Downloaded from https://academic.oup.com/mnras/article/506/2/2199/6276745 by University of Texas at Austin user on 01 June 2022

MNRAS **506**, 2199–2231 (2021) Advance Access publication 2021 May 17

STARFORGE: Towards a comprehensive numerical model of star cluster formation and feedback

Michael Y. Grudić [®], ¹★ Dávid Guszejnov [®], ² Philip F. Hopkins [®], ³ Stella S. R. Offner [®] and Claude-André Faucher-Giguère [®]1

Accepted 2021 May 2. Received 2021 April 16; in original form 2020 October 25

ABSTRACT

We present STARFORGE (STAR FORmation in Gaseous Environments): a new numerical framework for 3D radiation magnetohydrodynamic (MHD) simulations of star formation that simultaneously follow the formation, accretion, evolution, and dynamics of individual stars in massive giant molecular clouds (GMCs), while accounting for stellar feedback, including jets, radiative heating and momentum, stellar winds, and supernovae. We use the GIZMO code with the MFM mesh-free Lagrangian MHD method, augmented with new algorithms for gravity, time-stepping, sink particle formation and accretion, stellar dynamics, and feedback coupling. We survey a wide range of numerical parameters/prescriptions for sink formation and accretion and find very small variations in star formation history and the IMF (except for intentionally unphysical variations). Modules for mass-injecting feedback (winds, SNe, and jets) inject new gas elements on the fly, eliminating the lack of resolution in diffuse feedback cavities otherwise inherent in Lagrangian methods. The treatment of radiation uses GIZMO's radiative transfer solver to track five frequency bands (IR, optical, NUV, FUV, ionizing), coupling direct stellar emission and dust emission with gas heating and radiation pressure terms. We demonstrate accurate solutions for SNe, winds, and radiation in problems with known similarity solutions, and show that our jet module is robust to resolution and numerical details, and agrees well with previous AMR simulations. STARFORGE can scale up to massive (>10⁵ M_{\odot}) GMCs on current supercomputers while predicting the stellar (\gtrsim 0.1 M_{\odot}) range of the IMF, permitting simulations of both high- and low-mass cluster formation in a wide range of conditions.

Key words: MHD – radiative transfer – turbulence – methods: numerical – stars: formation – ISM: general.

1 INTRODUCTION

Many physical mechanisms are important in star formation (SF). It is initiated by the formation of radiatively cooled, gravitationally unstable cores of gas and dust from magnetized, supersonic, turbulent flows found in giant molecular clouds (GMCs; Larson 1981; Mac Low & Klessen 2004; McKee & Ostriker 2007; Girichidis et al. 2020). These cores collapse to protostars, and once formed, protostars and stars influence the surrounding gas flow in via feedback: the injection of mass, momentum, and energy into the interstellar medium (ISM) in the form of radiation, accretion-powered collimated bipolar outflows (hereafter simply 'jets'), radiatively driven stellar winds, and supernova (SN) explosions, which may ultimately limit the total stellar mass that can form. The accretion of individual stars is eventually truncated by feedback, gas exhaustion, or dynamical interactions with gas clumps or other stars, setting their final masses (Krause et al. 2020). Hence, the problem of SF is an intricate, tightly-coupled interaction of gravity, magnetohydrodynamics (MHD), atomic and molecular physics, radiation, stellar physics, and feedback.

A basic requirement of any star formation theory is to explain the hallmark phenomena of SF, including the stellar initial mass function (IMF), the (in-)efficiency of SF, and the properties of stellar clusters and associations (Krumholz 2014). These phenomena must emerge from the various processes at work in GMCs, so it is important to disentangle the physics' respective roles. This has yet to be accomplished, partly because the wide range of length-scales and multitude of physics involved make SF very challenging to model.

1.1 Requirements for a complete dynamical model of star formation and feedback

While some progress has been made with simpler models that consider only e.g. turbulence and gravity (Padoan, Nordlund & Jones 1997; Padoan & Nordlund 2002; Krumholz, McKee & Klein 2005; Hennebelle & Chabrier 2008; Padoan & Nordlund 2011; Hopkins 2012; Hennebelle & Chabrier 2013), other physics are likely to be important. In particular, *feedback* is important for understanding the end-point of star formation (the disruption of GMCs), and its implications for other questions such as the IMF and stellar multiplicity have only begun to be explored. Many analytic and semi-analytic calculations of the effects of feedback in GMCs have been performed

¹CIERA and Department of Physics and Astronomy, Northwestern University, 1800 Sherman Ave, Evanston, IL 60201, USA

²Department of Astronomy, The University of Texas at Austin, TX 78712, USA

³TAPIR, Mailcode 350-17, California Institute of Technology, Pasadena, CA 91125, USA

 $[\]hbox{* E-mail: mike.grudic@northwestern.edu}\\$

(for reviews, see McKee & Ostriker 2007; Krumholz, McKee & Bland-Hawthorn 2019), yielding useful dimensional arguments and analytic insights. But GMCs are complex, turbulent, inherently three-dimensional entities that evolve on their internal crossing time-scale (Larson 1981; Mac Low & Klessen 2004). Thus, even under the gross simplification of treating GMCs as isolated entities (i.e. neglecting galactic environment), (semi-)analytic predictions inevitably hinge on many highly uncertain assumptions. With so much parameter freedom it is difficult to say whether a given model is correct for the *right reasons*, limiting physical insight and ultimately predictive power. Direct numerical simulations of star formation are a necessary tool to resolve these uncertainties.

In the past two decades, great progress has been made incorporating stellar feedback into direct numerical simulations of star-forming GMCs (for reviews see Dale 2015; Krumholz et al. 2019). But these studies have shown that further progress on the key questions of star formation requires next-generation simulations that do all of the following:

- (i) Resolve individual star formation: Many simulations of star cluster formation do not attempt to resolve the formation and ongoing accretion of individual stars across the entire stellar mass range, instead relying on a sub-grid SF prescription that either enforces a certain underlying IMF directly or is fine-tuned to recover the observed one (Colín, Vázquez-Semadeni & Gómez 2013; Howard, Pudritz & Harris 2017; Kim et al. 2017; Sormani et al. 2017; Su et al. 2018; Grudić et al. 2018a; He, Ricotti & Geen 2019; Lahén et al. 2019; Wall et al. 2020). But there is an infinite number of ways to do this, each with different implicit assumptions about how star formation works, and the choice of prescription can have a major effect upon simulation results (Grudić & Hopkins 2019), limiting predictive power. Simulations should ideally attempt to resolve the formation and accretion of individual stars (or sink particles), and to recover a realistic IMF self-consistently from physical (not numerical) processes. This is obviously necessary anyway if one wishes to study the physical origins of the IMF and stellar multiplicity.
- (ii) Follow stellar dynamics: SF simulations that do not integrate stellar orbits explicitly generally discretize the stellar mass formed into a collisionless fluid represented by gravitationally softened particles (e.g. Grudić et al. 2018a; Lahén et al. 2019; Li et al. 2019), which can produce qualitatively correct star cluster density profiles (Grudić et al. 2018b; Lahén et al. 2020), but have the severe limitation that the collisionless description (and phase-space density conservation) breaks down on mass scales $M_{\star} \lesssim 100\,\mathrm{M}_{\odot}$, so if cluster formation is a hierarchical assembly from smaller masses (e.g. Bonnell, Bate & Vine 2003), then individual stellar dynamics is always important at some stage in the process. A simulation must also treat dynamics on the scale of binary separations to accurately predict stellar multiplicity, let alone phenomena such as common disc accretion (e.g. Lee et al. 2019; Muñoz, Miranda & Lai 2019; Duffell et al. 2020).
- (iii) Follow MHD, chemistry, and cooling: Obviously, following the dynamics of GMCs, star formation, and accretion requires gas-dynamical simulations, and stars cannot form if gas cannot radiatively cool. Moreover, the ISM is magnetized, and this fact can easily have important implications for star formation. The magnetic field can act as an additional source of support, potentially stabilizing otherwise-unstable cores (Chandrasekhar 1951; Mouschovias & Spitzer 1976), affecting the IMF (Price & Bate 2007; Guszejnov et al. 2020b), the rate of star formation (Federrath & Klessen 2012; Federrath 2015), and altering the pressure balance, morphology, growth of

instabilities, and transport of energy in feedback bubbles (Krumholz, Stone & Gardiner 2007; Offner & Liu 2018; Krumholz & Federrath 2019).

- (iv) Scale up to massive GMCs: Current star formation simulations that do both (i) and (ii) have focused upon lower mass systems, simulating gas masses of $100\text{--}1000\,\text{M}_\odot$ (Federrath 2015; Cunningham et al. 2018; Jones & Bate 2018; Lee & Hennebelle 2018; Li, Klein & McKee 2018; Wurster, Bate & Price 2019; Colman & Teyssier 2020), producing $\sim\!10\text{--}100\,\text{M}_\odot$ in stars. Low-mass clusters are important to model, as they can be readily compared to well-studied sites of star formation in the Solar neighbourhood (e.g. Evans, Heiderman & Vutisalchavakul 2014), but the overwhelming majority of SF in our Galaxy occurs in massive complexes with gas mass $\gg\!10^5\,\text{M}_\odot$ (McKee & Williams 1997; Murray & Rahman 2010). Simulated low-mass clusters are also less likely to host massive ($\gtrsim\!10\,\text{M}_\odot$) stars, and hence cannot be used to study massive SF. 1
- (v) Account for all major feedback channels: 3D MHD simulations of SF have not generally incorporated all known dynamically important feedback mechanisms (jets, winds, full-spectrum radiation, and SNe) acting in concert. A comprehensive treatment of feedback is needed because different feedback channels are effective on different scales, and can interact non-linearly. For example, direct radiation pressure from a massive star is ineffective if it couples deep within the star's potential well (Krumholz 2018), and radiation pressure in general may be subdominant to protostellar outflows for regulating the growth of individual massive stars (Rosen & Krumholz 2020). But by regulating accretion or punching optically thin holes, outflows could help photons to couple their momentum farther away from the star, eventually allowing them to disrupt the host GMC (Fall, Krumholz & Matzner 2010; Murray, Quataert & Thompson 2010; Hopkins, Quataert & Murray 2012; Raskutti, Ostriker & Skinner 2016; Kim, Kim & Ostriker 2018; Grudić et al. 2018a; Hopkins & Grudić 2019; Hopkins et al. 2020a). Meanwhile, jets can be a powerful feedback mechanism that can regulate star formation on the ≤ 1 pc scales of individual cores and dense clumps (Matzner & McKee 2000; Nakamura & Li 2007; Wang et al. 2010; Cunningham et al. 2011; Federrath 2015; Offner & Chaban 2017; Cunningham et al. 2018), but may have only weak effects upon the gas kinematics at larger (i.e. $\gtrsim 10\,\mathrm{pc}$) scales within the GMC (Murray, Goyal & Chang 2018). Many other synergies between feedback mechanisms can also be theorized.

