The momentum budget of clustered supernova feedback in a 3D, magnetised medium

Eric S. Gentry,[1][*] Mark R. Krumholz,[2] Piero Madau,[1,3] Alessandro Lupi[3]

[1]*Department of Astronomy and Astrophysics, University of California at Santa Cruz, 1156 High St., Santa Cruz, CA, 95064, USA*
[2]*Research School of Astronomy & Astrophysics, Australian National University, Canberra, ACT 2611, Australia*
[3]*Sorbonne Université, CNRS, UMR 7095, Institut d'Astrophysique de Paris, 98 bis bd Arago, 75014 Paris, France*

ABSTRACT

While the evolution of superbubbles driven by clustered supernovae (SNe) has been studied by numerous authors, the resulting radial momentum yield is uncertain by as much as an order of magnitude depending on the computational methods and the assumed properties of the surrounding interstellar medium (ISM). In this work, we study the origin of these discrepancies, and seek to determine the correct momentum budget. We carry out 3D hydrodynamic and magnetohydrodynamic simulations of clustered supernova explosions, using a Lagrangian method and checking for convergence as we increase the resolution. We find that the terminal momentum of a shell driven by clustered supernovae is dictated almost entirely by the rate of mixing across the contact discontinuity between the hot and cold phases, and that this mixing rate is dominated by numerical diffusion even at the highest resolution we can achieve, 0.03 $M_\odot$. Magnetic fields also reduce the mixing rate, so that MHD simulations produce higher momentum yields than HD ones at equal resolution. As a result, we obtain only a lower limit on the momentum yield from clustered supernovae. Combining this with our previous 1D results, which provide an upper limit because they allow almost no mixing across the contact discontinuity, we conclude that the momentum yield per SN from clustered SNe is bounded between 2×10^5 and 3×10^6 $M_\odot$ km s^{-1}. We urge that published claims to provide a converged value for this quantity, which are based on simulations with substantially higher rates of numerical diffusion than ours, be treated with caution.

Key words: hydrodynamics – magnetic fields – ISM: bubbles – ISM: supernova remnants

1 INTRODUCTION

Feedback from supernovae (SNe) is an important component of understanding the interstellar medium (ISM), galactic winds, and galactic evolution (e.g., McKee & Ostriker 1977; Dekel & Silk 1986; Murray et al. 2005; Ostriker & Shetty 2011; Kim et al. 2011; Jenkins & Tripp 2011; Hopkins et al. 2012; Dekel & Krumholz 2013; Faucher-Giguère et al. 2013; Creasey et al. 2013; Thompson & Krumholz 2016). Unfortunately, the processes governing the strength of SN feedback operate nonlinearly and at small scales. This makes it difficult to include in analytic models or large galactic simulations without a simplified prescription for SN feedback.

In the past, most investigations of the key factors governing SN feedback strength have focused on single, isolated SNe (e.g., Chevalier 1974; Cioffi et al. 1988; Thornton et al. 1998; Iffrig & Hennebelle 2015). In reality, however, core collapse SNe are clustered in space and time: massive stars are born in clusters, and explode after $\sim 3-40$ Myr, before these stars can significantly disperse. The few studies that have looked at the feedback from multiple clustered, interacting SNe have found conflicting results. Some authors have found that clustering can *increase* momentum yields (Keller et al. 2014; Walch & Naab 2015; Gentry et al. 2017), whereas others have found that it can slightly *decrease* momentum yields (Kim & Ostriker 2015) or leave them essentially unchanged (e.g., Kim et al. 2017). This effect was particularly large for Keller et al. (2014) and Gentry et al. (2017), who found that clustering could increase the average momentum per SN to 5-10 times greater than the traditional yields for isolated SNe.

It has been suggested that the differences in results for clustered SN simulations could stem from the different levels of mixing in the simulations, from both physical and

[*] E-mail: egentry@ucsc.edu

nonphysical sources. Unfortunately, each recent simulation was idealised in significantly different ways, which makes it difficult for us to directly isolate which aspects were the primary drivers of the differences. Our goal in this paper is to identify the causes of the discrepancies between different published results, and resolve whether, when including appropriate physics, clustering does in fact lead to significant changes in the terminal momentum of supernova remnants.

One of the key issues that we investigate is dimensionality and resolution. We found that clustering produces an order of magnitude enhancement in momentum (Gentry et al. 2017), but these results were based on 1D spherically-symmetric simulations. Assuming spherical symmetry is potentially misleading because we know of fluid instabilities (such as the Vishniac instabilities and the Rayleigh Taylor instability) that affect SNR morphologies (Vishniac 1983, 1994; Mac Low & McCray 1988; Mac Low & Norman 1993; Krause et al. 2013; Fierlinger et al. 2016). Even small perturbations can be amplified and noticeably change key properties of SNR evolution. For isolated SN simulations, 1D and 3D simulations do not produce significantly different terminal momenta (e.g., Martizzi et al. 2015, Kim & Ostriker 2015, and Walch & Naab 2015 all find differences of less than 60% between 1D and 3D), but it is worth re-investigating the issue for clustered SNe. It could be that the longer time frame allows the instabilities to grow to have larger effects.

Conversely, our 1D simulations achieved much higher resolution than in any of the 3D simulations found in the literature. We found that the terminal momentum for clustered supernovae did not converge until we reached peak resolutions better than 0.1 pc, far higher than the resolutions of published 3D simulations. Moreover, we achieved this convergence only by using pseudo-Lagrangian methods that minimised numerical diffusion across the contact discontinuity at the inner edge of the superbubble, whereas many of the published 3D results are based on Eulerian methods that, for fronts advecting across the grid at high speed, are much more diffusive. Indeed, it is noteworthy that the one published 3D result that finds a significant momentum enhancement from clustering (Keller et al. 2014) uses a Lagrangian method, while all the papers reporting no enhancement are based on Eulerian methods. Clearly, given the conjoined issues of resolution and dimensionality, further investigation is warranted.

Since the suppression and enhancement of mixing is a key unknown for the feedback budget of clustered SNe, we also explore the role of magnetic fields, which may reduce the amount of physical mixing. Our interest in this possibility comes primarily from the example of cold fronts in galaxy clusters, where magnetic fields draped across a shock front have been used to explain the stability of these cold fronts against fluid instabilities (Vikhlinin et al. 2001; Markevitch & Vikhlinin 2007, although see also Churazov & Inogamov 2004 who show that magnetic fields might not be necessary for stabilising cold fronts).

