HYDRODYNAMICAL COUPLING OF MASS AND MOMENTUM IN MULTIPHASE GALACTIC WINDS

EVAN E. SCHNEIDER¹ AND BRANT E. ROBERTSON²

¹ Steward Observatory, University of Arizona, 933 North Cherry Avenue, Tucson, AZ 85721, USA

² Department of Astronomy and Astrophysics, University of California, Santa Cruz, 1156 High Street, Santa Cruz, CA 95064, USA

Received 2016 July 6; revised 2016 August 18; accepted 2016 September 6; published 2017 January 11

ABSTRACT

Using a set of high-resolution hydrodynamical simulations run with the Cholla code, we investigate how mass and momentum couple to the multiphase components of galactic winds. The simulations model the interaction between a hot wind driven by supernova explosions and a cooler, denser cloud of interstellar or circumgalactic media. By resolving scales of $\Delta x < 0.1$ pc over > 100 pc distances, our calculations capture how the cloud disruption leads to a distribution of densities and temperatures in the resulting multiphase outflow and quantify the mass and momentum associated with each phase. We find that the multiphase wind contains comparable mass and momenta in phases over a wide range of densities and temperatures extending from the hot wind ($n \approx 10^{-2.5} \text{ cm}^{-3}$, $T \approx 10^{6.5} \text{ K}$) to the coldest components ($n \approx 10^2 \text{ cm}^{-3}$, $T \approx 10^2 \text{ K}$). We further find that the momentum distributes roughly in proportion to the mass in each phase, and the mass loading of the hot phase by the destruction of cold, dense material is an efficient process. These results provide new insight into the physical origin of observed multiphase galactic outflows and inform galaxy formation models that include coarser treatments of galactic winds. Our results confirm that cool gas observed in outflows at large distances from the galaxy ($\gtrsim 1$ kpc) likely does not originate through the entrainment of cold material near the central starburst.

Key words: galaxies: evolution – hydrodynamics – ISM: clouds – supernovae: general – turbulence

1. INTRODUCTION

Star-forming galaxies commonly feature a multiphase galactic wind, observed at a wide variety of densities, temperatures, and velocities (e.g., Lehnert & Heckman 1996; Martin 2005; Rupke et al. 2005; Strickland & Heckman 2007; Tripp et al. 2011; Rubin et al. 2014) and over a large range of redshifts (e.g., Weiner et al. 2009; Coil et al. 2011; Nestor et al. 2011; Bouché et al. 2012; Kornei et al. 2012; Bordoloi et al. 2016). Despite their ubiquity, fully characterizing these winds can prove difficult. Spatially resolved observations of the wind’s many phases remain challenging, even for the nearest star-forming systems (Shopbell & Bland-Hawthorn 1998; Westmoquette et al. 2009; Rich et al. 2010; Leroy et al. 2015). Different observational techniques and instruments are required for different phases, so amassing a complete picture for even a single galaxy represents a large coordinated effort. At higher redshifts, absorption-line studies that trace outflowing gas in and around star-forming galaxies can be challenging to interpret as they require making assumptions about the wind’s geometry (e.g., Rubin et al. 2011; Bouché et al. 2012). While much progress has been made in recent years thanks to the installation of the Cosmic Origins Spectrograph on the *Hubble Space Telescope*, large uncertainties still exist regarding the contributions of different phases of winds to the net mass, momentum, and energy content of outflows (Heckman et al. 2015).

Winds also play an important role in theoretical studies of galaxy evolution. Supernova-driven winds provide an attractive method of feedback in cosmological simulations, allowing galaxies to regulate their star formation rates and gas supply over cosmic time (e.g., Oppenheimer & Davé 2008; Davé et al. 2011; Faucher-Giguère et al. 2011; Dalla Vecchia & Schaye 2012; Muratov et al. 2015). Recent simulations have successfully reproduced the galaxy stellar mass function across a wide range of redshifts by including phenomenologically

motivated wind models (Vogelsberger et al. 2014; Schaye et al. 2015; Davé et al. 2016). However, the processes that launch winds and govern their evolution as they escape galaxies remain unresolved on the scale of cosmological simulations. We currently must turn to smaller-scale, higher-resolution simulations to learn more about the physical nature of the winds themselves.

On these smaller physical scales, idealized simulations of galactic winds have also presented a theoretical challenge. Both analytic studies and hydrodynamic simulations of winds have had difficulty accelerating cool gas to the velocities observed in winds, because the dense phases get destroyed by hydrodynamic instabilities too quickly (e.g., Scannapieco & Brüggén 2015; Zhang et al. 2015; Brüggén & Scannapieco 2016). Magnetic fields may play an important role in stabilizing the cool gas (McCourt et al. 2015), but without realistic comparisons to observations the most important physical processes at play in multiphase winds are difficult to ascertain. A detailed analysis of the momentum and energy budget of gas in different phases in these hydrodynamic simulations has not yet been conducted. These data would be valuable both for improving subgrid prescriptions of winds in cosmological simulations and for comparing with observations to better determine where our theoretical understanding of winds fails. However, such a study requires high resolution across a large simulation volume in order to track the gas in different phases for significant periods of time.

In this work, we aim to improve our theoretical understanding of multiphase galactic winds via high-resolution, idealized simulations. Using the recently released graphics processing unit (GPU) based code Cholla³ (Schneider & Robertson 2015), we can perform hydrodynamic simulations of

³ A public version of the Cholla code is available at <http://github.com/cholla-hydro/cholla>.

the interaction between cool and hot phases of a starburst-driven wind at high resolution (<0.1 pc) over a large volume (>100 pc). The code performs well enough to compute such simulations on a static mesh and thus capture the interaction between the different phases of gas across a much larger region than any previous study (e.g., Cooper et al. 2009; Scannapieco & Brügger 2015; Banda-Barragán et al. 2016). The ability to track gas in each phase over long periods of time allows a direct probe of the momentum coupling between the hot and cool phases of the wind. In addition, the calculations add an element of physical realism to the cool gas by changing the initial density structure of the multiphase clouds to better match the features seen in spatially resolved outflows of dense gas.

Our simulations model a multiphase galactic wind as cold, dense interstellar medium (ISM) or circumgalactic medium clouds embedded within a hot, rarified background flow driven by supernovae. Because the cool material starts at rest with respect to the background wind, the initial interaction between the two phases drives a shock into the dense cloud. While the current work focuses on the cloud densities, shock Mach numbers, and physical scales relevant to galactic winds, the adopted numerical setup allows for comparisons with previous investigations of cloud–shock interactions.

Because of its ubiquity in the ISM, the shock–cloud interaction problem has been studied by many authors. Early numerical work by Klein et al. (1994) investigated the case of a planar shock interacting with a spherical cloud using two-dimensional (2D), adiabatic simulations. Their work indicated that clouds encountering a shock typically survive for a few “cloud crushing times,” roughly the timescale for the initial shock to propagate through the cloud. For strong shocks, the cloud crushing time depends on the density contrast between the cloud and the ambient medium, the size of the cloud, and the speed of the shock. Earlier under-resolved numerical work came to similar conclusions (Nittmann et al. 1982; Bedogni & Woodward 1990). These studies found that shocked clouds travel ~ 8 cloud radii before mixing with the ambient medium as a result of hydrodynamic instabilities. Adiabatic three-dimensional (3D) simulations (Stone & Norman 1992; Xu & Stone 1995) corroborated the 2D results and additionally attempted to account for different cloud geometries. Cloud geometry and orientation in those simulations did not affect the timescale for cloud fragmentation, but did substantially affect the late-time morphology of the clouds before they were destroyed.

These early studies could reasonably ignore radiative cooling effects by limiting their studies to small clouds. In larger-scale problems where the cooling timescale is smaller than the dynamical timescale, thermal energy losses must be included. Many authors have investigated this regime (e.g., Mellema et al. 2002; Fragile et al. 2004; Melioli et al. 2005; Cooper et al. 2009) and demonstrated that radiative cooling inhibits destruction of the dense material and extends the lifetime of the cloud relative to the adiabatic case. Rather than efficiently mixing with the hot post-shock wind, radiatively cooling clouds tend to get strung out into filaments containing individual “cloudlets” of dense gas that can survive much longer. Other authors have investigated the effects of conduction (e.g., Marcolini et al. 2005; Orlando et al. 2005; Brügger & Scannapieco 2016) and magnetic fields (e.g., Mac Low et al. 1994; Fragile et al. 2005; Shin et al. 2008; McCourt et al. 2015; Banda-Barragán et al. 2016) on the cloud–shock

interaction, with varying results for the stabilization of the cloud.

While multiple previous works studied a range of potentially important physics, few explored the impact of the initial structure of the cloud on the results of cloud–shock interactions. Early work focused on modeling supernova remnants in the ISM, and a simple spherical cloud provided a sufficient approximation for the initial conditions. In radiatively cooling galactic winds, however, the initial morphology of the cloud may have a profound effect on its evolution. Only Cooper et al. (2009) previously studied how the internal structure might influence the cloud destruction, using a fractal cloud as a proxy for a realistic cloud in a galactic wind. They found that fractal clouds survived for less time than initially spherical clouds. More recently, Schneider & Robertson (2015) examined how a turbulent interior cloud structure can alter the cloud crushing timescale in adiabatic simulations.

Our current study aims to better quantify the differences in the physical picture for inhomogeneous clouds and more broadly describe the way the gas phases in the outflow evolve. Specifically, we attempt to capture the region of parameter space relevant for the cool ($\sim 10^4$ K) clouds observed in galactic winds near the disks of star-forming galaxies. In this regime, the wind can be adequately modeled as a hot ($\sim 10^6$ K), supersonic fluid containing a population of embedded clouds of denser, cooler, initially stationary material. Depending on the exact density contrast between the cool and hot phases, the cooling timescale may fall below the local dynamic timescale, and the simulations therefore should include radiative cooling. Other potentially relevant effects, such as conduction and magnetic fields, we leave for future study.

An outline of our paper follows. We describe in Section 2 the model used to study the interaction between the multiple phases of the wind. In Section 3 we explain the setup of our wind simulations. Section 4 presents the qualitative evolution of the wind–cloud interaction, including the impact of the initial surface density of the cool gas on the cloud evolution. In Section 5, we describe in detail the density and temperature structure of the multiphase outflow. In Section 6 we study the velocities of the gas and describe how momentum distributes between different phases of the wind. Section 7 presents a resolution study focused on increasingly small scale features in turbulent clouds. Section 8 contains our interpretation of these results, including a discussion of our findings in relation to previous work, possible effects of incorporating additional physical processes, and an analysis of the fate of dense gas within a gravitational potential. We summarize in Section 9.

2. A MULTICOMPONENT WIND MODEL

Theories of starburst galaxies have long suggested that the combination of stellar winds and supernovae should drive a hot ($\sim 10^8$ K) wind out of the starburst region (e.g., Chevalier & Clegg 1985). This hot wind fluid remains difficult to observe directly, requiring high spatial resolution X-ray spectra. In the nearby starburst galaxy M82, where such observations are possible, the detection of a diffuse $\sim 4 \times 10^7$ K plasma in the central region indicates the presence of a hot wind (e.g., Griffiths et al. 2000; Strickland & Heckman 2007). Our current study assumes that stellar feedback can drive such a wind, and that the coupling of energy and momentum from the hot wind fluid with cooler gas leads to the multiphase winds seen in many starburst systems (see the review by Veilleux et al. 2005).

In this work, we seek to better quantify the coupling between the hot, rarified phase of galactic winds and the cooler, denser outflowing gas that is nearly ubiquitously observed in rapidly star-forming systems.

We attempt to account for the origin of a multiphase wind by modeling the interaction between a hot fast outflow and the cold, dense “clouds” it may entrain. Adequately capturing the hydrodynamic processes occurring in such a scenario requires resolutions two orders of magnitude smaller than the size of the cool clouds in question. Even with a tool like Cholla, modeling clouds with sizes of ~ 10 pc requires an idealized set of simulations to probe sufficiently the interactions between the hot and cool gas phases in a wind. Correspondingly, our simulations examine dense clouds in a box with a background wind (see Figure 2 below), representing cool material exposed to the hot phase of a wind. In the following two subsections, we detail our models for both the hot background and cool cloud components of these multiphase winds.

2.1. Hot Wind Component

We seek to model the impact of a hot, supernova-driven outflow on cooler material. Given a small set of assumptions, including spherical symmetry and negligible radiative cooling, the hot phase of the wind can be modeled analytically at distances close to the plane of a galaxy. All of our simulations use an analytic model of the hot wind as a constant background state with properties set using the adiabatic wind model of Chevalier & Clegg (1985, hereafter CC85). The CC85 model envisions a hot wind driven by central energy and mass input from stellar feedback processes. Three input parameters determine the solutions to the model: the energy input as a function of time, $\dot{E}$, the mass input as a function of time, $\dot{M}$, and the size of the driving region within which energy and mass are injected, R_* . By $\dot{M}$ we mean the total mass injection rate to the wind due to supernovae and mass loading, *not* the star formation rate. With these parameters, a solution to the set of spherical hydrodynamic equations can be found that smoothly transitions from subsonic within the driving region to supersonic at further radii. The solutions cross the sonic point at $r = R_*$.