1.2 Enter STARFORGE

In this work, we introduce the STARFORGE (STAR FORmation in Gaseous Environments) project,² a new initiative to perform next-generation 3D star cluster formation simulations in massive GMCs. The STARFORGE numerical framework that we have implemented in the GIZMO code (Hopkins 2015, hereafter H15) simultaneously follows the formation, accretion, and dynamics of individual stars in massive GMCs, with optional physics modules capable of accounting for all of the most widely discussed stellar feedback mechanisms (jets, radiative heating and momentum,

¹Some works have simulated massive cluster and star formation with nominally individual star particles, incorporating SN (Padoan et al. 2017, 2020) and photoionization (Gavagnin et al. 2017) feedback, but at fairly modest (∼ 500−1500 au) resolution compared to state-of-the-art low-mass SF simulations. At this resolution it is only possible to follow the widest binaries, and the predicted IMF may suffer bias or low-mass incompleteness. ²http://www.starforge.space

stellar winds, and SNe), satisfying the requirements (i)–(v) laid out above. In Guszejnov et al. (2020b, hereafter Paper 0), we used numerical simulations (using an early version of the methods presented here) to show that the simple recipe of isothermal MHD turbulence and gravity fails to yield a realistic IMF and star formation history in Milky Way conditions, motivating the need for additional physics implemented in STARFORGE. In the present paper (Paper 1), we present and test the numerical methods of STARFORGE, permitting simulations that combine all of the important SF physics discussed here into a realistic simulation of GMC evolution and star cluster formation. In Guszejnov et al. (2021, hereafter Paper 2), we use the algorithms presented here to explore the effects of protostellar jets upon SF across an unprecedented parameter space of GMC properties and jet model parameters.

This paper is structured as follows. In Section 2, we present the 'core' algorithms that any 3D star cluster formation simulation must have in some form: MHD, gravity, sink particle methods, and an integration scheme that couples gas and stars stably and achieves acceptable accuracy in stellar dynamics. In Section 3, we describe the treatment of cooling, chemistry, and thermodynamics (treating the opacity limit and protostellar heating either with selfconsistent radiative transfer, or simple inexpensive approximations). In Section 4, we describe and test algorithms for the numerical coupling of stellar feedback in the form of iets, winds, SNe, and radiation. In Section 5, we explore various potential applications of these methods beyond isolated GMC simulations, and also enumerate the many caveats and uncertain assumptions inherent in simulating SF and feedback in this manner, motivating future work. In Section 6, we summarize our findings and outline the programme of the STARFORGE project.

2 CORE ALGORITHMS FOR STAR FORMATION

The STARFORGE framework is implemented in the GIZMO multimethod, multiphysics N-body and MHD simulation code (Hopkins 2015).3 GIZMO was selected for the project for several reasons. It implements second-order, Galilean-invariant, Lagrangian meshless finite-volume MHD methods (Hopkins & Raives 2016, hereafter HR16), which have several useful advantages for SF problems (discussed in Section 2.1). It includes a gravity solver (Section 2.2) that is spatially adaptive (solved consistently with the MHD discretization) with near-ideal scaling up to 10⁶ cores (Hopkins et al. 2018b). In addition to solving the MHD equations, GIZMO's meshless discretization and reconstruction schemes provide a flexible framework for solving additional, non-core physics such as diffusion, conduction, and non-ideal MHD terms (Hopkins 2017), radiative transfer (Hopkins & Grudić 2019; Hopkins et al. 2020a), and stellar feedback (Hopkins et al. 2018a). All equations are integrated according to a flexible, hierarchical powers-of-two time-stepping scheme (Section 2.3) that makes it possible to follow processes over a wide range of time-scales, from the $\sim 10\,\mathrm{Myr}$ lifetime of a GMC to a $\lesssim 1\,\mathrm{yr}$ binary orbit.

Conceptually, our approach follows previous Lagrangian 3D star formation simulations (Klessen & Burkert 2000; Bate, Bonnell & Bromm 2003): we discretize the mass of the GMC and the surrounding medium into discrete elements of mass Δm , and integrate their evolution in time according to the MHD equations.

Eventually the self-gravitating MHD equations can no longer be followed self-consistently at the centres of runaway core collapse, so we replace these centres with sink particles (Bate, Bonnell & Price 1995) nominally representing individual protostars. These sink particles interact with the gas via gravity, accretion, and optionally feedback, with feedback rates determined by a sub-grid model of (proto-)stellar evolution based upon that used in Offner et al. (2009). We target a MHD resolution scale on the order of a few 10 au, comparable to state-of-the-art low-mass star cluster formation simulations. We defer physics on smaller scales (e.g. disc formation, accretion, jet launching, and protostellar evolution) to a sub-grid approach, acknowledging the various caveats that this entails (Section 5.2).

We provide a glossary of the various numerical resolution-related quantities in Table 1.

2.1 Magnetohydrodynamics

The default MHD solver used by STARFORGE simulations is the Meshless Finite Mass (MFM) method presented in HR16,⁴ which we will briefly summarize. This method discretizes the fluid into a finite number of gas cells of mass ΔM_i , each representing a domain of volume $V_i = \Delta M_i / \rho_i$ as determined by the kernel⁵ volume partition described in H15. This partition defines the effective face areas $A_{gg'}$ between each interacting pair of gas cells g and g', 6 between which the conservative MHD equations are evolved in standard finite-volume fashion:

$$\frac{\mathrm{d}}{\mathrm{d}t} (VU)_g = \sum_{g'} A_{gg'} \cdot F_{gg'},\tag{1}$$

where $(V U)_g$ gives the usual conserved quantities (mass, momentum, energy, etc.) integrated over the volumetric domain of the cell, and $F_{gg'}$ is the tensor of their fluxes. The fluxes are obtained by solving the appropriate (HLLD) Riemann problem using the fluid states reconstructed at the interface according to a slope-limited, second-order least-squares gradient estimator, evaluated in the frame moving with the interface to ensure Galilean invariance. In MFM, the interfaces are defined and move such that the mass flux vanishes identically, so the method follows the motion of constant-mass, quasi-Lagrangian fluid elements. Cells exchange conserved quantities ensuring machine-precision conservation in this operation. Magnetic field divergence errors are controlled by augmenting equation (1) with the usual Powell et al. (1999) and Dedner et al. (2002) source terms and using the Hopkins (2016) constrained gradient method for obtaining the consistent fluid reconstruction operator that minimizes the numerically unstable terms.

Because the volume partition associated with each cell can have complicated shapes (see Hopkins 2015 for discussion), it is useful to define an effective cell size $\Delta x_g \equiv V_g^{1/3} \equiv (\Delta m_g/\rho_g)^{1/3}$ (the equivalent cell side length for a cubic cell of the same volume and

³http://www.tapir.caltech.edu/~phopkins/Site/GIZMO.html

⁴MFM is our method of choice, but all STARFORGE methods are compatible with any quasi-Lagrangian MHD method implemented in GIZMO, including MFV and SPMHD, enabling easy comparisons.

⁵Following HR16, we adopt the M4 cubic spline as the default kernel partition function $W_{gg'} = W(|\mathbf{x}_{g'} - \mathbf{x}_{g}|, H_{g})$, with kernel radius of compact support H_{g} defined recursively by $H_{g} = 2 \Delta x_{g}$ where Δx_{g} is the kernel-weighted mean cell separation: $\Delta x_{g} \equiv V_{g}^{1/3} = (\overline{n_{g}^{\text{cells}}})^{-1/3} = [\sum_{g'} W_{gg'}]^{-1/3}$.

⁶Throughout this work, we adopt index notation for gas cells and sinks where i and j denote any element regardless of type, g and g' denote gas cells, and s and s' denote sink particles.

Table 1. Glossary of numerical resolution-related quantities in STARFORGE simulations. ΔM_{-3} is the mass resolution Δm in units of the fiducial $10^{-3} M_{\odot}$ resolution, $n_3 = X_{\rm H} \rho / M_{\rm p} \approx 0.7 \rho / M_{\rm p}$ is the local number density of H in units of $10^3 \, {\rm cm}^{-3}$, and $c_{\rm s,\,0.2}$ is the minimum gas isothermal soundspeed in units of $0.2 \, {\rm km \, s}^{-1}$.

Symbol	Meaning	Notes or expression	Fiducial value
Δm	Normal gas cell mass resolution	Numerical parameter	$10^{-3}\mathrm{M}_\odot$
$\Delta M_{ m w}$	Wind cell mass resolution	Numerical parameter	$10^{-4}\mathrm{M}_\odot$
h	Volume-equivalent spherical cell radius	equation (2)	0.02 pc $\Delta m_{-3}^{1/3} n_3^{-1/3}$ 0.03 pc $\Delta m_{-3}^{1/3} n_3^{-1/3}$
Δx	Volume-equivalent Cartesian cell length	$(\Delta m/\rho)^{1/3}$	$0.03 \text{ pc } \Delta m_{-3}^{1/3} n_3^{-1/3}$
$N_{ m NGB}$	Effective neighbour number	Numerical parameter	32
Н	Kernel radius of compact support	$\left(\frac{3N_{\text{NGB}}\Delta m}{4\pi\rho}\right)^{1/3} \approx 2\Delta x$	$0.06 \mathrm{pc} \Delta m_{-3}^{1/3} n_3^{-1/3}$
$f_{ m J}$	Number of Jeans lengths per cell length	Equation (18)	$0.03 n_3^{1/6} c_{s,0.2}^{-1} \Delta m_{-3}^{1/3}$ 1/2
f _{J, max}	$f_{\rm J}$ value for marginal Jeans resolution	Heuristic; problem-dependent; see Section 2.4	1/2
$\Delta x_{\rm J}$	Minimum Jeans-resolved cell length	Equation (20)	$30 \text{ au} \times \Delta m_{-3} c_{s,0,2}^{-2}$
$ ho_{ m J}$	Maximum Jeans-resolved density	Equation (19)	$30 \text{ au} \times \Delta m_{-3} c_{\text{s},0.2}^{-2} \\ 3 \times 10^{-14} \text{g cm}^{-3} \times \Delta m_{-3}^{-2} c_{\text{s},0.2}^{-6}$
$t_{\rm acc}$	Accretion smoothing time-scale	Equation (33)	$500 \mathrm{yr} \times \Delta m_{-3}$
S_{\star}	Sink particle force softening compact support radius	Numerical parameter	$18 \mathrm{au} \times \Delta m_{-3}$
$R_{\rm sink}$	Sink particle maximum accretion radius	Equation (22)	$18 \mathrm{au} \times \Delta m_{-3}$

mass):

$$\Delta x_g \equiv \left(\frac{\Delta m_g}{\rho_g}\right)^{1/3} \approx 0.03 \,\mathrm{pc} \left(\frac{\Delta m}{10^{-3} \mathrm{M}_{\odot}}\right)^{1/3} \left(\frac{n_{\mathrm{H},g}}{10^3 \,\mathrm{cm}^3}\right)^{-1/3},$$
(2)

and volume-equivalent spherical radius

$$h_g \equiv \left(\frac{3\Delta m_g}{4\pi\rho_g}\right)^{1/3} \approx 0.02 \,\mathrm{pc} \left(\frac{\Delta m}{10^{-3} \mathrm{M}_{\odot}}\right)^{1/3} \left(\frac{n_{\mathrm{H},g}}{10^3 \,\mathrm{cm}^3}\right)^{-1/3}, \tag{3}$$

where the latter expressions are given in terms of the typical STARFORGE mass resolution of $10^{-3}\,\mathrm{M}_\odot$ and the number density of H atoms $n_{\mathrm{H,g}}\approx 0.7\rho_g/m_{\mathrm{p}}$. We emphasize, as discussed in HR16, that MFM has little in common with smoothed-particle MHD (SPMHD) – MFM is formally a member of the class of Arbitrary Lagrangian-Eulerian (ALE) finite-volume Godunov methods, much more closely related to Voronoi moving-mesh methods (e.g. Springel 2010; Duffell & MacFadyen 2011), and in fact reduces to a Voronoi-mesh method in the limit of sharply peaked kernel functions with exact volume quadrature.

