



Numerical simulations of radiatively driven dusty winds

Mark R. Krumholz¹★ and Todd A. Thompson²★

¹*Department of Astronomy and Astrophysics, University of California, Santa Cruz, CA 95064, USA*

²*Department of Astronomy and Center for Cosmology and Astro-Particle Physics, The Ohio State University, Columbus, OH 43210-1173, USA*

Accepted 2013 June 24. Received 2013 June 19; in original form 2013 February 18

ABSTRACT

Radiation pressure on dust grains may be an important mechanism in driving winds in a wide variety of astrophysical systems. However, the efficiency of the coupling between the radiation field and the dusty gas is poorly understood in environments characterized by high optical depths like those in ultraluminous infrared galaxies (ULIRGs) and massive dense star clusters. We present a series of idealized numerical experiments, performed with the radiation-hydrodynamic code *ORION*, in which we study the dynamics of such winds and quantify their properties. We find that, after wind acceleration begins, radiation Rayleigh–Taylor instability forces the gas into a configuration that reduces the rate of momentum transfer from the radiation field to the gas by a factor of ~ 10 – 100 compared to an estimate based on the optical depth at the base of the atmosphere; instead, the rate of momentum transfer from a driving radiation field of luminosity L to the gas is roughly L/c multiplied by one plus half the optical depth evaluated using the photospheric temperature, which is far smaller than the optical depth one would obtain using the interior temperature. When we apply our results to conditions appropriate to ULIRGs and star clusters, we find that the asymptotic wind momentum flux from such objects should not significantly exceed that carried by the direct radiation field, L/c . This result constrains the expected mass-loss rates from systems that exceed the Eddington limit to be of the order of the so-called ‘single-scattering’ limit, and not significantly higher. We present an approximate fitting formula for the rate of momentum transfer from radiation to dusty gas through which it passes, which is suitable for implementation in sub-grid models of galaxy formation. Finally, we provide a first map of the column density distribution of gas in a radiatively driven wind as a function of velocity, and velocity dispersion.

Key words: hydrodynamics – instabilities – ISM: jets and outflows – radiative transfer – galaxies: ISM – galaxies: star clusters.

1 INTRODUCTION

Dusty winds are ubiquitous in astrophysics: they are driven on scales ranging from single stars (e.g. Habing 1996) to star clusters (e.g. Lopez et al. 2011) to entire galaxies (e.g. Veilleux, Cecil & Bland-Hawthorn 2005). The driving mechanisms of these winds are diverse and in some cases uncertain, but one possible mechanism for many of them is the force exerted by radiation interacting with dusty matter. Photons moving through dusty gas can be scattered or absorbed by dust grains, transferring some of their momentum. The grains, in turn, transfer this momentum to the gas either through hydrodynamic drag or via magnetic fields, possibly giving rise to a wind. Radiation pressure on dust has been suggested as an important feedback mechanism in regulating star formation on the scales of galaxies (Scoville 2003; Thompson, Quataert &

Murray 2005; Andrews & Thompson 2011) and individual massive stars clusters (O’dell, York & Henize 1967; Scoville et al. 2001; Krumholz & Matzner 2009; Fall, Krumholz & Matzner 2010; Krumholz & Dekel 2010; Murray, Quataert & Thompson 2010), for driving dusty fountain flows in normal spirals (Chiao & Wickramasinghe 1972; Elmegreen 1983; Ferrara 1993) and for driving galaxy-wide superwinds (Murray, Quataert & Thompson 2005; Murray, Ménard & Thompson 2011; Hopkins, Quataert & Murray 2012).

However, assessing these claims has been difficult due to limited understanding of the radiation–matter interaction that drives the flow. For optically thin flows the problem is relatively simple, since the state of the radiation field is decoupled from the gas. For optically thick media, however, the problem is significantly more difficult, because the gas is capable of reshaping the radiation field. This can lead to a number of complex phenomena, such as photon bubbles (Blaes & Socrates 2003) and radiation Rayleigh–Taylor instability (RRTI; Krumholz et al. 2009; Jacquet & Krumholz 2011; Jiang, Davis & Stone 2013). The existence of these behaviours makes it

★E-mail: krumholz@ucolick.org (MRK); thompson@astronomy.ohio-state.edu (TAT)

non-trivial to calculate from first principles whether an object with a given set of properties can produce a radiatively driven dusty wind at all, and, if it does, what properties that wind is likely to possess. This problem has thus far prevented definitive identification of the driving mechanisms for winds observed in a variety of systems (e.g. Faucher-Giguère & Quataert 2012; Newman et al. 2012).

In Krumholz & Thompson (2012, hereafter Paper I), we addressed the first part of this problem: under what circumstances do we expect an object to launch a radiatively driven dusty wind? We developed an idealized model system that allowed us to extract the important dimensionless numbers governing wind launching, and we then conducted numerical experiments with the radiation-hydrodynamics code *ORION* to explore the non-linear behaviour of the system. The major results of Paper I are that the behaviour of gravitationally confined, dusty columns of matter subjected to radiative fluxes are governed primarily by two characteristic values: the dust optical depth and the Eddington ratio, both computed using the opacities that apply at the surface (i.e. the photosphere) of the dusty gas. For high optical depths and surface Eddington ratios close to but below unity, which may describe many galaxies and star clusters, we showed that radiation passing through the gas drives statistically steady turbulence with average Eddington ratio of unity, but not a wind.

This result, however, does not answer the question of what happens if a wind *is* launched – either because the surface Eddington ratio exceeds unity, or because some other mechanism is able to eject matter, by itself or in conjunction with radiation forces. This question is the main focus of our paper. We seek to determine at what rate the matter in a radiatively driven wind is able to extract momentum from the radiation field, and how this depends on properties such as the strength of the radiative driving and the optical depth of the matter. In addition to illuminating the physics of the winds, we also derive a rough fitting formula that can be used in numerical simulations that do not include radiation hydrodynamics, and instead treat radiative driving using sub-grid semi-analytic models (e.g. Oppenheimer & Davé 2006; Hopkins, Quataert & Murray 2011). Jiang et al. (2013) performed preliminary work on this problem in the context of winds where the dominant opacity source is Thompson scattering from free electrons, and concluded that RRTI would limit the wind mass and momentum flux. We seek to investigate whether the same is true for dusty winds, and to extend their results by drawing quantitative rather than qualitative conclusions about how the wind momentum depends on the properties of the system.

The remainder of this paper is as follows. In Section 2, we briefly review the basic equations and model system developed in Paper I, and consider how to extend them to the case of a dusty wind. In Section 3, we describe our numerical simulations, and in Section 4 we analyse the results they produce. In Section 5, we discuss the implications of our results and provide some caveats, and Section 6 summarizes our conclusions.

2 MODEL SYSTEM

2.1 Governing equations and model system

As in Paper I, we treat a section of a galactic disc or a young star cluster as an idealized model system consisting of a slab of gas with total surface density Σ filling the domain $z > 0$. A vertical radiation flux $F = F_0 \hat{z}$ enters the domain of interest at $z = 0$, and there are no radiation sources at $z > 0$ other than the thermal emission of the gas.

The slab of material is confined by a constant vertical gravitational force per unit mass $-g\hat{z}$; we neglect the self-gravity of the gas.

Since we are interested in cases where the gas layer is optically thick, we describe this system using the two-temperature flux-limited diffusion (2TFLD) approximation, in which we assume that the radiation spectrum is locally a Planck function at every point, but we do not require that the temperature T_r describing this Planck function be identical to the gas temperature T_g . We discuss the limitations of the 2TFLD method in Section 5.6. In this approximation, interaction of radiation and matter is governed by the Planck and Rosseland mean opacities κ_P and κ_R .

The equations governing this system are (Krumholz et al. 2007)

$$\frac{\partial}{\partial t} \rho = -\nabla \cdot (\rho \mathbf{v}) \quad (1)$$

$$\frac{\partial}{\partial t} (\rho \mathbf{v}) = -\nabla \cdot (\rho \mathbf{v} \mathbf{v}) - \nabla P - \lambda \nabla E - \rho g \hat{z} \quad (2)$$

$$\begin{aligned} \frac{\partial}{\partial t} (\rho e) = & -\nabla \cdot [(\rho e + P) \mathbf{v}] - \kappa_P \rho (4\pi B - cE) \\ & + \lambda \left(2 \frac{\kappa_P}{\kappa_R} - 1 \right) \mathbf{v} \cdot \nabla E - \frac{3 - R_2}{2} \kappa_P \rho \frac{v^2}{c} E \\ & - \rho g v_z \end{aligned} \quad (3)$$

$$\begin{aligned} \frac{\partial}{\partial t} E = & \nabla \cdot \left(\frac{c\lambda}{\kappa_R \rho} \nabla E \right) + \kappa_P \rho (4\pi B - cE) \\ & - \lambda \left(2 \frac{\kappa_P}{\kappa_R} - 1 \right) \mathbf{v} \cdot \nabla E \\ & + \frac{3 - R_2}{2} \kappa_P \rho \frac{v^2}{c} E - \nabla \cdot \left(\frac{3 - R_2}{2} \mathbf{v} E \right), \end{aligned} \quad (4)$$

where \mathbf{v} is the gas velocity, $P = \rho k_B T_g / \mu m_H$ is the gas pressure, μ is the mean molecular weight in hydrogen masses, $e = P/[(\gamma - 1)\rho] + v^2/2$ is the gas specific energy, μ is the mean mass per gas particle in units of the hydrogen mass m_H , γ is the gas ratio of specific heats, $B = caT_g^4/4\pi$ is the frequency-integrated Planck function, $E = aT_r^4$ is the radiation energy density, F is the radiation flux, λ is the flux limiter and R_2 is the Eddington factor. In the 2TFLD approximation, adopting the flux-limiter of Levermore & Pomraning (1981) and Levermore (1984), the radiation quantities are related by

$$\mathbf{F} = -\frac{c\lambda}{\kappa_R \rho} \nabla E \quad (5)$$

$$\lambda = \frac{1}{R} \left(\coth R - \frac{1}{R} \right) \quad (6)$$

$$R = \frac{|\nabla E|}{\kappa_R \rho E} \quad (7)$$

$$R_2 = \lambda + \lambda^2 R. \quad (8)$$

In this work, we adopt opacities

$$(\kappa_R, \kappa_P) = (10^{-3/2}, 10^{-1}) \left(\frac{T_g}{10 \text{ K}} \right)^2 \text{ cm}^2 \text{ g}^{-1}, \quad (9)$$

an approximation to the behaviour of dust opacity at temperatures below ~ 150 K (Semenov et al. 2003). For simplicity, as in Paper I,

we continue to use this power-law approximation even at higher temperatures, and note that this will generally lead us to overestimate the opacity and thus the strength of matter–radiation coupling.

