2	Theoretical Model for Cloud Growth	182
5.3	Methods	185
5.3.1	Setup	186
5.3.2	Source Terms	186
5.3.3	Initial Conditions	194
5.4	Results: Breakout and Wind Outflow	194
5.4.1	Superbubble Breakout	194
5.4.2	Asymmetric Outflows	199
5.4.3	Turbulent Winds	200
5.4.4	Mass and Energy Flux Phase Distribution and Cloud Entrainment	210
5.5	Results : Clouds	212
5.5.1	Cloud Identification	212
5.5.2	Measuring Cloud Properties	213
5.5.3	Cloud Model	217
5.5.4	Non-thermal Pressure Support	231
5.5.5	Cloud Origins: Passive Scalars	233
5.5.6	Cloud-Wind Alignment	235
5.6	Discussion	236
5.6.1	Connection to Small Scale Simulations	236
5.6.2	Implications for Galaxy/Cosmological Scale Simulations	238
5.6.3	Implications for Observations	240
5.6.4	Further Considerations	241
5.6.5	Looking Forward	243
5.7	Conclusions	244
6	Summary and Conclusions	248
A	Code Verification Test: 1D Diffusion Couple	254
B	Hydrostatic Equilibrium Test	257
C	Constrained Turbulence Test	259
D	Extended Townsend Algorithm	261
E	Estimating Cloud Growth Rates	269
	Bibliography	276

Chapter 1

Introduction

*There are more things in Heaven and Earth, Horatio,
Than are dreamt of in your philosophy.*

William Shakespeare, Hamlet

Galaxies are complex ecosystems—making sense of observations thus means making sense of the inherent nonlinear, dynamical structures that underlie them. A multitude of open problems still surround galaxy formation and evolution today. One that sits at the very heart is the challenge to understand the multiphase nature of not just the galaxies and their surrounding environment (circumgalactic medium (CGM)), but also that of the Cosmic Baryon Cycle. Galactic outflows driven by feedback mechanisms carry material outwards while inflowing gas provides fuel for new star formation. This cycling connects processes on stellar ($\sim \text{pc}$) scales to galactic ($\sim \text{kpc}$) and cosmological ($\sim \text{Mpc}$) scales, weaving them into a tightly interdependent network and creating a multiphysics, multiphase, multiscale ecosystem.

In recent years, a wealth of observations of the CGM ranging from quasar line spectroscopy to emission line mapping with integral field units have provided challenging tests

for galaxy formation theory (Tumlinson et al. 2017). One particularly surprising finding from such observations is the prevalent and ubiquitous presence of cold ($\sim 10^4$ K) gas in the CGM. The balance between gravity and thermal pressure suggests that gas in the CGM should exist at much higher temperatures ($> 10^6$ K). Instead, observations include small and dense cold ($\sim 10^4$ K) clouds (Hennawi et al. 2015) and intermediate temperature ($\sim 10^5$ K) gas which should be thermally unstable and cool rapidly (Tumlinson et al. 2011). Observations also suggest that cold clouds have significant sources of non-thermal pressure support (Werk et al. 2014) such as turbulence, magnetic fields and cosmic rays. How these all come together to impact cold clouds is still poorly understood. Furthermore, such clouds should be rapidly destroyed by hydrodynamic instabilities (Klein et al. 1994), so how do they survive? Recent work has shown that for outflowing winds, capturing the appropriate scales and boundary conditions is *critical* in answering this question. However, outflows are but one step in the cycling of baryons, and similar studies for the CGM and infalling clouds, and in general more realistic environments, are still lacking. Lastly, multiphase systems are notoriously difficult to study because they introduce new, much smaller scales that can significantly impact large scale behavior. For instance, the interactions at the interfaces between phases can alter bulk flow and phase properties. If we do not understand the small scales, we will not get the large scales right.

Where do we stand? Theories for the origin of cold gas in the CGM include formation via thermal instability in the hot halo gas (the precipitation model; Sharma et al. 2012; Voit et al. 2019) or transport outward via outflowing winds driven by stellar or AGN feedback. Empirical support for this idea come from the detection of outflowing galactic winds which contain both cold and hot gas (Veilleux et al. 2005; Steidel et al. 2010; Rubin et al. 2014; Heckman & Thompson 2017). These observations were puzzling, since any cold gas that is swept up should be destroyed in these winds. The timescale for the acceleration of a cold cloud is a factor of ~ 10 times longer than the timescale

for its destruction via Kelvin-Helmholtz and Rayleigh-Taylor instabilities (Zhang et al. 2017). This was verified in simulations (Cooper et al. 2009; Scannapieco & Brüggen 2015; Schneider & Robertson 2017). However, it was recently shown that under the right conditions, these cold clouds can not only survive but even grow (Gronke & Oh 2018). This is possible when the cooling time of mixed gas is shorter than the destruction timescale. The mechanism that drives this growth is the formation of a long tail structure of cold gas where mixing and cooling takes place efficiently in radiative turbulent mixing layers at the interface of the two gas phases (Ji et al. 2018; Fielding et al. 2020; Tan et al. 2021). Crucially, this requires long simulation domains that were not captured in earlier simulations. This has led to renewed interest in such simulations. As mentioned, winds are only one piece of the puzzle. For example, in precipitation models, the ratio of the cooling time to free-fall time (the key parameter in the model) is controlled via feedback processes. A crucial part of this process is the transport of cold gas to the disk itself. How this happens remains unclear.

A large focus of the research presented in this dissertation lies in studying the physics that govern the dynamic nature of these multiphase systems. Despite being ubiquitous, many uncertainties remain due to their surprisingly rich complexity. Combining analytic theory with numerical simulations, we delve into their inner workings so as to be able to understand and model them. In Chapter 2, we start at the smallest scales in the problem with a deep dive into the interfaces between phases and how they determine bulk evolution. We then explore the connection between these mixing layers and observables in Chapter 3. Applying these results to larger scales, we look at cold clouds moving through hot backgrounds, both infalling under gravity (Chapter 4) and in turbulent outflowing winds (Chapter 5). Lastly, we summarize and conclude in Chapter 6.

Chapter 2

Radiative Mixing Layers: Insights from Turbulent Combustion

*Some say the world will end in fire,
Some say in ice.
From what I've tasted of desire
I hold with those who favor fire.
But if it had to perish twice,
I think I know enough of hate
To say that for destruction ice
Is also great
And would suffice.*

Robert Frost, Fire and Ice

Radiative mixing layers arise wherever multiphase gas, shear, and radiative cooling are present. Simulations show that in steady state, thermal advection from the hot phase balances radiative cooling. However, many features are puzzling. For instance, hot gas entrainment appears to be numerically converged despite the scale-free, fractal structure of such fronts being unresolved. Additionally, the hot gas heat flux has a characteristic velocity $v_{\text{in}} \approx c_{\text{s,cold}}(t_{\text{cool}}/t_{\text{sc,cold}})^{-1/4}$ whose strength and scaling are not intuitive. We revisit these issues in 1D and 3D hydrodynamic simulations. We find that over-cooling only happens if numerical diffusion dominates thermal transport; convergence is still possible even when the Field length is unresolved. A deeper physical understanding of radiative fronts can be obtained by exploiting parallels between mixing layers and turbulent combustion, which has well-developed theory and abundant experimental data. A key parameter is the Damköhler number $\text{Da} = \tau_{\text{turb}}/t_{\text{cool}}$, the ratio of the outer eddy turnover time to the cooling time. Once $\text{Da} > 1$, the front fragments into a multiphase medium. Just as for scalar mixing, the eddy turnover time sets the mixing rate, independent of small scale diffusion. For this reason, thermal conduction often has limited impact. We show that v_{in} and the effective emissivity can be understood in detail by adapting combustion theory scalings. Mean density and temperature profiles can also be reproduced remarkably well by mixing length theory. These results have implications for the structure and survival of cold gas in many settings, and resolution requirements for large scale galaxy simulations.

2.1 Introduction

Multiphase media are ubiquitous in astrophysics. Interfaces between different phases are not infinitely sharp, but thickened by energy transport processes such as thermal conduction (Borkowski et al. 1990; Gnat et al. 2010) and collisionless cosmic ray (CR)

scattering (Wiener et al. 2017). Shear flows further structure the interface, by driving the Kelvin-Helmholtz (KH) instability, which seeds turbulence and fluid mixing. In ideal hydrodynamics, the KH instability is scale free. However, non-ideal processes, such as viscosity, can set a characteristic scale. Perhaps the most important of these non-ideal processes is radiative cooling, which typically is very strong in mixed gas at temperatures intermediate between the two stable phases. For instance, under coronal conditions the cooling curve peaks at $T \sim 10^5$ K, intermediate between the $T \sim 10^4$ K and $T \sim 10^6$ K phases. Radiative turbulent mixing layers (TMLs) then arise where the exchange of mass, momentum and energy between phases is governed by the interaction between turbulence and radiative cooling. This has many important physical and observational consequences. For example, in the circumgalactic medium (CGM), such physics governs the growth or destruction of cold clouds embedded in a hot wind (Klein et al. 1994; Mellema et al. 2002; Pittard et al. 2005; Cooper et al. 2009; Scannapieco & Brüggén 2015; Schneider & Robertson 2017; Gronke & Oh 2018, 2020a), and the survival of cold streams inflowing from cosmological accretion (Cornuault et al. 2018; Mandelker et al. 2020). TMLs also ‘host’ $T \sim 10^5$ K gas which could explain the abundance of OVI seen in galaxy halos, despite the fact that it is thermally unstable (Slavin et al. 1993). In addition, TMLs play crucial roles in the ISM (e.g., in supernova explosions), galaxy clusters (e.g., in the interface between optical filaments and the intracluster medium), AGN environments (e.g., chaotic cold accretion on the AGN, Gaspari et al. 2013); survival and stability of AGN jets (Hardee & Stone 1997), and many other astrophysical settings.

Despite their ubiquity and importance, radiative mixing layers have received relatively little attention compared to adiabatic simulations of the Kelvin Helmholtz instability. Begelman & Fabian (1990) wrote an early analytic paper suggesting that radiative mixing layers are characterized by a mean temperature $\bar{T} \sim (T_{\text{hot}}T_{\text{cold}})^{1/2}$ and width $l \sim v_t t_{\text{cool}}$, where v_t is the turbulent velocity and t_{cool} is the cooling time of the mixed gas. In a

series of papers (Kwak & Shelton 2010; Kwak et al. 2011; Henley et al. 2012; Kwak et al. 2015), Kwak and collaborators ran 2D hydrodynamic simulations and compared to observed line column densities and ratios, but not to analytic theory. Esquivel et al. (2006) ran 3D MHD simulations, but not for long enough for effective mixing (or a stable equilibrium) to develop.

More recently, Ji et al. (2019) performed 3D hydrodynamic and MHD simulations, including both photoionization and non-equilibrium ionization. Interestingly, they found strong discrepancies with analytic models even in the purely hydrodynamic regime, with characteristic inflow and turbulent velocities much less than the shear velocity, and of order the cold gas sound speed, $c_{s,\text{cold}}$. They also found the layer width $l \propto t_{\text{cool}}^{1/2}$ rather than $l \propto t_{\text{cool}}$, implying a surface brightness and mass entrainment velocity $Q, v_{\text{in}} \propto t_{\text{cool}}^{-1/2}$. They also found that previous analytic scalings (e.g., that the column density is independent of density or metallicity) do not agree with simulations. Subsequently, Gronke & Oh (2018, 2020a) looked at mass growth of cold clouds embedded in a wind and found similar inflow velocities $v_{\text{in}} \sim c_{s,\text{cold}}$, but with a different scaling $Q, v_{\text{in}} \propto t_{\text{cool}}^{-1/4}$, which has also been seen by Mandelker et al. (2020); Fielding et al. (2020). Fielding et al. (2020) ran a suite of 3D hydrodynamic simulations similar to Ji et al. (2019), and highlighted the fractal nature of the interface; they derived a formula for v_{in} based on this observation.

The situation is far from resolved. In our opinion, some of the biggest outstanding questions are:

- *Scalings.* In previous work (Gronke & Oh 2020a), we found:

$$v_{\text{in}} \approx 0.2c_{s,\text{cold}} \left(\frac{t_{\text{cool}}}{t_{\text{sc,cold}}} \right)^{-1/4} = 0.2c_{s,\text{cold}} \left(\frac{c_{s,\text{cold}}t_{\text{cool}}}{L} \right)^{-1/4}, \quad (2.1)$$

where L is a characteristic length scale, and $c_{s,\text{cold}}$ is the sound speed of the cold gas.

These scalings are not intuitive, and do not contain the shear velocity v_{shear} and overdensity χ which might be expected to play a role in the hot gas entrainment rate. What is their origin? And why are there discrepant scalings of $v_{\text{in}} \propto t_{\text{cool}}^{-1/2}$ (Ji et al. 2019) and $v_{\text{in}} \propto t_{\text{cool}}^{-1/4}$ (Gronke & Oh 2018, 2020a; Fielding et al. 2020)?

- *Energetics.* In steady state, cooling in the mixing layer is balanced by enthalpy flux from the hot gas, at a rate $\sim 5/2 P v_{\text{in}}$. To order of magnitude, inflow of the hot gas occurs at roughly the cold gas sound speed, $v_{\text{in}} \sim c_{\text{s,cold}}$, as seen in Eq. (2.1). This may seem surprising, since it is far below the maximum rate $\sim c_{\text{s,hot}}$ at which the hot gas can potentially deliver enthalpy. For instance, saturated thermal conduction has a heat flux $\sim P c_{\text{s,hot}}$. Why is turbulent heat diffusion so inefficient? The simulations of (Gronke & Oh 2018, 2020a) suggested that pulsations of the cold gas cloud (driven out of pressure balance with surroundings by radiative cooling) were responsible for drawing in hot gas, in which case $c_{\text{s,cold}}$ might be a natural velocity scale. However, it is not clear why $v_{\text{in}} \sim c_{\text{s,cold}}$ should be similar in a plane parallel shear layer, where the velocity shear drives mixing.
- *Robustness to Resolution.* Perhaps the most surprising feature of the simulations is the robustness of v_{in} (or equivalently, the surface brightness Q) to numerical resolution. Most cooling occurs in the thermal front, where the gas transitions between the thermally stable phases T_{cold} and T_{hot} . It is widely accepted that for numerical convergence, such transition layers must have finite thickness (by explicit inclusion of thermal conduction) and moreover that these fronts must be numerically resolved by at least 4 cells (Koyama & Inutsuka 2004). Otherwise, cooling gas fragments to the grid scale, and the total amount of cooling is resolution dependent (‘numerical overcooling’). Most simulations mentioned above do not include explicit thermal conduction and most of the emission occurs in zones ~ 1 cell

thick – yet the surface brightness Q appears numerically converged. Surprisingly, the value of v_{in} in the high resolution calculations of a single plane parallel mixing layer (Ji et al. 2019; Fielding et al. 2020) agree with the results of Gronke & Oh (2018, 2020a,b), which embed a macroscopic cloud in a wind. In the latter case, by necessity resolution is orders of magnitude worse and the entire mixing layer is essentially unresolved. Simulations of radiative cooling in a turbulent, thermally bistable medium also show convergence in global quantities such as the density PDF, despite no explicit thermal conduction and lack of convergence in cold gas morphology (Gazol et al. 2005). The morphology of the mixing region is a strong function of resolution. For instance, the area of the cooling surface increases with resolution, and recently Fielding et al. (2020) demonstrated that the area is a fractal with $A \propto \lambda^{-1/2}$, where λ is the smoothing scale. Since the volume of the cooling region scales as $\sim A\lambda \sim \lambda^{1/2}$, one would expect the total cooling rate to be resolution dependent. Somehow it is not, even when characteristic scales such as the cooling length $c_s t_{\text{cool}}$ are highly under-resolved. It is critical to understand this, particularly in the context of prescribing resolution requirements for larger scale simulations of galaxy formation. For instance, the circumgalactic medium (CGM) in present day state of the art galaxy simulations is unconverged, with HI column densities continually rising with resolution (van de Voort et al. 2019; Peeples et al. 2019; Mandelker et al. 2019; Hummels et al. 2019; Nelson et al. 2020).

In this paper, we exploit the close parallels between a two-phase turbulent radiative front and a turbulent combustion front to understand the above issues. In the parlance of combustion fronts, hot gas is the ‘fuel’ and cold gas is the ‘oxidizer’ which ‘burn’ to give ‘ashes’ (more cold gas). There is an extensive literature on combustion which not only has theoretical and computational underpinnings, but vast experimental backing as well

– a critical component in a situation where it is unclear whether numerical hydrodynamic codes can attain the required dynamic range. We explore the distinction between laminar and turbulent radiative fronts, with a particular focus on numerical convergence and robustness to resolution.

The structure of the paper is as follows. In §2.2, we detail the implementation of radiative cooling and thermal conduction in our simulations. In §2.3, we describe 1D simulations with radiative cooling and conduction which probe the dependence of laminar fronts to resolution. In §2.4, we explore parallels between radiative fronts and turbulent combustion, and review findings from the turbulent combustion literature. Based on this, we also develop an analytic model of radiative TMLs. In §2.5, we show 3D simulations which develop a turbulent mixing layer with radiative cooling via the Kelvin-Helmholtz instability. We compare our results to analytic predictions, and a 1D mixing length model. Finally, we conclude in §2.5.8.

2.2 Methods

We carry out our simulations using the publicly available MHD code Athena++ (Stone et al. 2020). All simulations are run on regular Cartesian grids and use the HLLC Riemann solver. The individual simulation setups of 1D laminar fronts and 3D turbulent fronts are described separately in §2.3.1 and §2.5.1. Here, we describe how we implement radiative cooling, present in all our simulations, and thermal conduction, present in all 1D simulations and a subset of 3D simulations (§2.5.5).




Figure 2.1: Heating and cooling rates as a function of temperature, with a temperature floor set at 10^4 K. The two stable phases are at 10^4 K and 10^6 K.

2.2.1 Radiative Cooling

The net cooling rate per unit volume is usually written as

$$\rho\mathcal{L} = n^2\Lambda - n\Gamma, \quad (2.2)$$

where Λ is the cooling function and Γ is the heating rate. For radiative cooling, we assume conditions of collisional ionization equilibrium and solar metallicity ($X = 0.7$, $Z = 0.02$). We obtain our cooling curve by performing a piece-wise power law fit to the cooling table given in Gnat & Sternberg (2007) over 40 logarithmically spaced temperature bins, starting from a temperature floor of 10^4 K. We then implement the fast and robust exact cooling algorithm described in Townsend (2009). We also add in a density dependent heating rate such that we have a thermally bistable medium. The cooling and heating curves that we used are shown in Fig. 2.1. While the inclusion of heating for a formally bistable medium is important in comparing to exact analytic solutions for the 1D front, it is inconsequential to the cooling rates in 3D simulations; we thus resort to a fixed temperature floor as well as setting the cooling rate in the hot medium ($T > 0.5T_{\text{hot}}$) to be zero in our 3D simulations. For some tests, we used a different shape of the cooling curve, which we specify in the relevant respective section.

2.2.2 Thermal Conduction

The conductive heat flux is $Q = -\kappa\nabla T$, where the thermal conductivity of an ionized plasma is given by Spitzer (1962):

$$\kappa_{\text{sp}} = 5.7 \times 10^{-7} T^{2.5} \text{ erg cm}^{-1} \text{ s}^{-1} \text{ K}^{-1}. \quad (2.3)$$

Instead of using the above temperature dependent conductivity, we assume a constant

conductivity equivalent to the value of κ_{sp} at the temperature of the warm gas, $T = 0.8 \times 10^5$ K. This is numerically convenient but does not significantly change the results presented. The conductivity we use where applicable is hence

$$\kappa = 10^6 \text{ erg cm}^{-1} \text{ s}^{-1} \text{ K}^{-1}. \quad (2.4)$$

As thermal conduction is a diffusive process, it is normally computationally expensive to implement. We employ a two moment approximation method for conduction similar to the approach used for implementing cosmic rays in Jiang & Oh (2018). This is done by introducing a second equation

$$\frac{1}{V_m^2} \frac{\partial Q}{\partial t} + \nabla E = -\frac{\rho Q}{(\gamma - 1)\kappa}, \quad (2.5)$$

with an effective propagation speed V_m . The latter represents the ballistic velocity of free electrons, which is $\sim \sqrt{m_p/m_e} \sim 45$ times larger than the gas sound speed¹. In the limit that V_m goes to infinity, the equation reduces to the usual equation for heat conduction. As long as V_m is large compared to the speeds in the simulation, the solution is a good approximation to the true solution. We check that our results are converged with respect to V_m (cf. Appendix A). The timestep of this approach scales as $O(\Delta x)$, compared to traditional explicit schemes which scale as $O(\Delta x^2)$. Implicit schemes which also have a linear scaling with resolution are constrained by the fact that they require matrix inversion over the whole simulation domain, which can be slow and hinders parallelization. The module employs operator splitting to compute the transport and source terms, using a two step van-Leer time integrator; the source term is added implicitly. We thank Y.F. Jiang for providing the code in advance of publication.

¹The analogous quantity in Jiang & Oh (2018) is the reduced speed of light for free-streaming cosmic rays.

2.2.3 Parameter studies

In our simulations, we vary thermal conduction and cooling strength, and will henceforth refer to them as constant multiples of the fiducial values described above. To adjust the cooling strength, we change the normalization of the cooling curve via multiplication by a constant prefactor Λ_0 . Physically, a change in the cooling time is usually due to a change in the ambient pressure; adjusting the normalization for the cooling curve achieves the same result and is more numerically convenient. In the stratified CGM, the cooling time is a function of radius. Similarly, to adjust conduction, we multiply the conductivity by a prefactor κ_0 . The cooling function and conductivity in a given simulation are thus given by

$$\Lambda(T) = \Lambda_0 \Lambda_{\text{fid}}(T) \quad (2.6)$$

$$\kappa = \kappa_0 \kappa_{\text{fid}}, \quad (2.7)$$

where Λ_{fid} and κ_{fid} are the fiducial cooling profile and conductivity given in Fig. 2.1 and Eq. (2.4) respectively.

In radiative mixing layers, radiative cooling is balanced by enthalpy flux (Ji et al. 2019),

$$Q \approx \frac{5}{2} P v_{\text{in}}, \quad (2.8)$$

where Q is the surface brightness. Hence, measuring Q or v_{in} are equivalent. We focus on measuring Q as it is a frame-independent quantity.

2.3 1D Simulations: Laminar Fronts

A large focus of this paper is on resolution requirements and convergence issues. As we shall see, the structure of the front depends strongly on whether the flow is laminar or turbulent, and on the dominant heat diffusion mechanism: thermal conduction, turbulence, or numerical diffusion. We first study the behavior of laminar flows in 1D simulations with thermal conduction and cooling. In the parlance of turbulent combustion discussed at length in §2.4, this gives us insight into the behavior of the laminar flame speed S_L and associated convergence issues. Conventional wisdom (e.g, Koyama & Inutsuka 2004) holds that it is necessary to (a) include explicit thermal conduction, and (b) resolve the *smallest* Field length in the problem (usually of the coldest gas), in order for calculations to be numerically converged. This is unequivocally true if we seek numerically converged temperature and density profiles. However, we shall see that if we merely seek numerical convergence in the mass flux j_x and hence the surface brightness Q , there are some subtleties which relax this requirement.

2.3.1 Setup

For time-steady thermal fronts, we can calculate the equilibrium solution by solving a set of coupled ODEs; this can then be compared to the time-dependent Athena++ simulations with varying resolution to understand convergence. For given boundary conditions, we can solve for the unique mass flux when we include both radiative cooling and conduction in the hydrodynamics equations (Kim & Kim 2013). We assume that

$\rho v^2 \ll P$, giving us the stationary state equations in the frame of the front:

$$j_x \equiv \rho v_x = \text{constant} \quad (2.9)$$

$$M_x \equiv P + \rho v_x^2 = \text{constant} \quad (2.10)$$

$$\kappa \frac{d^2 T}{dx^2} = j_x c_p \frac{dT}{dx} + \rho \mathcal{L}(T), \quad (2.11)$$

where we have defined the mass flux j_x and $c_p = \gamma(\gamma - 1)^{-1} k_B / \bar{m}$ is the specific heat at constant pressure. We now have a second order ordinary differential equation in Eq. (2.11) that can be solved numerically for the eigenvalue j_x , subject to the boundary conditions:

$$T_{-\infty} = T_1 \quad T_{+\infty} = T_2 \quad \frac{dT}{dx} \Big|_{\pm\infty} = 0. \quad (2.12)$$

Once we solve for j_x , we can confirm that the approximation $\rho v^2 \ll P$ holds.

Integrating Eq. (2.11) also yields a relationship between the mass flux and the cooling flux Q :

$$j_x = \frac{Q}{c_p(T_2 - T_1)}; \quad Q = - \int_{-\infty}^{\infty} \rho \mathcal{L} dx. \quad (2.13)$$

Equation (2.13) makes clear that the mass flux j_x depends on the detailed temperature and density profile within the front; thus requiring that the structure of the front be resolved. Whether the front condenses or evaporates is given by the sign of j_x , which in turn depends on the pressure of the system (Zel'Dovich & Pikel'Ner 1969). This is equivalent to a criterion on the cooling time, $t_{\text{cool}} \propto P^{-1}$. At some critical pressure P_{crit} , $j_x = 0$ and the front is static. For $P > P_{\text{crit}}$ ($P < P_{\text{crit}}$), cooling (heating) dominates and hence there is a net mass flux from the hot (cold) phase to the cold (hot) phase. We are

interested in cold gas mass growth, and so focus on the condensing case.

From Eq. (2.11), we can write down two relevant length scales set by conduction (Kim & Kim 2013) - the diffusion length, over which conduction balances mass flux, and the Field length, over which conduction balances radiative cooling (Begelman & McKee 1990). Figure 2.2 shows that the advective term is much smaller than the other two terms in Eq. (2.11), which balance one another. Thus the Field length

$$\lambda_F = \sqrt{\frac{\kappa T}{n^2 \Lambda}} \quad (2.14)$$

is the relevant scale here. It was previously found in studies of thermal instability with radiative cooling that this length scale needed to be adequately resolved in order for simulations to converge (Koyama & Inutsuka 2004; Kim & Kim 2013).

To verify the numerical solution for the steady front equilibrium and test for convergence with resolution, we set up the solution profile in a series of Athena++ simulations. We first initialize the simulation domain as a one dimensional box with $x = [-100, 300]$ pc. When we reduce resolution, we switch to a larger box with a range of $x = [-400, 1200]$ pc to avoid boundary effects. The front profile is generated by numerically solving the ODE for the steady state solution, and centered such that it has a temperature of 10^5 K at $x = 0$. The left side has an initial temperature of 10^4 K and a number density of 10^{-2} cm^{-3} , while the right side has an initial temperature of 10^6 K and a number density of 10^{-4} cm^{-3} . These correspond to the cold and hot stable equilibrium states respectively where the net cooling rate is zero. Outflowing boundary conditions are used at both ends. With the above setup, we perform a resolution study over four orders of magnitude in order to identify what scale lengths need to be resolved in the simulation. We perform three resolution sweeps, one with the fiducial cooling curve where $\Lambda_0 = \kappa_0 = 1$, one with very strong cooling where $\Lambda_0 = 128$, corresponding to the strongest cooling used in




Figure 2.2: *Lower panel:* Temperature profile solution of the thermal front. *Upper panel:* Corresponding conductive, advective and cooling terms in Eq. (2.11) across the thermal front.

§2.5, and one with weak conduction where $\kappa_0 = 0.1$. By varying thermal conduction or radiative cooling at similar Field lengths, we can probe how convergence changes when the Field length is under-resolved but the relative influence of numerical and explicit diffusion is different.

2.3.2 Results

The bottom panel of Fig. 2.2 shows the temperature profile of the solution obtained from the ODE in Eq. (2.11), which we solve via the shooting method. The final solution corresponds to an inflow velocity of 2 km s^{-1} in the hot gas. The top panel of Fig. 2.2 shows the relative importance of the conductive and advective terms in Eq. (2.11); conduction balances cooling over most of the front. By varying the parameters κ_0 and Λ_0 , we find that $j_x, Q \propto (\kappa/t_{\text{cool}})^{1/2}$, as expected from Eqs. (2.13) and (2.14). This provides a reference solution which we compare against the Athena++ results in our resolution study.

The results of the resolution study are shown in the top panel of Fig. 2.3. First consider the runs with conduction. With increasing resolution, we see convergence towards the mass flux computed from Eq. (2.13). At high resolution, when the Field length is resolved, the structure of the thermal front is resolved and agrees with the reference solution. However, as we lower the resolution, Q deviates from the reference solution, and increases steadily $Q \propto \sqrt{\Delta}$ (where Δ is the grid scale). These are marked by dashed lines in Fig. 2.3. In the runs with no thermal conduction, only numerical diffusion balances cooling. Convergence vanishes and throughout the entire range, $Q \propto \sqrt{\Delta}$; as $\Delta \rightarrow 0$, $Q \rightarrow 0$. This is in line with the expectation that for zero conduction, there should be a vanishing mass flux. All of the behavior in the fiducial case is in agreement with canonical expectations.

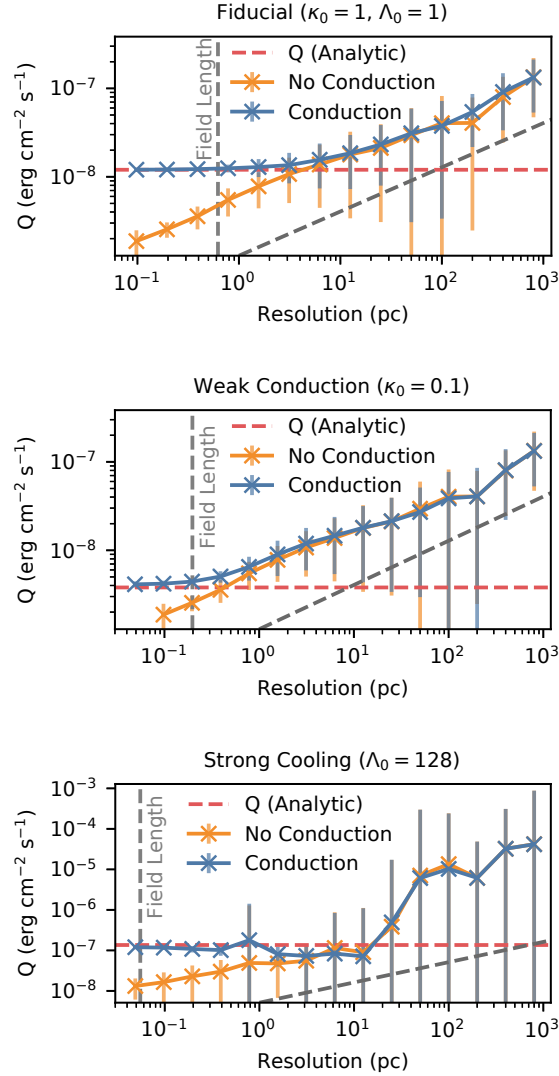


Figure 2.3: Resolution Study: Top panel shows 1D runs with fiducial cooling and conduction. The next two panels reduce the Field length λ_F , either by reducing conduction (middle panel), or increasing cooling (bottom panel). While the reduced conduction case shows resolution dependence once the grid scale $\Delta > \lambda_F$, the increased cooling case shows that Q remains converged over two orders of magnitude even when λ_F is unresolved. Clearly, the Field length is not the key criterion when determining convergence. In the text, we argue instead that it is the relative strengths of thermal and numerical diffusion that matters. Field lengths, analytic solutions and $\sqrt{\Delta}$ scalings (grey dashed lines) are also shown for comparison.

Since – in line with previous expectations – the Field length $\lambda_F \propto \sqrt{\kappa t_{\text{cool}}}$ appears to be the critical scale which must be resolved, we reduce it in two ways, either by reducing the cooling time t_{cool} or reducing the conductivity κ . We find that these two procedures do not give the same result for the same reduced Field length. If we keep cooling fixed but reduce conduction (middle panel of Fig. 2.3), then Q becomes resolution dependent once $\Delta \gtrsim \lambda_F$, as expected. By contrast, in the setup with strong cooling, (bottom panel of Fig. 2.3), we find that the mean cooling Q is slightly lower but still close to the converged value for $\Delta \lesssim 100\lambda_F$, even though the Field length is completely unresolved. Instead, lower resolution results in rapid temporal oscillations in Q , which increase in amplitude for lower resolution. Instead of an offset, Q simply oscillates about the correct equilibrium value. Similar behavior is also observed in the 3D simulations as described further below in §2.5.6 and shown in Fig. 2.17. Here, we plot the mean value of Q , while error bars indicate the standard deviation.

These results make clear that one must distinguish between errors due an unresolved front (stiff source terms) and errors due to numerical diffusion. In our case, an unresolved front contributes to the variance of the solution (numerical dispersion), but does not bias the solution. It can be beaten down by time averaging. Numerical diffusion, on the other hand, unavoidably biases the solution. The criterion for a converged solution is therefore not $\Delta < \lambda_F$, but rather $D_{\text{num}} < D_{\text{thermal}}$; i.e. that numerical diffusion is subdominant to thermal diffusion.

To expand on this point: the static radiative interface is a stiff problem where the source term (radiative cooling) defines a length scale (the Field length, over which thermal diffusion and radiative cooling balance) which is often much smaller than other scales of interest and can lie below the grid scale. It is well-known that hyperbolic systems with a stiff source term which is unresolved can have wave speeds which are either spurious (e.g., see Colella et al. 1986 for detonation waves), or still centered about the

correct value, albeit with a larger dispersion (e.g., LeVeque 2002, see §17.10-17.18). Relaxation systems are known to be well-behaved if certain subcharacteristic requirements are satisfied; although the reason is still not fully understood (Pember 1993). At least with Athena++, which uses a stable, second-order accurate modified Gudunov method for handling stiff source terms (Sekora & Stone 2010), and the two-moment conduction module we have used, radiative thermal fronts appear to fall into this class of problem, potentially because the sound speed of the cold gas sets a characteristic velocity scale. When the Field length is not resolved, cooling and conduction cannot balance exactly due to discretization errors in the temperature and its derivatives. Instead, they (and hence Q) oscillate about thermal balance and the true answer. While numerical diffusion creates systematic biases in the true steady state solution, numerical dispersion creates fluctuating errors which can be averaged out over a long time series. Of course, also the latter can only buy a limited amount of dynamic range before errors swamp the solution (in the example shown in the lower panel of Fig. 2.3, it is ~ 2 orders of magnitude). We quantify this effect below.

The Field length is a strong function of temperature, and despite the fact that (under isobaric conditions) volumetric emissivity peaks at few $\times 10^4$ K, the contribution to the integral surface brightness is more broadly distributed:

$$\begin{aligned}
 Q &= \int \rho \mathcal{L} dx = \int \rho \mathcal{L} \frac{T}{T'} \frac{dT}{T} \\
 &\approx \int \rho \mathcal{L} \lambda_F d(\log T) \propto \left(\frac{\Lambda(T) \kappa}{T} \right)^{1/2}, \tag{2.15}
 \end{aligned}$$

where we have used $nT \approx \text{const}$ and $T' \approx T/\lambda_F$. Figure 2.4 shows the integrand $\rho \mathcal{L} \lambda_F$. It has two distinct peaks at $T \sim 10^4$ and 10^5 K for the $\kappa = \text{const}$ case considered here, and is dominated by higher temperatures $T > 10^5$ K for the more realistic case of Spitzer conduction. Even though the volumetric emissivity peaks at $T \sim 10^4$ K, Q has




Figure 2.4: The distribution of net cooling with temperature. The solid line shows the integrand of Eq. (2.15), while the dashed line shows $Q(T_{\text{floor}} = T)$, where the lower limit of the integral in Eq. (2.15) is clipped at T_{floor} . Note that Q has contributions from a broad range of temperatures from $10^4 \text{K} < T < 10^5 \text{K}$, and does not plummet until $T_{\text{floor}} > 10^5 \text{K}$.

contributions from a broad range of temperatures, because the Field length is a strongly increasing function of temperature:

$$\lambda_F = \left(\frac{\kappa T}{n^2 \Lambda(T)} \right)^{1/2} \propto \left(\frac{T^{3+n}}{\Lambda(T)} \right)^{1/2}, \quad (2.16)$$

where $\kappa \propto T^n$ (and $n = 0$ for $\kappa = \text{const}$, $n = 5/2$ for Spitzer conduction). Thus, even when $\lambda_F(T \sim 10^4 \text{ K})$ is unresolved, Q will be approximately correct as long as $\lambda_F(T \sim 10^5 \text{ K})$ is resolved. This explains well the numerical results shown in the lower panel of Fig. 2.3, for which $\lambda_F(T \sim 10^5 \text{ K})/\lambda_F(T \sim 10^4 \text{ K}) \sim 30$.

In summary: in our 1D simulations, in the absence of thermal conduction, the surface brightness is resolution dependent $Q \propto \sqrt{\Delta}$. If explicit thermal conduction is included and larger than numerical diffusion, then Q is numerically converged, even if the Field length is unresolved. The unresolved Field length merely contributes to an increased variance. However, once $\Delta > \lambda_F(T \sim 10^5 \text{ K}) \sim 30\lambda_F(T \sim 10^4 \text{ K})$, the error bars grow rapidly.

2.4 Analytic Estimates from Turbulent Combustion

In 3D, turbulence in the mixing layer complicates matters considerably. In this section, we explore parallels between radiative and combustion fronts, and review findings from the turbulent combustion literature. Based on this, we also develop an analytic model of radiative TMLs.

2.4.1 Introduction & Terminology

There are close parallels between a two-phase radiative front and a combustion front. In a radiative front, the ‘fuel’ is hot gas, which ‘burns’ (i.e. cools radiatively) in a

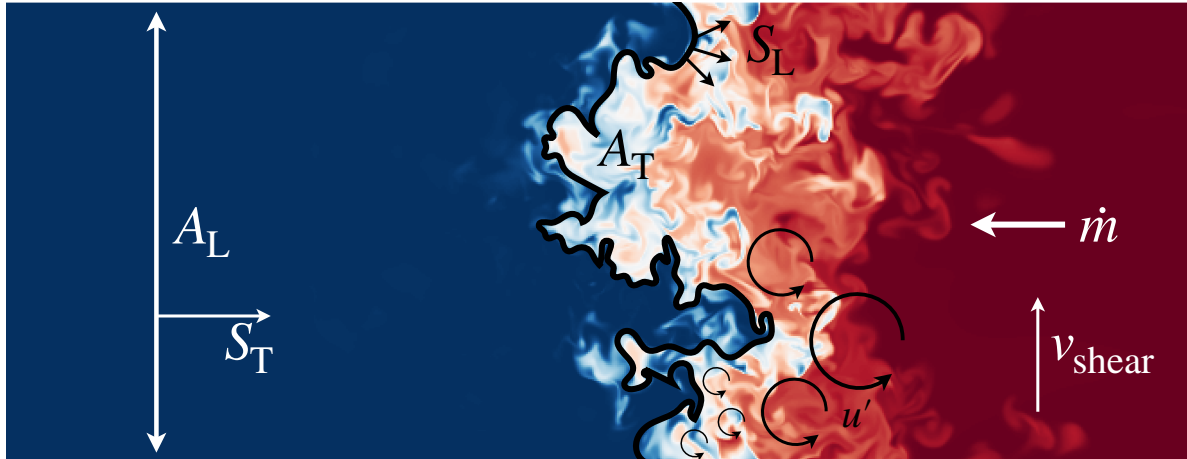


Figure 2.5: Slice through a mixing layer simulation with the relevant quantities from combustion theory marked.

temperature and density dependent manner upon mixing with ‘oxidizer’ (cold gas) to produce ‘ash’ (more cold gas). The one unusual characteristic in radiative fronts is that the end product is more ‘oxidizer’. Moreover, combustion fronts share the property that the rate of burning and hence speed of front propagation is determined by conditions within the front, which in general must therefore be resolved; it similarly reduces to a non-linear eigenvalue problem in 1D (Zeldovich et al. 1985). Due to the obvious bearing of turbulent combustion on gas chamber combustion engines in automobiles and jet engines, with implications for fuel efficiency and air pollution, the literature is vast (see Kuo & Acharya 2012 for a recent comprehensive textbook). Importantly, besides sophisticated high resolution numerical simulations, there is a plethora of experimental results. Within the astrophysical community, similar issues arise in thermonuclear burning fronts on carbon-oxygen white dwarfs, where conditions within the front determine the burning rate. The large scale separation ($\sim 10^7$) between the size of the white dwarf (which sets the outer scale of turbulence) and the front structure precludes direct calculation of the fronts in simulations; a subgrid model (e.g., Niemeyer & Hillebrandt 1995; Schmidt et al. 2006; Jackson et al. 2014) is necessary. Here, we draw upon this literature to provide an

intuitive physical picture for the puzzles described in §2.1.

In the language of combustion literature, in radiative TMLs the fuel and oxidizer are not perfectly pre-mixed before combustion. Instead, the two are initially separate. They are subsequently entrained and dispersed by large scale eddies with turbulent velocity u' . The turbulent cascade down to small scales leads to stretching, fragmentation and a vast increase in surface area, until at small scales diffusion mixes the fuel and oxidizer and allows combustion to take place. The net result is that fuel is consumed at a rate \dot{m} . The process is similar to how stirring enables a vast increase in surface area and the large mixing rate between coffee and cream, despite the long molecular diffusion time.

What is the characteristic velocity at which a thermal front propagates, and how do results depend on the source of thermal diffusion? Thermal diffusion is canonically due to thermal conduction. In the absence of turbulence, this gives rise to a ‘laminar flame speed’ $S_L \sim \sqrt{D_L/t_{\text{cool}}}$, where D_L is the customary diffusion coefficient with units of L^2T^{-1} . This can be seen by balancing thermal diffusion with cooling (Eq. (2.13) below), or simply from dimensional analysis. If, however, thermal conduction is not included, which is often the case in many numerical simulations, then all diffusion is numerical: gas mixing and thermal diffusion operate close to the grid scale. When numerical diffusion dominates, $D \sim v\Delta x$, where v is a characteristic velocity and Δx is the grid scale. Thus, in the absence of conduction, $S_L \propto \sqrt{\Delta x}$.

Turbulence gives rise to a large increase in surface area of the phase boundary $A_T \gg A_L$ which leads to a ‘turbulent flame speed’ $S_T \gg S_L$, where S_T (corresponding to v_{in} in Eq. (2.1)) characterizes the rate at which fuel is consumed and the overall front propagates. Our goal is to understand S_T , which sets the hot gas entrainment rate. In Fig. 2.5 we visualize the quantities introduced in this section.

In this paper, we consider the impact of shear driven turbulence on the radiative front. Evaporative fronts are subject to a well-known corrugational instability, the Darrieus-

Landau instability (Zeldovich et al. 1985; Landau & Lifshitz 1987; Inoue et al. 2006), which deforms the front and drives turbulence, which can also lead to increased surface area and accelerated reaction rates. Here, we focus only on condensation fronts.

2.4.2 Characteristic Regimes

Turbulent combustion is characterized by several dimensionless numbers (Kuo & Acharya 2012). These give rise to classification into many distinct burning regimes which are typically shown on a plot known as a Borghi diagram. For our purposes, the most useful parameter is the Damköhler number:

$$\text{Da} = \frac{\tau_{\text{turb}}}{\tau_{\text{react}}} = \frac{L}{u't_{\text{cool}}(T)}, \quad (2.17)$$

which gives the ratio of the eddy turnover time at the outer scale L , with turbulent velocity u' , to a reaction time². For us, it can also be viewed as $\text{Da} \sim L/L_{\text{cool}}(T)$, the ratio of the integral length scale of turbulence to the cooling length $L_{\text{cool}}(T) \sim u't_{\text{cool}}(T)$. Note that the cooling time $t_{\text{cool}}(T)$ varies with temperature through the front. The Damköhler number separates two asymptotic regimes, $\text{Da} \ll 1$ ('well-stirred reactor') and $\text{Da} \gg 1$ ('corrugated flamelets') which are universal across all classification schemes. The cooling time $t_{\text{cool}}(T)$ decreases continually across the front, as the temperature T declines from the hot to the cold gas temperature (T_h and T_c , respectively) and the cooling function peaks extremely close to T_c . Initially, close to the hot gas boundary, $\text{Da} < 1$. Turbulence cascades to the diffusion (grid) scale before the two components

²Another important parameter in combustion studies is the Karlovitz number $\text{Ka} = \delta^2/\eta^2$, where δ is the diffusive scale (corresponding to the Field length) and η is the Kolmogorov scale at which viscosity damps turbulence. It determines whether the propagation of the small scale interface is set by laminar burning, or whether turbulence alters the structure of the diffusive front. Since we do not have explicit thermal conduction or viscosity, the flamelet fronts are generally unresolved and we have $\delta \sim \eta \sim \Delta x$, i.e. $\text{Ka} \sim 1$ in our simulations.

react. Fuel and oxidizer are well mixed and thus the reaction rate is uniform across the entire volume. This is known as the ‘well-stirred reactor’ regime. In mixing length theory, this can be characterized by a turbulent diffusion coefficient $D_{\text{turb}} \sim u'L$. Thus, in the weak cooling ($\text{Da} < 1$) regime, we expect:

$$S_{\text{T}} \approx \left(\frac{D_{\text{turb}}}{t_{\text{cool}}} \right)^{1/2} \approx \left(\frac{u'L}{t_{\text{cool}}} \right)^{1/2}. \quad (2.18)$$

However, as we move toward the cold gas boundary, the temperature and cooling time fall. When $\text{Da} > 1$, burning proceeds before mixing is complete, and combustion thus takes place inhomogeneously. In our context, inhomogeneous cooling leads to fragmentation into a multiphase medium. The criterion $\text{Da} \sim 1$ corresponds to the transition between single and multiphase structure in the mixing layer. The steep temperature dependence of cooling means that most cooling takes place in thin unresolved fronts close to $T \sim 10^4\text{K}$ at the interface between cold and hot gas. The turbulent cascade wrinkles this interface and vastly increases its area, thus increasing the volumetric cooling rate. In a famous paper, Damköhler (1940) conjectured that the increase in surface area leads to a turbulent flame speed:

$$S_{\text{T}} \approx S_{\text{L}} \left(\frac{A_{\text{T}}}{A_{\text{L}}} \right), \quad (2.19)$$

where A_{T} and A_{L} are the turbulent and laminar flame areas. This comes from simply equating the mass flux through A_{L} at velocity S_{T} with the mass flux through A_{T} at velocity S_{L} , as illustrated in Fig. 2.5. This intuitive notion can be made more precise and proven (Bray & Cant 1991).

2.4.3 Scalings of the turbulent velocity S_T

To calculate S_T , we therefore need to know A_T/A_L . This has no single consensus answer; for instance, Table 5.1 of Kuo & Acharya (2012) lists 20 fits to S_T/S_L obtained from theory, simulation and experiment depending on geometry, boundary conditions and flame wrinkling process. One simple way of parametrizing most known scalings is to write:

$$\frac{S_T}{S_L} = \frac{A_T}{A_L} = 1 + \left(\frac{u'}{S_L}\right)^n \approx \left(\frac{u'}{S_L}\right)^n, \quad (2.20)$$

where the last equality holds for $u' \gg S_L$. The most well-known scaling is $n = 1$ (Damköhler 1940), which has substantial experimental support in a variety of settings. For instance, Libby et al. (1979); Clavin & Williams (1979); Peters (1988); Bray (1990); Bedat & Cheng (1995) obtain similar scalings in both theory and experiment. It implies

$$S_T \approx u', \quad (2.21)$$

i.e. that the combustion front simply propagates at the turbulent velocity. A useful geometrical interpretation comes from Damköhler (1940) and Shchelkin (1943) who considered the distortion of the flame-burning front into several ‘Bunsen cones’ – analogous to a Meker burner. A simplified version of his argument is as follows: consider a flat interface of area $A_L = L^2$. It propagates in a direction normal to the front at velocity S_L . Over a burning time $t_{\text{burn}} \sim L/S_L$, laminar burning will traverse a distance L , whereas turbulent motions traverse a distance $l_{\text{turb}} \sim u't_{\text{burn}} \sim (u'/S_L)L$, creating a wrinkled (conical) region with area $A_T \sim l_{\text{turb}}L \sim (u'/S_L)A_L$. Thus, $A_T/A_L \sim u'/S_L$. A more careful consideration of the conical geometry gives

$$\frac{A_T}{A_L} \sim \left(1 + \left(\frac{2u'}{S_L}\right)^2\right)^{1/2}, \quad (2.22)$$

which reduces to $A_T \propto u'$ for $u' \gg S_L$. Note that in our context, $S_T \sim u'$, independent of all other parameters, including S_L , which in general is resolution dependent.

Thus far, we have ignored the influence of other parameters. As previously mentioned, turbulent combustion is in fact characterized by at least two dimensionless numbers in a Borghi diagram, typically either (Re_L, Da) or $(\eta/\delta, u'/S_L)$, where δ is the thermal diffusive scale and η is the Kolmogorov scale. An important boundary in the Borghi diagram is the Klimov-Williams line, where $\text{Ka} = (\delta/\eta)^2 \sim 1$, where laminar flame scales and turbulent stretching scales become comparable. As noted earlier, in numerical codes where numerical diffusion is dominant, we expect $\delta \sim \eta \sim \Delta$, so that $\text{Ka} \sim 1$, and we are always in this regime. In a broad neighborhood of the Klimov-Williams line, flame propagation has been argued to obey the scaling (Gülder 1991):

$$\frac{S_T}{u'} = \text{Da}^{1/4} = \left(\frac{L}{u' \tau_{\text{react}}} \right)^{1/4}, \quad (2.23)$$

which fits a large body of burning velocity data (Gülder 1991; Zimont et al. 1995). Note that this is precisely the $v_{\text{mix}} \approx c_{\text{s,cold}} (t_{\text{cool}}/t_{\text{sc,cold}})^{-1/4}$ scaling previously reported (Gronke & Oh 2018, 2020a; Fielding et al. 2020) if we identify $u' \approx c_{\text{s,cold}}$ and $\text{Da} \sim L/(u' t_{\text{react}}) \sim t_{\text{sc,cold}}/t_{\text{cool}}$ – but backed up by experimental data. In this work, we will test and confirm the resolution independence of v_{in} (cf. §2.5.6), and thus continue with Eq. (2.23) as our ‘fiducial’ scaling in the strong cooling regime where $\text{Da} > 1$. Note that while combustion theory can provide a link between S_T and u' and helps us understand the core questions presented in §2.1, the scaling of u' with respect to the flow properties depend on the turbulent driving process and have to be found from numerical experiments (see §2.5.3).

2.4.4 Details of the Fiducial $S_T/u' = \text{Da}^{1/4}$ Scaling

What is the theoretical justification for Eq. (2.23)? Following Tennekes (1968); Kuo & Corrsin (1972), Gülder (1991) argues that turbulent vortex tubes should be separated by a distance of order the Taylor microscale³. Assuming that laminar burning fronts must cover a distance of order the Taylor microscale to complete burning, he arrives at Eq. (2.23). We do not recount his arguments here, but instead refer interested readers to the original paper.

It is not clear how applicable the Gülder (1991) argument is to our numerical simulations, which do not have explicit viscosity and a well defined Reynolds number, and thus do not have a well-defined Taylor microscale (which will vary with resolution). For us, the most important fact is that there is significant experimental evidence in turbulent combustion data for the scaling in Eq. (2.23), which do not suffer from the same limitations as our numerical simulations. Here, we propose a simpler alternative argument which gives similar results.

What is the effective cooling time $\tilde{\tau}_{\text{cool}}$ of an inhomogeneous medium where $\text{Da} > 1$? It is clearly not the standard cooling time t_{cool} , since only a small fraction of the medium is cooling. Consider cold gas as a scalar pollutant, which diffuses over scales λ on a timescale $t_D \sim \lambda^2/D_{\text{turb}}$. In the fast cooling ($\text{Da} \gg 1$) limit, all of the mixed gas will cool. Over a cooling time, the cold gas diffuses over a distance (setting $t_D \sim t_{\text{cool}}$):

$$\lambda_{\text{cool}} \sim \sqrt{D_{\text{turb}} t_{\text{cool}}} \sim \sqrt{Lu' t_{\text{cool}}} \quad (2.24)$$

Thus, after a cooling time t_{cool} , only a fraction $f_{\text{cool}} \sim \lambda_{\text{cool}}/L$ of the gas in an eddy has mixed and cooled; only after $N_{\text{cool}} \sim f_{\text{cool}}^{-1} \sim L/\lambda_{\text{cool}}$ cooling times does all the gas in the

³The Taylor microscale is a lengthscale which comes from a Taylor series expansion of flow correlations; it is the scale at which shear is maximized. While it is larger than the Kolomogorov scale, it can be thought to demarcate the end of the inertial range and the beginning of the dissipation range.

eddy cool. The effective cooling time is therefore:

$$\tilde{\tau}_{\text{cool}} \sim N_{\text{cool}} t_{\text{cool}} \sim \frac{L}{\lambda_{\text{cool}}} t_{\text{cool}} \sim \sqrt{\frac{L}{u'} t_{\text{cool}}}, \quad (2.25)$$

i.e., the geometric mean of the eddy turnover time and the cooling time. Equivalently, we can view Eq. (2.24) as the effective mean free path of a fluid element. The mean free time is therefore:

$$\tilde{\tau}_{\text{cool}} \sim \frac{\lambda_{\text{cool}}}{u'} \sim \sqrt{\frac{L}{u'} t_{\text{cool}}}, \quad (2.26)$$

which gives the same result. Hot gas in the multiphase, strong cooling region is converted to cold gas on a timescale $\tilde{\tau}_{\text{cool}}$, which is shorter than the mixing time L/u' , but longer than homogeneous cooling time t_{cool} , since only a small fraction of the volume is cooling.

Equation (2.25) is a common random walk result. For instance, consider a photon in a medium which scatters (with optical depth τ_s) and absorbs (with optical depth τ_a). Then the effective optical depth is $\tau_* \sim \sqrt{\tau_a \tau_s}$, with effective survival time $t_* \sim \sqrt{t_a t_s}$ (Rybicki & Lightman 1986). Similarly, when considering the competition between thermal conduction and cooling, the Field length:

$$\lambda_F \sim \sqrt{\frac{\kappa T}{n^2 \Lambda(T)}} \sim \sqrt{\lambda_e v_e t_{\text{cool}}} \quad (2.27)$$

(using $\kappa \sim P v_e \lambda_e / T$) is the geometric mean of the elastic (λ_e ; the Coulomb mean free path) and inelastic ($v_e t_{\text{cool}}$) mean free paths for a thermal electron. Considering the Field length as the effective mean free path for an electron, the mean free time is $t_e \sim \lambda_F / v_e \sim \sqrt{t_e t_{\text{cool}}}$, where $t_e \sim \lambda_e / v_e$. Eqs. (2.24) and (2.25) are the equivalent analogs for a turbulent eddy, with $\lambda_e \rightarrow L$, $v_e \rightarrow u'$. Note that the largest eddies dominate mixing, and so therefore all quantities related to turbulence are evaluated at the outer scale L .

Later in §2.5.4, we shall see from simulation results that (as assumed here) the turbulent diffusion coefficient $D_{\text{turb}} \sim u'L$ is relatively unaffected by cooling.

Substituting $\tilde{\tau}_{\text{cool}}$ (Eq. (2.25)) for the cooling time in the usual expression for the turbulent flame velocity (Eq. (2.18)), we obtain:

$$S_{\text{T}} \sim \left(\frac{D_{\text{turb}}}{\tilde{\tau}_{\text{cool}}} \right)^{1/2} \sim u' \left(\frac{L}{u't_{\text{cool}}} \right)^{1/4} \sim u' \text{Da}^{1/4}. \quad (2.28)$$

Beyond the turbulent flame speed (Eq. (2.28)), this ansatz makes predictions which are testable in the simulations:

- *Effective emissivity.* This model predicts an effective cooling time in the multiphase region given by Eq. (2.25), so that the effective emissivity is:

$$\tilde{\epsilon} \sim \frac{P}{\tilde{\tau}_{\text{cool}}} \sim P \left(\frac{u'}{t_{\text{cool}}L} \right)^{1/2}. \quad (2.29)$$

The $\tilde{\epsilon} \propto u'^{1/2}t_{\text{cool}}^{-1/2}$ scaling can be checked in the simulations.

- *Width of multiphase regions.* Equivalently, if Eqs. (2.28) and (2.29) hold, then we can use $Q \sim Pv_{\text{in}} \sim \tilde{\epsilon}h$ to find that the width h of the multiphase region scales as:

$$h \propto L \left(\frac{u't_{\text{cool}}}{L} \right)^{1/4} \propto \text{Da}^{-1/4}, \quad (2.30)$$

where the $h \propto (u')^{1/4}t_{\text{cool}}^{1/4}$ scaling can be tested in the simulations. Of course, of Eqs. (2.28), (2.29), and (2.30), only two are independent.