In this paper, we first test the effects of bringing our simulations from 1D to 3D and carry out a 3D convergence study, and then we test the effects of adding magnetic fields into our 3D simulations. In Section 2, we discuss our computational methods. In Section 3 we discuss the results of our simulations, with Section 3.1 detailing the results of our

3D hydrodynamic simulations, and Section 3.2 detailing the effect of adding magnetic fields. In Section 4 we discuss our conclusions and compare to other works.

2   COMPUTATIONAL METHODS

For this work we repeat one of the 1D simulations from Gentry et al. (2017), and conduct 3D simulations of the same setup at a range of resolutions and including or excluding magnetic fields. For the most part our 1D simulations reuse the code developed by Gentry et al. (2017), with minor changes that we discuss in Section 2.1. In Section 2.2, we discuss the methodology for our 3D simulations, for which we use the GIZMO code (Hopkins 2015; Hopkins & Raives 2016). We use GIZMO for this work because it has a Lagrangian hydrodynamic solver; in our previous 1D simulations, we found that Lagrangian methods were more likely to converge for simulations of clustered SNe (Gentry et al. 2017).

2.1   1D simulation

The methods for our 1D simulation are very similar to those used in our previous work (Gentry et al. 2017), with only slight modifications. First, we disable the injection of pre-SN winds, because injecting small amounts of mass over extended periods is impractical at the resolutions we are able to achieve in the 3D simulations. Second, we initialize the ISM to be at an equilibrium temperature ($T \sim 340$ K or a specific internal energy of $e_{int} \sim 3.5 \times 10^{10}$ erg g^{-1} for an initial ISM density of $\rho_0 = 1.33$ $m_{\rm H}$ cm^{-3} and gas-phase metallicity of $Z = 0.02$, rather than $T \sim 15000$ K).[1] This simplifies the analysis, as changes in energy now only occur in feedback-affected gas. Furthermore the initial temperature makes little difference as the gas would otherwise rapidly cool to its equilibrium state (the 15000 K gas had a cooling time of a few kyr). Using these modified methods we re-ran the most-studied cluster from our previous work, one that had a stellar mass of $M_\star = 10^3$ $M_\odot$ (producing 11 SNe) and was embedded in an ISM of initial density $\rho_0 = 1.33$ $m_{\rm H}$ cm^{-3} and an initial gas-phase metallicity of $Z = 0.02$.[2] These changes allowed for more direct comparison with our 3D simulations, and do not affect our final conclusions.

The remainder of our methodology is identical to that of Gentry et al. (2017), which we summarise here for convenience. To generate a star cluster of given mass, we used the SLUG code (da Silva et al. 2012; da Silva et al. 2014; Krumholz et al. 2015) to realistically sample a Kroupa (2002) IMF of stars. We assume every star with an initial mass above $8 M_\odot$ explodes as a core collapse SN. The lifetimes of these massive stars are computed using the stellar evolution tracks of Ekström et al. (2012); the SN mass and metal yields are computed using the work of Woosley & Heger (2007) while all SNe are assumed to have

---

[1] Throughout this paper we will quote temperatures calculated by GRACKLE which accounts for temperature dependence in the mean molecular weight, μ (Smith et al. 2017).

[2] This cluster can be found in the tables produced by Gentry et al. (2017) using the id 25451948-485f-46fe-b87b-f4329d03b203.

an explosion energy of 10^{51} ergs. This cluster of stars is embedded in an initially-static, homogeneous ISM, with each SN occurring at the same location. The resulting superbubble is evolved using a 1D, spherically-symmetric, Lagrangian hydrodynamic solver first developed by Duffell (2016). Cells are split (merged) when they become sufficiently larger (smaller) than the average resolution. Metallicity-dependent cooling (assuming collisional ionization equilibrium) is included using GRACKLE (Smith et al. 2017). The simulation is evolved until the radial momentum reaches a maximum, at which point it is assumed that the superbubble mixes into the ISM.

2.2 3D simulations

Rather than adapt our 1D code to work in 3D, we instead chose to use the GIZMO simulation code (Hopkins 2015; Hopkins & Raives 2016), which includes a Lagrangian hydrodynamic solver with additional support for magnetohydrodynamics (MHD). For all of our runs, we used the Meshless Finite Mass solver on a periodic domain, while ignoring the effects of gravity. We assume the gas follows an ideal equation of state with a constant adiabatic index $\gamma = 5/3$, as in our 1D simulation. When including magnetic fields, we used GIZMO's standard solver for ideal MHD, as detailed in Hopkins & Raives (2016).

We modify the standard GIZMO code in two ways.[3] First, we added metallicity-dependent cooling using GRACKLE (Smith et al. 2017). Second, we inject SN ejecta, distributed in time, mass, and metal content using the same realization of SN properties as our 1D simulation. At the time of each SN, we inject new gas particles (each with mass approximately equal to the average existing particle mass) at random locations using a spherical gaussian kernel with a dispersion of 2 pc centred on the origin. Each new particle has equal mass and metallicity, which are determined by the SN ejecta yields.[4] For simulations which include magnetic fields, we linearly interpolate the magnetic field vector from nearby existing particles to the origin, and initialize the new feedback particles with that interpolated magnetic field vector. This procedure does not exactly preserve $\nabla \cdot \boldsymbol{B} = 0$, but GIZMO's divergence cleaning procedure rapidly damps away the non-solenoidal component of the field produced by our injection procedure.

We initialize the 3D simulations with the same static ($\boldsymbol{v} = 0$) homogeneous ISM as our 1D simulations ($\rho = 1.33 m_{\rm H}$ cm^{-3}, $Z = 0.02$ and $T \sim 340$ K). For simulations with magnetic fields, we include a homogeneous seed field, with $\boldsymbol{B} = (0, 0, 5)$ μG (identical to Iffrig & Hennebelle (2015)), corresponding to a plasma $\beta \approx 0.05$. We place particles of mass Δm on an evenly spaced grid of spacing Δx_0, which extends for a box size of L. (Particle locations are perturbed on the mpc scale in order to avoid pathologies in GIZMO's density

solver caused by the symmetric grid.) In Table 1, we present the parameters of the initial conditions. We ran each 3D simulation for 40 Myr, by which point the radial momentum, the quantity of primary interest for our study, had stabilised.