We pause here to note that a few subtle physical assumptions built into the above equations. First, we have not distinguished between the temperatures of dust and gas, and thereby have implicitly assumed that they are the same. Secondly, the opacity should depend on the radiation spectrum and thus on the effective radiation temperature; in writing down the opacity law, equation (9), in terms of the gas temperature, we have implicitly assumed that the gas temperature and the radiation temperature are also tightly coupled, at least in the regime where the opacity is large enough to matter. Thirdly, we have also assumed that dust and gas are dynamically coupled. We show in Appendix A that these assumptions are well justified in the regime we consider.

In Paper I, we show that these equations are fully characterized by four non-dimensional parameters

$$f_{E,*} = \frac{\kappa_{R,*} F_0}{g c} \quad (10)$$

$$\beta_s = \frac{c_{s,*}}{c} = \frac{1}{c} \sqrt{\frac{k_B}{\mu}} \left(\frac{g}{a \kappa_{R,*}} f_{E,*} \right)^{1/8} \quad (11)$$

$$\tau_* = \Sigma \kappa_{R,*} \quad (12)$$

$$k_0 = \frac{\kappa_{P,*}}{\kappa_{R,*}}. \quad (13)$$

Here $\kappa_{R,*} = \kappa_R(\rho_*, T_*)$ and similarly for $\kappa_{P,*}$, and the starred quantities in turn are defined by

$$T_* = \left(\frac{F_0}{c a} \right)^{1/4} \quad \rho_* = \frac{\Sigma g \mu m_H}{k_B T_*}. \quad (14)$$

We may think of T_* as the characteristic temperature at the photosphere of the dusty gas, where the radiation escapes to infinity. The natural units of velocity, length and time for the problem are

$$c_{s,*} = \sqrt{\frac{k_B T_*}{\mu m_H}} \quad h_* = \frac{c_{s,*}^2}{g} \quad t_* = \frac{h_*}{c_{s,*}}. \quad (15)$$

In real astrophysical systems, β_s is always very small for any non-relativistic flow, and k_0 is always of the order of unity and probably varies little from one galaxy to another. Thus, in practice the quantities $f_{E,*}$ and τ_* determine the behaviour of the system.

2.2 Dimensionless numbers for winds

If gas has been launched into a wind, it has obviously overcome its initial gravitational confinement. As discussed in the Introduction, in Paper I we show that this occurs only for $f_{E,*} > 1$, i.e. only if the Eddington ratio at the dust photosphere exceeds unity. In principle, one wishes to determine the properties of winds launched at a range of $f_{E,*} > 1$. However, we focus on the asymptotic limit $f_{E,*} \rightarrow \infty$, corresponding to a freely accelerating wind with negligible gravitational confinement. Our reasons for doing so are threefold. First, this reduces the parameter space we must explore. Secondly, it is very likely that the case where there is no gravitational confinement will produce the largest possible wind momentum flux, and so it can serve as a useful upper limit. We will see below that even this upper limit is quite restrictive on the possible momentum of the wind. Thirdly, in the case of radiating optically thick discs (e.g. starburst galaxy disc or the discs around quasars), $f_{E,*}$ rises with

height above the disc, so winds at large distances will have larger $f_{E,*}$ values (Zhang & Thompson 2012).

For $f_{E,*} \rightarrow \infty$, or equivalently $g \rightarrow 0$, the quantities ρ_* , h_* and t_* cease to be well defined. It is therefore helpful to define alternative natural units in the freely accelerating wind case. The sound speed $c_{s,*}$ remains the natural unit of velocity, and to define a unit of time it is helpful to ask how long it would take the momentum carried by the direct radiation field to accelerate matter from rest to this speed. The momentum flux per unit mass of the injected radiation field is

$$f_{\text{rad,dir}} = \frac{F_0}{\Sigma c}, \quad (16)$$

and so we define the acceleration time as

$$t_a = \frac{c_{s,*}}{f_{\text{rad,dir}}} = \frac{\tau_* c_{s,*}}{\kappa_{R,*} F_0 / c} = \frac{\tau_*}{f_{E,*}} t_*. \quad (17)$$

If radiative trapping is significant, we expect the matter to increase its velocity by $c_{s,*}$ in a time significantly shorter than t_a . Finally, we can define characteristic length and density scales from the combination of t_a and $c_{s,*}$. These are

$$h_a = c_{s,*} t_a = \frac{\tau_*}{f_{E,*}} h_* \quad \rho_a = \frac{\Sigma}{h_a} = \frac{f_{E,*}}{\tau_*} \rho_*. \quad (18)$$

We report all results in this paper in units of ρ_a , h_a and t_a .

3 NUMERICAL SIMULATIONS

We solve equations (1)–(4) with $g = 0$ using the radiation-hydrodynamics code ORION. Our simulations are two dimensional, and take place in the (x, z) plane; a flux F of radiation is injected at the bottom of the computational box, $z = 0$. The boundary conditions are periodic in the x direction, impermeable at the lower z boundary and open at the upper z boundary. More details on the boundary conditions are given in Paper I. All other parameters of the simulations are also the same as in Paper I, except that we set the external gravitational field $g = 0$, so that $f_{E,*} \rightarrow \infty$.

To study the behaviour of winds in the limit $f_{E,*} \rightarrow \infty$ as a function of τ_* , we run four simulations, which we denote T3, T10LR, T10HR and T30; the first of these have $\tau_* = 3$, the second two have $\tau_* = 10$, and the third has $\tau_* = 30$. The two $\tau_* = 10$ runs are identical except in their resolutions and the sizes of the computational domains. We perform both runs as a check on the resolution dependence of our results. We summarize the physical and numerical parameters of the simulations in Tables 1 and 2.

In Paper I, we show that initially laminar, gravitationally confined gas slabs subjected to radiation fluxes such that $f_{E,*}$ is above a certain critical value, but less than unity, will develop radiation-driven turbulence. This occurs in a time $\lesssim 100 t_*$. (Note that values of t_* for our runs are given in Table 1, and are the same for every run; $100 t_* \approx 100$ kyr.) Since this is short compared to most astrophysical time-scales of relevance, it is reasonable to assume that gas being launched in a wind will be in a fully turbulent state. We therefore

Table 1. Simulation physical parameters.

Name	τ_*	Σ (g cm ⁻²)	t_* (kyr)	t_a (kyr)	$h_a/10^{-2}$ (pc)	$\rho_a/10^{-16}$ (g cm ⁻³)
T3	3	1.4	1.1	6.9	0.38	1.2
T10	10	4.6	1.1	23	1.3	1.2
T30	30	14	1.1	69	3.8	1.2

Notes. T10 describes both runs T10LR and T10HR, which have identical physical parameters but different resolutions and box size. All models have $T_* = 82$ K, $c_{s,*} = 0.54$ km s⁻¹.

Table 2. Simulation numerical parameters.

Name	IC	$N_x \times N_z$	$L_x \times L_z$	Δx	t_{run}
T3	T3F0.5	1024×16384	85.3×1365	0.083	28.0
T10LR	T10F0.5	512×32768	25.6×407	0.05	21.9
T10HR	T10F0.5	1024×16384	25.6×407	0.025	10.6
T30	–	1024×16384	27.3×437	0.027	3.6

Notes. The initial condition (IC) column gives the name of the corresponding run in Paper I used to produce the initial condition. $N_x \times N_z$ is the size of the computational domain in cells, $L_x \times L_z$ is the size of the computational domain, Δx is the size of a computational cell and t_{run} is the duration for which we run the simulation. All quantities are given in units of h_a and t_a . Note that all models were run during the phase where gravity was turned on with $\Delta x/h_* = 0.5$. Finally, for a description of the initial conditions for run T30, see the Appendix.

do not start our simulations with laminar gas layers. Instead, we use as initial conditions the end states of the simulations from Paper I, as summarized in Table 2. We modify these conditions only in that we place the gas in a computational box that is larger in the vertical direction, in order to accommodate vertical expansion of the gas layer once it is no longer gravitationally confined. We initialize computational cells that are outside the computational domain of the simulations of Paper I by giving them densities equal to the background density and temperature from Paper I, $10^{-10}\rho_*$ and T_* , respectively. In run T10LR, we also down-sample the resolution by a factor of 2. The exception to the above statements is run T30, for which we do not have a corresponding run from Paper I. We describe how we generate its initial conditions in the Appendix.

As the simulations proceed, when necessary we shift all velocities in the computational domain by a constant offset in order to bring the centre of mass velocity of the gas back to zero. Our method is simple: we have added an option to the ORION code that, upon restart from a checkpoint, calculates the centre of mass velocity of the computational domain in the z direction, then subtracts the corresponding velocity from all computational cells, altering the momenta and total energies appropriately. The calculation then restarts from the modified data. We apply this option whenever a visual inspection of the data indicates that the bulk of the mass is well away from the bottom boundary of the computational domain. This enables us to continue the simulations longer without the gas reaching the top of the computational box. We shift the velocities in this manner only when the vast majority of the gas is well away from the bottom of the computational box, so that there are no significant artificial forces exerted by the bottom of the computational box. In the analysis below, we remove these offsets and present the results as if the entire simulation had simply taken place in a larger box. In principle, we could shift the positions as well, but this is less convenient computationally, since it would require translating values from one cell to another, and filling in values of density, momentum and energy in the new cells added to the computational domain by any shifts.

4 RESULTS

4.1 Qualitative behaviour

Figs 1–8 shows a series of snapshots of the simulation density fields. As the plots show, the initial state that results from the RRTI (Jacquet & Krumholz 2011) acting on a gas confined by gravity, consists of a relatively horizontal, turbulent layer. In the absence of gravitational confinement, the radiation force rapidly drives the gas into a




Figure 1. Density distribution as a function of time in run T3. Snapshots are shown at intervals of $5t_a$, starting from $t = 0$, as indicated at the top of each panel. White bars indicate a region around the vertical centre of mass; we show a zoom-in of this region in Fig. 5. Note that the vertical extent shown does not necessarily match the size of the computational box given in Table 2, because we have compensated for the effects of our periodic shifts the centre of mass velocity of the entire computational domain in some runs – see Section 3 for details.




Figure 2. Same as Fig. 1, but for run T10LR. The zoomed-in region is shown in Fig. 6.

predominantly vertical, filamentary structure. In between the filaments of dense gas there are low-density channels. As the material is accelerated upward by the radiation field, the gas becomes more elongated and spread over a progressively larger vertical extent. We are eventually forced to halt our simulations primarily because the vertical extent of the gas becomes comparable to the vertical size of our computational domain.




Figure 3. Same as Fig. 1, but for run T10HR. The zoomed-in region is shown in Fig. 7.

Fig. 9 shows an example of the distribution of density, temperature, velocity and radiation flux in one of the runs once the channel structure has developed. The channels are characterized by several features. First, within them the gas is travelling at extremely large velocities relative the dense gas in the filaments. At the snapshot shown, the velocity difference approaches many tens of $c_{s,*}$. As a result of this velocity difference, the edges of the channels appear to




Figure 4. Same as Fig. 1, but for run T30. The zoomed-in region is shown in Fig. 8. Note that the first panel is not in fact empty – the gas at time 0 is simply compressed into an extremely thin layer whose width, on the scale plotted, is less than a single pixel.

be scalloped by Kelvin–Helmholtz instabilities. Secondly, because of their lower optical depths, the channels carry the great majority of the radiative flux. The flux within the channels approaches $10F_0$, while inside the filaments the flux is $\ll F_0$. Thus, the matter effectively collimates the radiation field, inducing a strong anticor-




Figure 5. Same as Fig. 1, but the panels show a zoom-in on a small region around the vertical centre of mass. The zoomed region is indicated by the white bars in Fig. 1.

relation between density and radiative flux. This anticorrelation is the main signature of the RRTI.