We choose the input parameters for our version of the CC85 model according to the fits derived by Strickland & Heckman (2009) using *Chandra* X-ray observations of the nearby starburst galaxy M82. In particular, we set $\dot{E} = 10^{42}$ erg s $^{-1}$, $\dot{M} = 2 M_\odot$ yr $^{-1}$, and $R_* = 300$ pc. In interpreting their results, we have made additional assumptions about the wind mass-loading factor β and the supernova thermalization fraction α , for which they give a range of correlated values. Here we are using the Strickland & Heckman (2009) interpretation of β , meaning the fraction of total mass injected into the wind as compared to the mass injected by supernovae and stellar winds, $\beta = \dot{M}/\dot{M}_{\text{SN+SW}}$. Likewise, our definition of α corresponds to their ϵ and refers to the fraction of the supernova energy that is deposited in low-density gas and does not suffer large radiative losses before being incorporated into the wind. We take $\beta = 1.42$ and $\alpha = 0.33$, values near the middle of the acceptable range of fits reported by Strickland & Heckman (2009). These choices give a central temperature of $T_c \gtrsim 10^{7.5}$ K in the driving region, consistent with estimates made from the X-ray emission (see Strickland & Heckman 2009, Table 2). The resulting values for number density, velocity, temperature,

Figure 1. Adiabatic wind model used in all simulations. The top three panels display physical values of number density, radial velocity, and temperature as a function of radius. The fourth panel shows the dimensionless Mach number of the wind, which crosses the sonic point at $r = R_* = 300$ pc in our model. The wind–cloud simulations use values at $r = 1$ kpc for the background wind, shown with the dashed vertical line in each panel.

and Mach number as a function of radius for this model are displayed in Figure 1.

For the background flow in our multiphase wind simulations, we use the physical parameters of the CC85 wind model shown in Figure 1 at $r = 1$ kpc. These values are

$$\begin{aligned} n_{\text{wind}} &= 5.2626 \times 10^{-3} \text{ cm}^{-3}, \\ v_{\text{wind}} &= 1.1962 \times 10^3 \text{ km s}^{-1} = 1.2225 \text{ pc kyr}^{-1}, \\ P_{\text{wind}}/k &= 1.9881 \times 10^4 \text{ cm}^{-3} \text{ K}, \end{aligned}$$

where k is the Boltzmann constant. At $r = 1$ kpc, the wind pressure corresponds to a temperature of $T_{\text{wind}} = 3.7778 \times 10^6$ K, as displayed in Figure 1. We choose a radius of 1 kpc to set the background wind properties in our simulations for several reasons. First, we wish to capture the mass loading of the wind outside the driving region, which restricts us to hot wind properties at radii $r > 300$ pc. Second, we will model the interactions between cool and hot material in a simulation volume with a physical length of 160 pc. Given that the wind properties in our simulations remain approximately constant across this volume and the hot wind density, temperature, and Mach number change most rapidly just outside the driving region (see Figure 1), we favored an initial radius of $r \sim 1$ kpc over smaller radii. Further, the best observations of a multiphase wind come from M82, where cool material clearly resides at $r > 1$ kpc above the disk (Leroy et al. 2015). The fits derived by Strickland & Heckman (2009) suggest that the hot wind should not yet have suffered serious radiative losses at $r \sim 1$ kpc, which would invalidate the CC85 model (Zhang et al. 2015; Thompson et al. 2016).

2.2. Cool Cloud Component

To capture the multiphase nature of galactic winds, our simulations also include a cool component representing interstellar or circumgalactic material. As with previous studies of galactic winds, we model this second component as dense clouds initially at rest with respect to the hot wind (e.g., Scannapieco & Brüggén 2015; Brüggén & Scannapieco 2016). Our aim is to extend previous studies by examining the detailed momentum and energy coupling between different wind phases, so we consider both idealized spherical clouds and more realistic turbulent clouds with a distribution of interior densities set by turbulent processes. Observations have revealed cool gas in outflows at a variety of densities and temperatures (see references in Section 1). By varying the median density of the cold gas, $\tilde{n}$, and its interior density structure, we are able to model a range of properties for the dense component of the wind. Both the median number density and the cloud morphology affect the integrated surface density of the cold gas, $\Sigma_{\text{cl}} = M_{\text{cl}}/\pi R_{\text{cl}}^2$, where M_{cl} is the total mass of the cloud and R_{cl} is the cloud radius. The surface density influences how momentum transfers from the hot wind to the cold component, as discussed below.

As the hot wind destroys the clouds, cool material will both heat and rarify. This material will accelerate as momentum transfers from the hot wind, and enter the outflowing wind with a range of velocities. To adequately track all of this material, we perform simulations using boxes with long aspect ratios. The simulation boxes feature transverse dimensions equal to $8R_{\text{cl}}$ and a long dimension along the wind direction of $32R_{\text{cl}}$.

Figure 2 displays an example of an entire box, showing a density projection of the $\tilde{n} = 1 \text{ cm}^{-3}$ turbulent cloud at time $t = 0$. The clouds are initially centered $2R_{\text{cl}}$ from the left boundary of the box to capture the resulting bow shock. The long dimension of the box enables the simulations to follow the bulk of the cloud as it accelerates and track material stripped from the main cloud body.

2.2.1. Cloud and Sphere Models

Each cloud initially sits at rest and in thermal pressure equilibrium with the surrounding hot wind. The clouds in our study are not dense enough to be gravitationally confined, allowing us to neglect self-gravity. The density of the cloud material therefore determines its initial temperature. We initialize the spherical clouds with constant interior temperature and initialize the interior temperatures of regions within the turbulent clouds along an appropriate isobar. We list the temperatures and median and mean densities of the cloud initial conditions in Table 1. We list the full range of temperatures for the turbulent clouds—their median temperature matches the spheres. For both the spheres and turbulent clouds we taper the densities at the edge, starting at a radius of 4.5 pc, such that the cloud density smoothly transitions from the median to the wind density. The density taper has the form

$$n(r)_{\text{cl}} = \tilde{n} \exp[\ln(n_{\text{wind}}/\tilde{n})/(R_{\text{cl}} - 4.5)][r - 4.5], \quad (1)$$

where r is the distance from the center of the cloud and $\tilde{n}$ is the median cloud density. $R_{\text{cl}} = 5$ pc for all of our clouds. To create the turbulent clouds, we excise a region from a Mach 5 isothermal turbulence simulation (Robertson & Goldreich 2012) and scale the density linearly to match the desired median. Figure 3 shows a histogram of the initial gas densities for the $\tilde{n} = 1 \text{ cm}^{-3}$ turbulent cloud simulation.

2.2.2. Cloud Crushing Time

To interpret our results in relation to previous studies of cloud–shock interactions, we make use of the “cloud crushing time” t_{cc} initially devised by Klein et al. (1994). When the hot wind first interacts with the cool cloud, a shock drives into the overdense gas (and a reverse shock reflects into the oncoming wind). The cloud crushing time estimates how long the initial shock takes to pass through and compress the cloud. This timescale can be calculated in terms of known quantities by relating the initial density contrast of the cloud to the wind $\chi = n_{\text{cl}}/n_{\text{wind}}$ along with the radius of the cloud R_{cl} and the wind velocity v_{wind} as

$$t_{\text{cc}} = R_{\text{cl}} \chi^{1/2} / v_{\text{wind}}. \quad (2)$$

We use the cloud crushing time as an evolutionary timescale for the clouds in our simulations and report t_{cc} for each of our simulations in Table 1. The derivation leading to Equation (2) assumes that the pre-shocked gas in the cloud evolves adiabatically, allowing for an estimate of the shock speed within the cloud from the density contrast and wind speed. In a radiatively cooling cloud the pre-shock conditions change as the cloud cools, decreasing the sound speed and slowing the shock. The cloud crushing times listed in Table 1 therefore represent imperfect estimates of the cloud compression timescale, and our simulations show that they underestimate the

Figure 2. Initial conditions for one of our cloud–wind simulations. Each simulation box is much larger than the initial size of the cloud, so we can track the long-term evolution of cloud material, even after it has been stripped from the main body of the cloud. This density projection shows the $\bar{n} = 1 \text{ cm}^{-3}$ turbulent cloud. The initial density distribution for the cloud material in this simulation is displayed in Figure 3.

duration of this phase. In our simulations, maximum compression is typically reached at $\sim 2t_{\text{cc}}$.

2.2.3. Cloud Cooling Time

The cooling time provides another important evolutionary timescale for the clouds we simulate. If the cloud cooling time greatly exceeds the cloud crushing time, we expect the clouds to evolve similarly to the well-studied adiabatic case (Klein et al. 1994; Xu & Stone 1995; Poludnenko et al. 2002; Nakamura et al. 2006). If the cloud crushing time grows much longer than the cooling time, the radiative loss of energy may affect the cloud properties substantially before the initial shock can disrupt it. We can estimate the cooling time of the clouds in our simulation as their thermal energy divided by the rate at which energy is lost owing to radiative cooling as

$$t_{\text{cool}} = \frac{3n_{\text{cl}}kT_{\text{cl}}}{2\Lambda(T_{\text{cl}})}, \quad (3)$$

where k is the Boltzmann constant and $\Lambda(T)$ is the cooling rate in $\text{erg s}^{-1} \text{ cm}^{-3}$, evaluated at the cloud temperature (see Appendix C for information on our calculation of cooling rates). In theory, we might like to use T_{cl} and n_{cl} post-shock to calculate the cooling time. Given the complex nature of the cooling function, these quantities often prove impossible to calculate analytically (Creasey et al. 2011). Additionally, the lognormal gas distribution in our turbulent clouds means that n and T vary for different parts of the cloud. Instead, we calculate the cloud cooling times using the median pre-shock density and temperature to compare most directly with the cloud crushing time. These cooling times appear in Table 1. Our simulations sample the transition from clouds dominated by adiabatic evolution to those in the radiatively cooling regime. As discussed in Section 4, our simulations reproduce this transition and thereby justify our assumptions used to compute the cooling time.

3. SIMULATIONS

We ran a total of six “production” simulations with initial parameters displayed in Table 1. The first parameter we varied in these simulations was the cloud density distribution. Half the simulations modeled constant-density spherical clouds to compare with previous studies. The other half modeled clouds with a lognormal density distribution appropriate for a turbulent gas (e.g., Padoan & Nordlund 2002; Kritsuk et al. 2007). We also varied the median density of the clouds

from $\bar{n} = 0.1 \text{ cm}^{-3}$ to $\bar{n} = 1.0 \text{ cm}^{-3}$ to sample a range of density and temperature phase space. At the low end, this density range samples the transition from adiabatic to radiative evolution of the cool gas. At the higher densities, the evolution of the clouds is very similar when scaled by the cloud crushing time (see Section 4), but higher-density clouds have increasingly lengthy lifetimes. As a result, our upper limit on cloud density is set by computational expense.

Throughout this paper, we refer to the production simulations by the names given in Table 1. Simulations with spherical clouds begin with an “S,” and those with turbulent clouds are denoted “T.” The remainder of the name references the median density of the cloud. For example, the simulation featuring a turbulent cloud with a median number density of $\bar{n} = 1 \text{ cm}^{-3}$ is “T1,” the spherical cloud with a median number density of $\bar{n} = 0.1 \text{ cm}^{-3}$ is “S01,” etc. We will sometimes refer to the spherical clouds as “spheres”; “cloud” alone will always refer to a turbulent cloud.

Each of the simulations listed in Table 1 was run in a volume with $N = 2048 \times 512 \times 512$ cells and physical dimensions of $160 \times 40 \times 40 \text{ pc}$, yielding a physical resolution for these simulations of $\Delta x = 0.07825 \text{ pc/cell}$. We also carried out a resolution study (see Section 7), with higher- and lower-resolution simulations for the $\bar{n} = 0.5$ turbulent cloud. Previous work on this problem has relied on techniques such as adaptive mesh refinement and cloud-tracking (moving the reference frame of the simulation to follow the main body of the cloud) to reduce computational costs. In contrast, we have constant resolution across our entire volume, and as a result we capture the evolution of both low- and high-density material. Our long boxes also allow us to track material at distances farther from the cloud than in previous studies.

We ran our simulations with the Cholla hydrodynamics code (Schneider & Robertson 2015), using piecewise parabolic reconstruction, an HLLC Riemann solver, a simple unsplit integrator, and a dual-energy scheme. Optically thin radiative cooling with a photoionizing UV background was assumed and implemented using precomputed Cloudy tables (Ferland et al. 2013). Details of the integration scheme, Riemann solver, dual energy, and radiative cooling appear in the appendices. We used outflow boundaries for all sides of the box except the left x -face, where the inflow was set according to the wind properties described in Section 2.1. All of the production simulations were run for a minimum of $25t_{\text{cc}}$. The ultra-high-resolution simulation used in our resolution study was only run for $12t_{\text{cc}}$, as it is extremely expensive.