We caution once again that there does not appear to be universally applicable turbulent velocity scalings in the literature, which tend to be situation dependent. Nonetheless, it is reassuring to see that the scalings we see in our numerical simulations with limited

dynamic range have also been seen in a large body of experimental data and have theoretical justification.

2.4.5 Implications for the fractal nature of mixing layers

These properties can also be related to the fractal nature of radiative mixing layers. Recently, Fielding et al. (2020) showed that the area of the cooling surface in radiative mixing layer simulations obeys a fractal scaling, with

$$\frac{A_T}{A_L} = \left(\frac{\lambda}{L}\right)^{2-D}, \quad (2.31)$$

where λ is the smoothing scale and $D = 2.5$ was the fractal dimension argued to hold by analogy with well-known fractals, and verified in their simulations. Turbulence combustion fronts are indeed well known to be fractals, due to the dynamical self-similarity of turbulence in the inertial range. Experimental measurements by e.g. instantaneous laser tomography have given values ranging from $D = 2.1 - 2.4$ in a variety of flow geometries, with a preferred value of $D = 2.35$ (Hentschel & Procaccia 1984; Sreenivasan et al. 1989); it has been argued that this fractal dimension is universal (Catrakis et al. 2002; Aguirre & Catrakis 2005). From Eq. (2.19), the fractal dimension can be used to calculate the turbulent flame speed (Gouldin et al. 1986; Peters 1988). The fractal scaling and consequent increase in area A_T should extend all the way down to the Gibson scale λ_G , which is defined to be the scale where the turbulent velocity equals the laminar flame speed, $v(\lambda_G) = S_L$. This is often unresolved in simulations. If we use the Kolmogorov scaling $v \propto \lambda^{1/3}$, then we obtain:

$$\frac{S_T}{S_L} = \frac{A_T}{A_L} = \left(\frac{\lambda_G}{L}\right)^{2-D} = \left(\frac{u'}{S_L}\right)^{3(D-2)}, \quad (2.32)$$

where we have used Eq. (2.31) and $v(\lambda_G) = S_L$. Thus, in Eq. (2.20), we have $n = 3(D-2)$. The experimental value of $D = 2.35$ gives $n = 1.05$, in good agreement with Damköhler's scaling, and fair agreement with the scaling in Eq. (2.23). The Fielding et al. (2020) value of $D = 2.5$ gives $n = 1.5$, or $S_T = u'(u'/S_L)^{1/2}$. If one uses the laminar $S_L \propto t_{\text{cool}}^{-1/2}$ from our static simulations, this would imply $S_T \propto t_{\text{cool}}^{1/4}$. However, in the Fielding et al. (2020) model, the speed at which a cooling layer advances is $S_L \propto t_{\text{cool}}^{1/2}$, so they end up with $S_T \propto t_{\text{cool}}^{-1/4}$ as well. The scalings are sensitive to the fractal dimension D and the measurement error on D obtained from the simulations is unclear at this point. In addition, the cutoff scale of turbulence may not be the Gibson scale. We caution that fractal arguments have not proven to be fully robust in the turbulent combustion context. For instance, the measured fractal parameters fluctuate depending on the extraction algorithm, and have not been able to correctly predict the turbulent burning velocity (Cintosun et al. 2007).

2.4.6 Implications for energetics and convergence criteria

The above considerations bear upon the two over-arching questions first raised in §2.1, which will be further addressed in the course of this paper.

Energetics. Why is $S_T \sim c_{s,\text{cold}}$? From Eq. (2.21), we have $S_T \sim u'$, i.e. of order the turbulent velocity at the outer scale. The timescale of the Kelvin-Helmholtz instability, which mixes the two fluids, is $t_{\text{KH}} \sim \sqrt{\chi}L/v_{\text{shear}}$; the characteristic turbulent velocity of the interface between hot and cold gas is $u' \sim v_{\text{shear}}/\sqrt{\chi} \sim \mathcal{M}_{\text{hot}}c_{s,\text{hot}}/\sqrt{\chi} \sim \mathcal{M}_{\text{hot}}c_{s,\text{cold}}$. If $\mathcal{M}_{\text{hot}} \sim 1$, as is true for many situations in the CGM (since the virial velocity is of order the virial sound speed), this reduces to $u' \sim c_{s,\text{cold}}$. We will study detailed scalings of u' in §2.5.3.

Resolution independence. Neither our fiducial scaling (Eq. (2.28)) nor the Damköhler scaling (Eq. (2.21)) depend on the diffusion coefficient, and thus are independent of resolution. Physically, this is because most radiative cooling takes place in the $Da > 1$ regime, when the cooling time is shorter than the eddy turnover time. When cooling is ‘fast’ compared to mixing, all gas which mixes cools – the rate limiting step is the rate at which turbulence cascades to diffusive scales, whereupon mixing and cooling happen on very short timescales. The time the turbulent cascade takes to reach small scales is simply $\tau_{\text{turb}} \sim L/u'$ the eddy turnover time at the outer scale, since in Kolmogorov turbulence, the eddy turnover time $\tau_1 \sim l/v_1 \propto l^{2/3}$ is a progressively smaller function of scale. The situation is similar to passive scalar mixing, except that here the passive scalar which is being advected is temperature. The rate at which coffee mixes with cream is given by the stirring time of the spoon, independent of the details of molecular diffusion. Similarly, the rate at which hot gas mixes with cold gas and subsequently cools is given by the eddy turnover time at the outer scale, independent of the details of thermal (numerical) diffusion, which set the structure of the (often unresolved) laminar thermal fronts. Thus, the important scale that needs to be resolved is the mixing due to turbulent eddies at the outer scale.

2.5 3D Simulations: Turbulent Fronts

We next turn to 3D simulations of radiative mixing layers. Due to an additional ingredient – turbulence – not present in 1D simulations, their properties are quite different. In this section, we compare the results of 3D simulations to the model discussed in the previous section.

2.5.1 Setup

Our setup closely follows the work of Ji et al. (2019). The coordinate system is as follows: y is the axis of shear flow, x is normal to the cold/hot interface (the principal direction of interest along which front properties vary), and z is the third remaining dimension. Boundary conditions are periodic along the y and z axes and outflowing along the x axis. The bounds of the x axis are $[-100, 200]$ pc and the bounds for the y and z axes are $[0, 100]$ pc. Cold 10^4 K gas is initially located in the negative x region and hot 10^6 K gas in the positive x region, separated by a smoothly varying front centered at $x = 0$ where $T = 10^5$ K. The initial front profile is obtained by solving for the 1D steady state solution as described previously. The initial gas density is set to $n_{\text{hot}} = 1.6 \times 10^{-4} \text{ cm}^{-3}$ and $n_{\text{cold}} = 1.6 \times 10^{-2} \text{ cm}^{-3}$ in the hot and cold gases respectively. We use a resolution of $384 \times 128 \times 128$ in the box, which corresponds to a cell length of 0.78 pc. This is approximately the minimum Field length in the simulation when thermal conduction is included. We also introduce a shear velocity profile across the front that takes the following form:

$$v_y = \frac{v_{\text{shear}}}{2} \tanh\left(\frac{x}{a}\right), \quad (2.33)$$

where we set the scale length $a = 5$ pc, and the shear velocity $v_{\text{shear}} = 100$ km/s, which is of order the sound speed of the hot medium. The profile is then perturbed as follows to induce the Kelvin Helmholtz instability:

$$\delta v_x = A \exp\left(-\frac{x^2}{a^2}\right) \sin(k_y y) \sin(k_z z), \quad (2.34)$$

where we set the perturbation amplitude A to be 1% of v_{shear} . We also set the perturbation wavelength $\lambda_i = 2\pi/k_i$ to be of order the box size, and set the ballistic speed of free

electrons to be $V_m \sim 15$ times the hot gas sound speed when thermal conduction is included. The latter pertains only to §2.5.5 – thermal conduction is not included in any of the other 3D simulations. We check that results are not sensitive to these choices. Unlike adiabatic mixing layers which continue to grow over time, our mixing layers appear stable after the initial onset and development of turbulence. All quantities presented were measured in the latter half of the simulations after the mixing layers had been given sufficient time to reach this stage. The exact time periods vary between simulations, but simulations were run sufficiently long to ensure that they span at least 20 Myr. Error bars reflect the standard deviation of the measured values. While the surface brightnesses Q were saved at very small time intervals, and hence have many measurements, the turbulent velocities u' were calculated from full simulation snapshots and have a smaller (~ 10) number of measurements per simulation.

2.5.2 Morphology of Mixing Layers: Transition from Single Phase to Multiphase

We begin by examining the morphology and slice averaged properties of the mixing layer, and how these vary with cooling (or equivalently, with Damköhler number Da). We shall soon see (§2.5.4) that temperature and velocity profiles can be calculated by judicious application of mixing length theory.

The upper panels in Fig. 2.6 show the normalized profiles for the emissivity, mean temperature, and volumetric fraction of gas in the cold phase⁴, for the weak ($\Lambda_0 = 1/4$), fiducial ($\Lambda_0 = 1$), and strong ($\Lambda_0 = 8$) cooling cases respectively. They also plot the Damköhler profile. In calculating the Damköhler number, we use a fixed length scale $L = 100$ pc (the box size in the direction of the flow), but use local values of the turbulent

⁴Defined to be $T < 5 \times 10^4$ K gas.




Figure 2.6: *Upper panels:* Normalized profiles for the mean emissivity, the mean temperature, and volumetric fraction of gas in the cold phase for various cooling strengths. The profile of the Damköhler number is also shown, denoting the region where mixing is more efficient than cooling. *Lower panels:* Mean shear and rms velocity profiles for the same selection of cooling strengths. The red dashed line shows the gradient of the shear velocity $\nabla(v_y)$ multiplied by a mixing length $\ell_{\text{mix}} = 15$ pc, which traces the rms velocity profile well.

velocity u' and cooling time t_{cool} measured from the simulation. The initial Damköhler number $\text{Da} = L/(u't_{\text{cool}})$ in the hot medium is small due to the extremely long cooling times. However, as mixing proceeds and the mean temperature falls in the mixing layer, the cooling time falls and the Damköhler number rises toward cooler regions. The fact that the Damköhler number is a function of position within the mixing layer is important for understanding some key properties. Note that the mixing layer has roughly constant pressure. There are small pressure fluctuations seeded by cooling which are compensated by increased turbulent pressure support, so that $P + \rho u'^2 \approx \text{const}$ (see Figure 8 of Ji et al. 2019), but these fluctuations are sufficiently small ($\delta P/P < 10\%$) that isobaric cooling is a good approximation.

The lower panels in Fig. 2.6 show the corresponding mean and rms velocity profiles. The rms velocity is calculated by first subtracting off the mean flow in both the y (flow) and x (normal to cold/hot interface) directions. We have explicitly checked that the velocity dispersion is roughly isotropic ($\sigma_x^2 = \sigma_y^2 = \sigma_z^2$), a sign of well-developed turbulence. While stronger shear flows do display more anisotropy, the difference stays within a factor of two. It is interesting such isotropy can arise, despite the strong anisotropy in mean flow. The velocity dispersion σ_z^2 is a particularly good indicator, since there is no mean flow in the z direction.

Figure 2.6 reveals a number of interesting properties:

- The criterion $\text{Da} = 1$ roughly controls the transition from single phase to multiphase gas, when the cold gas fraction first becomes non-zero. In the weak and fiducial cooling cases, the mean temperature falls substantially in the single phase regime (the ‘well-stirred reactor’, in the language of §2.4), before the gas turns multiphase. Thus, the mean temperature profile and the cold gas profile do not track one another. A substantial amount of the cooling flux is emitted in the single-

phase regime. However, for the strong cooling case, all cooling takes place in the multiphase regime. In this case, the mean temperature profile tracks the cold gas profile; $\bar{T} \approx f_{\text{cold}}T_{\text{cold}} + (1 - f_{\text{cold}})T_{\text{hot}}$.

- The turbulent velocity tracks the shear, $u' \propto \nabla v_y$. This is expected from mixing length theory, where $u' \approx l\nabla v_y$, and l is the mixing length. We discuss this further in §2.5.3.
- The normalized emissivity has an approximately Gaussian profile, as one would expect if cooling balances the divergence of turbulent diffusion. A diffusive process will of course have a Gaussian profile. For instance, in a multiphase medium, the fractal hot/cold gas boundary (where most of the cooling takes place in a thin sheet) has a Gaussian distribution of displacements from the mid-point, as expected for a random walk. The emissivity tracks the cold gas fraction rather than the mean temperature profile, peaking at $f_{\text{cold}} \approx 0.5$. This makes sense, since the surface area of the hot-cold interface (which dominates cooling) peaks when $f_{\text{cold}} = 0.5$. The emissivity profile becomes narrower in the strong cooling regime. Later, we shall see that the area under the blue curves $Q \propto t_{\text{cool}}^{-1/2}, t_{\text{cool}}^{-1/4}$ in the weak and strong cooling regimes respectively.
- In the strong cooling regime, the cooling emissivity and turbulence track one another closely. Both peak at the same spatial location (where $f_{\text{cool}} \approx 0.5$). This was predicted by Eq. (2.29), where $\tilde{\epsilon} \propto u'^{1/2}$. This is consistent with our model, where turbulent mixing regulates the fraction of gas available for cooling.
- The mean temperature profile and mean velocity profile also track one another quite closely, corresponding to $\mathcal{M} \approx 1$ in the shear layers (see also Figure 9 of Ji et al. 2019). Thus, for instance, cooling causes the shear profile to narrow in moving from

weak to fiducial cooling. This makes sense since turbulent diffusion governs both momentum and thermal transport. Furthermore, in the strong cooling case, the cooling emissivity also tracks the shear profile: $\tilde{\epsilon} \propto u^{1/2} t_{\text{cool}}^{-1/2} \propto (\nabla v_y)^{1/2} t_{\text{cool}}^{-1/2}$. This correspondence fails when sink/source terms in the energy equation which are not present in the momentum equation become dominant: (i) very strong cooling (see the low temperature portion of the strong cooling regime), or (ii) highly supersonic flow in the hot medium (not shown). In the latter case, shocks and turbulent dissipation heat the gas, and so the hot region remains hot even when significant cool gas is mixed in. These effects narrow the temperature profile relative to the velocity profile.

The distinction between the weak (single phase) and strong (multiphase) cooling regimes can be clearly seen in Fig. 2.7. At first glance, both cases appear to be similar, except that the weak cooling case has a broader mixing layer (top panel). However, it is already apparent that there is a lot more intermediate temperature ($T \sim 10^5\text{K}$) gas in the weak cooling case. We can also see this in the temperature slices, which only show the ‘multiphase’ portion of both cases (when f_{cold} is non-zero). For weak cooling, the ‘hot’ phase in this regime is significantly cooler than $T = 10^6\text{K}$, the initial temperature of the hot gas – it has cooled via efficient mixing in the single-phase regime. By comparison, the temperature contrast between the two phases is much higher in the strong cooling case, with a clearly bimodal temperature distribution centered at $T \sim 10^4\text{K}$ and $T \sim 10^6\text{K}$. In both cases, the amount of intermediate temperature ($T \sim 10^5\text{K}$) gas peaks when $f_{\text{cool}} \sim 0.5$, where the emissivity also peaks. In the strong cooling case, cooling is clearly dominated by the very thin (unresolved) interface between the phases, as can be seen in the emissivity slices (bottom panel). This is less true in the weak cooling case, where a larger fraction of the volume contributes to cooling (note the low filling factor of interface




Figure 2.7: Slices of temperature and cooling for the ‘low cooling’ (left) and ‘high cooling’ regime where the cooling function has been reduced or amplified by a factor of four, respectively. The top row shows a temperature slice orthogonal to the flow while the middle and lower panels show temperature and emissivity at three different depths in the mixing layer (marked in the top with grey dashed lines). The cooling slices have been normalized by the boost factor of $\Lambda_0 = 1/4$ and $\Lambda_0 = 4$ for the left and right panels, respectively, to allow comparison of interface widths. We observe that the gas is strongly multiphase, especially when cooling is stronger, and that cooling happens mostly at the interfaces between the two phases. When cooling is weaker, these interfaces are thicker. This is consistent with the idea that they are defined by a diffusive length $\lambda_F \propto t_{\text{cool}}^{1/2}$.

regions at $f_{\text{cold}} = 0.5$, when cooling peaks). Furthermore, the interface regions (which should scale as $l \propto t_{\text{cool}}^{1/2}$) are now broader and clearly resolved.

2.5.3 Scaling Relations

The key theoretical quantity of interest in radiative mixing layers is the hot gas entrainment rate v_{in} , or equivalently the surface brightness Q (assuming that hot gas enthalpy flux balances cooling). This determines the rate at which hot gas is converted to cold gas, which has many important implications, amongst them the ability of cold gas to survive in the face of hydrodynamic instabilities (Gronke & Oh 2018). In previous

work, we derived the mass growth rate shown in Eq. (2.1). However, this was performed at relatively low resolution. Higher resolution work similar to that done here (Ji et al. 2019) also obtained scaling relations, with some important differences. However, their results relied on a rather small number of simulations. Here, we clarify the nature of the scaling relations using a larger set of simulations, and thus put the results obtained from previous studies in a broader context. In particular, we explicitly test⁵ the predicted scaling relations Eqs. (2.18) and (2.28), which state that in the weak cooling regime, $Q \propto v_{\text{in}} \propto u^{1/2}(L/t_{\text{cool}})^{1/2}$, while in the strong cooling regime, $Q \propto v_{\text{in}} \propto u^{3/4}(L/t_{\text{cool}})^{1/4}$, with no additional dependence on other parameters such as overdensity χ and flow Mach number (relative to the hot gas sound speed) \mathcal{M} (§2.5.3). We then test scalings for emissivity, or equivalently for the width of the mixing layer (§2.5.3). Finally, we test how turbulent velocities u' vary with $\chi, \mathcal{M}, t_{\text{cool}}$ in our specific setup (§2.5.3).

Scaling Relations for Q

Dependence on Cooling. Figure 2.8 shows the surface brightness Q as a function of cooling strength Λ_0 . It is clear that $Q \propto \Lambda^{1/2} \propto t_{\text{cool}}^{-1/2}$ ($Q \propto \Lambda^{1/4} \propto t_{\text{cool}}^{-1/4}$) in the weak (strong) cooling regimes, as predicted by Eqs. (2.18) and (2.28) respectively. We have already seen that the Damköhler number varies spatially across a mixing layer. Here, it is useful to define a Damköhler number characterizing a single simulation as a whole. This provides a reference point for differentiating between the two regimes. Thus, for each simulation, we have to choose a single value for the spatially varying quantities u', t_{cool} . We choose the peak value of u' ; later, we shall see in §2.5.3 that this is insensitive to cooling. For t_{cool} , we use the cooling time of mixed intermediate $T = 2 \times 10^5$ K gas, which is $t_{\text{cool,mix}} = 10$ Myr for the cooling time in the fiducial simulation, and adjust

⁵In practice, we only vary t_{cool} when testing the scaling $Q \propto (L/t_{\text{cool}})^n$. Since the cooling length is the only scale in the problem (ideal hydrodynamics is scale free), varying L and t_{cool} at fixed L/t_{cool} are equivalent. We have checked this previously for cloud-crushing setups.

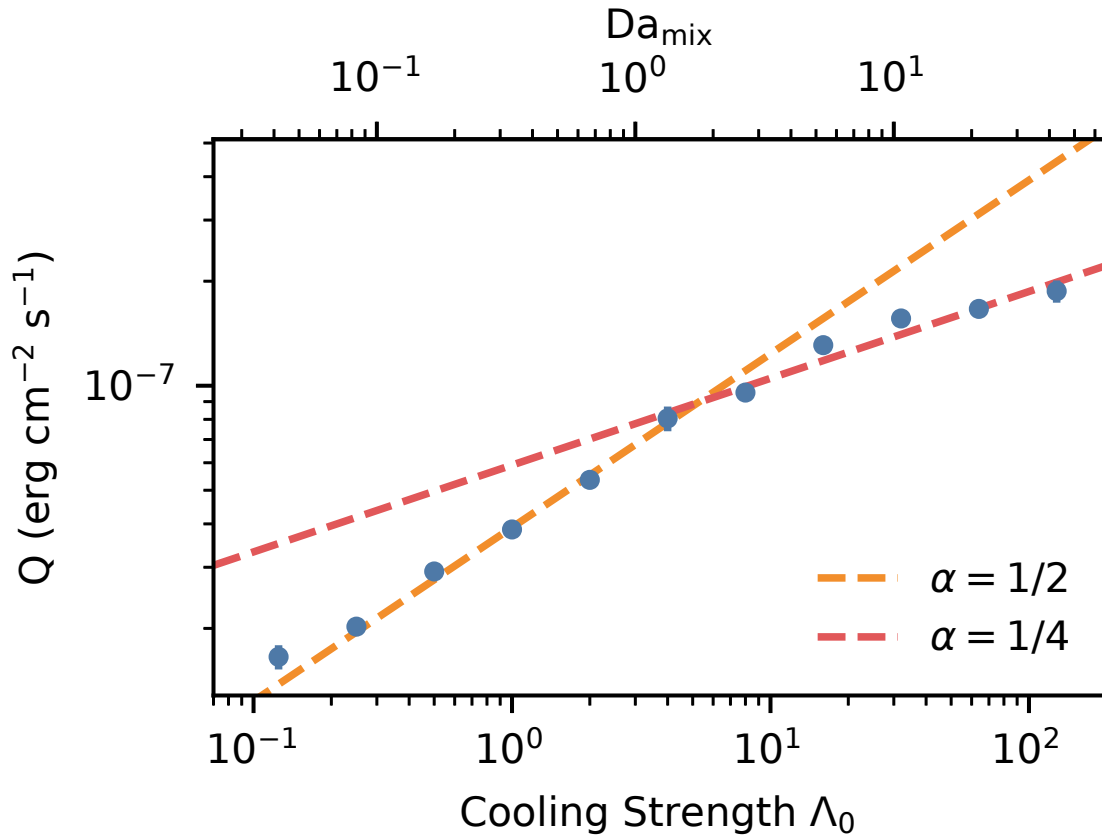


Figure 2.8: Surface brightness Q as a function of cooling strength. We see that Q scales as $t_{\text{cool}}^{-1/2}$ for weak cooling and $t_{\text{cool}}^{-1/4}$ for strong cooling. These two regimes can be characterized by Da_{mix} , which we show in the top axis for reference.




Figure 2.9: Surface brightness Q plotted against turbulent velocities for the two cooling regimes. Expected scalings are given by the orange dashed lines.




Figure 2.10: $Q/u^{3/4}$ is independent of both the shear velocity (upper panel) and the overdensity (lower panel) as expected from Eq. (2.28). Note that to change the overdensity we varied both the cold and hot gas temperatures (shown in blue and orange, respectively).

it accordingly in other simulations. We denote the resulting characteristic Damköhler number as Da_{mix} , which is shown in the top of Fig. 2.8 as a secondary axis. The turnover between the two scalings thus occurs where $\text{Da}_{\text{mix}} \sim 1$. When Da_{mix} is small, we are in the weak cooling regime and conversely, when Da_{mix} is large, we are in the strong cooling regime.

Dependence on Turbulence. Figure 2.9 shows the surface brightness Q as a function of the measured peak turbulent velocity u' , in the weak and strong cooling regimes respectively. The turbulent velocity was varied by changing the shear velocity ($u' \propto v_{\text{shear}}^{0.8}$; see §2.5.3). As given by Eqs. (2.18) and (2.28), $Q \propto u'^{1/2}$ and $Q \propto u'^{3/4}$ in the weak and strong cooling regimes, respectively. Interestingly, these relationships stay the same even when the flow is supersonic with respect to the hot gas.

No Hidden Parameters. As discussed in §2.4, radiative mixing layers are characterized by the dimensionless parameters $\text{Da} = \tau_{\text{turb}}/t_{\text{cool}}$, \mathcal{M} , χ . Above, we tested the dependence on Da via the dependence on τ_{turb} and t_{cool} . By contrast, our theoretical predictions for Q (Eqs. (2.18) and (2.28)) do not contain any explicit dependence⁶ on \mathcal{M} , χ . We confirm this by plotting Q divided by our fiducial scalings against χ and \mathcal{M} in Fig. 2.10. For the simulations that vary the overdensity, we assumed a flat cooling curve and adjusted Λ_0 to keep t_{cool} of the cold gas constant throughout. We see that our fiducial scalings are accurate, with no additional dependence on the parameters χ and \mathcal{M} across a wide dynamic range. While Fig. 2.10 shows the strong cooling case ($\Lambda_0 = 64$), we also check that this holds for the scaling in the fiducial regime.

In summary, our fiducial formula for Q is

$$Q = Q_0 \left(\frac{P}{160 k_{\text{B}} \text{cm}^{-3} \text{K}} \right) \left(\frac{u'}{30 \text{ km/s}} \right)^{1/2} \left(\frac{L}{100 \text{ pc}} \right)^{1/2} \left(\frac{t_{\text{cool, min}}}{0.03 \text{ Myr}} \right)^{-1/2} \quad (2.35)$$

⁶ Q does, however, have *implicit* dependence on \mathcal{M} , χ since $u' = u'(\mathcal{M}, \chi)$. See §2.5.3

in the slow cooling regime and

$$Q = Q_0 \left(\frac{P}{160 k_B \text{cm}^{-3} \text{K}} \right) \left(\frac{u'}{30 \text{ km/s}} \right)^{3/4} \left(\frac{L}{100 \text{ pc}} \right)^{1/4} \left(\frac{t_{\text{cool, min}}}{0.03 \text{ Myr}} \right)^{-1/4} \quad (2.36)$$

in the fast cooling regime, where t_{cool} is evaluated at the peak of the cooling curve and the scalings are normalized by Q at the turnover point in Fig. 2.8:

$$Q_0 \sim 8.8 \times 10^{-8} \text{ erg cm}^{-2} \text{ s}^{-1}. \quad (2.37)$$

Scaling relations for Effective Emissivity

The above simulations confirm our scalings for Q . However, we would like to test the theoretical ideas behind them. The formula in the single-phase regime (Eq. (2.18)) is a straightforward application of mixing length theory, entirely analogous to the thermal conduction case (§2.3). However, the formula for the multiphase regime, (Eq. (2.28)) is much less well-established. Its central claim is that the eddy lifetime is now the geometric mean of the eddy turnover time and the cooling lifetime (Eq. (2.25)). As discussed in §2.4.4, this can be tested by checking that the effective emissivity of the multiphase medium scales as $\tilde{\epsilon} \propto u'^{1/2} t_{\text{cool}}^{-1/2}$ (Eq. (2.29)), and that the width of the multiphase region scales as $h \propto \text{Da}^{-1/4} \propto u'^{1/4} t_{\text{cool}}^{1/4}$. Note that these two quantities are related by $Q \propto \tilde{\epsilon} h$, so only one of them constitutes an independent test. In Fig. 2.11, we show $\tilde{\epsilon}$ at the spatial location where it is maximized, as a function of u' and cooling strength $\Lambda \propto t_{\text{cool}}^{-1}$. It clearly conforms to the expected scalings. We also confirm that the FWHM of the multiphase region agrees with the predicted scalings, though this is a less well-defined and noisier quantity. In §2.5.4, we shall see that the predicted effective emissivity allows remarkably accurate predictions of mean temperature profiles.

In summary, our fiducial formula for effective emissivity in the strong cooling regime




Figure 2.11: The effective emissivity for various shear velocities with $\Lambda_0 = 64$ and for different cooling strengths. Expected scalings are given by the red dashed lines.

is the Gaussian model

$$\tilde{\epsilon} = c_p \frac{P}{\tilde{\tau}_{\text{cool}}} \mathcal{N}(0, \sigma^2), \quad (2.38)$$

where from Eq. (2.25):

$$\tilde{\tau}_{\text{cool}} = 2.5 \text{ Myr} \left(\frac{L}{100 \text{ pc}} \right)^{1/2} \left(\frac{u'}{30 \text{ km/s}} \right)^{-1/2} \left(\frac{t_{\text{cool, min}}}{0.03 \text{ Myr}} \right)^{1/2}, \quad (2.39)$$

where u' is evaluated at its peak (and follows the scaling relations obtained in §2.5.3), and t_{cool} is evaluated at the peak of the cooling curve. This is the form plotted as the red dashed lines in Fig. 2.11; note that it is a derived quantity with no free parameters. From Eqs. (2.36), (2.38) and (2.39), we obtain the width of the cooling region as

$$\sigma \sim 16 \text{ pc} \left(\frac{L}{100 \text{ pc}} \right)^{3/4} \left(\frac{u'}{30 \text{ km/s}} \right)^{1/4} \left(\frac{t_{\text{cool, min}}}{0.03 \text{ Myr}} \right)^{1/4}. \quad (2.40)$$

In the weak cooling (homogeneous reactor) regime, the emissivity is unchanged from the standard $\epsilon = n^2 \Lambda(T)$ form.

Scaling Relations for u'

We now consider how the turbulent velocity u' depends on other parameters in our simulation, specifically cooling, the overdensity χ and shear Mach number \mathcal{M} . We focus on the strong cooling regime, since that is of the most astrophysical interest (e.g., for cloud survival, Gronke & Oh 2020a) and less well-understood. We stress that these scalings are particular to our setup and not as general as the scalings for Q . They will differ depending on the source of turbulent driving. Thus, we do not invest the same effort in deriving and understanding them.

It is still useful to note some theoretical considerations. The free energy for driving

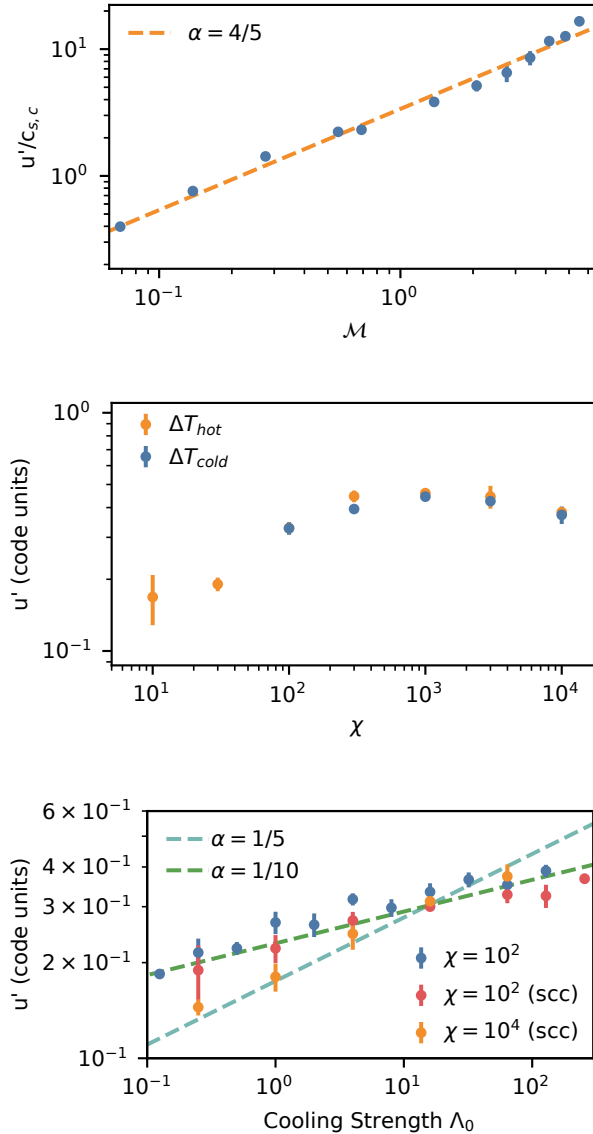


Figure 2.12: Dependence of turbulent velocity on shear velocity (using $\chi \sim 100$), overdensity and cooling. Overdensity and shear velocity are varied in the strong cooling regime ($\Lambda_0 = 64$). Scalings are represented by dashed lines. The middle panel and bottom panel includes simulations with a simplified cooling curve (see text).

turbulence in our mixing layer is shear, where:

$$u'^2 \approx l^2 \left(\frac{\partial v_y}{\partial x} \right)^2, \quad (2.41)$$

with l being the characteristic size of vortices. This is simply the statement that the vorticity of the eddies and the shear flow (from which the eddies derive their vorticity) are comparable. Since we volume average within slices through the mixing layer when measuring u' , the contribution is dominated by the turbulent velocity of the hot gas in the region where turbulence u' and emissivity peak ($f_{\text{hot}} \sim 0.5$). In our simulations, we indeed see $u' \propto \nabla v_y$ (Fig. 2.6). Cooling can play an important role in regulating the width of the mixing layer, and hence ∇v and u' (see Fig. 2.6). However, once deep in the multiphase regime, cooling only has a weak effect. We can see this from $\nabla v \sim v_{\text{shear}}/h$; if we combine our prediction for h , Eq. (2.30), and Eq. (2.41), we obtain:

$$u' \propto v_{\text{shear}}^{4/5} \left(\frac{L}{t_{\text{cool}}} \right)^{1/5}. \quad (2.42)$$

Note that for a given shear velocity and cooling time, there is no explicit overdensity dependence. In the multi-phase $\text{Da} \gg 1$ regime, t_{cool} is the cooling time of the cold gas.

Figure 2.12 show the scalings of u' with v_{shear} , χ , and Λ_0 . The upper panel includes the scalings from previously mentioned simulations that vary the shear velocity from the fiducial setup. We find, as in Eq. (2.42) that $u' \propto v_{\text{shear}}^{4/5}$.

Changing the overdensity involves a change in the temperature regime where cooling takes place. We want to do so while keeping t_{cool} constant. We adopt a simplified form of the cooling curve $\Lambda = \text{const}$ for $T < 10^4 \text{K}$ and $\Lambda \propto T^4$ for $T > 10^4 \text{K}$, such that $t_{\text{cool}} \propto T^{-2}$ and $t_{\text{cool}} \propto T^2$ below and above $T = 10^4 \text{K}$ respectively. This singles out the minimum in the cooling time to occur at $T \sim 10^4 \text{K}$ (as is true of more realistic cooling

curves). This cooling time is held constant. To vary χ , we keep either the hot or cold phase constant and change the temperature and density of the other phase.

In the middle panel of Fig. 2.12, we see that for $\chi \gtrsim 100$, u' behaves according to Eq. (2.42) with no dependence on χ , regardless of whether the hot or cold phase is being varied. On the other hand, at lower overdensities, u' does not scale as expected: it declines toward low χ instead. In this low overdensity regime, the temperature and velocity profiles decouple and no longer track one another. This was first noted by Fielding et al. (2020) (see their Figure 1), but we draw a slightly different conclusion from them: the decoupling of thermal and momentum profiles is not general but only happens at low χ . The reason is that the cooling time of mixed gas ($t_{\text{cool,mix}} \sim \chi t_{\text{cool,cold}}$) is still relatively short, where $T_{\text{mix}} \sim (T_{\text{hot}} T_{\text{cold}})^{1/2}$. The hot gas then rapidly cools after a small amount of mixing with cold gas. Radiative cooling outpaces momentum transport, which mostly takes place when the gas is already cold; the velocity shear and turbulence peak in the single phase regime. This vitiates the assumptions behind Eq. (2.42). Indeed, the assumption of a thin mixing layer is no longer valid. For the ratio of the thickness of the shear layer h and the the horizontal length L_y we can write to first order

$$\frac{h}{L_y} \sim \frac{v_{\text{in}}}{v_{\text{shear}}} \sim \frac{c_{\text{s,cold}}}{\mathcal{M} c_{\text{s,hot}}} \sim \frac{1}{\mathcal{M} \sqrt{\chi}}, \quad (2.43)$$

where the first equality comes from the continuity equation. Hence, the flow decelerates on a length scale comparable to the thickness of the mixing layer as the simulation proceeds. Overall, this regime holds less astrophysical significance: because of the location of the stable phases in the cooling curve, most situations of astrophysical interest involve density contrasts $\chi \gtrsim 100$, where Eq. (2.42) holds.

In the bottom panel of Fig. 2.12, we check the dependence of u' on t_{cool} . The $\chi = 10^4$ simulation follows the expected $u' \propto \Lambda^{0.2}$ scaling. However, the $\chi = 100$ simulation




Figure 2.13: A comparison between simulation profiles and a 1D mixing length model at two different cooling strengths shows good agreement.

follows a slightly weaker $u' \propto \Lambda^{0.1}$ scaling. This remains true for simulations which use the full (realistic) cooling curve (blue points). This is because the turbulent velocity approaches equation (2.42) asymptotically as χ increases. For instance, it only becomes fully independent of overdensity for $\chi \gtrsim 300$ (middle panel). In any case, the difference is small.

2.5.4 Comparing Simulations to 1D Mixing Length Models

In §2.3, we constructed time-steady 1D models of thermal interfaces, where thermal conduction balances radiative cooling. In the single-phase, weak cooling case, one can construct similar profiles by substituting turbulent heat diffusion for thermal conduction (as has been done for galaxy clusters, Kim & Narayan 2003; Dennis & Chandran 2005).




Figure 2.14: Measurement of turbulent diffusion through the mixing layer in a simulation without radiative cooling. The upper panel shows the corresponding average temperature. In the lower panel, the blue line shows κ_{turb} measured from the simulation via Eq. (2.48), while the orange line shows the mixing length approximation for the fiducial setup (cf. §2.5.4). The dashed line shows the value used for the profiles shown in Fig. 2.13

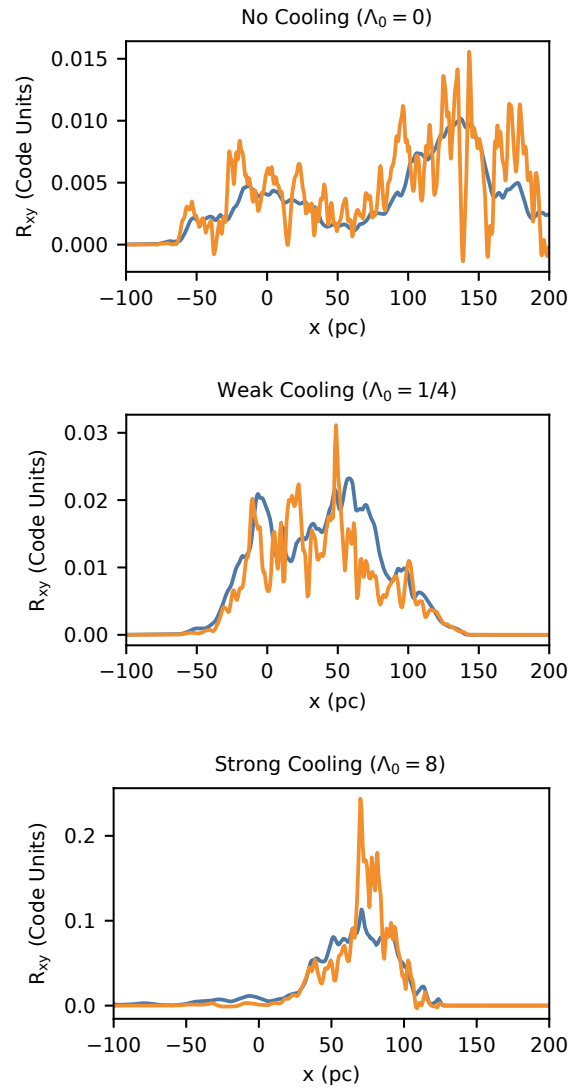


Figure 2.15: For adiabatic, weak and strong cooling, the Reynolds stress is shown in blue along with the mixing length model in orange. As in the case of turbulent heat transport, we find good agreement between them.

However, it may seem absurd to carry this over to the strong cooling regime, where the highly fluctuating, fractal and multiphase structure of the radiative front seems to preclude a simple mean-field approach. Here, we show that with judicious choice of the effective emissivity, such a model matches simulations surprisingly well.

We first establish that in an isobaric medium, we can write the turbulent heat flux in a form similar to that for thermal conduction, $F_{\text{turb}} = -\kappa_{\text{turb}}\nabla T$. In mixing length theory, the turbulent heat flux is proportional to the gradient of specific entropy (e.g., Dennis & Chandran 2005):

$$F_{\text{turb}} = -D_{\text{eddy}}\rho T\nabla s, \quad (2.44)$$

where $s = c_V \ln(p/\rho^\gamma)$ is the specific entropy, $\gamma = c_P/c_V = 5/3$ is the ratio of specific heats, $c_V = 3/2k_B/\bar{m}$ is the specific heat at constant volume, and D_{eddy} is the eddy diffusivity, with units $[D_{\text{eddy}}] = L^2T^{-1}$. However, under isobaric conditions $\nabla P = 0$, evaluating the above expression gives:

$$F_{\text{turb}} = -D_{\text{eddy}}\rho c_P \nabla T, \quad (2.45)$$

i.e., the turbulent heat flux is proportional to the temperature gradient. In an isobaric medium, one can equally well think of passive scalar advection of entropy or temperature⁷. For simplicity, and analogous to thermal conduction, we will consider $F_{\text{turb}} = -\kappa_{\text{turb}}\nabla T$, where:

$$\kappa_{\text{turb}} = D_{\text{eddy}}\rho c_P. \quad (2.46)$$

We then assume that the coefficient κ_{turb} is a constant independent of temperature. We can show after the fact that this is a reasonable assumption. As in §2.3, we then

⁷Of course, only entropy advection is correct in a stratified medium like a star or galaxy cluster, where mixing length theory is usually applied.

solve the 1D steady-state equation:

$$\kappa_{\text{turb}} \frac{d^2 T}{dx^2} = j_x c_p \frac{dT}{dx} + \rho \mathcal{L}. \quad (2.47)$$

In §2.3, the thermal conductivity κ was known and we solved for the mass flux $j_x = \rho v$ as an eigenvalue. Here, since the medium is multiphase, the emissivity is not the same as that of single-phase medium with the same mean temperature. Motivated by the scalings in §2.5.3, we model the emissivity as a Gaussian as specified in Eq. (2.38). This sets Q and hence the mass flux $j_x = Q/c_p(T_{\text{hot}} - T_{\text{cold}})$. Because our cooling is now a function of position and not temperature, we specify the value of κ_{turb} and solve for the profile via the shooting method, subject to the same boundary conditions as before (Eq. (2.12)). The results for two strong cooling ($\Lambda_0 = 4, 64$) cases are shown in Fig. 2.13, where we have used $\kappa_{\text{turb}} = 5 \times 10^6 \text{ erg cm}^{-1} \text{ s}^{-1} \text{ K}^{-1}$ for both cases. This is justified below. Such a simple mixing length model provides a remarkably good fit to the profiles seen in our simulations. This allows the construction of rapid semi-analytic models of radiative mixing layers, which is very useful when comparing against observations (e.g., matching line column density ratios) when the underlying model parameters such as u', L, t_{cool} are unknown and one has to search over parameter space. We can thus construct models for multiphase mixing layers with the same speed and ease as for thermal conduction.

We estimate the turbulent diffusion coefficient by applying mixing length theory to direct measurements from the simulation. The simplest way to do this is to first obtain κ_{turb} from adiabatic simulations where we can model the turbulence using the Reynold's Averaged Navier-Stokes (RANS) equations. This approach separates the flow into two components: a mean flow and a time-dependent varying flow. By representing initial variables such as temperature and velocity as $\phi = \bar{\phi} + \phi''$ where $\bar{\phi}$ is the density weighted mean value, we find an extra term of the form $\langle \rho v_x'' T'' \rangle$ which is modelled using a simple

gradient relation:

$$F_{\text{turb}} = \langle c_p \rho v_x'' T'' \rangle = -\kappa_{\text{turb}} \frac{d\bar{T}}{dx}. \quad (2.48)$$

Figure 2.14 shows the resulting measurement of κ_{turb} as a function of height in the mixing layer for a single time snapshot where the averages were taken along the y - z plane. Consistent with our assumptions, κ_{turb} is roughly constant, and the dashed line shows the value we adopt in the simple model above. The solid orange line in Fig. 2.14 shows the mixing length approximation:

$$\kappa_{\text{turb}} = \bar{\rho} c_p l u', \quad (2.49)$$

with a mixing length of $l = 4$ pc, which fits the result of Eq. (2.48) from the simulation remarkably well.

In mixing length theory, the mixing length l cannot be obtained from first principles, but must be calibrated from experiments or simulations. Nonetheless, the value we obtain is reasonable from order of magnitude considerations. Since $u' \approx l \nabla v_y \sim (l/L) v_{\text{shear}}$, we have:

$$l \sim \frac{u'}{v_{\text{shear}}} L \sim \frac{c_{\text{s,cold}}}{c_{\text{s,hot}}} L \sim \frac{L}{\sqrt{\chi}}, \quad (2.50)$$

which gives $l \sim 10$ pc for our setup.

Due to the multiphase structure of the mixing layers with strong cooling, it is not possible to use Eq. (2.48) to measure κ_{turb} there. It is interesting that κ_{turb} derived from adiabatic simulations provides a good fit when used to solve for temperature profiles in strong cooling simulations, and is consistent with the finding that cooling appears to have little effect on turbulence.

Instead, in cooling simulations we can focus on velocity profiles to verify the mixing length approach. In Fig. 2.15, we plot the Reynolds stress in adiabatic, weak, and strong

cooling simulations, and compare with the expectation from mixing length theory that

$$R_{xy} = -\langle \rho v_x'' v_y'' \rangle = \bar{\rho} \nu_T \nabla v_y, \quad (2.51)$$

where the turbulent viscosity $\nu_T = u_x' l$. The orange line shows the mixing length estimate from the right side of Equation (2.51) and is again a remarkably good fit, with a mixing length $l \sim 4$ pc throughout all simulations. Since the mixing length ansatz for Reynolds stress is a good approximation, we can also construct mean shear v_y and turbulent velocity profiles u' analytically as well, though we eschew this for the sake of brevity. This suggests that the turbulent Prandtl number (ν_T/D_{eddy}) is of order unity as typical in turbulent flows (Tennekes & Lumley 1972).

2.5.5 Thermal Conduction

We perform a quick assessment of the impact of isotropic thermal conduction. We defer anisotropic field-aligned conduction to future work. We use the same (constant, temperature independent) thermal conductivity as the 1D simulations, given in Eq. (2.4), which we vary in amplitude κ_0 . Note that for our fiducial case, $\kappa_{\text{turb}} \sim \kappa_{\text{cond}}$. The results are shown in Fig. 2.16. Conduction has no impact until $\kappa_{\text{cond}} > \kappa_{\text{turb}}$. At this point, Q falls back towards the laminar speed S_L , but S_L at this transition is already close to S_T , indicating that thermal conduction is strong enough to compete with turbulent diffusion as the main source of heat transport. Without thermal conduction, we assumed that the turbulent velocity u' was much larger than S_L , but this assumption breaks down for strong thermal conduction since that increases S_L . The scale at which $u' \sim S_L$ is known as the Gibson scale. In turbulent combustion, below this scale, the flames burn through the turbulent eddies within an eddy lifetime and hence the speed of the front is unaffected by the turbulence. The front is thus said to be ‘wrinkled’ by the turbulence,




Figure 2.16: Dependence of cooling in the mixing layer on thermal conductivity. The curves show runs where the fiducial κ is multiplied by a constant scaling factor κ_0 . Thermal conduction does not matter until it gets large enough to suppress the turbulence.

but not ‘corrugated’ due to the turbulent eddies. This is also known as the ‘wrinkled flames regime’ in the Borghi diagram. Conduction also suppresses the multiphase nature of the mixing layer by broadening the hot/cold gas interface. Our results are consistent with previous cloud-crushing studies which show that thermal conduction hinders hydrodynamic instability but otherwise has fairly mild effects for clouds large enough to resist thermal evaporation, in conditions typical of the CGM (Brüggen & Scannapieco 2016; Armillotta et al. 2017; Li et al. 2020). Overall, as long as turbulent diffusion dominates heat transport, conduction can be safely ignored.

2.5.6 Convergence

The convergence properties of this setup have been previously studied (Ji et al. 2019; Fielding et al. 2020). We therefore perform a restricted set of resolution studies for our 3D simulations, considering the fiducial case ($\Lambda_0 = 1$) and a case with strong cooling ($\Lambda_0 = 128$). The results are shown in Fig. 2.17. The resolutions go from a quarter to twice that of the fiducial resolution. Lower resolution runs are also shown for the run with strong cooling. Error bars are derived from fluctuations in Q in the time series. In the fiducial case, we find that we are well converged, with little difference in the mean Q and small error bars, indicating that the simulation is well resolved. However, for the case of strong cooling, oscillations are significantly large for lower resolutions, consistent with the 1D case, resulting in larger error bars. The cooling over time is shown in the middle and bottom panels of Fig. 2.17 for both cases, where we can see more clearly that in the case with fiducial cooling, the curves are generally smooth with no rapid oscillations. In the case with strong cooling, we see that as we lower the resolution, we see rapid oscillations with increasing amplitude. These oscillations have a period of roughly t_{cool} . In 3D, the oscillations cause apparent broadening of the interface, and

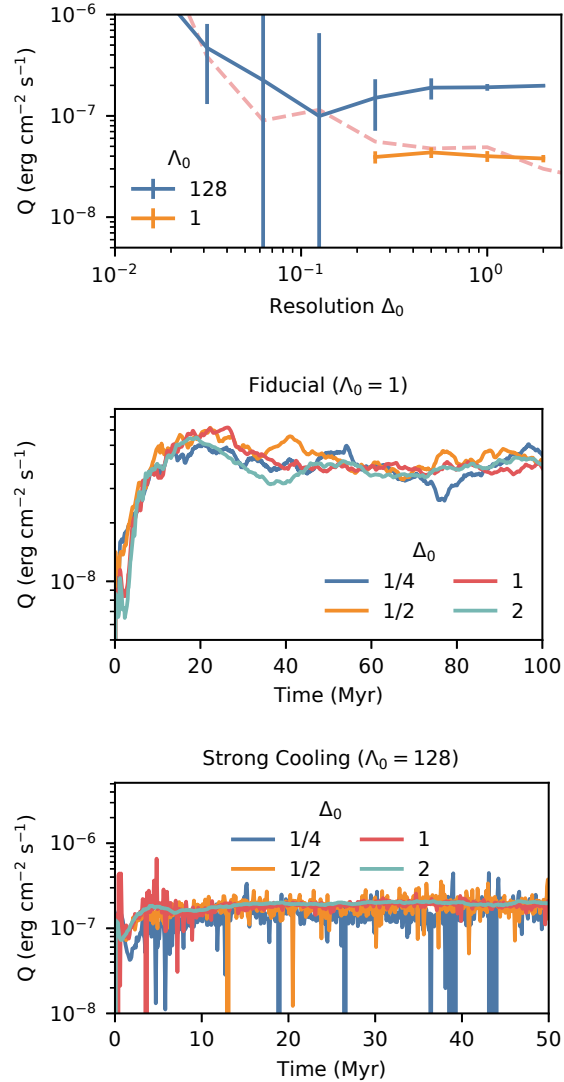


Figure 2.17: *Upper panel:* Result of varying the resolution of the 3D runs with the fiducial and strongest cooling rates. The fiducial resolution $\Delta_0 = 1$ is converged in both cases. The pink dashed line shows the result for the 1D simulations with strong cooling and no conduction for comparison (orange line in Fig. 2.3), reflecting where numerical diffusion becomes dominant. *Middle and Lower panels:* The time profiles of cooling for the various points shown in the top panel. The middle panel shows runs from the fiducial runs ($\Lambda_0 = 1$), while the bottom panel shows runs with the strongest cooling ($\Lambda_0 = 128$). The runs with strong cooling shows rapid oscillations that grow in amplitude as resolution is lowered.

the cooling surface appears to adjust with resolution. However, the mean Q still remain close to the converged value. As long as Q is averaged over a sufficient time interval, our fiducial resolution is sufficient, even for our strongest cooling case. This is consistent with previous results of larger scale simulations, and follows the expectations from the 1D results in §2.3.

2.5.7 Comparison with Previous Work

We now compare our results to recent work on radiative TMLs. We confine our comparisons to the formula for hot gas entrainment (Eq. (2.28)) and its physical justification.

Ji et al. (2019) was the first paper to confront analytic models of radiative TMLs with simulations. They pointed out that the inflow and turbulent velocities were much less than the shear velocity, that radiative cooling was balanced by enthalpy flux from the hot gas (rather than turbulent dissipation, as in e.g., White et al. 2016), and that contrary to the widely cited model of Begelman & Fabian 1990, the layer width does not scale as $h \propto v_t t_{\text{cool}}$. They obtained a scaling $h \propto t_{\text{cool}}^{1/2}$, $v_{\text{in}} \propto t_{\text{cool}}^{-1/2}$, which in hindsight is the scaling in the weak cooling regime; they did not run enough simulations to discern the change in slope to $v \propto t_{\text{cool}}^{-1/4}$ in the strong cooling regime. Interestingly, they noted that using standard emissivities, mixing length theory matches temperature/density profiles well in the weak, but not strong cooling regimes. We now know this is because emissivity changes in the multiphase regime (Eq. (2.38)).

Gronke & Oh (2020a) obtained a scaling relation (Eq. (2.1)) which is identical to our fiducial Eq. (2.28) in the strong cooling regime if $u' \sim c_{\text{s,cold}}$. They consider a cold cloud embedded in a hot wind, which grows in mass and entrains. It continues to grow even after it is fully entrained ($v_{\text{shear}} \rightarrow 0$); in fact, growth *peaks* at this point. The cloud pulsates due to loss of pressure balance seeded by radiative cooling; this in turn drives

turbulence and hot gas entrainment.

Both of these studies considered magnetic fields, which are ignored here. For instance, the plane parallel shear simulations of Ji et al. (2019) show that B-fields suppress turbulence and mixing due to magnetic tension, but the MHD cloud simulations of Gronke & Oh (2020a) nonetheless show cloud growth at the same rate as hydrodynamic simulations, despite very different cloud morphology in the two cases. The difference likely lies in the very different nature of turbulent driving in the shear flow and cloud scenarios, which also affects the growth in surface area. Given the substantial differences between hydrodynamic and MHD turbulence, it is important to eventually extend the arguments in this paper to the MHD case.

Both Ji et al. (2019) and Gronke & Oh (2020a) invoked low pressure due to fast cooling to seed turbulence and set the entrainment rate of the hot gas, rather than the Kelvin Helmholtz instability. In Ji et al. (2019), this was argued to be due to the constancy of $P + \rho u'^2$ across the mixing layer (so that pressure drops due to cooling boost turbulence), as well as the fact that v_{in} appeared only weakly dependent on v_{shear} and independent of χ , two factors which set the Kelvin Helmholtz timescale. In light of our larger suite of simulations, it is now clear that in fact entrainment rates do depend on v_{shear} . Turbulence seeded by cooling is also inconsistent with the very weak dependence of u' with cooling strength that we see here (Fig. 2.12). For this problem, it is important to have sufficient dynamic range and dense sampling to establish scaling relations (as we have seen in the u' vs χ relation; Fig. 2.12). In this paper we argue – consistent with results from the combustion literature, and as argued by Fielding et al. (2020) – that turbulence, rather than pressure gradients, is the primary driver of hot gas entrainment. This statement has to be qualified by the fact that in the cloud case, cooling-induced pressure gradients appear to be the primary driver of cloud pulsations and turbulence, so the end result can be the same. Thus, the Gronke & Oh (2020a) scaling for v_{in} still

holds, as potentially do their analytic arguments⁸. However, we await detailed study of u' scalings in this scenario to refine the model.

Fielding et al. (2020) ran simulations of radiative plane parallel mixing layers very similar to Ji et al. (2019) and this work. They rightly emphasize the fundamental role of turbulence in hot gas entrainment, and directly measure fractal properties in their simulations. They derive an analytic model whose scalings are similar to Gronke & Oh (2020a) and this work.

The analytic model of Fielding et al. (2020) states that:

$$v_{\text{in}} = \frac{w}{t_{\text{cool}}} \left(\frac{A_w}{A_L} \right) = \frac{w}{t_{\text{cool}}} \left(\frac{w}{L} \right)^{-1/2}, \quad (2.52)$$

where w is a length scale defined by $t_{\text{mix}} \sim w/v_{\text{turb,hot}}(w) \sim t_{\text{cool}}$ and the second equality arises from fractal scalings with fractal dimension $D = 2.5$, which they measure directly from their simulations. The first equality is very similar to Eq. (2.19), except that w/t_{cool} is substituted for S_L . However, at face value, this argument would seem to imply that if the length scale w is not resolved (and replaced by a resolution element Δ), then inflow becomes resolution dependent, $v_{\text{in}} \propto \Delta^{1/2}$. Neither Fielding et al. (2020) nor we see evidence for this, even in simulations where w is highly under-resolved.