3 RESULTS

In Table 2 we provide a summary of the key numeric results of each simulation. First, we extract the time of maximum momentum after the last SN (only accurate within about 0.5 Myr.) At that time we extract the effective radius of the region affected by the bubble (particles with speeds greater than 1 m s^{-1}):

$$R_{\rm eff} = \left(\frac{3}{4\pi} \sum_{i:|\boldsymbol{v}_i|>1{\rm m\ s}^{-1}} {\rm Volume}_i \right)^{1/3} \qquad (1)$$

and the total mass of those particles, $M_{\rm affected}$.[5] Next we extract the kinetic energy $E_{\rm kin}$ and the change in the internal energy $\Delta E_{\rm int}$ of the entire domain (which should be approximately equal to the values for the bubble-affected region). Finally, we extract the radial momentum using two approaches: one by simply measuring the radial momentum at the same time as the previous quantities (denoted $p_{\rm max}$), and another using a "ratchet" approach justified and explained in Section 3.2 (denoted $p_{\rm ratchet}$).

In the following subsections we discuss the results in greater detail. First, we compare our 1D and 3D results in Section 3.1. Second, we look at the effect of including magnetic fields in our 3D simulations in Section 3.2.

3.1 Hydrodynamic results and convergence study

In Figure 1 we show the radial momentum evolution of our median-resolution 3D simulation without MHD (3D_20_HD), and compare it to our 1D simulation. As can be seen in that figure, we observe a significant difference between the final momenta in our 1D and 3D simulations. While our 1D simulation of clustered SNe shows a large gain in momentum per SN compared to the isolated SN yield,[6] our 3D simulation shows no such gain. That discrepancy needs to be addressed.

This cannot be explained just by the fact that the 3D simulation has a lower initial resolution. In our previous work we tested the resolution dependence in our 1D simulations, and found that even with an initial spatial resolution of 5 pc, we still measured a terminal momentum yield roughly 10 times higher than what we find in our 3D simulation here as long as we ran our code in pseudo-Lagrangian rather than Eulerian mode (Gentry et al. 2017, Figure 14). So the problem is not convergence in our 1D simulation, but we have not yet shown whether our 3D results are converged.

---

[3] Our modifications of GIZMO and our analysis routines can be found at: github.com/egentry/gizmo-clustered-SNe.

[4] While this approach leads to a well-sampled injection kernel at our higher resolutions, the kernel is only sampled by about 5 new particles for each SN in our lowest resolution run, 3D_40_HD. This undersampling is not ideal and might slightly alter the bubble's evolution, but the stochasticity this introduces does not appear to affect our conclusions.

[5] The exact velocity threshold is somewhat arbitrary, leading to a roughly 10% uncertainty in the affected mass depending on the chosen threshold.
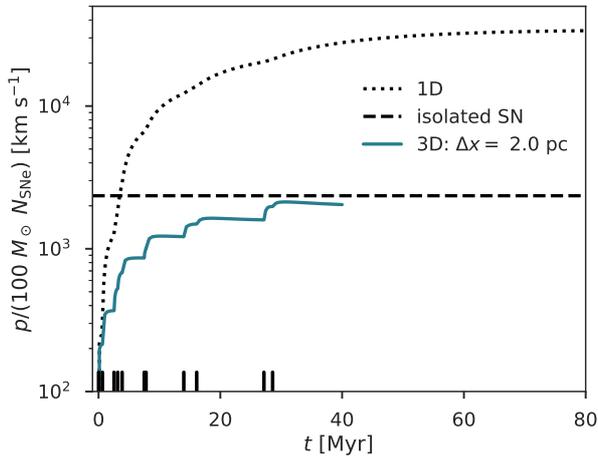
[6] We estimate the isolated SN momentum yield, 2.4×10^5 $M_\odot$ km s^{-1}, using the first SN of our 3D_20_HD simulation, although all of our 3D simulations would give the same value within a few percent. This is approximately consistent with previous single SN simulations (e.g. Martizzi et al. 2015; Kim & Ostriker 2015).

Table 1. Initial Conditions. The mass resolution Δm is not included for the 1D run, as it is neither constant in space nor time.

Name	1D/3D	$B_{z,0}$ (μG)	β	Δx_0 (pc)	Δm ($M_\odot$)	L (pc)
1D_06_HD	1D	0	∞	0.6		1200
3D_10_HD	3D	0	∞	1.0	0.03	600
3D_20_HD	3D	0	∞	2.0	0.26	600
3D_40_HD	3D	0	∞	4.0	2.10	600
3D_20_MHD	3D	5	0.05	2.0	0.26	1200

Table 2. Results.

Name	N_{SNe}	t (Myr)	R_{eff} (pc)	$M_{affected}$ ($10^6 M_\odot$)	p_{max}/N_{SNe} ($100 M_\odot$ km s^{-1})	$p_{ratchet}/N_{SNe}$ ($100 M_\odot$ km s^{-1})	E_{kin} (10^{49} erg)	ΔE_{int} (10^{49} erg)
1D_06_HD	11	94.8	552	23.2	33978	33978	65.0	26.3
3D_10_HD	11	30.7	218	1.7	2425	2474	8.7	0.8
3D_20_HD	11	30.7	200	1.5	2128	2182	7.4	0.8
3D_40_HD	11	31.7	209	1.8	2007	2039	6.8	7.0
3D_20_MHD	11	29.6	423	10.5	1213	2418	12.6	13.1

Figure 1. Comparison of the momentum evolution of 1D and 3D simulations of the same cluster (simulations 1D_06_HD and 3D_20_HD respectively). The "isolated SN" value is estimated using the first SN of the 3D_20_HD simulation, although it does not vary substantially between any of our 3D simulations.


Figure 2. Resolution study of our 3D simulations.

To test for convergence in our 3D simulations, we compare our 3 simulations which differ only in resolution (3D_10_HD, 3D_20_HD and 3D_40_HD); in Figure 2, we show the momentum evolution of each simulation. From that figure, we conclude that our 3D simulations do not appear converged, unlike our 1D simulations. The momentum yield is increasing monotonically as we increase the resolution, so our 3D results are converging in the direction of our 1D results, but even at the highest resolution we can afford the momentum yield remains well below the 1D case. Thus we do not know if the 3D results would converge to the same value as the 1D case, even with infinite resolution.