Table 1
Summary of Simulations Performed, and Physical Parameters of Interest for the Cloud in Each Simulation

Model Name	$\bar{n}_{\text{cl}}^{\text{a}}$ (cm^{-3})	$\bar{n}_{\text{cl}}^{\text{b}}$ (cm^{-3})	$\Sigma_{\text{cl}}^{\text{c}}$ ($M_{\odot} \text{ pc}^{-2}$)	N_{H}^{d} (cm^{-2})	T_{cl}^{e} (K)	M_{cl}^{f} ($M_{\odot}$)	χ^{g}	t_{cc}^{h} (kyr)	$t_{\text{cool}}^{\text{i}}$ (kyr)
S01	0.1	0.086	0.013	2.87×10^{18}	1.988×10^5	1.05	1.9×10^1	17.8	26.0
S05	0.5	0.45	0.064	1.43×10^{19}	3.976×10^4	5.04	9.5×10^1	39.8	3.96
S1	1.0	0.88	0.13	2.85×10^{19}	1.988×10^4	10.0	1.9×10^2	56.4	1.30
T01	0.1	0.16	0.022	2.05×10^{19}	2.2×10^3	1.73	1.9×10^1	17.8	26.0
T05	0.5	0.83	0.11	1.02×10^{20}	4.4×10^2	8.61	9.5×10^1	39.8	3.96
T1	1.0	1.7	0.24	2.23×10^{20}	2.0×10^2	18.8	1.9×10^2	56.4	1.30

Notes. Simulations with a spherical cloud begin with an “S”; those with a turbulent cloud begin with a “T.” All simulations listed in Table 1 are run in a volume with a numerical resolution of $2048 \times 512 \times 512$ cells, corresponding to a physical box size of $160 \times 40 \times 40$ pc.

^a Median density of the cloud, calculated including all material with a density greater than $1/10 \bar{n}$.

^b Mean density of the cloud, calculated including all material with a density greater than $1/10 \bar{n}$.

^c Surface density of the cloud, calculated as $M_{\text{cl}}/(\pi R_{\text{cl}}^2)$.

^d Maximum initial column density sight line through the cloud (excluding the hot wind).

^e Initial temperature of the cloud material, in pressure equilibrium with the wind. The median temperature is provided for the spherical clouds, and the temperature of the highest-density regions is provided for the turbulent clouds.

^f Initial mass of the cloud, calculated including all material with a density greater than $1/10 \bar{n}$.

^g Initial density contrast of the cloud with the background wind, calculated using $\chi = \bar{n}_{\text{cl}}/n_{\text{wind}}$.

^h Cloud crushing time, calculated using Equation (2) with the median density and a cloud radius of 5.0 pc.

ⁱ Cloud cooling time, calculated using Equation (3) with the median cloud density and temperature.

Figure 3. Normalized histogram of the initial density distribution for the $\bar{n} = 1$ turbulent cloud, shown in blue, fit with a Gaussian in log-space. Note: this distribution does not include the density taper at the edges of the cloud, as its purpose is to illustrate the lognormal distribution of the bulk of the cloud’s mass.

The six simulations listed in Table 1 have high enough resolution to adequately capture the hydrodynamic instabilities that disrupt the spherical clouds, $64 \text{ cells}/R_{\text{cl}}$. These simulations were used to produce the majority of the results in this paper. In addition, we have run low-resolution versions of each of the simulations in Table 1 in a volume with half as many cells ($N = 1024 \times 256 \times 256$). The primary purpose of these low-resolution runs was to verify initial conditions, estimate runtimes, and test for convergence. We have also run a low-resolution simulation with a different mean molecular weight (see Appendix C) to test the effect of the cooling implementation on our results.

In total, we have run 14 simulations for this paper. The seven low-resolution simulations were carried out on the El Gato cluster at the University of Arizona. The high-resolution production simulations and ultra-high-resolution comparison simulation were carried out on the Titan supercomputer at the

Oak Ridge Leadership Computing Facility via a director’s discretionary time allocation. In sum, these high-resolution simulations took ≈ 1.5 million Titan core-hr, which corresponds to 50,000 GPU-hr. The time each simulation required varied greatly based on the total length of cloud survival, from the lowest-density $\bar{n} = 0.1$ spherical cloud simulation that required ≈ 1500 GPU-hr to the $\bar{n} = 1.0$ spherical cloud simulation that required ≈ 9400 GPU-hr. The $\bar{n} = 0.5$ ultra-high-resolution turbulent cloud simulation required $\approx 14,000$ GPU-hr despite only running to $12t_{\text{cc}}$.

4. COOL CLOUD EVOLUTION

We begin our results with a general description of the evolution of the cool gas in the clouds, focusing on the $\bar{n} = 1$ turbulent cloud simulation (T1), our fiducial case. The evolution of cool material when exposed to a hot wind is of interest, given the uncertainty in the theoretical community about whether it is possible to accelerate cool material to the hot wind speed without destroying it (e.g., Scannapieco & Brüggén 2015; Zhang et al. 2015). If cool gas can be efficiently entrained, the process could provide an explanation for the presence of cool material observed at large distances from galaxies (e.g., McCourt et al. 2015). The efficiency with which cool material is destroyed will also affect the mass-loading factor of the hot wind, which may in turn affect the chances of thermal instability and rapid cooling of the hot wind (Wang 1995; Thompson et al. 2016). In this section, we describe the overall evolution of the cool gas in our simulations as a function of both cloud morphology and initial surface density. In particular, we focus on the destruction time of the cool material.

The early stages of adiabatic shock–cloud interactions have been thoroughly described in the literature, particularly beginning with the comprehensive study of Klein et al. (1994). Cloud evolution is often described in terms of the cloud crushing time, $t_{\text{cc}} = \chi R_{\text{cl}}/v_{\text{wind}}$ (see Section 2.2). In an adiabatic simulation, the initial compression of the cloud is followed by a downstream expansion, after which the cloud quickly fragments and dense material is destroyed. The entire cloud destruction process typically takes $4t_{\text{cc}}\text{--}5t_{\text{cc}}$ for spheres.

If the cloud crushing time is calculated using the median gas density, adiabatic turbulent clouds are destroyed in a similar number of crushing times (Schneider & Robertson 2015).

When dense clouds are able to cool radiatively, their evolution follows a different path. The early cloud crushing phase is much more effective as the heat generated by compression is radiated away. In a radiative cloud, gas densities can reach over an order of magnitude above their pre-shock levels. Though radiative clouds still get disrupted within $10t_{cc}$, the individual dense “cloudlets” that result can take as long as $40t_{cc}$ to mix into the hot wind (Scannapieco & Brüggén 2015).

4.1. Turbulent Clouds versus Spheres

In Figures 4 and 5 we illustrate the general evolution of the $\bar{n} = 1 \text{ cm}^{-3}$ turbulent and spherical cloud simulations (models T1 and S1; Table 1). These figures show the column density of gas in the box (including the hot wind), in both a y -axis projection with the hot wind entering the box from the left and an x -axis projection in the direction of the wind. The figures show snapshots of the simulations at five representative times—the initial conditions, and after $2t_{cc}$, $5t_{cc}$, $10t_{cc}$, and $20t_{cc}$.

In both simulations, the maximum average cloud density is reached at $t = 2t_{cc}$. The maximum average density is $\bar{n} = 5.2 \text{ cm}^{-3}$ for T1 and $\bar{n} = 6.8 \text{ cm}^{-3}$ for S1. (We calculate the mean density using material with density greater than $1/10$ the initial median density, as in Table 1.) Despite the similarity of this compression timescale, Figures 4 and 5 show a drastic difference in cloud morphology at $t = 2t_{cc}$. While the initial shock has propagated at different speeds through regions of different density for the turbulent cloud, the shock has very effectively compressed the sphere into a single flat pancake. As a result, the evolution of the two types of clouds diverges strongly at later times. Low-density material is quickly accelerated between $t = 5$ – $20t_{cc}$ in the turbulent cloud, leaving only a few high-density cloudlets at late times. We quantify this rapid mass loss in Figure 6, which shows the normalized cloud mass as a function of cloud crushing time. The mass is calculated using material at or above $1/3$ the initial median density. By $t = 20t_{cc}$, nearly 80% of the original mass of the turbulent cloud has been mixed into the hot wind and fallen below the density threshold of $n = 0.33 \text{ cm}^{-3}$.

The loss of material proceeds more rapidly for turbulent clouds than for spheres, even with the cloud crushing time calculated as a function of the median density. As can be seen in the third panel of Figure 5, after the initial pancake stage, the spherical cloud compresses into a single core, with a small surface area and high column density in the wind direction. This morphology is a result of the original shock moving in the wind direction combined with the compression from shocks on the sides of the cloud, which push material toward the center. This high-density core presents a smaller surface area for ablation of material than the many small high-density regions in the turbulent cloud (compare the third panels of Figures 4 and 5). The single bow shock at the front of the spherical cloud protects the core, while also resulting in a pressure gradient through the cloud. Cloud elongation in the direction of the wind gradually breaks up the original core into smaller fragments, each of which has an individual bow shock and loses material. However, as demonstrated in Figure 6, the overall process is much slower for spheres than for turbulent clouds. Only $\sim 40\%$ of the original spherical cloud material has been lost by $t = 20t_{cc}$.

4.2. Median Density and Cloud Lifetimes

Figures 4 and 5 show the evolution of clouds with a median density of $\bar{n} = 1 \text{ cm}^{-3}$. We have also run four high-resolution simulations with lower median densities. Models T05 and S05, the turbulent and spherical clouds with a median density of $\bar{n} = 0.5$, show a mass and morphology evolution similar to their higher-density counterparts. Figure 6 shows that the mass loss as a function of cloud crushing time is nearly identical for the $\bar{n} = 0.5 \text{ cm}^{-3}$ and $\bar{n} = 1 \text{ cm}^{-3}$ clouds out to $25t_{cc}$. However, as the original median density continues to decrease, the evolution begins to follow a qualitatively different path. The original shock that passes through the cloud does not compress the gas enough for it to reach densities where it can cool efficiently. The warm cloud gas quickly gets rarified and accelerated with the hot wind, causing the mass loss of the clouds to proceed much more rapidly at lower original median densities. This difference is clearly visible in the $\bar{n} = 0.1$ evolutionary tracks in Figure 6.

In model S01, almost none of the gas reaches the requisite density to cool. In fact, only a small ring of gas within the cloud reaches the threshold density. This ring is visible in Figure 7, which shows surface density projections of the $\bar{n} = 0.1$ sphere simulation at $2t_{cc}$. The ring structure is caused by the shock moving in the wind direction colliding with the oblique shocks coming in from the sides of the cloud. While the densities in this collision do not get large enough to result in a single compact core as seen in the higher-density models, this small amount of dense gas is enough to cause the extended tail seen in the mass evolution track of Figure 6. However, most of the cloud is destroyed in less than $5t_{cc}$, consistent with adiabatic simulations of cloud–shock interactions (Schneider & Robertson 2015).

The evolution of the $\bar{n} = 0.1 \text{ cm}^{-3}$ turbulent cloud (model T01) is less dramatically different as compared to its higher-density counterparts. Although the median cloud density is below that required for effective cooling, the lognormal density distribution of initial cloud material spans a range up to $n \approx 10 \text{ cm}^{-3}$. As a result, parts of the cloud that were initially above the median density are still able to cool effectively after being shocked. As Figure 6 shows, the post-shock mass loss for the $\bar{n} = 0.1$ turbulent cloud proceeds faster than the $\bar{n} = 1$ or $\bar{n} = 0.5$ clouds, but more slowly than the $\bar{n} = 0.1$ sphere. We expect that the speed of mass loss for turbulent clouds would continue to increase as the initial median density is lowered, until eventually the evolution proceeded on an adiabatic track (shown by the dotted black line in Figure 6).

5. PHASE STRUCTURE OF THE WIND

In this section, we investigate the physical state of the gas in our simulations. The typical temperature of gas of a given density is useful in interpreting the cloud evolution. The density threshold for rapid cooling is also an important feature that determines whether any dense gas will survive for many cloud crushing times. The high resolution of our simulations across the entire volume enables this study of the detailed phase structure of the gas as it evolves.

5.1. Density and Temperature Structure

Figure 8 shows density–temperature phase diagrams for the $\bar{n} = 1$ turbulent cloud at four different times in its evolution—the initial conditions, and $t = 5t_{cc}$, $10t_{cc}$, and $20t_{cc}$. Each bin is

Figure 4. Time series evolution of our fiducial simulation, the $\bar{n} = 1 \text{ cm}^{-3}$ turbulent cloud (model T1). Plots on the left show surface density projected along the y-axis, with the wind entering the box from the left, while the right column shows the surface density projected in the direction of the wind velocity. The scale of the axis ticks is 10 pc. Snapshots are shown at $t = 0, 2t_{cc}, 5t_{cc}, 10t_{cc}$, and $20t_{cc}$. The dashed circle in the right column shows the original extent of the cloud. Note: the y-projection at $t = 20t_{cc}$ has been shifted by 30 pc to recenter the cloud material.