2.5.8 Conclusions

Radiative mixing layers are closely analogous to turbulent combustion fronts: in both cases, the speed of front propagation v_{in} is determined by the temperature and density sensitive reaction rate, and thus conditions within the front itself. This is in contrast to shock propagation, where the shock speed and jump conditions are simply determined by

⁸E.g., they identify the timescale $t_{\text{sc,cold}} \sim H/c_{\text{s,cold}}$, where $H \sim (r_{\text{cl}}c_{\text{s}}t_{\text{cool}})^{1/2}$ is analogous to the length scale in Eq. (2.24). This is identical to the effective cooling time (Eq. (2.25)) which is critical to the model in this paper.

conservation laws, independent of the small-scale details of shock structure. To obtain v_{in} , it would seem that the structure of the front must be accurately resolved. Thus, it has long been thought that calculations of thermal fronts can only be converged if thermal conduction is included and the Field length is resolved (Koyama & Inutsuka 2004). Yet, recent simulations (Ji et al. 2019; Gronke & Oh 2018, 2020a; Mandelker et al. 2020; Fielding et al. 2020) show remarkable robustness to resolution, despite the absence of thermal conduction – even when the cooling front is one cell thick! They also show characteristic front propagation speeds of order the cold gas sound speed $v_{\text{in}} \sim c_{\text{s,cold}}$ (far less than the maximum possible $c_{\text{s,hot}}$) and scalings $v_{\text{in}} \propto (t_{\text{sc}}/t_{\text{cool}})^{-1/4}$ (where t_{sc} is a sound-crossing time) which are not trivial to understand. In this paper, we use models derived from the turbulent combustion literature to shed light on these issues.

We first examine the impact of resolution on laminar fronts. The restriction to laminar fronts allows the problem to be considered in 1D, where there are analytic solutions. In the absence of thermal conduction, there is clear resolution dependence, such that $v_{\text{in}} \sim \sqrt{D_{\text{num}} t_{\text{cool}}} \propto \Delta^{1/2}$. The numerical diffusion coefficient from truncation error is $D_{\text{num}} \sim v\Delta$, v is a characteristic velocity, and Δ is the grid resolution. The inclusion of conduction is indeed required for convergence. However, surprisingly it is not strictly necessary to resolve the Field length for convergence. Instead, the key requirement for convergence is that explicit thermal diffusion simply be larger than numerical diffusion: i.e., $D_{\text{cond}} > D_{\text{num}}$, where $D_{\text{cond}} \sim \kappa/(\rho c_{\text{P}})$, where κ is the standard conduction coefficient and c_{P} is the specific heat at constant pressure. If the Field length is unresolved, numerical dispersion increases, as the front structure is not accurately resolved and there are larger errors in the temperature derivatives and conductive heat flux. Nonetheless, the steady-state simulations oscillate about the correct answer. The error can be beaten down by time averaging. This is not unusual for a stiff problem where the smallest length scale remains unresolved.

We then examine the effects of turbulence. As in Ji et al. (2019), we simulate a plane-parallel shear layer where the Kelvin-Helmholtz instability drives turbulence and mixing. We find, consistent with previous findings, that the inflow velocity v_{in} and surface brightness Q are converged without thermal conduction. Heuristically, we argue that this is because as long as the turbulent driving scale L is well resolved $L \gg \Delta$, the turbulent diffusivity $D_{\text{turb}} \sim u'L$ is always larger than the numerical diffusivity $D_{\text{num}} \sim v\Delta$. Similar to our 1D results, lower resolution simply implies larger numerical dispersion and temporal oscillations in the front profile. We also find that thermal conduction has little effect unless it is larger than the turbulent diffusivity.

The front is characterized by the dimensionless parameters: overdensity χ , Mach number \mathcal{M} , and most importantly the Damköhler number $\text{Da} = \tau_{\text{turb}}/t_{\text{cool}} = L/(u't_{\text{cool}})$, where u' is turbulent velocity at the outer scale L , and t_{cool} is the local cooling time. The Damköhler number, which measures the relative importance of mixing and cooling, increases as temperature falls within the front. There are two distinct regimes:

- *Weak cooling* ($\text{Da} \gg 1$): the ‘well stirred’ regime. Since the cooling time is longer than the eddy turnover time, the gas entropy is set primarily by mixing. Thus, it forms a single phase gas with smoothly varying temperature within the front. The front structure is entirely analogous to a thermal conduction front, except that the conductive diffusivity D_{cond} is replaced by the turbulent diffusivity D_{turb} . This implies a front thickness $h \sim (D_{\text{turb}}t_{\text{cool}})^{1/2} \propto t_{\text{cool}}^{1/2}$ and an inflow velocity $v_{\text{in}} \sim (D_{\text{turb}}/t_{\text{cool}})^{1/2} \propto t_{\text{cool}}^{-1/2}$.
- *Strong cooling* ($\text{Da} \ll 1$): the ‘corrugated flamelet’ regime. In this limit, the cooling time is shorter than the mixing time, the gas entropy is set primarily by cooling, and the gas fragments into a multiphase medium. The interface between the two phases is highly corrugated, and has been shown to have a fractal structure (Fielding et al.

2020). This increase in surface area of the front boosts the mass flux across the front. The surface area increase is resolution dependent. Nonetheless, hot gas as a whole is consumed at a rate $v_{\text{in}} \sim u'$ (where u' is the turbulent velocity at the outer scale), independent of resolution. The rate limiting step in determining hot gas consumption is the turbulent mixing rate, which proceeds at the outer scale velocity u' . It is similar to how mixing time of a passive scalar (e.g., cream in coffee) is set by the eddy turnover time at the outer scale, independent of the details of molecular diffusivity. Rapid mixing in both cases depends on the large increase in surface area due to turbulence.

Our results are also of importance to the resolution requirements in larger scale simulations; e.g., cosmological simulations which are currently unconverged in the cold gas properties (Faucher-Giguère et al. 2016; Hummels et al. 2019). Ultimately, the physics of radiative TMLs explored here sets the mass and momentum transfer between the hot and the cold phase, and thus, impacts not only the morphology of the multiphase medium but also, for instance, the fuel supply for future star-formation. In this work, we showed that as long as numerical diffusion is not the dominant mixing mechanism, it is sufficient in the presence of turbulence to resolve the outer scale L of the turbulence to obtain a converged solution, and not the width of the laminar front, contrary to conventional wisdom. While in many astrophysical applications L is likely of the order of \sim parsecs (McCourt et al. 2018; Gronke & Oh 2018) and thus challenging to resolve directly in large scale simulations, our findings relax the resolution requirements by up to orders of magnitude.

At a more detailed level, one must still take into account the behavior of small scales in the strong cooling regime. We argue that there is a characteristic effective cooling timescale $\tilde{\tau}_{\text{cool}} \sim \sqrt{(L/u')t_{\text{cool}}}$. This effective cooling time is resolution independent.

The turbulent velocity u' measured in the simulations peaks in the multiphase region where the cold gas fraction is $\sim 50\%$, where cooling also peaks. A similar lifetime for eddies in combustion fronts was given by Gülder (1991), by assuming that the mixing front is corrugated down to the Taylor microscale. Thus, the front propagates at a velocity $v_{\text{in}} \sim (D_{\text{turb}}/\tilde{\tau}_{\text{cool}})^{1/2} \propto u'^{3/4}t_{\text{cool}}^{-1/4}$. Our fiducial scalings are given by Eqs. (2.35) and (2.36) in the slow and fast cooling regimes respectively. The slow cooling result is a straightforward application of 1D mixing length theory, while the fast cooling result agrees well with previous simulation work (Eq. (2.1)) if $u' \sim c_{\text{s,cold}}$ (see below).

We have verified directly in our simulations the scalings $v_{\text{in}}, Q \propto u'^{1/2}t_{\text{cool}}^{-1/2}$ and $v_{\text{in}}, Q \propto u'^{3/4}t_{\text{cool}}^{-1/4}$ in the weak and strong cooling regimes respectively (Figs. 2.8 and 2.9). We also show that there are no hidden parameters; and in particular no dependence on overdensity χ or Mach number. Astrophysically, the strong cooling regime is often of more interest. For instance, for clouds embedded in a hot wind to survive cloud-crushing instabilities, $t_{\text{cool}}(T_{\text{mix}}) < t_{\text{cc}} \sim L/u'$ (Gronke & Oh 2018), where $T_{\text{mix}} \sim (T_{\text{hot}}T_{\text{cold}})^{1/2}$, which implies that most emission is in the strong cooling regime. In this regime, we verified in our simulations that within the front, the cooling rate tracks the cold gas fraction (which tracks the surface area), peaking at $f_{\text{cold}} \sim 50\%$, and that the cooling flux has a Gaussian shape (Fig. 2.6), as expected for the front position if it undergoes a random walk. In addition, we show that the effective emissivity in the multiphase region of the simulations scales as $\tilde{\epsilon} \sim P/\tilde{\tau}_{\text{cool}} \propto \sqrt{u't_{\text{cool}}}$, in agreement with our model. The emissivity $\tilde{\epsilon} \propto u'^{1/2}$ tracks turbulence and hence the shearing rate. The width of the strong cooling region also obeys an analytic scaling relation Eq. (2.30) predicted by theory. If we use a turbulent diffusion coefficient and emissivity $\epsilon, \tilde{\epsilon}$ in the weak (strong) cooling regimes respectively, we can match temperature and density profiles in the simulations with mixing length theory remarkably well. The turbulent velocity follows mixing length scalings $u' \approx l\nabla v_y$, and the Reynolds stress is also accurately represented by mixing length the-

ory (Fig. 2.15). This allows for rapid construction of semi-analytic profiles of radiative mixing layers without recourse to simulations, which is very useful for comparing against observations.

All that remains is to specify the turbulent velocity at the outer scale, u' . Turbulence can arise either from external driving (in this paper, due to the shear flow), or be driven by radiative cooling itself (e.g., clouds with $r > c_s t_{\text{cool}}$ which lose sonic contact with their surroundings and begin to pulsate; Gronke & Oh 2020a,b). In this shear driven case, we have verified that $u' \propto \nabla v_{\text{shear}}$, as predicted by mixing length theory. To order of magnitude, $u' \sim c_{s,\text{cold}}$ for cloud pulsations or transonic $\mathcal{M} \sim 1$ shear flows, but here we find detailed parameter dependences. For the plane-parallel shear flow in these simulations, we find that $u' \propto v_{\text{shear}}^{0.8}$, almost no dependence on cooling time t_{cool} , and a non-monotonic dependence on overdensity (Fig. 2.12). For $\chi \gtrsim 100$, Eq. (2.42) is a reasonable approximation. These scalings will of course depend on the nature of turbulent driving. In the future, we plan to investigate turbulent scalings in pulsating clouds, and the effect of B-fields on these scalings. We stress, however, that Eq. (2.28) is general, independent of the source of turbulent driving.

In summary, the cold gas mass growth rates we find in our 3D simulations agree with our analytic model (§2.4.4) and read:

$$v_{\text{in}} \approx 11.3 \text{ km s}^{-1} \mathcal{M}_{\text{turb}}^{1/2} \left(\frac{L}{100 \text{ pc}} \right)^{1/2} \left(\frac{t_{\text{cool,cold}}}{0.03 \text{ Myr}} \right)^{-1/2} \quad (2.53)$$

for the $\text{Da} < 1$ ‘well stirred’ (slow cooling) regime, and

$$v_{\text{in}} \approx 9.5 \text{ km s}^{-1} \mathcal{M}_{\text{turb}}^{3/4} \left(\frac{L}{100 \text{ pc}} \right)^{1/4} \left(\frac{t_{\text{cool,cold}}}{0.03 \text{ Myr}} \right)^{-1/4} \quad (2.54)$$

for the $\text{Da} > 1$ ‘corrugated flame’ (fast cooling) regime. Here, $\mathcal{M}_{\text{turb}} \equiv u'/c_{s,\text{cold}}$ but

as stated above u' – unlike the v_{in} scalings – depends on the geometry employed. For shearing layers, we find (§2.5.3) that

$$u' \approx 50 \text{ km s}^{-1} \mathcal{M}^{4/5} \left(\frac{c_{\text{s,c}}}{15 \text{ km s}^{-1}} \right)^{4/5} \left(\frac{t_{\text{cool,cold}}}{0.03 \text{ Myr}} \right)^{-0.1}, \quad (2.55)$$

for $\chi \gtrsim 100$ and $\mathcal{M} \equiv v_{\text{shear}}/c_{\text{s,hot}}$ as used throughout.

Of course, at higher levels of precision, details of the interaction between turbulence, diffusion and cooling remain to be explored. Just as there are a plethora of models and computational algorithms for subgrid turbulent scalar transport (often used in simulations of metal mixing), there are a plethora of models for subgrid turbulent combustion (e.g., see Swaminathan & Bray 2011). The issues are more complex since combustion can backreact on the flow and change its properties. Such models have been used in simulations of thermo-nuclear burning on white dwarfs (Schmidt et al. 2006; Jackson et al. 2014), where the burning fronts are unresolved. Such sub-grid models would be an interesting avenue for future work, particularly in the context of cosmological simulations of galaxy formation, where the separation of scales is even more forbidding than in Type Ia supernova problem. Another avenue for more detailed future work is the inclusion of non-equilibrium chemistry. In this work, equilibrium cooling curves were used in all simulations. In reality, material will often be out of equilibrium, with recombination/ionization rates often having time scales longer than the cooling time. Metallicity differences between the different phases could also lead to further complications beyond the analysis in this paper. While non-equilibrium ionization/recombination was taken into account in (Kwak & Shelton 2010; Ji et al. 2019), to our knowledge there have not been any studies which incorporate non-equilibrium cooling. Such details could be important due to their bearing on the predictions of observables such as column densities. It is also important to continue to verify scalings for S_{T} in higher resolution simulations,

perhaps with thermal conduction so that S_L is well-defined and resolution-independent. One important limitation of current simulations is that the front width δ and Kolmogorov scale η are unresolved and simply equal to the grid scale, so that $Ka \sim (\delta/\eta)^2 \sim 1$. In practice, these scales could be sufficiently separated (with $Ka \gg 1$) that the arguments in §2.4.4 no longer apply. If so, turbulence can penetrate the conductive interface and affect conditions there, impacting the effective and total cooling rates. An intriguing approach in the spirit of 1D modeling in this paper, and useful for developing physical insight is the Linear Eddy Model (Kerstein 1988), which attempts to model the effects of turbulence in 1D so that extremely high resolution can be achieved, and has good support from experiments and direct numerical simulation. It has been successfully applied to the Type Ia supernova problem (Woosley et al. 2009). These are promising avenues for future work.

Acknowledgements

We thank Omer Blaes, Drummond Fielding, Yan-Fei Jiang, Chris White, and Nir Mandelker for helpful discussions. We also thank Yan-Fei Jiang for the use of his conduction module for Athena++, and Suoqing Ji for providing his setup in FLASH. We have made extensive use of the yt astrophysics analysis software suite (Turk et al. 2011), matplotlib (Hunter 2007), numpy (Van Der Walt et al. 2011), and scipy (Virtanen et al. 2020) whose communities we thank for continued development and support. We acknowledge support from NASA grants NNX17AK58G, 19-ATP19-0205, HST theory grants HST-AR-15039.003-A, HST-AR-15797.001-A, XSEDE grant TG-AST180036, and the NASA Hubble Fellowship grant HST-HF2-51409.

Chapter 3

A Model for Line Absorption and Emission from Turbulent Mixing Layers

*All the variety, all the charm, all the beauty
of life is made up of light and shadow.*

Leo Tolstoy, Anna Karenina

Turbulent mixing layers (TMLs) are ubiquitous in multiphase gas. They can potentially explain observations of high ions such as O VI, which have significant observed column densities despite short cooling times. Previously, we showed that global mass, momentum and energy transfer between phases mediated by TMLs is not sensitive to details of thermal conduction or numerical resolution. By contrast, we show here that observables such as temperature distributions, column densities and line ratios *are* sensitive to such considerations. We explain the reason for this difference. We develop a prescription for applying a simple 1D conductive-cooling front model which quantitatively

reproduces 3D hydrodynamic simulation results for column densities and line ratios, even when the TML has a complex fractal structure. This enables sub-grid absorption and emission line predictions in large scale simulations. The predicted line ratios are in good agreement with observations, while observed column densities require numerous mixing layers to be pierced along a line of sight.

3.1 Introduction

Observations of ultraviolet absorption lines of high ions like C IV, Si IV, N V and O VI trace intermediate temperature ($\sim 10^5$ K) gas, assuming collisional ionization equilibrium (CIE). They are widely observed in a range of astrophysical contexts, such as our own galactic disk and halo, external galaxies, and high velocity clouds. Observations of a significant column density of these ions is puzzling since the gas they trace should cool quickly. One possibility is that they exist in turbulent mixing layers (TMLs) between cold 10^4 K and hot 10^6 K gas. There, radiative cooling is balanced by enthalpy flux into the TML, reaching a steady state. Such TMLs are relevant to a host of issues such as the stability and survival of AGN jets (Hardee & Stone 1997), cold clouds in a hot wind (Scannapieco & Brüggén 2015; Schneider & Robertson 2017; Gronke & Oh 2018, 2020a), and cold streams inflowing from cosmological accretion (Mandelker et al. 2020). Observational diagnostics of TMLs could be very informative. For instance, the mass entrainment rate per unit area is directly proportional to the bolometric surface brightness (in the absence of radiative heating and scattering).

Models for absorption and emission in conductive-cooling fronts (McKee & Cowie 1977; Borkowski et al. 1990; Gnat et al. 2010) have predicted column densities that are too low, requiring many layers to be pierced along a sightline, and line ratios that do not match observations (Wakker et al. 2012). To our knowledge, the only equivalent work for

TMLs is Slavin et al. (1993), based on an analytic model by Begelman & Fabian (1990). The latter had many important physical insights, but its detailed predictions are not in agreement with 3D hydrodynamic simulations (Ji et al. 2019). However, Slavin et al. (1993) made the important early prediction that TML column densities are also too low. This has been borne out in simulations (Kwak & Shelton 2010; Ji et al. 2019), although Kwak & Shelton (2010) find line ratios in good agreement with observations (Wakker et al. 2012). Subsequent modeling of TMLs (Ji et al. 2019; Fielding et al. 2020; Tan et al. 2021) has focused on hot gas mass entrainment rates, which is crucial for the cold gas survival issues mentioned above. Such work has found that global mass, momentum and energy transfer between phases is not sensitive to details of thermal conduction or numerical resolution (Tan et al. 2021). By contrast, we show that the same is not true of observables such as temperature distributions, column densities and line ratios. We develop a prescription for how we can apply 1D conductive-cooling front models to 3D TMLs to obtain analytic predictions for these quantities, which we then verify with simulations.

3.2 Methods

We carry out our simulations using the publicly available MHD code Athena++ (Stone et al. 2020). For details of our setup and implementation, we refer the reader to Sections 2 and 5.1 of Tan et al. (2021). In brief, we simulate a shear layer between gas of temperatures $T_{\text{cold}} = 10^4$ K and $T_{\text{hot}} = 10^6$ K, with shear velocity $v_{\text{shear}} = 100$ km s⁻¹, and include radiative cooling along with isotropic thermal conduction. We use a conductivity $\kappa_{\text{cond}} = T_6^\alpha 10^6$ erg cm⁻¹ s⁻¹ K⁻¹, where $T_6 = T/10^6$ K, and a CIE cooling function Λ based on a broken power law fit to Gnat & Sternberg (2007). When we vary conduction and cooling in our simulations, we label them as κ_n and Λ_n respectively, where the sub-

script n denotes a constant prefactor multiplying the fiducial values as stated above. We also label α_m where m is the value of the exponent α . (The fiducial simulation is thus $\kappa_1 \Lambda_1 \alpha_0$.)

We calculate ion column densities along sightlines through simulations using Trident (Hummels et al. 2017), which generates synthetic spectra. For simplicity and consistency, we ignore photoionization and instead assume CIE ion fractions from Gnat & Sternberg (2007). Using Trident, we add fields for the ions we are interested in by post-processing snapshots from simulation data. We assume solar metallicity and zero redshift. Lastly, we use pyatomdb (Foster & Heuer 2020) to compute line emissivities.

3.3 1D Mixing Layer Models

Why should complex TMLs be amenable to 1D modeling? Here we justify this approach. TMLs can be characterized by their Damköhler number, $Da = t_{\text{turb}}/t_{\text{cool}}$, the ratio of the eddy turnover time of the largest eddies to the cooling time. While t_{cool} is temperature dependent, it has proven useful to evaluate Da at the temperature of mixed gas, $T_{\text{mix}} \sim (T_{\text{cold}} T_{\text{hot}})^{1/2}$, which we henceforth assume. TMLs can be either single phase ($Da < 1$), with temperature varying continuously with depth in the interface, or multiphase ($Da > 1$), with the slice-averaged cold gas fraction changing continuously (Tan et al. 2021). If multiphase, the interface has a large scale fractal structure down to the scale of the interface width (Fielding et al. 2020). In both regimes, numerically converged mass entrainment rates do not require the resolution of the Field length or even the interface, only the outer eddy scale of turbulence in cold gas (Tan et al. 2021). This is of order the box size here, and the cold cloud size in driven turbulence (Gronke et al. 2022).

That the temperature distribution of a single phase TML can be reproduced by a

1D model is reasonable. But doing so for a fractal, multiphase TML might appear implausible. It is useful to distinguish between macroscopic and microscopic heat diffusion. Macroscopic heat diffusion (such as turbulence) drives global energy transport, dictating the global structure of the TML and the *coarse-grained* temperature profile⁹. However, it does not drive actual energy exchange between the two phases; all that changes is the relative amount of hot and cold gas. Changing the *fine-grained*, thermodynamic temperature – which determines ionic abundances – requires microscopic heat transfer via explicit thermal conduction. Each segment of the fractal interface between hot and cold gas is a locally planar, laminar heat conduction front, whose temperature profile is set by the competition between explicit thermal conduction and cooling. All intermediate temperature gas lies in this interface, which *can* be modeled in 1D. Since the same universal interface profile holds at every segment of the fractal interface, it sets the temperature PDF. The global structure of the TML is immaterial. In §3.4, we verify this conjecture by comparing 1D models to 3D simulations.

Thus, if one is to accurately capture this temperature distribution in 3D simulations, then in contrast to mass entrainment, the interface (and hence the Field length) *must* be resolved. This is currently impossible in large-scale simulations. Moreover, explicit thermal conduction *must* be included.

3.3.1 1D Mixing Layers

In a 1D mixing layer, an equilibrium state can be reached between two stable phases with three ingredients: radiative cooling, thermal conduction, and enthalpy advection.

⁹Indeed, Tan et al. (2021) showed that the mean, coarse-grained temperature profile $\bar{T}(x) \approx f_c(x)T_{\text{cold}} + (1 - f_c(x))T_{\text{hot}}$, where $f_c(x)$ is the spatially varying mass fraction of cold gas, can be reproduced by a mixing length model.

In the frame of the front, this gives (?):

$$\frac{d}{dx} \left(\kappa \frac{dT}{dx} \right) = j_x c_p \frac{dT}{dx} + \rho \mathcal{L}(T), \quad (3.1)$$

where $j_x = \rho v_x$ is the constant mass flux and $c_p = \frac{\gamma}{\gamma-1} \frac{k_B}{m}$ is the specific heat at constant pressure. κ is the thermal conductivity and $\rho \mathcal{L} = n^2 \Lambda - n \Gamma$ is the net cooling rate per unit volume, where Λ is the cooling function and Γ is the heating rate. We assume that $\rho v^2 \ll P$ so pressure is constant, which can be verified in the solutions. Given the boundary conditions $T_{-\infty} = T_{\text{cold}}$, $T_{\infty} = T_{\text{hot}}$ and $\frac{dT}{dx} \Big|_{\pm\infty} = 0$, we can solve for the equilibrium solution numerically with j_x as an eigenvalue using the shooting method. We can also integrate Eq. (3.1) to give us the relationship between j_x and the surface brightness Q :

$$j_x = \frac{Q}{c_p(T_{\text{hot}} - T_{\text{cold}})}; \quad Q = - \int_{-\infty}^{\infty} \rho \mathcal{L} dx. \quad (3.2)$$

3.3.2 Temperature Distribution

For a given front solution $T(x)$, the volume weighted probability density function (PDF) of the temperature distribution is given by dx/dT multiplied by some normalization factor, henceforth referred to in our plots as just the probability density. In the lower panel of Fig. 3.1, we show in the shaded lilac histogram the temperature distribution for the solution to Eq. (3.1) with our fiducial parameters. We consider temperatures in the range from 10^4 to 10^6 K, excluding the boundary temperatures themselves. In the upper panel, we show the corresponding magnitudes of the advection, cooling and conduction terms in Eq. (3.1). Similar to the analysis in McKee & Cowie (1977) for spherical clouds, we identify three separate regions where the distribution can be understood via simplified versions of Eq. (3.1).




Figure 3.1: *Lower panel:* The lilac histogram shows the temperature distribution across the front. Three separate regions are identified where the distribution can be understood by simplifying Eq. (3.1), with colored curves showing resulting distributions. *Upper panel:* Corresponding magnitudes of the advection, cooling and conduction terms in Eq. (3.1).

- *Region A* : In the low temperature region, we can see from the upper panel of Fig. 3.1 that cooling dominates over advection and is balanced by the conduction term. Eq. (3.1) thus simplifies to

$$\frac{d}{dx} \left(\kappa \frac{dT}{dx} \right) = \rho \mathcal{L}. \quad (3.3)$$

The blue curve in the lower panel shows the temperature distribution of the solution to this simplified equation, in excellent agreement with the actual distribution in this region (normalizations of the colored lines in the lower panel have been adjusted for easy comparison). From Eq. (3.3), we can define a characteristic length scale known as the Field length (Begelman & McKee 1990),

$$\lambda_F = \sqrt{\frac{\kappa T}{n^2 \Lambda}}. \quad (3.4)$$

- *Region B* : At some intermediate temperature, the cooling and advection terms are equal and the conduction term is zero, hence:

$$\kappa \frac{dT}{dx} \approx \text{constant}. \quad (3.5)$$

This is a single inflection point separating regions A and C.

- *Region C* : In the high temperature region where cooling is weak, the advection term dominates and is balanced by the the conduction term. Eq. (3.1) thus simplifies to

$$\frac{d}{dx} \left(\kappa \frac{dT}{dx} \right) = j_x c_p \frac{dT}{dx}. \quad (3.6)$$




Figure 3.2: Temperature distributions remain unchanged for different values of κ , Λ and P relative to fiducial values. Subscripts denote multiples of the fiducial value. However, changing the metallicity can change the distribution by changing the shape of the cooling curve.

As with Region A, the orange curve in the lower panel of Fig. 3.1 shows the temperature distribution of the solution to Eq. (3.6). In fact, for a constant κ , we can solve Eq. (3.6) analytically, which gives an exponential temperature profile with a distribution that scales as $\lambda_D/(T_{\text{hot}} - T)$, where λ_D is a diffusion length scale:

$$\lambda_D = \frac{\kappa}{j_x c_p}. \quad (3.7)$$

3.3.3 Length Scales

In regions A and C, the temperature scale heights λ_F , λ_D are obtained by balancing conduction against cooling and enthalpy advection respectively (i.e., the second and third

term of Eq. (3.1)). Since there is no net heat flux into the TML ($\frac{dT}{dx} \pm \infty = 0$), the latter two balance across the front as a whole, i.e. enthalpy advection balances cooling, giving Eq. (3.2), which gives $j_x \propto Q \propto (n^2\Lambda)\lambda$ at a given temperature. Substituting into Eq. (3.7), this gives $\lambda_D \propto \sqrt{\kappa_0/(n^2\Lambda_0)} \propto \lambda_F$, where κ_0 and Λ_0 are constant prefactors multiplying $\kappa(T)$ and $\Lambda(T)$, i.e. both the Field length λ_F (Eq. (3.4)) and the diffusion length λ_D (Eq. (3.7)) share the same scalings with respect to κ_0 , Λ_0 , and $P \propto n$ (at fixed T). Since varying any of these rescales the solutions in regions A and C identically, the temperature distribution ($\propto dx/dT$) is thus *independent* of κ_0 , Λ_0 and P , unless they change Da sufficiently to affect the choice of κ (see §3.3.4). This is verified numerically in Fig. 3.2, where we show that the distribution remains unchanged whether we vary κ_0 , Λ_0 or P . Thus, a change in isobaric cooling time or reduced conduction due to tangled B-fields does not affect the temperature PDF. While changes in the normalization of cooling or conduction processes do not affect the temperature PDF, changes in their temperature dependence (e.g., via metallicity for cooling) do, as we now discuss.

3.3.4 Non-Constant Conductivity

Consider a temperature dependent conductivity $\kappa \propto T^\alpha$. What are relevant values of α ? For single phase TMLs ($\text{Da} < 1$), where the coarse-grained and fine-grained temperatures coincide, the temperature PDF is set by turbulent heat diffusion:

- $\alpha = -0.5$: Since conductivity and diffusivity $D \sim \nu L$ are related by $\kappa = D\rho c_p \propto DP/T$, this arises when $D \propto T^{0.5}$. This is seen in the low Da regime of TMLs, where turbulent diffusion scales with the local sound speed, $D_{\text{turb}} \propto c_s \propto T^{0.5}$ (see Figure 14 of Tan et al. 2021).

For multiphase TMLs ($\text{Da} > 1$), the temperature PDF is set by microscopic thermal conduction:




Figure 3.3: Distributions for various values of α . The scaling with temperature gets steeper as α increases.

- $\alpha = 2.5$: Spitzer conductivity. Spitzer (1962) gives the thermal conductivity of an ionized plasma as:

$$\kappa_{\text{sp}} = 5.7 \times 10^{-7} T^{2.5} \text{ erg cm}^{-1} \text{ s}^{-1} \text{ K}^{-1}. \quad (3.8)$$

- $\alpha = 0$: Constant conductivity. This was assumed in previous simulations of TMLs which included thermal conduction (e.g. Kim & Kim 2013; Tan et al. 2021), largely for numerical reasons.
- $\alpha = -1$: Constant diffusivity. This is a good approximation for numerical diffusion $D_{\text{num}} \sim v\Delta x$ in simulations without thermal conduction (e.g. Kwak & Shelton 2010; Ji et al. 2019).

In Fig. 3.3, we show the corresponding temperature distributions for each α . As α increases, the temperature scale height and PDF is weighted toward higher temperatures.

3.3.5 Column Densities and Line Ratios

Given a solution for the temperature profile, we can compute column densities. From the scalings in §3.3.3, and using $N_{\text{O VI}} = \int n_{\text{O VI}} dl \propto n_{\text{O VI}} \lambda$, we obtain:

$$N_{\text{O VI}} = \ell N_c \left(\frac{\kappa}{10^6 \text{ erg cm}^{-1} \text{ s}^{-1} \text{ K}^{-1}} \right)^{0.5} \left(\frac{\Lambda(T_{\text{O VI}})}{10^{-21.5} \text{ erg cm}^3 \text{ s}^{-1}} \right)^{-0.5} \left(\frac{Z}{Z_{\odot}} \right), \quad (3.9)$$

where N_c is a reference column density, κ is the conductivity, $\Lambda(T_{\text{O VI}})$ is the the cooling function at $T_{\text{O VI}} \sim 3 \times 10^5$ K (where O VI abundance peaks), and ℓ is a correction factor. Similar equations can be written for the other ions as well. Since $\lambda_F, \lambda_D \propto \sqrt{\kappa/(n^2 \Lambda(T))}$ have the same scalings, the scaling $N_i \propto \sqrt{\kappa/\Lambda}$ holds whether the ion peaks in region A or C. The value N_c depends on α . For example, in a turbulent single phase front where $\alpha = -0.5$, $N_c = 5.10 \times 10^{11} \text{ cm}^{-2}$, but in a front with just Spitzer conduction where $\alpha = 2.5$, $N_c = 2.1 \times 10^{12} \text{ cm}^{-2}$. N_c also depends on the shape of the cooling curve (and hence metallicity indirectly).

We include a path length correction factor ℓ to match 3D simulations, since sight-lines that intersect the mixing layer at an angle have longer path lengths. We estimate this to be a factor of $\sim \sqrt{2}$. This correction factor could also account for a sightline intersecting the interface multiple times in a fractal TML. However, because the mixing layer does not often ‘wrap around’ on large scales, we find in our 3D simulations that the sightlines usually only intersect the mixing layer 1-2 times. Similarly, a line through the fractal coastline on a map will typically intersect the water-land boundary once.

An important point in applying Eq. (3.9) is the choice of κ . The two candidates

are explicit thermal conduction κ_{cond} and turbulent conduction κ_{turb} . If $\kappa_{\text{cond}} > \kappa_{\text{turb}}$, then κ_{cond} should be used. However if $\kappa_{\text{cond}} < \kappa_{\text{turb}}$, then we have to consider Da of the system. For large Da (multiphase), κ_{cond} should be used, since the width of individual interfaces are governed by explicit thermal conduction. If $\text{Da} < 1$ (single phase), they are set by turbulent conduction, and κ_{turb} should be used instead.

The scalings of Eq. (3.9) are consistent with those of Equation 30 in Ji et al. (2019). They found that $N_{\text{OVI}} \propto Z^{0.8}$, but with the change in cooling function with metallicity folded in. In the single phase regime that their simulations fall in, $\kappa_{\text{turb}} = D_{\text{turb}}\rho c_p \propto P$ which then translates to $N_{\text{OVI}} \propto P^{0.5}$, as seen in their simulations.

Similarly, we can also compute a line surface brightness Q_i as

$$Q_i = f_i Q_{\text{total}} \quad ; \quad f_i \equiv \frac{\int n^2 \epsilon_i(T) dl}{\int n^2 \Lambda(T) dl}, \quad (3.10)$$

where ϵ_i is the line emissivity and Q_{total} is the total surface brightness modeled in Tan et al. (2021) (which for a fractal interface differs from Eq. 3.2). For example, using pyatomdb to compute the emissivity of [O III] 5008.24 Å, we obtain $f_i \sim 8.4 \times 10^{-3}$ for $\alpha = 0$ and $f_i \sim 5.2 \times 10^{-3}$ for $\alpha = -0.5$.

3.4 Results

3.4.1 Temperature PDFs

In Fig. 3.4, we compare 3D hydrodynamic simulations with 1D models. The distributions are from single snapshots where the mixing layer has fully developed. The following categories are represented:




Figure 3.4: Distributions of simulations (solid lines) compared to expected corresponding 1D models (dashed lines). Simulations vary cooling and conduction prefactors, along with α . Normalizations are adjusted for comparison.

- *Strong Conduction:* When $\kappa_{\text{cond}} > \kappa_{\text{turb}}$ (purple line $\Lambda_1 \kappa_8 \alpha_0$), mixing is set by conduction and not by turbulence, and hence the distribution follows the constant κ model ($\alpha = 0$, red dashed line).
- *Single Phase:* In the low Da regime, where mixing is faster than cooling, the gas is single phased. Our fiducial setup ($\Lambda_1 \kappa_1 \alpha_0$, orange line) lies in this region. Since κ_{turb} is larger than κ_{cond} and hence dominates mixing, we expect the distribution to follow the $\alpha = -0.5$ model (orange dashed line), as explained in §3.3.4.
- *Multiphase:* The rest of the simulations have strong cooling (Λ_{128}) and hence lie in the high Da regime. Although κ_{turb} is larger than κ_{cond} , the multiphase structure of the mixing layer means that the thickness of the interface locally is still set by κ_{cond} . We show simulations for a range of α , including one with no explicit conduction (blue line) which hence only has numerical diffusivity ($\alpha = -1$, blue dashed line). The simulation with a Spitzer scaling (pink line) differs from the $\alpha = 2.5$ model for $T < 3 \times 10^5$ K. This is because the Field length at lower temperatures is unresolved, as we now discuss.

3.4.2 Resolution

What resolution is required for convergence? It is usually thought that one needs to resolve the Field length λ_F (e.g. see Figure 7 of Kim & Kim 2013). Our results are consistent with this. In our highest resolution simulations, we are just able to resolve the smallest λ_F with fiducial cooling (Λ_1). However, in simulations with strong cooling (Λ_{128}), the Field length λ_F of gas below $T = 5 \times 10^4$ K remains unresolved. As a result, we see in Fig. 3.4 that these simulations all show a dip at $\sim 2 \times 10^4$ K, where λ_F is the smallest. In resolution tests, this feature becomes more prominent as we lower the resolution. However, if the lines we are interested in only trace gas at $T \sim 10^5$ K,




Figure 3.5: Column densities from simulation sightlines (solid circles) compared with their corresponding α models (hollow circles).

then it is sufficient to just have enough resolution to resolve λ_F at 10^5 K. The lowest resolution simulation also shows a drop at higher temperatures, as numerical diffusion starts to dominate over thermal conduction. This can also be seen in the simulation with Spitzer conductivity (pink line in Fig. 3.4), where the distribution switches over from the expected model to the one with numerical diffusion at lower temperatures where the Spitzer conductivity is small.

3.4.3 Column Densities and Line Ratios

Finally, we use Trident to compute column densities along sightlines through the mixing layer simulations. We run a hundred sightlines through the mixing layer in the simulation, randomly initializing the start and end points of the sightlines on the sides

of the box each time. These column densities are then summed. While there can be a large variation in column densities along a single sightline, this is greatly reduced when passing through many mixing layers, as required to match observed column densities. To estimate the variance, we repeat this process 75 times over several time snapshots.

The average column densities per mixing layer are plotted in Fig. 3.5 for the fiducial setup (Λ_1 , single phase) and one with much stronger cooling (Λ_{128} , multiphase). Each setup is also compared to the model predictions from §3.3 with $\ell = \sqrt{2}$. For the $\alpha = -0.5$ model, we used $\kappa_{\text{turb}} = 10^7 \text{ erg cm}^{-1} \text{ s}^{-1} \text{ K}^{-1} (T/10^4 \text{ K})^{-0.5}$ from Figure 14 of Tan et al. (2021). The models and simulations are in good agreement. This implies that the number of interface intersections per mixing layer in the multiphase regime is of order unity.

We show the corresponding model and simulation line ratios in Fig. 3.6. Observational values for the Milky Way were obtained from Wakker et al. (2012) and are also shown for comparison. We find that the single phase TML model is a good match with the observations, consistent with the 2D simulations of Kwak & Shelton (2010). While the multiphase TML simulation is also a good match, the constant conductivity used therein is not physically motivated, unless conduction is saturated (Cowie & McKee 1977). Assuming a Spitzer conductivity (pink hollow circles) instead leads to predictions which differ significantly. Such potential constraints on $\text{Da} = t_{\text{turb}}/t_{\text{cool}} < 1$ are interesting, particularly if one of these timescales can be independently estimated.

We also estimate line emission from the simulations above using `pyatomdb` and compare with Eq. 3.10. For $[\text{O III}] 5008.24 \text{ \AA}$, we find $Q_i \sim 4 \times 10^{-10} \text{ erg cm}^{-2} \text{ s}^{-1}$ and $Q_i \sim 4.5 \times 10^{-9} \text{ erg cm}^{-2} \text{ s}^{-1}$ for Λ_1 and Λ_{128} respectively. While these values are around a factor of 2 higher than the model values, the difference mostly comes from post-processing the line emission rather than tracking total emission in the simulations over a timestep as done for Q_{total} .




Figure 3.6: Line ratios from simulations (solid circles) and their corresponding α models (hollow circles). Grey bars represent observational data of the MW halo from Wakker et al. (2012).

3.5 Discussion

We have found that the thermodynamic temperature distribution in TMLs depends only on the local rather than global front structure. As shown in Fielding et al. (2020) and Tan et al. (2021), the global structure is set by turbulence, which can significantly deform the front in the fast cooling limit, leading to a complicated fractal structure. The interaction between turbulence and cooling sets the overall surface brightness and mass entrainment rate. It is not sensitive to the details of thermal diffusion, and only requires the outer eddy scale to be resolved. However, the local front structure is set by balance between radiative cooling, advection, and thermal diffusion (conduction, turbulence), requiring much smaller scales such as the Field length to be resolved. Fortunately, the temperature PDFs, column densities and line ratios in converged 3D simulations are surprisingly well matched by simple 1D models of local front structure. This is good news, because resolving the local front structure is currently impossible in larger galaxy scale simulations. Instead, the model we have presented can be inserted as a subgrid prescription, when calculating the contribution of TMLs to line emission or absorption.

We regard this as significant progress. At the same time, there are refinements we have ignored, which we leave to future work. Some of these include:

- *Photoionization; Non-equilibrium Ionization (NEI)*. The effects of photoionization and NEI were studied in 3D simulations by Ji et al. (2019). They have two effects: (i) the gas is over-ionized compared to CIE expectations, and thus has larger line column densities (typically by a factor of a few). (ii) Decreased radiative cooling efficiency due to over-ionization; leading to thicker mixing layers and larger column densities, although in practice this is a minor effect. Both of these effects can be captured in our analytic model, by altering the temperature dependence of the ionization fraction $x_i(T)$ and cooling function $\Lambda(T)$, as a function of radiation field

(Wiersma et al. 2009) or cooling history Gnat & Sternberg (2007), using lookup tables.

- *Multiple clouds; Kinematic Structure.* To account for the observations, a line of sight has to pass through $\sim 100 - 1000$ mixing layers; which is conceivable if the cold gas has a ‘fog-like’ structure (McCourt et al. 2018; Gronke & Oh 2020b). Cold gas in a turbulent medium acquires a wide-ranging, almost scale-free range of sizes, but the covering fraction is dominated by small cloudlets (Gronke et al. 2022). Our analytic model enables us to assign column densities and line ratios for these cloudlets, while kinematic structure due to turbulence can be obtained from the simulation.
- *Nonthermal Forces.* Our simulations are purely hydrodynamic. B-fields can suppress the Kelvin-Helmholtz instability, significantly reduce mass entrainment rates and column densities Ji et al. (2019). Non-thermal pressure support from cosmic-rays has similar effects.
- *Metallicity and Dust Depletion.* We have assumed solar and equal cold/hot gas metallicities/abundances. The first assumption is easily modified. The second can be handled by modeling the relative cold/hot gas fraction in the mixing layer.
- *Anisotropic Conduction.* In our simulations, we only model isotropic conduction. Anisotropic conduction along tangled B-fields potentially implies a reduction in κ . As long as $\kappa(T)$ can be calibrated from high resolution simulations, it can be used in our 1D model.

Acknowledgements

We thank M. Gronke and L. Lancaster for helpful comments. We acknowledge support from NASA grants NNX17AK58G, 19-ATP19-0205, HST-AR- 15797.001-A, NSF grant AST-1911198 and XSEDE grant TG-AST180036. This research was supported in part by the National Science Foundation under Grant No. NSF PHY-1748958 to KITP, and made use of `yt` (Turk et al. 2011).

Chapter 4

Cloudy with A Chance of Rain: Accretion Braking of Cold Clouds

*Body cells replace themselves every month.
Even at this very moment. Most everything
you think you know about me is nothing
more than memories.*

Haruki Murakami, A Wild Sheep Chase

Understanding the survival, growth and dynamics of cold gas is fundamental to galaxy formation. While there has been a plethora of work on ‘wind tunnel’ simulations that study such cold gas in winds, the infall of this gas under gravity is at least equally important, and fundamentally different since cold gas can never entrain. Instead, velocity shear increases and remains unrelenting. If these clouds are growing, they can experience a drag force due to the accretion of low momentum gas, which dominates over ram pressure drag. This leads to sub-virial terminal velocities, in line with observations. We develop simple analytic theory and predictions based on turbulent radiative mixing

layers. We test these scalings in 3D hydrodynamic simulations, both for an artificial constant background, as well as a more realistic stratified background. We find that the survival criterion for infalling gas is more stringent than in a wind, requiring that clouds grow faster than they are destroyed ($t_{\text{grow}} < 4t_{\text{cc}}$). This can be translated to a critical pressure, which for Milky Way like conditions is $P \sim 3000 k_B K \text{ cm}^{-3}$. Cold gas which forms via linear thermal instability ($t_{\text{cool}}/t_{\text{ff}} < 1$) in planar geometry meets the survival threshold. In stratified environments, larger clouds need only survive infall until cooling becomes effective. We discuss applications to high velocity clouds and filaments in galaxy clusters.

4.1 Introduction

The cycle of baryons – particularly that of cold gas, the fuel for star formation – is absolutely fundamental to galaxy formation and a crucial link between galactic and cosmological scales (Péroux & Howk 2020). This cycle can take various forms: (i) Outflows due to feedback processes (Thompson et al. 2016; Schneider et al. 2018). Observationally, cold gas is frequently seen outflowing at velocities comparable to virial/escape velocities (Veilleux et al. 2005; Steidel et al. 2010; Rubin et al. 2014; Heckman & Thompson 2017). (ii) Inflow of cold gas which forms via thermal instability in the halo (Joung et al. 2012; Sharma et al. 2012; Fraternali et al. 2015; Voit et al. 2019; Tripp 2022), or is supplied by direct cosmology accretion (cold streams; Kereš et al. 2005; Dekel & Birnboim 2006), and falls under gravity. (iii) Fountain recycling, which is a combination of these two processes. A useful analogy is the terrestrial water cycle, where evaporation, condensation and precipitation both play crucial roles.

All of these motions involve velocity shear between cold gas clouds and background hot gas. A long-standing problem has been to understand why clouds are not shredded

by hydrodynamic instabilities, particularly the Kelvin-Helmholtz instability. The hydrodynamic acceleration time for a cloud of radius r , overdensity χ embedded in a wind of velocity v_w is $t_{\text{acc}} \sim \chi r / v_w$, the timescale for the cloud to sweep up its own column density. By contrast, the cloud destruction (‘cloud crushing’) time is $t_{\text{cc}} \sim \sqrt{\chi} r / v_w$, i.e. of order the Kelvin-Helmholtz time, implying that $t_{\text{acc}} / t_{\text{cc}} \sim \sqrt{\chi}$, i.e. clouds should be destroyed before they can be accelerated (Klein et al. 1994; Zhang et al. 2017). Numerous simulation studies, including those with radiative cooling, concluded that cold clouds get destroyed before they can become entrained with the wind (e.g. Cooper et al. 2009; Scannapieco & Brüggen 2015; Schneider & Robertson 2017); magnetic fields can ameliorate but do not solve the problem (McCourt et al. 2015; Gronke & Oh 2020a).

In recent years, it was realized that there are regions of parameter space where the cooling efficiency of the mixed, ‘warm’ gas is sufficiently large to contribute new comoving cold gas which can significantly exceed the original cold gas mass, enabling the cloud to survive. Cloud growth is thus mediated by these turbulent mixing layers (Begelman & Fabian 1990; Ji et al. 2018; Fielding et al. 2020; Tan et al. 2021). The criteria for this to happen is $t_{\text{cool,mix}} / t_{\text{cc}} < 1$, where $t_{\text{cool,mix}}$ is the cooling time of the mixed warm gas (with $T_{\text{mix}} \sim (T_{\text{hot}} T_{\text{cold}})^{1/2}$) and t_{cc} is the cloud crushing time (Gronke & Oh 2018). This criterion is always satisfied for a large enough cloud $r > c_{\text{s,cold}} t_{\text{cool,mix}}$ (where $c_{\text{s,cold}}$ is the sound speed of the cold gas), which grows and entrains by gaining mass and momentum from cooling mixed hot gas. Thus, the cloud eventually comoves with the wind, with a cold gas mass which can be many times the original cloud mass. These conclusions have been borne out in many subsequent studies (e.g., Sparre et al. 2020; Li et al. 2020; Abruzzo et al. 2022b; Girichidis et al. 2021; Farber & Gronke 2022).

However, cold gas survival and growth has only been understood for part of the baryon cycle, galactic outflows. To date, there have only been a handful of studies studying cold gas survival and growth during *infall*, which is arguably even more fundamental to

processes such as star formation.

An important outstanding problem in galaxy evolution is that the observed star formation rates (SFRs) in galaxies at a range of redshift are unsustainable - they would rapidly deplete current existing gas reservoirs - and hence these galaxies require some form of continuous accretion to supply the necessary fuel (Erb 2008; Hopkins et al. 2008; Putman et al. 2009). For example, our Milky Way has a SFR of $\sim 2 M_{\odot} \text{ yr}^{-1}$ but only $\sim 5 \times 10^9 M_{\odot}$ of existing fuel, and would thus burn through this supply in just 2-3 Gyrs (Chomiuk & Povich 2011; Putman et al. 2012). Supplementary inflow must come in the form of low-metallicity ($Z < 0.1 Z_{\odot}$) gas, so as to satisfy constraints from disk stellar metallicities and chemical evolution models (Schönrich & Binney 2009; Kubryk et al. 2013).

At the same time, we see infall in the form of ‘high-velocity’ and ‘intermediate-velocity’ clouds (HVCs and IVCs; Putman et al. 2012) with relatively low metallicities, as well as a galactic fountain with continuous circulation of material between the disk and corona (Shapiro & Field 1976; Fraternali & Binney 2008). Fountain-driven accretion could supply the disk with gas for star formation, and explain the observed kinematics of extra-planar gas (Armillotta et al. 2016; Fraternali 2017). It is tempting to speculate from the results of wind tunnel simulations that star formation in the disk exerts a form of positive feedback: cold gas thrown up into the halo ‘comes back with interest’, by mixing with low metallicity halo gas which cools and increases the cold gas mass.

HVCs are also good candidates and could provide a significant amount of the necessary fuel for star formation, provided they survive their journey to the disk (Van Woerden et al. 2004; Putman et al. 2012; Fox et al. 2019). First detected in HI 21 cm emission by Muller et al. (1963), HVCs are gas clouds observed moving at high velocities relative to the local standard of rest. The traditional definition for HVCs is thus those clouds with velocities in the Local Standard of Rest frame $|v_{\text{LSR}}| \geq 90 \text{ km/s}$ (Wakker & van Woerden

1991) (although similar clouds whose velocities significantly overlap that of the disk may be missed; Zheng et al. 2015). They have been observed in all regions of the sky, and come in a range of sizes (Putman et al. 2012). Clouds are grouped into various complexes based on spatial and kinematic clustering but because of their proximity, precise distances to HVCs are difficult to measure. The main method of doing so is to use halo stars of known distances in the same sky region to bracket the cloud distance by looking for absorption lines (or lack thereof) in the stellar spectra. By determining if a HVC is in front of or behind each star, the HVC's distance can thus be effectively constrained. Most HVCs with distances measured as such are found between 2-15 kpc, with most heights above the disk < 10 kpc (Wakker et al. 2008; Thom et al. 2008). The head-tail morphology observed in many HVCs (Putman et al. 2011), along with observations that the majority of high velocity absorbers kinematically and spatially lie in the vicinity of HVCs (Putman et al. 2012), strongly suggest that the HVCs are mixing as they travel through the ambient medium. There is a wealth of literature on observations of HVCs – we refer the reader to reviews such as Putman et al. (2012) for a more comprehensive account.

As we have discussed, the survival of HVCs is inherently problematic, since they are vulnerable to hydrodynamic instabilities while travelling through the hot background (Klein et al. 1994; Zhang et al. 2017). Early theoretical efforts to model HVCs initially focused on predicting their velocity trajectories, without taking into consideration their mass evolution. These early models assumed that these HVCs fell ballistically (Bregman 1980) or reached a terminal velocity when eventually slowed by hydrodynamic drag forces (Benjamin & Danly 1997), and were used in evaluating the contributions of HVCs in larger feedback models (Maller & Bullock 2004). However, the decoupling of the velocity and mass evolution implied by this approach has been shown to be untenable for HVCs with the advent of high resolution hydrodynamical simulations, many of which show that the mass and morphology of the clouds evolve significantly (e.g., Kwak et al. 2011;

Armillotta et al. 2017; Gritton et al. 2017; Gronke & Oh 2020a). While wind-tunnel setups are numerous, the number of 3D simulations of clouds falling under the influence of gravity and including radiative cooling is more limited (Heitsch & Putman 2009; Heitsch et al. 2022; Grønnow et al. 2022). The survival criterion for infalling clouds has not been quantified, and analytic models for mass and velocity evolution which match simulations do not yet exist. We will tackle these challenges in this paper.

Presumably, similar considerations apply, with a minimum cloud size $r_{\text{crit}} \sim c_{\text{s,cold}} t_{\text{cool,mix}}$ required for survival and growth. However, this ignores a crucial distinction between outflowing and infalling cold gas clouds. Outflowing gas clouds gradually entrain, so destruction processes become weaker as the velocity shear is reduced. The cloud only has to survive until it becomes comoving with the hot gas, at which point hydrodynamic instabilities are quenched (and mass growth peaks). Indeed, wind tunnel simulations (particularly for clouds with sizes just above r_{crit}) often show clouds which initially break up into small fragments, with a significant amount mixed into the hot medium, but eventually survive as the fragments entrain and grow. The cold fragments then coalesce – the cloud ‘rises from the dead’ to a peaceful environment. In contrast, infalling clouds accelerate under the action of gravity, with continually *increasing* velocity shear, and consequently increasing cloud destruction rate, which is maximized at the cloud terminal velocity. Thus, the cloud instead is exposed to continually worsening conditions, and somehow has to survive an unrelenting hot wind. Moreover, the properties of the wind change with time, as the cloud falls through a background stratified hot medium.

The survival and growth of a cold cloud under such conditions is the focus of this paper. We develop simple analytic scalings which we test in 3D hydrodynamic simulations. Unsurprisingly, several important aspects, such as cloud survival criteria, are quite different from the wind tunnel case.

What is at stake? As previously mentioned, if clouds can survive and grow, the ulti-

mate fuel supply for star formation could simply be coronal gas, whose condensation is triggered by star formation feedback and Galactic fountain recycling. During this process, cold gas also exchanges angular momentum with coronal gas, which links fountain circulation to the observable kinematics of coronal gas. More broadly, the physics of radiative turbulent mixing layers is complex, and theoretical studies demand empirical tests. Unlike clouds embedded in galactic winds, which lie at extra-galactic distances and are difficult to resolve, there is a plethora of spatially and kinematically resolved observational data for intermediate and high velocity clouds in the Milky Way. There is also ample similar data for infalling filaments in galaxy clusters (e.g. Russell et al. 2019). Such systems can be used as laboratories for the interaction between multiphase gas, mixing, and radiative cooling, which is also critical to galactic winds but difficult to test there. We shall see that we predict sub-virial terminal velocities at odds with standard predictions (which balance hydrodynamic drag with gravity) but in much better agreement with observations. Moreover, the predicted terminal velocity from the model is an observable that can be tested, at least on a statistical basis (given observational uncertainties and degeneracies). Such empirical tests have thus far been sorely lacking in cloud physics models.

The outline of this paper is as follows. In Section 4.2, we outline analytic theory and predictions for the dynamics, growth and survival of infalling cold clouds. In Section 4.3, we describe our simulation setup. In Sections 4.4 and 4.5, we describe simulation results, both for an artificial constant background (which allows us to test analytic scalings), as well as a more realistic stratified background. In Section 4.6, we discuss applications to the Milky Way (HVCs) and galaxy clusters (infalling filaments). Lastly, we summarize and conclude in Section 4.7.

4.2 Dynamics of Infalling Clouds

4.2.1 Cloud Evolution and Terminal Velocities

A falling cloud growing via accretion can be described by the following set of differential equations:

$$\frac{dz}{dt} = v \quad (4.1)$$

$$\frac{d(mv)}{dt} = mg - \frac{1}{2}\rho_{\text{hot}}v^2C_0A_{\text{cross}} \quad (4.2)$$

$$\frac{dm}{dt} = \frac{m}{t_{\text{grow}}} \quad (4.3)$$

where z , v and m represent the distance fallen, velocity, and mass of the cloud respectively, $t_{\text{grow}} \equiv m/\dot{m}$ is the growth timescale (which we discuss in Section 4.2.2), g is the gravitational acceleration, C_0 is the drag coefficient (geometry dependent; of order unity here), ρ_{hot} is the density of the background medium, and A_{cross} is the cross-sectional area which the cloud presents to the background flow. We shall see that it is important to distinguish A_{cross} from A_{cloud} , the overall surface area of the cloud. We shall also see that t_{grow} is roughly independent of mass growth, so that from equation (4.3), mass growth is nearly exponential. Note that equation (4.3) assumes steady growth and omits terms which contribute to cloud destruction. Thus, it does not apply to clouds which are losing rather than gaining mass. In this paper, we focus on scenarios where clouds survive and grow, which is the novel feature in our new model (previous works, e.g. Afruni et al. 2019, have looked at scenarios with significant mass loss). In Section 4.2.3 we will quantify the criterion for cloud survival. In this work, we only consider the hydrodynamic case and leave investigation of other factors such as magnetic fields, externally driven turbulence, and cosmic rays to future work.