4.2 Radiative trapping

The development of vertical filamentary structure and the resulting collimation of the radiation field has profound effects on its ability to trap the radiation field and extract momentum from it. To quantify the rate at which the gas takes up momentum from the radiation field, it is helpful to examine the z component of the momentum equation including radiation and gravitational forces; this is

$$\frac{d}{dt}(\rho v_z) = -\nabla \cdot (\rho \mathbf{v} v_z) - \frac{dP}{dz} + \kappa_R \rho \frac{F_z}{c} - \rho g, \quad (19)$$

where F_z is the z -component of the radiation flux and the use of κ_R in the equation implicitly equates the flux-mean and Rosseland-mean opacities, as is appropriate in the optically thick regime. In the flux-limited diffusion approximation this equation is equivalent to equation (2), as shown by Krumholz et al. (2007), but the analysis is more transparent when the equation is written in the form above. If we integrate this equation over the entire computational domain, and ignore the small terms that arise from forces and fluxes across the top and bottom boundaries of the computational domain, the first two terms on the right-hand side vanish and we are left with

$$\frac{d}{dt} \langle \rho v_z \rangle = \left\langle \kappa_R \rho \frac{F_z}{c} \right\rangle - \langle \rho \rangle g, \quad (20)$$

where for any quantity q we defined the volume average by

$$\langle q \rangle = \frac{1}{L_x L_z} \int_{-L_x/2}^{L_x/2} \int_0^{L_z} q \, dz \, dx. \quad (21)$$

Dividing both sides by $\langle \rho \rangle$ and noting that the mass in the computational domain is very close to constant (since we are careful to ensure there is no significant mass-loss from the top of the computational box) gives

$$\frac{d\bar{v}_z}{dt} = \frac{1}{c} \frac{\langle \kappa_R \rho F_z \rangle}{\langle \rho \rangle} - g, \quad (22)$$




Figure 6. Same as Fig. 5, but for run T10LR.

where $\bar{v}_z = \langle \rho v_z \rangle / \langle \rho \rangle$ is the mass-weighted mean z velocity of the gas. We use the first term on the right-hand side to define the mean radiation force per unit mass,

$$f_{\text{rad}} = \frac{1}{c} \frac{\langle \kappa_R \rho F_z \rangle}{\langle \rho \rangle}. \quad (23)$$

Equivalently, we may think of this term as describing the mass-weighted mean radiation force. Based on our observation that density and flux are strongly anticorrelated, we expect that to be much

less than the volume-weighted mean radiation force $\langle \kappa_R F_0 \rangle / c$ would be. The second term on the right-hand side is simply the gravitational force per unit mass.

At this point, it is useful to rewrite the equation by multiplying through by a factor of $t_a / c_{s,*}$ to non-dimensionalize. Doing so gives

$$\frac{t_a}{c_{s,*}} \frac{d\bar{v}_z}{dt} = \frac{f_{\text{rad}}}{f_{\text{rad,dir}}} - \frac{\tau_*}{f_{E,*}}. \quad (24)$$

Following Krumholz & Matzner (2009) and Paper I, we define the trapping factor by

$$1 + f_{\text{trap}} = \frac{f_{\text{rad}}}{f_{\text{rad,dir}}}. \quad (25)$$

Physically, the trapping factor is simply the factor by which the radiation force is amplified by trapping of the radiation field by the gas. The quantity $1 + f_{\text{trap}}$ is equivalent to the amplification factor τ_{IR} defined by Thompson et al. (2005), although we refer to it as f_{trap} here because, as we will see, its relationship to optical depth is not trivial. Using equation (24) to rewrite the equation above, we obtain

$$f_{\text{trap}} = \frac{t_a}{c_{s,*}} \frac{d\bar{v}_z}{dt} + \frac{\tau_*}{f_{E,*}} - 1. \quad (26)$$

The quantity on the right-hand side is directly measurable from our simulations (and $f_{E,*} = \infty$ in the absence of gravity, so the term $\tau_*/f_{E,*} = 0$), so our simulations provide us with a direct measurement of f_{trap} as a function of time. However, we must make one important modification to equation (26), which comes from a limitation of our numerical method. Because we are using flux-limited diffusion, we do not properly capture the interaction of the gas with the direct, beamed radiation field produced by stars. Instead, we are treating the radiation field only after this first absorption. Since the final -1 represents the contribution from this direct radiation field, we do not subtract it off when computing f_{trap} from the simulations. This is likely conservative, since our method does capture some of the effects of the first absorption, in which case the results we obtain should be upper limits on f_{trap} . However, we cannot completely rule out the possibility that inclusion of the direct radiation force would somehow change the structure of the gas and indirectly increase the trapping of the reradiated field.

Fig. 10 shows the gas mean velocity and trapping factor as a function of time in each of our simulations. For constant f_{trap} , the gas velocity should increase linearly with time. Instead, we see that the velocity increase is steep at first and then becomes much shallower, and this is reflected in the plots of f_{trap} , which are large at first and then decline over a few t_a . The initially high values are easy to understand given our starting conditions. When the gas is confined by gravity and there is no wind, both the time-averaged value of \bar{v}_z and its rate of change must be zero. Consulting equation (26), this requires that $f_{\text{trap}} = \tau_*/f_{E,*}$, again omitting the -1 because our simulation does not properly model the direct radiation force. The initial value of f_{trap} we measure is indeed close to this, though the match is not exact because $d\bar{v}_z/dt$ is not precisely zero at all times in the gravity-confined state; instead, it oscillates about zero.

Once the gravitational confinement is removed, however, the gas morphology changes from predominantly horizontal to predominantly vertical, and f_{trap} drops. At late times f_{trap} oscillates up and down about a value well below the initial one. Upward and downward oscillations of f_{trap} correspond to variations in the gas morphology. At times, for example at times $t/t_a = 5$ and 10 in run T10HR (see Fig. 7), the filaments formed by the radiation




Figure 7. Same as Fig. 5, but for run T10HR.




Figure 8. Same as Fig. 5, but for run T30. As with Fig. 4, note that the first panel is not empty, but it appears so because the gas at time 0 is simply compressed into an extremely thin layer whose width, on the scale plotted, is less than a single pixel.

are fairly coherent leave fairly large vertical channels unobstructed, and at these times f_{trap} is low. At other times, such as $t/t_a = 7.5$ in run T10HR, the filaments are more fragmented and cover more or the domain horizontally, giving rise to larger values of f_{trap} .

Finally, comparing runs T10LR and T10HR suggests that are results are relatively well-converged. Since the gas is turbulent for RRTI in its fully developed state, the flow is chaotic and we do not expect either morphologies or exact values of f_{trap} as a function of time to be resolution-independent. However, examining the results in Fig. 10, we see that values of f_{trap} versus time produced in the two runs are qualitatively similar, and that quantitatively their means are well within the level of variance in f_{trap} we measure in each run as the flow varies chaotically. This suggests that our values of f_{trap} are converged.

We summarize our results for f_{trap} in Table 3, where we report the minimum, mean and maximum values of f_{trap} we measure in each of our simulations once the initial transient phase ends. For comparison, we also report $\tau_*/f_{E,*} \approx f_{\text{trap, grav}}$, the mean value of f_{trap} in the initial, turbulent, gravity-confined state and $\kappa(T_{\text{mp}})\Sigma \approx \tau_{\text{IR}}$, the optical depth computed by multiplying the column density by the opacity evaluated using the mid-plane temperature T_{mp} . The latter has been used as an approximate value for f_{trap} by a number of authors, as we discuss in more detail below. Clearly none of these values are equal; instead $f_{\text{trap}} \ll f_{\text{trap, grav}} \ll \tau_{\text{IR}}$.




Figure 9. A section from run T10LR at $t = 12.4t_a$ showing the density and velocity distribution (colours and vectors in the left-hand panel), and the gas temperature and radiative flux distribution (colours and vectors in the right-hand panel). The region shown is centred in the vertical centre of mass of the gas at this time, $z_{\text{cm}} = 298h_a$, and the velocities shown are relative to the vertical centre of mass velocity of the section shown, $v_{z,\text{cm}} = 77.0c_{s,*}$.




Figure 10. Mean gas velocity versus time (top row) and trapping factor versus time (bottom row) for runs T3 (first column), T10 (second column) and T30 (third column). In all panels, solid lines show the results of the simulations; for T10, blue shows T10LR and green T10HR. In the upper panels, black dashed lines show the change in velocity versus time that would be expected for trapping factors $f_{\text{trap}} = 1, 10, 100$ and 1000 , from shallowest line to steepest. In the bottom panel, dashed black lines show $f_{\text{trap}} = 1$ and $f_{\text{trap}} = \tau_*/f_{E,*}$, with the value of $f_{E,*}$ computed before gravity is turned off – i.e. the value of $f_{E,*}$ that was used in the simulation from Paper I from which we take our initial conditions. In the absence of gravity, as is the case for the simulations shown here, $f_{E,*} = \infty$.

Table 3. Simulation results.

Name	t_{\min}	$\text{Min}(f_{\text{trap}})$	$\langle f_{\text{trap}} \rangle$	$\text{Max}(f_{\text{trap}})$	$\tau_*/f_E, *$ $\approx f_{\text{trap, grav}}$	$\kappa(T_{\text{mp}})\Sigma$ $\approx \tau_{\text{IR}}$	$\langle \tau \rangle$
T3	5	1.0	1.6 ± 0.4	2.5	6	15	6.2 ± 0.6
T10LR	3	2.6	6.0 ± 1.3	9.9	20	120	32 ± 3
T10HR	3	2.2	4.9 ± 2.4	8.9	20	120	30 ± 5
T30	1	5.3	12.3 ± 5.7	23.3	600	2000	132 ± 11

Notes. For each run, $\text{min}(f_{\text{trap}})$, $\langle f_{\text{trap}} \rangle$ and $\text{max}(f_{\text{trap}})$ give the minimum, time-averaged and maximum values of f_{trap} that occur in the simulation after time t_{\min} . The error bars given on $\langle f_{\text{trap}} \rangle$ represent the 1σ range measured from the simulations at times $> t_{\min}$. For comparison, $\tau_*/f_E, * \approx f_{\text{trap, grav}}$ gives the time-averaged value of f_{trap} in the steady-state gravitationally confined configuration from which we start, while $\kappa(T_{\text{mp}})\Sigma \approx \tau_{\text{IR}}$ is the average optical depth at the start of the calculation, computed using the mass-weighted mean mid-plane temperature as in Paper I. The quantity $\langle \tau \rangle$ is the mass-weighted mean optical depth of the computational domain computed using the temperatures from the simulations; see equation (37). As with $\langle f_{\text{trap}} \rangle$, the value shown is the time average at times $> t_{\min}$, and the error bars give the 1σ range over this time.