Figure 5. Time series evolution of the $\tilde{n} = 1 \text{ cm}^{-3}$ sphere simulation (model S1). As in Figure 4, the left column shows surface density projected along the y-axis, and the right column shows the surface density projected along the wind axis. Snapshots are shown at $t = 0$, $2 t_{cc}$, $5 t_{cc}$, $10 t_{cc}$, and $20 t_{cc}$. Note: the reference frame of the y-projection at $t = 20 t_{cc}$ has been shifted by 40 pc to recenter the cloud.

Figure 6. Mass evolution (cloud mass divided by initial cloud mass as a function of time) of each of our high-resolution simulations. Cloud mass is calculated as the sum of all gas in the simulation with a density greater than $1/3$ the initial median density. Tracks for turbulent clouds are shown as solid lines, while spherical clouds are shown with dashed lines. The mass-loss track for a $\tilde{n} = 0.1$ adiabatic cloud simulation is also plotted (dotted line) for comparison.

Figure 7. Surface density projections of the $\tilde{n} = 0.1$ sphere simulation at $t = 2t_{\text{cc}}$. The left panel shows a y-projection, and the right panel shows an x-projection, as in Figures 4 and 5. A small ring of high surface density material is formed by the collision of shocks within the cloud, but most of the cloud remains at densities too low to cool effectively. The dashed circle in the right panel shows the original extent of the cloud.

colored according to the total mass it contains, to better focus attention on the locations with the majority of the cloud material. Throughout the simulation, most of the mass is in the hot wind, visible as the red region at the upper left of the distribution in each panel (and as a single bin in the initial conditions). At $t = 0$, the cloud's mass is distributed across a range of densities and temperatures, with most of the mass in regions at or above the median density of $\tilde{n} = 1 \text{ cm}^{-3}$. Each panel also contains an isobar showing the original pressure of the gas (recall that the cloud's pressure is matched to the wind in the initial conditions), as well as a dashed line that shows the equilibrium location between heating and cooling in our simulations.

After the cloud has been shocked, the density–temperature phase diagram takes on a characteristic shape. At early times, most of the gas is at higher pressure than the initial conditions. As the simulation proceeds, the cloud gas slowly evolves back

Figure 8. Mass-weighted density–temperature phase diagrams at $t = 0, 5t_{\text{cc}}, 10t_{\text{cc}},$ and $20t_{\text{cc}}$ for the $\tilde{n} = 1$ cloud simulation (model T1). The initial equilibrium pressure is plotted as a solid line, and the temperature at which heating equals cooling as a function of density is plotted as a dotted line.

Figure 9. Mass-weighted density–temperature phase diagrams for the $\tilde{n} = 0.1$ turbulent (top) and spherical (bottom) cloud simulations at $2t_{cc}$. While much of the mass in the turbulent cloud has crossed the $\log(n) = 0$ threshold and is able to cool efficiently, much of the mass of the spherical cloud remains below this density. The low-density gas is unable to cool and is quickly mixed into the hot wind.

toward thermal pressure equilibrium with the incoming wind. Cloud material quickly fills out the entire range of densities between the original cloud density and the wind. A large amount of mass has been compressed to high densities, much of it over an order of magnitude higher than the initial densities. The high-density material cools effectively, with much of the cloud mass located just above the equilibrium cooling temperature, at the lower right in each panel. There also exists a mass concentration around $\log(n) = 0.5$ and $\sim 2 \times 10^4$ K. This buildup reflects the shape of the cooling curve, with maximally efficient cooling around 10^5 K and a steep falloff around 10^4 K. These features remain throughout the subsequent evolution, as the mass in high-density bins is gradually reduced.

The phase diagrams for simulations T01 and S01 yield an estimate of the density threshold required for efficient cloud cooling in our simulations. Figure 9 shows the mass-weighted density–temperature phase diagrams for both simulations at $t = 2t_{cc}$. The turbulent cloud, displayed in the left panel, has a large amount of mass in the cooling sweet spot just above $\log_{10}(n) = 0$. In contrast, most of the shocked sphere material never crosses this density threshold, and as a result, the gas remains at high temperatures and low densities, where it

quickly mixes with the hot wind. Animations of the phase diagrams show clearly the path that the cool gas takes in the turbulent cloud simulation. Cloud gas originally above the median density shocks to densities just above $\log_{10}(n) = 0$, after which the gas begins to cool rapidly down to 10^4 K. Only a small fraction of the sphere material reaches densities of $\log_{10}(n) = 0$, and as a result, most of the mass of the cloud is lost at early times.

6. MOMENTUM COUPLING

While the phase diagrams of the cloud gas are enlightening, in this section we seek to quantify several less obvious aspects of the cloud–wind interaction. The first area we address is cool cloud entrainment—the acceleration of cool material by the wind. We investigate entrainment in our simulations by studying 2D density–velocity histograms, to determine the typical velocities attained by gas of a given density. We follow our discussion of entrainment with an investigation of the detailed coupling of momentum between the hot wind and the cool cloud material. We assess how cloud mass transitions from one phase to another and how the momentum in different phases changes over time. As in previous sections, we will focus on the fiducial $\tilde{n} = 1 \text{ cm}^{-3}$ turbulent cloud simulation.

6.1. Cool Cloud Entrainment

While cool clouds have been observed in galactic outflows at a variety of distances and velocities, the primary physical mechanism responsible for fast-moving cool gas continues to be debated. In this section, we investigate the velocity of the cool gas in our simulations and show that in the purely hydrodynamic case, hot winds are unlikely to accelerate cool gas to the high velocities seen in many outflows.

Figure 10 shows mass-weighted density–velocity histograms for model T1. The figure shows two representative snapshots, at $t = 5t_{cc}$ and $15t_{cc}$, with the velocity in the wind direction plotted on the y-axis. As in the density–temperature diagrams, the hot wind appears as a mass concentration at the upper left in each panel, with a small spread around the initial wind density and velocity. While cloud material has clearly acquired a range of densities, only low-density material travels at high speeds. Only 3% of the cloud mass above the original median density of $\tilde{n} = 1 \text{ cm}^{-3}$ is moving faster than 200 km s^{-1} by $15t_{cc}$. The average velocity of the dense material ($n > 1 \text{ cm}^{-3}$) is only $v_x \approx 120 \text{ km s}^{-1}$.

One of the goals of this work was to investigate how different cloud density structures change the effectiveness of the hot wind in accelerating high-density material. This question is addressed in Figure 11, which compares the density–velocity histograms for the $\tilde{n} = 1 \text{ cm}^{-3}$ turbulent cloud at $t = 10t_{cc}$ to the sphere. This figure indicates that accelerating high-density material proves more difficult in the turbulent cloud case. This difference results from the different initial column densities of the cloud material, as we describe below.

In a simple model, the ram pressure of the hot wind will accelerate cold cloud material (we neglect other forces in the following analysis). The ram pressure of the hot wind scales as

$$P_{\text{ram}} = n_{\text{wind}} v_{\text{wind}}^2, \quad (4)$$

which gives a constant ram pressure per unit mass of $P_{\text{ram}}/m_H = 7.2 \times 10^{13} \text{ cm}^{-1} \text{ s}^{-2}$ for our background wind conditions. The associated acceleration of the cloud, g_{cl} , will

Figure 10. Mass-weighted density–velocity phase diagrams at $5t_{cc}$ and $15t_{cc}$ for the $\bar{n} = 1 \text{ cm}^{-3}$ turbulent cloud simulation. High-density material is not effectively accelerated in turbulent clouds.

then be

$$g_{cl} = \frac{P_{ram}}{m_H} \frac{m_H}{\Sigma_{cl}} = \frac{P_{ram}}{m_H} \frac{1}{N_H}, \quad (5)$$

where Σ_{cl} is the average surface density of the cloud, and $N_H = n_{cl}L$ is the column density of a particular region of the cloud. If the cloud is a constant-density sphere, then L can be approximated as twice the cloud radius. For the $\bar{n} = 1 \text{ cm}^{-3}$ spherical cloud with a radius of $r \approx 5 \text{ pc}$, the column density is approximately $N_H \approx 3 \times 10^{19} \text{ cm}^{-2}$ across the whole cloud, and the resulting acceleration is

$$g_{cl} \sim 2.3 \times 10^{-6} \text{ cm s}^{-2} \left[\left(\frac{n_{cl}}{1 \text{ cm}^{-3}} \right) \left(\frac{L}{10 \text{ pc}} \right) \right]^{-1}, \quad (6)$$

or

$$g_{cl} \sim 0.75 \text{ km s}^{-1} \text{ kyr}^{-1} \left[\left(\frac{n_{cl}}{1 \text{ cm}^{-3}} \right) \left(\frac{L}{10 \text{ pc}} \right) \right]^{-1}. \quad (7)$$

The cloud crushing time for the $\bar{n} = 1$ clouds is $t_{cc} \approx 56.4 \text{ kyr}$. The acceleration in Equation (7) gives a cloud velocity of $v_{cl} \sim 85 \text{ km s}^{-1}$ after $2t_{cc}$ (113 kyr), which is roughly consistent with our results for the velocities of the densest gas after $2t_{cc}$, shown in the right panel of Figure 11. In fact, the average velocity of gas denser than $n = 500 \text{ cm}^{-3}$ at $2t_{cc}$ for the $\bar{n} = 1$ sphere is 77 km s^{-1} . After the cloud has been crushed,

Figure 11. Mass-weighted density–velocity phase diagrams for the $\bar{n} = 1$ cloud (left) and sphere (right) at $2t_{cc}$. Acceleration of dense gas is much more effective for spherical clouds, which has lower initial column density than the dense regions of turbulent clouds with the same initial median density.

Equation (7) is no longer an adequate estimate of the cloud acceleration, because the column densities have increased by over an order of magnitude (compare the first and second panels in Figure 5).

In contrast, at $t = 0$ the densest sight lines through the $\bar{n} = 1 \text{ cm}^{-3}$ turbulent cloud have column densities as high as $N_H = 2 \times 10^{20} \text{ cm}^{-2}$, almost an order of magnitude larger than the sphere (see the first panel in Figure 4). At these column densities, the cloud acceleration is only $g_{cl} \approx 0.11 \text{ km s}^{-2} \text{ kyr}^{-1}$, yielding a dense gas velocity of $v_{cl} \sim 12 \text{ km s}^{-1}$ at $t = 2t_{cc}$. As a result, the regions of high column density are accelerated less in the turbulent cloud, and the high-density gas at $t = 2t_{cc}$ is traveling considerably slower than in the spherical cloud case. The average velocity of gas with $n > 500 \text{ cm}^{-3}$ is only $v \approx 16 \text{ km s}^{-1}$ at $t = 2t_{cc}$.

In both the turbulent cloud and sphere simulations, much less acceleration occurs at late times than predicted by Equation (7) (see right panel of Figure 10). As mentioned previously, this minimal acceleration can be explained by the increased column densities that result from the cloud crushing. In both the turbulent and spherical clouds, compression toward the center caused by the initial shock results in column densities at $t = 2t_{cc}$ that are an order of magnitude higher than the original values (compare the first and second panels in both Figures 4 and 5). Thus, we conclude that entrainment of cool material to high velocities in a hot wind remains very

inefficient under these circumstances, and the problem only compounds for clouds with more realistic shapes and density distributions.

These simulations also show that the average surface density of the cloud, $\Sigma_{\text{cl}} = M_{\text{cl}}/\pi R_{\text{cl}}^2$, as given in Table 1 fails to adequately predict how much dense gas will be accelerated. This fact can be demonstrated by comparing the dense gas acceleration for the $\tilde{n} = 0.5 \text{ cm}^{-3}$ turbulent cloud (not shown) and the $\tilde{n} = 1 \text{ cm}^{-3}$ sphere. The $\tilde{n} = 0.5 \text{ cm}^{-3}$ turbulent cloud originally has a lower average surface density than the $\tilde{n} = 1$ sphere— $\Sigma_{\text{cl}} = 0.11$ and $\Sigma_{\text{cl}} = 0.13 M_{\odot} \text{ pc}^{-2}$, respectively. However, the spatial coherence of the densest regions within the turbulent cloud leads to individual column densities several times higher than the average column density of the sphere, around $N_{\text{H}} \approx 10^{20} \text{ cm}^{-2}$ as compared to $N_{\text{H}} = 3 \times 10^{19} \text{ cm}^{-2}$. As a result, the average velocity of the densest material in the $\tilde{n} = 0.5 \text{ cm}^{-3}$ turbulent cloud simulation is $v_x \approx 30 \text{ km s}^{-1}$, less than half the speed of the dense gas in the $\tilde{n} = 1 \text{ cm}^{-3}$ spherical cloud simulation.