The terms on the right hand side in the momentum equation (equation (4.2)) represent the gravitational and hydrodynamic drag forces. In standard models, these two terms are assumed to balance one another in steady-state, giving the hydrodynamic drag terminal velocity

$$v_{\text{T,drag}} = \sqrt{\frac{2mg}{\rho_{\text{hot}}C_0A_{\text{cross}}}} \simeq \sqrt{\frac{2\chi Lg}{C_0}} \quad (4.4)$$

for a falling cloud with volume $\sim A_{\text{cross}}L$ and $\chi = \rho_{\text{cloud}}/\rho_{\text{hot}}$. The hydrodynamic drag time (momentum divided by the drag force) is given by $t_{\text{drag}} \sim \chi L/v$. In fact, *this gives the terminal velocity only if the left hand side of equation (4.2) vanishes, $\dot{p} = m\dot{v} + \dot{m}v = 0 \Rightarrow \dot{v} = 0$, which is correct only if cloud mass does not evolve so $\dot{m} = 0$* . If $\dot{m} > 0$, i.e. the cloud grows by accreting mass from the background, then from momentum conservation, since the background gas is at rest and has zero initial momentum, this will slow down the cloud. In the limit that the hydrodynamic drag term is small compared to $\dot{m}v$,

$$v = \frac{m}{\dot{m}}(g - \dot{v}). \quad (4.5)$$

Thus, if $\dot{v} \ll g$, there is a second terminal velocity

$$v_{\text{T,grow}} = \frac{mg}{\dot{m}} = gt_{\text{grow}}. \quad (4.6)$$

The two terminal velocities in equations (4.4) and (4.6) represent regimes where the cloud acceleration under gravity is predominantly balanced by either hydrodynamic drag or the momentum transfer from background accretion respectively. We can separate them by considering the ratio $t_{\text{grow}}/t_{\text{drag}}$. When this ratio is large, gravity is balanced by drag. Conversely, when this ratio is small, gravity is balanced by accretion. The transition between the two is marked by where $t_{\text{grow}} \sim t_{\text{drag}}$. We can illustrate this by solving




Figure 4.1: The terminal velocity (normalized by the drag terminal velocity) as a function of the ratio of the growth time t_{grow} and the drag time t_{drag} . The dashed lines show the corresponding values for the terminal velocities when assumed to be set by either drag or growth. As the ratio increases, there is a smooth transition from $v_{T,grow}$ to $v_{T,drag}$.

equations (4.1) – (4.3) numerically for a constant t_{grow} . The result is shown in Fig. 4.1, where we plot the terminal velocity as a function of $t_{\text{grow}}/t_{\text{drag}}$. We can see that when $t_{\text{grow}} \ll t_{\text{drag}}$, the terminal velocity follows $v_{\text{T,grow}}$, and when $t_{\text{grow}} \gg t_{\text{drag}}$, the terminal velocity follows $v_{\text{T,drag}}$ as expected.

Which regime is more realistic? Let us first explore this for the idealized case of spherical clouds. For a spherical cloud, the growth time is given by (Gronke & Oh 2020a):

$$t_{\text{grow}} \equiv \frac{m}{\dot{m}} \sim \frac{\rho_{\text{cold}} r^3}{\rho_{\text{hot}} A_{\text{cloud}} v_{\text{mix}}} \sim \chi \frac{r}{v_{\text{mix}}}. \quad (4.7)$$

This seems long: if $v_{\text{mix}} \sim c_{\text{s,cold}}$ (a reasonable estimate; see Section 4.2.2 of Gronke & Oh 2020a and Sections 4.6 & 5.3.3 of Tan et al. 2021), then $t_{\text{grow}} \sim \chi t_{\text{sc}}$, where t_{sc} is the sound crossing time across the cloud. By contrast, the hydrodynamic drag time for a spherical cloud (as mentioned previously) is:

$$t_{\text{drag}} \sim \chi \frac{r}{v}, \quad (4.8)$$

which is much shorter, since $t_{\text{drag}}/t_{\text{grow}} \sim v_{\text{mix}}/v \sim c_{\text{s,cold}}/c_{\text{s,hot}} \sim \chi^{-1/2} \ll 1$, if we assume the virial velocity to be a characteristic infall speed, $v \sim v_{\text{vir}} \sim c_{\text{s,hot}}$. The fact that $t_{\text{drag}} \ll t_{\text{grow}}$ makes physical sense. The hydrodynamic drag time is also the timescale for a cloud to sweep up its own mass in hot gas ($\rho_{\text{hot}} A_{\text{cross}} v t_{\text{drag}} \sim \rho_{\text{hot}} r^3 \chi \sim m$). Even if all this mass is incorporated into the cloud, then at best $t_{\text{grow}} \sim t_{\text{drag}}$. In fact, only a small fraction of this gas is actually incorporated into the cloud, so that $t_{\text{grow}} \gg t_{\text{drag}}$. This suggests that hydrodynamic drag is the main drag force, which results in a terminal velocity given by equation (4.4).

However, as previously mentioned, clouds in a shearing wind do *not* remain spherical; they develop extended cometary tails (as seen both in simulations and observa-

tions). This change in geometry – and in particular the large increase in surface area – is *crucial* for enabling momentum transfer via mass growth. In hydrodynamic drag, $F_{\text{drag}} \sim \rho_{\text{hot}} v^2 A_{\text{cross}}$, the area $A_{\text{cross}} \approx \pi r^2$ is the cross-sectional area the cloud presents to the wind. Thus, F_{drag} remains roughly *constant* during cloud evolution. By contrast, in $\dot{m} \sim \rho_{\text{hot}} A_{\text{cloud}} v$, the area A_{cloud} is the surface area of the cloud available for mixing. In a cometary structure, this is dominated by the sides of the cylinder, so that $A_{\text{cloud}} \sim 2\pi r L$, where L is the length of the tail. Thus, $\dot{m} \propto A_{\text{cloud}} \propto L$ *increases* as a cloud develops a cometary tail. It is this increase in \dot{m} , and thus the effective momentum transfer rate $F_{\text{grow}} \sim \dot{p}_{\text{grow}} \sim \dot{m} v$, compared to a constant F_{drag} , which causes mass growth to dominate momentum transfer: $t_{\text{grow}} \sim \chi r / v_{\text{mix}}$ is roughly constant, while $t_{\text{drag}} \sim m v / F_{\text{drag}} \sim \rho_{\text{cloud}} A_{\text{cross}} L v / (\rho_{\text{hot}} A_{\text{cross}} v^2) \sim \chi L / v$ increases as the mass of the cloud increases. In particular,

$$\frac{t_{\text{grow}}}{t_{\text{drag}}} \sim \frac{r}{L} \frac{v}{v_{\text{mix}}} \sim \frac{r}{L} \chi^{1/2}. \quad (4.9)$$

In cloud crushing simulations, the tail grows during the process of entrainment to a length $L \sim v t_{\text{drag}} \sim \chi r$ during the ‘tail formation’ phase (?), so that $t_{\text{grow}}/t_{\text{drag}} \sim \chi^{-1/2} \ll 1$. The continuous shear for infalling clouds can lead to even more extended tails since the cloud does not entrain, so $t_{\text{grow}}/t_{\text{drag}} \ll 1$ is easily satisfied¹⁰.

Finally, it is important to realize that there is a third timescale in the problem, the free-fall time $t_{\text{ff}} \sim v_{\text{vir}}/g$. This sets the evolutionary lifetime available to clouds, before they fall to the halo center. Clouds will not grow significantly (and reach the terminal velocity $v_{\text{T,grow}}$ given by equation (4.6)), unless $t_{\text{grow}} < t_{\text{ff}}$. Indeed, $t_{\text{grow}} < t_{\text{ff}}$ is required for a subvirial terminal velocity. We can show this by recalling that $F_{\text{grav}} \sim mg \sim m v_{\text{vir}}/t_{\text{ff}}$, while the drag force from mass growth is $F_{\text{grow}} \sim \dot{m} v \sim m v/t_{\text{grow}}$. At the

¹⁰Shorter entrainment times than t_{drag} have been observed in cloud crushing simulations (e.g., Gronke & Oh 2020a; Farber & Gronke 2022).

terminal velocity v_T , we have $F_{\text{grav}} \sim F_{\text{grow}}$, so that:

$$f_{\text{sub-vir}} \equiv \frac{v_{T,\text{grow}}}{v_{\text{vir}}} \sim \frac{v_T}{c_{s,\text{hot}}} \sim \frac{t_{\text{grow}}}{t_{\text{ff}}}. \quad (4.10)$$

This is useful because $f_{\text{sub-vir}}$ – infall velocities, normalized to the virial velocity – can be measured observationally. Indeed, $f_{\text{sub-vir}} < 1$, sub-virial infall velocities, is commonly observed in LRGs (Huang et al. 2016; Zahedy et al. 2019) and galaxy clusters (Russell et al. 2016), much lower than predicted terminal velocities from hydrodynamic drag models (Lim et al. 2008). Our models can explain these puzzling observations, as we describe in Section 4.6.2. It also allows for testable predictions. Since $f_{\text{sub-vir}}$ is measured and t_{ff} is known from the density profile, we can predict $t_{\text{grow}} \approx f_{\text{sub-vir}} t_{\text{ff}}$ from kinematic observations, assuming that clouds have reached terminal velocity. This can be compared with predictions for t_{grow} from equations (4.22) and (4.23), given measured or inferred cloud and background hot gas properties. Lastly, the mass growth that a cloud experiences is $m/m_0 \sim \exp(t_{\text{ff}}/t_{\text{grow}}) \sim \exp(f_{\text{sub-vir}}^{-1})$. Thus, a measurement of sub-virial velocities directly constrains the degree to which mixing and cooling enhances cool gas infall to the central galaxy. Significantly sub-virial infall implies that cold clouds grow considerably before reaching the halo center. These analytical estimates can be compared to measurements of the mass infall rate (e.g., Fraternali & Binney 2006; Fox et al. 2019). In Section 4.5, we will also show the rather remarkable result that in an isothermal atmosphere with constant gravity, $f_{\text{sub-vir}}$ is *fixed* by geometry, specifically the scaling between cloud mass and area (α in equation (4.21)), *independent* of all other properties of the system. For our infalling clouds, we find $f_{\text{sub-vir}} \approx 0.6$.

4.2.2 Cloud Growth

Previous models of infalling clouds have considered the interplay between gravity and hydrodynamic drag forces, assuming a fixed cloud mass (Benjamin & Danly 1997). However, a fixed cloud mass is unrealistic due to various processes that trigger mixing with the hot background gas or shred the cloud. Mass evolution therefore cannot remain static; clouds should either be destroyed ($\dot{m} < 0$), or grow ($\dot{m} > 0$) over time.

In the absence of cooling, clouds moving relative to a background medium are destroyed by hydrodynamic instabilities on the cloud crushing timescale (Klein et al. 1994; Scannapieco & Brügger 2015)

$$t_{\text{cc}} \sim \sqrt{\chi} \frac{r}{v}, \quad (4.11)$$

where χ is the ratio of the cloud density to the background density, r is the cloud radius, and v is the magnitude of the relative velocity between the cloud and the background. This cloud crushing timescale reflects the destruction of the cloud via internal shocks induced inside the cloud due to its velocity with respect to the medium it is moving through (assuming that this velocity is supersonic with respect to the sound speed within the cloud), and is roughly the same timescale on which surface instabilities such as the Kelvin-Helmholtz and Rayleigh-Taylor instabilities grow to the cloud scale (Klein et al. 1994). This destructive fate can however be counteracted by mass growth due to cooling. In wind tunnel simulations of ‘cloud crushing’, Gronke & Oh (2018) found that in order for cold gas to survive, cooling needs to be strong enough to satisfy the criterion

$$t_{\text{cool,mix}} < t_{\text{cc}}, \quad (4.12)$$

where $t_{\text{cool,mix}}$ is the cooling time of the *mixed* gas, defined as $T_{\text{mix}} \sim \sqrt{T_{\text{cloud}} T_{\text{hot}}}$ (in

the spirit of Begelman & Fabian 1990, see also Hillier & Arregui 2019 for an alternative derivation). That is, if the cooling time of the mixed gas is shorter than the *initial* cloud crushing time, then cold gas survives and is eventually entrained in the hot background wind.

However, infalling clouds have an important aspect that differentiates them from clouds in a wind – gravity. Clouds encountering a hot wind gradually entrain in the wind, so that shear eventually drops to zero if the cloud manages to survive until entrainment. The cloud thus encounters destructive forces for a limited period of time. By contrast, clouds in a gravitational field will always keep falling and shearing against the background gas. Thus, the survival criterion is different, and more stringent; we discuss this in Section 4.2.3.

Assuming cloud survival, let us quantify the timescale on which clouds grow. We first derive some scaling relations, before deriving numerical expressions. For now, we ignore fudge factors (due to geometry, etc.) which can be up to an order of magnitude. As in equation (4.7), the mass growth rate of a cloud can be written as

$$\dot{m} \sim \rho_{\text{hot}} A_{\text{cloud}} v_{\text{mix}}, \quad (4.13)$$

where ρ_{hot} is the density of the hot background medium, A_{cloud} is the effective surface area of the cloud¹¹ and v_{mix} is the the velocity corresponding to the mass flux from the hot background onto this surface. As above, if we write $\dot{m} \sim \rho_{\text{cold}} A_{\text{cloud}} r$, this gives:

$$t_{\text{grow}} \sim \chi \frac{r}{v_{\text{mix}}}. \quad (4.14)$$

¹¹The effective surface area corresponds to the (smoothed) enveloping area of the cloud and not the (non-convergent) surface area of the cold gas. See Gronke & Oh (2020a) for further discussion of this distinction.

Plane parallel simulations of mixing layers (Tan et al. 2021) show:

$$v_{\text{mix}} \sim u'^{3/4} \left(\frac{r}{t_{\text{cool}}} \right)^{1/4} \sim v_{\text{shear}}^{3/5} v_0^{3/20} \left(\frac{r}{t_{\text{cool}}} \right)^{1/4} \quad (4.15)$$

where t_{cool} is the cooling time in *cold* gas (the minimum cooling time in the mixing layer, a convention we adopt henceforth) and u' is the peak turbulent velocity in the mixing layer (usually in intermediate temperature gas). Note that while the first step in equation (4.15), i.e., $v_{\text{mix}}(u')$, is generally valid, we have used the scaling $u' \propto v_{\text{shear}}^{4/5}$ for relating u' to the parameters of the setup. This scaling was found numerically in Tan et al. (2021) for plane-parallel mixing layers and we have written it here as $u' \sim v_{\text{shear}}^{4/5} v_0^{1/5}$ to preserve dimensionality (v_0 simply encodes normalization). If we set $v_{\text{shear}} \sim v_{\text{T,grow}} \sim gt_{\text{grow}}$, this yields:

$$t_{\text{grow}} \sim \chi^{5/8} \frac{r^{15/32} t_{\text{cool}}^{5/32}}{g^{3/8} v_0^{3/32}}; \quad v_{\text{T,grow}} < c_{\text{s,hot}}. \quad (4.16)$$

While the above scalings focus on the subsonic and transonic cases, large enough clouds can reach velocities exceeding the sound speed of the hot gas. In such a case, the turbulent mixing velocity saturates and stops scaling with the cloud velocity (Yang & Ji 2023), changing the above scalings. In this case, from equations (4.14) and (4.15), we obtain:

$$t_{\text{grow}} \propto \frac{\chi r^{3/4}}{c_{\text{s,hot}}^{3/5} t_{\text{cool}}^{1/4}}. \quad (4.17)$$

We now give numerical expressions, which are calibrated to simulations. For cooling dominated regimes (defined below), Tan et al. (2021) found that v_{mix} in turbulent mixing layers follows

$$v_{\text{mix}} \approx 9.5 \text{ km s}^{-1} \left(\frac{u'}{50 \text{ km s}^{-1}} \right)^{3/4} \left(\frac{L_{\text{turb}}}{100 \text{ pc}} \right)^{1/4} \left(\frac{t_{\text{cool}}}{0.03 \text{ Myr}} \right)^{-1/4}, \quad (4.18)$$

where L_{turb} is the outer scale of the turbulence. Note that equation (4.18) only applies in the ‘fast cooling’ ($\text{Da}_{\text{mix}} \equiv L_{\text{turb}}/(u' t_{\text{cool,mix}}) > 1$, where Da_{mix} is the Damköhler number; Tan et al. 2021) regime, where the cooling time is much smaller than the turbulent mixing time L_{turb}/u' . As we will discuss below, however, this is always true for surviving clouds.

Tan et al. (2021) note that u' is geometry dependent, but find for shearing layers that

$$u' \approx 50 \text{ km s}^{-1} \mathcal{M}^{4/5} \left(\frac{c_{\text{s,hot}}}{150 \text{ km s}^{-1}} \right)^{4/5} \left(\frac{t_{\text{cool}}}{0.03 \text{ Myr}} \right)^{-0.1}, \quad (4.19)$$

for $\chi \gtrsim 100$ and $\mathcal{M} \equiv v_{\text{shear}}/c_{\text{s,hot}}$. From equation (4.18), we can approximate $v_{\text{mix}} \sim c_{\text{s,cold}}$ for quick estimates. While Tan et al. (2021) only considered mixing layers with subsonic to transonic velocity shears, Yang & Ji (2023) found that beyond $\mathcal{M} = 1$, u' in the mixing region stops scaling with \mathcal{M} and saturates. We include this in our model by setting $\mathcal{M} \rightarrow \min(1, \mathcal{M})$. We find good evidence for this in our simulations.

Equations (4.18) and (4.19) assume fully developed turbulence. When a cloud falls from rest however, there is a transient period when turbulence is developing. We hence set a time dependent weight factor $w_{\text{kh}}(t)$ to account for the initial onset of turbulence. Turbulence develops over the timescale for the development of the Kelvin-Helmholtz instability; on the scale of the cloud $t_{\text{kh}} = f_{\text{kh}} t_{\text{cc}}$ where f_{kh} is some constant of proportionality (Klein et al. 1994). We use the simplest ansatz that

$$v_{\text{mix}} \rightarrow w_{\text{kh}}(t) v_{\text{mix}}; \quad w_{\text{kh}}(t) = \min \left(1, \frac{t}{f_{\text{kh}} t_{\text{cc}}} \right), \quad (4.20)$$

which amounts to v_{in} growing linearly with time over the instability growth time, until fully developed and capped at unity. We will justify this ansatz in our simulations. Since t_{cc} is changing over time, we note that $t/t_{\text{cc}} \propto vt \sim z$, where z is distance the cloud has fallen. We find in our simulations that $f_{\text{kh}} \sim 5$ for a constant background and ~ 1 for

a stratified background. In a more realistic setting with less idealized initial conditions, this time-dependent weight factor might not be necessary as the initial mixing can be already seeded from the outflowing section (assuming $v < v_{\text{esc}}$), extrinsic turbulence, or cooling induced pulsations (Gronke & Oh 2020b, 2022).

What is an appropriate scaling relation for the effective cloud surface area A_{cloud} ? In cloud crushing simulations, areal growth follows two phases (Gronke & Oh 2020a; Abruzzo et al. 2022b). In the ‘tail-formation’ phase, surface area growth is dominated by the formation of a cometary tail, with $A_{\text{cloud}} \propto L \propto m$, where L is the length of the tail. The stretching of the cloud means that the area to mass ratio $A_{\text{cloud}}/m \approx \text{constant}$, rather than $A_{\text{cloud}}/m \propto m^{-1/3}$, as for fixed geometry. Once the tail grows to a length $L \sim \chi r$ (the hydrodynamic drag length), the cloud becomes entrained in the wind from efficient momentum transfer, and due to lack of shear the tail no longer grows. The cloud surface area thereafter scales roughly as $A_{\text{cloud}} \propto (m/\rho_{\text{cloud}})^{2/3}$, as one would expect for a monolithic cloud.

However, our falling clouds do not get entrained - rather the opposite in fact, as they start at rest and accelerate until reaching some terminal velocity. This means they start ‘entrained’ and then begin to shear against background gas. They never leave the ‘tail-formation’ phase, since there is a constant velocity difference between the cloud and background medium. The cloud sees a continuous headwind which drives turbulence, mixing, and lengthening. Instead of $A_{\text{cloud}} \propto m/\rho_{\text{cloud}}$ or $A_{\text{cloud}} \propto (m/\rho_{\text{cloud}})^{2/3}$, we assume that $A_{\text{cloud}} \propto (m/\rho_{\text{cloud}})^\alpha$, where α is a growth scaling exponent between 2/3 and 1. Physically, this is because both mass growth onto the surface of the cloud and a lengthening tail are concurrent processes. We will demonstrate that this is a good assumption for the mass growth of the falling clouds in our simulations, where we find

$\alpha \approx 5/6$. The cloud surface area is thus

$$A_{\text{cloud}} \approx A_{\text{cloud},0} \left(\frac{m}{m_0} \frac{\rho_{\text{cloud},0}}{\rho_{\text{cloud}}} \right)^\alpha, \quad (4.21)$$

where $A_{\text{cloud},0}$, $\rho_{\text{cloud},0}$ and m_0 are the initial cloud surface area, density and mass respectively. Note that since $\dot{m} \propto m^\alpha$ where $\alpha = 5/6$ is close to 1, the growth is close to exponential¹². The cloud density ρ_{cloud} changes because the ambient pressure increases as the cloud falls in a stratified medium, compressing the cloud.

Using equations (4.13), (4.18), (4.20) and (4.21), we can write the growth time $t_{\text{grow}} \sim m/\dot{m}$ as

$$t_{\text{grow}} = \frac{t_{\text{grow},0}}{w_{\text{kh}}(t)} \left(\frac{c_{s,150}}{v} \right)^{3/5} \left(\frac{t_{\text{cool}}}{t_{\text{cool},0}} \right)^{1/4} \left(\frac{m}{m_0} \frac{\rho_{\text{hot},0}}{\rho_{\text{hot}}} \right)^{1-\alpha}, \quad (4.22)$$

where $c_{s,150} = 150 \text{ km s}^{-1}$ is the sound speed of gas at 10^6 K and the initial growth time $t_{\text{grow},0}$ is given by

$$t_{\text{grow},0} \approx 35 \text{ Myr} \left(\frac{f_A}{0.23} \right) \left(\frac{\chi}{100} \right) \left(\frac{r}{r_{100}} \right) \left(\frac{L_{\text{turb}}}{L_{100}} \right)^{-1/4} \left(\frac{t_{\text{cool},0}}{0.03 \text{ Myr}} \right)^{1/4}. \quad (4.23)$$

where $r_{100} = L_{100} = 100 \text{ pc}$ and r is the initial cloud size. We will assume generally that $L_{\text{turb}} \sim r$ (since the hydrodynamic instabilities which drive turbulence and mixing have an outer scale set by cloud size). We have included an unknown normalization factor f_A to account for uncertainties arising from geometrical differences between the single mixing layers in Tan et al. (2021) and our cloud setup here, the use of the initial size of the sphere as a characteristic scale (see discussion at the end of Section 4.6), and any

¹²Similar scalings $\dot{m} \propto A \propto m^\alpha$, where $\alpha \approx 0.8$, are seen in simulations of cloud growth when clouds are embedded in a turbulent medium (Gronke et al. 2022). This super-Euclidean scaling can be understood as the outcome of the fractal nature of the mixing surface, where area $A \propto m^{D/3}$, where D is the fractal dimension (Barenblatt & Monin 1983). In their mixing layer simulations, Fielding et al. (2020) measure $D \approx 2.5$, which gives $\alpha \approx D/3 = 5/6$, consistent with the above.

other simplifying assumptions we might have made. We find in our simulations that $f_A \sim 0.23$. We can simplify equations (4.22) and (4.23) by ignoring the weak mass and hot gas density dependence, and setting $L_{\text{turb}} \sim r$, to obtain:

$$t_{\text{grow}} = \frac{35 \text{ Myr}}{w_{\text{kh}}(t)} \left(\frac{f_A}{0.23} \right) \left(\frac{c_{\text{s},150}}{v} \right)^{3/5} \left(\frac{\chi}{100} \right) \left(\frac{r}{r_{100}} \right)^{3/4} \left(\frac{t_{\text{cool}}}{0.03 \text{ Myr}} \right)^{1/4}. \quad (4.24)$$

Equations ((4.22)) and (4.23) should be used when evaluating t_{grow} if the velocity $v(t)$ varies with time (i.e., when solving equations (4.1) – (4.3). However, a key quantity is the growth time at the terminal velocity $v = gt_{\text{grow}}$, which we shall see determines whether the cloud can survive (Section 4.2.3). Inserting $v = gt_{\text{grow}}$ into equation (4.22), setting $w_{\text{kh}}(t) = 1$, and using $f_A = 0.23$, we obtain the numerical version of equation (4.16):

$$t_{\text{grow}} = 40 \text{ Myr} \left(\frac{g}{g_{\text{fid}}} \right)^{-3/8} \left(\frac{\chi}{100} \right)^{5/8} \left(\frac{r}{r_{100}} \right)^{15/32} \left(\frac{t_{\text{cool}}}{0.03 \text{ Myr}} \right)^{5/32}, \quad (4.25)$$

where $g_{\text{fid}} = 10^{-8} \text{ cm s}^{-2}$. On the other hand, for supersonic speeds, as we have discussed, the turbulent mixing velocity saturates and stops scaling with the cloud velocity (Yang & Ji 2023). Setting $v \sim gt_{\text{grow}}$ to $v \sim c_{\text{s,hot}}$ instead in equation (4.24), we find the numerical version of equation (4.17):

$$t_{\text{grow}} = 35 \text{ Myr} \left(\frac{c_{\text{s,hot}}}{c_{\text{s},150}} \right)^{-3/5} \left(\frac{\chi}{100} \right) \left(\frac{r}{r_{100}} \right)^{3/4} \left(\frac{t_{\text{cool}}}{0.03 \text{ Myr}} \right)^{1/4}. \quad (4.26)$$

4.2.3 Cloud Survival

The model we have presented only accounts for mass growth of the cloud and does not include processes that result in mass loss. In addition, the initial onset and development of turbulence is only very crudely incorporated. The absence of these refinements mean that we should expect differences between model predictions and simulations, certainly

for clouds that are losing mass, and at early times even for clouds that do survive and grow. We leave the inclusion and refinement of these components for future work, as we find that the model as presented works well for surviving clouds. Since the key assumption of our model is that the cloud is growing, we now discuss when this is a valid assumption.

As we previously discussed, clouds placed in a wind tunnel encountering a hot wind can survive if $t_{\text{cool,mix}} < t_{\text{cc}}$ (equation (4.12)). Physically, $t_{\text{cool,mix}}$ can be understood as the time it takes gas to cool in the downstream tail region of the cloud. Even if the initial pristine cloud material does not survive, if *mixed* gas can cool and survive, then the cold gas mass will increase. Since this mixed gas in the tail is cooling from the background, it is much more entrained in the wind than the initial cloud and hence able to survive – once the cold gas is entrained, it is no longer subject to destruction by shear.

The ‘usual’ survival criterion $t_{\text{cool,mix}} < t_{\text{cc}}$ above is certainly a necessary condition for survival. If no gas can cool before the cloud is completely disrupted, the cloud cannot survive. However, this criterion is not a sufficient one. This is because the physical process associated with t_{cc} is not simply surface evaporation. If this were so, then the above criterion would indeed be sufficient as any mixing would lead to a net increase in cloud mass. Instead, the entire cloud is disrupted (i.e. the cloud is broken up into smaller fragments; Klein et al. 1994; Schneider & Robertson 2017). Hence, as we shall see, it is not enough that mixed gas can cool faster than the cloud crushing time.

Compared to a wind tunnel setup, the considerations for an infalling cloud are different. Since the cloud’s velocity increases instead, and there is no entrainment, t_{cc} decreases over time. The only way for cold gas to survive is if it is produced at a rate faster than

it is destroyed:

$$t_{\text{grow}} < f_{\text{S}} t_{\text{cc}}, \quad (4.27)$$

where f_{S} is some constant¹³ factor of order unity, which we shall calibrate in simulations. It encodes the fact cloud destruction takes place over several cloud crushing times (Klein et al. 1994; Scannapieco & Brügger 2015). In evaluating $t_{\text{cc}} \sim \chi^{1/2} r/v$, the cloud radius is evaluated at its initial value. As in wind tunnel experiments, this turns out to be a very good approximation, since the cloud grows mostly in the streamwise direction. If the velocity is evaluated at the terminal velocity $v_{\text{T}} \sim gt_{\text{grow}}$, then equation (4.27) is equivalent to:

$$\frac{gt_{\text{grow}}^2}{\chi^{1/2} r} < f_{\text{S}}. \quad (4.28)$$

As we have seen, there are two regimes for t_{grow} , subsonic and supersonic infall. The criterion for subsonic infall is $t_{\text{grow}} < t_{\text{ff}}$ (equation (4.10)). Using equation (4.25), and assuming $t_{\text{ff}} \sim c_{\text{s,hot}}/g$, this can be rewritten as $r < r_{\text{sonic}}$, where

$$r_{\text{sonic}} \sim 150 \text{ pc} \left(\frac{t_{\text{cool}}}{0.03 \text{ Myr}} \right)^{-1/3} \left(\frac{g}{g_{\text{fid}}} \right)^{-4/3} \left(\frac{\chi}{100} \right)^{-4/3} \left(\frac{c_{\text{s,hot}}}{c_{\text{s,150}}} \right)^{32/15}. \quad (4.29)$$

Thus, clouds must be *smaller* than some critical radius to fall at sub-virial velocities. In this regime, ($v_{\text{T}} \lesssim c_{\text{s,hot}}$), t_{grow} is given by equation (4.25), and the survival criterion, equation (4.28), becomes:

$$t_{\text{cool}} < 5 \times 10^{-3} \text{ Myr} \left(\frac{f_{\text{S}}}{2} \right)^{16/5} \left(\frac{r}{r_{100}} \right)^{1/5} \left(\frac{g}{g_{\text{fid}}} \right)^{-4/5} \left(\frac{\chi}{100} \right)^{-12/5}. \quad (4.30)$$

¹³Although we find that a constant factor is sufficient for our purposes, this coefficient has been found to vary in supersonic flows. For example, Scannapieco & Brügger (2015) found that in the cloud crushing setup with a supersonic wind, f_{S} scales as $\sqrt{1 + \mathcal{M}_{\text{hot}}}$ where \mathcal{M}_{hot} is the Mach number of the hot medium (see also Li et al. 2020; Bustard & Gronke 2022, for alternative scalings). However, we mostly probe the subsonic to transonic regime.

Note that equation (4.30) is almost independent of cloud size. Indeed, $t_{\text{grow}}/t_{\text{cc}} \propto gt_{\text{grow}}^2/r \propto r^{-1/16}$, i.e., a very weak scaling. We shall verify this in Section 4.4.4.

Is it possible for clouds to survive in the supersonic regime ($r > r_{\text{sonic}}$)? This requires $t_{\text{ff}} < t_{\text{grow}} < f_S t_{\text{cc}}$. This in turn requires that clouds be smaller than some critical size r_{SS} , since $t_{\text{grow}}/t_{\text{cc}} \propto gt_{\text{grow}}^2/r \propto r^{1/2}$ in the supersonic regime (using $t_{\text{grow}} \propto r^{3/4}$ from equation (4.26)). Thus, supersonic infall and survival requires:

$$r_{\text{sonic}} < r < r_{\text{SS}}, \quad (4.31)$$

where r_{SS} is given by:

$$r_{\text{SS}} = 100 \text{ pc} \left(\frac{t_{\text{cool}}}{0.03 \text{ Myr}} \right)^{-1} \left(\frac{g}{g_{\text{fid}}} \right)^{-2} \left(\frac{\chi}{100} \right)^{-2} \left(\frac{c_{\text{s,hot}}}{c_{\text{s,150}}} \right)^{12/5}. \quad (4.32)$$

Note that equation (4.31) can only be fulfilled if $r_{\text{sonic}}/r_{\text{SS}} < 1$, where:

$$\frac{r_{\text{sonic}}}{r_{\text{SS}}} \sim 1.5 \left(\frac{t_{\text{cool}}}{0.03 \text{ Myr}} \right)^{2/3} \left(\frac{g}{g_{\text{fid}}} \right)^{2/3} \left(\frac{\chi}{100} \right)^{2/3} \left(\frac{c_{\text{s,hot}}}{c_{\text{s,150}}} \right)^{-4/15}. \quad (4.33)$$

Figure 4.2 shows the survival criteria above (equations (4.30) and (4.32)) for $g = g_{\text{fid}}$, $\chi = 100$, $c_{\text{s,hot}} = c_{\text{s,150}}$, and $f_S = 2$. It is clear that survival is mostly independent of cloud size and depends instead on the cooling time.

In practice, the subsonic case is of most interest. There, clouds must satisfy $t_{\text{grow}} < \min(t_{\text{ff}}, t_{\text{cc}})$, which translates into a maximum allowed cloud size (equation (4.29)) and a maximum allowed cooling time in cold gas (equation (4.30)). The latter criterion is quite stringent. Since the dependence on size in equation (4.31) is weak, under isobaric conditions $t_{\text{cool}} \propto 1/P$, we can translate equation (4.30) into a critical pressure. For

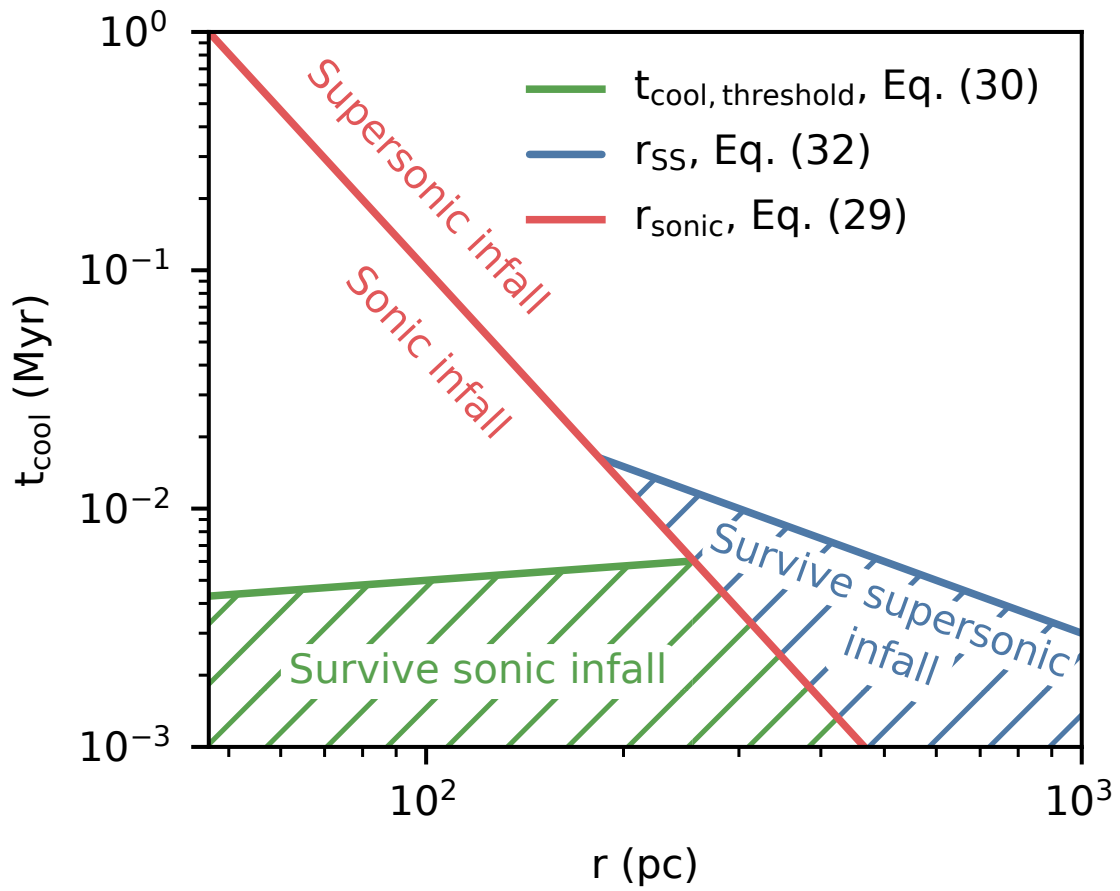


Figure 4.2: Cloud survival for subsonic and supersonic infall for different cloud sizes and cooling times. Survival is mostly sensitive to the latter.

$f_S = 2$, and ignoring the size dependence, equation (4.30) is equivalent to:

$$P > 3000 \text{ k}_B \text{ K cm}^{-3} \left(\frac{g}{g_{\text{fid}}} \right)^{4/5} \left(\frac{\chi}{100} \right)^{12/5}, \quad (4.34)$$

where the RHS is the critical pressure P_{crit} above which a falling cloud can survive. We can also write equation (4.30) in terms of the cooling time of the hot gas $t_{\text{cool,hot}} \sim \chi^2 t_{\text{cool}}[\Lambda(T_{\text{cold}})/\Lambda(T_{\text{hot}})]$ and the free fall time $t_{\text{ff}} \sim c_{s,\text{hot}}/g$ to obtain:

$$\frac{t_{\text{cool,hot}}}{t_{\text{ff}}} \lesssim 1 \left(\frac{\chi}{100} \right)^{-2/5} \left(\frac{\Lambda(T_{\text{cold}})}{\Lambda(T_{\text{hot}})} \right), \quad (4.35)$$

where we have ignored the weak dependence on g , $t_{\text{cool}}/t_{\text{ff}} \propto g^{1/5}$. This is similar to the criterion ($t_{\text{cool,hot}}/t_{\text{ff}} < 1$, where t_{cool} is evaluated at one scale height) for precipitation out of a thermally unstable background medium in a plane parallel atmosphere¹⁴ (McCourt et al. 2012). Since all our analytics and simulations are in the framework of plane parallel systems, the numerical factor in equation (4.35) will likely change in spherical systems. Equation (4.35) has the very interesting implication that clouds which condense via thermal instability are able to survive subsequent infall, as long as they are below the critical size given by equation (4.29). Note that the physics of stratified thermal instability which leads to the $t_{\text{cool,hot}}/t_{\text{ff}} < 1$ criterion – overstable gravity waves driven by cooling – is quite different from what we have discussed here, so it is non-trivial (perhaps coincidental) that both thermal instability and falling cloud survival have similar criteria.

¹⁴It is somewhat more stringent than the requirement for thermal instability in spherical systems ($t_{\text{cool,hot}}/t_{\text{ff}} < 10$; Sharma et al. 2012), where the gravitational acceleration g and hence t_{ff} varies as a function of radius. However, it has been shown that there is no geometrical difference in cold gas condensation in plane parallel and spherical geometries; the apparent difference arises from definitional differences in where $t_{\text{cool}}/t_{\text{ff}}$ is evaluated and cold gas is located (Choudhury & Sharma 2016).

4.3 Methods

We carry out our simulations using the publicly available MHD code Athena++ (Stone et al. 2020). All simulations are run in 3D on regular Cartesian grids using the HLLC approximate Riemann solver and Piecewise Linear Method (PLM) applied to primitive variables for second order spatial reconstruction. By default, we use the second-order accurate van Leer predictor-corrector scheme for the time integrator, but switch to the third-order accurate Runge-Kutta method when the former is not stable enough, in particular for simulations with a constant background where the cooling time is extremely short throughout the entire simulation.

Our simulation setups consist of rectangular boxes with identical x, y dimensions and an extended vertical z axis. They are filled with static hot $T_{\text{hot}} = 10^6$ K gas with initial density $n_0 = 10^{-4} \text{ cm}^{-3}$ at $z = 0$. A cold $T_{\text{cold}} = 10^4$ K spherical cloud, initially at rest, is also inserted, usually a quarter box height from the bottom. This placement allows us to follow the development of a cometary tail behind the cloud as it falls. The initial cloud density is perturbed at the percent level randomly throughout the cloud to reduce numerical artifacts arising from the initial symmetry. We use outflowing boundary conditions, except at the bottom of the box (negative z) where the background profile is enforced in the ghost cells and the velocity is set to be that of the frame velocity. This is valid as long as cloud material does not interact with this bottom boundary. The frame velocity is based on a cloud-tracking scheme we implement where we continuously shift into the reference frame of the center of mass of the cold gas, defined as gas below a temperature of $T \sim 2 \times 10^4$ K, an approach widely used in similar falling cloud simulations (Heitsch et al. 2022) and wind tunnel simulations (McCourt et al. 2015; Gronke & Oh 2018, 2020a). This scheme allows our simulation box to ‘track’ the cloud as it falls and hence reduces computational costs. The fiducial resolution of the boxes are $256^2 \times 2048$

(see Section 4.5.4 for a resolution test). The dimensions of the boxes are $10^2 \times 80 r_{\text{cloud}}$. This translates to r_{cloud} being resolved by ~ 25 cells.

During the simulations, the clouds are allowed to fall freely under gravity. We assume a constant gravitational acceleration $\mathbf{g} \equiv -g\hat{z}$, with $g = 10^{-8} \text{ cm s}^{-2}$, as appropriate for the Milky Way, taken from the fit in Benjamin & Danly (1997) for distances between 1 and 10 kpc. We discuss the impact of a more realistic gravitational profile and apply them within the scope of our model in Section 4.6.

In our implementation of radiative cooling, we assume collisional ionization equilibrium (CIE) and solar metallicity ($X = 0.7$, $Z = 0.02$)¹⁵. We obtain our cooling curve by performing a piece-wise power law fit to the cooling table given in Gnat & Sternberg (2007) over 40 logarithmically spaced temperature bins, starting from a temperature floor of 10^4 K , which we also enforce in the simulation. We then implement the fast and robust exact cooling algorithm described in Townsend (2009). For this cooling curve, the cooling time in the cold gas is $t_{\text{cool}} \sim 0.15 \text{ Myr}$. To emulate the effect of heating and to prevent the background medium from cooling over simulation timescales, we cut off any cooling above $5 \times 10^5 \text{ K}$. The particular choice of this value is unimportant (Gronke & Oh 2018, 2020a; Abruzzo et al. 2022b).

We run two different sets of simulations with different static background profiles. The first set has gravity acting on a cloud which is embedded in a constant background, i.e. constant hot gas temperature, density and pressure. This is obviously unphysical, since there are no pressure gradients in the background to counteract gravity. However, it is very useful for understanding the underlying physical mechanisms which affect the cloud, without the confounding effects of the varying background which a cloud falling through a stratified medium experiences. To prevent the background from falling under

¹⁵We phrase our results in terms of cooling times, so they can easily be scaled for different cooling curves. We note however, that the minimum cooling time at $T \sim 1.5 \times 10^4 \text{ K}$, which is dominated by hydrogen cooling, is relatively insensitive to metallicity.

gravity, we introduce an artificial balancing force $\rho_{\text{hot}}g$ upwards. The hot background thus feels a net zero force from gravity, while the cold cloud is negligibly affected. For this set of simulations, we also vary the cooling time by changing the normalization of the cooling function by a constant factor Λ_0 . For example, $\Lambda_0 = 100$ would be a case where cooling is a hundred times stronger than the fiducial value, corresponding to cooling an environment where $n_{\text{hot}} = \Lambda_0 n_0 = 10^{-2} \text{ cm}^{-3}$, or $nT = 10^4 \text{ K cm}^{-3}$, a relatively high pressure. For the constant background, we adopt $\Lambda_0 = 100$ as a default, so that cooling is extremely strong and cloud growth is guaranteed. We emphasize that the constant background is simply used to provide a clean test of our analytic model, so that (for instance) the cooling time is not a function of position, as in a stratified atmosphere.

The second, more realistic, setup is that of a hydrostatic isothermal halo. The density profile of the background is thus:

$$n(z) = n_0 \exp\left(-\frac{gm_H}{k_B T_{\text{hot}}} z\right), \quad (4.36)$$

where n_0 is the midplane density, z is the height above the disk and $H \equiv k_B T_{\text{hot}}/gm_H = 2.8 \text{ kpc}$ is the isothermal scale height (assuming the mean molecular weight $\mu = 1$). This is a simplified model that is likely to break down close to the disk below 2 kpc, where it likely underestimates the background density, since the background gas is cooler. However, this simple model allows us to study the effects of both a changing background profile and the resultant decrease in cooling time as the cold gas falls inwards. Besides the initial setup of the background profile, since we are employing a cloud-tracking scheme, the boundary cells are set accordingly throughout the course of the simulation using this background profile and the current height of the cloud, which we also track.

Our cloud chambers are somewhat artificial in that they are arbitrarily long. Thus, for instance, in the stratified case, the cloud can fall through an unrealistically large

number of scale heights (well beyond when the plane parallel approximation is valid). In practice, transition to a spherical gravitational potential with declining gravitational acceleration g means that even if clouds fall ballistically, they will only accelerate to transonic velocities, rather than fall supersonically. However, our setup is a clean probe of the underlying physics. In all the cases we care about, where the cloud survives, infall is subsonic.

In order to evaluate the cold gas mass m as well as other related quantities such as the mass growth rate, we use a temperature threshold of $T \sim 2 \times 10^4$ K below which we define the gas to be ‘cold’. No magnetic fields are included in our simulations. We leave the exploration of the MHD case to future work.

4.4 Results : Constant Background

Our first objective is to test our semi-analytic model for falling clouds (equations (4.1) – (4.3)) against full 3D simulations. Hence, the first set of our simulations are set up with a constant background, where the properties of the background medium are held unchanged as the cloud falls. We use this setup as a simple way to explore and test our model in an environment where the cooling time is kept constant. This allows us to test the various components of our model by adjusting individual parameters, *ceteris paribus*.

4.4.1 Time Evolution

In order to understand the dynamical evolution of a falling cloud, we first present the time history of various quantities of interest, both as predicted by the model and as seen in the simulations. Note that the model (equations (4.1) – (4.3)) predicts $m(t)$, $v(t)$, and $z(t)$ independently, without using any input from the simulations. Figure 4.3 shows the evolution of these quantities over the course of a simulation with an initial cloud radius

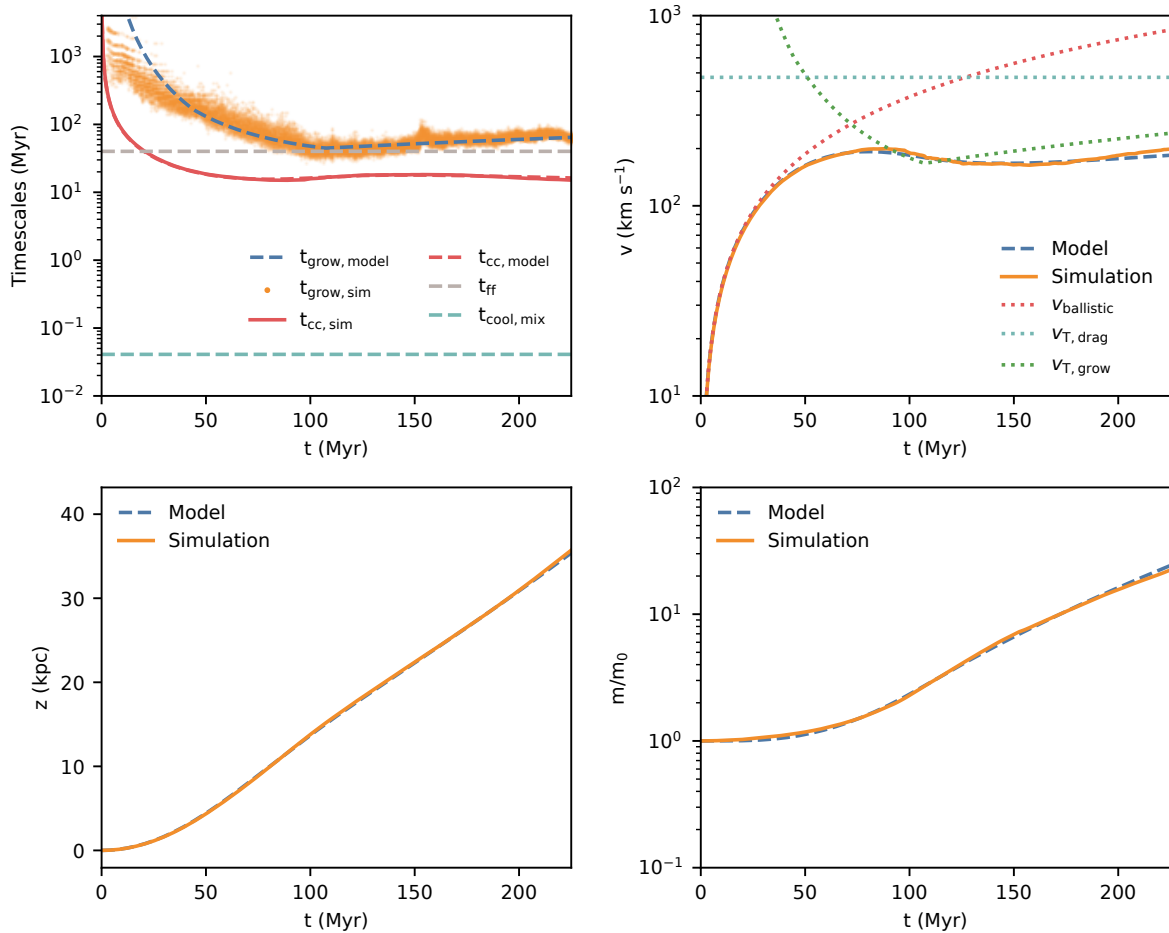


Figure 4.3: Time evolution of various quantities for a $r = 300$ pc cloud falling in a constant background. From left to right, top to bottom, the panels compare the growth time t_{grow} , the velocity v , the distance fallen z , and the cold gas mass m of the cloud in the simulation versus the model. The upper panels also include comparison with other quantities of interest. Model predictions are in good agreement with simulations results.

$r = 300$ pc. These are, from left to right and top to bottom - timescales, cloud velocity, distance fallen, and the total mass of cold gas. The simulation runs for over 200 Myr, which is between 10 to 15 cloud crushing times.

The various timescales shown in the upper left panel of Fig. 4.3 are as follows: The cooling time of the *mixed* gas $t_{\text{cool,mix}}$, where mixed gas is defined as gas at $T_{\text{mix}} \sim \sqrt{T_{\text{hot}}T_{\text{cold}}} \sim 10^5$ K, the free-fall time $t_{\text{ff}} = c_{\text{s,hot}}/g$, the cloud crushing time $t_{\text{cc}} = \sqrt{\chi}r/v$, which uses the *initial* cloud radius r and the instantaneous cloud velocity, and the instantaneous cloud growth time $t_{\text{grow}} = m/\dot{m}$, computed using the mass of cold gas (defined as gas with $T < 2 \times 10^4$ K). For the latter two timescales (t_{cc} and t_{grow}), both model and simulation results are shown for comparison. While wind tunnel setups define t_{cc} using the initial wind velocity, we use the instantaneous cloud velocity (defined as the center of mass velocity of cold gas) instead. This changes with time - it is initially infinitely long since the cloud starts at rest, but decreases as the cloud accelerates. Similarly, t_{grow} is initially infinite, since there is no turbulence at the start of the simulation (any mixing would be due to numerical diffusion, since we do not implement physical diffusion). Mass growth then begins with the initial onset of turbulence, which we have included in the model via the weight term $w_{\text{kh}}(t)$. Our crude model for $w_{\text{kh}}(t)$ means that our analytic model for t_{grow} is less accurate at these times. However, since t_{grow} is in any case long in these stages, with mass increasing very slowly, inaccuracy in modeling the growth of turbulence fortunately has little impact on $m(t)$ (and by extension $v(t)$ and $z(t)$). The model performs well at matching the simulation results for both t_{cc} and t_{grow} . Since $t_{\text{grow}} \sim t_{\text{ff}}$, the terminal velocity of the cloud here is roughly the sound speed of the hot gas, as expected from equation (4.10). For all simulations, $t_{\text{cool,mix}} \ll t_{\text{cc}}$, as required to be in the fast cooling regime.

The upper right panel of Fig. 4.3 shows the velocity evolution of the cloud, as measured by the center of mass velocity of the cold gas. We also show the velocity trajectory

from the model, along with three other characteristic velocities. These are the ballistic velocity $v_{\text{ballistic}} = gt$ and the ‘terminal’ drag and growth velocities $v_{\text{T,drag}}$ and $v_{\text{T,grow}}$ respectively, as given by equations (4.4) and (4.6). The terminal velocities¹⁶ are computed using the size of the initial cloud, and we can see that $v_{\text{T,grow}} < v_{\text{T,drag}}$, as expected. The ram pressure drag experienced by the cloud is thus much weaker than the mixing-cooling induced drag due to momentum transfer as hot surrounding gas is accreted onto the cloud (as expected from the estimates presented in Section 4.2.1). The relative contribution of ram pressure drag can be seen in the small deviation of the model (which includes both effects) from $v_{\text{T,grow}}$. The cloud initially accelerates ballistically, before reaching a high enough velocity where the cooling drag force kicks in and slows the cloud down. Since the cooling drag force operates on a timescale t_{grow} , the cloud remains ballistic until $t \sim t_{\text{grow}}$. This progression means that the cloud can experience a phase where its velocity is *decreasing* as it falls. While not strongly apparent in this setup, this effect can be pronounced when the background is not constant, which we discuss in the following section. The model does an excellent job at matching the evolution of the cloud velocity over time, and in particular the cloud reaches the asymptotic velocity $v_{\text{T,grow}} \approx gt_{\text{grow}}$ predicted by the mixing and cooling induced accretion of hot background gas.

The remaining two lower panels of Fig. 4.3 show the distance the cloud has fallen and the total mass of cold gas. Of course, the two quantities are not independent from the upper panels: we expect to predict $z(t)$ accurately since we predict $v(t)$ accurately, and we expect to predict $m(t)$ accurately since we could predict t_{grow} accurately. Overall, it is remarkable how well our simple model of ‘accretion braking’ matches the simulations. We now explore how it performs in different regions of parameter space.

¹⁶While we use the terminology of a ‘terminal’ velocity, $v_{\text{T,grow}} \approx gt_{\text{grow}}$ is in fact time dependent here since t_{grow} has a mass dependence.

4.4.2 Area Growth Rate

We first investigate the areal growth scaling in equation (4.21), where we stated that we expect the value of α to lie between $2/3$ and 1 . Equation (4.22) can be rewritten as

$$\dot{m} = \frac{m_0}{t_{\text{grow},0}} \left(\frac{v}{c_{s,6}} \right)^{3/5} \left(\frac{m}{m_0} \right)^\alpha. \quad (4.37)$$

Figure 4.4 shows the mass growth rate of cold gas \dot{m} as a function of the cold gas mass m normalized by the initial cloud mass m_0 in three simulations with $r = 100, 300$ and 1000 pc. We expect from our model that past the turbulent onset and acceleration phases, the cloud should reach terminal velocity and its mass growth rate should thus follow lines with slope α . The dashed lines in Fig. 4.4 show mass growth rate curves from our model with $f_A = 0.23$ and $\alpha = 5/6$. These choice of values give a good match to the mass growth rate curves from simulations represented by the solid lines, which are obtained by smoothing the instantaneous values of \dot{m} represented by the grey points. The slopes are initially steeper as the cloud accelerates. As discussed in the Section 4.2, we find that $\alpha \sim 5/6$ seems to be an good fit to simulation data – supporting the idea that both processes of cloud growth on the surface ($\alpha \sim 2/3$) and in a lengthening tail ($\alpha \sim 1$) are at play (or that the effective surface area scales in a fractal manner).

As noted above, we also observe a ‘burn-in phase’, where the mass growth is initially low because turbulence is developing around and behind the cloud due to instabilities, then ramps up quickly due to both turbulent onset and a rapid increase in surface area. Small sudden drops are associated with cold mass that exits the simulation box due to its fixed size, which are likely to occur at late times in our simulations. The computational cost of tracking cloud growth over longer periods of time increases significantly as the clouds keep growing in size and length which require increasingly larger boxes to contain. For the large 1 kpc radius cloud, we were unable to run the simulation for a sufficient




Figure 4.4: The mass growth rate as a function of cold gas mass for clouds of different initial sizes. Curves are labelled by the initial cloud radius and whether they represent model solutions (M) or simulations (S), which are shown as dashed and solid lines respectively. Using a scaling of $\alpha = 5/6$ in the model matches the mass growth rate in the simulations well.

time to see the mass growth rate reach the same steady growth as convincingly as the smaller clouds, but nevertheless the mass growth is in line with model predictions for all growing clouds.

4.4.3 Scalings

To verify our analytic scalings for t_{grow} in the subsonic and supersonic regimes, equations (4.25) and (4.26), we vary each parameter to test the scalings explicitly. However, the parameters cannot be arbitrarily varied – they are limited to the region of parameter space where the clouds survive. This is given by equation (4.30) and (4.31) for subsonic and supersonic infall respectively.

Scaling With Cloud Size

We first vary the initial cloud size r . The upper plot of Fig. 4.5 shows t_{grow} as a function of time for the range of cloud sizes, while the lower plot shows the scaling of t_{grow} with r , measured at the times indicated by the black circles in the upper plot where the weight function in the model reaches unity, or in other words, turbulence and mixing has fully developed. In the upper plot, simulation results are represented by the small points colored by cloud size. Solid lines show model predictions. In the lower plot, the orange line represents the model predictions while the analytic scalings of $r^{15/32}$ and $r^{3/4}$ derived above (before and after saturation of turbulent velocities for subsonic and supersonic infall respectively) are plotted as dashed lines. The simulation results match the model and analytic scalings.