Meshless, Lagrangian, Galilean-invariant MHD methods have several advantages for simulating SF in GMCs with feedback. In Lagrangian methods, Galilean invariance implies that the time-step required does not depend upon the bulk flow velocity v as $\Delta t \propto \Delta x/v$ (as is required for stable advection in Eulerian fixed-grid methods), so larger time-steps are possible in the highly supersonic flows of GMCs, and the presence of very high ($\gtrsim 100~{\rm km~s^{-1}}$) bulk velocities in accretion flows or winds does not incur such a high cost, an issue often encountered by Eulerian simulations combining high velocities with high-spatial refinement levels.

Galilean and rotational invariance also ensure that structures formed in the simulation (e.g. dense cores and clumps) to evolve internally in a manner independent of their mean bulk velocity and orientation with respect to the coordinate axes (to machine precision). Numerical errors are velocity-independent, and a significant source numerical diffusivity in supersonic flows in Eulerian methods (the grid advection operation, Robertson et al. 2010; Pontzen et al. 2021) is absent. These advantages can enable more rapid convergence of phenomena involving highly supersonic flows, large density con-

trasts, angular momentum conservation, and coupling to self-gravity, all of which are highly relevant in SF. We refer the reader to H15, HR16, and Hopkins (2016) for demonstrations of the performance of the MFM MHD method in a wide variety of standard test problems.

2.1.1 Non-ideal MHD terms

By default, STARFORGE simulations solve the equations of ideal MHD, but it is also possible to include additional terms in the momentum, energy, and induction equations, including *Spitzer* anisotropic conduction and Braginskii viscosity (e.g. Su et al. 2017; Hopkins et al. 2020b), Ohmic resistivity, ambipolar diffusion, and the Hall effect (e.g. Hopkins 2017). These terms are implemented numerically by operator-splitting with the ideal MHD update cycle, as described in Hopkins (2017). The effects of these non-ideal terms in star formation will be the subject of a future study.

2.1.2 Coupled dust-gas dynamics

Physically, dust grains are coupled to gas aerodynamically and hence do not necessarily move with the gas (Draine & Salpeter 1979), and this can have important effects for GMC physics and star formation (Hopkins 2014; Hopkins & Lee 2016). Our default set-up assumes a constant dust-to-metals ratio for cooling, radiative transfer, etc., but compatible with STARFORGE modules are fully compatible with GIZMO's dust dynamics module (Hopkins & Lee 2016; Lee, Hopkins & Squire 2017; Moseley, Squire & Hopkins 2019), which follows dust tracer particles in a Monte Carlo sampling of phase space and grain size, with an arbitrary grain-size distribution, including Stokes, Epstein, and Coulomb drag, Lorentz forces with collisional, photoelectric, and cosmic ray charging, and gas back-reaction. The effect of these physics on star formation will be the subject of future work.

2.2 Gravity

We compute the gravitational field $g = -\nabla \Phi$, tidal tensor $T = -\nabla \nabla \Phi$ (where $\Phi = \nabla^{-2} 4\pi G \rho$ is the gravitational potential) and

the gravitational jerk j at the location of every gas cell and sink particle in the simulation using a modified version of the massively parallel, approximate tree-force algorithm introduced in Springel (2005, hereafter S05). This algorithm recursively subdivides the simulation domain into an oct-tree structure, and uses the monopole approximation of the field contribution of the contents of a tree node, unless an the opening criterion is satisfied, in which case the opening criteria are re-evaluated recursively for all sub-nodes and forces evaluated accordingly. We use the acceleration-based opening criterion introduced in S05 (which requires that the quadrupole error term $a_Q \sim GML^2/r^4$ of a node is less than a specified fraction of the total field g), but also always require the original Barnes & Hut (1986) opening criterion: a tree node is always opened if it subtends an angle $\theta \equiv \frac{L}{r} > \Theta$, where L is the side length of the node, r is the distance between the node centre of mass and the target point for field evaluation, and $\Theta = 0.5$ is the maximum opening angle. This ensures that a dense sub-system of a hierarchically structured system that dominates its own field (e.g. a dense clump or a binary) still has some control over the accuracy of the force contribution from surrounding material, which is still needed to predict its centreof-mass motion (Grudić et al. 2020). T and j are computed in the same pass through the gravity tree as g, summing the respective monopole contributions of tree nodes and particles according to the same opening criterion (Vogelsberger et al. 2008; Grudić & Hopkins 2020). Gravitational forces are updated for gas cells only as frequently as required per the Grudić (2020) adaptive criterion, using j to construct a predictor of g between updates. This generally decreases overall cost of calling the gravity solver by a factor of 2 or

2.2.1 Softening

We use a softened form of the gravitational force law for sink particles or gas cells that fall within each other's respective softening radii S_i and S_j , i.e. $r_{ij} < \max(S_i, S_j)$, ensuring that all interactions are antisymmetric, conserving total linear and angular momentum. Softening for gravitational interactions between gas cells is fully adaptive, i.e. we set $S_i = H_i$ so that the gravitational force resolution is always scaled to be consistent with the cell volume partitioning assumed when solving the MHD equations. This prevents various unphysical effects that are seen if the hydro and gravitational resolution scales are mismatched (Bate & Burkert 1997). At the second-order consistency of our MHD method, H15 showed that this can be done by using the same spherically symmetric compact spline softening scheme as smoothed particle hydrodynamics (SPH), with the same additional terms and symmetrization to ensure conservation of momentum and energy (Price & Monaghan 2007). The full form of the pairwise force law between gas cells is given in H15 (equations H8-H10).

For interactions between sink particles, it would be ideal to use the full, unsoftened $1/r^2$ force law to be able to follow stellar dynamics on all scales down to stellar radii. Unfortunately, this is not presently possible for our code, because binaries with arbitrarily-close separations (down to surface contact) can form and harden dynamically (Heggie 1975), potentially requiring very short time-steps. In theory, this workload could be accomplished by our hierarchical individual time-stepping scheme (Section 2.3), but in practice the massively parallel architecture of GIZMO is such that global overheads tend to eventually bottleneck time-steps involving a small subset of the total particle number (e.g. two sink particles in a very short-period binary). Therefore, we adopt a finite, *fixed* softening

radius S_{\star} for sink particles in a given simulation, allowing us to follow collisional dynamics accurately on spatial scales $\gtrsim S_{\star}$, while limiting the effective hardness of binaries having separation $\lesssim S_{\star}$.

Lastly, softened interactions between fixed-softening sinks and adaptively-softened gas cells must be handled specially, because the respective softenings can differ by orders of magnitude - e.g. if a star with $S_{\star} \sim 20$ au is moving through a diffuse part of the GMC where $n_{\rm H} \sim 10\,{\rm cm}^{-3}$, and hence a gas cell would have size $\sim 0.1\,{\rm pc}$ (equation 2). In such a case using the same Price & Monaghan (2007) symmetrization as gas-gas interactions (averaging the forces) would result in unphysical noise, because the interaction between the gas and star on the scale of S_{\star} is totally unresolved, but the back-reaction on the star depends sensitively upon its position with respect to the cell centre. Instead, we take the *maximum* of S_{\star} and the gas kernel radius H as the softening radius, in both directions of the pairwise interaction (thus conserving momentum). A natural choice for the fixed S_{\star} is to match it to the finest possible Jeans-resolved spatial MHD resolution, which we will show to be $\sim 20 \, \text{AU}$ in Section 2.4 for our fiducial mass resolution of 10^{-3} M_{\odot}.

In all pairwise interactions, the tidal tensor T and jerk j (for sinks) are summed using spatial derivatives of the same softened force kernel that is used for g, with the same symmetrization scheme used for that particular interaction.

2.3 Time-stepping

Gas cells and sink particles are advanced in time in a hierarchical powers-of-two individual block-time-stepping scheme (S05). To compute the time-step taken by an element, we compute numerous time-step criteria for capturing the various physical processes in the simulation, take the smallest of these, and round it down to the next step in the powers-of-two hierarchy. Individual time-steps are essential because the shortest time-steps required are typically on the order of a few days, requiring $\sim\!10^9$ time-steps over the $\sim\!10\,\mathrm{Myr}$ lifetime of a GMC, but the vast majority of elements in the simulation require much less-frequent updates.

2.3.1 Time-step criteria

Gas cells obey all of the standard local, Galilean-invariant, MHD-specific time-step criteria given in HR16, except that we neglect the gravitational component of the acceleration in the Power et al. (2003) acceleration criterion. Instead, both gas cells and sink particles obey the tidal time-step constraint (Grudić & Hopkins 2020):

$$\Delta t_i < \Delta t_{\text{tidal}} = \sqrt{\eta} \left(\frac{\|\boldsymbol{T}\|^2}{6} \right)^{-1/4}, \tag{4}$$

where ||T|| denotes the Frobenius norm of the tidal tensor and η is a tolerance parameter controlling the overall accuracy of integration. In Grudić & Hopkins (2020), we showed that ||T|| encodes a reliable estimate of the local gravitational dynamical time $t_{\rm dyn} \sim \Omega^{-1} = \sqrt{r^3/GM}$, respecting the equivalence principle (invariance to the addition of a uniform external field g') and interpolating between appropriate limits for a continuous mass distribution and in the vicinity of a point mass (e.g. sinks) more accurately and robustly than the usual acceleration-based criterion.

Sink particles also obey their own special time-step criterion for ensuring orbital integration accuracy (von Hoerner 1960):

$$\Delta t_{\rm s} < \Delta t_{\rm 2-body} = \sqrt{\eta} \left(t_{\rm c,min}^{-1} + t_{\rm dyn,min}^{-1} \right)^{-1},$$
 (5)

where

$$t_{c,\min} = \min_{s' \neq s} \frac{\sqrt{r_{ss'}^2 + \epsilon_{\star}^2}}{v_{ss'}} \tag{6}$$

and

$$t_{\rm dyn,min} = \min_{s' \neq s} \sqrt{\frac{\left(r_{ss'}^2 + \epsilon_{\star}^2\right)^{3/2}}{G\left(m_s + m_s'\right)}},\tag{7}$$

where j runs over all other sink particles, $\epsilon_{\star} = h_{\star}/2.8$ is the Plummerequivalent sink softening radius, and $r_{ss'}$, $v_{ss'}$, M_s , $m_{ss'}$ are the separation, relative velocity, and respective masses of sink particles s and s'. Note that $\Delta t_{2\text{-body}}$ is simply the harmonic mean of a kinematic orbital crossing time-scale $\sim r/v$ and an orbital dynamical timescale $t_{\rm dyn} = \Omega^{-1} \sim \sqrt{r^3/GM_{\rm tot}}$, but replacing r with a softened version $\sqrt{r^2 + \epsilon_{\star}^2}$. We treat this as a single time-step criterion using the harmonic mean of the two time-scales because the smooth interpolation in the regime $t_{\rm c,\,min} \sim t_{\rm dyn,\,min}$ yields slightly better integration accuracy for eccentric binary orbits. Unlike the tidal criterion (equation 4), $\Delta t_{2\text{-body}}$ is symmetric between pairs of sinks, ensuring that binaries are updated synchronously when this is the dominant time-step constraint, which can give better conservation of orbital parameters. The global min operations can be evaluated efficiently in the pass through the gravity tree, combining stellar masses that exist within the same tree node if it is not opened (consistent with the force approximation).