To get a hint at why our 3D simulations are behaving differently from our 1D simulation, we can look at the morphology of the density field for a slice through the centre of simulation 3D_20_HD at t = 7.53 Myr shown in Figure 3. In that figure we clearly see anisotropies in the bubble's

dense shell, which cannot not be captured by our 1D code that enforces spherical symmetry. These anisotropies are the result of physical instabilities (such as the Vishniac instabilities and the Rayleigh Taylor instability) amplifying numerical inhomogeneities in the background ISM (Vishniac 1983, 1994; Mac Low & McCray 1988; Mac Low & Norman 1993; Krause et al. 2013; Fierlinger et al. 2016; Yadav et al. 2017). Many similar numeric studies of superbubbles within a homogeneous ISM were also unable to achieved converged results for key quantities (e.g., Krause et al. 2013; Fierlinger et al. 2016; Yadav et al. 2017), so it is not surprising that we are unable to achieve converged results. Using Lagrangian methods may help, but they are not enough even at these resolutions.

These physical instabilities provide a direct mechanism for affecting the shell morphology, but the mechanism by which this affects the momentum yield is less obvious. In general the decreased momentum is assumed to be due to the breakdown of the contact discontinuity (see e.g., Fierlinger et al. 2016), which in the case of perfect spherical symmetry would separate the bubble interior (which dom-
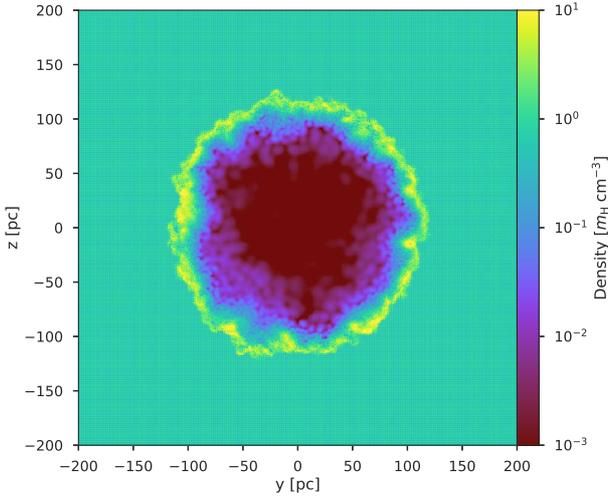
Figure 3. Reference density slice of the 3D simulation with median resolution (`3D_20_HD`) at $t = 7.53$ Myr, approximately 0.03 Myr after the 6th SN.

inates the thermal energy of the system) from the bubble shell (which dominates the mass of the system). If there is little transport or mixing between these two components, then the cooling time remains relatively long: the thermal energy cannot get to the efficiently-cooling material. But when forms of mixing and transport are added (either implicitly due to numerical errors, or explicitly through prescriptions for additional physics (e.g. Duffell 2016; Fierlinger et al. 2016)) the cooling time can decrease dramatically, sapping energy earlier, leading to a decreased momentum yield.

To understand how much mixing is present in our 3D simulations, and how it depends on resolution, we turn to density-temperature phase diagrams which are shown in Figure 4. These phase diagrams correspond to a time soon after the 6th SN, with a delay long enough to allow the injected energy to spread throughout the bubble but sufficiently short to avoid significant energy losses due to cooling in any simulation. (All 3D simulations have about 1.1×10^{51} ergs more total energy than the start of the simulations, but `1D_06_HD` retains more energy from previous SNe, and contains about 2.7×10^{51} ergs of total energy relative to the simulation start.) When we look at the mass-weighted phase diagram for our highest resolution simulation (`3D_10_HD`), we see that the mass is dominated by a cold dense shell, with a minority of mass in less-dense warm and hot phases ($> 10^3$ K).[7] Even when we vary the resolution, we only find negligible changes in the fraction of mass in the cold phase; the cold phase ($T < 10^3$ K) contributes 99.2% of the mass in every simulation of our resolution study (`3D_10_HD`, `3D_20_HD` and `3D_40_HD`). What does change is the density and temperature distribution of the warm/hot gas ($T > 10^3$ K). As we increase the resolution, the warm/hot gas shifts to lower and lower densities. This effect is very apparent for gas near the peak of the cooling curve at $T \sim 10^5$ K, which has a

mass-weighted median density of $\sim 10^{-0.5}$ m_H cm^{-3} in our lowest resolution run, and $\sim 10^{-2}$ m_H cm^{-3} in our highest resolution run.

This has a significant impact on the overall cooling times of the simulations. The right column of Figure 4 shows that while the cold, dense phase dominates the mass, the minority of mass in the warm/hot phases dominates the cooling rate. This is important because resolution primarily affects these warm/hot phases, and it affects those phases by shifting them to higher densities at lower resolution, causing each particle to become more efficient at cooling. This results in significantly shorter cooling times: from 27 Myr at the best resolution to 0.3 Myr at the worst resolution, nearly two orders of magnitude difference. This increase primarily occurs in the warm/hot phases; at all resolutions gas warmer than 10^3 K constitutes slightly less than 1% of the total mass, but this mass is responsible for 81% of the cooling at our highest resolution, and $> 99\%$ of the cooling at our lowest resolution.

When we look at the phase diagrams for our 1D simulation, we see significant differences in the distributions of mass and cooling rate, leading to the very different behavior of the 1D simulation. In particular, the 1D simulation completely lacks material at intermediate densities ($\sim 10^{-2} - 10^0$ m_H cm^{-3}) due to how well the 1D simulation retains the contact discontinuity. The diffuse bubble-dense shell transition occurs within only a few cells, and the entire dense shell is resolved by just 5-10 cells. In our 3D simulations, these intermediate densities contribute a negligible amount of mass, but are responsible for much of the cooling. Without this intermediate phase material, almost all of the cooling in the 1D simulation occurs in the dense shell, powered by the small amount of thermal energy that is leaking out of the bubble interior. While this appears to provide a cooling time faster than the highest resolution 3D simulations, that is only because the 3D simulations have not yet reached their peak cooling rates after this SN. (Waiting until the time of peak cooling would mean each 3D simulation would no longer contain roughly the same amount of energy. It would not affect the conclusions on the resolution dependence of our results, but it would complicate the analysis of the phase diagrams.) When we compare the minimum overall cooling times after this SN, we find that `1D_06_HD` indeed has a longer cooling time than even `3D_10_HD` (0.57 Myr and 0.24 Myr respectively).