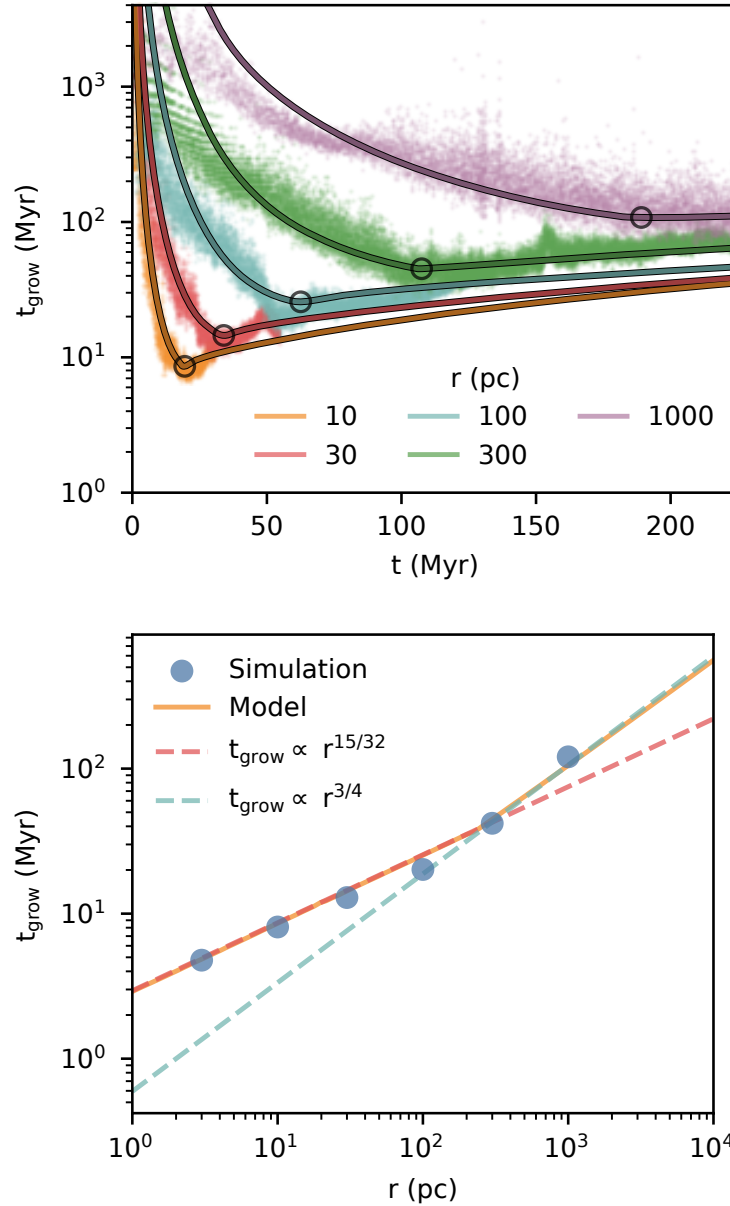


Figure 4.5: *Upper panel:* The growth time as a function of time for clouds of different sizes in the effective cooling regime ($\Lambda_0 = 100$). All clouds shown here are growing and survive. Solid lines show model predictions, while colored points represent simulation results. *Lower panel:* The growth time where turbulence is fully developed ($w_{\text{kh}}(t) = 1$) as a function of cloud size. Dashed lines show expected analytical scalings in the subsonic ($t_{\text{grow}} \propto r^{15/32}$) and supersonic ($t_{\text{grow}} \propto r^{3/4}$) regimes, while the solid orange line shows the model predictions. Both are in agreement.




Figure 4.6: The growth time for different cooling strengths Λ_0 , which modify the cooling time $t_{\text{cool}} \propto \Lambda_0^{-1}$. Dashed lines show expected analytical scalings, while the solid orange line shows the model predictions. As expected, the dependence of t_{grow} on t_{cool} is weak.

Scaling With Cooling

Next, we vary the cooling strength parameter Λ_0 by a factor of 3 above and below the fiducial value. Figure 4.6 shows the scaling of t_{grow} with $\Lambda_0 \propto 1/t_{\text{cool}}$, along with the simulation and model results as before. The simulations are in agreement with the weak t_{cool} scaling. Despite this, as we will see later, survival is sensitive to cooling time rather than size, and hence it is difficult to probe the scaling to weaker cooling. Unfortunately, reducing the cooling strength further leads to cloud destruction. Higher cooling strengths require shorter timesteps and larger boxes, and are hence numerically challenging. While we vary the cooling strength explicitly here, strong cooling also corresponds to denser environments where higher densities lead to shorter cooling times.

Scaling With Gravity

We also vary the gravitational strength g from 0.1 to 3 times the fiducial value. Figure 4.7 shows the scaling of t_{grow} with g . As before, we also plot the model and the expected $g^{-3/8}$ and g^0 scalings, for subsonic and supersonic infall respectively. Simulation results are consistent with the model in both cases.

Scaling With Density Contrast/Hot Gas Temperature

Lastly we vary χ by changing the background temperature. Figure 4.8 shows the scaling of t_{grow} with χ . Unlike the previous sections, we do not see the expected $\chi^{5/8}$ scaling. This can be understood by the scaling of the turbulent velocity u' with χ ; in our derivation, we assumed u' is independent of χ . As seen in the middle panel of Fig. 12 of Tan et al. (2021), this is true for $\chi \gtrsim 100$, but for $\chi \lesssim 100$, then $u' \propto \sqrt{\chi}$. If we put in this scaling $u' \propto \sqrt{\chi}$, we see that the χ dependence of t_{grow} becomes weaker and better matches the simulation results. We expect the our predicted $t_{\text{grow}} \propto \chi^{5/8}$ scaling to hold




Figure 4.7: The growth time for different gravitational fields. Dashed lines show expected analytical scalings $t_{\text{grow}} \propto g^{-3/8}, g^0$ for subsonic and supersonic infall respectively, while the solid orange line shows the model predictions.




Figure 4.8: The growth time for different overdensities. Dashed lines show analytical scalings $t_{\text{grow}} \propto \chi^{5/8}, \chi$ for subsonic and supersonic infall respectively. At low overdensities ($\chi \lesssim 100$), the simulations differ from the expected scalings, which we attribute to lower turbulent velocities in mixing layers. If this is taken into account, simulations and models (dotted green line) match. We also test one case at high overdensity $\chi \sim 1000$ for cluster-like parameters, where multiple parameters were varied. The model and simulations match well.

at higher χ , but the simulations required to probe this regime in detail require very long boxes and are beyond the numerical scope of this work. We also plot a single simulation, along with the model expectation, where multiple parameters were varied, not just χ , so as to sample a different region of parameter space with higher χ . These are plotted as standalone points. For this particular simulation, the parameters we have used are $r = 300 \text{ pc}$, $\chi = 1000$, $g = 4 \times 10^{-8} \text{ cm/s}^2$ and $n = 1 \text{ cm}^{-3}$. Cooling here is not boosted since we use a high density instead (i.e. $\Lambda_0 = 1$). We find that the growth time for this simulation remains in line with the model.

4.4.4 Survival

Since we are primarily interested in modeling clouds which are growing, it is useful to determine when we are in such a growth regime. In Section 4.2.3, we argued that this criterion is given by $t_{\text{grow}} < f_S t_{\text{cc}}$, where f_S is some constant factor of order unity. We now test this by running a number of simulations to explore the parameter space, varying the initial cloud radius between 3 pc and 3 kpc, and the cooling time between the fiducial value and 100 times shorter. Figure 4.9 shows¹⁷ the fate of simulated clouds for various cloud sizes and cooling times. Solid lines denote a contour of constant cooling strength, while the vertical axis shows the ratio of the growth time to the cloud crushing time $t_{\text{grow}}/t_{\text{cc}}$. These timescales are calculated by evaluating the model where our weight factor $w_{\text{kh}}(t) = 1$. Physically, this is where turbulence has fully developed and t_{grow} stabilizes. Alternatively, evaluating $t_{\text{grow}}/t_{\text{cc}}$ at some time αt_{cc} yields the same result, but can change the normalization of $t_{\text{grow}}/t_{\text{cc}}$ (this ratio gets larger as α gets smaller

¹⁷Question marks denote simulations where it is unclear what the fate of the cloud is. For example, the cloud might break up, with one portion accelerating and getting destroyed, while leaving some much slower falling material behind it that possibly survives and grows. The cold material then hits the boundary of the box at the top or bottom and we cannot track further evolution. This seems to happen near our survival boundary, where the long term fate of the cloud can be sensitive to cloud dynamics. It also happens for the largest clouds.




Figure 4.9: Overview of the fate of cold gas in the ‘constant background’ case as a function of cloud size and different cooling strengths. Points denote whether clouds in the corresponding simulations are growing in mass or losing mass; question marks denote cases where the fate is uncertain. The breaks correspond to where the turbulent velocity u' saturates when the cloud velocity reaches the sound speed of the hot background $c_{s,\text{hot}}$. This causes $t_{\text{grow}}/t_{\text{cc}}$ to increase with cloud size instead. In the simulations marked with ‘?’, the final fate of the cold gas is unclear.

since $w_{\text{kh}}(t) < 1$). The implication here is that the threshold value of f_S is depends on when $t_{\text{grow}}/t_{\text{cc}}$ is evaluated.

In general, the results are in line with criterion $t_{\text{grow}}/t_{\text{cc}} \lesssim f_S \sim 4$ for survival, and the discussion in Section 4.2.3. Rather than being sensitive to cloud size, clouds get destroyed when cooling is weak, and only survive when cooling is strong enough. Cloud size does begin to play a role when $t_{\text{grow}} > t_{\text{ff}}$, so that infall velocities become supersonic, and $t_{\text{grow}}/t_{\text{cc}} \propto r^{1/2}$. As discussed in Section 4.2.3, this happens when $r > r_{\text{sonic}}$, (equation (4.29)); $r_{\text{sonic}} \sim 200\text{pc}$ in our models, where we see the change to a $t_{\text{grow}}/t_{\text{cc}} \propto r^{1/2}$ scaling. The low mass growth rates at high Mach number means that it is harder for clouds to fall supersonically and still survive; it is only possible in a limited size range $r_{\text{sonic}} < r < r_{\text{SS}}$ (where r_{SS} is given by equation (4.32)).

To reinforce the point that $t_{\text{grow}}/t_{\text{cc}} < f_S$ is a more stringent survival criteria than others, in Fig. 4.10 we show the boundaries in the $r - \Lambda_0$ plane for two other possible criteria: (i) $t_{\text{cool,mix}} < t_{\text{cc}}$, which is the criterion for cloud survival in a wind, (ii) $\text{Da}_{\text{mix}} \equiv L/(u't_{\text{cool,mix}}) > 1$, which is the criterion for a multi-phase medium in the presence of turbulence and radiative cooling (Tan et al. 2021). The two criterion are closely related. In Fig. 4.10, we see that clouds which satisfy these criterion are nonetheless destroyed, while the more restrictive criterion $t_{\text{grow}}/t_{\text{cc}} < f_S$ straddles the boundary between destruction and survival. Note that for sufficiently small clouds, $t_{\text{grow}} < t_{\text{cool,mix}}$ (blue dashed line) instead of the other way round. However, this lies in the cloud destruction regime and thus is irrelevant.

4.4.5 Growth and Free-fall Timescales

In Section 4.2, we saw that if the drag force from mass accretion balances gravity such that $F_{\text{grav}} \sim F_{\text{grow}}$, then we expect that $t_{\text{grow}}/t_{\text{ff}} \sim v_{\text{T,grow}}/c_{\text{s,hot}}$. We show that we do




Figure 4.10: Comparison of various survival criteria (dashed lines) to the simulation results as a function of cloud radius and cooling strength. The criteria are satisfied above the respective lines. The symbols indicate whether a cloud grows or gets destroyed (as in Fig. 4.9).




Figure 4.11: Evolution of the falling velocity of the cloud as a function of (evolving) t_{grow}/t_{ff} for different cloud sizes. Black triangles indicate the direction of evolution at $t = t_{ff}$. As the cloud accelerates, the growth time is decreasing until it stabilizes at the growth terminal velocity $v_{T,grow}/c_{s,hot} \sim t_{grow}/t_{ff}$.

indeed see this in our simulations in Fig. 4.11. The blue dotted line shows the equality, while the colored points are simulation results for various cloud sizes over time. Solid lines show the model values for the same time range as the corresponding simulations. Initially, t_{grow} is large as turbulence develops, but once they reach the terminal velocity $v_{\text{T,grow}} \sim gt_{\text{grow}}$, falling clouds indeed obey the scaling $v_{\text{grow}} \sim c_{\text{s,hot}}(t_{\text{grow}}/t_{\text{ff}}$, as seen from the fact that the clouds evolve to the blue dotted line and stays there.

4.5 Results : Stratified Background

In our second set of simulations, we consider a more realistic setup of a cloud falling through an isothermal hydrostatic background. This means that $P, \rho \propto \exp\left(\frac{z}{H}\right)$, where z is the vertical height the cloud has fallen and H is the scale height of the background medium. As mentioned in Section 4.3, the density profile of the background is thus:

$$n(z) = n_0 \exp\left(\frac{z}{H}\right), \quad (4.38)$$

where $n_0 = 10^{-4} \text{ cm}^{-3}$ is the initial background density, z is the height the cloud has fallen and $H \equiv k_B T_{\text{hot}}/gm_H = 2.8 \text{ kpc}$ is the isothermal scale height (assuming the mean molecular weight $\mu = 1$). We define our origin where the cloud begins to fall, hence density increases rather than decreases exponentially with z . While the use of a constant gravitational acceleration g is not in general a realistic assumption, this simplification helps in isolating the relevant physics.

4.5.1 Time Evolution

We now present the time evolution of a simulation where the cloud comfortably survives, along with the model predictions for various quantities. Unlike the constant

background setups, we do not artificially boost the cooling function in these simulations. Instead, the cooling time naturally varies with density and hence height. Fig. 4.12 shows the evolution of these quantities over the course of a simulation with an initial cloud radius $r = 1 \text{ kpc}$ and $g = g_{\text{fid}} = 10^{-8} \text{ cm s}^{-2}$. As before, these are, from left to right and top to bottom - timescales, cloud velocity, distance fallen, and the total mass of cold gas.

The upper left panel of Fig. 4.12 shows the same timescales as in Fig. 4.3: The cooling time of the *mixed* gas $t_{\text{cool,mix}}$, which decreases as the clouds falls, the free-fall time $t_{\text{ff}} = c_{\text{s,hot}}/g$, the cloud crushing time $t_{\text{cc}} = \sqrt{\chi}r/v$, which uses the *initial* cloud radius r and the instantaneous cloud velocity, and the instantaneous cloud growth time $t_{\text{grow}} = m/\dot{m}$, computed using the mass of cold gas. For the latter two timescales (t_{cc} and t_{grow}), both model and simulation results are shown for comparison. We have adjusted the value of f_{kh} in the weight term $w_{\text{kh}}(t)$ to be 1 for the stratified background as that is more in line with simulation results. This suggests a more rapid onset of turbulence for clouds that are falling into a denser background (this parameter is of course, only a crude approximation of the relevant processes involved). The model performs well at matching the simulation results for both t_{cc} and t_{grow} , although marginally less so than for the constant background. This can be attributed to the cloud initially travelling through a region of parameter space where it is not in the growth regime. Since our model does not include cloud destruction, this leads to a deviation of the simulation from the model. The velocity evolution of the cloud is shown in the upper right panel. The cloud initially accelerates ballistically, before the cooling drag force kicks in and slows the cloud down. Since the cooling drag force operates on a timescale t_{grow} , the cloud remains ballistic until $t \sim t_{\text{grow}}$. During this time, the cloud can reach velocities greater than the eventual terminal velocity $v_{\text{T,grow}} = gt_{\text{grow}}$. The subsequent deceleration due to cooling slows the cloud down such that the velocity turns around and starts to *decrease*. This has implications for cloud survival which we discuss further on. At late

times the cloud velocity approaches a roughly constant value. We now delve into this further.

4.5.2 Terminal Velocity

Previously, we argued that the terminal velocity should approach a value $v_{T,\text{grow}} \approx gt_{\text{grow}}$ (equation (4.6)). Indeed, it does so, after some ‘overshoot’ as described above. However, as apparent from equation (4.22), t_{grow} itself is a function of parameters such as $t_{\text{cool}}(t)$, $m(t)$, $\rho_h(t)$ which change with time as the cloud falls through a stratified atmosphere. Thus, one might expect t_{grow} and consequently $v_{T,\text{grow}}$ to vary with time as the hot plasma surrounding the cloud increases in density. Instead, what is surprising from Fig. 4.12 is that t_{grow} asymptotes to a *constant* value. Indeed, it does so quite early, before $v \rightarrow v_{T,\text{grow}}$. How can we understand this?

From equation (4.22), and using $t_{\text{cool}} \propto 1/n \propto \exp(\frac{-z}{H})$, we can write:

$$t_{\text{grow}}(t) \propto v(t)^{-3/5} \left(\frac{m(t)}{m_0} \right)^{1-\alpha} \exp\left(-\left(\frac{5}{4} - \alpha\right) \frac{z(t)}{H}\right). \quad (4.39)$$

as a time-dependent quantity. The rate at which t_{grow} changes is:

$$\frac{d \ln t_{\text{grow}}}{dt} = \frac{\dot{t}_{\text{grow}}}{t_{\text{grow}}} = -\frac{3}{5} \frac{\dot{v}}{v} + \frac{(1-\alpha)}{t_{\text{grow}}} - \left(\frac{5}{4} - \alpha\right) \frac{v}{H} \quad (4.40)$$

From equation (4.6), this can be contrasted with the rate at which v evolves:

$$\frac{d \ln v}{dt} = \frac{\dot{v}}{v} = \frac{g}{v} - \frac{1}{t_{\text{grow}}} \quad (4.41)$$

We can make two observations. Firstly, equation (4.40) has terms of opposing sign. Thus, it is possible that $\dot{t}_{\text{grow}} \rightarrow 0$, i.e. $t_{\text{grow}} \approx \text{const}$, rather than evolving with background

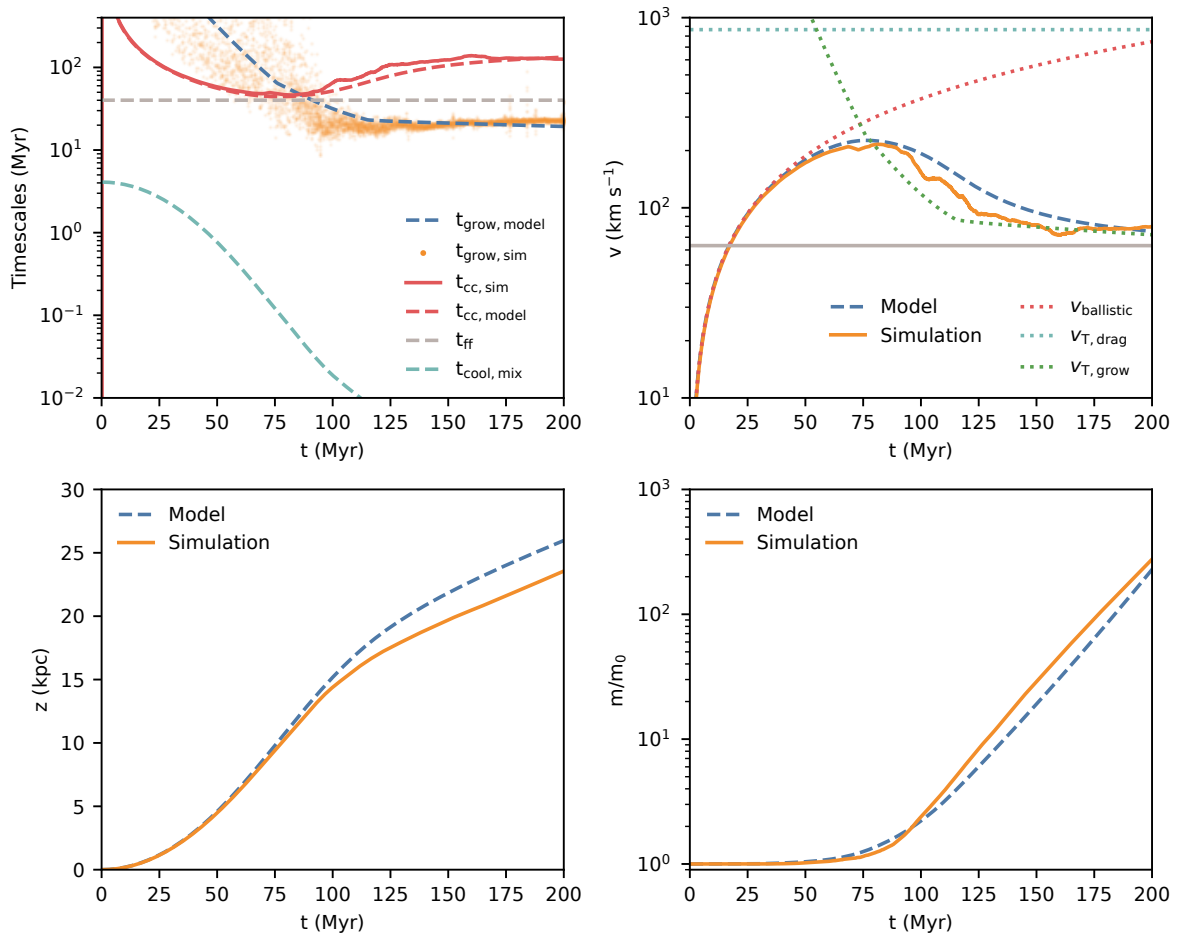


Figure 4.12: Time evolution of various quantities for a $r = 300$ pc cloud falling in a stratified background. From left to right, top to bottom, the panels compare the growth time t_{grow} , the velocity v , the distance fallen z , and the cold gas mass m of the cloud in the simulation versus the model. The upper panels also include comparison with other quantities of interest. Model predictions are in good agreement with simulations results.

quantities. Physically, this is because of a negative feedback loop. Suppose t_{grow} decreases as a cloud falls into denser surroundings. The subsequent increase in mass causes t_{grow} to increase (from equation (4.39)). The opposite is also true: if t_{grow} is large, the cloud will fall faster (due to weaker accretion drag) into denser regions, reducing t_{grow} . Secondly, by comparing terms on the right-hand side of equations (4.40) and (4.41), the timescale on which t_{grow} equilibrates to its steady-state value is comparable to the timescale on which v equilibrates to its steady-state value¹⁸, $v_{\text{T,grow}} = gt_{\text{grow}}$. Thus, $\dot{v}, \dot{t}_{\text{grow}} \rightarrow 0$ on similar timescales. From setting equations (4.40) and (4.41) to zero, the steady-state value of t_{grow} , and hence $v_{\text{T,grow}}$, is given by:

$$v_{\text{T,grow}} = gt_{\text{grow}} \approx \sqrt{\frac{1-\alpha}{\frac{5}{4}-\alpha}} Hg \approx \sqrt{\frac{2}{5}} c_{\text{s,hot}}, \quad (4.42)$$

where in the last step we use $\alpha = 5/6$ and $g \approx c_s^2/H$ for an isothermal atmosphere. This velocity is shown by the grey line in Fig 4.12. This then has the remarkable implication that in an isothermal atmosphere with constant gravity, $f_{\text{sub-vir}} = v_{\text{T}}/c_{\text{s,hot}} = t_{\text{grow}}/t_{\text{ff}}$ (equation (4.10)) of a cloud where accretion induced drag dominates is *independent* of all properties of the system except cloud geometry, specifically α . For our measured value of $\alpha = 5/6$ from infalling clouds with cometary tails, we predict $f_{\text{sub-vir}} = [(1-\alpha)/(5/4-\alpha)]^{1/2} \approx 0.6$. In Fig. 4.13, we compare velocity evolution in our model (equations (4.1) – (4.3)), to the asymptotic velocities from equation (4.42), for different cloud sizes and gravitational fields. Equation (4.42), which only depends on α , correctly predicts the asymptotic velocity. Note, however, that reaching the asymptotic velocity requires falling through many scale heights, and a planar $g \approx \text{const}$ isothermal atmosphere may not be realistic over such lengthscales. ‘Velocity overshoot’ also implies that large clouds (which exhibit stronger overshoot) might be seen to fall faster than predicted. In systems

¹⁸Indeed, because of ‘velocity overshoot’, t_{grow} equilibrates first.




Figure 4.13: Velocities in a stratified environment converge to a constant value that only depends on α (where area $A \propto M^\alpha$), independent of all other properties such as cloud size, or gravity. Curves show velocity evolution in our model (equations (4.1) – (4.3)), while solid lines give the asymptotic velocities from equation (4.42).




Figure 4.14: The mass growth rate as a function of cold gas mass for clouds of different initial sizes and different gravitational strengths. Curves are labelled by the initial parameters and whether they represent model solutions (M) or simulations (S), which are shown as dashed and solid lines respectively.

with varying $g(r)$ and $T(r)$ (and thus non-constant scale heights), the result can be more complex, and the most straightforward way to arrive at predictions is to simply integrate the set of ODEs, equations (4.1) – (4.3). We will show an example in Section 4.6.2.

4.5.3 Scaling With Cloud Size and Gravity

In Fig. 4.14, we compare the mass growth rates as a function of mass for simulations with varying initial cloud sizes and gravitational strengths to model predictions. Varying




Figure 4.15: The mass growth rate as a function of cold gas mass for a $r = 300$ pc cloud with $g = g_{\text{fid}}$ at different resolutions ($8\times$ higher and lower mass resolution than in the fiducial run, respectively). The simulations are relatively well converged.

g allow us to test the model for different scale heights. We can see that the model predictions are in good agreement with simulations results. In all cases, the simulations converge to the $1/t_{\text{grow}}$ slope predicted by the model. The divergence at early time is due to the fact that for this setup, the clouds start in a destruction regime since cooling is relatively weak.

4.5.4 Resolution Convergence

To test if our results for mass growth rates are converged. we run a $r = 300$ pc cloud with $g = g_{\text{fid}}$ at various resolutions, varying the fiducial resolution by a factor of 2. Fig. 4.15 shows that the three resolutions show little difference in mass growth rates and that the simulation appears to be converged, although the higher resolution simulation matches the model slightly better – the cloud is disrupted less initially and reaches the model growth rate more rapidly.

4.5.5 Survival in a stratified background

For a cloud falling in a constant background we found that the survival criterion was given by a competition between the growth and destruction timescales of the cloud: $t_{\text{grow}} < f_S t_{\text{cc}}$ where f_S is a constant factor. We wish to ascertain if the same condition applies to clouds falling in a stratified background.

In the case of a constant background, t_{grow} changes very little over time (once turbulence has developed), with only a very weak scaling with mass, and cooling is strong enough so v approaches gt_{grow} without ‘overshooting’, something we noted in Fig. 4.12 above. For a stratified background, both these assumptions do not hold - t_{grow} changes continuously with background density, and an overshoot is often observed. Since our initial conditions are in the regime where clouds do not survive, surviving clouds are those that are able to survive long enough to enter the growth zone.

One ansatz would be to use the asymptotic value of t_{grow} and v that we derived above in equation (4.42) and evaluate the survival criteria there. This gives:

$$r > \frac{v_{\text{T,grow}}^2}{gf_S\sqrt{\chi}} \quad (4.43)$$




Figure 4.16: The fate of clouds of different size falling in stratified backgrounds with different gravitational strengths. Survival criterion evaluated at different times are shown. The best survival criterion is given by the teal curve, i.e. equation (4.27) evaluated at the maximum velocity, for $f_s = 3$.

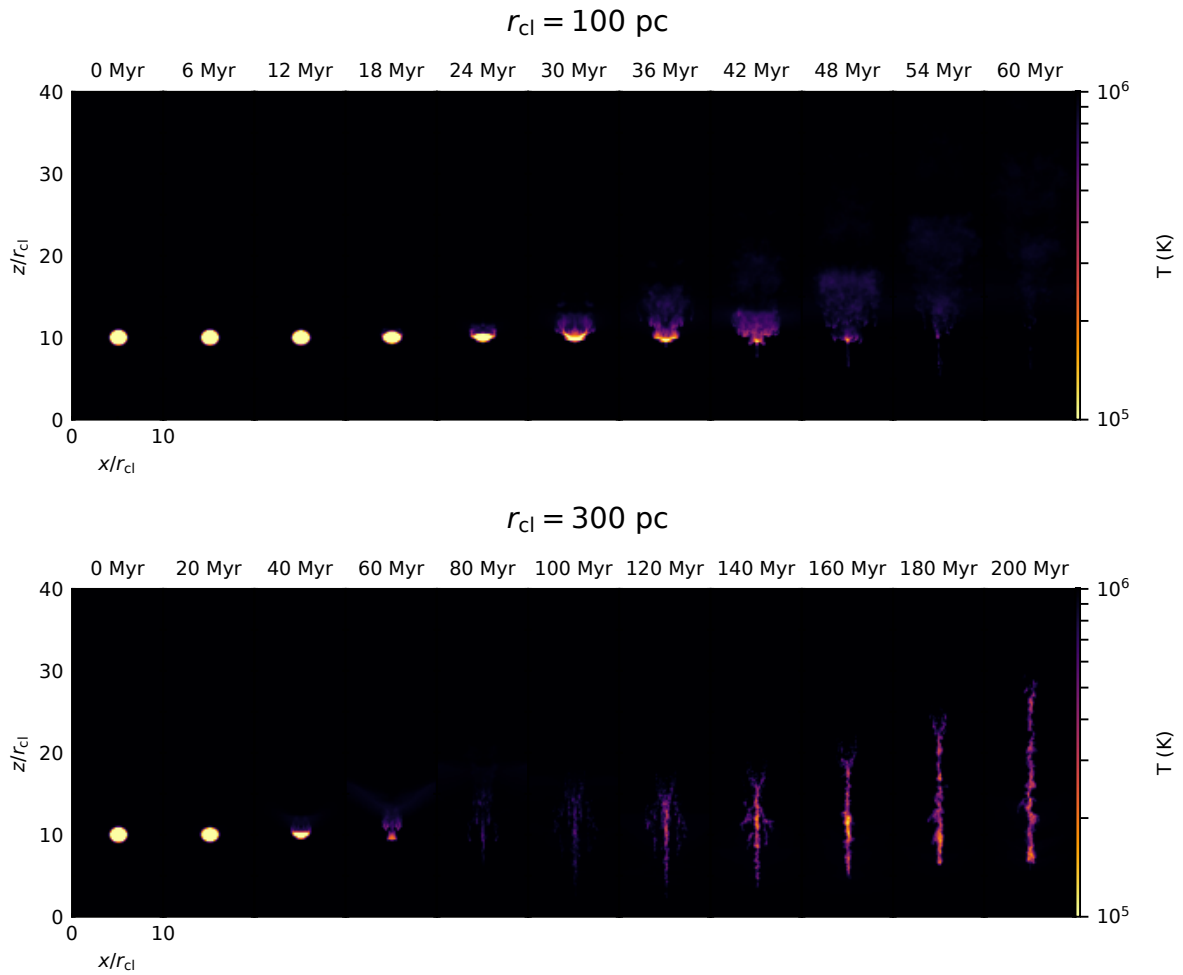


Figure 4.17: Snapshots of the projected density weighted temperature through the box (temperature here is hence just an indication of the amount of cold gas when projected along the y-axis) for a 100 pc and a 300 pc cloud at various points in their evolution. The former is disrupted completely while the latter reaches the survival zone and grows. x and z here simply reflect the size of the box along the respective axes normalized by cloud size.

This condition is given by the blue dashed line in Fig. 4.16. Note that it is a *lower* bound on r , since $v_{T,\text{grow}}$ is independent of r . It has the right qualitative behavior as a survival criterion, but does not seem to match the survival thresholds seen in the simulations. Clouds have to fall many scale heights to reach the asymptotic velocity given by equation (4.42) – often survival is determined much earlier. Indeed, the falling clouds often overshoot this asymptotic velocity as they initially fall ballistically, as seen in Fig. 4.16. We can estimate the time where gravity and cooling balance:

$$\dot{m}v \sim \frac{m}{t_{\text{grow}}}v \sim mg\frac{t}{t_{\text{grow}}} \quad (4.44)$$

assuming the cloud is falling ballistically in this initial phase. Hence, $t \sim t_{\text{grow}}$ is the time where the cloud is slowed from its ballistic free falling trajectory. If we evaluate equation (4.27) at this time in the simulation, we can solve numerically for some r_{crit} . Of course, this only makes sense if $v(t = t_{\text{grow}}) > v_{T,\text{grow}}$, i.e. there is an overshoot so t_{cc} is shorter. The larger the difference in the two velocities, the more likely the cloud is to be destroyed in this overshooting phase. In Fig. 4.16, we show this limit in the orange dashed line. We see that this matches the simulation results more closely for larger values of g , where the clouds accelerate to higher velocities. Ultimately, it is the *maximum* velocity that determines if a cloud survives. We thus show in the red and teal curves in Fig. 4.16 the survival criterion evaluated at $v = v_{\text{max}}$ from the model. The red curve use $f_{\text{S}} = 4$ as in the previous section, while the teal curve has $f_{\text{S}} = 3$, which seems to be a better match to the simulation results. It is unsurprising that we find a different value of f_{S} here, since we are evaluating our quantities at a different time.

In Fig. 4.17, we show the evolution of the 100 pc and the 300 pc cloud for $g = g_{\text{fid}}$. The 100 pc cloud does not survive and is disrupted completely, while the 300 pc starts to get disrupted but survives long enough to reach the zone of growth and then grows.

Note the tail growth in the surviving case. To summarize, we have looked at clouds that start outside the growth zone in a stratified medium, and find that in order to survive, the cloud has to make it to the growth zone. Since the cloud is accelerating ballistically before it reaches high enough pressures where cooling is efficient enough for it to grow and slow down, only large clouds can survive this infall. We explore the implications of the survival conditions in this and the previous section on astrophysical systems of interest in the following section.

4.6 Discussion

4.6.1 High Velocity Clouds

3D simulations of clouds falling under gravity with mixing and cooling processes included have only been studied to a limited extent previously. Heitsch & Putman (2009) concluded that clouds below $10^{4.5} M_{\odot}$ are disrupted within 10 kpc. Notable differences in setup include a smaller box length along the tail direction and starting initially with colder clouds, as their temperature range extended down to 100 K. Heitsch et al. (2022) focused on metallicity measurements, tracing original versus accreted cloud material. They found that most of the original cloud material does not survive and is instead replaced by accreted gas which mostly happens in the tail. Grønnow et al. (2022) observed cloud growth in MHD simulations but did not follow the clouds for many cloud crushing times. We have followed up by providing a model for the mass growth of such clouds based on the underlying process of turbulent mixing and cooling, so as to tackle the key questions of when HVCs can survive, how much mass they accrete, and how fast they travel. We then tested the model against a suite of numerical simulations. What then are the implications for HVCs?




Figure 4.18: Survivability of HVCs in Milky Way conditions with a size r_{cl} and dropping height d_0 . For clouds outside the survival region given by equation (4.34), the color coding corresponds to the ratio t_{grow}/t_{cc} evaluated at the maximum velocity along the cloud trajectory. The horizontal white dashed line shows where the survival criterion equation (4.27) is satisfied for $f_S = 4$. Large clouds that fall in from large distances can still survive as they are not destroyed before reaching the survival region.

In Fig. 4.18, we show our estimates for cloud survival in a Milky Way like profile in the cloud size-initial height parameter space. Specifically, we employ the profiles from Salem et al. (2015) who combine the density profile of Miller & Bregman (2015) with a temperature profile mapped from a NFW halo (Navarro et al. 1997), which we also use to set the gravitational profile. In the region of interest, $T \sim 10^6$ K. Figure 4.18 shows the ratio of the growth time and the cloud crushing time $t_{\text{grow}}/t_{\text{cc}}$ evaluated at the maximum velocity the cloud reaches along its trajectory. We also show the threshold of survival (equation (4.27)) at ~ 4 from the previous section. The analytic expectation (equation (4.34)) for where cooling is strong enough for clouds to survive regardless is demarcated by the white dashed line. Outside this region, larger clouds can survive falling from further out, simply from the fact that $t_{\text{cc}} \propto r_{\text{cl}}$.

More generally however, Fig. 4.18 shows that except for these larger ($\gtrsim 100$ pc) clouds, HVCs in the Milky Way should only survive if they start at an initial height of $d_0 \lesssim 10$ kpc. Most HVC complexes detected do indeed fall within this regime – with the notable detection of the ones associated with the LMC and its Leading Arm located at $\gtrsim 20$ kpc (Richter et al. 2017).

While this prediction seems to explain the observed survival of most HVCs, we want to highlight that due to the mass transfer from the hot to the cold medium, most surviving clouds in the Milky Way in our model would fall at $v_{\text{GSR}} \sim 70$ km/s (equation (4.42)) and might thus have velocities v_{LSR} which are too low to be classified as HVCs. Such a population of intermediate to low velocity clouds is of course to be expected even from simply studying the velocity distribution of HVCs and “filling in” the gap at $v_{\text{LSR}} \sim 0$, and has been the subject of several theoretical studies (e.g. Peek et al. 2007; Zheng et al. 2020) – as well as observational attempts to locate them (e.g. Peek et al. 2009; Bish et al. 2021). Thus far, there does not seem to be a firm conclusion on the existence of such a low-velocity population. Our work provides a theoretical foundation for the existence

of such clouds and predicts an overabundance of them in the Milky Way halo at lower heights ($\lesssim 10$ kpc).

An interesting example of a nearby HVC is the Smith Cloud (Smith 1963), lying only 3 kpc below the galactic plane with a metallicity of $\sim 0.5 M_{\odot}$, and which is falling towards the galactic plane at velocity $v_z \sim 70$ km/s (Fox et al. 2016). A longstanding mystery has been explaining the survival of the Smith Cloud at its current location. A simple ballistic analysis suggests that the cloud might have already passed through the disk (Lockman et al. 2008) and should hence have been disrupted, in which case some mechanism is needed to explain its survival, such as the cloud being embedded in a dark matter sub-halo, which would shield the gas and extend its lifetime (Nichols & Bland-Hawthorn 2009). It is possible that the relative high metallicity and survival of the Smith Cloud can be potentially explained instead by accretion of ambient material driven by turbulent mixing and cooling. Henley et al. (2017) ran a wind tunnel simulation with the aim of reproducing a Smith cloud like setup, and found entrainment of the background gas largely in the tail of the cloud. Galyardt & Shelton (2016) ran simulations of the Smith Cloud with gravity and in a stratified background. They concluded that if the Smith Cloud was in a dark matter sub-halo, it would comprise gas accreted only after the sub-halo passed through the disk. Alternatively, if the Smith Cloud was not accompanied by such a sub-halo, then it must be on first approach, since the cloud would not survive its journey through the Galactic disk. Our model could naturally explain the survival of a Smith Cloud that was on first approach, as it fulfills the survival criterion Eq. (4.27), i.e., it falls within the ‘survival zone’ of the Milky Way’s halo. The trajectory in this case would be very different from the ballistic one since the accretion dynamically affects the cloud.

Since the terminal velocity is independent of the cloud size, one would expect no observable relationship between, for instance, cloud column density and infall velocity,

although there may be significant scatter since this requires the cloud velocity to ‘turn around’ and reach asymptotic terminal velocity. This is consistent with observations (Westmeier 2018).

We have thus far considered clouds that are infalling from large distances and potentially feed the disk. In our model, HVCs and IVCs can continually grow in mass once they are near enough to the disk. It therefore also gives credence to the notion that fountain-driven accretion can supply the disk with fuel for star formation: cold gas thrown up into the halo ‘comes back with interest’, by mixing with low metallicity halo gas which cools and increases the cold gas mass (Armillotta et al. 2016; Fraternali 2017). Such low metallicity gas is required to satisfy constraints from disk stellar metallicities and chemical evolution models (Schönrich & Binney 2009; Kubryk et al. 2013). The equations for mass transfer and velocity derived in this work can also be incorporated into semi-analytic ‘fountain flow’ models and checked against observations.

4.6.2 Clusters

Galaxy clusters are amongst the largest virial systems in the universe and thus present opportune test beds for the comparison of observations and theoretical models of galactic properties and evolution. The hot intracluster medium (ICM) in such environments reaches temperatures in the range of 10^7 – 10^9 K which can be probed observationally via X-ray emission originating from the thermal bremsstrahlung radiation of this hot diffuse plasma (Sarazin 1986). However, the ICM does not exist simply in a single phase. Observations from measurements of carbon monoxide (CO) which traces cold molecular gas find an abundance in these central cluster galaxies, with molecular gas mass correlating with X-ray gas mass (Pulido et al. 2018). One theory for the origin of the cold molecular gas is that they develop from thermal instabilities triggered in

the wakes of cooling updrafts of radio bubbles that rise and lift low entropy X-ray gas (McNamara et al. 2016). These form the cold filaments observed to trace the streamlines around and behind the bubbles, which should eventually decouple from the velocity structure of the hot flow and fall back towards the galaxy center (Russell et al. 2019).

A particularly interesting conundrum is the low observed velocities of the molecular gas measured by CO line emission in ALMA target systems (McNamara et al. 2014; Russell et al. 2016; Pulido et al. 2018; Russell et al. 2019). They are significantly smaller ($< 100 \text{ km s}^{-1}$) than both stellar velocity dispersions ($200\text{--}300 \text{ km s}^{-1}$) and galaxy escape velocities ($\sim 1000 \text{ km s}^{-1}$), implying that the molecular gas is tightly bound to the galaxy and should be expected to be infalling. Even initially outflowing gas should at some point stall and fall back inwards. These low velocities are puzzling as models of free falling clouds in cluster potentials have estimated that they can be accelerated to hundreds of kilometers per second after falling just a few kpc (Lim et al. 2008; Russell et al. 2016). The large density contrast between the molecular gas and the hot background in the ICM means that ram pressure should do little to slow down these falling clouds, which would rapidly accelerate to high velocities. Small velocities would require the observed cold gas to have been falling gravitationally for only a short amount of time. While this could be explained if the infalling cold gas observed was mostly recently decoupled from the hot gas, there is no reason to suggest that this should be the case. Furthermore, the rapid acceleration means we should see steep velocity gradients in these filaments. However, we observe shallower velocity gradients that are inconsistent with free-fall (Russell et al. 2016, 2017). Some observations find that free-fall models can match observations in outer filaments, but break down for inner regions (Lim et al. 2008; Vantyghem et al. 2016). One caveat here is that increasing the spatial resolution of observations can reveal more complex spatial and velocity structures (Lim et al. 2008). Lastly, if the molecular gas was free-falling, we would expect to generally detect higher velocities at smaller radii, but

there is no evidence for this. A large influx of cold gas implies that circumnuclear disks should be more common in comparison with filaments, while the opposite is observed (Russell et al. 2019).

The conclusion then is that the picture of free falling clouds fails to explain a large number of observations with regard to these filaments, which suggest that the infalling cold gas has to be slowed by some alternative process other than ram pressure drag. One possibility which has been previously proposed is that magnetic stresses slow the clouds' descent, since it has been suggested that the cold filaments are significantly magnetically supported (Fabian et al. 2008). However, the magnetic pressure that would be required to slow such a filament's infall along its length requires a strong non-radial magnetic field component with $\beta \sim 0.1$ (Russell et al. 2016).

Our results suggest an alternative explanation that naturally addresses the above issues. As noted above, observations of the prevalence of molecular gas are closely tied to systems with shorter cooling times. As shown in the previous section, the filamentary mass growth driven by turbulent mixing and cooling of these infalling cold filaments serve as a braking mechanism via accretion induced drag. This can significantly reduce the acceleration of the cold gas when the cooling time of mixed gas is short. To illustrate this point, we compare our model to the free fall model used in Lim et al. (2008) in their analysis of observed filaments in the cD galaxy NGC 1275 (Perseus A) located in the Perseus cluster. For simplicity, we follow the approach of Lim et al. (2008) and adopt an analytic model of the mass density and gravitational potential of the form from Hernquist (1990). The mass density and gravitational potential as a function of radial distance are

thus given by:

$$\rho(r) = \frac{M a}{2\pi r} \frac{1}{(r+a)^3} \quad (4.45)$$

$$\phi(r) = -\frac{GM}{r+a} \quad (4.46)$$

where M is the total galactic mass, r is the radial distance, and a is a scale length. We also use the same values they deduce from luminosity observations of Smith et al. (1990) and an estimated mass to light ratio, with $M = 8.3 \times 10^{11} M_{\odot}$ and $a = 6.8$ kpc. We use the number density profile given in Churazov et al. (2003) for the Perseus cluster, which is mostly a constant $n = 4 \times 10^{-2} \text{ cm}^{-3}$ below 30 kpc and adopt a constant temperature profile of $T = 10^7$ K.

Figure 4.19 shows the observational contours of velocity as a function of radial distance from the center of Per A for the outer western filament as shown in Figure 10 of Lim et al. (2008). In Fig.4.19, we have also reproduced the free-fall trajectories used in Lim et al. (2008), where they include one for galactic masses of $M = 8.3 \times 10^{11} M_{\odot}$ (M8b) and $M = 3.4 \times 10^{11} M_{\odot}$ (M3b), both starting from a radius of 8.5 kpc. The free-fall model that assumes the $M = 8.3 \times 10^{11} M_{\odot}$ mass deduced from luminosity observations is unable to produce a good fit to observations, and hence the mass needs to be tuned to $M = 3.4 \times 10^{11} M_{\odot}$ to fit a free-fall model to the observed contours. This tuning of mass and drop height is sensitive to both these factors, mainly due to the rapid acceleration by gravity in free-fall. In comparison, we show the same curve for $M = 8.3 \times 10^{11} M_{\odot}$ but using our model (M8c) (equations (4.1)-(4.3)) which includes the braking effect due to growth from mixing and cooling. This shows the trajectory for a cloud where $r_{\text{cl}} = 50$ pc, assuming that $L/r \sim 100$. We see that if the cloud initially falls from even a radial distance of 15 kpc, it matches the observations well without changing the galaxy mass. Clouds can thus fall from a further distance out than observed. It should be noted that




Figure 4.19: Observed velocity contours of the outer western filament in Per A from Lim et al. (2008) are shown in purple. Ballistic trajectories are shown by the blue and orange lines for galactic masses of $M = 8.3 \times 10^{11} M_{\odot}$ (as observed) and $M = 3.4 \times 10^{11} M_{\odot}$ (tuned to obtain the correct infall velocities) respectively. The red line shows the trajectory of a cloud in our model with a galactic mass of $M = 8.3 \times 10^{11} M_{\odot}$ but which is experiencing accretion drag. In the latter case, tuning of galactic mass is not necessary to explain observations.

the conditions here are on the boundary of the survival criterion from equation (4.30), due to its strong scaling with χ .

In Fig. 4.20, we show velocity trajectories for clouds dropped from various heights given by solid lines, with the 8 kpc distance used as a lower bound. We find that even clouds that are dropped much further away do not accelerate as rapidly to high velocities as in the ballistic case. Ballistic trajectories for clouds dropped from the same heights are shown for comparison by the dashed lines, and can be seen to rapidly accelerate past observed velocities. On the other hand, the clouds in our model are slowed and stay within the range of observed velocities for much longer times. Hence, we are much less sensitive to the exact distance at which the cold gas first begins to fall. Our results are consistent with the lower velocities and shallower velocity gradients observed relative to what would be expected from free-fall without requiring that the observed cold gas had only just recently cooled, or that magnetic drag from a strongly magnetized background must be present. In addition, the survival of cold gas and the filamentary morphology can also be naturally explained by cooling tails.

4.6.3 Other Implications

We have found that it is more difficult for infalling cold material to survive, compared to their outflowing counterparts, which are eventually entrained and do not experience further shear forces thereafter. This conclusion has a range of wider implications which we will now touch on.

Assuming isobaric conditions, our survival criterion is equivalent to $t_{\text{cool,hot}}/t_{\text{ff}} \lesssim 1$ (equation (4.35)), which is equivalent to the criterion for linear thermal instability in a plane parallel atmosphere. As previously remarked, this has the interesting implication that cold gas which forms via thermal instability should be able to survive infall, though




Figure 4.20: Observed velocity contours of the outer western filament in Per A from Lim et al. (2008) are shown in purple. Trajectories starting from different initial heights in our model are shown by the solid lines. Dashed lines show ballistic trajectories with the same starting point. The velocities we predict are much less sensitive to the initial drop height compared to the ballistic model.

this should be re-examined in a spherical potential, where the threshold for cold gas condensation changes, and t_{ff} (and gravitational acceleration) vary with radius. This is an interesting avenue for future work.

Our results imply that clouds which grow in mass when they fall should undergo accretion-induced braking, a prediction which can be tested in larger scale simulations with more realistic set-ups. Nelson et al. (2020) find an abundance of cold clouds of sizes 1 kpc and smaller around the CGM of ‘red and dead’ intermediate redshift elliptical galaxies in the TNG50 simulation. These clouds are mostly infalling, with the radial velocity distribution peaking at just one third of the virial velocity. They also find that the clouds are accreting and growing. They are long lived, surviving for cosmological timescales. This appears to be consistent with predictions from our model - that infalling cold clouds are growing and thus slowed to sub-virial velocities. It would be interesting to directly compare growth times t_{grow} , and infall velocities, to see if the expectation $v_{\text{T}} \sim gt_{\text{grow}}$ is fulfilled.

Similarly, our results will affect the dynamics of cosmic cold streams feeding galaxies at high- z (Dekel & Birnboim 2006; Kereš et al. 2005). Thus far, Mandelker et al. (2020) has found that the survival criterion for cold clouds seems to be able to translate relatively well to stream survival¹⁹. However, in their studies they used a constant hot gas velocity – similar to the outflowing cloud simulations – which implies that their shear declines rapidly in the simulation. Since in reality cosmic streams are also accelerated by gravity, the situation for streams is likely closer to the setup studied here. This would imply that (a) an equivalently more stringent survival criterion would apply to streams, and (b) their terminal velocity is given by $\sim gt_{\text{grow}}$. Indeed, unlike idealized simulations, cosmological simulations find that streams reach a roughly constant terminal velocity (Dekel et al.

¹⁹There is some controversy regarding the destruction timescale but for the relevant $\chi \sim 100 - 1000$ the different possibilities agree (cf. discussion in Bustard & Gronke 2022)

2009; Goerdt & Ceverino 2015); a result which has not been quantitatively explained. These implications directly affect the cold gas mass supplied towards the inner galaxies in dark matter halos.

Interestingly, coronal rain in our Sun is also observed to fall below free-fall velocities – on average falling with only $\sim 1/3 - 1/2$ of the ballistic value (see review by Antolin & Froment 2022). While the temperatures and resulting overdensities are for coronal rain similar to what has mostly been considered here, the main difference is the strong magnetic field. Thus, most studies within the solar community have focused on magnetic fields as explanation of the slowdown and it has in fact been shown (using mostly one and two-dimensional simulations) that coronal rain can be efficiently decelerated due to a buildup of pressure in front of the cold cloud (Oliver et al. 2014; Martínez-Gómez et al. 2020). Clearly, the magnetic fields do play a major role here and affect the dynamics. However, it is noteworthy that mass transfer can also lead to significant slowdown. Plugging typical values found observationally ($n \sim 10^{11} \text{ cm}^{-3}$, $T \sim 2 \times 10^4$, $r \sim 1 \text{ Mm}$, $\chi \sim 100$, $g = 274 \text{ m s}^{-2}$; Antolin & Froment 2022) into Eq. (4.25) yields $v_{\text{term,drag}}/v_{\text{term,grow}} \sim 0.45$. Thus, the ‘accretion braking’ process described in this work might be another important drag force at play; an interesting avenue for future work.

4.6.4 Further Considerations

While the model we have presented explores and captures the core physics at play, simplifications and assumptions have been made along the way. We discuss several considerations which could provide interesting avenues in order to expand and build on this model.

Additional Physics

There are a number of physical processes whose impact and importance we have not touched on in this work, but which could lead to complications and should be studied in future work. One such component is magnetic fields. Magnetic fields have been shown to significantly affect the morphology of clouds in both the wind tunnel and falling cloud setups, while their effect on mass growth is still uncertain (Grønnow et al. 2017, 2018; Gronke & Oh 2020a; Grønnow et al. 2022; Abruzzo et al. 2022b). For example, magnetic fields can suppress the KH instability, reducing mass entrainment rates (Ji et al. 2019; Grønnow et al. 2022), although mass growth rates in some full cloud simulations appear minimally impacted (Gronke & Oh 2020a). Another source of non-thermal physics that could be important to study in this context is cosmic rays (Huang et al. 2022; Armillotta et al. 2022). Self-gravity has been found to matter for compact HVCs (Sander & Hensler 2021). We have also not included explicit viscosity and thermal conduction (although we point out that for turbulent mixing layers the mass transfer is generally dominated by turbulent diffusion, Tan et al. 2021).

Initial Cloud Morphology

There is some uncertainty regarding an appropriate choice for the initial structure of the cloud. A spherical cloud is clearly an idealized choice. Instead of a uniform density sphere, smoothly varying density and temperature profiles connecting the two phases have been used for more realism (Heitsch & Putman 2009; Kwak et al. 2011; Gritton et al. 2014; Sander & Hensler 2021). Furthermore, Cooper et al. (2009) found that fractal clouds were destroyed faster as compared to uniform spheres due to more rapid cloud breakup. Schneider & Robertson (2017) similarly found that an initially turbulent structure within the cloud would enhance fragmentation and ultimately facilitate cloud

destruction. However, the above are all concerned with cloud destruction, where the clouds are in a regime where the cloud is ultimately destroyed over time ($t_{\text{cool,mix}} > t_{\text{cc}}$ for wind tunnel setups). The importance of the initial cloud structure can thus be understood in the context that it determines how the cloud is destroyed as it fragments and breaks up. However, if we are in the regime where one is concerned about cloud growth instead, then this dependence on the initial setup seems to matter less. Gronke & Oh (2020a) found that in the regime of cloud growth, there was little difference in either the mass growth or velocity evolution between an initially turbulent or uniform cloud. In fact, the turbulent case actually grew slightly faster, since it had a larger surface area at the start. Still, this suggests that the initial morphological evolution of the sphere does have some dependence on the choice of the initial structure of the cold gas cloud. In terms of numerical values, this creates some amount of uncertainty in our model, in particular with regards to the initial values of the cloud surface area and its initial evolution, which Heitsch et al. (2022) refers to as the ‘burn-in’ phase. In our model, this uncertainty is folded in by calibrating a constant prefactor of order unity to the results from our simulations. It is possible that the precise value of this factor might vary depending on setup and initial cloud structure.

Temperature Floor and Self Shielding

In our simulations, we have assumed a temperature floor of $T \sim 10^4$ K. However, it would be useful to understand the phase structure of cold neutral gas that provides an additional layer of structure to the clouds (Girichidis et al. 2021; Farber & Gronke 2022) and how this might impact cloud growth and dynamics. This is especially for comparison with observations, which often detect warm gas surrounding cold cores. On a related note, we have assumed that all our clouds are optically thin. However, self-shielding could be important for the more massive clouds.

Infall Conditions

We have assumed our clouds fall directly towards the disk. However it is likely that most clouds will have some sort of rotational velocity component and hence fall inwards on some orbit trajectory. As mentioned in Heitsch & Putman (2009), this component is more akin to the wind tunnel setups since net acceleration is reduced. We have also assumed a quiescent background - realistic environments are likely subject to large scale turbulence (Gronke et al. 2022). This could affect mixing rates or significantly lengthen infall times and introduce a large stochastic variability in the infall velocity, much the same way a leaf falling to the ground in a windy environment follows a much longer trajectory. How this might affect cloud growth and dynamics is a natural follow up to this work.

Metallicity

We have assumed solar metallicity everywhere in both phases. Depending on the origin of the cold cloud, it is possible that the metallicity of the original cloud and the background differ significantly. Gritton et al. (2014) and Heitsch et al. (2022) have showed that there is significant mixing of metals in such a case, with important implications for observables.

4.7 Conclusions

The growth and survival of infalling cold clouds has received considerably less attention compared to their outflowing counterparts. While the two appear to be similar problems at first glance, they have in fact a crucial difference between them, which is that infalling clouds continuously feel the force of gravity. This leads to very different dynamical evolution of the infalling clouds, and also a more stringent criterion for sur-

vival. Using 3D hydrodynamical simulations, we have studied the growth and survival of such clouds, considering both a constant background as well as a more realistic stratified background. We have also developed a model for the dynamical evolution of these clouds based on turbulent mixing layer theory, and shown that they are able to predict the results of the simulations. These also agree well with analytical estimates. Our main findings are:

- *Not a Wind Tunnel:* Infalling clouds do not correspond to wind tunnel setups, where the velocity shear is initially large and decreases as the cloud gets entrained. Instead, the velocity shear is initially small but increases as the cloud accelerates. This means that criteria such as $t_{\text{cool,mix}} < t_{\text{cc}}$ for survival are not applicable.
- *Modelling Cloud Growth:* An important component determining how fast the cloud grows is the surface area of the cloud. We find that $A \propto m^{5/6}$. This is consistent with either a mix between surface and tail growth or a fractal surface area. Combining this with models of the inflow velocity allow us to model the growth time of the clouds, as given in equations (4.22) and (4.23). We can hence evolve equations (4.1) – (4.3) to model the evolution of cloud properties.
- *Accretion Drag:* Clouds falling due to gravity can experience an alternative form of drag if they are growing via turbulent radiative mixing layers, since they are effectively accreting low momentum gas. This drag is dominant over the usual ram pressure drag as the clouds develop long tails along the direction of infall. This leads to much lower predicted infall velocities compared to models which only consider ballistic trajectories or ram pressure drag. In particular, the terminal velocity $v_{\text{T}} \approx gt_{\text{grow}}$, where $t_{\text{grow}} = m/\dot{m}$ is given by equation (4.25) for subsonic infall.