4.3 Wind velocity distribution

Fig. 10 and Table 3 describe the mean velocity and mass-averaged momentum transfer from radiation to gas. However, it is also interesting to look at the distribution of matter velocities. In Fig. 11, we show mass-weighted velocity probability distribution functions (PDFs) for each of our runs at several times. In the initial con-




Figure 11. Velocity distribution functions for each of our simulations at times $t/t_a = 0, 10$ and 28 (for runs T3, T10LR and T10HR) and $t/t_a = 0$ and 3 (for run T30), as indicated in the legend. In each panel, the histogram shows the fraction of the mass in the simulation that falls into a given bin of z velocity at the indicated time. Vertical dashed lines indicate the mass-weighted mean velocity \bar{v}_z at that time.

dition, the distribution of z velocities is, as one might expect for a roughly hydrostatic atmosphere, symmetric about $v_z = 0$. The width is fairly small, reflecting the relatively low Mach numbers we obtained for turbulent atmospheres in Paper I. At times $\gg t_a$, not only is the mean velocity \bar{v}_z larger, the spread of velocities is larger as well. At late times the PDF for all the runs is slightly asymmetric, with the majority of the mass residing at velocities slightly below the mean, and a tail extending well above the mean.

The division between high- and low-velocity material corresponds to the division between material in the opaque filaments and material in or at the edges of the radiation-dominated channels, as illustrated in Fig. 12. To construct this figure, along every vertical line of sight we measure the column density

$$\Sigma(x) = \int_0^{L_z} \rho(x, z) dz. \quad (27)$$

We then assign every cell a column density $\Sigma(x)$ corresponding to the value at its x position, and construct the two-dimensional PDF of $\Sigma(x)$ and v_z . From the 2D PDF, we see the same asymmetry as in Fig. 11, where the velocity distribution extends further from the mean in the positive direction than the negative direction. In the 2D PDF, it is clear that the high-velocity material consists preferentially of gas with low $\Sigma(x)$. The correlation is relatively weak, and the overall range in $\Sigma(x)$ is relatively small, because the filaments are not perfectly vertical. Thus, most of the time a given vertical line of sight will intersect both dense filaments and low-density channels, rather than looking straight down the barrel of a channel. Nonetheless, this column density–velocity anticorrelation represents a possible observable signature of radiation pressure-driven dusty winds.

The overall width of the velocity distribution, including both low- and high-speed components, is of the order of $\sim 20c_{s,*}$ in all the runs; the dispersion of horizontal velocities is substantially smaller. The dispersion does not appear to increase substantially over the time interval shown, and thus $\sim 20c_{s,*}$ is likely the steady-state value, at least over the range of τ_* values that we have explored. This corresponds to a one-dimensional Mach number in the vertical direction of the order of 20 – not exactly 20, since much of the gas is somewhat warmer than T_* and thus has a sound speed greater than $c_{s,*}$. For our fiducial choice of dimensional scaling ($c_{s,*} = 0.54 \text{ km s}^{-1}$), this gives a physical velocity dispersion of roughly 10 km s^{-1} in the wind, compared to bulk velocities




Figure 12. Two-dimensional velocity-column density distributions in the simulations at the latest time slices shown in Fig. 11 ($t/t_a = 28$ for runs T3 and T10LR, $t/t_a = 10$ for run T10HR and $t/t_a = 3$ for run T30). Each pixel shows the logarithm of the fraction of the simulation mass in the indicated bin of v_z and $\Sigma(x)$, normalized so that the most massive bin has a value of unity. Note that both the x and y axes are offset such that material at the mass-weighted mean column density and velocity would appear at the position $(0, 0)$.

of $\sim 100 \text{ km s}^{-1}$ at the same time. While this Mach number and velocity dispersion are larger than we found in Paper I for radiation Rayleigh–Taylor (RRT)–unstable atmospheres that do not drive a wind, they are still close to an order of magnitude smaller than the values observed in the most vigorously star-forming ultraluminous infrared galaxies (ULIRGs) (e.g. Downes & Solomon 1998). Finally, we note that the results for runs T10LR and T10HR are qualitatively similar at equal times, suggesting at least rough convergence.

5 DISCUSSION

5.1 Fitting formulae for radiation trapping

By combining the results of this paper with those of Paper I, we are now in a position to provide a fitting formula for the value of f_{trap} in optically thick radiation pressure-driven atmospheres and winds. Such a formula is useful in simulations or analytic calculations that seek to include radiation pressure effects, but that do not properly capture the radiation-hydrodynamic behaviour associated with the RRTI. Examples include one-dimensional models (e.g. Krumholz & Matzner 2009; Murray et al. 2010), models that include radiation pressure only via a subgrid prescription rather than by solving the equation of radiative transfer (e.g. Hopkins et al. 2011; Agertz et al. 2012), models that solve the transfer equation only for the direct and not the dust-reprocessed radiation field (and therefore implicitly set $f_{\text{trap}} = 0$; e.g. Peters et al. 2010; Wise et al. 2012; Kim et al. 2013a,b) and models that solve the transfer equation in one dimension under an assumption of spherical symmetry and therefore miss RRT effects (e.g. Novak, Ostriker & Ciotti 2012).

The value of f_{trap} for a radiation pressure-dominated wind or atmosphere is a function of the two main dimensionless parameters for the problem, $f_{E,*}$ and τ_* . We have sampled this parameter space quite coarsely, but we can nonetheless provide a rough fit that captures the results of our simulations, and which is an improvement over simple prescriptions. In Paper I, we explored the regime $f_{E,*} < 1$, and found that for a given τ_* there exists a critical $f_{E,*}$ above which instability sets in. At values of $f_{E,*}$ below this value the atmosphere is supported predominantly by gas pressure over most of its height, and radiation pressure is dynamically unimportant. Above the critical $f_{E,*}$ values, we found that RRTI causes the value of f_{trap} to self-adjust so that the radiation force exactly balances gravity without producing a wind. This is $f_{\text{trap}} \approx \tau_*/f_{E,*} - 1$. To extend this to lower τ_* than we have sampled, we simply impose the requirement that f_{trap} cannot be less than 0. Thus, our approximation for $f_{E,*} < 1$ is

$$f_{\text{trap,lo}} \approx \max \left(\frac{\tau_*}{f_{E,*}} - 1, 0 \right). \quad (28)$$

In the regime $f_{E,*} \rightarrow \infty$ that we explore in this paper there is a wind, but the rate at which it takes up momentum from the radiation field is limited. Fig. 13 shows our estimated values of f_{trap} as a function of τ_* , together with a crude linear fit that is consistent with the simulations:

$$f_{\text{trap,hi}} \approx 0.5\tau_*. \quad (29)$$

This is only a ‘by-eye’ fit, but it describes the data very well, and given the small number of simulations and the error bars on each one, a more sophisticated fitting procedure does not seem justified.

To combine the two cases $f_{E,*} < 1$ and $f_{E,*} \rightarrow \infty$, we hypothesize that f_{trap} will obey equation (28) up to the point where $f_{E,*} = 1$. At this point, we will have $f_{\text{trap}} \approx \tau_* - 1$ for $\tau_* \gg 1$. Beyond this point, as $f_{E,*}$ increases f_{trap} will smoothly decrease on to the fit given by equation (29) in the limit $f_{E,*} \rightarrow \infty$. Since we have not mapped out the intermediate $f_{E,*}$ regime, obviously the functional form of this decrease is not well-constrained, and we cannot rule out the possibility that f_{trap} behaves non-monotonically over this range, for example developing a peak at a special value of $f_{E,*}$. However, there is no good reason to believe that such a phenomenon should occur, and in its absence the functional form we adopt to interpolate between the behaviour at $f_{E,*} < 1$ and $f_{E,*} \rightarrow \infty$ matters little, since




Figure 13. Values of f_{trap} as a function of τ_* measured in the simulations presented in this paper. The points represent the time-averaged value, thick error bars show the standard deviation, and thin error bars show the range from minimum to maximum; all values are as given in Table 3. For runs T10LR and T10HR, the points are offset slightly from $\tau_* = 10$ for clarity. The dashed black line is $f_{\text{trap}} = 0.5\tau_*$.

f_{trap} only changes by a factor of 2 between those cases. We therefore adopt a very simple interpolation between the two cases,

$$f_{\text{trap}} \approx f_{\text{trap,hi}} + \frac{f_{\text{trap,lo}} - f_{\text{trap,hi}}}{\max(f_{E,*}, 1)} \quad (30)$$

in the regime where RRTI occurs. This fitting formula has the correct asymptotic behaviour in the limits $\tau_* \rightarrow 0$, $f_{E,*} < 1$ and $f_{E,*} \rightarrow \infty$, and is consistent with the simulations over the range of $f_{E,*}$ and τ_* we have explored. We can also use this formula to determine the ratio of mass-averaged radiation force to gravitational force. In Paper I, we showed that this is

$$\langle f_E \rangle = (1 + f_{\text{trap}}) \frac{f_{E,*}}{\tau_*}. \quad (31)$$

If a wind is launched, from equation (24), we can see that the net rate at which it gains momentum including both radiative acceleration and gravitational deceleration is

$$\frac{dp_{\text{wind}}}{dt} = (1 + f_{\text{trap}}) \left(\frac{\langle f_E \rangle - 1}{\langle f_E \rangle} \right) \frac{L}{c}. \quad (32)$$

Fig. 14 shows the results of our fit for f_{trap} , $\langle f_E \rangle$ and $(dp_{\text{wind}}/dt)/(L/c)$ as a function of τ_* and $f_{E,*}$. In the plot we can see the three dominant regimes identified in our simulations. Below the critical line there is no RRTI, and radiation is dynamically subdominant. For $f_{E,*}$ large enough to turn on RRTI but still below about unity, $\langle f_E \rangle$ is fixed to unity, and f_{trap} self-adjusts to compensate, decreasing as $f_{E,*}$ increases at fixed τ_* . Finally, at $f_{E,*}$ greater than about unity, a wind appears. In this regime f_{trap} is a function primarily of τ_* , and is quite insensitive to $f_{E,*}$. On the other hand $\langle f_E \rangle$ increases with $f_{E,*}$, indicating that gravity is becoming progressively weaker relative to radiation. As a result, the wind strength is monotonically increasing with $f_{E,*}$, but only slowly, since gravity is relatively unimportant once $f_{E,*}$ is even a slightly above unity. We stress that the exact location of the wind-launching line should not be taken too seriously, particularly at $\tau_* < 1$, given the sparsity with which we have sampled the parameter space. Nonetheless, the qualitative result that for $\tau_* > 1$ a wind appears only for $f_{E,*} \gtrsim 1$ should be robust.