Sink particles also observe various time-step constraints derived from local gas conditions, to ensure the stability of local gas interactions occurring within the hydrodynamic stencil, such as accretion and feedback injection. First, it cannot time-step more than $4 \times$ the smallest time-step of a gas neighbour:

$$\Delta t_s < \Delta t_{\text{ngb}} = \min_{g} 4\Delta t_g, \tag{8}$$

where g runs over all overlapping, potentially interacting gas neighbours, i.e. $r_{sg} < \max(H_s, H_g)$. This is analogous to the constraint imposed for neighbouring gas cells in GIZMO, following Saitoh & Makino (2009). A sink particle's time-step is also constrained to anticipate the infall and/or orbital motion of surrounding gas, via a gas free-fall time criterion:

$$\Delta t_s < \Delta t_{\rm ff} = \sqrt{\frac{\eta \left(\max\left(\epsilon_\star, \Delta x_s\right)\right)^3}{Gm_s}},$$
 (9)

where η is the parameter controlling overall integration accuracy and Δx_s is the effective gas cell size in the vicinity of the sink. Sinks also obey a local Courant–Friedrichs–Lewy (CFL)-type time-step constraint:

$$\Delta t_{s} < \Delta t_{\text{CFL},\star}$$

$$= C_{\text{CFL}} \frac{\Delta x_{s}}{\max\left(\sqrt{c_{\text{s,s}}^{2} + v_{\text{A,s}}^{2} + |\boldsymbol{v}_{\text{s}} - \boldsymbol{v}_{\text{gas,s}}|^{2}}, \tilde{c}, v_{\text{fb,s}}\right)},$$
(10)

where v_s is the velocity of the sink, $c_{s,s}$, $v_{A,s}$, and $v_{gas,s}$ are the local gas sound speed, Alfvén speed, and gas velocity (reconstructed using a simple kernel-weighted interpolation), \tilde{c} is the (possibly reduced) speed of light (only included if radiative transfer is enabled), and v_{fb} is an estimate of the maximum velocity of gas emerging from the sink due to feedback:

$$v_{\text{fb}} = \max(v_{\text{SN}}, \min(v_{\text{wind}}, v_{\text{shell}})), \tag{11}$$

where $v_{\rm SN}$ is the SN ejecta velocity given by equation (47) (or 0 if the sink is not currently going SN), $v_{\rm wind}$ is the stellar wind

velocity (equation 45), and $v_{\rm shell}$ is the greater of the velocity of an energy-conserving (Weaver et al. 1977) or momentum-conserving (Steigman, Strittmatter & Williams 1975) shell as its radius reaches the resolution scale Δx :

$$v_{\text{shell}} = \max\left(0.38 \left(\frac{L_{\text{wind}}}{\rho \Delta x_s^2}\right)^{1/3}, \left(\frac{0.053 \left(L/c + \dot{M}_{\text{wind}} v_{\text{wind}}\right)}{\rho \Delta x_s^2}\right)^{1/2}\right), \tag{12}$$

where $\dot{M}_{\rm wind}$ is the sink's wind mass-loss rate (equation 32), $v_{\rm wind}$ is the wind velocity (equation 45), $L_{\rm wind} = \frac{1}{2} \dot{M}_{\rm wind} v_{\rm wind}^2$ is the mechanical luminosity of the wind, L is the bolometric luminosity of the sink, ρ is the local gas density. We have found that including something like the $v_{\rm fb}$ term in $\Delta t_{\rm CFL}$, \star can be important to prevent the sink particle from 'overshooting' the amount of feedback it injects, i.e. injecting feedback over a time-step longer than the time required for the local gas cells to react to it, leading to an unstable solution.⁷

Reciprocally to the v_{fb} term in equation (10), gas cells also obey a time-step constraint to anticipate the arrival of feedback from a star:

$$\Delta t_g < \Delta t_{\text{fb}} = C_{\text{CFL}} \min_{s \in \text{sinks}} \frac{\sqrt{r_{gs}^2 + \max\left(\epsilon_{\star}, \Delta x_g\right)^2}}{v_{\text{fb.s}}}, \tag{13}$$

where *s* runs over all sink particles and the min can be evaluated efficiently in the gravity tree pass. If a hyperbolic RT solver (e.g. M1) is enabled, we also enforce a local radiation CFL condition:

$$\Delta t_g < \Delta t_{\text{CFL,rad}} = C_{\text{CFL}} \frac{2h_g}{\tilde{c}}.$$
 (14)

Likewise, we enforce appropriate local time-step criteria in the relevant methods papers for the various optional physics (e.g. non-ideal MHD and dust) described above.

2.3.2 Time integration

Given a choice of individual time-step as in Section 2.3.1, we require a time-integration scheme that achieves acceptable truncation error. The error budget of a multiphysics SF simulation is dominated by errors in MHD, radiative transfer, stellar evolution/feedback, and gravity, all of which are necessarily approximate and/or have large modelling uncertainties. Some errors may not even converge away in the limit of infinite resolution: e.g. moments-based radiative transfer methods will not converge to the exact radiative transfer solution in general. Hence, for gas, high-order integration schemes are unlikely beat down the leading-order error terms in the global GMC and star cluster evolution. For all gas cells, and as a robust fall-back option for stars in special circumstances, we use the standard second-order Kick–Drift–Kick (KDK) integrator (S05):

$$\mathbf{v}_{i} \mapsto \mathbf{v}_{i} + \frac{1}{2} \Delta t_{i} \, \mathbf{a}_{i},
\mathbf{x}_{i} \mapsto \mathbf{x}_{i} + \Delta t_{i} \, \mathbf{v}_{i},
\mathbf{v}_{i} \mapsto \mathbf{v}_{i} + \frac{1}{2} \Delta t_{i} \, \mathbf{a}_{i}, \tag{15}$$

where a_i is the total gravitational+MHD+radiative acceleration of cell/particle i, which is re-evaluated after every initial half-step kick.

⁷This is only required to stabilize feedback mechanisms using weighted local injection within the hydrodynamic stencil: the component of radiation pressure due to unresolved absorption, and stellar winds with unresolved free expansion (Section 4). Feedback mechanisms that resolve the ejecta self-consistently (resolved winds, jets, and SN) are stable because the high-velocity ejecta 'wake up' ambient gas cells as they approach, bringing them down to the necessary time-step automatically (Saitoh & Makino 2009).

Some additional control on orbital integration accuracy for stars is needed to preserve the properties of binaries once formed. Any numerical integration scheme will incur a certain fractional energy error $\Delta E/E$ per orbit (as well as an angular momentum error $\Delta J/J$ and phase error $\Delta \phi/\phi$, but here we use $\Delta E/E$ as an overall proxy for integration error, as is standard). A true symplectic integrator such as the leapfrog with constant time-steps will preserve orbital energy and angular momentum on average, but true symplecticity is lost once the adaptive KDK version is adopted and errors accumulate over time, causing the semimajor axis to change with each orbit (S05). If a fairly typical $\sim 100 \, \text{yr}$, $e = 0.9 \, \text{binary}$ is to survive the $\sim 10 \, \mathrm{Myr}$ lifetime of its host GMC in a simulation, then we require $|\Delta E/E| \ll 1$ and hence an energy error per orbit of $\ll 10^{-5}$. In Fig. 1, we show that this would require $\gg 2000$ time-steps per orbit with the KDK integrator (and would demand a minimum time-step at periastron of $\lesssim 1$ d), which we have found to demand an excessivelylarge overhead. Instead, we adopt a modified version of the fourthorder Hermite integrator (Makino & Aarseth 1992) for stars. At the beginning of the time-step, we evaluate the initial accelerations $a_{s,0}$ and jerk $j_{s,0}$ of all sinks in a special sinks-only gravity tree pass. We then perform the initial prediction step:

$$\mathbf{x}_{s} \mapsto \mathbf{x}_{s,0} + \Delta t \, \mathbf{v}_{s,0} + \frac{1}{2} \Delta t^{2} \mathbf{a}_{s,0} + \frac{1}{6} \Delta t^{3} \mathbf{j}_{s,0},
\mathbf{v}_{s} \mapsto \mathbf{v}_{s,0} + \Delta t \, \mathbf{a}_{s,0} + \frac{1}{2} \Delta t^{2} \mathbf{j}_{s,0}, \tag{16}$$

re-evaluate a_s and $j_{,s}$ using the new positions and velocities, and then perform the correction step:

$$v_{s} \mapsto v_{s,0} + \frac{1}{2} \Delta t \left(a_{s} + a_{s,0} \right) + \frac{1}{12} \Delta t^{2} \left(j_{s,0} - j_{s} \right),$$

$$x_{s} \mapsto x_{s,0} + \frac{1}{2} \Delta t \left(v_{s} + v_{s,0} \right) + \frac{1}{12} \Delta t^{2} \left(a_{s,0} - a_{s} \right),$$
(17)

where the order here matters because the update to x_s requires the updated version of v_s . This is subtly different from the original implementation in Makino & Aarseth (1992), in that we perform two force/jerk evaluations per time-step, one at the beginning of the timestep and one after the prediction step, whereas the original Hermite scheme only re-evaluates the force/jerk after the prediction step. We discovered serendipitously that this small modification gives a scheme that converges at the same order, but can give order-ofmagnitude smaller energy errors at fixed time-step size in binary integration (Fig. 1). In a typical direct N-body application the entire cost of the simulation is force/jerk evaluation and there is not much parallelization overhead, so this advantage would be nullified by simply taking 2× smaller time-steps at equal cost. In GIZMO, the force/jerk comes relatively cheaply, but there can be significant global overheads involved in taking smaller time-steps, so our modified Hermite scheme is more suitable. For our standard choice of $\eta =$ 0.01, this method achieves a relative energy error of $<10^{-6}$ per orbit for an e = 0.9 binary (and this decreases steeply for smaller e).

In a given time-step, a sink particle is first provisionally time-stepped according to the KDK scheme, co-evolving it alongside the gas update cycle so that the gas-star coupling seen by the gas is unaltered by the Hermite integration (but saving the initial state of the time-step). At the end of the time-step, the sink particle is eligible to accrete gas cells. If it does, low-order integration errors are introduced and j ceases to be well-defined, so we simply keep the more-robust KDK result for that time-step. If it does not accrete, it is eligible to take a Hermite time-step, updating via equation (16) using the saved values $x_{s,0}$, $v_{s,0}$, $a_{s,0}$, and $j_{s,0}$, and performing the subsequent force evaluation and correction step (equation 17). Given




Figure 1. Relative energy error accumulated per orbit integrating e=0.9 binary motion with the second-order Kick–Drift–Kick (KDK) (S05), fourth-order Hermite (Makino & Aarseth 1992) and our modified Hermite integrator (equations 16–17), as a function of the time-step tolerance parameter η , which controls the number of steps taken per orbit per our adaptive time-stepping scheme (equations 4 and 5 and powers-of-two block scheduling, Section 2.3). Conserving binary properties in a $\sim 10\,\mathrm{Myr}$ GMC simulation ($\gtrsim 10^4-10^7$ binary orbits) is only practical with a higher order scheme. Both Hermite schemes happen to converge at fifth order in this problem when using our time-step criteria, and our modified version performs better at fixed η , for an extra force and jerk evaluation per time-step.

the order of the MHD reconstruction, and the inability to define the jerk given e.g. shocks, there would be no gain from using this integrator for gas.

2.4 Isothermal hydro+gravity tests and resolution requirements

Before discussing sink particles, it is useful to first examine the behaviour of our methods in test problems involving simple isothermal hydrodynamics and gravity.