- *Relationship between Speed and Growth Rates:* The balance between gravity and growth results in $v_T/c_{s,\text{hot}} \sim t_{\text{grow}}/t_{\text{ff}}$. That is, the ratio of the terminal velocity and the virial velocity is also the ratio of the the growth time to the free-fall time. This is useful since infall velocities are measured observationally. The growth rate of the cloud can then be deduced. We expect sub-virial velocities ($v_T < c_{s,\text{hot}}$) to be indicative of considerable mass growth ($t_{\text{grow}} < t_{\text{ff}}$) in clouds. Observed sub-virial infall velocities are otherwise difficult to explain with existing models. In an isothermal atmosphere with constant gravity, we predict $v_T \approx 0.6c_{s,\text{hot}}$, independent of all other properties of the system, although convergence to this asymptotic velocity can be slow.
- *Criterion for Cloud Survival:* The criterion for clouds to survive and grow is $t_{\text{grow}} < 4t_{\text{cc}}$ (equation (4.27)). The most important factor in determining cloud survival is the cooling time. We find that the ratio of $t_{\text{grow}}/t_{\text{cc}}$ is almost independent of cloud size (within a large practical range of parameter space). Hence, in order to survive and grow, clouds need only be within regions where densities/pressure are high enough such that cooling times are sufficiently short. For $\chi = 100$, this criterion can be written as

$$P > 3000 \text{ k}_B\text{K cm}^{-3} \left(\frac{g}{10^{-8} \text{ cm s}^{-2}} \right)^{4/5} \quad (4.47)$$

- *Stratified Backgrounds and Cloud Size:* In stratified environments, clouds that start their infall beyond such survival ‘zones’ can still survive provided they are not completely destroyed before reaching these zones. This favors larger clouds which have longer cloud crushing times. Larger clouds are hence more likely to be observed at distances where the above criterion is not satisfied.

In summary, we have identified a new mechanism for the deceleration of clouds that has not been considered in existing models, with important bearings on cloud survival, growth, and dynamics. We have presented a model for cloud growth (equations (4.1) – (4.3)), evolution (equations (4.22) and (4.23)), and survival (equation (4.27)) that agree well with simulations. These results can be applied to range of systems with infalling cold gas such as HVCs and clusters, and addresses important questions of survival, growth, and sub-virial velocities that have been highlighted by observations. Future work will refine this model with additional physics such as magnetic fields, cosmic rays and self-shielding, as well as allowing the gas to cool down to lower temperatures.

Acknowledgements

We thank Greg Bryan, Yakov Faerman, Drummond Fielding and Yong Zheng for helpful discussions. We have made use of the yt astrophysics analysis software suite (Turk et al. 2011), matplotlib (Hunter 2007), numpy (Van Der Walt et al. 2011), and scipy (Virtanen et al. 2020) whose communities we thank for continued development and support. We acknowledge support from NASA grants NNX17AK58G, 19-ATP19-0205, HST-AR-15797.001-A, NSF grant AST-1911198 and XSEDE grant TG-AST180036. MG thanks the Max Planck Society for support through the Max Planck Research Group. This research was supported in part by the National Science Foundation under Grant No. NSF PHY-1748958 to KITP.

Chapter 5

Cloud Atlas: Navigating the Multiphase Landscape of Tempestuous Galactic Winds

*My life amounts to no more than one drop in
a limitless ocean. Yet what is any ocean, but
a multitude of drops?*

David Mitchell, Cloud Atlas

Galaxies comprise intricate networks of interdependent processes which together govern their evolution. Central among these are the multiplicity of feedback channels, which remain incompletely understood. One outstanding problem is the understanding and modeling of the multiphase nature of galactic winds, which play a crucial role in galaxy formation and evolution. We present the results of three dimensional magnetohydrodynamical tall box interstellar medium patch simulations with clustered supernova driven outflows. Fragmentation of the interstellar medium during superbubble breakout seeds

the resulting hot outflow with a population of cool clouds. We focus on analyzing and modeling the origin and properties of these clouds. Their presence induces large scale turbulence, which in turn leads to complex cloud morphologies. Cloud sizes are well described by a power law distribution and mass growth rates can be modelled using turbulent radiative mixing layer theory. Turbulence provides significant pressure support in the clouds, while magnetic fields only play a minor role. We conclude that many of the physical insights and analytic scalings derived from idealized small scale simulations translate well to larger scale, more realistic turbulent magnetized winds, thus paving a path towards their necessary yet challenging inclusion in global-scale galaxy models.

5.1 Introduction

Galaxies are complex ecosystems—extracting the most out of our ever growing body of increasingly detailed observations requires understanding the inherently nonlinear, dynamical structures that underlie them. A multitude of open problems still surround galaxy formation and evolution today. One that sits at the very heart is the challenge to understand the multiphase nature of not just the gas within galaxies and their surrounding environment (circumgalactic medium (CGM)), but also the flow of this material in and out of galaxies and how this regulates galaxy evolution (often dubbed the Cosmic Baryon Cycle). Galactic outflows driven by feedback mechanisms carry material outwards while inflowing gas provides fuel for new star formation. This cycling connects processes on stellar ($\sim \text{pc}$) scales to galactic ($\sim \text{kpc}$) and cosmological ($\sim \text{Mpc}$) scales, weaving them into a tightly interdependent and highly multiscale tapestry.

Galactic winds, which play a central role in regulating the baryon cycle, have garnered significant attention in recent years due to the clear evidence of their influence on galaxy evolution. They are thought to be primarily driven by feedback processes originating

from either massive stars (for lower mass haloes) or active galactic nuclei (AGN; at the higher end of the mass spectrum), and are necessary ingredients for any realistic model (see Somerville & Davé 2015; Naab & Ostriker 2017, for detailed reviews). The result is the formation of large-scale outflows that transport material out of the galaxy while also shaping the CGM (Tumlinson et al. 2017).

Observations of galactic winds reveal them to be cosmically ubiquitous across star forming galaxies (Martin 1999; Rubin et al. 2014). They also exhibit a complex, multiphase structure consisting of cold ($T \lesssim 10^2$ K), cool ($T \sim 10^4$ K), warm ($T \sim 10^5$ K), and hot ($T \gtrsim 10^6$ K) gas components, with the various phases spanning a range of different properties and dynamics (Veilleux et al. 2005; Strickland & Heckman 2009; Steidel et al. 2010; Rubin et al. 2014; Heckman & Thompson 2017; Bolatto et al. 2021). A better understanding of the formation, survival, and growth of cool gas clouds which make up the most readily observable phase, as well as their interactions with the other gas phases, is needed in order to provide vital insights into the mass transport mechanisms and overall energetics of these outflows. This in turn is essential for constructing a comprehensive picture of the multiphase nature of galactic winds and their impact on galaxy evolution.

In parallel, theoretical modeling and simulations of galactic winds have become increasingly sophisticated, aiming to reproduce and understand the multiphase nature and complex dynamics observed in these outflows. Chevalier & Clegg (1985) introduced an early analytic model where mass and energy injection into a spherically symmetric region powered a hot outflowing wind. Building on this simple model, recent theoretical work has included radiative cooling and gravity (Thompson et al. 2016), non-spherical expansion (Nguyen & Thompson 2021), more realistic injection (Bustard et al. 2016; Nguyen et al. 2023), and frameworks for coupled multiphase evolution (Huang et al. 2020; Fielding & Bryan 2022).

In this work, we focus on winds driven by stellar feedback, primarily in the form

of energy released by core-collapse supernovae (SNe). While large scale cosmological simulations are now approaching high enough resolutions where multiphase structures are seen to develop self consistently in galactic winds (e.g., in TNG50 Nelson et al. 2019 and FIRE Anglés-Alcázar et al. 2017; Pandya et al. 2021), the detailed properties of multiphase outflows are still most clearly seen in high resolution galaxy scale simulations (e.g., Schneider et al. 2020; Steinwandel et al. 2022; Rey et al. 2023), or interstellar medium (ISM) patch simulations (e.g., TIGRESS Kim et al. 2017; Kim & Ostriker 2018; SILCC Gatto et al. 2017). The \sim pc-scale resolution of these simulations makes it possible to resolve the energy injection by SNe and the resulting interactions with the surrounding ISM (e.g., de Avillez 2000; Joung & Mac Low 2006; Hill et al. 2012). Such simulations have also revealed a trove of complications surrounding the launching of these winds. For example, the efficacy of this process is sensitive to the spatial distribution of SNe (Creasey et al. 2013; Martizzi et al. 2015; Li et al. 2017; Smith et al. 2021) as well as spatiotemporal clustering (Kim & Ostriker 2017; Fielding et al. 2018) and self-consistency of treatments of star formation and feedback (Kim et al. 2017; Kim & Ostriker 2018).

The survival and growth of cool gas clouds entrained in these winds has also been a hotbed of analytic and numerical studies in recent years. Theoretically, this picture was initially problematic as these cool clouds should be quickly destroyed by hydrodynamic instabilities during this process (Klein et al. 1994). The timescale for the acceleration of a cool cloud is a factor of $\gtrsim 10$ times longer than the timescale for its destruction via Kelvin-Helmholtz (KH) and Rayleigh-Taylor (RT) instabilities (Zhang et al. 2017). This was also verified in simulations (Cooper et al. 2009; Scannapieco & Brüggén 2015; Schneider & Robertson 2017). However, it was recently shown that under the right conditions, these cool clouds can not only survive but even grow within the hot wind (Armillotta et al. 2017; Gronke & Oh 2018, 2020a). This is possible when the cooling time of mixed gas is shorter than the destruction timescale. Subsequent studies have shown that exact

characterizations of this parameter space is complex (Li et al. 2020; Sparre et al. 2020; Kanjilal et al. 2021; Farber & Gronke 2022; Abruzzo et al. 2022b). The mechanism that drives this growth is the formation of a long tail structure of cool gas where mixing and cooling takes place efficiently in turbulent radiative mixing layers at the interface of the two gas phases (Ji et al. 2018; Fielding et al. 2020; Tan et al. 2021). Crucially, this requires long simulation domains that were not captured in earlier simulations, highlighting the importance of capturing the appropriate scales and boundary conditions that define such problems.

The interplay between the cool and hot gas phases, and the role of various physical processes such as radiative cooling, turbulence, and magnetic fields, are hence crucial in determining the fate of cool gas clouds and their impact on the overall dynamics of galactic winds (Fielding & Bryan 2022). Despite these advances, many aspects of galactic wind simulations still rely on simplified models and assumptions, such as idealized cloud geometries or time-constant wind properties. The next step forward is to study the problem under more realistic conditions, incorporating the diverse range of physical processes and interactions that govern the evolution of galactic winds and their multiphase structures.

To that end, in this work we connect results from small scale idealized cloud crushing simulations to large scale multiphase galactic winds by studying how clouds form, evolve and interact with the wind. This mesoscale approach bridges the microscale phenomena controlling individual clouds to the macroscale phenomena controlling galactic/cosmological scale evolution by simultaneously resolving the lifespan of thousands of individual clouds as well as the driving of the highly turbulent, temporally variable, hot wind that they interact with. The basis of our study builds upon a series of simulations following a similar design to those presented in Fielding et al. (2018), which capture the dynamics in a vertically stratified patch of a galactic disc. In these simulations, clustered




Figure 5.1: Volume rendering of the main simulation used in this work that demonstrates the presence of numerous cool clouds (shown in opaque blue) embedded within the hot outflow (shown in transparent orange) powered by the SNe (shown as the light source) exploding within the disc mid-plane.

SNe lead to the formation of superbubbles which propagate quickly enough to break out of the galactic discs, providing a pathway for the energy and momentum released to vent into the halo in the form of a strong galactic wind. The shredding and entrainment of cool dense clumps are seen in these winds, and are expected to be important in the influencing the structure and dynamics of the outflow. We extend the box asymmetrically so as to follow the clouds in the outflow out to a larger height, as well as include additional physics such as self-developed magnetic fields in the turbulent ISM. Figure 5.1 shows a rendering from our simulation of cool structures (in blue) in the hot outflowing wind (in orange). Our main focus is on analyzing the clouds in these winds. More specifically, we investigate how they are formed during superbubble breakout and their properties relative to predictions from idealized small scale simulations of turbulent mixing layers and individual clouds.

The outline of this paper is as follows. In Section 5.2, we propose an analytic model for the distribution and growth of cool clouds in an outflowing wind. In Section 5.3, we provide a detailed description of our simulation setup. In Sections 5.4 and 5.5, we present simulation results, with the former focusing on bulk properties of the outflow and the latter on analysis of the clouds. Lastly, we discuss some implications and complications in Section 5.6 before summarizing and concluding in Section 5.7.

5.2 Seeding Clouds In Outflows

Despite the evergrowing body of work on small and large scale multiphase systems, equivalent studies of intermediate scales remain lacking. There is therefore an imperative to develop and validate models which bridge this gap. Here, we outline such a model describing a population of cool clouds in hot outflowing winds, building on results from work at smaller scales, which we then compare to simulations in following sections. In

particular, we focus on two main components—the distribution of cloud sizes and the rate at which they grow. Together, they enable us to better build and refine analytic frameworks for modelling these systems. We will discuss each in turn, beginning with the distribution of cloud sizes.

5.2.1 Cloud Size Distribution

Inverse power law distributions, where the probability of a quantity taking on some value varies inversely with the power of that value, are widely observed in both the physical and social sciences, with a large existing body of empirical evidence supporting their existence (Newman 2005). They are commonly referred to as Zipf’s law in the discrete case, or Pareto distributions in the continuous case. Their theoretical origins are, however, much less certain. In particular, it is difficult to explain their seeming universality. A ~ -2 exponent, which indicates that there is a constant contribution from each logarithmic bin, is ubiquitous, even within astronomy (Guszejnov et al. 2018). Take, for example, the stellar initial mass function, which has a roughly $dN(m)/dm \propto m^{-2.35}$ scaling at higher masses (Krumholz 2014). This has been attributed to turbulent fragmentation (Padoan & Nordlund 2002; Hopkins 2013), where initially turbulent gas fragments into clumps. In this case, fragmentation is driven primarily by gravitational collapse.

Gronke et al. (2022) find in simulations of turbulent multiphase boxes that the cool gas breaks up into droplets which follow a similar distribution. This setup is akin to a population of clouds fully entrained in a turbulent background wind. They argue that this power law, and in particular the -2 exponent, arises from competition between fragmentation and growth, where growth is provided from both a multiplicative source (cooling induced growth) as well as an additive source (coagulation of multiple smaller

clumps with a larger clump).²⁰ Fragmentation, on the other hand is driven solely by the turbulent field. Likewise, Fielding et al. (2022) find the same mass distribution of cool clumps in turbulent magnetized boxes where the multiphase medium is formed via thermal instability.

While our setup differs from these more idealized ones, many features are similar. Instead of a single large cloud in a turbulent velocity field, we have an ISM that is broken up by an expanding hot superbubble. While the ISM itself is initially turbulent, the fractal expansion of the bubble surface through the multiphase ISM (Lancaster et al. 2021) fragments the swept up ISM in a scale free manner, eventually launching asymmetric winds above and below the disc. It is the fragmentation of the initial ISM that determines the distribution of cloud sizes, at least when the outflowing wind first breaks out. The initial cloud mass population in the hot wind is thus seeded by the fragmented ISM that the superbubble which birthed the wind itself has broken up. This leads to a distribution that scales as:

$$\frac{dN}{dm} \propto m^{-2}. \tag{5.1}$$

We will show that our simulations support this picture of cloud genesis via swept up ISM shell fragmentation during superbubble breakout. This picture is fundamentally distinct from the naive model in which preexisting cold clouds in the ISM are accelerated out of the galaxy while remaining mostly intact.

What sets the lower and upper limits on cloud sizes in this model? Since the clouds are formed during the breakup of the ISM, the largest cloud size is simply constrained by the initial scale of the ISM, i.e., the scale height H of the disc. The scale of the smallest clouds is a more complex question. While fragmentation can lead to arbitrarily small clouds,

²⁰In Gronke & Oh (2022), they explore the effects of coagulation further and identify a critical Mach number below which coagulation dominates. This critical Mach number is small in our simulations, suggesting that coagulation should not dominate.

they should be rapidly destroyed in the hot turbulent wind below some critical size, where growth due to cooling becomes ineffective. This survival threshold r_{crit} presents a natural choice for the lower limit. However, the problem of what determines r_{crit} is a thorny one, and there has been no shortage of recent investigations into the matter. We will briefly review some of them.

For a cloud of size r_{cl} with overdensity χ relative to its background and temperature T_{cl} at rest embedded in a hot wind of temperature T_{hot} moving at a velocity v_{wind} , Gronke & Oh (2018) identified a characteristic size $r_{\text{crit,cc}}$ above which the clouds survive till entrainment and grow. This size corresponds to where the cloud crushing time $t_{\text{cc}} \sim \sqrt{\chi} r_{\text{cl}} / v_{\text{wind}}$, the timescale on which a cloud is destroyed by hydrodynamic instabilities (Klein et al. 1994), becomes longer than the cooling time of mixed gas $t_{\text{cool,mix}} \equiv t_{\text{cool}}(T_{\text{mix}})$, where $T_{\text{mix}} \sim \sqrt{T_{\text{cl}} T_{\text{hot}}}$ (Begelman & Fabian 1990). From their turbulent box simulations, Gronke et al. (2022) found that there is also an additional empirically determined Mach number dependence of the functional form $f(M_{\text{turb}}) \sim 10^{0.6 M_{\text{turb}}}$. They attribute this additional Mach dependence to increased turbulent disruption, in contrast to the increased survival times seen in cloud crushing simulations with high Mach winds which stems from cloud compression (Scannapieco & Brüggén 2015; Bustard & Gronke 2022). The critical size for survival is thus given by

$$r_{\text{crit,cc}} = 2 \text{ pc} \frac{T_{\text{cl},4}^{5/2} M_{\text{wind}}}{P_3 \Lambda_{\text{mix},-21.4}} \frac{\chi}{100} f(M_{\text{turb}}), \quad (5.2)$$

where $T_{\text{cl},4} \equiv (T_{\text{cl}}/10^4 \text{ K})$, $P_3 \equiv nT/(10^3 \text{ cm}^{-3} \text{ K})$, $\Lambda_{\text{mix},-21.4} \equiv \Lambda(T_{\text{mix}})/(10^{-21.4} \text{ erg cm}^3 \text{ s}^{-1})$, and M_{wind} is the Mach number of the wind.

Alternatively, Li et al. (2020) and Sparre et al. (2020) lay out a survival criterion $t_{\text{cool,wind}} < t_{\text{life}}$, where $t_{\text{cool,wind}}$ is the cooling time of the *hot* background wind and $t_{\text{life}} =$

$\bar{f}t_{cc}$ is the predicted cloud lifetime. \bar{f} is an empirically calibrated function, given by

$$\bar{f} = 10 \left(\frac{r_{cl}}{1 \text{ pc}} \right)^{0.3} \left(\frac{n_{hot}}{0.01 \text{ cm}^{-3}} \right)^{0.3} \left(\frac{T_{hot}}{10^6 \text{ K}} \right) \left(\frac{v_{wind}}{100 \text{ km s}^{-1}} \right)^{0.6}, \quad (5.3)$$

where n_{hot} is the density of the hot gas. At $\chi = 100$, the difference between the two criteria is small, and when comparing the two, attributed to different ways of defining cloud survival (Kanjilal et al. 2021). At larger χ however, these criteria differs by orders of magnitude. Abruzzo et al. (2022b) find that at $\chi > 100$, the Li/Sparre criterion agrees much better with simulation results, but also point out that $t_{cool,wind}$ cannot be physically important since their results are unchanged when they artificially shut off wind cooling. They instead propose a new criterion $t_{cool,mix} < t_{shear}$ which captures the empirically derived χ scaling of r_{crit} in a physically motivated model (Matthew Abruzzo et al. in prep.). Here $t_{shear} \sim r_{cl}/v_{wind}$ is the wind crossing time over the cloud, the idea being that a cloud survives if mixed gas on the cloud surface is able to cool before being stripped away. This criterion gives a cloud survival size

$$r_{crit,shear} \equiv \sqrt{\chi} r_{crit,cc}. \quad (5.4)$$

5.2.2 Theoretical Model for Cloud Growth

Assuming that the conditions are right for a cloud to survive and grow in a wind, we can write its mass growth rate as

$$\dot{m} \sim \rho_{wind} v_{in} A_{cloud}, \quad (5.5)$$

where ρ_{wind} is the density of the hot wind, v_{in} is the inflow velocity corresponding to the mass flux from the hot background onto the cloud, and A_{cloud} is the effective surface area

of the cloud.

The first ingredient in this cloud growth model is A_{cloud} . It is important here to distinguish this effective surface area from the non-convergent surface area of the cool gas (Fielding et al. 2020; Abruzzo et al. 2022b), which corresponds more closely to the temperature isosurfaces we measure in our simulations below. Instead, a more accurate characterization of A_{cloud} is an effective surface area corresponding to some smooth envelope around the cloud (Gronke & Oh 2020a). In cloud crushing simulations, surface area has an initial rapid growth stage associated with tail formation, followed by slower isotropic growth (Gronke & Oh 2020a; Abruzzo et al. 2022b). This corresponds to A_{cloud} going from initially scaling as $\propto V$ (linearly with volume), to scaling as $\propto V^{2/3}$. The initial exponential growth can be attributed to a tail formation phase where the high shear mainly leads to mass growth in a growing tail downstream of the cloud. Eventually, the cloud becomes entrained and the surface area scales much as a monolithic cloud would. However, it was found in Gronke & Oh (2022) and Tan et al. (2023) that for clouds which never get entrained (either because of large scale turbulence or gravitational acceleration), and are hence subject to continuous shear, the cloud surface area instead scales as $A_{\text{cloud}} \propto V^{5/6}$. This scaling can arise from the fractal nature of the mixing surface, which has been measured in mixing layer simulations to have a fractal dimension $D \sim 2.5$ (Fielding et al. 2020), which would mean $A \propto V^{D/3} \sim V^{5/6}$. If we assume that the smallest surviving clouds are the most spherical, we can model the effective cloud surface area as

$$A_{\text{cloud}} = A_{\text{crit}} \left(\frac{V}{V_{\text{crit}}} \right)^{5/6}, \quad (5.6)$$

where A_{crit} and V_{crit} are the spherical area and volumes respectively corresponding to the critical cloud survival size r_{crit} , which we discuss below.

Next we discuss the second ingredient in the model, the inflow velocity v_{in} . In plane

parallel simulations of radiative turbulent mixing layers (Tan et al. 2021; Fielding et al. 2020), v_{in} is found to scale as

$$v_{\text{in}} \sim v_{\text{turb}}^{3/4} \left(\frac{L_{\text{turb}}}{t_{\text{cool},\text{min}}} \right)^{1/4}, \quad (5.7)$$

where v_{turb} is the turbulent velocity in the mixing layer, L_{turb} is the largest mixing scale (typically the outer scale of the turbulence), and $t_{\text{cool},\text{min}}$ is the minimum cooling time in the simulation, which scales inversely with pressure ($t_{\text{cool},\text{min}} \propto P^{-1}$). Additionally, these simulations find that the turbulent velocity scales with the relative shear velocity in these setups, such that $v_{\text{turb}} \sim f_{\text{rel}} v_{\text{shear}}$, where f_{rel} is some constant of proportionality (Tan et al. (2021) report $f_{\text{rel}} \sim 0.3$ and Fielding et al. (2020) find $f_{\text{rel}} \sim 0.1 - 0.2$) that varies with the geometry of the setup (e.g., Mandelker et al. (2020) find $f_{\text{rel}} \sim 0.2 - 0.3$ in supersonic streams). However, in cloud crushing setups, clouds eventually entrain, i.e., v_{shear} goes to zero, yet these clouds continue to grow (Gronke & Oh 2020a), suggesting that this relationship breaks down at lower values of v_{shear} . This is likely due to continued driving of mixing due to cooling, either via induced pulsations (Gronke & Oh 2022) or is self-sustained as the cooling velocity drives ongoing turbulence after entrainment (Abruzzo et al. 2022b). To account for this saturation at low shear velocities, we can hence express the turbulent velocity as

$$v_{\text{turb}} \approx \max(f_{\text{rel}} v_{\text{shear}}, c_{\text{s,cold}}). \quad (5.8)$$

In our simulations, we find that $f_{\text{rel}} \sim 1/15$.

For strong cooling regimes, Tan et al. (2021) calibrate the equations above to simu-

lations of turbulent mixing layers and find that

$$v_{\text{in}} \approx 10 \text{ km s}^{-1} \left(\frac{v_{\text{turb}}}{50 \text{ km s}^{-1}} \right)^{3/4} \left(\frac{L_{\text{turb}}}{100 \text{ pc}} \right)^{1/4} \left(\frac{t_{\text{cool,min}}}{0.03 \text{ Myr}^{-1}} \right)^{-1/4}. \quad (5.9)$$

The above scalings assume that we are in the subsonic to transonic regime where the relative velocity between the hot wind and the cloud does not exceed the hot gas sound speed. The turbulent mixing velocity saturates as clouds become supersonic (Yang & Ji 2023), which would change the scalings above. We also assume that we are in the rapid cooling regime, where the cooling time is much shorter than the turbulent mixing time L/v_{turb} , i.e., the outer scale eddy turb over time (Damköhler number $Da \equiv L/(v_{\text{turb}}t_{\text{cool,mix}}) > 1$; Tan et al. 2021).

To summarize, we have presented a model for seeding a population of growing clouds into an outflowing wind. In this work, we will show that we can apply this to simulations of a population of clouds in a more realistic outflowing wind environment that is turbulent and also includes magnetic fields. We find that contrary to initial expectations, this simple model is surprisingly effective.

5.3 Methods

In this work, we carry out 3D magnetohydrodynamical (MHD) simulations using the publicly available code Athena++ (Stone et al. 2020) aimed at capturing the multiphase dynamics in a vertically stratified patch of a galactic disc. SNe corresponding to a star cluster are seeded by hand in the middle of the disc, which leads to the formation of a superbubble that propagates rapidly enough to break out of the disc and release energy and momentum into the halo in the form of a galactic wind. This builds upon the design of a similar suite of simulations presented in Fielding et al. (2018) where additional

details and discussion can be found. In our analyses in the following sections, we focus on studying the formation and dynamics of a resulting population of cool material that gets sown into the wind. Before that, we describe in this section our simulation setup, implementations of various physics, and the initial conditions we adopt.

5.3.1 Setup

All simulations are run in three dimensions on Cartesian grids using the HLLD approximate Riemann solver and the third-order accurate Runge-Kutta time integrator. Our simulation setup consists of a rectangular box with dimensions $512 \times 512 \times 2048$ pc and a fiducial resolution of 2 pc (we also discuss in this work a simulation run with a higher resolution of 1 pc but which fails to drive a successful outflow). The box is asymmetric—the mid-plane is located a quarter box length from the bottom—so as to follow the outflow on the upper side out to a greater height. We use periodic boundary conditions on the sides of the box, along the x and y directions. We adopt outflowing boundary conditions (zero gradient with inflow explicitly disallowed) along the z direction at the top and bottom. We also implement a density floor of $n = 10^{-8} \text{ cm}^{-3}$.

5.3.2 Source Terms

We include a background gravitational profile, driven turbulence, optically thin radiative cooling, photoelectric heating, and SNe injection. Here, we provide more details about the implementation of each of these in turn. We do not include any explicit viscosity or thermal conduction. A $\gamma = 5/3$ equation of state is adopted throughout.

Gravity

We first consider the density profile of a thin isothermal disc in hydrostatic equilibrium. We begin by assuming that the self-gravity of the gas is a subdominant component of the total gravitational potential and is hence unimportant (i.e., $\sim 2\pi G\Sigma_{\text{gas}}z$ is small, which is true for $\Sigma_{\text{gas}} \ll 1000M_{\odot}/\text{pc}^2$). We hence neglect self gravity in our simulations. The spherical potential at a vertical height z above the disc and a radial distance R from the center is

$$\Phi(R, z) = -\frac{GM}{r} = -\frac{GM}{(R^2 + z^2)^{1/2}}, \quad (5.10)$$

and hence the gravitational acceleration is

$$g(R, z) = -\frac{\partial\Phi}{\partial z} = -\frac{GM}{(R^2 + z^2)^{3/2}}. \quad (5.11)$$

For a thin disc, $z \ll R$,

$$g(R, z) \approx -\frac{GM}{R^3}z = -\Omega^2z, \quad (5.12)$$

where $\Omega = v_{\text{cir}}/R$ is the angular velocity. Our equation of hydrostatic equilibrium is thus

$$\frac{dP}{dz} = \rho g z = -\rho\Omega^2z. \quad (5.13)$$

With $P = c_i^2\rho$ (c_i is the isothermal sound speed, related to the adiabatic sound speed by $c_s^2 = \gamma c_i^2$), we can easily solve this for the density profile:

$$\rho(z) = \rho_0 \exp\left[-\frac{z^2}{2H^2}\right], \quad (5.14)$$

where $H \equiv c_i/\Omega$ is the scale height and ρ_0 is the mid-plane density at $z = 0$.

We can however drop the thin disc assumption and solve for the profile out to larger z . We now have

$$g(R, z) = -\frac{GM}{R^3} \frac{R^2}{(R^2 + z^2)^{3/2}} z = -\Omega^2 \frac{R^2}{(R^2 + z^2)^{3/2}} z. \quad (5.15)$$

Solving this gives us:

$$\rho = \rho_0 \exp \left[\frac{R^2}{H^2} \left(\frac{R^2}{(R^2 + z^2)^{1/2}} - 1 \right) \right]. \quad (5.16)$$

These are the profiles we adopt in our simulations. We test systems in hydrostatic equilibrium with this gravitational profile in Appendix B. Note we have assumed an external spherical potential for simplicity. One could in principle construct the gravitational potential with a disc component and a dark matter halo component along with self-gravity for a more realistic treatment, which is important if one is going to distances $\gg 1$ kpc (Li et al. 2017). However, our simplified assumption is sufficient for our box size.

Stratified Turbulence

We drive turbulence on large scales so that the mass-weighted velocity dispersion $\delta v \equiv \langle v^2 \rangle^{1/2} \approx 10$ km/s, which is roughly the sound speed of 10^4 K gas and consistent with observed velocity dispersions in the ISM. The turbulent kinetic energy injection rate is thus $\dot{E}_{\text{turb}} \approx \rho \delta v^3 L_{\text{box}}^2$ where L_{box}^2 is the horizontal box length. The turbulence is driven on the scale of the disc scale height with power distributed evenly between wave numbers $k = 2$ to $k = 3$. The turbulence is constrained to follow a Gaussian profile with scale height $H \sim 80$ pc (we test this implementation of constrained turbulence in Appendix C). We drive the turbulence every 5×10^{-3} crossing times t_{cross} . The driven




Figure 5.2: Cooling curves along with our fit and heating curves for $\langle n_H \rangle = 100 \text{ cm}^{-3}$ at several different pressures (e.g., P_4 is $P = 10^4 k_B \text{ cm}^{-3} \text{ K}$).

turbulence in Athena++ adopts an Ornstein-Uhlenbeck process to smoothly vary the velocity perturbations over some correlation time (Lynn et al. 2012), which we set to 8 Myr. Turbulence along with heating and cooling leads to the formation of a turbulent multiphase ISM which we allow to form for 60 Myrs (roughly a turnover time) prior to any supernova explosions.

Cooling/Heating

The net cooling rate per unit volume is formulated as

$$\rho\mathcal{L} = n^2\Lambda - n\Gamma, \quad (5.17)$$

where Λ is the cooling function and Γ is the heating rate. Our cooling function $\Lambda(T)$ combines the collisionally ionized equilibrium (CIE) cooling curve in Gnat & Sternberg (2007) for $T \geq 10^4$ K with the cooling function for $T \leq 10^4$ K in Koyama & Inutsuka (2002). We obtain our cooling curve by performing a piece-wise power law fit over ~ 50 logarithmically spaced temperature bins, starting from a temperature floor of 10 K up to a maximum temperature of 10^{10} K. We also include a photoelectric heating (PEH) rate $\Gamma = 10^{-26}(\langle n_H \rangle / \text{cm}^{-3}) \text{ erg s}^{-1}$, where we have scaled the solar value by the average mid-plane density to approximate the scaling of PEH with higher star formation rates in denser regions. This approach is similar to those used in Kim & Ostriker (2017) and Fielding et al. (2018), which run similar setups. We assume solar metallicity ($X = 0.7$, $Z = 0.02$). We also use a fixed mean molecular mass μ , for a fully ionized plasma. This has a ≤ 2 factor of error in the temperature of neutral/partially ionized gas below 10^4 K, but should not affect our overall conclusions. Figure 5.2 shows the aforementioned cooling curves along with our fit and heating curves for $\langle n_H \rangle = 100 \text{ cm}^{-3}$ at several different pressures (e.g., P_4 is $P = 10^4 k_B \text{ cm}^{-3} \text{ K}$).

We also adopt an additional constraint on the simulation timestep over the standard CFL constraint where we require that the timestep is less than or equal to one quarter of the shortest single-point cooling time t_{cool} across the whole domain. This is to ensure good coupling between cooling and the hydrodynamical evolution.

Finally, we develop and implement an extended version of the fast and robust exact cooling algorithm described in Townsend (2009) to include heating (See Appendix D for




Figure 5.3: The SNe rate as as function of the stellar population age. This consists of the core-collapse SNe rate (< 44 Myr) and the Ia rate (> 44 Myr).

full details).

Supernova Injection

We include both core-collapse supernova and Ia rates using the piece-wise power law fits given in Hopkins et al. (2023). The combined SNe rate is

$$\frac{R_{\text{SN}}/M_*}{\text{Gyr}^{-1}M_{\odot}^{-1}} = \begin{cases} 0 & (t < t_1) \\ a_1(t/t_1)^{s_1} & (t_1 \leq t \leq t_2) \\ a_2(t/t_2)^{s_2} & (t_2 \leq t \leq t_3) \\ a_3(t/t_3)^{s_3} & (t > t_3), \end{cases} \quad (5.18)$$

where $s_1 \equiv \ln(a_2/a_1)/\ln(t_2/t_1)$, $s_2 \equiv \ln(a'_2/a_2)/\ln(t_3/t_2)$, $s_3 = -1.1$, $(a_1, a_2, a'_2, a_3) = (0.39, 0.51, 0.18, 0.0083)$ and $(t_1, t_2, t_3) = (3.7, 7.0, 44)$ Myr. The time of the first Ia is set to be after the time of the last CCSNe ($t = t_3$). For $M \equiv M_*/M_{\odot}$, the SNe number per initial solar mass $n_{\text{SN}} \equiv N_{\text{SN}}/M$ is thus given by

$$n_{\text{SN}}(t) = \begin{cases} 0 & (t < t_1) \\ \frac{a_1}{s_1+1} \frac{t_1}{1\text{Gyr}} [(t/t_1)^{s_1+1} - 1] & (t_1 \leq t \leq t_2) \\ n_{\text{SN}}(t_2) + \frac{a_2}{s_2+1} \frac{t_2}{1\text{Gyr}} [(t/t_2)^{s_2+1} - 1] & (t_2 \leq t \leq t_3) \\ n_{\text{SN}}(t_3) + \frac{a_3}{s_3+1} \frac{t_3}{1\text{Gyr}} [(t/t_3)^{s_3+1} - 1] & (t > t_3). \end{cases} \quad (5.19)$$

The time where n supernovae have gone off is hence

$$t_{\text{SN}}(M, n) = \begin{cases} t_1 \left[\frac{n(s_1+1)}{Ma_1(t_1/1\text{Gyr})} \right]^{1/(s_1+1)} & (n \leq n_1) \\ t_2 \left[\frac{(n-n_1)(s_2+1)}{Ma_2(t_2/1\text{Gyr})} \right]^{1/(s_2+1)} & (n_1 \leq n \leq n_2) \\ t_3 \left[\frac{(n-n_2)(s_3+1)}{Ma_3(t_3/1\text{Gyr})} \right]^{1/(s_3+1)} & (n \geq n_2). \end{cases} \quad (5.20)$$

We deposit 10^{51} ergs per SN as pure thermal bombs. While one could determine the amount of thermal and kinetic energy injected by each SN (see Martizzi et al. (2015), who calibrated their model to high resolution simulations of individual SN remnants), this would matter only for the first couple of SNe. Since the SNe are clustered tightly both in space and time, most of them barring the first few occur in the hot and low density remnant of the previous SNs, assuming that $\Delta t_{\text{SN}} < t_{\text{PE}}$ (t_{PE} being the timescale over which the SNR reaches pressure equilibrium or mixes with the ambient ISM) so that a coherent bubble can be driven (Fielding et al. 2018). Hence their cooling radii r_{cool} are around an order of magnitude larger than the injection radius r_{inj} , which means that most of the energy from the SN is transferred to the surrounding gas.

We assume a spherical geometry for the injection site with radius $r_{\text{inj}} = 20 \text{ pc}$. We assume a stellar ejecta mass of $8.72 M_{\odot}$ for core-collapse SN and $1.4 M_{\odot}$ for Ia. We also assume that all SN go off at the origin, since the cluster radius $R_{\text{cl}} = 10 \text{ pc}$ is relatively small.

While out of the scope of this work, we note that in reality, stars far from the disc center could have a disproportionate effect in driving winds. For instance, Ia at late times from stars that settle high above the disc or OB (binary) runaways (e.g., Steinwandel et al. 2022). This is mainly due to the dual effects of a larger r_{cool} in the lower density region along with the lower scale height contributing to a large r_{cool}/H ratio.

For each cell with center within the sphere radius, we then inject energy from the SN uniformly. Note that there will be some error due to resolving the volume of the sphere, but the error is small as long as we resolve the injection radius by more than a few cells.

5.3.3 Initial Conditions

We adopt a circular velocity $v_{\text{cir}} = 175 \text{ km/s}$ and galactic radial distance $R = 1 \text{ kpc}$, which corresponds to a scale height $H = 66 \text{ pc}$. The box is initially filled with $T = 10^4 \text{ K}$ gas with a mid-plane density $n = 100 \text{ cm}^{-3}$ at $z = 0$. The average density within one scale height is $\langle n \rangle = 86 \text{ cm}^{-3}$. Using $\langle n \rangle \equiv \Sigma_g / (2H\mu m_p)$, this corresponds to a gas surface density $\Sigma_g \approx 175 M_\odot \text{ pc}^{-2}$. For a star formation efficiency $\epsilon_* \equiv M_{\text{cl}} / \pi h^2 \Sigma_g \sim 0.1$ (as found in Grudić et al. (2018)), this corresponds to a star cluster mass $M_{\text{cl}} \sim 2.5 \times 10^5 M_\odot$. For comparison, the simulation from the suite of TIGRESS simulations under the SMAUG project with the highest star formation rate (R2) has $\Sigma_g \approx 74 M_\odot \text{ pc}^{-2}$ (Kim et al. 2020a). We initialize a magnetic field aligned along the x-direction with plasma beta $\beta \equiv P_{\text{gas}} / P_B = 100$ everywhere. However, through cooling and turbulent dynamo amplification, $\beta \sim 1$ is rapidly achieved in the ISM.

5.4 Results: Breakout and Wind Outflow

In this section, we provide context for the winds that contain our clouds, setting the stage for our main focus—a closer look at cloud properties—in the next section. There is already a large body of existing work studying these SN driven outflows—we concentrate mainly on analysis meant to highlight the most salient points that are relevant to these clouds.

5.4.1 Superbubble Breakout

Supernovae (SNe) are powerful energetic events that release a tremendous amount of energy ($\sim 10^{51} \text{ ergs}$) and mass into their surrounding environments. However, individual SNe are ineffective at driving galactic winds because most of this energy is radiated

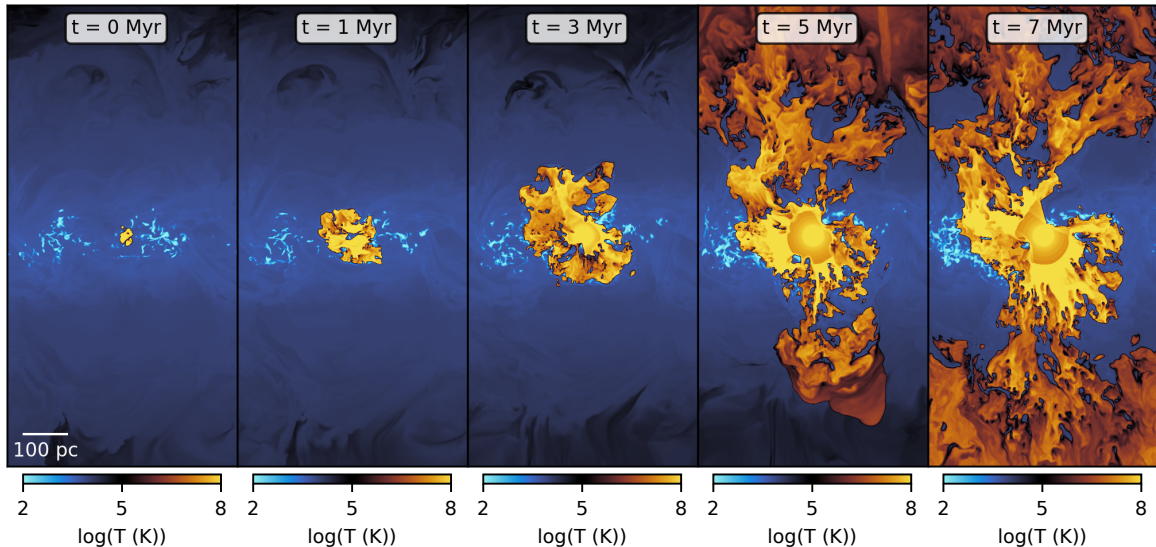


Figure 5.4: The various stages of a superbubble breakout, driven by clustered SNe. As the first SNe go off, a hot bubble expands outwards, propagating through the multiphase ISM, leading to a complex fractal structure. As the superbubble eventually breaks out of the disc, the ISM fragments and breaks up, thus seeding the hot outflow with cool clouds.

away before it can break out of the disc. Spatiotemporal clustering of SNe resolves this problem, and is motivated by the fact that stars form primarily in clusters. When multiple supernova explosions occur in the same region of space and over a relatively short period of time, they can overlap and form superbubbles. These superbubbles are able to drive the expansion of the material outwards from the cluster, breaking out of the disc before the cluster runs out of SNe. After superbubble breakout the hot material inside is able to vent its energy into the halo and power an outflowing galactic wind. As the superbubble propagates through the clumpy multiphase environment, it sweeps around higher-density regions, resulting in a surface that is fractal in nature. This complex interplay leads to the fragmentation and break up of the ISM, seeding the hot wind with a population of cool clouds. Figure 5.4 illustrates this by showing temperature slices at various stages of the process described above.

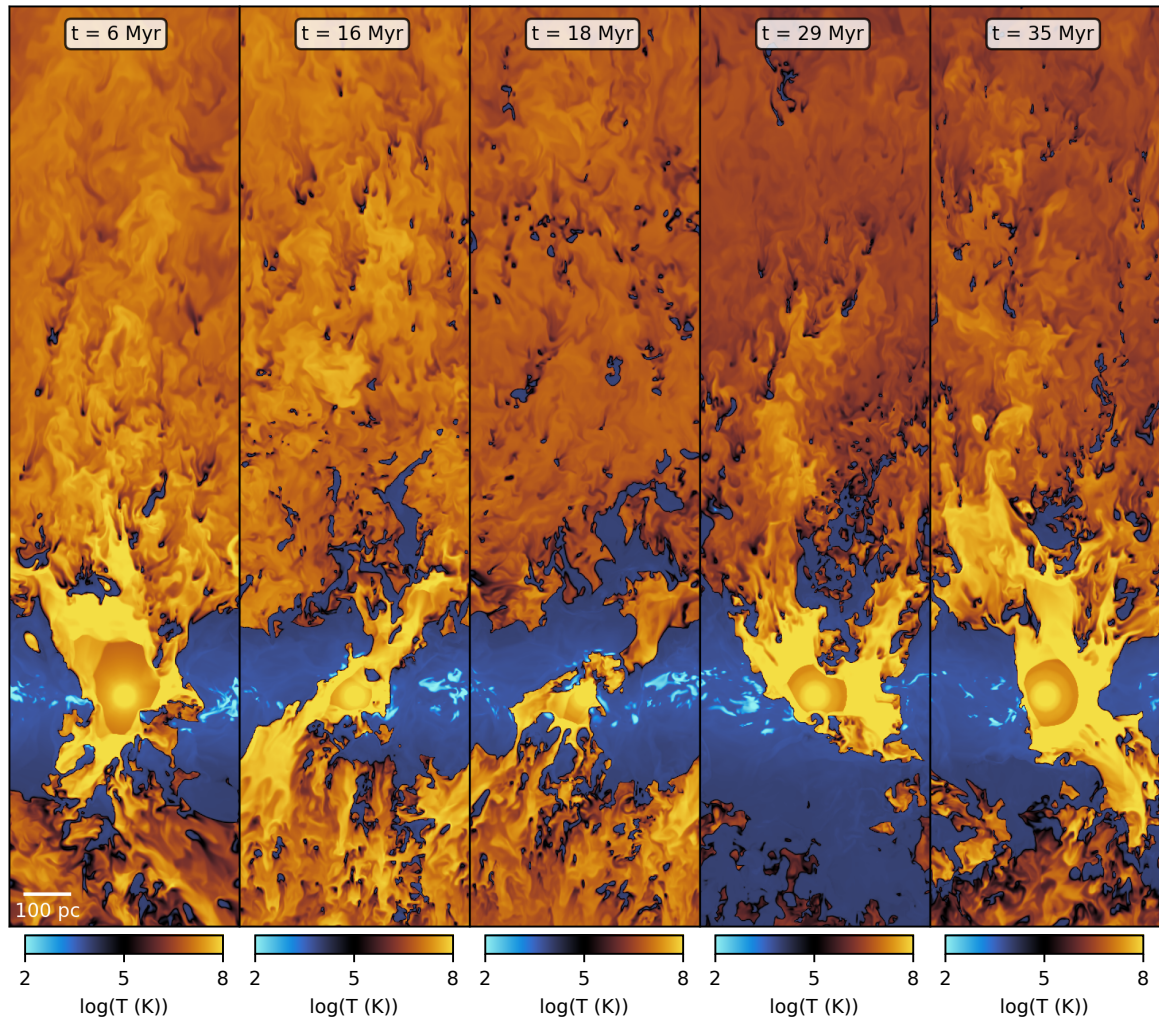


Figure 5.5: Time slices of temperature from our fiducial simulation where the breakout successfully launches a wind above the disc, despite the outflow being interrupted in the middle.




Figure 5.6: Time slices of temperature from a simulation where the breakout is not successful in launching a wind on the top side and is disrupted too early by cold gas in the ISM. This causes the breakout to cool and fall back towards the disc, a cycle which is repeated several times.




Figure 5.7: The outwards mass flux over time at four different heights showing the initial breakout and outflow above the disc. This then reverses sides twice over the timespan of the SNe. Dashed line represents negative values, while the red shaded region highlights the time window after the SNe rate peaks in which we analyze wind embedded clouds in the next section.

5.4.2 Asymmetric Outflows

By virtue of having to make its way through a multiphase ISM, the expansion and eventual break out of the superbubble can exhibit great asymmetry in terms of driving outflows above and below the disc. This can be self-reinforcing—whichever side breaks through first is then free to vent into the low density halo, creating a channel of lower resistance for the hot gas to funnel through and leading to a weak or non-existent outflow on the opposite side. In addition, turbulent motions in the ISM can exacerbate or reverse this asymmetry. For example, if cold material in the ISM moves over the cluster region, it ‘caps’ and inhibits the path of the hot expanding wind, much like turning off a valve, resulting in a redirection towards the other side of the disc. In this manner, an outflowing wind can be quenched at the base on one side midway during the lifetime of the cluster SNe. While we did not investigate this in further detail, we observed that this was a common occurrence in our simulations—due to the nature of the setup, a wind that breaks out on one side pushes the ISM radially outwards in the plane of the disc. Because of the periodic boundary conditions, this leads to a build up of ISM material on the sides which then falls back towards the cluster. This sloshing motion often eventually halts the wind and drives it to the opposing side. In a more realistic larger scale setup, this indicates that dynamical interactions between different star clusters could complicate wind driving, either reinforcing or disrupting outflows depending on their spatiotemporal separations.

Figures 5.5 and 5.6 show slices from two very similar simulations differing primarily in the random seeding of turbulent driving in which the breakout succeeds and fails in driving an outflowing wind, respectively. In Figure 5.5, we see that the wind breaks out on the topside and that fragments of the ISM become entrained in the hot outflowing wind, whereas it fails to do the same on the bottom side, which is apparent when looking

at the two panels at later times. Meanwhile, in Figure 5.6, the wind initially break outs out but is first briefly interrupted, causing the hot wind to cool and fall back towards the disc, before breaking out again and then being stopped a second time, again leading to an overall cooling of the hot material which falls back disc-wards. This demonstrates the way winds can fail to be launched.

We can look at this more quantitatively by tracking how the mass flux in the simulation evolves over time. In Figure 5.7, we show the outwards mass flux per unit area at four different heights in our box as a function of time for our fiducial simulation shown in Figure 5.5. The blue and orange lines show the mass fluxes below and above the disc respectively at 0.5 kpc, while the red and teal lines are at 1 and 1.5 kpc above the disc. The dashed orange line shows when the mass flux becomes negative, i.e., material is falling back to the disc. We can see that the breakout is initially on the top side which launches a strong outflowing wind. As the outflow is interrupted at the base, the wind loses its power source, and begins to slow and dissipate. Meanwhile, the outflow is directed to the bottom side of the disc. This eventually reverses again nearing the end of the cluster SN period. The red region highlighted in Figure 5.7 shows the time range in which we analyze clouds in Section 5.5. This choice corresponds to a time period right after the peak of the SNe rate (see Figure 5.3).

5.4.3 Turbulent Winds

Figures 5.8, 5.9, and 5.10 show the range of quantities we track in our simulations, and allow us to make some general statements about the outflowing winds. Figure 5.8 shows slices of temperature, outflow velocity, number density, pressure and magnetic plasma beta for a single time snapshot corresponding to the middle panel of Figure 5.6. This shows the general picture of clustered SNe driving a hot outflowing wind which contains a




Figure 5.8: Slices of temperature, outflow velocity, number density, pressure and magnetic plasma beta for a single time snapshot. Clustered SNe drive a hot outflowing wind which contains a population of cool clouds in this wind.




Figure 5.9: Slices of the velocities in the three axis-aligned directions. The vertical velocities v_z reflect the bulk outflow velocity, while the velocities v_x and v_y parallel to the disc plane reflect the strong turbulence in the wind.




Figure 5.10: Slices of the passive scalars we have used in this simulation. We use a total of 5 passive scalars. The first two are introduced right before the first SN explosion. All cold gas ($T < 5 \times 10^3 \text{ K}$) is dyed with passive scalar s_0 , and all cool gas ($5 \times 10^3 \text{ K} < T < 2 \times 10^4 \text{ K}$) is dyed with passive scalar s_1 . Passive scalars s_2 and s_3 track how much of the gas in a cell was *ever* hot or cold. Lastly, passive scalar s_4 tracks all SN injected gas. Beyond the disc and the injection site, s_2 and s_3 are close to unity everywhere, indicating that there is strong mixing between phases.




Figure 5.11: 2D volume weighted histograms of temperature, pressure, magnetic plasma beta, and turbulent mach number with height 14 Myr after the first SN. These reflect wind properties: $T \sim 10^7$ K, $P/k_b \sim 10^{4.5}$ cm $^{-3}$ K, $\beta \sim 10^2 - 10^3$, and $M_{\text{turb}} \sim 0.25$.

population of cool clouds, which we focus our analysis on in the next section. Figure 5.11 is composed of 2D histograms that show how these properties vary with height.

These are taken from a single time slice 14 Myr after the first SN, in the middle of the red shaded region shown in Figure 5.7. The histograms are volume weighted and hence largely reflect wind properties, especially at heights above 0.5 kpc. The top-most panel shows that the wind has a temperature of $T \sim 10^7$ K, along with cool clouds that have a temperature of $T \sim 10^4$ K. In the second panel from the top, it can be seen that pressure drops off rapidly away from the disc. The pressure in the wind is roughly $P/k_b \sim 10^{4.5} \text{ cm}^{-3} \text{ K}$. In terms of non-thermal components in the wind, both magnetic fields and turbulence are small compared to thermal pressure. We can see from the third panel that although the disc is strongly magnetized with plasma beta $\beta \sim 1$, the wind itself has a very weak magnetic field with $\beta \sim 10^2 - 10^3$. Lastly, the bottom-most panel shows that above 0.5 kpc, the turbulent mach number in the wind $M_{\text{turb}} \sim 0.25$. This brings us to a discussion of the nature of turbulence and mixing in the wind, where we will also define how M_{turb} is computed.

Dispersed multiphase flow systems such as particles or droplets in liquid or gaseous flows are common in various engineering fields (Balachandar & Eaton 2010). When the mass fraction of the dispersed phase is comparable to that of the carrier phase, the back reaction is non-negligible and the system is said to have two-way coupling. The nature of this multiphase turbulence is still considered an open problem due to the involved complexities. For example, it is unclear how particles can affect turbulence in the carrier phase—studies have shown generally that small particles can attenuate turbulence by dissipating energy for turbulent eddies, while larger particles can amplify turbulence through mechanisms such as vortex shedding (Gai et al. 2020). Likewise, it is apparent in our simulations that the presence of cloud fragments induces turbulence in the wind. This arises firstly during the breakout phase as the superbubbles propagates through the

multiphase ISM. Once this reaches the outflow phase, the flow of the hot wind around the population of clouds similarly induces a large scale turbulent field on the scale of the disc scale height. Figure 5.9 shows slices of the velocities in the three axis-aligned directions. The vertical velocities v_z reflect the bulk outflow velocity, while the velocities v_x and v_y parallel to the disc plane reflect the turbulence in the wind. In this particular slice, we can see that at low heights, the direction of the jet dominates the velocity field for v_x , but cascades to smaller scales at the wind flows out.

To examine the turbulence more closely, we apply a Gaussian filter to estimate the turbulent velocity v_{turb} . This method involves approximately removing large scale variations related to the bulk flow. This is done by convolving each velocity component with a Gaussian kernel to estimate the bulk velocity along that component. Subtracting this thus leaves the turbulent velocity of that component. While exact results depend on the kernel size, qualitative conclusions remain unchanged (Abruzzo et al. 2022b). We choose to use 16 pc, which sits comfortably within the range of sizes spanned by the cool clouds. The i -th component of the turbulent velocity is hence defined here as

$$v_{i,\text{turb}}(\mathbf{x}) = v_i(\mathbf{x}) - \frac{\iiint f_\sigma(\mathbf{x} - \mathbf{r})\rho(\mathbf{r})v_i(\mathbf{r})d^3\mathbf{r}}{\iiint f_\sigma(\mathbf{x} - \mathbf{r})\rho(\mathbf{r})d^3\mathbf{r}}, \quad (5.21)$$

where $f_\sigma(\mathbf{x})$ is the three-dimensional Gaussian filter. We weight the filter by density so that the velocities within the cloud are not dominated by the that of the hot gas. The right most panel of Figure 5.9 shows how the turbulent mach number using the above definition for turbulent velocity ($M_{\text{turb}} \equiv v_{\text{turb}}/c_s$) varies in the slices. In general, the turbulent mach number in the wind is high only close to the disc, and quickly weakens away from it. On the other hand, the cool clouds have transonic internal turbulence. This is shown more quantitatively in Figure 5.12, which shows a 2D histogram of M_{turb} and temperature T at the same time snapshot as Figure 5.11, but only considering heights




Figure 5.12: 2D histogram of turbulent mach number M_{turb} at heights above 0.5 kpc estimated using a filtering method to account for bulk flows and temperature T . The orange dashed line shows the average M_{turb} (~ 1) within the orange shaded region while the red dashed line shows the average M_{turb} (~ 0.25) in the red shaded region. These probe the cool clouds and the hot wind respectively.

above 0.5 kpc. The orange dashed line shows the average M_{turb} (~ 1) within the orange shaded region which represents the cool clouds, while the red dashed line similarly shows the average M_{turb} (~ 0.25) in the red shaded region which represents the hot wind. In all, there is significant turbulence both within the clouds and in the outflowing wind. This means that these environments are more similar to the turbulent boxes studied in (Gronke et al. 2022) rather than idealized laminar wind tunnel setups. This is critical because it is turbulence that drives mixing between the phases and hence their coupling.