Figure 14. Values of $\log f_{\text{trap}}$ (top), $\log \langle f_E \rangle$ (middle) and $\log[(dp_{\text{wind}}/dt)/(L/c)]$ (bottom) as a function of $f_{E,*}$ and τ_* , computed using the fitting formula given by equation (30). In each panel, contours lines appear at values of $-2, -1, 0, 1$ and 2 . The thick black line in the top two panels shows the critical curve below which RRTI shuts off (white region). The thick black line in the bottom panel shows the critical value below which no wind is launched.

Readers may note that, for $\tau_* \ll 1$, it is possible for there to be a wind even when $f_{E,*} < 1$ and RRTI does not occur. Physically, this corresponds to a medium that is optically thin to dust-reprocessed radiation, but it still absorbs the direct radiation field. In this case, the radiation is absorbed once, is re-emitted, and then immediately escapes, so the RRTI that we see in our simulations does not occur. However, if the direct radiation field carries enough momentum, this single absorption may still be sufficient to overcome gravity and launch a wind. Simulations by Kuiper et al. (2012) suggest that RRTI does not occur in this case, which is not surprising, since RRTI relies upon the ability of the gas to shape the radiation field. That cannot happen if the radiation is only absorbed once.

5.2 Implications for star-forming systems

Given our results for f_{trap} and $\langle f_E \rangle$, it is interesting to ask what our models predict for star-forming systems, which have been posited to be regulated by radiation pressure. We consider two types of objects: proto-star clusters, which we approximate as spherical, and galactic discs, which we approximate as planar. For a spherical object of total (gas plus stellar) mass M , gas mass fraction f_g , stellar mass fraction $f_* = 1 - f_g$, and radius R , within which the stars have a light-to-mass ratio Ψ , we have a central luminosity $L = \Psi f_* M$, gas surface

density $\Sigma = (1 - f_*)M/4\pi R^2$, surface gravitational acceleration $g = GM/R^2$, surface flux $F = L/4\pi R^2$ and surface temperature $T_* = (L/4\pi R^2\sigma)^{1/4}$. Thus, for the rough fit to the Semenov et al. (2003) opacity model given by equation (9), our estimates of the key dimensionless parameters near the surface become

$$\tau_* = 2.6 f_{g,0.5} f_{*,0.5}^{1/2} \Psi_3^{1/2} \Sigma_1^{3/2}, \quad (33)$$

$$f_{E,*} = 0.079 f_{*,0.5}^{3/2} \Psi_3^{3/2} \Sigma_1^{1/2}, \quad (34)$$

where $f_{*,0.5} = f_*/0.5$, $f_{g,0.5} = (1 - f_*)/0.5$, $\Sigma_1 = \Sigma/10 \text{ g cm}^{-2}$ and $\Psi_3 = \Psi/10^3 L_\odot/M_\odot$; for a zero-age stellar population, $\Psi = 1140 L_\odot/M_\odot$ (Fall et al. 2010), so our normalization should be appropriate for a young cluster. In computing τ_* , we have assumed that the gas is arranged in a thin shell (as it must be if it is to be ejected), so $\tau_* = \kappa_* M/4\pi R^2 = \kappa_* \Sigma/4$, with the factor of 4 arising from the difference between the surface density Σ as measured by an external observer (which is what appears in the above equations) and the surface density as seen by radiation escaping from the centre of the star cluster; if we instead adopt a uniform density sphere geometry, τ_* will be larger by a factor of 3. Similarly, variations in the dust opacity per unit gas mass could plausibly increase τ_* and $f_{E,*}$ by as much as a factor of a few, and downward by much larger factors in low-metallicity systems. Note that, as pointed out by Fall et al. (2010), M and R enter only through the combination Σ . The value to which we have normalized Σ , 10 g cm^{-2} , is roughly the maximum observed value for stellar systems anywhere in the Universe (Hopkins et al. 2010), and thus the values of τ_* and $f_{E,*}$ above should be regarded as upper limits. Note further that, even though $f_{E,*} < 1$ in equation (34), this does not imply that radiation is unimportant to the dynamics of the dusty gas. Indeed, radiation pressure should drive strong turbulence as in our simulations presented in Paper I.

We can perform a similar calculation for a galactic disc with total surface density Σ and gas mass fraction f_g . Such a disc has a surface gravitational force $g = 2\pi G \Sigma$. For a stellar population older than ~ 4 Myr, the light to star formation rate ratio approaches a roughly constant value $\Phi = 6.1 \times 10^{17} \text{ erg g}^{-1} \text{ s}^{-1} = 1.0 \times 10^{10} L_\odot/(M_\odot \text{ yr}^{-1})$ (calculated using starburst99 – Leitherer et al. 1999; Krumholz & Tan 2007; Krumholz & Dekel 2010). It is therefore convenient to write the radiative flux as $F = \Phi \dot{\Sigma}_*$, where $\dot{\Sigma}_*$ is the star formation rate per unit area. Plugging this flux into our scalings for the dimensionless parameters gives

$$\tau_* = 0.67 f_{g,0.5} \Phi_{10}^{1/2} \dot{\Sigma}_{*,3}^{1/2} \Sigma_0, \quad (35)$$

$$f_{E,*} = 0.43 \Phi_{10}^{3/2} \dot{\Sigma}_{*,3}^{3/2} \Sigma_0^{-1}, \quad (36)$$

where $\Sigma_0 = \Sigma/1 \text{ g cm}^{-2}$, $\dot{\Sigma}_{*,3} = \dot{\Sigma}_*/10^3 M_\odot \text{ pc}^{-2} \text{ Myr}^{-1}$ and $\Phi_{10} = \Phi/10^{10} L_\odot/M_\odot$. The normalizations of Σ and $\dot{\Sigma}_*$ here have been chosen to match those of the most vigorously star-forming galaxies observed. Indeed, none of the galaxies in the large sample compiled by Krumholz, Dekel & McKee (2012) exceed this star formation rate. Since observations indicate that $\dot{\Sigma}_* \propto \Sigma^p$ with $p \approx 1-1.5$, using a normalization for galaxies of lower star formation rates and surface densities would lead to lower values of τ_* and $f_{E,*}$. Thus, the values above are, as in the case of single clusters, upper limits for star-forming systems. However, we note that quasar discs on $\sim 1-50 \text{ pc}$ scales can and do exceed these limits (Sirko & Goodman 2003; Thompson et al. 2005).

Given these numbers for star clusters and galactic discs, we can draw a few conclusions. The first, already suggested in Paper I,

is that, in the absence of additional forces, the dust-reprocessed radiation field cannot launch winds or eject mass from the great majority of star clusters and galaxies. This is because we find that winds are only launched when $f_{E,*} > 1$. For star clusters even our upper limit is well below this value, and for galaxies only the most extreme systems approach it, while galactic winds are inferred to be ubiquitous (e.g. Veilleux et al. 2005). This is not to say that radiation pressure is not important. As discussed above, if the radiation force is sufficiently strong and $\tau_* < 1$, it may be possible for the *direct* radiation field to eject matter, particularly in systems where gravity is already partially offset by magnetic fields, turbulent motions or some other force (e.g. Murray et al. 2005; Krumholz & Matzner 2009; Fall et al. 2010; Murray et al. 2010; Hopkins et al. 2011). Indeed, Krumholz & Matzner (2009) compile a sample of super-star clusters, and show that for some of them the direct radiation force, combined with the momentum of line-driven winds, is likely to be able to eject matter even without significant radiative trapping. Even in somewhat lower luminosity systems where it cannot eject the bulk of the gas, radiation pressure may still be able to drive small amounts of mass to speeds above the escape speed and eject it, as happens for example with massive stars. Nonetheless, our results show that ejecting mass from star-forming systems via radiation pressure is significantly more difficult than many models assume.

A second implication of our work is that, if radiation pressure does launch winds, and if it were the sole driving mechanism, those winds are not likely to carry a momentum flux much larger than a few times L/c . Equations (33) and (35) show that τ_* on galactic scales never much exceeds unity, and that even for the densest clusters it is < 10 ; a more typical value for massive clusters would give $\tau_* \sim 1$. Since we find that $f_{\text{trap}} \approx 0.5\tau_*$, this means that we cannot expect winds accelerated primarily by radiation to have f_{trap} larger than ~ 1 . Thus, radiation pressure-driven winds from star-forming systems should not carry a momentum flux that exceeds L/c by more than a few tens of percent. The best-fitting values of the momentum fluxes of the winds produced by giant star-forming clumps at $z \sim 2$ exceed this limit (Genzel et al. 2011; Newman et al. 2012), which taken at face value would suggest that they cannot be primarily radiation driven (consistent with Krumholz & Dekel 2010). However, we caution that there are very significant uncertainties on these measurements, and for most sources any reasonable estimate of the error bars does not exclude a momentum flux that is close to L/c . Moreover, since the winds we observe now were launched some time ago, it is entirely possible that the present-day luminosity we measure for these sources is smaller than it was when the winds were launched. Given the observational uncertainties, we cannot conclude that giant clump winds cannot be radiatively driven, only that, if they are, either their momentum fluxes must have been overestimated or their luminosities at the point of wind launching underestimated.

It is interesting to ask how our conclusions compare with those of prior authors, and why they differ. Most notably, our conclusion that dust-reprocessed radiation is unlikely to be a significant factor in launching winds or disrupting massive clusters is inconsistent with those of a number of authors, including Murray et al. (2010), Murray et al. (2011), Hopkins et al. (2011, 2012) and Genel et al. (2012). We can understand the difference by examining how the radiation force imparted to the matter is computed in these models versus in our simulations. Murray et al. (2010, 2011) treat the matter as a one-dimensional thin shell, compute the resulting gas temperature and thus the opacity, and compute the radiation force by multiplying the flux by that opacity. Similarly, Hopkins et al. (2011, 2012) adopt a uniform, high dust opacity ($\kappa = 5 \text{ cm}^2 \text{ g}^{-1}$)

and compute the radiation force by multiplying that opacity by the total radiation flux. Genel et al. (2012) use a subgrid model that does not explicitly account for radiation forces or any other feedback effect, but manually injects winds at a strength that is tuned to match the results of the Murray et al. and Hopkins et al. models. Our numerical results show that these approaches can produce a large overestimate of the radiation force and the trapping factor. In light of this finding, it seems necessary to recompute the models of Murray et al. (2010), Murray et al. (2011), Hopkins et al. (2011, 2012) and Genel et al. (2012) using the approximate fitting formula for f_{trap} that we have derived.

5.3 Physical origins of weak trapping

Our simulations show that simple estimates of radiative trapping in dusty gas, such as those adopted by many previous authors, tend to produce values of the asymptotic momentum of the dusty gas that are substantially too large. We can identify two significant effects that contribute to this error. First, these models did not take into account the anticorrelation between radiation flux and gas density produced by RRTI, which causes the flux seen by the bulk of the matter to be significantly smaller than the volume-averaged flux. We illustrate this effect in Fig. 15, which shows the distribution of density and flux in the simulation volume for a particular time in run T10LR. From the figure, it is clear that dense matter generally has a much lower radiation flux passing through it than more diffuse matter. Secondly, they miss the effect that, because the flux is low within the bulk of the matter, the gas and radiation temperatures are also low. This reduces the opacity within the bulk of the matter, further weakening the matter–radiation coupling.