2.4.1 Existing tests

The standard MHD and gravity algorithms in GIZMO have been extensively tested and applied to hundreds of different problems in the literature, so we will not repeat these. We do note these tests have demonstrated that our default implementation can simultaneously accurately evolve phenomena including gas in regular or warped Keplerian discs, strong interacting shocks, current sheets and flux tubes, supersonic *and* subsonic turbulence, fluid mixing instabilities (Kelvin–Helmholz, Rayleigh Taylor, etc.), multifluid dustgas dynamics, collisional+collisionless gravitational dynamics, and reproduces the correct linear growth rates of the magnetorotational instability (MRI) and non-ideal Hall MRI and anisotropic MHD instabilities (magnetothermal, heat–flux–bouyancy) (Hopkins 2015; Hopkins & Raives 2016; Zhu & Li 2016; Lupi, Volonteri & Silk 2017; Deng et al. 2019a, b; Moseley et al. 2019; Rennehan et al. 2019; Hu &

Chiang 2020; Panuelos, Wadsley & Kevlahan 2020). Tests of idealized problems involving self-gravitating MHD including the Evrard (1988) problem (spherical collapse of a self-gravitating polytrope), the MHD Zel'dovich (1970) pancake (self-gravitating collapse of an initially linear density perturbation along one dimension in a 3D Hubble flow) demonstrate that the MFM/MFV methods in GIZMO (as well as related moving-mesh methods) converge much more rapidly than popular AMR or SPH methods applied to the same problem (Hopkins 2015; Hopkins & Raives 2016; Hubber, Rosotti & Booth 2018).

Several studies have argued that the most notable advantages of MFM compared to SPH or AMR methods may come in astrophysical discs, which are crucial for the physics of stellar accretion but are often marginally resolved in our simulations (meaning that more-rapid convergence at fixed resolution is especially valuable). For example, (1) MFM accurately conserves angular momentum and prevents both unphysical disc 'spreading' and/or clumping/fragmentation via artificial viscous instabilities in SPH or catastrophic angular momentum loss from spurious coordinatealignment torques which are inescapable in AMR (Hopkins 2015; Few et al. 2016; Zhu & Li 2016; Lupi et al. 2017; Panuelos et al. 2020; Deng, Ogilvie & Mayer 2021). (2) Few et al. (2016) found MFM more rapidly converges to correct linear growth rates for spiral arms and other disc instabilities, compared to AMR or SPH, while Deng et al. (2021) found a similar result for physical parametric instabilities of warped discs. (3) Deng, Mayer & Meru (2017) showed MFM was the only method surveyed which exhibited convergence to exact solutions for gravitoturbulent fragmentation in cooling discs. (4) MFM, at fixed resolution, has been shown to more accurately capture boundary-layer mixing in discs (especially those formed via collisions), avoiding artificial suppression of subsonic turbulence and mixing common in SPH (Zhu & Li 2016; Deng et al. 2019b). (5) Hubber et al. (2018) demonstrated that MFM simulations of 'gap opening' via massive planets or binaries in discs converged more rapidly and maintained the gaps more accurately than equivalent SPH or AMR simulations (which tended to produce artificially-high torques and therefore stellar accretion rates in this regime).

2.4.2 Isothermal collapse tests

Next, we consider a variant of the 'isothermal test case' from Boss & Bodenheimer (1979): a uniform-density, un-magnetized, spherical solar-mass core with initial radius 5×10^{16} cm, in uniform rotation with $\Omega = 7.2 \times 10^{-13} \text{rad s}^{-1}$ and a 10 per cent m = 2 azimuthal density perturbation with an isothermal equation of state $P = c_s^2 \rho$, $c_s = 0.166 \text{ km s}^{-1}$ (Burkert & Bodenheimer 1993; Bate & Burkert 1997; S05). We use the MFM hydrodynamics solver with the default STARFORGE gravity and time-stepping set-up (Sections 2.2–2.3), and initialize the cells in a glass configuration with the density perturbation imposed by rescaling cell masses. In Fig. 2, we plot the maximum density in the simulation as a function of time while varying the average cell mass Δm from $10^{-3}-10^{-7}\,\mathrm{M}_{\odot}$, and compare with SPH results from Bate & Burkert (1997) and S05. The solution appears to converge to an answer in fair agreement with the highest resolution SPH results in S05. Moreover, our solution with $N_{\rm gas} = 10^4$ cell resolution is closer to its respective converged limit than SPH simulations with 3.34×10^4 and 8×10^4 particles, respectively. However, at low enough resolution numerical effects become apparent, as evidenced by the ~ 10 per cent delay of the collapse of our lowest resolution run with only 10³ gas cells.




Figure 2. Evolution of the maximum density $\rho_{\rm max}$ in the standard isothermal test problem (Boss & Bodenheimer 1979), with time in units of the global cloud free-fall time $t_{\rm ff}$. We plot MFM results for various mass resolutions $\Delta m = 10^{-7} - 10^{-3} {\rm M}_{\odot}$ and compare with SPH results from Bate & Burkert (1997) and S05.

It is important to assess the effects of resolution upon SF simulations, as this will inform our sink particle prescription. A common convergence parameter for self-gravitating isothermal hydrodynamics simulations is the number of Jeans lengths per cell (Bate & Burkert 1997; Truelove et al. 1997; Hubber, Goodwin & Whitworth 2006):

$$f_{\rm J} \equiv \frac{\Delta x}{\lambda_{\rm J}} = \sqrt{\frac{G}{\pi c_{\rm s}^2}} (\Delta m)^{1/3} \rho^{1/6}$$

$$\approx 0.03 \left(\frac{\Delta m}{10^{-3} {\rm M}_{\odot}}\right)^{1/3} \left(\frac{n_{\rm H}}{10^3 {\rm \, cm}^{-3}}\right)^{1/6} \left(\frac{c_{\rm s}}{0.2 {\rm \, km \, s}^{-1}}\right)^{-1},$$
(18)

using $\Delta x = \left(\frac{\Delta m}{\rho}\right)^{1/3}$ and $\lambda_{\rm J} = c_{\rm s} \sqrt{\frac{\pi}{G\rho}}$. The consequences of underresolving λ_J (i.e. allowing $\Delta x \gg \dot{\lambda}_J$) vary from method to method, and have been the subject of extensive study. Truelove et al. (1997, hereafter T97) found that Eulerian grid simulations that do not enforce $f_{\rm J} < \frac{1}{4}$ are subject to artificial fragmentation (AF), wherein fragments of unphysical origin can form even in a smooth, symmetric collapse. A similar effect is seen in SPH simulations if care is not taken to match the gravitational resolution to the hydrodynamic resolution (Section 2.2.1), e.g. adopting a constant gas softening length that is much smaller than the particle spacing (Bate & Burkert 1997). Clearly, AF is undesirable, so a variety of approaches have been developed to prevent it, e.g. by fine-tuning the sink particle formation, accretion, and merger criteria in conjunction with the refinement scheme (e.g. Krumholz, McKee & Klein 2004; Haugbølle, Padoan & Nordlund 2018). AF does not occur in SPH simulations that maintain consistency between gravitational and hydrodynamic resolution (Bate & Burkert 1997; Whitworth, Boffin & Francis 1998; Hubber et al. 2006), and more recently it has been confirmed that this is true for MFM as well in the linear Jeans problem (Hubber et al. 2018; Yamamoto et al. 2021 in preparation). With these methods, fragments that should physically collapse but are insufficiently Jeansresolved either do not collapse, or simply collapse more slowly.




Figure 3. Effect of (under-)resolving the Jeans length in isothermal collapse with the Meshless Finite Mass (MFM) method with adaptive gravitational softening. We plot a surface density map of the filament formed in the T97 self-gravitating isothermal hydrodynamics test problem, at the time that the maximum density is $\rho_{\text{max}} = 10^{-9.5} \text{g cm}^{-3}$ (cf. T97 Fig. 4). From top left to bottom right, we vary the mass resolution Δm from $8 \times 10^{-5} - 1.6 \times 10^{-7} \, \text{M}_{\odot}$, which varies the maximum number of Jeans wavelengths per cell $f_J = (\Delta x/\lambda_J)_{\text{max}} \approx \sqrt{G} c_s^{-1} \, (\Delta m)^{1/3} \, \rho_{\text{max}}^{1/6}$ from 1.1 to 0.14. Failure to resolve the Jeans length simply coarse-grains the structure of the filament – there is no evidence of artificial fragmentation when the Jeans length is poorly resolved.

Here, we also check for AF in the exact test problem simulated in T97, a variant of the Boss & Bodenheimer (1979) problem, using an initial Gaussian density profile. With a 10 per cent m=2initial density perturbation, T97 found that the converged solution is the formation of a single collapsing filament, but if the Jeans resolution criterion $\Delta x > \lambda_{\rm I}/4$ was violated then they would obtain an unphysical solution containing two filaments instead. In Fig. 3, we plot the structure formed in the simulation at the time that the maximum density exceeds $10^{-9.5}$ g cm⁻³, at a variety of mass resolutions such that the T97 criterion is strongly violated at our lowest resolution (8 × 10⁻³ M_{\odot} , $\Delta x \approx 1.1 \lambda_{\rm J}$), and is well satisfied at our highest (1.56 \times 10⁻⁷ M_{\odot} , $\Delta x \approx 0.14 \lambda_J$). No additional filament or fragment forms even when the T97 criterion is strongly violated – the effect of poor resolution appears consistent with a simple spatial coarse graining of the structure of the filament. T97 also found that the version of the problem with no initial density perturbation resulted in the formation of numerical fragments, unless $f_{\rm I}$ < 1/4 was enforced. We have verified that this is not the case for MFM: axisymmetry is preserved accurately throughout the collapse, even when the Jeans resolution criterion is strongly violated.

Our findings for MFM appear consistent with previous results in the linear Jeans problem (Hubber et al. 2018; Yamamoto et al. in preparation): unstable scales that are well-resolved ($f_J \ll 1$) collapse as they should, and scales that should be stable are stable. Marginally resolved ($f_J \sim 1$) unstable wavelengths are either artificially stabilized, or collapse more slowly than is physical (e.g. the lowest resolution run in Fig. 2), and these effects converge away with sufficient resolution.

2.4.3 Resolution criteria

What density- and length-scales should then be considered 'resolved' in isothermal self-gravitating flows? This depends on what threshold value of f_J is considered acceptable for the question at hand, which is generally problem-dependent with no one straightforward answer. But assuming we do adopt a certain maximum $f_{J, \max}$ to demarcate the boundary of 'trusting' results in a certain problem, the maximum Jeans-resolved density is,

$$\rho_{\rm J} \equiv f_{\rm J,max}^6 \frac{\pi^3 c_{\rm s}^6}{G^3 \Delta m^2}
\approx 3 \times 10^{-14} {\rm g \, cm}^{-3} \left(\frac{f_{\rm J,max}}{0.5}\right)^6 \left(\frac{\Delta m}{10^{-3} {\rm M}_{\odot}}\right)^{-2} \left(\frac{c_{\rm s}}{0.2 {\rm \, km \, s}^{-1}}\right)^6,$$
(19)

and the minimum Jeans-resolved cell length is

$$\Delta x_{\rm J} \equiv f_{\rm J,max}^{-2} \frac{G \Delta m}{\pi c_{\rm s}^2} \approx 30 \,\text{au} \left(\frac{f_{\rm J,max}}{0.5}\right)^{-2} \left(\frac{\Delta m}{10^{-3} \text{M}_{\odot}}\right)$$

$$\times \left(\frac{c_{\rm s}}{0.2 \,\text{km s}^{-1}}\right)^{-2}.$$
(20)

We caution that direct comparisons of the 'resolved' density- or length-scale between SF simulations should ideally be made at fixed $f_{\rm J,\,max}$ (i.e. correcting by appropriate $f_{\rm J,\,max}$ factors), and that even then there can be many confounding factors when comparing across different methods.