In summary, we are unable to achieve converged momentum yields in our 3D simulations due to physical instabilities that lead to resolution-dependent numerical mixing between phases, causing resolution-dependence in the cooling rates of our superbubbles. Due to this numerical mixing, the cooling rates are almost certainly still being overestimated, and thus the terminal momentum underestimated, even in our highest resolution 3D Lagrangian simulations. Our 1D simulations completely lack intermediate-temperature material, and thus likely provide an upper limit on the terminal momentum, but we do not know how close to this upper limit the 3D simulations would get if we could run at much higher resolution.

---

[7] We also see a negligible amount of mass at unusually low temperatures, < 100 K. These particles are SN ejecta, which have very high metallicities that have been frozen-in due to the Lagrangian nature of the code.
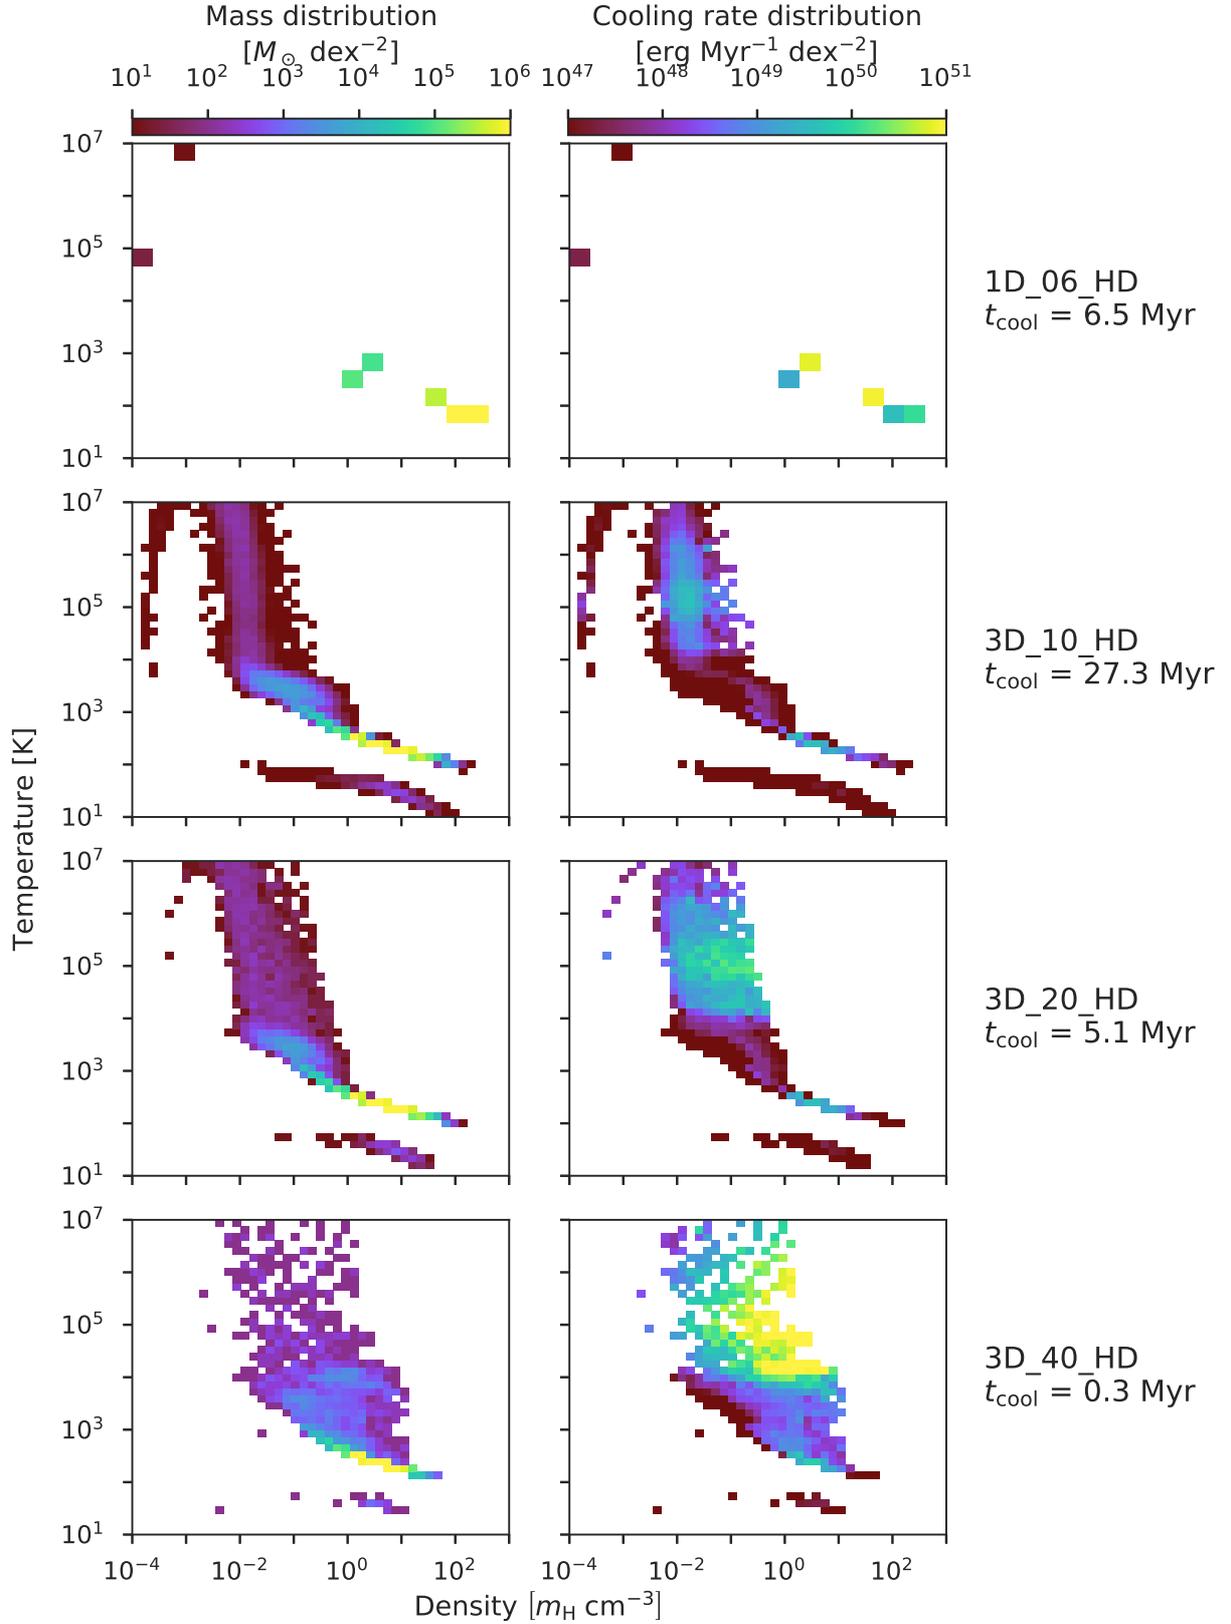
Figure 4. Phase diagrams for our 3D simulations at $t = 7.53$ Myr, about 0.03 Myr after the 6th SN, when all simulations still retain almost all of the energy from the most recent SN. The left column shows the distribution of mass within temperature-density space, and the right column shows the cooling rate distribution within the space. The rows show the non-magnetised simulations with initial resolution worsening from top to bottom. To the right of each row, we give the cooling time of each simulation, $t_{\rm cool} \equiv E_{\rm int}/\dot{E}_{\rm cool}$, for reference.
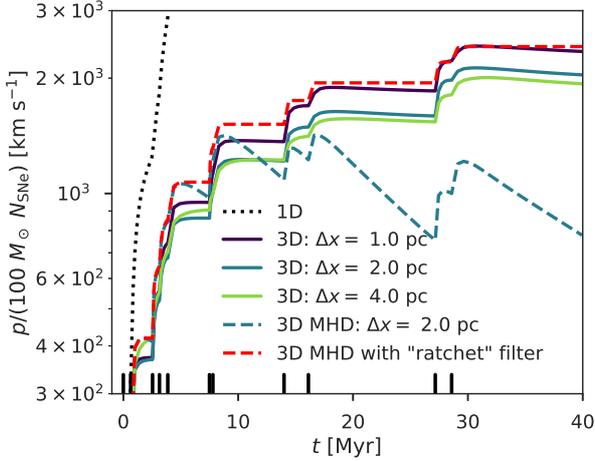
Figure 5. Same as Figure 2, except now including the momentum evolution of our MHD simulation (`3D_20_MHD`; blue dashed curve), as well its "ratchet"-filtered momentum evolution, p_{ratchet} (red dashed curve).
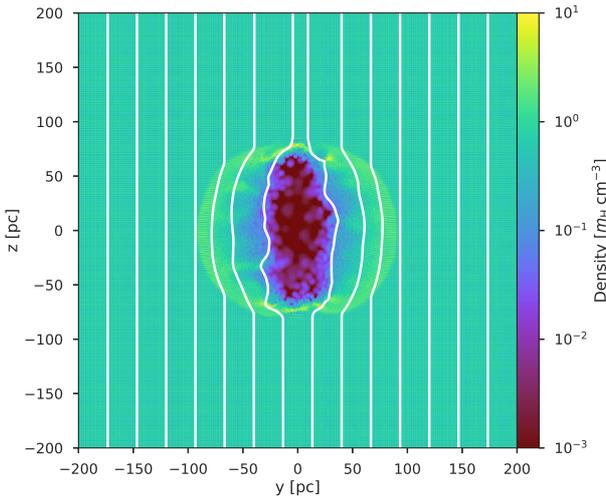
Figure 6. Same as Figure 3, except now for simulation `3D_20_MHD` with *approximate* magnetic field lines overplotted, and at an earlier time ($t = 2.56$ Myr; approximately 0.02 Myr after the 3rd SN).