It is interesting to investigate further the relative importance of these two effects, particularly because this has implications for how general our results should be. The flux-density anticorrelation would seem likely to be quite independent of the particular opacity law, and thus potentially important even in systems where the opacity has




Figure 15. Distribution of gas density and radiation flux throughout the computational volume, for run T10LR at $t = 12.4t_a$, the same run and time shown in Fig. 9. The colour in each pixel of $(\rho/\rho_a, F/F_0)$ indicates the fraction of the simulation volume has the indicated density and total radiation flux. The colour scale is normalized so that the value in the most populated bin is unity. White pixels correspond to values of density and flux that are not populated.




Figure 16. Values of $\langle\tau\rangle$ as a function of t for all simulations (solid lines). For comparison, we also show τ_{IR} (dotted lines) and f_{trap} (dashed lines; identical to the values shown in the lower panels of Fig. 10).

a different functional form (e.g. electron scattering opacity, which is independent of temperature). The reduction in opacity due to a reduction of the temperature in the bulk of the matter, on the other hand, relies on the particular property that the opacity rises with temperature in dusty gas, and it is unclear how this effect might change for different opacity laws.

To address this question, we compute the quantity

$$\langle\tau\rangle \equiv \langle\kappa_R\rho\rangle L_z, \quad (37)$$

from our simulations. Physically, this is the optical depth computed using the mass-weighted mean opacity. If the flux were uniform (i.e. $F_z = F_0$ everywhere), it is easy to show that we would have $f_{\text{trap}} = \langle\tau\rangle$. Thus, the value of $\langle\tau\rangle$ provides some insight into the relative importance of the non-uniformity of the flux. If the flux-density anticorrelation is the dominant effect in reducing f_{trap} , since $\langle\tau\rangle$ excludes this effect we should find $\langle\tau\rangle \sim \tau_{\text{IR}}$, where τ_{IR} is the expected trapping factor for a laminar medium as estimated for example from the models of Thompson et al. (2005), Murray et al. (2010) and Murray et al. (2011). On the other hand, if f_{trap} is reduced primarily due to the way RRTI reduces the temperature and thus the opacity, then we would expect $\langle\tau\rangle \sim f_{\text{trap}}$ instead.

We show $\langle\tau\rangle$ as a function of time in Fig. 16, where for comparison we also plot f_{trap} and τ_{IR} . We also report the time-averaged value of $\langle\tau\rangle$, computed in the same manner as for f_{trap} , in Table 3.¹ We find that $\langle\tau\rangle$ is, in a logarithmic sense, roughly half-way between τ_{IR} and the actual value of f_{trap} we measure in our simulations. This indicates that the flux-density anticorrelation and the reduction in opacity due to non-uniform temperatures are about equally important in reducing the rate of momentum transfer from radiation to gas in our dusty flows. That there is at least some reduction in the momentum transfer to the gas even due to flux-density anticorrelation even when the opacity is grey is qualitatively consistent with the findings of Jiang et al. (2013), who find RRTI operating in a medium with a grey opacity.

Based on this result, we tentatively predict that even for an opacity that depends on a different positive power of the temperature, the flux-density anticorrelation induced by RRTI should produce some reduction in f_{trap} compared to what one would estimate using a model in which the density and radiation flux are taken to be uniform. However, if the opacity is grey this reduction should be smaller

¹ Careful readers may notice that, for run T30, the plot of $\langle\tau\rangle$ does not extend quite as far in time as the plot of f_{trap} , and in some places is more smoothly interpolated in time; this is due to a coding error that caused some portion of the output from this run to be lost, preventing us from performing the post processing required to calculate $\langle\tau\rangle$ at the corresponding times. This issue did not affect f_{trap} because the values required for it were calculated as the simulation ran.

than what we find here by roughly a factor of 2 in a logarithmic sense. These issues are particularly important in the case of dust opacities since it is only for $T < 150 - 200$ K that $\kappa_R \propto T^2$. At higher temperatures, the dust opacity is approximately constant with temperature (Semenov et al. 2003). Thus, in some regimes we would expect the flux-density anticorrelation to dominate. Compared to our $\kappa_R \propto T^2$ model, a dust opacity that flattened at high temperatures would probably produce slightly lower f_{trap} (since the decrease in opacity should at least marginally reduce the net radiation force) but much lower τ_{IR} (since the opacity at the mid-plane would be greatly reduced), and thus less of a difference between τ_{IR} and f_{trap} .

It is unclear what would happen in the case of an opacity that falls with temperature, as is the case for a Kramers opacity law, but it is conceivable that in this case f_{trap} might actually be closer to the naive laminar estimate, since for such an opacity the temperature-dependence induced by the non-uniform radiation field would tend to increase rather than decrease the strength of matter–radiation coupling. However, we emphasize that all of these conclusions are tentative, since it is entirely possible that the non-linear development of the RRTI, and thus the amount of flux-density anticorrelation it produces, itself depends on the opacity law.

5.4 Relation to dusty star winds

Although our work is focused on the problem of star cluster and galactic winds, the general problem of radiative driving of dusty gas also arises in the context of winds from dusty late-type stars (e.g. Goldreich & Scoville 1976; Habing 1996). Before proceeding, it is important to point out a significant way in which this problem differs from our work here and in Paper I. At the low temperatures typical of interstellar gas even in intensely star-forming galaxies, the opacity roughly $\kappa \propto T^2$ (e.g. Semenov et al. 2003), as we use in our models. This means that the opacity general drops monotonically with height in an atmosphere. In the case of dusty stellar winds, on the other hand, where temperatures are near the grain sublimation temperature, the opacity is much more complex and non-monotonic, both due to grain formation, destruction and drift relative to the gas, and because even for a constant grain population the opacity varies are roughly T^0 (i.e. constant) rather than T^2 at temperatures close to the grain sublimation temperature. As discussed in the previous section, for such an opacity law we tentatively expect that there will be some reduction in f_{trap} compared to the expected values for a uniform medium, but that it will be weaker than what we find here.

With this caveat aside, we note that one-dimensional models have been reasonably successful at reproducing many observations of massive star winds (e.g. Ivezić & Elitzur 1995; Elitzur & Ivezić 2001; Ivezić & Elitzur 2010), implying that RRTI may not be critical for these stars. However, the observationally inferred momenta of dusty star winds are usually below L/c (e.g. Groenewegen et al. 2009), and that even the highest inferred momenta are no more than $\sim 10L/c$ (e.g. Elitzur & Ivezić 2001), which in turn suggests that these stars are not in the regime where we require strong amplification of the force by radiative trapping that might be inhibited by RRTI. On the other hand, one might legitimately worry that even the more modest reduction in optical depth we expect for a grey opacity might make it difficult to accelerate winds to $10L/c$. We tentatively conclude that there might be some tension between our numerical results and the current generation of asymptotic giant branch wind models, and potentially even some tension with the observations. However, the problem clearly needs further investigation, since our estimate of the factor by which RRTI should reduce f_{trap} for complex opacity laws is highly uncertain.

5.5 Limitations due to geometric simplifications of the simulations

Our simulations represent an idealized numerical experiment with a simple geometry. It is therefore interesting to ask how a more realistic setup would likely affect our results. One obvious simplification in our simulations is that they are two rather than three dimensional. The implications of this are discussed extensively in Paper I, and we refer readers to the discussion there.

A second simplification is that we have assumed a planar geometry, whereas a real wind will generally approach a spherical geometry, at least once it is far from its launch point. A small section of a spherical shell of wind material may be treated as planar, and so our planar results should continue apply locally. The main difference between planar and spherical geometries, therefore, is that in planar geometry τ_* and $f_{E,*}$ are fixed, whereas for a spherical wind they will vary as the wind expands. This variation is caused by two effects. First, as a spherical shell of constant mass expands in radius R , its surface density drops as R^{-2} , which reduces τ_* . Secondly, the gravitational force g and the flux F_0 encountering a spherical shell also both drop as R^{-2} . This means that the ratio F_0/g remains constant; however, the drop in F_0 reduces T_* as $R^{-1/2}$ and thus $\kappa_{R,*}$ as R^{-1} , which affects both $f_{E,*}$ and τ_* . The net effect is that, for a spherical shell of fixed mass and radius R , $f_{E,*} \propto R^{-1}$ and $\tau_* \propto R^{-3}$, and thus an expanding spherical shell traces a line of slope 3 in the $(\log f_{E,*}, \log \tau_*)$ plane depicted in Fig. 14. Systems start at the upper right of the plane, then move down and to the left as they expand. Examining the figure, we see that such a trajectory will result in a value of f_{trap} that decreases with time, and that it approaches an asymptotically constant value of $\langle f_E \rangle$. This is not surprising; it is simply a statement that, as a shell expands, its optical depth drops and thus the dust-reprocessed radiation field becomes less and less important compared to the direct one, which gives constant $\langle f_E \rangle$.

A third simplification of our simulations is that we have assumed a constant flux as would be produced in a galaxy with all the stars at the mid-plane, or by a star cluster with all the stars concentrated at the centre and the mass at fixed radius. In reality, the sources of radiation are intermixed with the gas being launched in a wind. This may well result in a significant reduction of the direct radiation force due to geometric cancellations. However, it should not substantially affect our results for f_{trap} , simply because an appreciable value of f_{trap} requires that the radiation field be trapped and therefore isotropized, forgetting its original direction. Thus, our results for f_{trap} should be robust against a change from planar or point-like sources to distributed sources.

5.6 Limitations of the flux-limited diffusion approximation

Our simulations make use of the 2TFLD approximation, which is an approximate treatment of radiative transfer. The approximation is highly accurate when the optical depth is very high, but at low optical depth it loses information about the directionality of the radiation field, leading to an inaccurate treatment of shadowing and similar effects. It is therefore important to ask to what extent our results might be affected by the limitations of our numerical method. While a full answer to this question can only come from repetition of the simulation with a more accurate radiation transport method, we can make a few general observations here.

Fig. 17 shows the flux limiter λ in one of our calculations at the same time as shown in Fig. 9. Values of λ close to $1/3$ correspond to locations where the optical depth is high, and the 2TFLD




Figure 17. Values of the flux limiter λ at the same time and for the same run as shown as in Fig. 9: $t = 12.4t_a$ in run T10LR. Values of $\lambda = 1/3$ indicate optically thick regions, while values near 0 indicate optically thin ones. The left-hand panel shows the entire computational domain, while the right-hand panel shows the same zoomed-in region as in Fig. 9, which is centred on the vertical centre of mass. Note that, in the left-hand panel, the x direction has been stretched by a factor of 16 relative to the z direction in order to render the image readable. In the left-hand panel, the white horizontal solid lines indicate the upper and lower boundaries of the zoomed-in region on the right. The two white horizontal dashed line indicates the values of z below which 90 per cent and 99 per cent of the mass lie.

approximation should be very accurate, while values of λ close to 0 correspond to optically thin regions. The white horizontal dashed lines indicate the values of z below which 90 per cent and 99 per cent of the mass in the computational domain lie. From the figure, it is clear that the great majority of the mass lies in a region where λ is very close to $1/3$. There are smaller values of λ within the channels through which radiation is escaping, but even there λ only falls to ~ 0.25 – 0.3 except in tiny volumes. It seems that the channels that carry most of the flux are optically thin only in a relative sense, i.e. they have much lower optical depth than most of the higher density gas around them, but they still have optical depths ~ 1 , not $\ll 1$. In contrast, λ does approach 0, as expected, high above where most of the matter is located.