This discussion of Jeans resolution neglects magnetic fields, which can supplement thermal pressure as a source of support against gravi-

tational collapse. For the purposes of gravitational stability analyses, it effectively adds the Alfvén speed $v_{\rm A}=|\pmb{B}|/\sqrt{\mu_0\rho}$ to the thermal sound speed $c_{\rm s}$ in quadrature, i.e. $c_{\rm s}\mapsto\sqrt{c_{\rm s}^2+v_{\rm A}^2}=\sqrt{1+\frac{2}{\beta}}c_{\rm s}$, modulo some geometry-specific O(1) factors in the β term

modulo some geometry-specific O(1) factors in the β term (Chandrasekhar 1951; Mouschovias & Spitzer 1976). Assuming that the convergence parameter for isothermal, self-gravitating MHD is instead the *magnetic* Jeans number obtained by substituting the above into equation (18), as has been argued in various works (Federrath et al. 2010; Myers et al. 2013), our assessment of the resolving power of the simulations (equations 19 and 20) is overly conservative. However, the densest gas in isothermal MHD core collapse attracts towards $\beta \sim 1$ (Mocz et al. 2017; Wurster et al. 2019; Guszejnov et al. 2020b), so the corrections to our analysis from magnetic fields are expected to be modest for the present purposes. Even if not, this would merely make our effective f_J threshold more conservative, so e.g. our sink algorithm would not follow gas as far into the marginally resolved regime, and our simulations are better-resolved than as quoted in equations (19) and (20).

2.5 Sink particles

We use sink particles to model the accretion, dynamics, and feedback of individual stars and protostars (e.g. Bate et al. 1995; Krumholz et al. 2004; Federrath et al. 2010; Hubber, Walch & Whitworth 2013; Bleuler & Teyssier 2014). A sink particle represents a designated region in the domain of the simulation in which physical processes are considered unresolved, and are relegated to sub-grid prescriptions. The general strategy is to put a sink in the centre of a collapsing core once the collapse process can no longer be followed self-consistently by the MHD scheme, and to allow this sink to accrete subsequent infalling material according to certain physically motivated rules.

Our sink implementation formally distinguishes between resolved accretion, i.e. the actual mass transfer from the gas in the simulation domain to the sink particle, and unresolved accretion: the transfer of mass from the sink's internal gas reservoir (comprising unresolved gas in the envelope or the protostellar disc) on to the protostar itself (and potentially into the protostellar outflow). Other works equate the two types of accretion, often assuming that gas removed from the simulation domain arrives at the protostar immediately (e.g. Krumholz et al. 2004), or using a detailed sub-grid model to decide how rapidly resolved accretion should occur (Hubber et al. 2013). For us it is important to model accretion on to the protostar distinctly from resolved accretion into the sink region, because we discretize resolved accretion into quanta - the mass resolution Δm - but would like a continuous estimate of the protostellar accretion rate for modelling the protostellar evolution and accretion luminosity.8 Our algorithm is most similar to that of Bate et al. (1995), with some additional rules for formation

⁸We have experimented with our own implementation of the algorithm of Hubber et al. (2013), which uses an estimate of \dot{M}_{\star} that interpolates between disc-like and Bondi-like regimes based on local gas kinematics, and uses that estimator to directly determine how much gas should be removed. However, we have found that in some problems the estimator of \dot{M}_{\star} used to set the rate of resolved accretion can underestimate the actual accretion rate of the surrounding flow, so mass piles up within the softening radius of the sink particle and the actual accretion rate ends up being set by the need to remove gas cells on too small a time-step (circumventing the normal criteria), defeating the purpose of trying to estimate and enforce the proper accretion rate as determined by physical processes. One potential issue is that the α -disc parameter used in the disc-like regime must be known a priori, otherwise the accretion rate will not match the boundary flow. This will generally vary with




Figure 4. Diagram of the flow of mass due to accretion and feedback, as managed by our sink particle algorithm. We follow gas cells of mass Δm until they satisfy all sink particle accretion criteria (Section 2.5.2) and they are transferred to the sink's accretion reservoir representing the envelope or disc gas mass present on scales $< R_{\rm sink}$. Mass is accreted from the reservoir towards the protostar according to the smoothed accretion prescription (equation 32), and if protostellar jets are enabled a fraction of this mass $f_{\rm w}$ is diverted to the jet reservoir. The rest arrives at the star, and mass is transferred from the star to the wind reservoir according to the wind mass loss rate (which is set to an extremely large value (with appropriate velocity) if the star goes SN, equations 48 and 47). The jet and wind reservoirs return gas to the simulation domain via their respective feedback channels (waiting until a sufficient mass is available to inject or spawn their respective mass quanta).

and accretion, and some additional modelling of protostellar accretion and feedback. We sketch the flow of mass dictated by our algorithm in Fig. 4, and describe the algorithm in detail in this section.

2.5.1 Formation

A gas cell is eligible to turn into a sink particle if and only if it satisfies the following criteria:

- (i) *Density threshold*: The gas cell is denser than a density threshold ρ_{th} , which we take to be the maximum density of marginal Jeans resolution, ρ_{J} (equation 19), assuming $f_{J} = 1/2$.
- (ii) Density maximum/no overlapping sink: The gas cell is the densest of all neighbouring gas cells or sink particles with overlapping kernel radii (with $r_{gi} < \max{(H_g, H_i)}$). For the purposes of this criterion we take sinks to have infinite density, i.e. overlapping with a pre-existing sink always prevents sink formation.
- (iii) *Increasing density*: The gas cell's density is increasing: $\nabla \cdot \mathbf{v} < 0$, according to the same least-squares matrix gradient estimator of $\nabla \mathbf{v}$ used for reconstructing fluid quantities for the MHD solver.
- (iv) *Virial/Jeans criterion*: The gas cell is gravitationally unstable/self-gravitating at the resolution scale, as determined by a local Jeans analysis including contributions from thermal pressure, magnetic fields, and velocity dispersion (Federrath et al. 2010; Hopkins, Narayanan & Murray 2013a). We evaluate a local virial parameter for the gas cell:

$$\alpha_g = \frac{\frac{2\pi^2}{\Delta x^2} \left(c_s^2 + v_A^2\right) + \|\nabla \boldsymbol{v}\|^2}{4\pi G \rho},\tag{21}$$

where $\Delta x = (\Delta m/\rho)^{1/3}$ is the local cell length, and $\|\cdot\|$ denotes the Frobenius norm. We permit sink formation only if $\alpha_{\nu} < 2$. It is easy to

turbulent and numerical viscosity, magnetic torques, gravitational torques, etc., and cannot generally be fit by a single choice of α .

verify that this reduces to the usual requirement that the cell is Jeans-unstable at the resolution limit where kinetic energy is negligible, and recovers the Hopkins et al. (2013a) kinematic virial criterion when the $\|\nabla v\|$ term dominates (e.g. preventing sink formation in Toomre-stable flows stabilized by shear near a star).

- (v) Tidal criterion: The tidal tensor T at the position of the gas cell is fully compressive (possesses three negative eigenvalues). Note that the linearization of the gravitational field about a point x_i is $g(x) - g(x_i) \approx T \cdot (x - x_i)$, so if T has three negative eigenvalues then the gravitational force seen in the frame comoving with a ballistic trajectory starting at x_i will always be directed back towards the origin. This is similar in motivation to the potential minimum requirement in Federrath et al. (2010) and the Hill sphere criterion in Hubber et al. (2013), intended to pick out actual centres of collapse from the shape of the gravitational landscape. Unlike a potential minimum criterion, the tidal criterion respects the equivalence principle, i.e. it is invariant to the transformation $g \mapsto g + g'$ for a spatially uniform g', which should not physically affect the internal dynamics of the simulation in any way, but would displace the location of a potential minimum. However, it is less strict than a potential minimum criterion, e.g. it is satisfied at every point in a uniform sphere (in which T is constant and negative-definite), whereas the potential minimum criterion is satisfied at one point, or none if the external field g' exceeds the internal field.
- (vi) Can collapse before accretion: The local gas free-fall time $t_{\rm ff} = \sqrt{\frac{3\pi}{32G\rho_{\rm g}}}$ is shorter than both the time-scale for approaching a sink particle and the orbital time-scales around that sink particle, as estimated by evaluating equations (6) and (7) for gas cells.

When a gas cell is converted to a sink particle, it is removed from the simulation domain, and the volume it occupied is reassigned to surrounding cells when they re-compute their volume partitions.

2.5.2 Accretion criteria

Gas cells are accreted by a sink particle if they satisfy the following criteria:

(i) Sink radius: A gas cell is only eligible for accretion if its centre approaches within sink radius $R_{\rm sink}$. We take $R_{\rm sink}$ to be the greater of the sink particle softening radius S_{\star} or volume-equivalent radius of a gas cell at the density of marginal Jeans resolution $\rho_{\rm J}$ (equation 19, assuming $f_{\rm J}=1/2$):

$$R_{\text{Sink}} = \max \left(S_{\star}, 0.79 \frac{G \Delta m}{c_{s}^{2}} \right)$$

$$= \max \left(S_{\star}, 18 \text{ au} \left(\frac{\Delta m}{10^{-3} \text{M}_{\odot}} \right) \left(\frac{c_{s}}{0.2 \text{ km s}^{-1}} \right) \right), \quad (22)$$

where c_s denotes the isothermal sound speed at sink formation (i.e. it is set at formation, and kept constant thereafter).

(ii) Boundedness criterion: The gas cell satisfies

$$2u_g + v_{A,g}^2 + |\mathbf{v}_g - \mathbf{v}_s|^2 < v_{\text{esc}}^2 = -2\Phi\left(r_{gs}\right), \tag{23}$$

where u_g is the specific internal energy of the gas, $v_{A,g}$ is its Alfvén speed, and $\Phi(r_{gs})$ is the softened gravitational potential of the sink, evaluated at the separation between the gas and sink r_{gs} . This checks that the sink-gas system is gravitationally bound, and could not escape to infinity in isolation.

(iii) Angular momentum criterion: The gas cell possesses less angular momentum than a circular Keplerian orbit around the sink at r_{gs} (Bate et al. 1995):

$$|(\boldsymbol{x}_{g} - \boldsymbol{x}_{s}) \times (\boldsymbol{v}_{g} - \boldsymbol{v}_{s})|^{2} < Gm_{s}r_{gs}. \tag{24}$$

In the limit of ballistic flow, this ensures that the orbit of the gas cell lies within R_{sink} (so we do not capture e.g. a gas cell that only makes a single close passage but then interacts outside R_{sink} and escapes).

(iv) **Size/density criterion**: The volume of the gas cell is less than the volume within R_{sink} :

$$V_g = \frac{m_g}{\rho_g} < \frac{4\pi}{3} R_{\rm sink}^3.$$
 (25)

This has the effect of ensuring that only gas having spatial resolution on the scale of $R_{\rm sink}$ can be accreted, which may be necessary for the other criteria to be reliable predictors of the gas's dynamics (and whether it is legitimately being accreted). In any true resolved accretion flow, gas will pile up around the sink until this criterion is eventually satisfied. It is analogous to maintaining the maximum refinement level in the vicinity of a sink in an AMR simulation (Krumholz et al. 2004).