Finally, Figure 5.10 shows slices of the passive scalars we have used in this simulation, and also illustrates the amount of mixing between phases happening in these winds. We use a total of 5 passive scalars. These passive scalars track gas with certain properties. The first two are introduced right before the first SN explosion. All cold gas ($T < 5 \times 10^3$ K) is dyed with passive scalar s_0 , and all cool gas (5×10^3 K $< T < 2 \times 10^4$ K) is dyed with passive scalar s_1 . This is to track how much cloud material was originally cold and cool ISM material. Passive scalars s_2 and s_3 track how much of the gas in a cell was *ever* hot or cold, similarly starting from right before the first SN, so as to quantify the amount of mixing. Lastly, passive scalar s_4 tracks all SN injected gas, which similar to s_0 and s_1 , tracks how much cloud material was from the SN. In the following section, we use s_0 , s_1 , and s_4 to quantify the original temperatures of cloud material. Here, we will just point out how beyond the disc and the base of the jet, s_2 and s_3 are close to unity everywhere. This means that all gas not close to the injection point or part of the ISM has at some point been cold or hot.

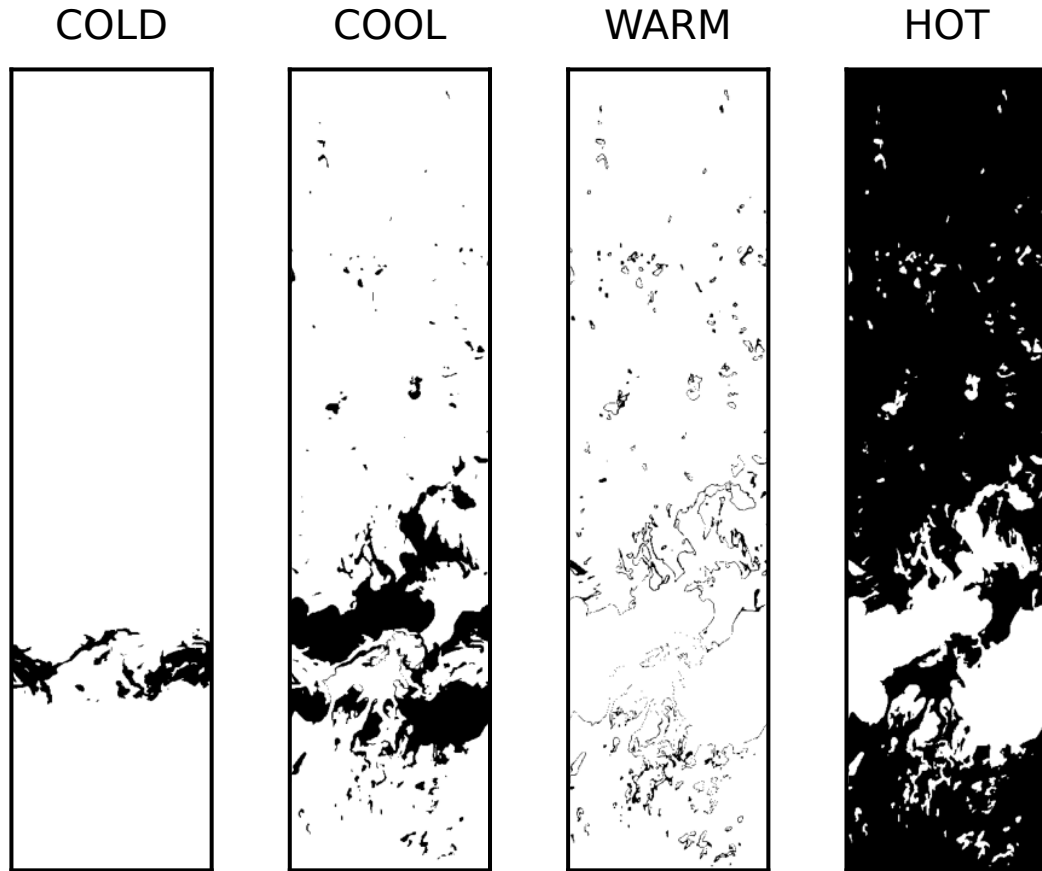


Figure 5.13: Distribution of the following gas phases: hot gas ($T > 5 \times 10^5$ K), warm gas (2×10^4 K $< T < 5 \times 10^5$ K), cool gas (5×10^3 K $< T < 2 \times 10^4$ K) and cold gas ($T < 5 \times 10^3$ K). Cold gas is located within the ISM while cool gas extends to clouds that are contained in the wind, which is made up of the volume filling hot phase. Warm gas exists mainly in interfacial mixing layers.

5.4.4 Mass and Energy Flux Phase Distribution and Cloud Entrainment

Finally, we can look at the fraction of the wind’s mass and energy flux carried by the different phases. To do so, we first concretely define four phases of interest. These phases are the hot gas ($T > 5 \times 10^5$ K), the warm gas (2×10^4 K $< T < 5 \times 10^5$ K), the cool gas (5×10^3 K $< T < 2 \times 10^4$ K) and the cold gas ($T < 5 \times 10^3$ K). Figure 5.13 shows the general location of each of these phases. Cold gas is only located within the ISM. Cool gas is found in the ISM and the clouds that are contained in the wind, which is made up of the volume filling hot phase beyond the disc. Warm gas is found in the interfacial mixing layers.

Figure 5.14 shows the temperature velocity distributions at different heights above the disc over the time period identified in Figure 5.7, weighted by mass flux in the top row and energy flux in the bottom row. The PDFs $f_M(u, w)$ and $f_E(u, w)$ are weighted by mass and energy flux respectively, where $u \equiv \log(v_{\text{out}})$ and $w \equiv \log(T)$ in units of km/s and K. The PDFs f_q for each quantity q is defined as $\frac{1}{\langle q \rangle} \frac{d^2 q}{du dw}$ where $\langle q \rangle$ is the average over both time and the horizontal slice, such that $\int f_q du dw = 1$. We define the mass flux to be ρv_z and the energy flux to be $E_z v_z$ where $E_z = \rho(\frac{1}{2}v_z^2 + e_{\text{th}})$ and e_{th} is the specific internal energy. The horizontal grey dashed lines demarcate the boundaries between the phases defined above. Consistent with previous work, the hot phase dominates the energy loading while there is significant mass loading by the cool phase (Fielding et al. 2018; Kim et al. 2020a,b). The grey dotted line shows the sound speed of the gas at each temperature—the outflow velocity of the wind is transonic with a Mach number close to unity. We also see that the cool material gets more entrained at increasing heights, indicating that there is momentum transport from the hot to the cool phase, which has been found to be driven by mixing rather than ram pressure acceleration (Melso et al.

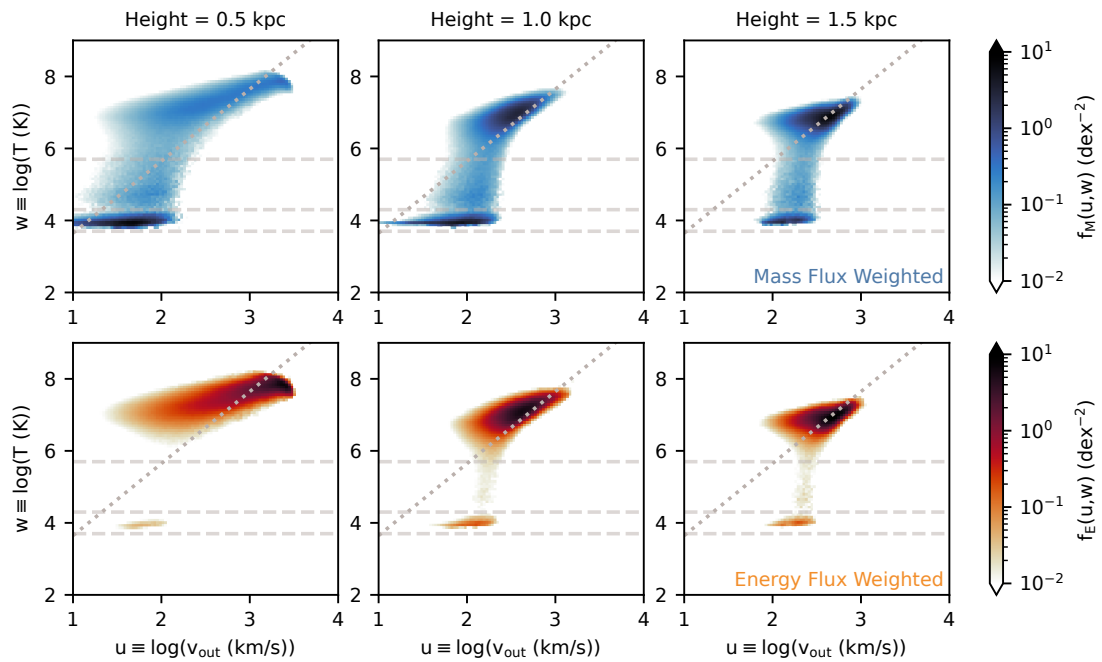


Figure 5.14: Temperature-velocity joint PDFs at different heights above the disc, weighted by mass flux in the top row and energy flux in the bottom row. The horizontal grey dashed lines demarcate the boundaries between the phases defined above. The hot phase dominates the energy transport while there is significant mass contained in the cool phase. The grey dotted line shows the sound speed of the gas at each temperature—the wind is transonic with a Mach number close to unity. Cool material gets more entrained at increasing heights.

2019; Schneider et al. 2020; Tonnesen & Bryan 2021).

5.5 Results : Clouds

In this section we focus our attention on analysis of the clouds embedded within the hot wind and their properties. We first discuss how we define and identify clouds, followed by how we measure related cloud properties. We then compare these measurements directly with the model discussed in Section 5.2. This includes size distributions and various scaling relations associated with cloud growth rates. Finally, we explore other interesting aspects of cloud properties relating to non-thermal support, material origins, and wind alignment.

5.5.1 Cloud Identification

In order to map out and catalogue the properties of clouds embedded in the hot outflowing wind, we must first define them. We identify clouds by applying a threshold at $T = 10^{5.5}$ K, represented by the red dashed line the upper panel of Figure 5.15. Any gas below this threshold is thus considered ‘cloud material’. The lower center panel of Figure 5.15 shows this cool material identified from the corresponding temperature slice. While we use a fixed threshold, we find that cloud identification in our simulations is generally insensitive to the exact value used.²¹ We can understand this physically—most cloud gas is at a temperature $\sim T = 10^4$ K, while most of the wind is at $T > 10^6$ K. Gas at intermediate temperatures mainly exists only at the thin interfaces between these two phases in turbulent radiative mixing layers. This can be seen visually in the temperature

²¹Instead of arbitrarily choosing a threshold value, we also explored using a non parametric approach such as Otsu’s image segmentation algorithm (Otsu 1979), which determines a threshold for separating two classes by maximizing (minimizing) the variance between (within) them, in our case for the distribution of $\log(T)$. This approach gives a similar value close to $\log(T(K)) = 5.5$, but the exact value varies with different time slices.

slice in the lower left panel of Figure 5.15 (also see the ‘warm’ panel in Figure 5.13). It is hence not surprising that the identification of clouds is understandably robust to the choice of threshold provided that this threshold temperature lies in the mixing region. The simple approach of using a fixed threshold is thus as effective as it is because of the strong biphasic nature of the wind and the clouds. This can be seen in the upper panel of Figure 5.15 which shows the temperature distribution in the wind at a single time snapshot, along with the threshold value used. Using a constant threshold also allows for more consistency over multiple time slices. We can thus define a single cloud as a collection of cells that fall below the threshold and that are interconnected (where connected is defined as sharing at least one corner). Except for the analysis on cloud size distributions, where we account for clouds that intersect the periodic dimensions and hence wrap around, we ignore clouds that are touching the borders of our simulation. This means we naturally exclude disc material. We also only consider clouds that have a volume greater than 16 cells. The bottom right panel of Figure 5.15 shows an example of these clouds arbitrarily colored by their cloud identification number. We also compile our catalog over multiple time snapshots (a total of 10 spanning 8 Myr).

5.5.2 Measuring Cloud Properties

For each identified cloud, we measure a number of properties as follows. The position of the cloud is computed as the volume weighted centroid of cloud material. Individual cloud properties such as velocity \mathbf{v}_{cl} , density ρ_{cl} , thermal and magnetic pressures P_{cl} and $P_{\text{B,cl}}$ are volume averaged over cool ($< 1.01 \times 10^4$ K) gas in the cloud. Temperature is density weighted, so that $P_{\text{cl}} = \rho_{\text{cl}} T_{\text{cl}}$. Passive scalars are also density weighted. We define the local background environment of a cloud as any hot ($> 10^6$ K) wind gas within a cube centered on the cloud that has a length on each side equal to twice that of the




Figure 5.15: *Top panel:* Histogram of gas cell temperatures from a single snapshot. Most of the gas is in the volume filling hot wind, with a peak at $T = 10^4$ K corresponding to cool clouds. The red dashed line represents our threshold value for identifying cloud material. *Bottom Panel:* From left to right: slices plots of temperature, cloud material, and identified clouds.

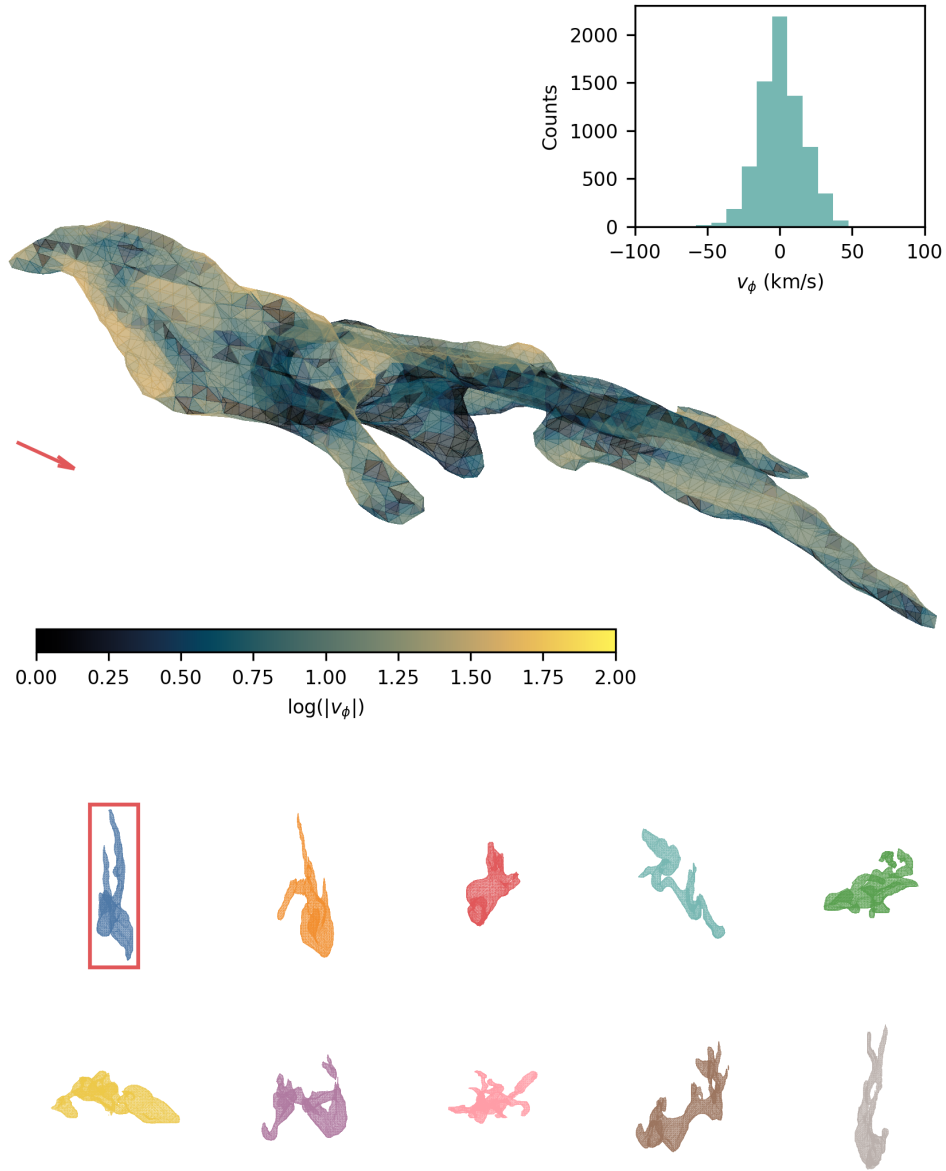


Figure 5.16: *Bottom Panel:* Meshes representing the temperature isosurface of a range of identified clouds. The cloud in the top left is shown in greater detail in the top panel. *Top panel:* Zoom-in of a cloud surface colored by \mathbf{v}_ϕ . The wind direction relative to the cloud is denoted by a red arrow. The inset shows the surface area weighted histogram of \mathbf{v}_ϕ . v_{turb} is the standard deviation of this distribution.

longest axis of the tight bounding box that contains the cloud, and also require that there are at least 16 such cells in this volume. Corresponding wind properties such as the background wind velocity \mathbf{v}_w are averaged in the same way as cloud properties. The turbulence in the wind $\mathbf{v}_{\text{turb},w}$ is taken to be the rms velocity in the wind frame in this environment.

We adopt a geometric approach towards characterizing the turbulence on the surface of a cloud, which involves measuring quantities on a temperature isosurface constructed using the marching cubes algorithm described in Lewiner et al. (2003). This isosurface is represented by a mesh of triangular faces. Figure 5.16 shows examples of this mesh in the lower panel for various clouds. The cloud in the upper left is shown in greater detail in the upper panel.

To measure the turbulence in the mixing layer on this surface, we use the method described in Abruzzo et al. (2022b), which they demonstrate provides a measure of turbulence that is comparatively robust to gradients in the laminar component of the background wind as compared to the other approaches of characterizing turbulence that they explored (filtering methods and velocity structure functions). The outline of the method is as follows: Consider a point on the temperature isosurface that is at the centroid of one of the triangular faces which constitute the mesh. Working in the cloud frame, we linearly interpolate the velocity \mathbf{v} and the logarithmic density ρ at that point. We define a set of axes $\hat{\mathbf{n}}, \hat{\mathbf{w}}$ and $\hat{\phi}$, where $\hat{\mathbf{n}}$ is the inwards direction normal to the surface, $\hat{\mathbf{w}}$ is the direction of the component of the background wind v_{wind} in the surface plane ($\hat{\mathbf{w}} \equiv \mathbf{v}_w / |\mathbf{v}_w|; \mathbf{v}_w \equiv \mathbf{v}_{\text{wind}} - (\mathbf{v}_{\text{wind}} \cdot \hat{\mathbf{n}})\hat{\mathbf{n}}$), and $\hat{\phi}$ is the direction on the surface orthogonal to both $\hat{\mathbf{n}}$ and $\hat{\mathbf{w}}$, i.e., $\hat{\phi} \equiv \hat{\mathbf{n}} \times \hat{\mathbf{w}}$. These definitions are physically motivated—the normal direction captures turbulent radiative mixing layer driven accretion, while the wind direction captures the coherent flow of the background wind. Picking out the $\hat{\phi}$ direction thus allows us to disentangle the turbulence from these flows. The velocity component

is hence simply $\mathbf{v}_\phi \equiv \mathbf{v} \cdot \hat{\phi}$. We thus estimate v_{turb} to be the area weighted standard deviation of \mathbf{v}_ϕ . The upper panel of Figure 5.16 shows how \mathbf{v}_ϕ varies over the surface of a cloud. The top right inset shows a surface area weighted histogram of \mathbf{v}_ϕ on this surface. As expected, the distribution is normal and centered on zero. v_{turb} for this cloud is simply the standard deviation of this distribution. The inflow velocity through this surface is similarly given by $\mathbf{v}_{\text{in}} \equiv \mathbf{v} \cdot \hat{\mathbf{n}}$.

In Appendix E, we detail how we estimate cloud growth rates \dot{m} , comparing the effectiveness of various methods of doing so. Ultimately, we use the net inwards mass flux on the temperature isosurface ($\dot{m}_{\text{surface}} \equiv \sum_{\text{surface}} \rho v_{\text{in}} \Delta A$) with a scaling factor calibrated to match the total cooling luminosity, so $\dot{m} \sim f_{\text{scale}} \dot{m}_{\text{surface}}$. Appendix E also presents an important caveat that these methods fail when the cloud is not actually growing, and thus do not reflect the transition to the cloud destruction regime. Instead, we show that we can estimate this scale from the cloud size distribution.

5.5.3 Cloud Model

Now that we have discussed how we identify clouds and measure their properties, we are ready to make specific comparisons between these measured properties with the components of the model as outlined in Section 5.2. We first compare the analytic formulation for the distribution of cloud sizes with that of the simulations. We then compare the cloud growth model with the predicted scalings for the surface area and associated mass flux as given in equations (5.1) – (5.8).

Distribution of Cloud Sizes

We begin by looking at the distribution of cloud sizes in the wind. In Section 5.2, we discuss how we expect a distribution of $dN/dm \propto m^{-2}$ due to turbulent fragmentation.




Figure 5.17: Cumulative cloud volume distribution at various time slices. The top panel is from the fiducial run, while the bottom panel is from the failed wind run. The pink regions indicate where clouds are resolved by less than 16 cells. Grey regions show the expected power law scaling of -1 . Grey dashed lines show expected lower (for survival) and upper bounds on cloud sizes.

This is equivalent to a cumulative distribution that scales as $N(> m) \propto m^{-1}$. Instead of a mass scaling, we consider instead the volume scaling $dN(> V)/dV$ since the fragmentation process is not driven by a density dependent process, although the difference made by this choice does not affect the results.

Figure 5.17 shows the cumulative size distribution of clouds at different times in the simulation. We show these distributions for two different runs—our fiducial run and the run with the failed wind discussed in the previous section. As a reminder, the latter has a higher resolution of 1 pc. Pink regions indicate where clouds are resolved by less than 16 cells (and hence not included in the catalog of cloud properties). Grey regions show the -1 power law scaling for comparison.

The upper panel shows distributions that start from when the superbubble first breaks out from the disc. The number of clouds in the wind increases over time as the disc continues to fragment. We note that the $N(> V) \propto V^{-1}$ power law distribution is seen even at early times, supporting the argument that this is a result of the disc fragmentation process during the process of superbubble breakout. The lower panel covers a narrower time range that corresponds to the analysis window we identified in the previous section where the wind briefly breaks out. Here, the cumulative size distribution appears slightly steeper and is more consistent with $N(> V) \propto V^{-1.2}$. It is possible that this slightly steeper slope is related to the higher resolution and presence of magnetic fields, but more likely it is due to the brief and weak nature of the wind—these clouds show very different growth scalings which we discuss below, and as seen in the previous section, the entire wind quickly cools and falls back to the disc. We have checked that these distributions do now show any consistent trends with height in these simulations.

Both distributions exhibit a flattening at small scales somewhat above the resolution limit, which is expected for clouds below the critical radii for sustained growth discussed in Section 5.2. To estimate these critical radii we take a typical set of parameters for these

winds. For example, in our fiducial simulation, these are $M_{\text{wind,rel}} \sim 0.4$, $M_{\text{turb,wind}} \sim 0.5$, $P_3 \sim 30$ and $\chi \sim 400$. Figure 5.18 shows the distributions of $M_{\text{wind,rel}}$ and χ along with their means. This gives a cloud survival size of $r_{\text{crit, shear}} \sim 5$ pc. We do the same for the simulation with the failed wind. These nicely match the scales below which we see the distributions flatten. The vertical grey dashed lines in Figure 5.17 show the upper and lower bound volume estimates corresponding to $r_{\text{crit, shear}}$ and the disc scale height $H \sim 66$ pc respectively. Ideally, the smallest surviving clouds should be resolved by multiple cells to achieve first order convergence in terms of total cloud mass. With a resolution of 2 pc and 1 pc, we resolve $r_{\text{crit, shear}}$ by a factor of ~ 2 and ~ 4 .

Cloud Surface Area

The isosurface method detailed above allows us to compare the scaling of the cloud surface area as a function of cloud size to the model scaling in equation (5.6). We take the cloud surface area A_{cl} to be the area of the isosurface. We define this to be with respect to a measurement scale of 2 pc when constructing the isosurface, since the isosurface area varies with this choice of scale (Fielding et al. 2020). We do not expect this area to be the same as A_{cloud} from equation (5.6), since its magnitude in part depends on the choice of scale and temperature of the isosurface (refer back to Section 5.2 for more discussion), but we do expect the scaling with size to remain the same. Figure 5.19 shows a 2D histogram of A_{cl} and cloud radius r_{cl} (the radius of a sphere with equivalent volume), along with the $A_{\text{cl}} \propto V^{5/6} \propto r_{\text{cl}}^{5/2}$ scaling from equation (5.6). We find that most of the clouds agree with this scaling, except for the largest clouds which appear to have a steeper slope closer to $A_{\text{cl}} \propto r_{\text{cl}}^{11/4}$. We discuss this further below.




Figure 5.18: Distributions of the Mach number of the wind relative to the cloud (top) and cloud overdensity relative to the wind (bottom). Dashed lines show population means.




Figure 5.19: Surface area of clouds plotted against their size. Most clouds follow a $V^{5/6}$ or $r_{cl}^{5/2}$ scaling denoted by the dashed orange line, although this relation appears to steepen for a handful of the largest clouds.

Cloud Growth

We can also see how well the scalings in equations (5.7) and (5.9) hold up. For reference, we expect the inflow velocity v_{in} to scale as

$$v_{\text{in}} \propto v_{\text{turb}}^{3/4} \left(\frac{L_{\text{turb}}}{t_{\text{cool}}} \right)^{1/4} \quad (5.22)$$

in the strong cooling regime where the mixing time $L_{\text{turb}}/v_{\text{turb}}$ is longer than the cooling time of mixed gas $t_{\text{cool,mix}}$ (Tan et al. 2021). For a typical cloud, $T_{\text{mix}} \sim 10^{5.5}$ K and $P_3 \sim 30$, giving $t_{\text{cool,mix}} \sim 0.15$ Myr. In the mixing layers, $v_{\text{turb}} \sim c_{\text{s,cool}} \sim 15$ km/s. This suggests that we are in the strong cooling regime if $L_{\text{turb}} > 2$ pc. We discuss this in further detail shortly, but L_{turb} is usually taken to be the outer scale of the turbulence, where mixing is the most effective, hence we expect our clouds to comfortably fall within this strong cooling regime. In fact, this criterion from Tan et al. (2021) is likely oversimplified since the cooling curve drops off sharply above 10^5 K—it is reasonable to expect that the requirement of L_{turb} might be even smaller.

To compare cloud properties with these scalings, we define

$$v_{\text{in}} \equiv \dot{m}/(\rho_{\text{wind}}A_{\text{cl}}) \quad (5.23)$$

following equation (5.5) but using the isosurface area and the associated mass flux through it. Figure 5.20 demonstrates the v_{in} scaling with, r_{cl} , v_{turb} , and P_{w} respectively. Here we use P_{w} as proxy for the characteristic mixing layer cooling time t_{cool} since it scales inversely with P_{w} , and we will discuss how r_{cl} relates to L_{turb} . More specifically, each panel show the scaling of v_{in} divided by the expected scalings of the other variables so as to remove any cross-dependencies. Each grey point in the panels represents individual clouds, while red crosses mark a corresponding binned scatter plot to more easily visualize




Figure 5.20: Scalings of the inflow velocity v_{in} with various parameters—the cloud size, turbulent velocity, and pressure (as a proxy for the cooling time). Observed scalings appear to match theoretical expectations.

trends in the data (Cattaneo et al. 2019). Orange lines show the scalings we expect from our model in Section 5.2.

The top panel of Figure 5.20 shows $v_{\text{in}}/(v_{\text{turb}}^{3/4}P_w^{1/4})$ plotted against r_{cl} . In previous work, it has been common to set the outer turbulent length scale L_{turb} to be either the box size (in the case of turbulent mixing layers) or the cloud size (for cloud simulations). In the latter, L_{turb} is in fact set to the *initial* cloud size and fixed for the rest of the simulation, even if the cloud grows in mass and volume by over an order of magnitude or fragments into many pieces. Here, we expect that if $L_{\text{turb}} \sim r_{\text{cl}}$, then we should see that $v_{\text{in}}/(v_{\text{turb}}^{3/4}P_w^{1/4}) \propto r_{\text{cl}}^{1/4}$. However we see in Figure 5.20 that there is no scaling with r_{cl} , save for the slight drop-off at large cloud sizes due to the steeper scaling of the surface area there.

Despite this, we find that v_{in} still scales with cloud size, but this is because in our measurements, $v_{\text{in}} \propto v_{\text{turb}}^{3/4} \propto r_{\text{cl}}^{1/4}$. Figure 5.21 shows this scaling, which is consistent with subsonic Kolmogorov turbulence. There are two possible interpretations here. The first is that our method of measuring v_{turb} is sensitive to the scale of the cloud, and that the scaling comes from measuring velocity differences at the surface on cloud scale separations. The second, which we lean towards, is that the turbulent length scale here is not set independently by each individual cloud, but is instead a common scale over all the clouds. This is motivated by the observation that all the clouds are evolving in the same common turbulent wind, where the turbulent velocity field has been driven on the outer scale of the entire system of clouds (\sim the disc scale height H). The idea that L_{turb} is set by the outer turbulent scale and not the size of each individual cloud is also consistent with previous work. For example, Gronke et al. (2022) use the initial cloud size as L_{turb} for a turbulent box. They note that this choice becomes ambiguous in their simulations due to mass growth and cloud fragmentation/coagulation—in principle, at late times, L_{turb} approaches the driving scale of turbulence in their simulation L_{box} .




Figure 5.21: The turbulent velocity measured on the cloud surface scales with the cloud size in a manner that is consistent with subsonic Kolmogorov turbulence.

Abruzzo et al. (2022b) directly measure velocity structure functions around wind tunnel clouds and also find that $L_{\text{turb}} \sim$ the initial cloud size even as their clouds grow.

The remaining panels in Figure 5.20 hence assume that there is no explicit dependence of v_{in} on r_{cl} besides that from v_{turb} . The middle panel shows the scaling of $v_{\text{in}}/P_{\text{w}}^{1/4}$, which is consistent with the expectation of scaling as $v_{\text{turb}}^{3/4}$. Again, there is a drop-off at higher values of v_{turb} , but this is similarly attributed to the higher surface area for the largest clouds that we divide by to get v_{in} (leading to a slope shallower by a power of $-1/4$ since v_{turb} scales with cloud size). The lowest panel of Figure 5.20 shows the scaling of $v_{\text{in}}/v_{\text{turb}}^{3/4}$ with P_{w} , where we see the $1/4$ scaling with cooling time characteristic of turbulent mixing layer growth in the strongly cooling regime.

Finally, combining all of the above, we can put in some numbers and directly compare to the model presented in Section 5.2 for the mass growth rate of clouds (equations (5.5), (5.6) and (5.9)). The orange dashed line in Figure 5.22 shows our model with the typical values in our wind: $P_3 \sim 30$, $T \sim 10^7$ K, $L_{\text{turb}} \sim H \sim 66$ pc, and $r_{\text{crit, shear}} \sim 5$ pc. For v_{turb} , we use the orange dashed line from Figure 5.21, so $v_{\text{turb}} \sim 10(r_{\text{cl}}/10 \text{ pc})^{1/3}$ km/s. Note there are no scaling factors here. We find that the model agrees extremely well with the estimated mass growth rates of the clouds in the simulation. Here, it is important to highlight some complications in the simulation data here that are not immediately obvious. Our model predicts an overall scaling of $\dot{m} \propto r_{\text{cl}}^{11/4}$, arising from $A_{\text{cloud}} \propto r_{\text{cl}}^{5/2}$ and $v_{\text{in}} \propto r_{\text{cl}}^{1/4}$. However, for large clouds, we have seen that the entirety of this r_{cl} scaling is encapsulated by the surface area scaling. The model still working well here suggests that either the clouds are growing slower than expected given their surface area, or that our measurement of the surface area of these largest clouds overestimates the effective surface area. In addition, we have assumed that ρ_{wind} is independent of r_{cl} , but we find that there is a selection effect stemming from clouds in higher pressure environment being found in patches of the wind which are being compressed. Since these patches are smaller




Figure 5.22: Mass growth rates compared to the analytic model presented in Section 5.2, which is shown by the orange dashed line, show good agreement.




Figure 5.23: The inflow velocity here is shown to scale linearly with pressure in the simulation run which fails to generate a strong wind.

in scale, corresponding clouds are likewise smaller. This artificially makes the slope of \dot{m} in Figure 5.22 visibly shallower at small cloud sizes.

‘Growth’ in a Failed Wind

It is also interesting to see what happens when the superbubble breakout fails to generate a strong wind. As seen in Figure 5.6, after the superbubble breaks out briefly and vents the hot gas, the hot material quickly begins to cool and falls back to the disc. Repeating the analysis above, we find that while the surface area to volume scaling above




Figure 5.24: The turbulent velocity on the cloud surface scales with the relative velocity between the cloud and the wind when shear is high, but levels off as clouds get entrained to the sound speed of the cool gas.

still holds, v_{in} in the cooling wind is independent of both cloud size and turbulent velocity, suggesting that the cool clouds are not growing via the same turbulent mixing mechanism. In Figure 5.23, we show that instead, v_{in} scales linearly with P_w (i.e., inversely with the cooling time), consistent with the picture above that the growth is instead directly driven by large scale cooling and condensation of the wind.

Dependence of Turbulent Mixing on Wind Shear

Lastly, we look at the scaling between the turbulent velocity v_{turb} in the mixing layer on the cloud surface with the relative shear velocity between the cloud and the surrounding wind v_{rel} . As discussed in Section 5.2, plane parallel turbulent mixing layer simulations find that $v_{\text{turb}} \propto v_{\text{rel}}$, close to linear. However, cloud crushing simulations observe that turbulence persists even when the cloud is entrained. We find results consistent with both findings and well represented by equation (5.8). This is shown in Figure 5.24. At high shear (high v_{vel}), there is a strong scaling of v_{turb} with v_{rel} shown by the orange dashed line, consistent with $v_{\text{turb}} \sim f_{\text{rel}} v_{\text{rel}}$ with $f_{\text{rel}} \sim 1/15$. However, as v_{rel} gets smaller, v_{turb} levels off at roughly the sound speed of the cool gas, shown by the teal dashed line. This supports the picture that the KH instability is but one way of generating turbulence in the mixing layer, and that turbulence can continue to persist and drive mixing even when clouds become entrained in the wind.

5.5.4 Non-thermal Pressure Support

Figure 5.25 illustrates the lack of thermal pressure balance between the clouds and the background wind and the contributions of the two main sources of non-thermal pressure support in the clouds: magnetic pressure and turbulent pressure. The top left panel shows a slice of thermal pressure normalized by the mean over the entire slice. The clouds appear generally as under-pressured regions relative to their surroundings. The top right panel shows this more clearly, where the orange dashed line shows the mean pressure of cool gas in the orange shaded region, and the red dashed line represents the mean pressure in the hot gas marked by the red shaded region. The strong dip in the intermediate region where the cooling time is short indicates that the mixing layer itself is not well resolved (Fielding et al. 2020). From left to right, the bottom row of




Figure 5.25: *Top Row:* On the left, we show a slice of pressure normalized by the mean. We see that clouds appear generally under-pressured as compared to their surroundings. This is shown more clearly on the right, where the orange dashed line shows the mean pressure of cool gas in the orange shaded region, and the red dashed line represents the mean pressure in the hot gas marked by the red shaded region. *Bottom Row:* From left to right: histograms of the ratios of (i) the ratio of thermal pressures in the cloud and the surrounding wind, (ii) the ratio of magnetic pressure and thermal pressure in the cloud, and (iii) the ratio of turbulent pressure and thermal pressure in the cloud.

histograms show ratios of (i) thermal pressures in the cloud and the surrounding wind, (ii) magnetic pressure and thermal pressure in the cloud, and (iii) turbulent pressure and thermal pressure in the cloud. From the upper right and lower left panels, we see that just looking at thermal pressures, clouds are under-pressured relative to their environment by a factor of 2. Magnetic pressure support is not large enough to make up the different, with the magnetic plasma beta only being ~ 4 . This is despite having $\beta \sim 1$ in the ISM prior to the onset of the SN driven wind. Instead, the missing pressure support is provided by turbulent pressure, which we define as $\rho_{\text{cl}}\langle v^2 \rangle_{\text{cl}}$. The turbulent pressure is roughly equal too the thermal pressure (which is equivalent to having a turbulent velocity equal to the sound speed of the cool gas). Hence, pressure support in the cloud is provided mostly in equal parts by thermal and turbulent pressure support, with a minor contribution from magnetic pressure. Taken together the *total* pressure of the clouds is, on average, equal to the hot wind pressure.

5.5.5 Cloud Origins: Passive Scalars

The passive scalars we have employed reveal some interesting points about the origins of these clouds. Figure 5.26 shows a histograms of concentrations of passive scalars in our clouds which track the amount of cloud material that was originally (right before the first SN) cold ISM gas, cool ISM gas, or mass injected by SN events. We find that most of the cloud is comprised of gas that was originally part of the cool ISM, supporting the fragmentation origin of the cloud, with a small fraction originating from the cold ISM gas. It should also be noted that in both cases, this gas likely mixed and cooled to different temperatures, either in the wind or cloud, as evidenced by the other passive scalars in Section 5.4 which showed that almost all gas had at some point been either cold or hot.




Figure 5.26: Histogram showing that most clouds are comprised of gas that was originally part of the cool ISM pre-SN, with a small amount from the cold ISM. The contribution from SN injected mass is negligible.




Figure 5.27: The angular separation between cloud and wind velocities is well described by an exponential distribution with a scale angle of $\sim 25^\circ$.

5.5.6 Cloud-Wind Alignment

Idealized setups involving clouds accelerated in a laminar wind or infalling under gravity often show an initial spherical cloud developing a pronounced head-tail structure, a morphological prediction that observations have widely been compared to (e.g., Putman et al. 2012). Clouds being accelerated in a highly turbulent wind exhibit far more complex structures, with the turbulence generating ‘tails’ that result in a rich diversity in cloud morphologies, as can be seen in the lower panel of Figure 5.16, where all clouds are pictured with the bulk flow in the upwards direction. Only a minority show similarities

to classic head tail structures. In general, the head-tail paradigm serves only as a first order indicator of the direction of bulk velocities. We can quantify the alignment between clouds and a turbulent wind by looking at the angular separation between cloud and local wind velocities. This is shown in Figure 5.27, where we find that this angular separation follows an approximately exponential distribution with a scale angle of $\sim 25^\circ$. While most clouds are well-aligned with the wind, a handful are significantly misaligned. This is consistent with observations of clouds in the Fermi Bubbles (Di Teodoro et al. 2020 for example look at two aligned clouds, but to our knowledge there have not been any studies of the misaligned ones).

5.6 Discussion

5.6.1 Connection to Small Scale Simulations

It is exciting that the physical insights and analytic scalings derived from small scale idealized simulations translate well to larger scales. This was first demonstrated in Tan et al. (2023) for individual infalling clouds, and now for a population of clouds in turbulent magnetized outflows in this work. In particular, these models build from the ground up upon the back of work on the physics of interfacial turbulent radiative mixing layers (e.g., Ji et al. 2019; Fielding et al. 2020; Tan et al. 2021), minimizing the need for empirical tuning of parameterized models. This success is promising for the future use of physically informed subgrid models in larger scale simulations. Nonetheless, some nebulous points remain—we briefly discuss issues of cloud survival/destruction and magnetic fields here, and limitations and caveats in a later part of this section.

In systems where the density contrast χ is high ($\geq 10^3$), there is still remaining uncertainty around the cloud survival criterion. In this work, we have presented a criterion

based on the idea that the characteristic cooling time of the mixing layers must be significantly shorter than the shear timescale that is motivated by and consistent with results of wind tunnel simulations at high χ (Abruzzo et al. 2022a). This criterion predicts clouds must be larger than the Gronke et al. (2022) criteria by a factor of $\sqrt{\chi}$ in order to grow, and is consistent with phenomenological criteria (e.g., Sparre et al. 2020; Li et al. 2020). However, there remains further investigation required to pin down the relevant physical processes that determine cloud survival in high χ environments like multiphase outflows.

We do not investigate cloud mass loss rates in this work. The methods we use to estimate cloud growth rates break down for clouds that are getting destroyed (see Appendix E). Incorporating cloud destruction rates into this model would improve the coupling between the phases, and is especially important in parameter regimes where clouds are not expected to grow. Similarly, this also requires more work on the small scale simulation side.

A final point on the topic of cloud survival, while not directly relevant to this work but important to keep in mind, is that Farber & Gronke (2022) show that the story is more complicated when $T_{\text{cloud}} < 10^4$ K. This is mainly due to the shape of the cooling curve, which Abruzzo et al. (2022a) also showed plays an important role. However, the survival criteria they find is equivalent to the Gronke & Oh (2018) criterion when $T_{\text{cloud}} \gtrsim 10^4$ K as it is in our simulations.

What about magnetic fields? This perennial question has been the subject of many wind tunnel simulations of cool clouds (e.g., Gregori et al. 1999; McCourt et al. 2015; Grønnow et al. 2018; Hidalgo-Pineda et al. 2023). Despite that, their effect remains unclear. For example, Gregori et al. (1999) find that destruction can be enhanced by more rapid acceleration while McCourt et al. (2015) find that they aid survival via magnetic draping inhibiting shear instabilities (see also Banda-Barragán et al. 2018; Ji et al. 2019; Grønnow et al. 2022). Sparre et al. (2020) find some enhancement in cloud survival for

$\beta \sim 10$, while Li et al. (2020) find no effect for $\beta \sim 10^6$. Hidalgo-Pineda et al. (2023) find that for small β (~ 1), much smaller clouds are able to survive. However, β is large in our winds and only significant in the disc—we hence conclude that it is unlikely that magnetic fields play a significant role in influencing cloud acceleration or survival in multiphase outflows. Magnetic fields don’t seem to inhibit mass growth rates either (Gronke & Oh 2020a; Hidalgo-Pineda et al. 2023), although morphologically clouds are reported to have very different filamentary structures as compared to their hydrodynamical counterparts (e.g., Tonnesen & Stone 2014; Gronke & Oh 2020a; Jung et al. 2023). Our model does not include magnetic fields, supporting their lack of impact on growth rates. Understanding why this is so is an interesting avenue for future work. While we do not compare to a hydrodynamical run without magnetic fields, we do not observe clear filamentary structures—which we attribute to a combination of a weakly magnetized wind and turbulence. In general, turbulence in the wind is the main generator of complex morphologies seen in the clouds.

5.6.2 Implications for Galaxy/Cosmological Scale Simulations

Having demonstrated that much of the insight garnered from small-scale simulations translates to more realistic larger-scale systems we can now address how these processes fit into the overall landscape of galaxy formation and global-scale simulations. The impact of capturing these multi-scale multiphase effects has been seen in isolated galactic scale simulations with self-consistently generated multiphase winds, which find that properties of the hot wind including temperature, density, and pressure fall off slower than expected with distance, and also travel slower than single-phase adiabatic winds (Fielding et al. 2017b; Schneider et al. 2020). These effects are consistent with expectations of mass, momentum and energy exchange between cool clouds and the surrounding wind (Fielding

& Bryan 2022). The total cool gas mass flux in Schneider et al. (2020) decreases with distance, suggesting that the cool gas is being destroyed and mass loading the hot phase, even for largest clouds. This may be understood by applying the $r_{\text{crit, shear}}$ criterion for cloud growth as opposed to $r_{\text{crit, cc}}$. In addition, they find that cool clouds are underpressurized by up to a factor of 10. While we find that turbulent pressure support is significant, this cannot fully account for the large factor. Given that we also find that regions of phase space where the cooling time is shortest are the most underpressurized, this can likely be attributed to lower resolutions (Fielding et al. 2020). Additional detailed studies of multiphase winds at even higher resolutions (or with conditions in which it is easier to resolve $r_{\text{crit, shear}}$) will be helpful in shedding further light on how the multiphase interactions shape the overall evolution of the winds.

Recent large cosmological simulations also exhibit galactic winds with self-consistently generated cool phases. While most have insufficient resolution to resolve any clouds, certain zoom-in and high resolution simulations can do so, albeit marginally. For example, multiphase galactic winds are seen in TNG50 (Nelson et al. 2019) and FIRE-2 (Pandya et al. 2021). The latter characterized the multiphase nature of the outflow by analyzing the contribution to fluxes from each phase and found results broadly consistent with our findings and past tall-box ISM patch simulations (e.g., Fielding et al. 2018; Kim et al. 2020a).

Correctly capturing the multiphase nature of galactic winds is not only essentially for accurately modeling the winds themselves, but as recent work has shown, may also be essential for capturing the correct regulation of star formation and thus galaxy evolution. When winds are able to separate into multiple phases, the hot phase, which has very high specific energy, heats and stirs the CGM, which prevents new star forming material from entering the galaxy (Fielding et al. 2017a), and the cold phase ejects material directly out of the ISM. Standard galactic feedback models cannot capture the high specific energy

phase in particular and as a result may be missing important regulation mechanisms, as was recently shown using regulator and semi-analytic models (Carr et al. 2022; Pandya et al. 2022), as well as isolated galaxy simulations (Smith et al. 2023). It is uncertain if cosmological simulations will ever achieve the resolutions necessary to properly capture the physical scales relevant to multiphase wind launching and dynamics. In which case, simulations and models such as those presented in this work are important in being able to bridge the gap towards achieving the capability to include realistic multiphase outflows using subgrid techniques.

5.6.3 Implications for Observations

A key missing piece in our understanding of galaxy evolution is the amount of mass and energy carried by winds from the ISM into the CGM and beyond. Most observations of galactic winds come from probes that are sensitive to cool gas with $T \sim 10^4$ K (e.g., Heckman et al. 1990; Martin 1999; Rubin et al. 2014), although in some rare nearby cases the hot phase is observable in X-ray (e.g., Lopez et al. 2020, 2023). Translating from observed quantities to an inferred mass flux is a difficult and uncertain process, however in almost all cases the inferred mass flux is orders of magnitude smaller than what standard theoretical models predict (e.g., McQuinn et al. 2019; Concas et al. 2022). Resolving this profound conflict between theory and observations requires a better treatment and understanding of multiphase outflows on both sides. Here, we have shown that the nature of feedback is likely to be dramatically different from standard single-phase galactic feedback models, which in part helps to relieve this tension. We can, however, also use this multiphase wind picture to help refine our understanding and modeling of observations.

In order to make the most of galactic wind observations, particularly new and planned spatially resolved emission observations (e.g., Reichardt Chu et al. 2022), new observa-

tional modeling paradigms that take the multiphase nature of galactic winds into account are required. In particular, our finding is that the vast majority of the readily observable cool gas is in the form of clouds with a fairly well-understood size distribution and a relatively small volume-filling fraction. Furthermore, because we have shown the properties of these cool clouds are closely coupled to the energy containing hot phase, with the two phases shaping each other's properties, future multiphase models may be able to constrain not only the mass flux (cool phase) but also the energy flux (hot phase).

Our findings also provide insight into the nature of the CGM. Observations in the CGM suggest that the cool phase is not in pressure equilibrium with the hot phase (Werk et al. 2014). One way of accounting for this discrepancy is the addition of non-thermal components. As previously discussed, our results suggest that turbulence within clouds is a significant contributor to pressure support, while magnetic fields are a minor actor, although other candidates such as cosmic rays may also provide further cold phase pressure support (Butsky et al. 2020).

5.6.4 Further Considerations

Despite modelling and simulating a more realistic turbulent magnetized multiphase system, we ultimately still make simplifying assumptions. The following are some such limitations and caveats, many of which are each interesting enough in their own right to warrant exploration in future work.

Additional Physics

The most obvious and direct of these is the inclusion of additional underlying physics which were not incorporated into this work, but could possibly affect such systems, such as significantly modifying the dynamics or thermodynamics of the outflow. Detailed

treatments and investigations of these processes are thus needed to accurately assess their impact and importance, in order that they can be properly accounted for when modelling the behavior of multiphase winds. One such source of uncertainty that could potentially have a large impact is cosmic rays, which have been found to modify the behavior of multiphase winds (Huang et al. 2022; Armillotta et al. 2022). They can provide an additional source of non-thermal pressure support and affect cooling (Butsky et al. 2020), or accelerate clouds directly (Wiener et al. 2019). Explicit viscosity (e.g., Li et al. 2020; Jennings & Li 2021) and thermal conduction (e.g., Brüggén & Scannapieco 2016; Li et al. 2020) have also not been included here (although we expect that in turbulent mixing layers, mixing is generally dominated by turbulent diffusion (Tan et al. 2021), explicit conduction can affect observables (Tan & Oh 2021)). The inclusion of non-equilibrium cooling as well as more sophisticated non-equilibrium chemical evolution models could potentially be important, especially for predictions of observables such as ion column densities (Sarkar et al. 2022). We assume solar metallicity and abundances everywhere. This is clearly an oversimplification. When metallicities of clouds and their environment differ, the mixing can have significant implications for observables (Gritton et al. 2014; Heitsch et al. 2022). Dust survival or depletion in these systems is also important to understand, in particular at lower temperatures (Farber & Gronke 2022). Closely tied to this is self-shielding. While we assume that the whole box is optically thin, leading to clouds with $T \sim 10^4$ K, more massive clouds with $N_{\text{HI}} > 10^{17} \text{ cm}^{-2}$ are likely able to self-shield and hence possess cold cores (Farber & Gronke 2022). In fact, molecular gas is observed both in the Milky Way (Di Teodoro et al. 2020; Noon et al. 2023) and other systems such as M82 (Walter et al. 2002).

Geometry

Besides being limited in terms of spatial extent, the nature of our setup lacks the correct geometry to track the outflowing wind further into the halo. While wind properties are often converged with resolution, local box simulations like this are limited by the Cartesian geometry of the setup which restrict the reliability of quantitative predictions of wind properties, which are often not converged with respect to box height instead. In particular, such a geometry which lacks the inverse squared distance scaling does not allow the adiabatic expansion and corresponding subsonic to supersonic transition of steady state winds that are a hallmark of predictions from analytic galactic wind models like Chevalier & Clegg (1985) (Martizzi et al. 2016; Fielding et al. 2017b). This in part motivated our focus on a cloud-centric analysis, rather than an extended look at associated wind properties such as mass and energy loading that are typically done in such setups.

5.6.5 Looking Forward

A natural next step is to apply and test cloud models further downstream in the wind. Analytic models of cloud growth/destruction as they are carried out in the wind such as the framework outlined in Fielding & Bryan (2022) can easily be extended and applied to populations of clouds in the manner we have done here. While we have determined a physically motivated initial distribution, validating the time evolution of such a distribution coupled to an expanding wind is the natural next step (Anthony Chow et al., in prep). Additionally, the scope of the simulations can be expanded by exploring the effects of the different mechanisms listed above (such as cosmic rays) on the outflow. This will provide us with a more comprehensive understanding of the physical processes that drive multiphase galactic winds. Finally, understanding the impact of varying the

properties of the galactic environment, such as gas surface density and metallicity, on the outflow will allow us to test the robustness of our results and determine the degree to which they are dependent on specific initial assumptions. In the long term, such models can inform subgrid approaches to modelling unresolved multiphase outflows in large scale galactic and cosmological simulations, allowing us to study the macro impacts on galaxy formation and evolution.

5.7 Conclusions

At the frontier of the field of galaxy formation and evolution lies the challenge to understand the multiphase nature of the environment and its implications. Galactic winds are a key component of these complex ecosystems and can be driven by stellar feedback channels. Observations reveal them to be common and also multiphase in structure. Theoretical modeling and simulations of galactic winds have become increasingly sophisticated, aiming to reproduce and understand these outflows. In this work, we have built on turbulent radiative mixing layer theory and applied this to understanding and modelling the formation of multiphase outflows. We have run 3D MHD tall box patch simulations with a clustered SN driven wind outflow, with a focus on analyzing and modelling the properties of the resulting seeded cloud population. The main findings are summarized as follows:

- *Seeded by Fragmentation:* During the breakout phase, the hot expanding bubble propagates through a multiphase ISM. The multiphase nature of the ISM tends to lead to asymmetric breakouts, and causes the outflows to fluctuate in power and direction. More importantly, it also leads to the fragmentation of the ISM during the breakout, which seeds the resulting hot outflowing wind with a population of cool clouds. Consistent with this formation history, we find that clouds are mostly

comprised of gas that was originally part of the cool ($T \sim 10^4$ K) ISM, rather than from cold ($T \lesssim 10^2$ K) or hot material ($T \gtrsim 10^6$ K).

- *Turbulent Clouds & Winds:* The uneven breakout and the presence of these clouds induces large scale turbulence in the wind. The clouds gradually get entrained via mixing induced accretion of momentum. The turbulent mach number in the hot phase of the wind is ~ 0.3 , and magnetic pressure is extremely weak, meaning thermal pressure is dominant. However, clouds have turbulent mach numbers ~ 1 internally and at their surfaces where mixing occurs.
- *Clouds Exhibit Complex Morphologies:* This turbulent environment naturally leads to complex cloud morphologies that do not always conform to the head-tail description. In some cases, clouds can appear to be extremely misaligned with the bulk flow.
- *Cloud Size Distribution:* Cloud sizes are well described by a power law distribution of $dN/dm \propto -2$. While this has been observed in previous works, we find that this scaling seems to hold even with the inclusion of magnetic fields in the ISM. The lower and upper scale cutoffs are consistent with estimates of the cloud survival radius and the scale height of the disc, respectively. We find that this scaling appears early on during the SN breakout stage—consistent with the source of clouds being the process of fragmentation of the ISM.
- *Cloud Survival:* The critical radius below which clouds can survive in a turbulent wind are consistent with $r_{\text{crit, shear}}$ which is given by equation (5.4). This criterion is motivated by a combination of turbulent box simulations from Gronke et al. (2022) and cloud simulations with high χ in Abruzzo et al. (2022b), who proposed that the survival criteria is set by $t_{\text{cool, mix}} < t_{\text{shear}}$.

- *Cloud Growth*: By combining analytic models for the surface area to volume relationship of clouds ($A_{\text{cl}} \propto V^{5/6}$) and mixing layer theory for the mass inflow velocity v_{in} , we can model the growth rate of clouds. We find that the model predictions are a good match to what we observe in the simulations, including predicted scalings and estimated \dot{m} .
- *Role of Wind Shear*: At high relative shear velocities between clouds and the surrounding wind, there is a strong relationship between the shear and the turbulent velocities which drive mixing and hence growth. However, as clouds get entrained, the turbulence becomes independent of the shear and is roughly the sound speed of the cool gas. The KHI is thus not the sole driver of turbulent mixing.
- *Non-Thermal Pressure Support*: Because turbulent velocities within these clouds are high (roughly the sound speed within the cloud), turbulent pressure provides as much support as thermal pressure. Surprisingly, even though the initial ISM had plasma beta ~ 1 , clouds are much more weakly magnetized, likely due to significant mixing with the hot high beta wind. These sources of non-thermal pressure support bring the clouds into *total* pressure equilibrium with their surroundings, despite having a factor of 2 lower thermal pressure.

In summary, we have shown that many of the physical insights and analytic scalings from idealized small scale simulations translate well to larger scale, more realistic turbulent magnetized winds. The multiphase component of these winds (a population of cool embedded clouds) can hence be well modelled, allowing for physics informed subgrid prescriptions which account for unresolved coupling between the various phases to be utilized in galactic and cosmological simulations where resolution limits are prohibitive. While refinements are required (e.g., the survival of molecular gas), moving forward, proper treatment of the small scales in this manner promises to pave the way towards

tackling burning questions that remain regarding the role of feedback in galaxy formation and evolution.

Acknowledgements

We thank Peng Oh, Greg Bryan, Yan-Fei Jiang, and Matthew Abuzzo for helpful discussions. We have made use of the yt astrophysics analysis software suite (Turk et al. 2011), matplotlib (Hunter 2007), numpy (Van Der Walt et al. 2011), scipy (Virtanen et al. 2020), CMasher (van der Velden 2020) and Blender (Blender Online Community 2023) whose communities we thank for continued development and support. BT acknowledges support from NASA grant 19-ATP19-0205 and NSF grant AST-1911198. BT was supported partly by the Simons Foundation through the Flatiron Institute’s Predoctoral Research Fellowship. DBF is supported by the Simons Foundation through the Flatiron Institute.

Chapter 6

Summary and Conclusions

The most important step a man can take.

It's not the first one, is it? It's the next one.

Always the next step, Dalinar.

Brandon Sanderson, Oathbringer

In broad strokes, the research presented in this dissertation has explored the physics of the multiphase environments in and around galaxies. This encompasses the galaxies themselves, the vast reservoir of material that surrounds them (known as the circumgalactic medium (CGM)), and the flows that connect the two. These flows govern galaxy growth and evolution by coupling the galaxies and their CGM. Outflows transport material away from the galaxy, while inflows set the fuel budget for future star formation. They hence comprise much of the machinery by which galaxies self-regulate.