6.2. Integrated Mass and Momentum

Having established that in our simulations momentum does not transfer efficiently enough to accelerate the dense phase to the wind velocity, we would like to better quantify the momentum gained by other phases of the outflow. Figure 6 showed the evolution of cloud mass for each of our high-resolution simulations as an integrated quantity, the total sum of material above a given density threshold. To better understand the way that gas evolves within the cloud, we now divide this evolution into multiple density bins. Figure 12 shows this binned mass evolution plot for the $\tilde{n} = 1$ cloud. The black line matches the one displayed in Figure 6, comprising all the material above a density threshold of $1/3 \tilde{n}$ but no longer normalized by the initial cloud mass. Colored lines in Figure 12 show the evolution of total cloud mass in four density bins: low, from $0.02 < n < 0.2 \text{ cm}^{-3}$; medium, from $0.2 < n < 2.0 \text{ cm}^{-3}$; high, from $2.0 < n < 20 \text{ cm}^{-3}$, and very high, $n > 20 \text{ cm}^{-3}$. The other three panels show one-dimensional (1D) histograms of the mass as a function of density. Integrating any of the colored regions in the histogram yields the value represented as a circle on the line of the same color in the top left panel.

Much of the evolution in Figure 12 takes place early on, with the most drastic shift occurring before $2t_{\text{cc}}$ as the cloud is crushed. Initially, much of the cloud mass ($12 M_{\odot}$) is in the high-density bin, $2.0 < n < 20 \text{ cm}^{-3}$, with a significant amount ($5.9 M_{\odot}$) also in the medium-density bin that surrounds the initial median density of $n = 1 \text{ cm}^{-3}$. Only $0.64 M_{\odot}$ is at $n > 20.0 \text{ cm}^{-3}$ initially. As the cloud is crushed, almost all of the material increases by about an order of magnitude in density, such that at $2t_{\text{cc}}$ the medium-density bin has only $0.85 M_{\odot}$, the high-density bin has $6.0 M_{\odot}$, and $11.6 M_{\odot}$ is in the very high density bin. These values are highlighted by circles at $2t_{\text{cc}}$ in the first panel of Figure 12. Very little mass is in the low-density bin at this time, and no cloud mass has been lost.

After $2t_{\text{cc}}$, the evolution proceeds more gradually, with material slowly moving from higher-density bins to lower-density ones. At $5t_{\text{cc}}$ mass is concentrated in the same locations evident in the density–temperature phase diagram—there is a significant amount of mass ($10.0 M_{\odot}$) at very high densities, as well as a bump around $\log_{10}(n) = 0.5$ as a result of the shape of the cooling curve. These features are still present at $10t_{\text{cc}}$,

though the mass in the highest-density bin has been substantially reduced. At all times in our simulation after $2t_{\text{cc}}$, the majority of the cloud mass is in the highest-density bin. Eventually we expect this mass to shift to lower-density bins as the last of the high-density gas is destroyed.

In the same way that we have integrated the total cloud mass, we can integrate total cloud momentum in the simulations. Figure 13 shows how the total momentum is divided among the same four density bins. The top left panel shows the evolution of the total momentum in the wind direction, computed as

$$p_{\text{tot}} = \sum_{n=n_{\text{low}}}^{n=n_{\text{high}}} M_i v_{x,i} \quad (8)$$

for each cell, i , in the simulation, where M_i is the integrated mass in that cell and v_x is the direction of the wind. The sum is taken over the same density ranges shown in Figure 12. In the other three panels, we show histograms of the momentum as a function of density at $t = 2t_{\text{cc}}$, $5t_{\text{cc}}$, and $10t_{\text{cc}}$. Integrals of the histograms are again shown as circles on the colored lines in the top left panel. At $t = 2t_{\text{cc}}$ the momentum is distributed in a similar way to the mass, with most of the momentum in the two highest-density bins. At $t = 5t_{\text{cc}}$, however, the distribution of total momentum has shifted. The three lower-density bins have continued to gain momentum, but the highest-density bin actually loses some, despite the fact that the highest-density bin still contains the majority of the cloud mass. At this point, the total momentum in the medium- and high-density bins ($0.2 \text{ cm}^{-3} < n < 20 \text{ cm}^{-3}$) is $p_{\text{tot}} = 495 M_{\odot} \text{ km s}^{-1}$, 50% more than is in the highest-density bin ($n > 20 \text{ cm}^{-3}$), $p_{\text{tot}} = 315 M_{\odot} \text{ km s}^{-1}$.

This shift indicates that momentum is transferred efficiently from the hot wind to lower-density gas. Gas with densities between $1 \text{ cm}^{-3} < n < 10 \text{ cm}^{-3}$ is in a relatively long-lived phase, giving it more time to gain momentum from the hot wind. The total momentum in the warm phase likely also increases as gas from higher-density bins moves to lower density. Gas in the highest-density bin with the most momentum may be the most likely to decrease in density. Without tracer particles, this shift is difficult to quantify in our simulations, but the highest-density bin clearly loses mass at every time $t > 2 t_{\text{cc}}$ and at least some high-density material with significant momentum shifts to the lower-density bins. Overall, the distribution of total momentum across the different gas phases in the cloud is remarkably equal. The density–velocity histograms shown in Figure 10 also indicate this equality, showing that lower-density material moves more quickly than higher-density material. At very late times, $t > 17 t_{\text{cc}}$, the highest-density bin again has the most total momentum—this reflects the lack of mass in the lower bins by this time.

We can also compare the distribution of momentum between the turbulent and spherical clouds by looking at the integrated momentum in the wind direction, a “momentum surface density”:

$$\Sigma_p = \sum_{i=0}^{i=N_x} n_i v_{x,i}, \quad (9)$$

where N_x is the total number of cells in the wind direction, n_i is the number density of the gas in cell i , and $v_{x,i}$ is the velocity in the wind direction in cell i . We show projections of this quantity for the $\tilde{n} = 1$ turbulent cloud and sphere at $t = 2t_{\text{cc}}$ in Figure 14. As one would expect given the differences in the

Figure 12. Evolution of cloud mass in different density bins for the $\bar{n} = 1 \text{ cm}^{-3}$ turbulent cloud. In the upper left panel we show the total mass in different density bins; colors of the lines in the top left panel correspond to density bins in the 1D histograms shown in the other three panels. Circles in the first panel correspond to the times displayed in the other three panels. The 1D histograms are normalized such that the integral over each density bin yields the total mass indicated by the circles.

density–velocity phase diagrams, the momentum surface density is much higher for the sphere, reaching a maximum of $2.4 \times 10^{23} \text{ km s}^{-1} \text{ cm}^{-2}$ in the highest column. The momentum surface density of the background wind is $3.1 \times 10^{21} \text{ km s}^{-1} \text{ cm}^{-2}$. In contrast, the momentum has a larger spatial extent and is spread over a wider range of gas column densities in the turbulent cloud, with the maximum column of $4.3 \times 10^{22} \text{ km s}^{-1} \text{ cm}^{-2}$. Thus, even though the average momentum surface density at this time is similar between the two simulations, the momentum is clearly much more concentrated in the high-density gas in the spherical cloud simulation. This result is consistent with the total momentum being shared across every density bin in Figure 13.

7. RESOLUTION EFFECTS

At our fiducial resolution of 64 cells/ R_{cl} , we expect the global properties of cloud evolution to be qualitatively correct (Gregori et al. 2000; Poludnenko et al. 2002; Melioli et al. 2005). These properties include the morphologies seen in the cloud destruction in Section 4, the general features seen

in the phase diagrams in Section 5, and the multiphase distribution of momentum demonstrated in Section 6.2. However, many authors have suggested that a higher resolution of at least 128 cells/ R_{cl} is required for convergence of quantitative measurements of properties like the mass-loss rate and destruction time (Klein et al. 1994; Nakamura et al. 2006; Scannapieco & Brüggén 2015). The effects of resolution may be especially important for turbulent cloud simulations, where increasingly dense structures are captured as the resolution of the simulation is increased. Therefore, we have carried out a numerical study to test the dependence of our results on the resolution of our simulations. We emphasize that this is *not* a convergence study in the classical sense, because the initial conditions of the cloud change with resolution (as described below).

The study compares three versions of the $\bar{n} = 0.5$ turbulent cloud simulation: an ultra-high-resolution simulation run with 128 cells/ R_{cl} (hereafter R_{128}), the production version with 64 cells/ R_{cl} (R_{64}), and a low-resolution version with 32 cells/ R_{cl} (R_{32}). Because of the large computational expense of the R_{128} simulation, we use a physical box of size $80 \times 30 \times 30 \text{ pc}$, as

Figure 13. Evolution of total momentum in the wind direction in different density bins for the $\tilde{n} = 1$ turbulent cloud. The top left panel shows the total momentum as a function of cloud crushing time, calculated by summing M_{v_x} over all cells with a density in the given range (see Equation (8)). Colors represent the same density bins shown in Figure 12. Other panels show 1D histograms of momentum as a function of density at $2t_{\text{cc}}$, $5t_{\text{cc}}$, and $10t_{\text{cc}}$. Integrating under the histograms gives the values shown as circles in the top left panel.

Figure 14. Comparison of the momentum column density ($\Sigma_p = m v_x$ integrated in the wind direction) for the $\tilde{n} = 1$ turbulent and spherical cloud simulations at $5t_{\text{cc}}$. The scale of the axis ticks is 10 pc. White dashed circles indicate the original extent of the cloud. Much higher momentum column densities are visible for the spherical cloud, which has been crushed to a single high-density core.

compared to the production simulations, which ran in boxes of size $160 \times 40 \times 40$ pc. The R_{128} simulation volume contains $2048 \times 784 \times 784$ cells, which yields a resolution of $\Delta x = 0.039$ pc. The reduced time step required by the Courant condition and the computational expense of the additional cells mean that we can only afford to run the simulation for $\sim 12t_{\text{cc}}$. In practice, using the shorter box length results in the material leaving the computational domain at later times. The R_{32} simulation is run in a box with the same physical size as the production simulation.

Each of the three simulations uses a cloud drawn from the same region of a Mach 5 isothermal turbulence simulation (Robertson & Goldreich 2012). To create the lower-resolution clouds, the highest-resolution simulation was resampled using a cubic spline interpolation. The bulk physical properties of the initial conditions are identical, including the median number density and total mass. As the resolution increases, we allow the cloud initial conditions to include progressively more small-scale structures and higher-density regions. The initial conditions of the higher-resolution simulations therefore do not

Figure 15. Comparison at $10t_{cc}$ between $\tilde{n} = 0.5$ turbulent cloud simulations with three different resolutions: 32 cells/ R_{cl} , 64 cells/ R_{cl} , and 128 cells/ R_{cl} . The intensity in each panel corresponds to the projected number density, and color reflects the temperature of the gas (purple is cold, green is warm, and red is hot). As the resolution of the simulations is increased, the high-density features are resolved into smaller structures. As a result, the densest gas in the cloud is accelerated less efficiently with increasing resolution.

simply reflect a better resolved version of low-resolution run initial conditions. The comparison presented below does not aim to act as a convergence study, but instead attempts to

capture how increasingly smaller scale features of the cloud initial conditions might influence the evolution of the wind-cloud system.

Figure 16. Cloud mass is displayed as a function of time for the three $\bar{n} = 0.5$ turbulent clouds in the resolution study. The resolution of each simulation in terms of cells/cloud radius is displayed in the lower left. Mass loss initially proceeds more quickly for the higher-resolution simulations.

Figure 15 shows snapshots of the three simulations at $10t_{cc}$, with the simulation time limited by the expense of the R_{128} run. The intensity of the image in each panel scales logarithmically with the projected number density, and the color reflects the gas temperature. The R_{32} and R_{64} simulations have been cropped to display the same region as the R_{128} simulation, for which the full box is shown. As expected, with each increase in resolution, finer-scale structure emerges. While in the R_{32} simulation only ~ 10 cloudlets form, the increased resolution and the more detailed initial conditions of the R_{128} simulation result in far more. At low resolutions the wind has successfully pushed the dense gas further, as evidenced by the bulk of the cloud material shifting further to the right in the upper panels. In the R_{128} simulation, some of the dense gas has traveled less than $2R_{cl}$ in $10t_{cc}$.

For a more quantitative measure of the difference between these simulations, we plot in Figure 16 the mass evolution of each cloud. Even at our highest resolution, the results have not yet converged. Figure 16 shows a general trend toward more efficient mass loss as the resolution is increased. To better understand this trend, we examine the mass evolution in the separate density bins used in Section 6.2. The resulting mass-loss curves for the R_{128} and R_{64} simulations are shown in Figure 17. At early times, the primary difference between the two simulations appears in the two highest-mass bins. The R_{128} gains slightly more mass in the highest-density bin ($n > 20 \text{ cm}^{-3}$), but contains significantly less mass in the bin with $2 \text{ cm}^{-3} < n < 20 \text{ cm}^{-3}$ than the R_{64} simulation.