It is possible for a gas cell to satisfy all of these criteria for more than one sink. In such instances, the gas is accreted by the sink s with which it has shortest mutual dynamical time $t_{\rm dyn} = \Omega^{-1} = 0$

$$\sqrt{\frac{r_{gs}^3}{G(m_g+m_s)}}$$
.

The quantization of resolved accretion into parcels of mass Δm has certain important limitations. Clearly, a sufficiently slow accretion flow with $\dot{M} \ll \Delta m/t$, where t is some time-scale of interest, cannot be captured. In the limit of a ballistic, Bondi-like flow, we can take $t=t_{\rm dyn}\,(< R)=\sqrt{R^3/GM_\star}$ at some radius of interest R. Then, assuming the physical accretion flow has a certain \dot{M}_\star , the most optimistic radius down to which the flow can be resolved is

$$R_{\min} = 73 \text{ au} \left(\frac{\Delta m}{10^{-3} M_{\odot}}\right)^{2/3} \left(\frac{\dot{M}_{\star}}{10^{-5} M_{\odot} \text{ yr}^{-1}}\right)^{-2/3} \left(\frac{M_{\star}}{1 M_{\odot}}\right)^{1/3},$$
(26)

where we insert typical values for Δm , \dot{M}_{\star} , and M_{\star} . Hence, the accretion flow becomes less well-resolved for smaller accretion rates and greater stellar masses. This may impose some numerical bias towards higher accretion rates in the accretion histories of sink particles, and underestimate more extended periods Bondi-Hoyle accretion from low-density gas. However, the effect does converge to the correct solution with sufficient mass resolution.

2.5.3 Updating conserved quantities

When a gas cell is accreted, it is deleted from the simulation domain and the volume partition of neighbouring gas cells is re-computed. Its mass, first mass moment $m_{\rm g} {\bf x}_{\rm g}$, momentum, and angular momentum are added to the sink:

$$m_s \mapsto m_s + m_g,$$
 (27)

$$\boldsymbol{x}_{\mathrm{S}} \mapsto \frac{m_{\mathrm{s}}\boldsymbol{x}_{\mathrm{s}} + m_{\mathrm{g}}\boldsymbol{x}_{\mathrm{g}}}{m_{\mathrm{s}} + m_{\mathrm{g}}} = \boldsymbol{x}_{\mathrm{s}}',\tag{28}$$

$$p_{s} \mapsto p_{s} + p_{g} = p'_{s}, \tag{29}$$

$$J_s \mapsto J_s + (p_s \times x_s + p_g \times x_g - p_s' \times x_s'),$$
 (30)

conserving mass, centre of mass, momentum, and angular momentum, respectively. The stored value of J_s does not necessarily correspond to the physical angular momentum of the star, merely the angular momentum within the sink (consisting of the star and a

presumed surrounding gas disc or envelope). Within the sink, the accreted mass is initially stored in the sink's accretion reservoir:

$$M_{\text{acc,s}} \mapsto M_{\text{acc,s}} + m_{\text{g}}.$$
 (31)

Note that our implementation does not address the long-standing issue of violating conservation of magnetic flux when a Lagrangian gas cell is deleted (e.g. Price & Bate 2007). The removal of a gas cell will also generally create a $\nabla \cdot \boldsymbol{B}$ error, and we rely upon our divergence cleaning scheme to damp it away. However, in Section 4.2.3 we show that the main quantities of interest that we wish to predict (star formation histories and the IMF) are in good agreement with results from a constrained-transport AMR code, which does not accrete magnetic flux and enforces $\nabla \cdot \boldsymbol{B}$ to machine precision. ¹⁰

2.5.4 Accretion from reservoir on to protostar

To model the continuous accretion of the protostar for the purposes of modelling protostellar evolution and feedback, we use a simple prescription:

$$\dot{M}_{\star,s} = (1 - f_{\rm w}) \frac{M_{\rm acc,s}}{t_{\rm acc}},$$
 (32)

where $\dot{M}_{\star,s}$ is rate at which mass arrives at the protostar, $f_{\rm w}$ is the fraction of gas mass transferred into the protostellar outflows instead (Section 4.2), and $t_{\rm acc}$ is the accretion time-scale. Both $f_{\rm w}$ and $t_{\rm acc}$ are variable, prescription-dependent quantities (we discuss $f_{\rm w}$ further in Section 4.2 and in Paper 2), but by default we take $t_{\rm acc}$ to be the mean time interval between the arrival of gas cells of mass Δm , assuming the accretion rate is $c_{\rm s}^3/G$:

$$t_{\rm acc} = \frac{G\Delta m}{c_{\rm s}^3} = 530 \,\text{yr} \left(\frac{\Delta m}{10^{-3} \text{M}_{\odot}}\right) \left(\frac{c_{\rm s}}{0.2 \,\text{km} \,\text{s}^{-1}}\right)^{-3},$$
 (33)

which is dimensionally the same as the free-fall time at the maximum Jeans-resolved density, $t_{\rm J} \sim (G\rho_{\rm J})^{-1/2}$. This feeds the protostar with an exponentially smoothed version of the discrete resolved accretion rate, with a 1/e-folding time equal to $t_{\rm acc}$. For the smallest plausible continuous accretion rate in the initial core collapse, $\dot{M}_{\star} \sim c_{\rm s}^3/G$ (Shu 1977), our choice of $t_{\rm acc}$ is simply the mean time interval between the accretion of mass quanta Δm , which guarantees that it limits unphysical discreteness noise without 'oversmoothing' accretion.

Note that the prescription in equation (32) is not meant to model the physical accretion rate at the protostellar surface in detail, and is merely a numerical scheme to obtain a continuous version of the resolved accretion rate with a smoothing time-scale adapted to the mass resolution. If the accretion flow is a direct radial infall (e.g. Bondi accretion) then the relevant physical accretion time-scale is the free-fall time (generally shorter than $t_{\rm acc}$). In the regime where the gas hits an angular momentum barrier before reaching

the protostar, accretion will generally take many orbits, and might be better described by e.g. a Shakura & Sunyaev (1973) α -disc type model (in which the dimensionless parameter α encodes the net effect of gravitational torques, magnetic fields, outflows, and viscosity upon angular momentum transport). In principle, our continuous accretion rate estimator could be fed into a physical model to obtain a more realistic estimate of the rate at which mass arrives at the protostar. However, protostellar accretion on sub-10 au scales is subject to a wide variety of poorly understood complex microphysics (e.g. making the specific choice of α an open problem), so we do not attempt to model such processes here.

2.5.5 Stellar evolution

In simulations with feedback, it is necessary to model the evolution of the protostar or star in the sink particle to inform the emergent luminosity, spectral energy distribution (SED), mass-loss rate, and wind/outflow velocity. We model the star or protostar according to a one-zone sub-grid model whose sole input is the present protostellar mass and the accretion rate (equation 32), originally following Nakano et al. (2000) and based upon the particular implementation of Offner et al. (2009). The model evolves the protostellar radius R_{\star} explicitly, and is calibrated to recover the results of detailed numerical simulations of individual protostellar evolution. This model has been used in many subsequent works by different groups with different codes (e.g. Myers et al. 2014; Federrath, Krumholz & Hopkins 2017; Murray et al. 2018), so we describe it only briefly and refer the reader to Offner et al. (2009) for full details. The evolution is separated into distinct phases:

- (i) *Pre-collapse*: If $M_{\star} < 0.01 \, \rm M_{\odot}$ then the protostar is presumed to be a ≈ 4 au first Larson core that has yet to undergo the second collapse phase (Larson 1969; Masunaga, Miyama & Inutsuka 1998).
- (ii) No burning: Once $M_{\star} > 0.01\,\mathrm{M}_{\odot}$ the core undergoes the second collapse to protostellar density, but deuterium has yet to ignite.
- (iii) Deuterium burning at fixed core temperature: D burning has started, fixing the core of the protostar at $\approx 1.5 \times 10^6$ K.
- (iv) *Core burning at variable core temperature:* The core temperature has begun to rise and D is convected to the core on short time-scales (burning it roughly as rapidly as it arrives at the protostar).
- (v) *Shell deuterium burning*: If D is still arriving rapidly enough, the protostar swells and forms an outer convective zone where the D ignites.
- (vi) Main sequence: The star has reached a central core temperature sufficient to ignite H.

At each time-step, the state of the protostar is updated based upon the present mass, accretion rate, and evolutionary phase, and dictates the evolution of the stellar radius R_{\star} and the emergent luminosity L_{\star} (which includes terms from accretion, Kelvin–Helmholtz contraction, D burning, and H burning, as given in Offner et al. 2009). We use the Tout et al. (1996) fits for the mass-dependent zero-age main sequence luminosity $L_{\rm MS}$ and radius $R_{\rm MS}$, and neglect the effects of stellar evolution beyond the main sequence (apart from modelling a Wolf –Rayet phase for winds, Section 4.3, and an eventual SN for >8 M_{\odot} stars, Section 4.4). For the purposes of modelling SNe, stars >8 M_{\odot} have a mass-dependent stellar lifetime:

$$t_{\star} = 9600 \,\text{Myr} \left(\frac{M_{\star}}{M_{\odot}}\right) \left(\frac{L_{\star}}{L_{\odot}}\right)^{-1} + 3.4 \,\text{Myr},\tag{34}$$

⁹The raw accreted angular momentum of a sink particle is typically of order $\sqrt{GM_sR_{\rm sink}}$, which depends on the numerical parameter $R_{\rm sink}$, and is typically orders of magnitude greater than the angular momentum of a star (Hubber et al. 2013). To determine the actual stellar angular momentum evolution one must model unresolved AM transfer processes.

¹⁰One possible solution for Lagrangian MHD codes (not explored here) would be to introduce a numerical resistivity $\eta_{\rm sink}$ that interpolates between \sim 0 when $r \gg R_{\rm sink}$ and $\eta_{\rm sink} \sim \sqrt{GM_{\star}R_{\rm sink}}$ when $r \sim R_{\rm sink}$, which would diffuse flux away from the star as mass is carried into the sink, modelling the physical non-ideal processes that occur near protostars.




Figure 5. Stellar properties as a function of zero-age main-sequence mass $M_{\rm ZAMS}$ used to model feedback and stellar evolution in STARFORGE. We plot the main-sequence luminosity $L_{\rm MS}$, radius $R_{\rm MS}$, and resulting effective temperature $T_{\rm eff}$ from Tout et al. (1996), the stellar lifetime t_{\star} per equation (34), the flux of H-ionizing > 13.6eV photons $Q_{\rm HI}$ (assuming a black-body spectrum of temperature $T_{\rm eff}$, Section 4.5), the wind mass-loss rates $\dot{M}_{\rm wind}$ and $\dot{M}_{\rm wind,WR}$ for main-sequence and Wolf–Rayet stars, and the wind velocity $v_{\rm wind}$ (with wind quantities assuming solar metallicity, see Section 4.3).

which interpolates between $\sim 10\,\mathrm{Gyr}$ for solar-type stars and $\sim 40\,\mathrm{Myr}$ for $8\,\mathrm{M}_\odot$ stars, and asymptotes to $\sim 3.7\,\mathrm{Myr}$ for the most massive ($\gtrsim 100\,\mathrm{M}_\odot$) stars. In Fig. 5, we plot L_MS , R_MS , t_\star , and various other useful derived quantities for stellar feedback (Section 4) as a function of the zero-age main-sequence mass M_ZAMS .