3.2 Magnetic fields

In Section 3.1 we showed that our numerical methods and resolution are not sufficient to achieve converged values of final radial momentum and other key parameters due to physical instabilities that develop within the superbubble shell. As described in Section 1, we expect that magnetic fields might affect the growth of physical instabilities, so we also run an MHD simulation as described in Section 2.2 to test the impact of magnetic fields on the final momentum. While the more-standard method of extracting momentum, p_{max}, quoted in Table 2 appears to show that the inclusion of magnetic fields significantly decreases the final momentum, in this subsection we show that that method for estimating the asymptotic momentum (finding the maximum

momentum following the last SN) is an oversimplification for simulations with magnetic fields. When we better isolate the momentum added by SNe, we find that adding magnetic fields can actually increase the momentum yield at fixed resolution. Indeed, our $\Delta x = 2.0$ pc MHD run produces a larger momentum injection than our $\Delta x = 1.0$ pc HD run.

First, to illustrate why the interpretation of the MHD simulation is more complex, in Figure 5 we compare its momentum evolution to those of the non-magnetised simulations. The MHD simulation initially shows an *increased* momentum yield relative to the corresponding simulation without magnetic fields at the same resolution (`3D_20_MHD`), but then the momentum decreases due to magnetic tension forces. The reason for this is obvious if we examine a density slice at an earlier time, [8] as shown in Figure 6: the expanding shell bends magnetic field lines outward, and the field lines exert a corresponding magnetic tension that reduces the radial momentum of the expanding shell. This effect is so strong that the momentum peaks after just 7 SNe; the remaining 4 SNe clearly add momentum but not enough to overcome the steady decline.

Due to this effect, the quantity p_{max} (the maximum momentum after the last SN) that we have used to characterise the hydrodynamic simulations is somewhat misleading, since our goal is to study the momentum injected by supernovae, not the combined effects of supernovae and magnetic confinement. To avoid this, we define an alternative quantity p_{ratchet}. To compute this quantity we sum any positive changes in radial momentum between snapshots, while ignoring any negative changes. We plot p_{ratchet} in Figure 5, and report the final value in Table 2. As expected, for the non-magnetic runs p_{ratchet} and p_{max} are essentially the same, and thus examining p_{ratchet} allows us to make an apples-to-apples comparison between the magnetic and non-magnetic results.