At early times ($t < 2t_{cc}$) the R_{128} simulation generates less high-density gas ($n > 1 \text{ cm}^{-3}$)—the density threshold required for efficient cooling. A higher fraction of the cloud gas in the R_{128} simulation resides at densities and temperatures susceptible to quick destruction by the hot wind. This process of incorporation into the hot wind happens quickly, as evidenced by the lack of mass in the low-density, slow ($v < 200 \text{ km s}^{-1}$) region in the density–velocity diagrams (see Figure 10). Because cloud gas does not spend much time in this region of phase space, the low-density curves in Figure 17 look similar. At later times, mass loss proceeds similarly for the

Figure 17. Evolution of cloud mass in four density bins for the R_{128} (dashed) and R_{64} (solid) simulations. The decreased mass in the $2 \text{ cm}^{-3} < n < 20 \text{ cm}^{-3}$ density bin results in faster mass loss at early times for the higher-resolution simulation, as the lower-density cloud gas is quickly accelerated by the wind.

high-resolution simulations (compare the slope of the blue and green curves after $\sim 5t_{cc}$ in Figure 16).

8. DISCUSSION

The results presented in the previous sections have important implications for the current theoretical understanding of galactic winds. First, we have showed that the structure of the dense clouds plays a significant role in their evolution when exposed to a hot wind. Clouds that start with lognormal density distributions set by turbulent processes are mixed more quickly into the hot wind than spherical clouds with similar masses or median densities. This result implies that the mass loading of galactic winds in the regions near the base of a hot outflow is an efficient process that likely results in significant mass-loading factors when hot gas interacts with denser clouds.

While our simulations indicate that mass loading likely proves more efficient when the dense clouds have a turbulent structure, we have also demonstrated that this same structure in clouds tends to inhibit dense gas entrainment. The total acceleration of cloud material by the hot wind closely relates to the initial column density of the cool gas. In a turbulent cloud, the spatial coherence of high-density regions leads to large column densities along individual sight lines, which makes dense gas difficult to accelerate. While individual dense clumps can remain long-lived as a result of efficient cooling, dense gas ($n > 1 \text{ cm}^{-3}$) in our simulations rarely achieves velocities higher than $\sim 200 \text{ km s}^{-1}$. Cloudlets tend to get ablated and mixed into the hot wind before traveling more than $\sim 30R_{cl}$.

These findings support a picture of galactic winds where mass loading into the hot phase operates efficiently near the base of an outflow, and any surviving dense gas accelerates very little. However, this efficient mass loading will result in hot winds that are more susceptible to thermal instability as they expand and cool. As calculated in Thompson et al. (2016), high mass-loading factors will decrease the radius at which gas can cool out of the hot wind. If the gas that cools out of the hot phase retains significant velocity, it could explain the high-velocity neutral gas seen at distances 1–100 kpc from starbursting galaxies without a need to resort to entrainment of cool gas directly associated with the ISM.

Furthermore, we have calculated in detail the distribution of mass and momentum associated with the clouds in our simulations. Far from being a simple two-phase medium, these winds are characterized by gas that spans a large range of densities and temperatures. Momentum from the hot wind couples to these phases with different efficiencies, such that the total momentum in each phase tends to be similar even if the mass is not. In cosmological simulations, where resolution limits the ability to capture these features of the multiphase galactic outflows, our calculations could be leveraged to improve treatments of the temperature and momentum distribution of the wind phases.

We now compare our findings to similar work in the literature. While many authors have studied the cloud–shock problem, relatively few have investigated the parameter space relevant to galactic winds, and we will focus our attention on these results. We also discuss the potential effects that additional physics could have on our results, including different cooling rates, conduction, and magnetic fields. Finally, we include an analysis of the fate of the dense gas in our simulations in the presence of a gravitational potential.

8.1. Cloud Structure

Cooper et al. (2009) presented the only previous study that has investigated the destruction of clouds with a density distribution that is not symmetric along multiple axes. In that work, the authors compared the destruction of a radiatively cooling fractal cloud to that of a radiatively cooling spherical cloud. The background hot wind properties in Cooper et al. (2009) resembled ours. Our results agree well with theirs, in that they found that fractal clouds were more susceptible to fast destruction than spheres. However, their computational volume was too small to follow the evolution of the gas for many cloud crushing times, and their resolution was relatively poor, and thus the fate of the small dense clumps that result from the destruction of an inhomogeneous cloud remained unclear. In our work, we find that while these small cores can survive for tens of cloud crushing times as a result of efficient cooling, they are very difficult to accelerate due to their high column density. Additionally, higher resolution tends to amplify these effects. As a result, the cloudlets in our simulations are gradually destroyed over the course of $t \sim 20t_{cc}$ as the dense gas gets eaten away by the hot wind.

8.2. Entrainment and Mass Loading

In the past several years, several studies have investigated the ability of hot winds to entrain cool gas and carry it large distances from a galaxy. Scannapieco & Brügger (2015) studied cloud–wind interactions using adaptive mesh refinement simulations with a maximum resolution equivalent to the fixed resolution of our simulations. The primary purpose of their work was to explore the effects that different background wind parameters had on cool cloud lifetimes and acceleration. Using spherical clouds of $T \approx 10^4$ K, the authors derived a scaling for cloud destruction time that only depended on the Mach number of the hot wind, with different density contrasts accounted for in the cloud crushing time. Here, we compare our results to that scaling at t_{50} , defined as the time when 50% of the cloud material is at or above $1/3$ the initial cloud density.

Scannapieco & Brügger (2015) find

$$t_{50} = 4t_{cc} \sqrt{1 + M_{wind}}, \quad (10)$$

which for our background wind would correspond to $t_{50} = 10t_{cc}$. Looking at Figure 6, we see that this actually fits the turbulent clouds quite well, but our spherical clouds have a much longer lifetime.

At first glance, this result may seem contradictory. However, the longer spherical cloud lifetimes may be explained by the different treatment of cooling in our simulations. As explained in Appendix C, we allow gas in our simulations to cool to temperatures as low as $T = 10$ K using cooling rates calculated for solar-metallicity gas. We also include the effects of a photoionizing background. The resulting heating keeps low-density gas ($n < 1 \text{ cm}^{-3}$) warm, but is much less effective at higher densities. In contrast, Scannapieco & Brügger (2015) simulated larger-scale clouds with the assumption of complete ionization and therefore only allowed cooling above $T = 10^4$ K. The ability of gas to cool to temperatures well below $T = 10^4$ K in our simulations results in greater cloud compression. The resulting smaller surface area decreases the efficiency with which material ablates and correspondingly increases the cloud lifetime.

We find further evidence for this explanation by performing simulations with a different mean molecular weight μ . All the simulations presented in this paper used $\mu = 1$ when converting from mass density to number density,

$$n = \rho / \mu m_p, \quad (11)$$

where m_p is the mass of a proton. However, we also performed a low-resolution turbulent cloud simulation with $\mu = 0.6$, the value appropriate for fully ionized solar-metallicity gas. In comparing the lifetime of the cloud in this simulation with the same cloud in the standard simulation, we find that the $\mu = 0.6$ cloud is destroyed more slowly, though the effect is small. This slight difference can be explained by the higher number densities for a given mass density in the $\mu = 0.6$ simulation. These higher number densities lead to more efficient cooling, which in turn leads to higher average densities at early times. The resulting high-density clumps of gas prove more difficult for the hot wind to destroy than in the equivalent $\mu = 1$ model. Hence, the $\mu = 0.6$ cloud lives longer.

Scannapieco & Brügger (2015) also find higher final velocities for their clouds than we do, even when comparing only spherical clouds. For clouds with similar background wind conditions they find velocities $v \sim 200\text{--}300 \text{ km s}^{-1}$ at $t = 15t_{cc}$ (see their Figure 7), while we never see velocities above 200 km s^{-1} . Again, this difference is likely a result of lower-temperature gas in our simulations. Because our clouds can compress to higher densities in the initial stages of the wind interaction, their column densities increase more and acceleration is more difficult. This effect is compounded in the turbulent cloud case with higher initial column densities.

8.3. Additional Physics

Other studies of cool gas in the context of galactic winds have investigated the effects of additional physics in cloud–wind interaction simulations. Brügger & Scannapieco (2016) perform a similar study to Scannapieco & Brügger (2015) but incorporate the effects of conduction. They find that conduction

can result in complete evaporation of the cloud at early times in cases where the column density of cloud material is below $N \simeq 1.5 \times 10^{18} \text{ cm}^{-2}$. The clouds we simulate do not start with column densities this low, and we therefore do not expect conduction to result in their rapid, complete destruction. We note that the initial conditions of our hot wind most resemble their Mach 6.4 run that shows the least amount of difference in cloud evolution between the conduction and nonconduction models.

Brüggen & Scannapieco (2016) demonstrate that when clouds in their simulations do survive, conduction tends to decrease the cross section of the cloud presented to the hot wind. The decreased cross section provides less surface area for ram pressure to act on the dense gas, which in turn decreases the efficiency of cloud acceleration. This effect is qualitatively similar to the impact of inhomogeneous column densities for turbulent clouds described in Section 6.1, though the origin is completely different. We suggest that the inhomogeneous cloud structure may work in concert with conduction to further increase the early destruction of low column density material and decrease the efficiency with which high column material accelerates.

Magnetic fields may also play a role in cool cloud evolution. A number of cloud–shock studies investigated the effects of planar magnetic fields in spherical clouds, with inconclusive results regarding whether the presence of field lines increased or decreased the time until cloud destruction (Gregori et al. 2000; Fragile et al. 2005; Shin et al. 2008). More recently, McCourt et al. (2015) performed wind–cloud simulations testing the effects of tangled magnetic fields incorporated within a spherical cloud. They showed that the presence of the magnetic field drastically increased the lifetime of the resulting cloudlets and increased their acceleration, allowing them to reach the hot wind speed without destruction.

The McCourt et al. (2015) simulations help motivate a future extension of our high-resolution wind–cloud simulations with realistic initial conditions to include MHD. The simulations by McCourt et al. (2015) use a resolution of 32 cells/ R_{cl} . Scannapieco & Brüggen (2015) demonstrated that spherical cloud simulations with resolution less than 64 cells per cloud radius tend to prolong cloud lifetimes. Our simulations of turbulent clouds show a trend toward faster mass loss with increasing resolution that has not yet converged at 128 cells/ R_{cl} . In addition, we have shown that spherical clouds have longer lifetimes than those with more realistic internal structure. While the results of McCourt et al. (2015) suggest that including a tangled magnetic field would likely prolong the lifetime of the dense components of the turbulent clouds we simulate, the combined effects of resolution and density structure make a quantitative prediction difficult. Incorporating magnetic fields in wind simulations with turbulent clouds is an avenue that we aim to pursue in future work.

In addition to potentially important additional physics, parameters such as the metallicity of the cloud gas and nature of the background wind in our simulations will affect the quantitative results we have presented. In this work we have focused exclusively on the effects of cloud structure and surface density. The primary effect of changing the metallicity of the gas would be to change the cooling rates. As noted in Section 8.2, tests with a different value of μ indicate that more rapid cooling leads to longer cloud survival, and our comparison with the simulations of Scannapieco & Brüggen

(2015) indicates that less efficient cooling reduces cloud survival time. However, this is not an order-of-magnitude effect, and we therefore do not expect a change in metallicity to drastically alter our results.

On the other hand, as Figure 1 shows, the state of the background wind in the Chevalier & Clegg (1985) model changes rapidly with radius as the wind escapes the galaxy. Given our choice of a single background wind state, we do not regard our simulations as completely generic with respect to galactic winds, though we do expect the general result of more rapid destruction of turbulent clouds to hold. Scannapieco & Brüggen (2015) demonstrated a scaling with background wind Mach number that indicates that clouds survive longer with increasing Mach number (see their Equation (22)). In the future, we would like to investigate the relationship between turbulent cloud lifetime and the background wind parameters, sampling a variety of distances from the galaxy that would inform the initial conditions for clouds at each distance.

8.4. Ram Pressure versus Gravity

As a final note, we consider the final fate of the dense gas in our simulations. Our simulations do not include gravity, and if the simulations continued running indefinitely, eventually all of the cool material would mix into the outflowing hot wind. However, we can estimate the effect of the host galaxy’s gravitational potential. Using an analysis similar to that in Section 6.1, we can compare the expected acceleration of dense cloud regions owing to the wind’s ram pressure, P_{ram} , to their expected deceleration owing to gravity.