Based on this figure, it seems likely that our treatment of radiation transport is reasonably accurate for the great majority of the mass. Where one might legitimately be concerned with the accuracy of the 2TFLD approximation is in the intermediate optical depth regime that lies, roughly, in the range from $z/h_a \approx 500 - 1000$ in the plot. However, as the dashed horizontal lines indicate, this regions contains no more than a few per cent of the mass in the simulation domain. In essence, we find that almost all the gas is in an optically thick shell where $\lambda \approx 1/3$ and the 2TFLD approximation should

be accurate, while a very small fraction is in a very extended atmosphere where its accuracy might be poor. However, since we are primarily concerned with the momentum imparted to the bulk of the mass, this extended atmosphere should have little effect on our main result, which is that this momentum is quite limited. Nonetheless, we caution that the precise details of the structure of the atmosphere, or the interiors of the low optical depth channels further within the shell, might be affected by the limitations of the 2TFLD method. These should therefore be treated with some caution.

Finally, we note that Jiang et al. (2013) have investigated the problem of RRTI in an optically thick medium for the case where the opacity is dominated by electron scattering rather than dust absorption, and have done so using a variable Eddington tensor (VET) method that should be more accurate than 2TFLD. They compare their VET results to those obtained using Eddington approximation, which is closely related to the diffusion approximation we use. The results of their comparison are consistent with our conclusions in this section. They find that the primary effect of going from the Eddington approximation to VET is to suppress some of the small-scale structure induced by RRTI, leading to larger scale filaments and channels (their fig. 9). This also has the effect of delaying the initial onset of instability some, because in the linear regime small modes grow fastest, and suppression of such modes leads to somewhat slower growth at first. However, once larger wavelength modes go non-linear, the difference in growth rates is small (their fig. 8). Since the time spent during the linear phase of instability growth for small modes is astrophysically negligible for the parameters describing real ULIRGs and massive star clusters (see the discussion in Paper I), it seems unlikely that a more accurate radiation transport method would change our results qualitatively.

6 SUMMARY

In this paper, we analyse the properties of optically thick radiation pressure-driven dusty winds. We consider the idealized problem of a column of material through which a specified radiation flux is passed. We first show that such a system is characterized by a single dimensionless number, τ_* , the optical depth of the matter computed using the opacity at the dust photosphere, and that this parameter will determine the rate at which the matter column absorbs momentum from the radiation field. We then use radiation-hydrodynamic simulations to measure this momentum transfer rate. We find that, after one to a few dynamical times, RRTI drives the gas into a configuration where most of the matter is in dense filaments aligned along the direction of the radiation flux, while most of the radiation flux passes through channels of reduced optical depth between the filaments. This configuration minimizes matter–radiation interaction, and thus limits the rate at which matter can take up momentum from the radiation field.

We combine this result with the result from Paper I, where we considered irradiated columns of matter confined by gravity, to produce a fitting formula for the behaviour of irradiated, gravity-confined dusty gas layers. The behaviour of these structures is determined by τ_* and by $f_{E,*}$, the ratio of radiative and gravitational forces at the dust photosphere. We identify three regimes of behaviour depending on the values of these parameters. At a given τ_* , there is a critical value of $f_{E,*}$ below which radiation is dynamically unimportant. For values of $f_{E,*}$ above the critical value but below unity, RRTI sets in and makes the gas turbulent, but does not produce a wind. Only for $f_{E,*} \gtrsim 1$ is there a wind, and even in the limit where gravity provides negligible confinement of that wind, the wind momentum flux is roughly $1 + 0.5\tau_*$ times the radiation momentum flux.

We then consider the implications of these results for star-forming clusters and galaxies. For observed clusters and galaxies, our results suggest that dust-reprocessed radiation is unlikely to be able to drive winds and eject matter. The direct radiation field may still be able to launch winds, but only in systems where its momentum alone is sufficient to overcome gravity, without significant amplification by radiative trapping.

ACKNOWLEDGEMENTS

We thank the anonymous referee for a helpful report. MRK acknowledges support from the Alfred P. Sloan Foundation, the NSF through grant CAREER-0955300, and NASA through Astrophysics Theory and Fundamental Physics Grant NNX09AK31G, and a Chandra Space Telescope Grant. TAT acknowledges support from the Alfred P. Sloan Foundation and NASA grant NNX10AD01G. This research was supported by allocations of advanced computing provided by the National Science Foundation and by the NASA High-End Computing Program. Simulations were performed on Kraken at the National Institute for Computational Sciences (<http://www.nics.tennessee.edu/>), Ranger at the Texas Advanced Computing Center at The University of Texas at Austin (<http://www.tacc.utexas.edu>) and Pleiades at the NASA Advanced Supercomputing (NAS) Division at Ames Research Center.

REFERENCES

- Abdo A. A. et al., 2010, *ApJ*, 709, L152
 Agertz O., Kravtsov A. V., Leitner S. N., Gnedin N. Y., 2012, *ApJ*, 770, 1
 Andrews B. H., Thompson T. A., 2011, *ApJ*, 727, 97
 Blaes O., Socrates A., 2003, *ApJ*, 596, 509
 Chiao R. Y., Wickramasinghe N. C., 1972, *MNRAS*, 159, 361
 Downes D., Solomon P. M., 1998, *ApJ*, 507, 615
 Elitzur M., Ivezić Ž., 2001, *MNRAS*, 327, 403
 Elmegreen B. G., 1983, *MNRAS*, 203, 1011
 Fall S. M., Krumholz M. R., Matzner C. D., 2010, *ApJ*, 710, L142
 Faucher-Giguère C.-A., Quataert E., 2012, *MNRAS*, 425, 605
 Ferrara A., 1993, *ApJ*, 407, 157
 Genel S. et al., 2012, *ApJ*, 745, 11
 Genzel R. et al., 2011, *ApJ*, 733, 101
 Goldreich P., Scoville N., 1976, *ApJ*, 205, 144
 Goldsmith P. F., 2001, *ApJ*, 557, 736
 Groenewegen M. A. T., Sloan G. C., Soszyński I., Petersen E. A., 2009, *A&A*, 506, 1277
 Habing H. J., 1996, *A&AR*, 7, 97
 Hopkins P. F., Murray N., Quataert E., Thompson T. A., 2010, *MNRAS*, 401, L19
 Hopkins P. F., Quataert E., Murray N., 2011, *MNRAS*, 417, 950
 Hopkins P. F., Quataert E., Murray N., 2012, *MNRAS*, 421, 3522
 Indriolo N., McCall B. J., 2012, *ApJ*, 745, 91
 Ivezić Z., Elitzur M., 1995, *ApJ*, 445, 415
 Ivezić Ž., Elitzur M., 2010, *MNRAS*, 404, 1415
 Jacquet E., Krumholz M. R., 2011, *ApJ*, 730, 116
 Jiang Y.-F., Davis S. W., Stone J. M., 2013, *ApJ*, 763, 102
 Kim J.-H., Krumholz M. R., Wise J. H., Turk M. J., Goldbaum N. J., Abel T., 2013a, *ApJ*, preprint (arXiv:1210.3361)
 Kim J.-H., Krumholz M. R., Wise J. H., Turk M. J., Goldbaum N. J., Abel T., 2013b, *ApJ*, preprint (arXiv:1210.6988)
 Krumholz M. R., 2013, *MNRAS*, preprint (arXiv:1304.2404)
 Krumholz M. R., Dekel A., 2010, *MNRAS*, 406, 112
 Krumholz M. R., Matzner C. D., 2009, *ApJ*, 703, 1352
 Krumholz M. R., Tan J. C., 2007, *ApJ*, 654, 304
 Krumholz M. R., Thompson T. A., 2012, *ApJ*, 760, 155 (Paper I)
 Krumholz M. R., Klein R. I., McKee C. F., Bolstad J., 2007, *ApJ*, 667, 626

- Krumholz M. R., Klein R. I., McKee C. F., Offner S. S. R., Cunningham A. J., 2009, *Sci.*, 323, 754
 Krumholz M. R., Dekel A., McKee C. F., 2012, *ApJ*, 745, 69
 Kuiper R., Klahr H., Beuther H., Henning T., 2012, *A&A*, 537, A122
 Lacki B. C., Thompson T. A., Quataert E., 2010, *ApJ*, 717, 1
 Leitherer C. et al., 1999, *ApJS*, 123, 3
 Levermore C. D., 1984, *J. Quant. Spectrosc. Radiat. Transfer*, 31, 149
 Levermore C. D., Pomraning G. C., 1981, *ApJ*, 248, 321
 Lopez L. A., Krumholz M. R., Bolatto A. D., Prochaska J. X., Ramirez-Ruiz E., 2011, *ApJ*, 731, 91
 Murray N., Quataert E., Thompson T. A., 2005, *ApJ*, 618, 569
 Murray N., Quataert E., Thompson T. A., 2010, *ApJ*, 709, 191
 Murray N., Ménard B., Thompson T. A., 2011, *ApJ*, 735, 66
 Narayanan D., Krumholz M., Ostriker E. C., Hernquist L., 2011, *MNRAS*, 418, 664
 Narayanan D., Krumholz M. R., Ostriker E. C., Hernquist L., 2012, *MNRAS*, 421, 3127
 Newman S. F. et al., 2012, *ApJ*, 761, 1
 Novak G. S., Ostriker J. P., Ciotti L., 2012, *MNRAS*, 427, 2734
 O'dell C. R., York D. G., Henize K. G., 1967, *ApJ*, 150, 835
 Oppenheimer B. D., Davé R., 2006, *MNRAS*, 373, 1265
 Peters T., Klessen R. S., Mac Low M., Banerjee R., 2010, *ApJ*, 725, 134
 Scoville N., 2003, *J. Korean Astron. Soc.*, 36, 167
 Scoville N. Z., Polletta M., Ewald S., Stolovy S. R., Thompson R., Rieke M., 2001, *AJ*, 122, 3017
 Semenov D., Henning T., Helling C., Ilgner M., Sedlmayr E., 2003, *A&A*, 410, 611
 Sirko E., Goodman J., 2003, *MNRAS*, 341, 501
 Thompson T. A., Quataert E., Murray N., 2005, *ApJ*, 630, 167
 Veilleux S., Cecil G., Bland-Hawthorn J., 2005, *ARA&A*, 43, 769
 Wise J. H., Abel T., Turk M. J., Norman M. L., Smith B. D., 2012, *MNRAS*, 427, 1
 Zhang D., Thompson T. A., 2012, *MNRAS*, 424, 1170

APPENDIX A: TEMPERATURE AND DYNAMICAL APPROXIMATIONS

Here, we justify three physical approximations we have made in our formulation of the governing equations: first, that the dust and gas temperatures are tightly coupled; secondly, that, at least in regions of high opacity, the dust and radiation temperatures are as well; thirdly, that the dust and gas are dynamically coupled and can be treated as a simple fluid.