This comparison is revealing, in that it shows that our simulation with magnetic fields (`3D_20_MHD`) injects *about 10% more* momentum than the analogous simulations without magnetic fields (`3D_20_HD`). To understand why this occurs, we can compare phase diagrams for the resolution-matched magnetised and non-magnetised runs in Figure 7, shown at the same time ($t = 2.56$ Myr) as Figure 6. There we see that magnetic fields have an effect similar to that of increasing resolution in Figure 4: both suppress the growth of fluid instabilities, causing the material near the peak of the cooling curve to stay at lower densities where it cools less efficiently. For gas near the peak of the cooling curve, $10^4 - 10^6$ K, the median density of the non-magnetised run is 2.6 m_{H} cm^{-3}, while in the magnetised run it is 1.3 m_{H} cm^{-3}. As a result the overall cooling time is about 2 times longer in the MHD run. Thus by suppressing the growth of instabilities, the inclusion of magnetic fields results in a longer overall cooling time and a higher yield of momentum.

---

[8] We chose to look at an earlier snapshot, when the magnetisation has only perturbed the bubble structure, rather than the later time shown in Figure 3, when the magnetisation would have caused a strong, non-linear change in the structure which could not be treated as a perturbation. In both cases the magnetic tension is present, but the earlier time makes it more straightforward to compare to the non-magnetised runs.
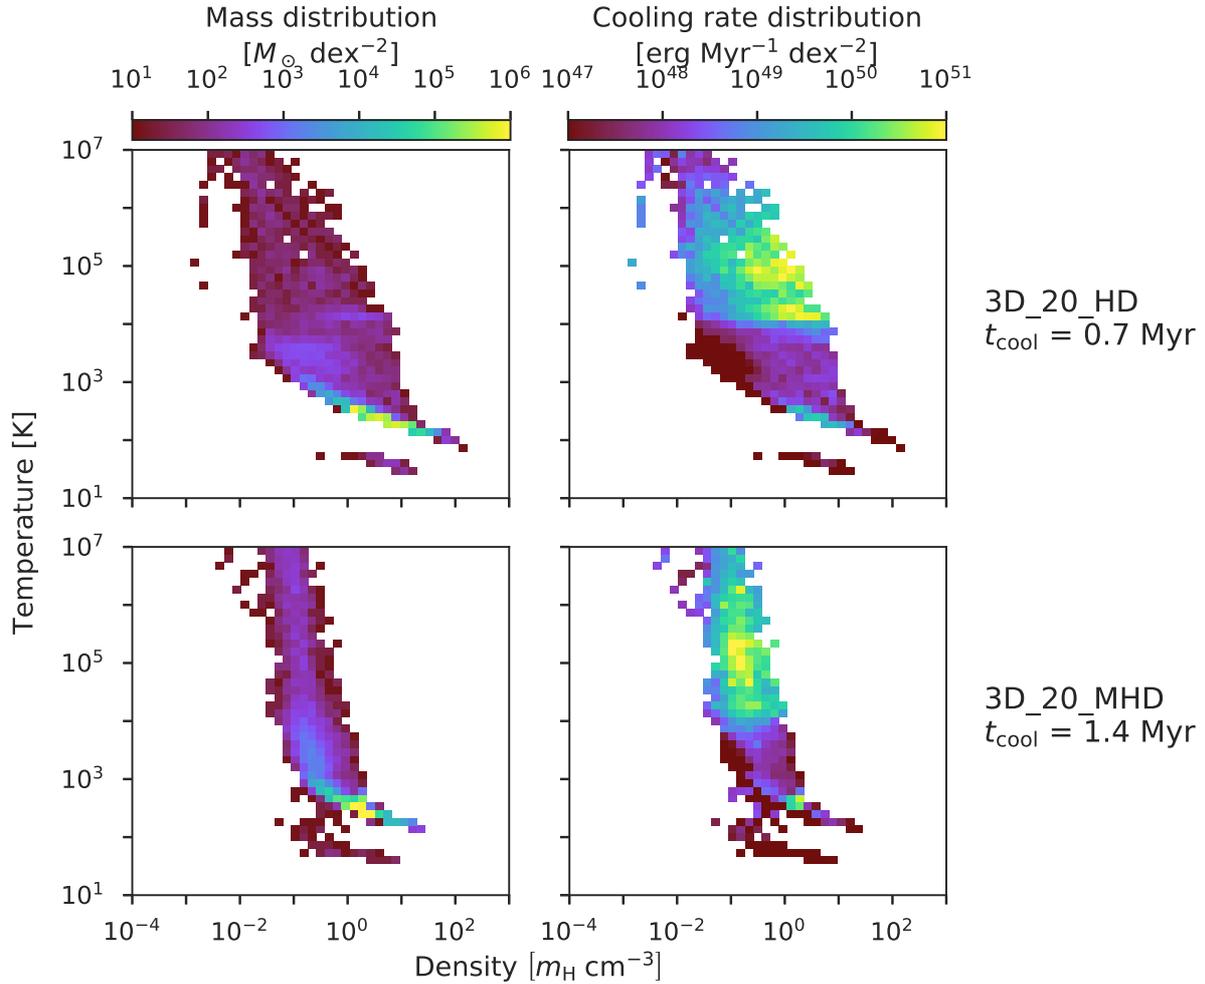
Figure 7. Same as Figure 4, except at the earlier time shown in Figure 6 (t = 2.56 Myr), and only showing the magnetised simulation (`3D_20_MHD`) and the corresponding non-magnetised simulation with the same resolution (`3D_20_HD`).

4   CONCLUSIONS

In this paper we revisit the question of whether clustering of supernovae leads to significant differences in the amount of momentum and kinetic energy that supernova remnants deliver to the ISM. This question is strongly debated in the literature, with published results offering a menu of answers that includes a decrease in momentum due to clustering (e.g., Kim & Ostriker 2015), no effect (e.g., Kim et al. 2017), a modest increase (e.g., Walch & Naab 2015), or a large increase (e.g., Keller et al. 2014; Gentry et al. 2017). We investigate whether this discrepancy in results is due to numerical or physical effects, and to what extent it might depend on whether the flow is modelled as magnetised or non-magnetised.

Our results offer some encouragement and also some unhappy news regarding the prospects for treating supernova feedback in galactic and cosmological simulations. The encouraging aspect of our findings is that we have identified the likely cause of the discrepancy between the published results. We find that the key physical mechanism driving the differences between our runs, and almost certainly between other published results, is the rate of mixing across the contact discontinuity between the hot interior of a superbubble

and the cool gas in the shell around it. Our 1D Lagrangian results (Gentry et al. 2017) maintain the contact discontinuity nearly perfectly, and this explains why they produce large gains in terminal momentum per supernova due to clustering. However, these results likely represent an upper limit on the momentum gain, because they do not properly capture instabilities that mix mass and thermal energy across the contact discontinuity.