We can estimate the ram pressure acceleration as a function of column density, $N_{\text{H}} = n_{\text{cl}}L$, via

$$g_{\text{ram}} \sim 2.3 \text{ km s}^{-1} \text{ kyr}^{-1} \left(\frac{N_{\text{H}}}{10^{19} \text{ cm}^{-2}} \right)^{-1}.$$

(This expression is equivalent to Equation (7).) Similarly, we can estimate the gravitational acceleration as a function of column density. At 1 kpc, the gravitational acceleration of M82 is

$$g_{\text{grav}} \sim 1.4 \times 10^{-7} \text{ cm s}^{-2} \left(\frac{M_{\text{M82}}}{10^{10} M_{\odot}} \right) \left(\frac{R}{1 \text{ kpc}} \right)^{-2},$$

$$g_{\text{grav}} \sim 0.044 \text{ km s}^{-1} \text{ kyr}^{-1}. \quad (12)$$

Given these estimates, we would expect cloud regions with column densities greater than $N_{\text{H}} \approx 5 \times 10^{20} \text{ cm}^{-2}$ to begin to fall back toward the galaxy. None of the clouds in our simulations initially have column densities quite this high—the densest sight lines in the $\tilde{n} = 1 \text{ cm}^{-3}$ turbulent cloud reach $N_{\text{H}} \approx 3 \times 10^{20} \text{ cm}^{-2}$. Nonetheless, the acceleration due to ram pressure and deceleration due to gravity for these column densities are very similar, so in the presence of gravity the densest gas in our turbulent cloud simulations would likely be accelerated very little or possibly fall back toward the central galaxy.

9. SUMMARY AND CONCLUSIONS

In this work, we have modeled the hydrodynamic evolution of radiatively cooling clouds in the context of galactic winds with very high numerical resolution. Our study investigated two main parameters relevant to cold cloud survival—the

initial structure of the cool gas and the median density of the cloud. We varied the cloud structure in our simulations between a lognormal density distribution with large-scale structure as set by turbulent processes and an idealized spherical distribution of gas. The median densities of our clouds were in the range of $\bar{n} = 0.1\text{--}1.0\text{ cm}^{-3}$. The median density affects the overall destruction time of the cool gas via the cloud crushing time, as well as the efficiency of cooling within the cloud.

We find that clouds with a turbulent density structure are destroyed more quickly than clouds with a homogeneous spherical density distribution. This efficient destruction results in faster mass loading of the hot wind, as intermediate- and low-density regions of turbulent clouds are quickly heated, rarified, and accelerated to the hot wind velocity. The entrainment of dense gas within cool turbulent clouds proves extremely inefficient, and much less efficient than for idealized spherical initial conditions. The varying column densities present in turbulent clouds result in very little acceleration of the densest regions, which are the only regions that survive for many cloud crushing times. These effects are amplified as the resolution of the simulations is increased and the clouds are allowed to become increasingly realistic. We therefore conclude that entrainment of turbulent ISM clouds in hot supernova winds does not explain the neutral gas observed at large distances from starburst galaxies, unless other physical processes (such as magnetic fields) substantially alter the results from the hydrodynamic case.

We have also provided an extensive description of the phase structure of the gas in the wind. Shortly after being shocked, the gas associated with the turbulent clouds spreads over a large range of densities and temperatures, with the densest regions cooling down to temperatures of $T \sim 100\text{ K}$. Each phase of gas remains close to thermal pressure equilibrium with the hot ($\gg 10^6\text{ K}$) wind. Interestingly, though the majority of the mass remains in the densest phases ($n > 20\text{ cm}^{-3}$) for much of the cloud evolution, the total momentum distributes fairly evenly across densities. Roughly the same amount of momentum transfers to cold neutral (hundreds of kelvin), cool ionized ($\sim 10^4\text{ K}$), and warm ionized ($\sim 10^5\text{ K}$) gas.

We acknowledge inspiration from the 2014 and 2016 Simons Symposia on Galactic Superwinds: Beyond Phenomenology. Some simulations from this work were performed on the EL GATO system at University of Arizona, funded by National Science Foundation grant no. 1228509 (PI Brant Robertson). This research also used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725, via an award of computer time on the Titan system provided by the INCITE program (AST 109 and AST 117). We thank the anonymous referee for a constructive report. We also acknowledge Romeel Davé, J. Xavier Prochaska, Evan Scannapieco, and Todd Thompson for discussions and comments that improved this work.

APPENDIX A INTEGRATION METHOD AND HLLC RIEMANN SOLVER

In Schneider & Robertson (2015), we presented the Cholla implementation of the six-step corner transport upwind (CTU) integration scheme originally described in Gardiner &

Stone (2008). The unsplit CTU integrator preserves symmetry and minimizes numerical diffusion, making the scheme a good choice for many magnetohydrodynamics problems (see Stone et al. 2008). Despite performing well in the standard Liska & Wendroff (2003) tests, we found that the transverse flux correction step of CTU led to estimates for the 1D density and energy fluxes that produced negative densities and pressures after the final update, particularly when used to model multidimensional strong radiative shocks. In fact, our implementation of the CTU integration method always fails for the simulations in this work, regardless of the choice of interface reconstruction method or Riemann solver. Therefore, we instead employed a very simple, robust integration scheme for the simulations presented in this paper, described below.

The integrator we use follows the initial steps of CTU, including piecewise parabolic reconstruction of the interface values in the primitive variables, characteristic time evolution of those interface values, and 1D Riemann solutions at each interface to calculate fluxes. We implemented each of these initial steps as described in Schneider & Robertson (2015). However, rather than updating the interface values using the transverse fluxes, we simply skip to the final update and use only the 1D fluxes to evolve the conserved quantities. Because the fluxes do not contain corrections for transverse directions, we found that this method requires a very low CFL number to be stable—we use $cfl = 0.1$ for all the simulations in this work. The low CFL number makes the integrator expensive despite its simplicity, but when combined with a dual-energy formalism the method is quite robust. Figure 18 shows the results of this new integration scheme as compared to the CTU method on the implosion test described in Liska & Wendroff (2003). As the implosion test shows, this integrator is slightly more diffusive than CTU. The additional diffusion allows this simple integrator to succeed in circumstances where CTU fails.

We also employ an HLLC Riemann solver for this work (see Harten et al. 1983; Toro et al. 1994). The Cholla implementation of the HLLC solver follows the description in Batten et al. (1997), which revises the original HLLC solver presented in Toro et al. (1994) with updated maximum and minimum wave speeds that account for the possibility of colliding shocks within a cell. The HLLC solver has several advantages over the exact and Roe Riemann solvers presented in Schneider & Robertson (2015). Unlike the exact solver, the HLLC solver does not require an iterative procedure to calculate the intermediate-state pressure. This feature makes it a more natural fit to the GPU thread model, which works best on algorithms without thread divergence (i.e., fewer “if” statements). Also, like all Riemann solvers in the HLL family, the HLLC is positive-definite (Batten et al. 1997), meaning that it never produces negative pressures or densities in the intermediate-state solution (Einfeldt et al. 1991). These unphysical solutions can arise using linearized solvers like the Roe, which require “fallback” techniques in failure scenarios (see, e.g., Lemaster & Stone 2009). Standard hydrodynamic tests show that the HLLC solver provides a comparable level of accuracy to the Roe and exact solvers, including problems that involve contact discontinuities (Batten et al. 1997). We qualitatively illustrate this in Figure 19, which compares well-resolved contact interfaces in Kelvin–Helmholtz simulations using the HLLC solver versus an exact solver.

Figure 18. 2D implosion test, as presented in Schneider & Robertson (2015). A colored pressure map is overlaid with 36 density contours ranging from 0.125 to 1.0. On the left are the results at $t = 2.5$ for the integration scheme used for the simulations in the paper; on the right are the results at the same time using the CTU integrator. Both simulations used a CFL number of 0.1, piecewise parabolic reconstruction, and an HLLC Riemann solver. The additional diffusion in the integration scheme presented on the left prevents as clear a jet from forming, but also allows the turbulent cloud simulations presented in this work to run without failing (which proved impossible with CTU).

Figure 19. We compare the results of a 2D Kelvin-Helmholtz simulation, starting with a discontinuous interface, on the left evolved using an HLLC Riemann solver, and on the right using an exact solver. Although the HLLC solver is an approximation, the contact discontinuity is well resolved with both methods.

APPENDIX B DUAL ENERGY IN Cholla

The conservative nature of Godunov methods enables shock capturing and allows them to correctly model shock heating, both highly desirable qualities in a hydrodynamics modeling scheme. However, in scenarios where the kinetic energy of the gas is large relative to the internal energy, the total energy formulation can lead to negative values of the internal energy. These scenarios may arise in the physical models we investigate here, where high Mach number turbulent flows interact with rapidly cooling gas. A negative value of the internal energy does not necessarily affect the dynamic behavior of the gas in such scenarios, which is dominated by the much larger kinetic energy. However, an accurate estimate of the internal energy is required to correctly calculate the gas temperature, which is needed to determine the correct radiative cooling source terms. Therefore, in such scenarios, a modification of the total energy update may improve the physical realism of the model. So-called “dual-energy” schemes that evolve both the internal energy and total energy simultaneously have proven to be a robust approach to fixing this problem (Ryu et al. 1993; Bryan et al. 1995). We employ a dual-energy formulation in Cholla that roughly follows the

descriptions in Bryan et al. (2014) and Teyssier (2015). Below, we briefly outline the 1D version of the dual-energy update.

Without the dual-energy model, Cholla stores and evolves the conserved hydrodynamic variables: density ρ , the three components of the momentum vector, $\rho \mathbf{v} = (\rho v_x, \rho v_y, \rho v_z)^T$, and the total energy, $E_{\text{tot}} = \frac{1}{2} \rho \mathbf{v}^2 + e$, where e is the internal energy. With the dual-energy model, Cholla also explicitly tracks an estimated internal energy. For the majority of steps in the hydrodynamic update, including interface reconstruction and Riemann solutions, the internal energy is treated as a passively advected scalar in the same manner as the transverse velocities. However, during the final update of the conserved variables Cholla evolves the separately tracked internal energy in cell i according to the following nonconservative equation:

$$e_i^{n+1} = e_i^n + \frac{\delta t}{\delta x} \left[(F(e)_{i-\frac{1}{2}} - F(e)_{i+\frac{1}{2}}) + \frac{1}{2} P_i (v_{i-1} - v_{i+1}) \right]. \quad (13)$$

The second term on the right-hand side of this update captures the flux of internal energy at the cell interfaces. The third term on the right-hand side encompasses the change in internal energy due to pressure forces. We use $1/2(v_{i-1} - v_{i+1})/\delta x$ as a 1D estimate of

the velocity derivative. Both the pressure P and velocities v_{i-1} and v_{i+1} used in Equation (13) are calculated at time n , making the update only first-order accurate in time. However, in cases where the conservative total energy update results in a negative estimate of the internal energy, this first-order estimate is more accurate.

Once the new total energy and internal energy have been calculated at time $n + 1$, we must determine which of the two internal energy calculations to use, the conservative estimate obtained by subtracting the kinetic energy from the total energy or the nonconservative estimate obtained with Equation (13). In addition, the internal energy must be synchronized with the total energy after the update. Here we follow the decision tree outlined in Bryan et al. (2014). At the end of the hydrodynamic update, the conservative estimate of the internal energy in cell i , $e_{\text{cons}} = E_{\text{tot}} - \frac{1}{2}\rho v^2$, is compared to the total energy in that cell. If the conservatively calculated internal energy is large enough (i.e., $e_{\text{cons}} > 0.001E_{\text{tot}}$), we use the conservative estimate for the updated internal energy: $e^{n+1} = e_{\text{cons}}$. In addition, to prevent the use of the nonconservative update in shocks, we compare the conservatively calculated internal energy to the maximum nearby total energy, $E_{\text{max}} = \max(E_{i-1}, E_i, E_{i+1})$. If $e_{\text{cons}} > 0.1E_{\text{max}}$, we again use the conservative estimate for the internal energy update, $e^{n+1} = e_{\text{cons}}$. However, if neither condition is met, we keep the nonconservative estimate for the separately tracked internal energy following Equation (13).

The last step in the dual-energy formulation is to synchronize the updated total energy with the updated internal energy. If the nonconservative estimate for e^{n+1} is used, the value of the total energy must be corrected by subtracting off the old conservatively calculated energy, and adding the new non-conservative estimate, i.e., $E_{\text{tot}}^{n+1} = E_{\text{tot}}^{n+1} - e_{\text{cons}} + e^{n+1}$.

APPENDIX C

OPTICALLY THIN RADIATIVE COOLING IN Cholla

The addition of radiative cooling requires the introduction of source terms to the right-hand side of the energy equation

$$\frac{\delta E}{\delta t} + \nabla \cdot [\mathbf{v}(E + P)] = \Gamma - \Lambda. \quad (14)$$

Here, $E = \frac{1}{2}\rho v^2 + e$ is again the total fluid energy per unit volume of the gas, ρ is the density, $\mathbf{v}$ is the velocity vector, and P is the gas pressure related to the internal energy e via an ideal gas equation of state, $P = (\gamma - 1)e$ with adiabatic index γ . The quantity Γ is a source term that accounts for heating of the gas, and Λ represents radiative losses (see, e.g., Katz et al. 1996). In the following subsections, we describe how Cholla couples the source terms to the adiabatic equations and how we calculate Γ and Λ .