A large focus of our work has been studying how the multiphase nature of these systems, which observations have steadily uncovered, shape them. Despite being ubiquitous, many uncertainties remain due to their surprisingly rich complexity. Combining analytic theory with numerical simulations, we delved into their inner workings so as to




Figure 6.1: Snapshots from our simulations showing the hierarchy of scales in this summary.

understand and model them. We adopted a systematic approach, starting at the smallest scales in the problem with a deep dive into the interfaces between phases and how they determine bulk evolution. We then explored the connection between these mixing layers and observables. Applying these results to larger scales, we looked at cold clouds moving through hot backgrounds, both infalling under gravity and in turbulent outflowing winds. Figure 6.1 depicts this laddering of scales in our simulations.

When first reviewing the literature on multiphase systems, we identified two important areas that had not yet been addressed. Firstly, there was a lack of consensus on the scalings of hot gas entrainment, stemming from a lack of understanding of the underly-

ing physics at play. Secondly, cooling in the simulations was taking place in unresolved interfaces. These interfaces were often only resolved by a single cell! Yet surprisingly, simulations appeared to be remarkably robust to resolution.

We hence ran high resolution 1D and 3D Athena++ simulations zooming in on these interfaces, known as radiative turbulent mixing layers (RTMLs). These layers arise wherever multiphase gas, turbulence, and radiative cooling are present. In steady state, thermal advection from the hot phase balances radiative cooling. We found that over-cooling only happens if numerical diffusion dominates thermal transport, which in turn is set by turbulent mixing; convergence is still possible even when the Field length (the lengthscale on which conduction and cooling balance) is unresolved, explaining why simulations were not sensitive to resolution. We showed that by exploiting parallels between mixing layers and turbulent combustion, which has well-developed theory and abundant experimental data, a deeper physical understanding of these fronts can be obtained. The key parameter in these mixing layers is the the ratio of the outer eddy turnover time to the cooling time. When the cooling time is shorter, the front fragments into a multiphase medium. Just as for scalar mixing, the eddy turnover time sets the mixing rate, independent of small scale diffusion. This is akin to stirring milk into your coffee, where the size and speed of the stirring sets the mixing rate, rather than the molecular diffusion rate. For this reason, thermal conduction often has limited impact on the net cooling in these interfaces. We also showed that the entrainment rate and the effective emissivity can be understood in detail by adapting combustion theory scalings. Mean density and temperature profiles can also be reproduced remarkably well by mixing length theory.

These results have implications for the structure and survival of cold gas in many settings and the resolution requirements for large scale galaxy simulations. This is a concern as cosmological simulation are currently unconverged in cold gas properties. Ultimately, the physics of these mixing layers set the mass and momentum transfer

between the hot and the cold phase and thus impact not only the morphology of the multiphase medium but also determine the fuel supply for future star-formation. Our results indicate it is sufficient in the presence of turbulence to resolve just the driving scales, hence relaxing resolution requirements significantly.

Although we showed that global mass, momentum and energy transfer between phases mediated by RTMLs does not require extremely high resolutions or proper treatment of thermal conduction for convergence, one issue remained with low resolution. To return to the coffee analogy, even though molecular diffusion is not important in setting the overall mixing rate, it is still the actual physical process responsible for mixing itself. This slow molecular diffusion rate is compensated for by the large increase in interface surface area during the stirring process. Similarly, for RTMLs, while hot gas entrainment rate may not be sensitive to resolution, the actual interface structure is. This has important bearings when computing observables such as line absorption and emission from these simulations. Since RTMLs are ubiquitous in multiphase gas, they can potentially explain observations of ions such as OVI, which have significant observed column densities despite short cooling times. It is hence important to get this right in order to be able to test our models with observational data.

We have shown that properties such as temperature distributions, column densities and line ratios in simulations *are* sensitive to resolution and thermal conduction. This is because they probe the interface structure of the mixing layer. We hence develop a prescription for applying a simple 1D conductive-cooling front model which quantitatively reproduces 3D hydrodynamic simulation results for column densities and line ratios, even when the RTML has a complex fractal structure. This enables subgrid absorption and emission line predictions in large scale simulations. The predicted line ratios are in good agreement with observations, while observed column densities require numerous mixing layers to be pierced along a line of sight.

After laying the groundwork for understanding and modelling RTMLs, we looked towards applying them to larger scale systems. While there has been a plethora of work on ‘wind tunnel’ simulations that study cold gas clouds in winds, the infall of this gas under gravity is at least equally important, and fundamentally different since cold gas can never entrain. Instead, velocity shear increases and remains unrelenting. We found that there was a lack of work on the infalling case, despite the survival, growth and dynamics of cold gas feeding the galaxy being fundamental to galaxy formation.

Hence, we looked at cold clouds falling through a hot background under gravity, both in constant and stratified environments. We found that if these clouds are growing, they can experience a drag force due to the accretion of low momentum gas, which dominates over ram pressure drag. This leads to sub-virial terminal velocities, in line with observations. We developed simple analytic theory and predictions that built upon our earlier work. Importantly, we found that the survival criterion for infalling gas is more stringent than in a wind, requiring that clouds actually grow faster than they are destroyed (whereas in winds only the cooling time of the mixed gas needs to be shorter than the destruction time). Additionally, in stratified environments, larger clouds need only survive infall long enough until cooling becomes effective. We applied these conclusions to realistic environments such as high velocity clouds in the MW and filaments in clusters.

Lastly, wind tunnel simulations have been used to derive models for cloud survival and growth, which in turn have been used to construct subgrid prescriptions for clouds in galactic winds. However, current simulation work are highly simplified and assume that (i) the cloud is a simple spherical cow, and that (ii) the hot gas wind properties are laminar and time-constant. Hence, we looked at how well these models hold up in *realistic* turbulent winds that are driven from the galaxy by clustered supernovae and that evolve dynamically. Fragmentation of the interstellar medium during superbubble

breakout seeds the resulting hot outflow with a population of cool clouds. We analyzed and modelled the origin and properties of these clouds. We found that cloud sizes are well described by a power law distribution and mass growth rates can be well modelled using turbulent radiative mixing layer theory. Turbulence provides significant pressure support in the clouds, while magnetic fields only play a minor role. We concluded that many of the physical insights and analytic scalings derived from idealized small scale simulations translate well to larger scale, more realistic turbulent magnetized winds, thus paving a path towards their necessary yet challenging inclusion in global-scale galaxy models.

Recent work has shown that capturing the appropriate scales and boundary conditions is *critical* in studying multiphase systems. While much of this work has been in the area of galactic winds, outflows are but one step in the cycling of baryons, and similar studies for the CGM and infalling clouds, and in general more realistic environments, are still lacking. As evidenced by our work, multiphase systems are notoriously difficult to study because they introduce new, much smaller scales that can significantly impact large scale behavior. For instance, the interactions at the interfaces between phases can alter bulk flow and phase properties. If we do not strive to understand the small scales, we will never get the large scales right. In short, the multiphysics, multiscale and multiphase nature of the environments around cold clouds render their exploration incredibly challenging. This remains a rich area for future work.

Appendix A

Code Verification Test: 1D Diffusion Couple

We test the conduction module by considering the interface between two regions of different temperatures that is initially represented by a step function located at $x = x_0$ when $t = 0$. The left side is at temperature T_1 , and the right side is at T_2 . The analytical solution is then given by:

$$T(x, t) = T_0 + \frac{\Delta T}{2} \operatorname{erf}\left(\frac{x - x_0}{2\sqrt{\alpha t}}\right), \quad (\text{A.1})$$

where T_0 is the mean of T_1 and T_2 , and $\Delta T = T_2 - T_1$. This solution assumes that density is fixed, so we turn off the hydrodynamics updates to the density field and velocity fields, and only let the energy of the simulation cells evolve.

We choose the set of parameters listed in Table A.1, and ensure that the chosen value

Resolution	T_1	T_2	x_0	γ	ρ	κ	V_m
512	100	500	40	5/3	0.75	1.5	10

Table A.1: Parameters used for the thermal conduction test.




Figure A.1: Results of the implementation test for thermal conduction. The data points from simulation at various times are shown to match the analytic solution, given by the solid lines.

of V_m is sufficiently high for a well converged solution. The results are shown in Fig. A.1, which show that the simulation data is a good match to the analytical solution. The code is also verified for a case where the density is not held constant in the resolution tests for 1D thermal fronts described in §2.3, where the integrated cooling over a steady thermal front is shown to converge to the expected analytical result.

Appendix B

Hydrostatic Equilibrium Test

Figure B.1 is a simple test of the setup in hydrostatic equilibrium for two boxes at different low resolutions (below fiducial). The box is at $T = 10^4$ K throughout and a mid-plane density of 10 cm^{-3} . There is a density floor at 10^{-6} cm^{-3} . In both cases, the error over time is smaller, and smaller for the higher resolution box. These tests are simple sanity checks of the initial conditions. In practice, outflowing gas quickly becomes the dominant effect.

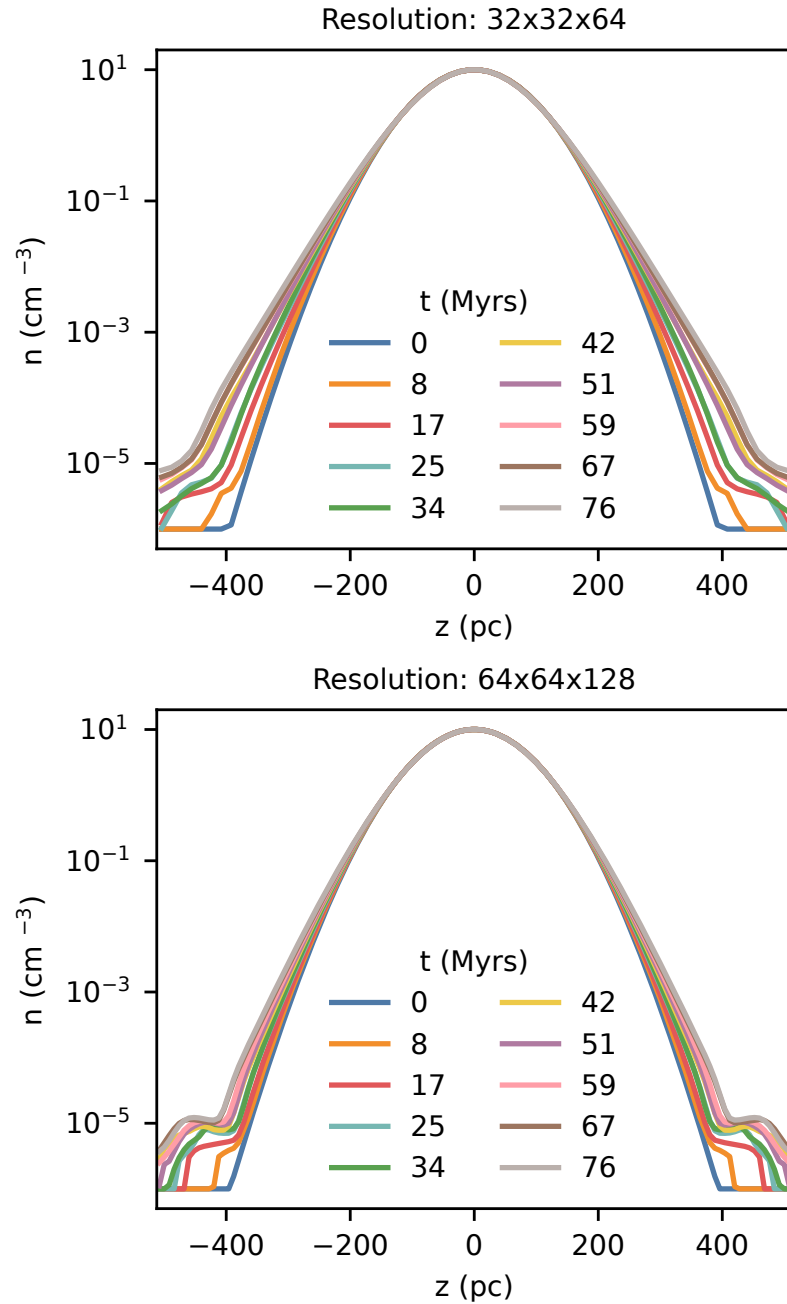


Figure B.1: A test of the initial hydrostatic profile for two low resolution boxes. We initialize 3D boxes with the analytical profiles at $T = 10^4 \text{ K}$ and see how well hydrostatic equilibrium is maintained. The error gets smaller at higher resolutions and is negligible. There is a density floor at 10^{-6} cm^{-3} .

Appendix C

Constrained Turbulence Test

For our implementation of turbulence, we add a turbulent scale height constrain z_{turb} . This is done by applying a Gaussian weight $\exp\{(-z/z_{\text{turb}})^2\}$ to the velocity perturbations in the turbulent driver. We test our implementation of z_{turb} by injecting decaying turbulence with some scale height into a uniform periodic box, as shown in Fig. C.1. The orange dashed line marks the expected kinetic energy as a function of height in the box at the time of injection centered at zero (in practise due to zeroing the total momentum, this has some shift), while the blue and red lines show early and late time values from the simulation. At late times, there is some decay in kinetic energy since we do not continuously drive, and the turbulence is uniform through the box as expected.




Figure C.1: A test of constrained turbulence. The turbulence is injected at the start of the simulation into a uniform box with periodic boundary condition. The orange dashed line marks the expected kinetic energy as a function of height in the box at the time of injection, while the blue and red lines show early and late time values from the simulation.

Appendix D

Extended Townsend Algorithm

Exact Integration Scheme

We first review the original Townsend algorithm (Townsend 2009). This is an algorithm for computing the change in temperature due to isochoric radiative cooling over some given time interval. It is based on an exact solution by noting that if one knows the cooling function, one can in principle simply integrate this function over the given time interval. The algorithm is as follows. We want to solve the following cooling equation:

$$\frac{dT}{dt} = -\frac{(\gamma - 1)\rho\mu}{k_B\mu_e\mu_H}\Lambda(T) \quad (\text{D.1})$$

We can define the dimensionless Temporal Evolution Function (TEF) as follows:

$$Y(T) \equiv \frac{\Lambda(T_{\text{ref}})}{T_{\text{ref}}} \int_T^{T_{\text{ref}}} \frac{dT}{\Lambda(T)} \quad (\text{D.2})$$

In principle, this only requires that $1/\Lambda(T)$ be analytically integrable. This is basically a normalized measure of the time taken to cool/heat from T_{ref} to T . We can then integrate

the cooling equation over a timestep such that the integrated cooling function becomes:

$$\frac{T_{\text{ref}}}{T^n} \frac{\Lambda(T_n)}{\Lambda(T_{\text{ref}})} [Y(T^n) - Y(T^{n+1})] = -\frac{\Delta t}{t_{\text{cool}}} \quad (\text{D.3})$$

where t_{cool} is the single point cooling time defined as

$$t_{\text{cool}} = \frac{k_{\text{B}}\mu_e\mu_{\text{H}}T}{(\gamma - 1)\rho\mu\Lambda(T)} \quad (\text{D.4})$$

and hence we can update the temperature as

$$T^{n+1} = Y^{-1} \left[Y(T^n) + \frac{T_n}{T_{\text{ref}}} \frac{\Lambda(T_{\text{ref}})}{\Lambda(T^n)} \frac{\Delta t}{t_{\text{cool}}} \right] \quad (\text{D.5})$$

In practice, this is done in 3 steps - computing the TEF, evolving it over the required timestep, and then transforming back to the new updated temperature. In the next section, we compute the TEF for the two most useful cases, piecewise power laws and piecewise linear functions.

Temporal Evolution Functions

Piecewise Power Laws

Physically realistic cooling function often come in the form of piecewise power laws that have been fitted to more complicated underlying models. The construction of the TEF and its inverse for piecewise power laws is given in the appendix of Townsend (2009), which are as follows. We assume the following functional form:

$$\Lambda(T) = \Lambda_k \left(\frac{T}{T_k} \right)^{\alpha_k} \quad (\text{D.6})$$

for some temperature bin $T_k \leq T \leq T_{k+1}$. The TEF is then

$$Y(T) = Y_k + \begin{cases} \frac{1}{1-\alpha_k} \frac{\Lambda_{\text{ref}}}{\Lambda_k} \frac{T_k}{T_{\text{ref}}} \left[1 - \frac{T_k}{T} \alpha_k^{-1} \right] & \alpha_k \neq 1 \\ \frac{\Lambda_{\text{ref}}}{\Lambda_k} \frac{T_k}{T_{\text{ref}}} \ln \left(\frac{T_k}{T} \right) & \alpha_k = 1 \end{cases} \quad (\text{D.7})$$

The constraint that $Y(T)$ is continuous leads to the recurrence relation

$$Y_k = Y_{k+1} - \begin{cases} \frac{1}{1-\alpha_k} \frac{\Lambda_{\text{ref}}}{\Lambda_k} \frac{T_k}{T_{\text{ref}}} \left[1 - \frac{T_k}{T_{k+1}} \alpha_k^{-1} \right] & \alpha_k \neq 1 \\ \frac{\Lambda_{\text{ref}}}{\Lambda_k} \frac{T_k}{T_{\text{ref}}} \ln \left(\frac{T_k}{T_{k+1}} \right) & \alpha_k = 1 \end{cases} \quad (\text{D.8})$$

with $Y_{\text{ref}} = Y(T_{\text{ref}}) = 0$. For cooling(heating), T_{ref} is any temperature higher(lower) than the current temperature and we construct the TEF for decreasing(increasing) temperatures. Hence for heating, we express the above recurrence relation as $Y_{k+1} = Y_k + \dots$, starting from $T_{\text{ref}} < T$. The inverse TEF is thus

$$Y^{-1}(Y) = \begin{cases} T_k \left[1 - (1 - \alpha_k) \frac{\Lambda_k}{\Lambda_{\text{ref}}} \frac{T_{\text{ref}}}{T_k} (Y - Y_k) \right]^{1/(1-\alpha_k)} & \alpha_k \neq 1 \\ T_k \exp \left[-\frac{\Lambda_k}{\Lambda_{\text{ref}}} \frac{T_{\text{ref}}}{T_k} (Y - Y_k) \right] & \alpha_k = 1 \end{cases} \quad (\text{D.9})$$

Piecewise Linear Function

We can likewise compute the TEF and its inverse for piecewise linear functions of the form

$$\Lambda(T) = \Lambda_k + \alpha(T - T_k), \quad \alpha = \frac{\Lambda_{k+1} - \Lambda_k}{T_{k+1} - T_k} \quad (\text{D.10})$$

where α is the slope in the temperature bin $T_k \leq T \leq T_{k+1}$. The TEF is then:

$$Y(T) = \begin{cases} Y_k + \frac{\Lambda_{\text{ref}}}{T_{\text{ref}}} \frac{1}{\alpha} \log \left(\frac{\Lambda_k}{\Lambda_k + \alpha(T - T_k)} \right) & \alpha \neq 0 \\ Y_k + \frac{\Lambda_{\text{ref}}}{T_{\text{ref}}} \frac{T_k - T}{\Lambda_k} & \alpha = 0 \end{cases} \quad (\text{D.11})$$

and its inverse is:

$$Y^{-1}(Y) = \begin{cases} T_k + \frac{\Lambda_k}{\alpha} \left(\left(\exp \left[\frac{\alpha T_{\text{ref}}}{\Lambda_{\text{ref}}} (Y - Y_k) \right] \right)^{-1} - 1 \right) & \alpha \neq 0 \\ T_k + \frac{T_{\text{ref}}}{\Lambda_{\text{ref}}} \Lambda_k Y - Y_k & \alpha = 0 \end{cases} \quad (\text{D.12})$$

The Extended Algorithm

The main benefit of the Townsend Algorithm is naturally that it is based on an exact solution and hence not sensitive to errors associated with temporal resolution. This makes it the best choice when using a simple treatment of cooling/heating that is only a function of temperature and does not have a functional dependence on quantities such as ionic abundances. One caveat is that thermal and hydrodynamical evolution are decoupled over the course of a single timestep. It is still recommended to include a further constraint on the timestep past the CFL condition that addresses this issue (for example requiring that the timestep be below some fraction of the cooling time).

In Townsend (2009), the algorithm was only outlined for cooling. Including heating is non-trivial because heating scales differently with density compared to cooling. This means that the net cooling/heating function becomes density dependent, which means that we can no longer optimize by pre-computing all TEFs. Another challenge is the behavior of power laws near equilibrium points. As far as we know, the only attempt to include heating and tackle this problem is in Zhu et al. (2017). In their implementation, they assume piecewise linear functions and only one equilibrium point. We generalize

this further to generally work with piecewise power laws, as most cooling functions are represented as such, and an arbitrary number of equilibrium points.

When calculating the net cooling/heating at a given density, we assume that the resulting table itself represents a piecewise power law, instead of being a linear combination of two piecewise power laws. The exception is that for bins that have a zero crossing, we assume they are piecewise linear functions. One advantage of the piecewise functions above is that we can use linear functions only for bins with equilibrium temperatures to interpolate smoothly across the bin, where a power law formulation breaks down. This is fine as long as we dot our i's and cross our t's when computing the TEF and its inverse. Note that by construction, the TEF for any point with net zero cooling/heating will have an infinite TEF. Hence in practise one should set the cooling/heating to some non-zero tiny number with the appropriate sign. Furthermore, for cooling/heating, we use the next higher/lower bin as the reference temperature and only calculate the TEF down/up to the next equilibrium temperature, as we must always remain within a cooling/heating region if we begin there.

In detail, we implement a new cooling module as a class that is initialized at the beginning of the simulation with some specified cooling and heating table. It is assumed that both cooling and heating are represented as piecewise power laws. The class gives the user the ability to call a cooling function for some input temperature, density and timestep and returns the new temperature in cgs units. The class also allows the user to query the single point cooling time at some given temperature and density, along with the minimum/maximum temperatures. For the cooling implementation, we use two preallocated scratch arrays, one to hold temperatures and one to hold the net cooling. At the start of the function, we first check that the temperatures are within the bounds of the provided table. If it is not, we return the lower/upper bound instead. We then populate the scratch arrays with a copy of the temperature table and the net cooling

table for the given density. Next, we figure out which temperature bin we are in and check if there is net cooling or heating in that bin. If the net value is zero, we return the current temperature. We now case on whether there is net cooling or heating. For either case, the steps are similar, but the directions in which we compute quantities are opposite, along with the choice of some indexes. We first use the next bin over as the reference bin, but check for the edge case where we are in an equilibrium bin. In that case we modify the next bin to be at the equilibrium temperature and set the net cooling in that bin to be zero. If we assume in a bin that cooling and heating have the following power law forms:

$$\Lambda = \Lambda \left(\frac{T}{T_i} \right)^{\alpha_i} \qquad \Gamma = \Gamma_i \left(\frac{T}{T_i} \right)^{\beta_i} \qquad (\text{D.13})$$

then the equilibrium temperature in a bin that transitions from heating to cooling or vice versa is then

$$T_{\text{eq}} = T_i \left(\frac{\Lambda_i}{n\Gamma_i} \right)^{1/(\alpha_i - \beta_i)} \qquad (\text{D.14})$$

We also flag the bin to be linear. We then calculate the TEF recursively downwards, until either we reach the bounds or we reach another equilibrium temperature, in which case we again modify the bin and flag it as a linear bin. We then compute the current TEF and then the TEF after the timestep. Since the bins with equilibriums are linear, there is some finite value of Y that if exceeded we immediately return the equilibrium temperature there. We then compute the inverse TEF, accounting for the fact that we might have crossed several bins. This gives us the new temperature.




Figure D.1: The cooling and heating functions used in our test case, with $n = 2$. The grey points show the values in each bin.

Testing

Figures D.1 and D.2 show a test of the above algorithm, where we have set up a simple cooling and heating curve with $n = 2$ and computed the calculate new temperature as a function of timestep for a range of starting temperatures. The black dashed lines in Fig. D.2 mark the expected equilibrium temperatures.




Figure D.2: The computed heating/cooling over time for a range of starting temperatures. Black dashed lines show the expected equilibrium temperatures.

Appendix E

Estimating Cloud Growth Rates

Given the large time interval between full data outputs, and the resulting difficulty in matching and tracking individual clouds across multiple snapshots, we would like to be able to estimate the mass growth rate of an individual cloud based solely on its properties from a single time snapshot alone.

Here we compare several methods of doing this, which we test by applying to a controlled setup containing only a single cloud and where the mass growth rate of the cloud is tracked explicitly with high time resolution. This provides a ‘ground truth’ which we can use to test our various estimators against. The setup consists of a cool cloud, initially at rest, falling under gravity in a constant hot background as detailed in Section 3 of Tan et al. (2023) (more specifically this setup corresponds to the $\Lambda_0 = 30$ and $r_{\text{cl}} = 300$ pc run). At fiducial resolution here, r_{cl} is resolved by ~ 12 cells.

We estimate the mass growth rate in three different ways. The first method is by estimating the total cooling luminosity $Q \equiv \int \rho \mathcal{L} dV$. If we then assume that radiative cooling balances enthalpy flux (i.e., in the sub to transonic regimes; Ji et al. 2019), we can estimate the mass growth rate as $\dot{m}_Q \sim Q/c_p T_{\text{hot}}$. This also assumes that the bulk of the cooling luminosity comes from the mixed gas in the turbulent mixing layer and




Figure E.1: The mass growth rate as a function of mass from the simulations (grey points; smoothed result in blue) and from estimates using (i) total cooling luminosity Q (both recorded and estimated) and (ii) mass flux through the bounding box of the cloud.




Figure E.2: The model applied to different choices of temperature for the isosurface T_{iso} . The scaling factor f_{iso} is seen to be dependent on T_{iso} .




Figure E.3: The mass loss rate for a cloud being destroyed, compared with the various methods to estimate mass growth rate. We find that these methods do poorly when the cloud is not actually growing.




Figure E.4: The ratio of \dot{m}_Q and \dot{m}_{surface} as a function of cloud size. The value of 0.05 given by the orange dashed line is the final scaling value we use in our main analysis.

that the contribution from the cool and hot gas components is negligible. In Figure E.1, we show how this stacks up against the actual mass growth rate. The mass growth from the simulation is represented by the solid blue line, which is obtained by smoothing the instantaneous values of \dot{m} represented by the grey points. In this simulation, the cooling in each cell per timestep is explicitly tracked, which using the method above corresponds to a mass growth rate given by the red dashed line. However, this quantity is not similarly tracked in our main simulations, and hence must be estimated instead by computing Q_{est} using the density and temperature of each cell, which gives us an estimate shown by the orange dashed line. This does even better than Q_{sim} because we explicitly avoid contributions from cells within a percent of the temperature floor, to account for the lack of explicit heating in this simulation.

The second method involves computing the mass flux through the bounding box of the cool gas. By computing the net mass influx into the volume defined by the bounding box, and assuming that this balances the mass flux from the hot to cool phase, we can estimate the mass growth rate. This is shown by the teal dashed line in Figure E.1. This estimate does surprisingly well despite its simplicity. While it works well in this controlled experiment, it fails in our main simulations due to the turbulent nature of the environment leading to fluctuating mass fluxes through the bounding box.

The third and last method is to compute the total mass flux through a temperature isosurface \dot{m}_{surface} . This is the isosurface we construct using the marching cubes algorithm as described in the previous section. In order to measure the mass flux through this surface, we interpolate the velocity and log density at the centroid of each triangular face in the isosurface mesh. The mass flux through each face is given by the component of this velocity that is normal to the face and in the inward mesh direction, multiplied by the density and the area of the face (while this also includes the velocity of the surface itself, this contribution should average to zero in the cloud frame). This gives us a total

mass influx rate over the entire isosurface. However, just this quantity alone significantly overestimates the mass growth rate. This is likely because translating this instantaneous quantity into a mass growth rate directly assumes that all the flux becomes cool gas, which does not hold here because our isosurface is itself sensitive to the velocity field. To account for this, we determine a normalization constant f_{iso} by calibrating to Q_{est} . The constant of proportionality $f_{\text{iso}} = 0.2$ for $T_{\text{iso}} = 5.0$, which is our default choice. Figure E.2 shows this model for isosurfaces of various temperatures, with $\log(T) = 4.2, 5.0$ and 5.5 . We have used $f_{\text{mix}} = 0.17, 0.2$ and 0.3 for them respectively. We expect f_{iso} to vary with the choice of isosurface temperature and the velocity field.

What if the cloud is not growing? Figure E.3 shows a run with weaker cooling ($\Lambda_0 = 1$) where the cloud is getting destroyed. Here, we see a drawback of the various methods presented—they do not account for mass loss. Hence, these methods should not generally extend to clouds that do not survive. The bounding box approach appears to work well at later times, but this is because the accreting cloud is able to form a steady inflow towards the cloud surface in this setup. In more turbulent environments, the signal from this method is washed out by the much higher turbulent velocities in the background.

Ultimately, the three different methods of estimating \dot{m} from instantaneous cloud properties presented above give similar results that are good estimates of the actual cloud mass growth rate, assuming that the clouds are growing. It should be stressed that these are estimates—having multiple methods provides a cross check that allows us to make this estimate of \dot{m} more reliably.

When measuring mass growth rates of the clouds in our main simulations, we thus use the mass flux through a temperature isosurface \dot{m}_{surface} , but normalized to match mass growth rates estimated using the cooling luminosity \dot{m}_{Q} . This choice allows for more consistency when discussing the scalings related to both the inflow velocity through this

isosurface and its computed surface area.

The normalization is done by finding the constant of proportionality f_{iso} discussed above. In other words, we calibrate a scaling factor f_{iso} such that $\dot{m} \sim f_{\text{iso}} \dot{m}_{\text{surface}}$, where $f_{\text{iso}} \equiv \langle \dot{m}_Q / \dot{m}_{\text{surface}} \rangle$. Figure E.4 shows $\dot{m}_Q / \dot{m}_{\text{surface}}$ as a function of cloud size for the clouds used in the analysis in Section 5.5, with $f_{\text{iso}} = 0.05$ represented by the orange dashed line being the scaling factor we adopt.

Bibliography

- Abruzzo M. W., Fielding D. B., Bryan G. L., 2022a, arXiv e-prints, p. arXiv:2210.15679
- Abruzzo M. W., Bryan G. L., Fielding D. B., 2022b, ApJ, 925, 199
- Afruni A., Fraternali F., Pezzulli G., 2019, A&A, 625, A11
- Aguirre R. C., Catrakis H. J., 2005, Journal of Fluid Mechanics, 540, 39
- Anglés-Alcázar D., Faucher-Giguère C.-A., Kereš D., Hopkins P. F., Quataert E., Murray N., 2017, MNRAS, 470, 4698
- Antolin P., Froment C., 2022, Frontiers in Astronomy and Space Sciences, 9, 820116
- Armillotta L., Fraternali F., Marinacci F., 2016, Monthly Notices of the Royal Astronomical Society, 462, 4157
- Armillotta L., Fraternali F., Werk J. K., Prochaska J. X., Marinacci F., 2017, MNRAS, 470, 114
- Armillotta L., Ostriker E. C., Jiang Y.-F., 2022, ApJ, 929, 170
- Balachandar S., Eaton J. K., 2010, Annual Review of Fluid Mechanics, 42, 111
- Banda-Barragán W. E., Federrath C., Crocker R. M., Bicknell G. V., 2018, MNRAS, 473, 3454
- Barenblatt G. I., Monin A. S., 1983, Proceedings of the National Academy of Science, 80, 3540
- Bedat B., Cheng R., 1995, Combustion and Flame, 100, 485
- Begelman M. C., Fabian A. C., 1990, MNRAS, 244, 26P
- Begelman M. C., McKee C. F., 1990, ApJ, 358, 375
- Benjamin R. A., Danly L., 1997, ApJ, 481, 764
- Bish H. V., Werk J. K., Peek J., Zheng Y., Putman M., 2021, The Astrophysical Journal, 912, 8
- Blender Online Community 2023, Blender - a 3D modelling and rendering package. Blender Foundation, Blender Institute, Amsterdam, <http://www.blender.org>
- Bolatto A. D., et al., 2021, ApJ, 923, 83
- Borkowski K. J., Balbus S. A., Fristrom C. C., 1990, ApJ, 355, 501
- Bray K. N. C., 1990, Proceedings of the Royal Society of London Series A, 431, 315

Bray K. N. C., Cant R. S., 1991, *Proceedings of the Royal Society of London Series A*, 434, 217

Bregman J. N., 1980, *ApJ*, 236, 577

Brüggen M., Scannapieco E., 2016, *ApJ*, 822, 31

Bustard C., Gronke M., 2022, *ApJ*, 933, 120

Bustard C., Zweibel E. G., D’Onghia E., 2016, *ApJ*, 819, 29

Butsky I. S., Fielding D. B., Hayward C. C., Hummels C. B., Quinn T. R., Werk J. K., 2020, *ApJ*, 903, 77

Carr C., Bryan G. L., Fielding D. B., Pandya V., Somerville R. S., 2022, *arXiv e-prints*, p. arXiv:2211.05115

Catrakis H. J., Aguirre R. C., Ruiz-Plancarte J., 2002, *Journal of Fluid Mechanics*, 462, 245

Cattaneo M. D., Crump R. K., Farrell M. H., Feng Y., 2019, *arXiv e-prints*, p. arXiv:1902.09608

Chevalier R. A., Clegg A. W., 1985, *Nature*, 317, 44

Chomiuk L., Povich M. S., 2011, *AJ*, 142, 197

Choudhury P. P., Sharma P., 2016, *MNRAS*, 457, 2554

Churazov E., Forman W., Jones C., Böhringer H., 2003, *ApJ*, 590, 225

Cintosun E., Smallwood G. J., Gülder Ö. L., 2007, *AIAA journal*, 45, 2785

Clavin P., Williams F. A., 1979, *Journal of Fluid Mechanics*, 90, 589

Colella P., Majda A., Roytburd V., 1986, *SIAM Journal on Scientific and Statistical Computing*, 7, 1059

Concas A., et al., 2022, *MNRAS*, 513, 2535

Cooper J. L., Bicknell G. V., Sutherland R. S., Bland-Hawthorn J., 2009, *ApJ*, 703, 330

Cornuault N., Lehnert M. D., Boulanger F., Guillard P., 2018, *A&A*, 610, A75

Cowie L. L., McKee C. F., 1977, *ApJ*, 211, 135

Creasey P., Theuns T., Bower R. G., 2013, *MNRAS*, 429, 1922

Damköhler G., 1940, *Zeitschrift für Elektrochemie und angewandte physikalische Chemie*, 46, 601

Dekel A., Birnboim Y., 2006, *MNRAS*, 368, 2

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

Dennis T. J., Chandran B. D. G., 2005, *ApJ*, 622, 205

Di Teodoro E. M., McClure-Griffiths N. M., Lockman F. J., Armillotta L., 2020, *Nature*, 584, 364

Erb D. K., 2008, *ApJ*, 674, 151

Esquivel A., Benjamin R. A., Lazarian A., Cho J., Leitner S. N., 2006, *ApJ*, 648, 1043

Fabian A. C., Johnstone R. M., Sanders J. S., Conselice C. J., Crawford C. S., Gallagher J. S. I., Zweibel E., 2008, *Nature*, 454, 968

Farber R. J., Gronke M., 2022, *MNRAS*, 510, 551

Faucher-Giguère C.-A., Feldmann R., Quataert E., Kereš D., Hopkins P. F., Murray N., 2016, *MNRAS*, 461, L32

Fielding D. B., Bryan G. L., 2022, *ApJ*, 924, 82

Fielding D., Quataert E., McCourt M., Thompson T. A., 2017a, *MNRAS*, 466, 3810

Fielding D., Quataert E., Martizzi D., Faucher-Giguère C.-A., 2017b, *MNRAS*, 470, L39

Fielding D., Quataert E., Martizzi D., 2018, *MNRAS*, 481, 3325

Fielding D. B., Ostriker E. C., Bryan G. L., Jermyn A. S., 2020, *ApJ*, 894, L24

Fielding D. B., Ripperda B., Philippov A. A., 2022, arXiv e-prints, p. arXiv:2211.06434

Foster A. R., Heuer K., 2020, *Atoms*, 8, 49

Fox A. J., et al., 2016, *ApJ*, 816, L11

Fox A. J., Richter P., Ashley T., Heckman T. M., Lehner N., Werk J. K., Bordoloi R., Peeples M. S., 2019, *ApJ*, 884, 53

Fraternali F., 2017, in Fox A., Davé R., eds, *Astrophysics and Space Science Library Vol. 430, Gas Accretion onto Galaxies*. p. 323 (arXiv:1612.00477), doi:10.1007/978-3-319-52512-9_14

Fraternali F., Binney J. J., 2006, *MNRAS*, 366, 449

Fraternali F., Binney J. J., 2008, *MNRAS*, 386, 935

Fraternali F., Marasco A., Armillotta L., Marinacci F., 2015, *MNRAS*, 447, L70

Gai G., Hadjadj A., Kudriakov S., Thomine O., 2020, *Theoretical and Applied Mechanics Letters*, 10, 241

Galyardt J., Shelton R. L., 2016, *ApJ*, 816, L18

Gaspari M., Ruszkowski M., Oh S. P., 2013, *MNRAS*, 432, 3401

Gatto A., et al., 2017, *MNRAS*, 466, 1903

Gazol A., Vázquez-Semadeni E., Kim J., 2005, *ApJ*, 630, 911

Girichidis P., Naab T., Walch S., Berlok T., 2021, *MNRAS*, 505, 1083

Gnat O., Sternberg A., 2007, *ApJS*, 168, 213

Gnat O., Sternberg A., McKee C. F., 2010, *ApJ*, 718, 1315

Goerdt T., Ceverino D., 2015, *MNRAS*, 450, 3359

Gouldin F., Schefer R., Johnson S., Kollmann W., 1986, *Progress in energy and combustion science*, 12, 257

Gregori G., Miniati F., Ryu D., Jones T. W., 1999, *ApJ*, 527, L113

Gritton J. A., Shelton R. L., Kwak K., 2014, *ApJ*, 795, 99

Gritton J. A., Shelton R. L., Galyardt J. E., 2017, *ApJ*, 842, 102

Gronke M., Oh S. P., 2018, *MNRAS*, 480, L111

Gronke M., Oh S. P., 2020a, *MNRAS*, 492, 1970

Gronke M., Oh S. P., 2020b, *MNRAS*, 494, L27

Gronke M., Oh S. P., 2022, arXiv e-prints, p. arXiv:2209.00732

Gronke M., Oh S. P., Ji S., Norman C., 2022, *MNRAS*, 511, 859

Grønnow A., Tepper-García T., Bland-Hawthorn J., McClure-Griffiths N. M., 2017, *ApJ*, 845, 69

Grønnow A., Tepper-García T., Bland-Hawthorn J., 2018, *ApJ*, 865, 64

Grønnow A., Tepper-García T., Bland-Hawthorn J., Fraternali F., 2022, *MNRAS*, 509, 5756

Grudić M. Y., Hopkins P. F., Faucher-Giguère C.-A., Quataert E., Murray N., Kereš D., 2018, *MNRAS*, 475, 3511

Gülde Ö. L., 1991, in *Symposium (International) on Combustion*. pp 743–750

Guszejnov D., Hopkins P. F., Grudić M. Y., 2018, *MNRAS*, 477, 5139

Hardee P. E., Stone J. M., 1997, *ApJ*, 483, 121

Heckman T. M., Thompson T. A., 2017, in Alsabti A. W., Murdin P., eds, , *Handbook of Supernovae*. p. 2431, doi:10.1007/978-3-319-21846-5_23

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

Heitsch F., Putman M. E., 2009, *ApJ*, 698, 1485

Heitsch F., Marchal A., Miville-Deschênes M. A., Shull J. M., Fox A. J., 2022, *MNRAS*, 509, 4515

Henley D. B., Kwak K., Shelton R. L., 2012, *ApJ*, 753, 58

Henley D. B., Gritton J. A., Shelton R. L., 2017, *ApJ*, 837, 82

Hennawi J. F., Prochaska J. X., Cantalupo S., Arrighi-Battaia F., 2015, *Science*, 348, 779

Hentschel H. G. E., Procaccia I., 1984, *Phys. Rev. A*, 29, 1461

Hernquist L., 1990, *ApJ*, 356, 359

Hidalgo-Pineda F., Farber R. J., Gronke M., 2023, arXiv e-prints, p. arXiv:2304.09897

Hill A. S., Joung M. R., Mac Low M.-M., Benjamin R. A., Haffner L. M., Klingenberg C., Waagan K., 2012, *ApJ*, 750, 104

Hillier A., Arregui I., 2019, *ApJ*, 885, 101

Hopkins P. F., 2013, *MNRAS*, 430, 1653

Hopkins A. M., McClure-Griffiths N. M., Gaensler B. M., 2008, *ApJ*, 682, L13

Hopkins P. F., et al., 2023, *MNRAS*, 519, 3154

Huang Y.-H., Chen H.-W., Johnson S. D., Weiner B. J., 2016, *MNRAS*, 455, 1713

Huang S., Katz N., Scannapieco E., Cottle J., Davé R., Weinberg D. H., Peebles M. S., Brüggén M., 2020, *MNRAS*, 497, 2586

Huang X., Jiang Y.-f., Davis S. W., 2022, *ApJ*, 931, 140

Hummels C. B., Smith B. D., Silvia D. W., 2017, *ApJ*, 847, 59

Hummels C. B., et al., 2019, *ApJ*, 882, 156

Hunter J. D., 2007, *Computing in Science & Engineering*, 9, 90

Inoue T., Inutsuka S.-i., Koyama H., 2006, *ApJ*, 652, 1331

Jackson A. P., Townsley D. M., Calder A. C., 2014, *ApJ*, 784, 174

Jennings R. M., Li Y., 2021, *MNRAS*, 505, 5238

Ji S., Oh S. P., McCourt M., 2018, *MNRAS*, 476, 852

Ji S., Oh S. P., Masterson P., 2019, *MNRAS*, 487, 737

Jiang Y.-F., Oh S. P., 2018, *ApJ*, 854, 5

Joung M. K. R., Mac Low M.-M., 2006, *ApJ*, 653, 1266

Joung M. R., Bryan G. L., Putman M. E., 2012, *ApJ*, 745, 148

Jung S. L., Grønnow A., McClure-Griffiths N. M., 2023, *MNRAS*, 522, 4161

Kanjilal V., Dutta A., Sharma P., 2021, *MNRAS*, 501, 1143

Kereš D., Katz N., Weinberg D. H., Davé R., 2005, *MNRAS*, 363, 2

Kerstein A. R., 1988, *Combustion Science and Technology*, 60, 391

Kim J.-G., Kim W.-T., 2013, *ApJ*, 779, 48

Kim W.-T., Narayan R., 2003, *ApJ*, 596, 889

Kim C.-G., Ostriker E. C., 2017, *ApJ*, 846, 133

Kim C.-G., Ostriker E. C., 2018, *ApJ*, 853, 173

Kim C.-G., Ostriker E. C., Raileanu R., 2017, ApJ, 834, 25

Kim C.-G., et al., 2020a, ApJ, 900, 61

Kim C.-G., et al., 2020b, ApJ, 903, L34

Klein R. I., McKee C. F., Colella P., 1994, ApJ, 420, 213

Koyama H., Inutsuka S.-i., 2002, ApJ, 564, L97

Koyama H., Inutsuka S.-i., 2004, ApJ, 602, L25

Krumholz M. R., 2014, Phys. Rep., 539, 49

Kubryk M., Prantzos N., Athanassoula E., 2013, MNRAS, 436, 1479

Kuo K. K.-y., Acharya R., 2012, Fundamentals of turbulent and multiphase combustion. John Wiley & Sons

Kuo A. Y.-S., Corrsin S., 1972, Journal of Fluid Mechanics, 56, 447

Kwak K., Shelton R. L., 2010, ApJ, 719, 523

Kwak K., Henley D. B., Shelton R. L., 2011, ApJ, 739, 30

Kwak K., Shelton R. L., Henley D. B., 2015, ApJ, 812, 111

Lancaster L., Ostriker E. C., Kim J.-G., Kim C.-G., 2021, ApJ, 914, 89

Landau L. D., Lifshitz E. M., 1987, Fluid Mechanics. Butterworth-Heinemann

LeVeque R. J., 2002, Finite Volume Methods for Hyperbolic Problems. Cambridge Texts in Applied Mathematics, Cambridge University Press, doi:10.1017/CBO9780511791253

Lewiner T., Lopes H., Vieira A. W., Tavares G., 2003, Journal of graphics tools, 8, 1

Li M., Bryan G. L., Ostriker J. P., 2017, ApJ, 841, 101

Li Z., Hopkins P. F., Squire J., Hummels C., 2020, MNRAS, 492, 1841

Libby P. A., Bray K., Moss J., 1979, Combustion and Flame, 34, 285

Lim J., Ao Y., Dinh-V-Trung 2008, ApJ, 672, 252

Lockman F. J., Benjamin R. A., Heroux A. J., Langston G. I., 2008, ApJ, 679, L21

Lopez L. A., Mathur S., Nguyen D. D., Thompson T. A., Olivier G. M., 2020, ApJ, 904, 152

Lopez S., Lopez L. A., Nguyen D. D., Thompson T. A., Mathur S., Bolatto A. D., Vulic N., Sardone A., 2023, ApJ, 942, 108

Lynn J. W., Parrish I. J., Quataert E., Chandran B. D. G., 2012, ApJ, 758, 78

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

Mandelker N., van den Bosch F. C., Springel V., van de Voort F., 2019, ApJ, 881, L20

Mandelker N., Nagai D., Aung H., Dekel A., Birnboim Y., van den Bosch F. C., 2020, MNRAS, 494, 2641

Martin C. L., 1999, ApJ, 513, 156

Martínez-Gómez D., Oliver R., Khomenko E., Collados M., 2020, Astronomy and Astrophysics, 634, 1

Martizzi D., Faucher-Giguère C.-A., Quataert E., 2015, MNRAS, 450, 504

Martizzi D., Fielding D., Faucher-Giguère C.-A., Quataert E., 2016, MNRAS, 459, 2311

McCourt M., Sharma P., Quataert E., Parrish I. J., 2012, MNRAS, 419, 3319

McCourt M., O’Leary R. M., Madigan A.-M., Quataert E., 2015, MNRAS, 449, 2

McCourt M., Oh S. P., O’Leary R., Madigan A.-M., 2018, MNRAS, 473, 5407

McKee C. F., Cowie L. L., 1977, ApJ, 215, 213

McNamara B. R., et al., 2014, ApJ, 785, 44

McNamara B. R., Russell H. R., Nulsen P. E. J., Hogan M. T., Fabian A. C., Pulido F., Edge A. C., 2016, ApJ, 830, 79

McQuinn K. B. W., van Zee L., Skillman E. D., 2019, ApJ, 886, 74

Mellema G., Kurk J. D., Röttgering H. J. A., 2002, A&A, 395, L13

Melso N., Bryan G. L., Li M., 2019, ApJ, 872, 47

Miller M. J., Bregman J. N., 2015, ApJ, 800, 14

Muller C. A., Oort J. H., Raimond E., 1963, Academie des Sciences Paris Comptes Rendus, 257, 1661

Naab T., Ostriker J. P., 2017, ARA&A, 55, 59

Navarro J. F., Frenk C. S., White S. D. M., 1997, ApJ, 490, 493 (NFW)

Nelson D., et al., 2019, MNRAS, 490, 3234

Nelson D., et al., 2020, MNRAS, 498, 2391

Newman M. E. J., 2005, Contemporary Physics, 46, 323

Nguyen D. D., Thompson T. A., 2021, MNRAS, 508, 5310

Nguyen D. D., Thompson T. A., Schneider E. E., Lopez S., Lopez L. A., 2023, MNRAS, 518, L87

Nichols M., Bland-Hawthorn J., 2009, ApJ, 707, 1642

Niemeyer J. C., Hillebrandt W., 1995, ApJ, 452

Noon K. A., Krumholz M. R., Di Teodoro E. M., McClure-Griffiths N. M., Lockman F. J., Armillotta L., 2023, arXiv e-prints, p. arXiv:2304.06356

Oliver R., Soler R., Terradas J., Zaqarashvili T. V., Khodachenko M. L., 2014, *The Astrophysical Journal*, 784, 21

Otsu N., 1979, *IEEE Transactions on Systems, Man, and Cybernetics*, 9, 62

Padoan P., Nordlund Å., 2002, *ApJ*, 576, 870

Pandya V., et al., 2021, *MNRAS*, 508, 2979

Pandya V., et al., 2022, *arXiv e-prints*, p. arXiv:2211.09755

Peek J. E. G., Putman M. E., Sommer-Larsen J., Heiles C. E., Stanimirovic S., Douglas K., Gibson S., Korpela E., 2007, in *American Astronomical Society Meeting Abstracts*. p. 14.08

Peek J. E. G., Heiles C., Putman M. E., Douglas K., 2009, *The Astrophysical Journal*, 692, 827

Peeples M. S., et al., 2019, *ApJ*, 873, 129

Pember R. B., 1993, *SIAM Journal on Applied Mathematics*, 53, 1293

Péroux C., Howk J. C., 2020, *ARA&A*, 58, 363

Peters N., 1988, in *Symposium (International) on Combustion*. pp 1231–1250

Pittard J. M., Dyson J. E., Falle S. A. E. G., Hartquist T. W., 2005, *MNRAS*, 361, 1077

Pulido F. A., et al., 2018, *ApJ*, 853, 177

Putman M. E., et al., 2009, *ApJ*, 703, 1486

Putman M. E., Saul D. R., Mets E., 2011, *MNRAS*, 418, 1575

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

Reichardt Chu B., et al., 2022, *MNRAS*, 511, 5782

Rey M. P., Katz H. B., Cameron A. J., Devriendt J., Slyz A., 2023, *arXiv e-prints*, p. arXiv:2302.08521

Richter P., et al., 2017, *A&A*, 607, A48

Rubin K. H. R., Prochaska J. X., Koo D. C., Phillips A. C., Martin C. L., Winstrom L. O., 2014, *ApJ*, 794, 156

Russell H. R., et al., 2016, *MNRAS*, 458, 3134

Russell H. R., et al., 2017, *ApJ*, 836, 130

Russell H. R., et al., 2019, *MNRAS*, 490, 3025

Rybicki G. B., Lightman A. P., 1986, *Radiative Processes in Astrophysics*

Salem M., Besla G., Bryan G., Putman M., van der Marel R. P., Tonnesen S., 2015, *ApJ*, 815, 77

Sander B., Hensler G., 2021, *MNRAS*, 501, 5330

Sarazin C. L., 1986, *Reviews of Modern Physics*, 58, 1

Sarkar K. C., Sternberg A., Gnat O., 2022, ApJ, 940, 44

Scannapieco E., Brüggem M., 2015, ApJ, 805, 158

Schmidt W., Niemeyer J. C., Hillebrandt W., 2006, A&A, 450, 265

Schneider E. E., Robertson B. E., 2017, ApJ, 834, 144

Schneider E. E., Robertson B. E., Thompson T. A., 2018, ApJ, 862, 56

Schneider E. E., Ostriker E. C., Robertson B. E., Thompson T. A., 2020, ApJ, 895, 43

Schönrich R., Binney J., 2009, MNRAS, 396, 203

Sekora M. D., Stone J. M., 2010, Journal of Computational Physics, 229, 6819

Shapiro P. R., Field G. B., 1976, ApJ, 205, 762

Sharma P., McCourt M., Quataert E., Parrish I. J., 2012, MNRAS, 420, 3174

Shchelkin K., 1943, Zhurnal Eksperimental'noi i Teoreticheskoi Fiziki, 13

Slavin J. D., Shull J. M., Begelman M. C., 1993, ApJ, 407, 83

Smith F. J., 1963, Plan. Space Sci., 11, 1126

Smith E. P., Heckman T. M., Illingworth G. D., 1990, ApJ, 356, 399

Smith M. C., Bryan G. L., Somerville R. S., Hu C.-Y., Teyssier R., Burkhart B., Hernquist L., 2021, MNRAS, 506, 3882

Smith M. C., et al., 2023, arXiv e-prints, p. arXiv:2301.07116

Somerville R. S., Davé R., 2015, ARA&A, 53, 51

Sparre M., Pfrommer C., Ehlert K., 2020, MNRAS, 499, 4261

Spitzer L., 1962, Physics of Fully Ionized Gases

Sreenivasan K. R., Ramshankar R., Meneveau C., 1989, Proceedings of the Royal Society of London Series A, 421, 79

Steidel C. C., Erb D. K., Shapley A. E., Pettini M., Reddy N., Bogosavljević M., Rudie G. C., Rakic O., 2010, ApJ, 717, 289

Steinwandel U. P., Kim C.-G., Bryan G. L., Ostriker E. C., Somerville R. S., Fielding D. B., 2022, arXiv e-prints, p. arXiv:2212.03898

Stone J. M., Tomida K., White C. J., Felker K. G., 2020, ApJS, 249, 4

Strickland D. K., Heckman T. M., 2009, ApJ, 697, 2030

Swaminathan N., Bray K. N. C., 2011, Turbulent premixed flames. Cambridge University Press

Tan B., Oh S. P., 2021, MNRAS, 508, L37

Tan B., Oh S. P., Gronke M., 2021, MNRAS, 502, 3179

Tan B., Oh S. P., Gronke M., 2023, MNRAS, 520, 2571

Tennekes H., 1968, Physics of Fluids, 11, 669

Tennekes H., Lumley J. L., 1972, First Course in Turbulence

Thom C., Peek J. E. G., Putman M. E., Heiles C., Peek K. M. G., Wilhelm R., 2008, ApJ, 684, 364

Thompson T. A., Quataert E., Zhang D., Weinberg D. H., 2016, MNRAS, 455, 1830

Tonnesen S., Bryan G. L., 2021, ApJ, 911, 68

Tonnesen S., Stone J., 2014, ApJ, 795, 148

Townsend R. H. D., 2009, ApJS, 181, 391

Tripp T. M., 2022, MNRAS, 511, 1714

Tumlinson J., et al., 2011, Science, 334, 948

Tumlinson J., Peebles M. S., Werk J. K., 2017, ARA&A, 55, 389

Turk M. J., Smith B. D., Oishi J. S., Skory S., Skillman S. W., Abel T., Norman M. L., 2011, ApJS, 192, 9

Van Der Walt S., Colbert S. C., Varoquaux G., 2011, Computing in Science & Engineering, 13, 22

Van Woerden H., Wakker B. P., Schwarz U. J., de Boer K. S., eds, 2004, High Velocity Clouds Astrophysics and Space Science Library Vol. 312, doi:10.1007/1-4020-2579-3.

Vantghem A. N., et al., 2016, ApJ, 832, 148

Veilleux S., Cecil G., Bland-Hawthorn J., 2005, ARA&A, 43, 769

Virtanen P., et al., 2020, Nature Methods, 17, 261

Voit G. M., Donahue M., Zahedy F., Chen H.-W., Werk J., Bryan G. L., O'Shea B. W., 2019, ApJ, 879, L1

Wakker B. P., van Woerden H., 1991, A&A, 250, 509

Wakker B. P., York D. G., Wilhelm R., Barentine J. C., Richter P., Beers T. C., Ivezić Ž., Howk J. C., 2008, ApJ, 672, 298

Wakker B. P., Savage B. D., Fox A. J., Benjamin R. A., Shapiro P. R., 2012, ApJ, 749, 157

Walter F., Weiss A., Scoville N., 2002, ApJ, 580, L21

Werk J. K., et al., 2014, ApJ, 792, 8

Westmeier T., 2018, MNRAS, 474, 289

White M. C., Bicknell G. V., Sutherland R. S., Salmeron R., McGregor P. J., 2016, MNRAS, 455, 2042

Wiener J., Oh S. P., Zweibel E. G., 2017, MNRAS, 467, 646

Wiener J., Zweibel E. G., Ruszkowski M., 2019, MNRAS, 489, 205

Wiersma R. P. C., Schaye J., Smith B. D., 2009, MNRAS, 393, 99

Woosley S. E., Kerstein A. R., Sankaran V., Aspden A. J., Röpke F. K., 2009, ApJ, 704, 255

Yang Y., Ji S., 2023, MNRAS, 520, 2148

Zahedy F. S., Chen H.-W., Johnson S. D., Pierce R. M., Rauch M., Huang Y.-H., Weiner B. J., Gauthier J.-R., 2019, MNRAS, 484, 2257

Zel'Dovich Y. B., Pikel'Ner S. B., 1969, Soviet Journal of Experimental and Theoretical Physics, 29, 170

Zeldovich A. B., Barenblatt G. I., Librovich V. B., Makhviladze G. M., 1985, Mathematical theory of combustion and explosions. Consultants Bureau

Zhang D., Thompson T. A., Quataert E., Murray N., 2017, MNRAS, 468, 4801

Zheng Y., Putman M. E., Peek J. E. G., Joung M. R., 2015, ApJ, 807, 103

Zheng Y., et al., 2020, ApJ, 896, 143

Zhu Q., Smith B., Hernquist L., 2017, MNRAS, 470, 1017

Zimont V., Lipatnikov A., et al., 1995, Chem. Phys. Reports, 14, 993

de Avillez M. A., 2000, MNRAS, 315, 479

van de Voort F., Springel V., Mandelker N., van den Bosch F. C., Pakmor R., 2019, MNRAS, 482, L85

van der Velden E., 2020, The Journal of Open Source Software, 5, 2004