A1 Thermal behaviour

To check the first two of these assumptions, we must compare the rate of dust–gas energy exchange via collisions with other rates of heating and cooling in the problem. Specifically, we must check that (1) the rate of energy exchange between the dust and the radiation field greatly exceeds the rate of collisional energy exchange between the dust and gas, ensuring that the dust is always well coupled the radiation field and (2) the rate of collisional energy exchange between the dust and gas is much larger than any other rate of heating or cooling for the gas alone, so that the gas is forced to the dust temperature. In what follows, unless stated otherwise, the rates of all heating and cooling processes follow the approximations described in Krumholz (2013).

The rate of dust–gas energy exchange per unit volume is

$$\Psi_{\text{gd}} = \alpha_{\text{gd}}(\rho/\mu_{\text{H}})^2 T_{\text{g}}^{1/2} (T_{\text{d}} - T_{\text{g}}) \quad (\text{A1})$$

$$= 5.8 \times 10^{-16} \rho_{-16}^2 T_{\text{g},2}^{1/2} (T_{\text{d}} - T_{\text{g}})_2 \text{ erg cm}^{-3} \text{ s}^{-1}, \quad (\text{A2})$$

where $\alpha_{\text{gd}} \approx 3.2 \times 10^{-34} \text{ erg cm}^3 \text{ K}^{-3/2}$ is the dust–gas collisional coupling coefficient, $\mu_{\text{H}} \approx 2.3 \times 10^{-24} \text{ g}$ is the gas mass per H nucleus, T_{g} is the gas temperature, T_{d} is the dust temperature, $\rho_{-16} = \rho/10^{-16} \text{ g cm}^{-3}$, $T_{\text{g},2} = T_{\text{g}}/100 \text{ K}$ and $(T_{\text{d}} - T_{\text{g}})_2 = (T_{\text{d}} - T_{\text{g}})/100 \text{ K}$. The scalings we have chosen here for ρ and T are intentionally close to ρ_{a} and T_* , so the numerical coefficients should be roughly representative of the actual values in the simulation, at least during the initial phases. The value to which we scale $T_{\text{d}} - T_{\text{g}}$ is intended to be representative of what would be required to generate a relative dust–gas temperature difference of order unity. The rate of dust energy exchange with the radiation field per unit volume is

$$\Psi_{\text{dr}} = \kappa_{\text{p}} \rho (4\pi B - cE) = \kappa_{\text{p}} \rho c a (T_{\text{d}}^4 - T_{\text{r}}^4) \quad (\text{A3})$$

$$= 2.3 \times 10^{-13} \rho_{-16} (T_{\text{d}} - T_{\text{r}})_2^4 \text{ erg cm}^{-3} \text{ s}^{-1}, \quad (\text{A4})$$

where we have written $E = aT_{\text{r}}^4$ for convenience.

For the gas, the competing heating processes are cosmic ray heating and shock heating, while the competing cooling process is molecular line emission. The cosmic ray heating rate is

$$\Gamma_{\text{CR}} = \zeta q_{\text{ion}} \rho / \mu_{\text{H}} \quad (\text{A5})$$

$$= 1.0 \times 10^{-17} \zeta_{-14} \rho_{-16} \text{ erg cm}^{-3} \text{ s}^{-1}, \quad (\text{A6})$$

where ζ is the primary ionization rate per H nucleus from cosmic rays, $\zeta_{-14} = \zeta/10^{-14} \text{ s}^{-1}$, and $q_{\text{ion}} \approx 15 \text{ eV}$ is the thermal energy added per primary ionization. The cosmic ray ionization rate to which we have scaled here is roughly 100 times a typical Milky Way value (e.g. Indriolo & McCall 2012), consistent with the values predicted by Lacki, Thompson & Quataert (2010) for starburst galaxies. There is also direct evidence for enhanced cosmic ray fluxes in starburst galaxies from γ -ray emission (Abdo et al. 2010). It is conceivable that the cosmic ray flux could be even higher than this, which would have important implications of this for the thermodynamics of the gas in ULIRGs, see Narayanan et al. (2011, 2012). The shock heating rate of course varies with position and time, but we can obtain a crude order of magnitude estimate by positing that the full kinetic energy per unit volume of the gas is dissipated into shocks each crossing time. This implies a heating rate

$$\Gamma_{\text{shock}} \approx \rho \sigma^3 / h_{\text{g}} \quad (\text{A7})$$

$$= 3.2 \times 10^{-17} \rho_{-16} \sigma_{6,0}^3 h_{\text{g},0}^{-1} \text{ erg cm}^{-3} \text{ s}^{-1}, \quad (\text{A8})$$

where σ is the gas velocity dispersion and h_{g} is the characteristic size of the region occupied by the gas and $h_{\text{g},0} = h_{\text{g}}/1 \text{ pc}$. The scaling here is again chosen to be close to the characteristic size of the structures seen in our simulations, which is $\sim 10 - 100 h_{\text{a}}$. Finally, the rate of molecular line cooling is the most difficult to estimate, since it depends on the complex physics of optically thick line emission. We use the `DESPOTIC` package (Krumholz 2013) to compute the cooling rate via CO emission for a cloud of volume density $10^{-16} \text{ g cm}^{-3}$, column density $\Sigma = 5 \text{ g cm}^{-2}$ and gas temperature 100 K, with a CO abundance of 1×10^{-4} per H nucleus, and obtain

$$\Lambda_{\text{CO}} = 3.5 \times 10^{-17} \text{ erg cm}^{-3} \text{ s}^{-1}. \quad (\text{A9})$$

Based on these calculations, we see that, for material with density $\sim \rho_{\text{a}}$, temperature $\sim T_*$ and velocity dispersion $\sim 10 \text{ km s}^{-1}$ (comparable to what we find in our simulations), we expect $|\Psi_{\text{dr}}| \gg |\Psi_{\text{g}}| \gg |\Gamma_{\text{CR}}| + |\Gamma_{\text{shock}}| + |\Lambda_{\text{CO}}|$. The former inequality ensures that the dust and radiation field will be forced to the same temperature, while the latter ensures that the dust and gas will be forced to the same temperature. Our result is not surprising in light of earlier work. Goldsmith (2001) shows that dust and gas as well coupled once the density exceeds $\sim 10^4 - 10^5 \text{ cm}^{-3}$, and Narayanan et al. (2011, 2012) conclude that strong dust–gas coupling dominates the thermodynamics of gas in ULIRGs.

However, we do caution that the gap between the dust–gas energy exchange rate and the dominant gas heating and cooling terms is only a factor of ~ 10 . At late times in the simulations, the density drops significantly as material spreads out, and once this happens the gas and dust should thermally decouple. Since the gap between the gas–dust energy exchange rate and the dust–radiation energy exchange rate is much larger, the dust would remain thermally coupled to the radiation field, but the gas would not. In practice, this means that the error in our assumption that the gas is at the same temperature as the dust manifests not as an error in the opacity (which would still be determined by the coupled dust and radiation temperatures, consistent with our assumptions), but instead as an error in the gas thermal pressure. This is unlikely to affect the dynamics significantly, because the gas velocities are generally highly supersonic, and thus thermal pressure is weak compared to ram pressure.

A2 Dynamical behaviour

We now investigate the final assumption, which is that dust and gas are dynamically well coupled, and do not drift significantly relative to one another. If the only mechanism capable of exchanging momentum between gas and grains is collisions, then the characteristic free-streaming length of a dust grain of mass m_{D} with respect to the gas is

$$\lambda_{\text{dg}} \approx \frac{m_{\text{D}}}{\mu} \frac{1}{n_{\text{g}} \sigma}, \quad (\text{A10})$$

where μ is the mean mass per gas particle, n_{g} is the number density of gas particles and σ is the cross-section for grain–gas collisions. This equation holds as long as the grain is small compared to the particle mean-free path, which will be the case for typical interstellar grains. If we take the grain to be a sphere of radius a and density ρ_{D} , and the cross-section $\sigma = \pi a^2$, then we can rewrite this equation as

$$\lambda_{\text{dg}} \approx \frac{4}{3} a \frac{\rho_{\text{D}}}{\rho} \quad (\text{A11})$$

$$= 1.3 \times 10^{-7} a_{-1} \rho_{\text{D},3} \rho_{-16} \text{ pc}, \quad (\text{A12})$$

where $a_{-1} = a/0.1 \mu\text{m}$, $\rho_{\text{D},3} = \rho_{\text{D}}/3 \text{ g cm}^{-3}$, and the values to which we have scaled in the numerical evaluation are typical for interstellar silicate grains.

In comparison, the characteristic length-scale for our problem is $h_{\text{a}} \sim 10^{-2} \text{ pc}$, and thus the characteristic grain streaming length is a factor of $\sim 10^5$ smaller. We can therefore safely ignore grain streaming at the stages of our simulation where the structure is developing. As with thermal coupling, at very late times once the density drops by ~ 5 orders of magnitude, grain streaming with respect to the matter can become significant, and the approximations we use may begin to break down. However, we note that equation (A11) represents a maximum possible grain streaming length,

because it ignores magnetic coupling between grains and gas. If the grains are charged, then magnetic forces will tie them to the gas more strongly than hydrodynamic collisions and make free streaming less important.

APPENDIX B: INITIAL CONDITIONS FOR RUN T30

For run T30, we do not have initial conditions from Paper I because we did not perform any runs with $\tau_* = 30$. To generate such conditions, we run a simulation with $\tau_* = 30$ and $f_{E,*} = 0.05$ (i.e. with gravity turned on) following the same procedure as for all other runs described in Paper I. We refer to this run as T30F0.05. As in the other runs from Paper I, we perform the simulation at a resolution $\Delta x = 0.5h_*$, in a computational domain of 1024×16384 cells, corresponding to a size of $512h_* \times 8192h_*$. We run the simulation for a time $t = 75t_*$, by which point a turbulent flow is fully devel-

oped. To produce initial conditions for run T30 in this paper, we must rescale the results of run T30F0.05, because $h_a/h_* = 600$, so that the resolution of run T30F0.05 is $\Delta x = 8.3 \times 10^{-4}h_a$. This is so high that it would be impossible to advance the run for a time comparable to t_a . We therefore downsample the output at the final time in run T30F0.05 by a factor of 32, producing a resolution of $\Delta x = 0.027h_a$. We also replicate the density, velocity, gas temperature and radiation energy density fields 32 times in the horizontal direction; since run T30F0.05 has periodic boundary conditions, this is fully self-consistent. The result is a cube of initial conditions that is 1024×512 cells in size, at a resolution $\Delta x = 0.027h_a$, corresponding to a physical size $27.3h_a \times 13.6h_a$. We use this state as the initial condition for run T30, extending the computational domain in the vertical direction exactly as for the other runs described in Section 3.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.