In 3D, instabilities produce intermediate temperature gas that radiates rapidly and saps the energy of the superbubble, lowering the terminal momentum. However, we are unable to obtain a converged result for the strength of this effect, because we find that the terminal momentum continues to increase with resolution even at the highest resolution (1 pc initial linear resolution, 0.03 $M_\odot$ mass resolution) that we achieve. The physical cause of this effect is clear: as we increase the resolution, we find that the mean density and total mass of gas near the peak of the cooling curve continuously decreases, and this leads to a monotonic decrease in the amount of energy lost to radiation. Consequently, we are forced to conclude that even at this resolution mixing across the contact discontinuity is dominated by numerical rather than physical diffusion, and that our estimate of the

momentum per supernova is only a lower limit. Our tests with magnetic fields reinforce this conclusion. We find that magnetic fields suppress the growth rate of instabilities, so that an MHD run injects more momentum per supernova than a non-magnetised simulation at the same resolution. In the real ISM, magnetic fields are doubtless present, so this effect cannot be neglected.

Our findings cloud the prospects for obtaining a good first-principles estimate of the true supernova momentum, and cast doubts on previous claims of convergence in the literature. Our peak resolution is higher than that achieved in previous 3D simulations, and our Lagrangian method suffers much less numerical diffusion across the contact discontinuity than the Eulerian methods used in most previous works, yet we are unable to reach convergence. Given this finding, it seems reasonable to investigate whether apparent convergence in Eulerian codes (e.g., Kim et al. 2017) is an illusion due to the fact that these codes suffer from high levels of numerical mixing when contact discontinuities advect supersonically across the grid, and that increasing the resolution reduces this diffusivity very slowly or not at all (e.g., Hopkins 2015). We are forced to conclude that *the true momentum yield from clustered supernovae remains uncertain at the order of magnitude level.*

We conclude by noting that we have not thus far investigated the effects of using a realistically-turbulent, multiphase interstellar medium. The presence of density inhomogeneities could well lead to higher rates of mixing across the contact discontinuity, and thus a reduction in the supernova momentum yield. However, we urge extreme caution in interpreting the results of any investigations of these phenomena, since we have shown that even state of the art simulation methods operating at the highest affordable resolutions cannot reach convergence in what should be substantially simpler problems. It is conceivable that the more complex density field of a realistic ISM might make it easier to reach convergence, but such a hope would need to be demonstrated rigorously using convergence studies in multiple numerical methods. Moreover, any such simulations would need to include magnetic fields in order to produce credible results for the real ISM.

ACKNOWLEDGEMENTS

REFERENCES

Chevalier R. A., 1974, ApJ, 188, 501

Churazov E., Inogamov N., 2004, MNRAS, 350, L52
Cioffi D. F., McKee C. F., Bertschinger E., 1988, ApJ, 334, 252
Creasey P., Theuns T., Bower R. G., 2013, MNRAS, 429, 1922
Dekel A., Krumholz M. R., 2013, MNRAS, 432, 455
Dekel A., Silk J., 1986, ApJ, 303, 39
Duffell P. C., 2016, ApJ, 821, 76
Ekström S., et al., 2012, A&A, 537, A146
Faucher-Giguère C.-A., Quataert E., Hopkins P. F., 2013, MNRAS, 433, 1970
Fierlinger K. M., Burkert A., Ntormousi E., Fierlinger P., Schartmann M., Ballone A., Krause M. G. H., Diehl R., 2016, MNRAS, 456, 710
Gentry E. S., Krumholz M. R., Dekel A., Madau P., 2017, MNRAS, 465, 2471
Hopkins P. F., 2015, MNRAS, 450, 53
Hopkins P. F., Raives M. J., 2016, MNRAS, 455, 51
Hopkins P. F., Quataert E., Murray N., 2012, MNRAS, 421, 3522
Iffrig O., Hennebelle P., 2015, A&A, 576, A95
Jenkins E. B., Tripp T. M., 2011, ApJ, 734, 65
Keller B. W., Wadsley J., Benincasa S. M., Couchman H. M. P., 2014, MNRAS, 442, 3013
Kim C.-G., Ostriker E. C., 2015, ApJ, 802, 99
Kim C.-G., Kim W.-T., Ostriker E. C., 2011, ApJ, 743, 25
Kim C.-G., Ostriker E. C., Raileanu R., 2017, ApJ, 834, 25
Krause M., Fierlinger K., Diehl R., Burkert A., Voss R., Ziegler U., 2013, A&A, 550, A49
Kroupa P., 2002, Science, 295, 82
Krumholz M. R., Fumagalli M., da Silva R. L., Rendahl T., Parra J., 2015, MNRAS, 452, 1447
Mac Low M.-M., McCray R., 1988, ApJ, 324, 776
Mac Low M.-M., Norman M. L., 1993, ApJ, 407, 207
Markevitch M., Vikhlinin A., 2007, Phys. Rep., 443, 1
Martizzi D., Faucher-Giguère C.-A., Quataert E., 2015, MNRAS, 450, 504
McKee C. F., Ostriker J. P., 1977, ApJ, 218, 148
Murray N., Quataert E., Thompson T. A., 2005, ApJ, 618, 569
Ostriker E. C., Shetty R., 2011, ApJ, 731, 41
Smith B. D., et al., 2017, MNRAS, 466, 2217
Thompson T. A., Krumholz M. R., 2016, MNRAS, 455, 334
Thornton K., Gaudlitz M., Janka H.-T., Steinmetz M., 1998, ApJ, 500, 95
Vikhlinin A., Markevitch M., Murray S. S., 2001, ApJ, 549, L47
Vishniac E. T., 1983, ApJ, 274, 152
Vishniac E. T., 1994, ApJ, 428, 186
Walch S., Naab T., 2015, MNRAS, 451, 2757
Woosley S. E., Heger A., 2007, Phys. Rep., 442, 269
Yadav N., Mukherjee D., Sharma P., Nath B. B., 2017, MNRAS, 465, 1720
da Silva R. L., Fumagalli M., Krumholz M., 2012, ApJ, 745, 145
da Silva R. L., Fumagalli M., Krumholz M. R., 2014, MNRAS, 444, 3275

This paper has been typeset from a TeX/LaTeX file prepared by the author.