C.1. Coupling of Source Terms

We implement radiative cooling in Cholla using an operator-split approach. After the hydrodynamic quantities have been updated for a given time step according to the ideal hydrodynamic equations, we add a source term to the internal energy to account for losses (or gains) as a result of radiative cooling (or heating) of the gas, such that

$$e^{n+1} = \tilde{e}^n + \dot{e}\Delta t_{\text{ad}}, \quad (15)$$

where $\dot{e}$ is the rate of change of the internal energy and Δt_{ad} is the hydrodynamic time step. Here, $\tilde{e}^n$ represents the updated internal energy after the hydrodynamic time step. The internal

energy is either calculated from the gas pressure assuming an ideal gas equation of state or tracked directly via the dual-energy formalism. Because the adiabatic time step can be large compared to the radiative cooling time $e/\dot{e}$, the rate of change in internal energy is often a nonlinear function of the temperature over the course of a single Δt_{ad} . Thus, we employ the common approach of subcycling the radiative cooling steps (e.g., Smith et al. 2008; Gray & Scannapieco 2010), calculating a new $\dot{e}$ many times over the course of a single adiabatic time step so as to limit the change in internal energy for a radiative substep, Δt_{rad} , to less than 5% of the current internal energy:

$$\frac{\Delta e}{e} < 0.05. \quad (16)$$

In practice, this update means that at the end of each hydrodynamic time step, we calculate the temperature for a given cell's number density and internal energy according to the ideal gas law

$$T = \frac{P}{nk}, \quad (17)$$

where n is the number density of the gas in cm^{-3} calculated assuming $n = \rho/(\mu m_p)$, m_p is the mass of a proton, and k is Boltzmann's constant. We have assumed a mean molecular weight of $\mu = 1$ for all calculations and have verified that this does not influence our results compared with values of $\mu = 0.6$, as described in Section 8.

We next look up the tabulated net cooling rate associated with this density and temperature using a bilinear interpolation (described below) and calculate the resulting change in temperature over the full adiabatic time step given the cooling rate, $\Delta T = \dot{T} \Delta t_{\text{ad}}$. If the change in temperature is greater than 5%, we shrink the radiative substep such that Δt_{rad} results in a temperature change of 5%. We then update the temperature with this smaller time step and repeat the process until we have synchronized with the full adiabatic time step.

C.2. Calculating Cooling and Heating Rates

We tabulate the cooling and heating rates per unit volume, Λ and Γ (measured in $\text{erg s}^{-1} \text{cm}^{-3}$), using the Cloudy code, version 13.03 (Ferland et al. 2013). For the calculations presented in this work, we assume that the gas is in photoionization equilibrium, subject to the $z = 0$ cosmic microwave background and the HM05 UV/X-ray background, as described in Hazy, the Cloudy documentation. The gas metallicity is assumed to be solar, using the GASS10 abundances in Cloudy (Grevesse et al. 2010). Examples of the absolute cooling rates for gas of several different densities are shown in Figure 20.

Our cooling and heating rates are tabulated on a grid, with points calculated at equally spaced logarithmic intervals of 0.1 between $10 \text{ K} < T < 10^9 \text{ K}$ and $10^{-6} \text{ cm}^{-3} < n < 10^6 \text{ cm}^{-3}$. To perform the interpolation necessary to calculate Λ and Γ at an arbitrary n and T , we take advantage of a novel feature of GPUs—texture memory. This special memory space allows a user to copy a 1D, 2D, or 3D array, or “texture,” onto the GPU, and then use built-in GPU functions to quickly retrieve arbitrary values from that texture using bilinear interpolation. We have tested this method against a similar CPU-based method using GSL bilinear interpolation functions and find that the GPU texture memory approach typically speeds up the radiative

Figure 20. Examples of cooling and heating rates as a function of temperature for solar-metallicity gas calculated using Cloudy. The cooling function (dashed line), heating function (dotted line), and net energy gain or loss (solid) are shown for gas of three different densities. The gas is assumed to be in photoionization equilibrium while being exposed to the $z = 0$ CMB and a Haardt & Madau UV/X-ray background. At low densities, photoionization heating leads to an equilibrium temperature of $\sim 10^4$ K, while at higher densities heating is less effective.

cooling calculation by orders of magnitude as compared to a CPU function. In fact, when implemented using the operator-splitting approach described above for simulations run with $\sim 128^3$ cells/GPU, the time spent calculating radiative cooling is completely negligible compared to the hydrodynamic calculations.

REFERENCES

- Banda-Barragán, W. E., Parkin, E. R., Federrath, C., Crocker, R. M., & Bicknell, G. V. 2016, *MNRAS*, **455**, 1309
- Batten, P., Clarke, N., Lambert, C., & Causon, D. M. 1997, *SIAM Journal on Scientific Computing*, **18**, 1553
- Bedogni, R., & Woodward, P. R. 1990, *A&A*, **231**, 481
- Bordoloi, R., Rigby, J. R., Tumlinson, J., et al. 2016, *MNRAS*, **458**, 1891
- Bouché, N., Hohensee, W., Vargas, R., et al. 2012, *MNRAS*, **426**, 801
- Brüggen, M., & Scannapieco, E. 2016, *ApJ*, **822**, 31
- Bryan, G. L., Norman, M. L., Stone, J. M., Cen, R., & Ostriker, J. P. 1995, *CoPhC*, **89**, 149
- Bryan, G. L., Norman, M. L., O’Shea, B. W., et al. 2014, *ApJS*, **211**, 19
- Chevalier, R. A., & Clegg, A. W. 1985, *Natur*, **317**, 44
- Coil, A. L., Weiner, B. J., Holz, D. E., et al. 2011, *ApJ*, **743**, 46
- Cooper, J. L., Bicknell, G. V., Sutherland, R. S., & Bland-Hawthorn, J. 2009, *ApJ*, **703**, 330
- Creasey, P., Theuns, T., Bower, R. G., & Lacey, C. G. 2011, *MNRAS*, **415**, 3706
- Dalla Vecchia, C., & Schaye, J. 2012, *MNRAS*, **426**, 140
- Davé, R., Oppenheimer, B. D., & Finlator, K. 2011, *MNRAS*, **415**, 11
- Davé, R., Thompson, R. J., & Hopkins, P. F. 2016, *MNRAS*, **462**, 3265
- Einfeldt, B., Roe, P. L., Munz, C. D., & Sjogreen, B. 1991, *JCoPh*, **92**, 273
- Faucher-Giguère, C.-A., Kereš, D., & Ma, C.-P. 2011, *MNRAS*, **417**, 2982
- Ferland, G. J., Porter, R. L., van Hoof, P. A. M., et al. 2013, *RMxAA*, **49**, 137
- Fragile, P. C., Anninos, P., Gustafson, K., & Murray, S. D. 2005, *ApJ*, **619**, 327
- Fragile, P. C., Murray, S. D., Anninos, P., & van Breugel, W. 2004, *ApJ*, **604**, 74
- Gardiner, T. A., & Stone, J. M. 2008, *JCoPh*, **227**, 4123
- Gray, W. J., & Scannapieco, E. 2010, *ApJ*, **718**, 417
- Gregori, G., Miniati, F., Ryu, D., & Jones, T. W. 2000, *ApJ*, **543**, 775
- Grevesse, N., Asplund, M., Sauval, A. J., & Scott, P. 2010, *Ap&SS*, **328**, 179
- Griffiths, R. E., Ptak, A., Feigelson, E. D., et al. 2000, *Sci*, **290**, 1325
- Harten, A., Lax, P. D., & van Leer, B. 1983, *SIAMR*, **25**, 35
- Heckman, T. M., Alexandroff, R. M., Borthakur, S., Overzier, R., & Leitherer, C. 2015, *ApJ*, **809**, 147
- Katz, N., Weinberg, D. H., & Hernquist, L. 1996, *ApJS*, **105**, 19
- Klein, R. I., McKee, C. F., & Colella, P. 1994, *ApJ*, **420**, 213
- Kornei, K. A., Shapley, A. E., Martin, C. L., et al. 2012, *ApJ*, **758**, 135
- Kritsuk, A. G., Norman, M. L., Padoan, P., & Wagner, R. 2007, *ApJ*, **665**, 416
- Lehnert, M. D., & Heckman, T. M. 1996, *ApJ*, **462**, 651
- Lemaster, M. N., & Stone, J. M. 2009, *ApJ*, **691**, 1092
- Leroy, A. K., Walter, F., Martini, P., et al. 2015, *ApJ*, **814**, 83
- Liska, R., & Wendroff, B. 2003, *SIAM Journal of Scientific Computing*, **25**, 995
- Mac Low, M.-M., McKee, C. F., Klein, R. I., Stone, J. M., & Norman, M. L. 1994, *ApJ*, **433**, 757
- Marcolini, A., Strickland, D. K., D’Ercole, A., et al. 2005, *MNRAS*, **362**, 626
- Martin, C. L. 2005, *ApJ*, **621**, 227
- McCourt, M., O’Leary, R. M., Madigan, A.-M., & Quataert, E. 2015, *MNRAS*, **449**, 2
- Melioli, C., de Gouveia dal Pino, E. M., & Raga, A. 2005, *A&A*, **443**, 495
- Mellema, G., Kurk, J. D., & Röttgering, H. J. A. 2002, *A&A*, **395**, L13
- Muratov, A. L., Kereš, D., Faucher-Giguère, C.-A., et al. 2015, *MNRAS*, **454**, 2691
- Nakamura, F., McKee, C. F., Klein, R. I., & Fisher, R. T. 2006, *ApJS*, **164**, 477
- Nestor, D. B., Johnson, B. D., Wild, V., et al. 2011, *MNRAS*, **412**, 1559
- Nittmann, J., Falle, S. A. E. G., & Gaskell, P. H. 1982, *MNRAS*, **201**, 833
- Oppenheimer, B. D., & Davé, R. 2008, *MNRAS*, **387**, 577
- Orlando, S., Peres, G., Reale, F., et al. 2005, *A&A*, **444**, 505
- Padoan, P., & Nordlund, Å. 2002, *ApJ*, **576**, 870
- Poludnenko, A. Y., Frank, A., & Blackman, E. G. 2002, *ApJ*, **576**, 832
- Rich, J. A., Dopita, M. A., Kewley, L. J., & Rupke, D. S. N. 2010, *ApJ*, **721**, 505
- Robertson, B., & Goldreich, P. 2012, *ApJL*, **750**, L31
- Rubin, K. H. R., Prochaska, J. X., Koo, D. C., et al. 2014, *ApJ*, **794**, 156
- Rubin, K. H. R., Prochaska, J. X., Ménard, B., et al. 2011, *ApJ*, **728**, 55
- Rupke, D. S., Veilleux, S., & Sanders, D. B. 2005, *ApJ*, **632**, 751
- Ryu, D., Ostriker, J. P., Kang, H., & Cen, R. 1993, *ApJ*, **414**, 1
- Scannapieco, E., & Brüggen, M. 2015, *ApJ*, **805**, 158
- Schaye, J., Crain, R. A., Bower, R. G., et al. 2015, *MNRAS*, **446**, 521
- Schneider, E. E., & Robertson, B. E. 2015, *ApJS*, **217**, 24
- Shin, M.-S., Stone, J. M., & Snyder, G. F. 2008, *ApJ*, **680**, 336
- Shopbell, P. L., & Bland-Hawthorn, J. 1998, *ApJ*, **493**, 129
- Smith, B., Sigurdsson, S., & Abel, T. 2008, *MNRAS*, **385**, 1443
- Stone, J. M., Gardiner, T. A., Teuben, P., Hawley, J. F., & Simon, J. B. 2008, *ApJS*, **178**, 137
- Stone, J. M., & Norman, M. L. 1992, *ApJL*, **390**, L17
- Strickland, D. K., & Heckman, T. M. 2007, *ApJ*, **658**, 258
- Strickland, D. K., & Heckman, T. M. 2009, *ApJ*, **697**, 2030
- Teyssier, R. 2015, *ARA&A*, **53**, 325
- Thompson, T. A., Quataert, E., Zhang, D., & Weinberg, D. H. 2016, *MNRAS*, **455**, 1830
- Toro, E. F., Spruce, M., & Speares, W. 1994, *ShWav*, **4**, 25
- Tripp, T. M., Meiring, J. D., Prochaska, J. X., et al. 2011, *Sci*, **334**, 952
- Veilleux, S., Cecil, G., & Bland-Hawthorn, J. 2005, *ARA&A*, **43**, 769
- Vogelsberger, M., Genel, S., Springel, V., et al. 2014, *Natur*, **509**, 177
- Wang, B. 1995, *ApJ*, **444**, 590
- Weiner, B. J., Coil, A. L., Prochaska, J. X., et al. 2009, *ApJ*, **692**, 187
- Westmoquette, M. S., Smith, L. J., Gallagher, J. S., & Exter, K. M. 2009, *Ap&SS*, **324**, 187
- Xu, J., & Stone, J. M. 1995, *ApJ*, **454**, 172
- Zhang, D., Thompson, T. A., Quataert, E., & Murray, N. 2015, arXiv:1507.01951