We do not model the end of life of stars less massive than $8\,M_\odot$ (i.e. planetary nebulae), but this could be important for calculations that run for much longer than a GMC lifetime (e.g. Section 5.1.2). We also presently neglect the formation of relic compact objects, but this would be a trivial modification to the inputs of the SN/wind module (simply reserving a certain relic mass), given a more-detailed stellar evolution prescription.

2.5.6 Merging criteria

In the code, sink particles are allowed to merge if they have a binary semimajor axis $< R_{\rm sink}$ and the secondary has a mass $< 10 \Delta m$. In theory this helps eliminate unphysical, spurious low-mass sinks that may form in proximity to legitimate sinks, or \sim few-au clumps of mass $< 0.01 \, {\rm M}_{\odot}$ that would physically be accreted by a protostar (Offner et al. 2009). In practice, this merger condition is not satisfied in most simulations, and generally only a few times (out of $\gtrsim 1000$ stars) if so. Hence, our results are not sensitive to our sink particle merging strategy. It is possible that physical stellar mergers are a channel for the formation of very massive stars in the centres of dense star clusters (Portegies Zwart et al. 1999; Bonnell & Bate 2005; Shi, Grudić & Hopkins 2020), but we generally run stellar softenings significantly larger than physical stellar radii and hence cannot

follow mergers self-consistently without some kind of sub-resolution modelling. 11

2.5.7 Singular isothermal collapse test

We first validate the formation and resolved accretion criteria of sink algorithm in the Shu (1977) singular isothermal sphere problem, the collapse of a core with an initial density profile $\rho = \frac{Ac_s^2}{4\pi r^2}$, where Aparametrizes the family of solutions and collapse occurs for A > 2. This problem possesses a single central singularity (to be represented by the sink), and admits a semi-analytic, spherically-symmetric solution for all fluid quantities [from the numerical solution of Shu (1977) equations 11 and 12]. This unambiguous reference solution allows it to quickly expose numerical quirks and bugs, whereas testing the sink particle algorithm on e.g. a full turbulent GMC collapse problem is both more expensive and less conclusive because the 'correct' solution (or whether it exists for a given physics setup) is unknown a priori. An insufficiently-strict sink formation prescription (or an overly strict accretion prescription) can result in the formation of multiple spurious sinks when there should be a single singularity. Errors in momentum conservation or gravity can cause the sink to drift from the centre of collapse, causing subsequent gas to arrive off-centre and form spurious discs or sinks. Passing this test does not prove that a sink algorithm is valid for all problems, but failing this test is a strong indicator that the algorithm is flawed.

For reliable test results, the initial conditions should represent the analytic initial density field with equal-mass elements as we use in our simulations, but this is non-trivial. For MFM, we initialize 125 000 equal-mass gas cells on a uniform radial grid (producing the desired r^{-2} density profile) with random initial angular positions, and relax the resulting Poisson sampling noise in the IC to a glass by reversing the sign of gravity and allowing cells to slide around on their respective initial radial shells, with an artificial drag force to damp out the motions towards equilibrium. We then rescale to survey various values of A. Exactly one sink forms in each test, and we plot its mass accretion rate in Fig. 6 for A values ranging from 3 to 1000. Agreement with the semi-analytic solution is excellent across the entire range of A surveyed. We also verify Galilean invariance by running a version of the A = 4 setup with a velocity boost of $100c_s$: even at this extreme bulk Mach number, the solution is preserved owing to the machine-precision Galilean invariance of the hydro, gravity, and sink particle algorithms. The error bars in Fig. 6 plot the $\pm \sigma$ variations of our continuous accretion rate estimator (equation 32), showing that its average error is at most a factor of \sim 2 for the lowest A values and accretion rates, and generally much less for higher accretion rates.

2.5.8 Effect of sink prescriptions

Because we wish to use the properties of sink particles to predict the IMF that emerges from a given set of physics, it is important to ensure that the results of STARFORGE simulations are insensitive to the specific parameter choices made in our sink algorithm for e.g. the density threshold and sink radius, and ideally have some robustness to the specific choice of sink formation and accretion

¹¹One possible approach to stellar mergers is to use the orbital energy and angular momentum (which are conserved absent perturbations) of stellar pairs passing within their respective softening radii to determine the physical, un-softened periastron radius, and hence whether the stars should physically merge.




Figure 6. Simulated accretion rate in the Shu (1977) singular isothermal collapse problem, in units of c_s^3/G , as a function of the instability parameter A, such that the initial density profile is $\rho=\frac{Ac_{\rm s}^2}{4\pi Gr^2}$ and A=2 is the threshold of stability. To the analytic solution (dashed) we compare results simulated with our default hydro, gravity, and sink particle algorithms simulated with the ICs at rest (squares) and with the initial gas moving at Mach 100 to verify Galilean invariance (star). Error bars indicate the $\pm \sigma$ quantiles of the accretion rate estimator used to feed the sub-grid protostar (equation 32) - variance is driven by the discretization of resolved accretion into chunks of mass Δm . Agreement with the analytic solution is excellent, and in all instances exactly one sink particle is formed (replacing the central singularity).

criteria as well. For this we re-run the M2e4 GMC set-up in Paper 0, a $2 \times 10^4 \,\mathrm{M}_{\odot}$ initially spherical GMC of radius 10 pc at $10^{-3} \,\mathrm{M}_{\odot}$ resolution with numerous variations from the prescription described in this section, listed in full in Appendix A. By including only minimal physics (isothermal MHD and gravity), the incremental effects of sink numerics are expected to be more pronounced than in a more complex setup with realistic thermodynamics and feedback. In this sense this test could be considered a worst-case assessment of the sensitivity of STARFORGE results to sink prescriptions.

The reference numerical parameters in this setup are $\Delta m =$ $10^{-3} \,\mathrm{M}_{\odot}, \ R_{\rm sink} = S_{\star} = 18 \,\mathrm{au}, \ \mathrm{and} \ \rho_{\rm J} = 2.6 \times 10^{-14} \mathrm{g \, cm}^{-3}, \ \mathrm{and}$ some tests vary these quantities (Appendix A). We plot the results of this sink parameter study in Fig. 7: the SFE, number of sinks N_* , mass-weighted median sink mass M_{50} , median sink mass M_{med} , and maximum sink mass M_{max} . Our SFE and IMF results are remarkably robust to wide variations in the sink particle prescription and parameters, including ρ_{th} over a factor of 10^6 , and R_{sink} and S_{\star} over a factor of 200. The SFE is particularly robust, with very good agreement across all tests (the one outlier is consistent with a simple delay in accretion). The only setups that produced markedly different results in the IMF were ones with obvious flaws, such as ignoring the density maximum criterion (making the choice of which gas cell to turn into a sink generally non-unique), making R_{sink} much smaller than S_{\star} (making the accretion criteria unreasonably difficult to satisfy, because gravity is unresolved at the sink radius), and increasing R_{sink} and S_{\star} by a factor of >100 (a gross mismatch with the simulation resolution). Moreover our results do not hinge on any one particular ingredient or assumption – neglecting each formation criterion in our prescription in turn had small effects (apart from the density maximum criterion). Stripping down our accretion criteria to

simpler versions also made a negligible difference. Hence, in practice there is a fair amount of redundancy between the different elements of our prescription.

We conclude from this experiment that the results of STARFORGE simulations are unlikely to have any strong dependence upon the details of our sink implementation, at least within the space of Bate et al. (1995)-like algorithms we have explored. Hence, to our knowledge, our sink implementation lacks parameter freedom for 'fine tuning' to ensure a particular desired result - rather, simulation results are mainly sensitive to physical processes as desired. We generally recommend that $\rho_{\rm th}$ and $R_{\rm sink}$ be matched to the nominal density and spatial resolution limits of the simulations (e.g. equations 19 and 20), but e.g. the exact numerical prefactors we have adopted for these quantities are not important, within reasonable limits. Our results hint that simpler prescriptions can perform just as well as ours, but we err on the side of redundancy because the cost of evaluating the various sink formation and accretion criteria is small, and no criterion appears to be unreasonably strict (or else the sink algorithm would allow the simulation to crash due to a runaway gas pile-up).

3 THERMODYNAMICS

Although often idealized as such, GMCs are not truly isothermal, and many potentially important effects in star formation require an explicit treatment of the thermal structure of the ISM, such as the dynamics of fragmentation (Bate et al. 2003; Larson 2005) and the evolution of bubbles driven by wind, radiative, and SN feedback. We evolve the internal energy of the gas according to the MHD equations explicitly (HR16), accounting for all gravitational and MHD work terms with heating and cooling. We explicitly evolve the species Z, He, C, N,O, Ne, Mg, Si, S, Ca, and Fe, and by default assume initial solar abundances (Z, He, C, N,O, Ne, Mg, Si, S, Ca, Fe) = (0.02, 0.28, $3.26\times10^{-3},\,1.32\times10^{-3},\,8.65\times10^{-3},\,2.22\times10^{-3},\,9.31\times10^{-4},\,1.08\times10^{-3},\,6.44\times10^{-4},\,1.01\times10^{-4},\,1.73\times10^{-3}),$ and re-scale abundances appropriately to the desired initial metallicity (but this can be freely varied).

We use a gas equation of state (EOS) with a variable adiabatic index γ , to account for variations in the equilibrium mixture of paraand ortho-hydrogen and the collisional dissociation of H₂ above $\sim 2000\,\mathrm{K}$ (Vaidya et al. 2015). However, we do not roll the heat of ionization into the EOS as in Vaidya et al. (2015), because this is handled separately by our cooling/chemistry solver. We fit to the values of γ given in Vaidya et al. (2015) (neglecting the feature corresponding to ionization) as a function of internal energy:

$$\gamma = \frac{5}{3} + \sum_{k=1}^{5} \delta_k S\left(a_k (\log_{10} u - b_k)\right), \tag{35}$$

where $S(x) = \frac{1}{2} (1 + \frac{x}{\sqrt{1+x^2}})$ is a $\delta_k = (-0.38, 0.22, -0.068, -0.42, 0.65), \quad a_k = (5.95, 6.18,$ 10.26,7.71,98.87), $b_k = (9.25, 9.89,10.24,11.13, 14.28)$, and uis the specific internal energy in cm² s⁻².

We operator-split the adiabatic MHD evolution with a standard implicit cooling algorithm, which solves for equilibrium internal energy, temperature, net cooling/heating rate, mean molecular weight, and ionization state of the gas (treating the adiabatic heating rate from the MHD solver as an additional heating term). Our treatment of cooling and heating terms largely follows the FIRE-2 simulations (described fully in Hopkins et al. 2018b Appendix B), in accounting for free-free, photoionization/recombination, Compton, photoelectric, metal-line, molecular, fine-structure, dust collisional,




Figure 7. Effect of variations in sink particle formation and accretion prescriptions and properties upon the results of a simulation of a $2 \times 10^4 \, M_{\odot}$ GMC of radius 10 pc, including just gravity and isothermal MHD (i.e. re-simulating M2e4 from Guszejnov et al. (2020b) at the same $10^{-3} \, M_{\odot}$ resolution). We plot the SFE (top left), number of sinks N_{\star} (top centre), and the mass-weighted median (M_{50}), number-weighted median (M_{med}), mean (M_{mean}), and maximum (M_{max}) mass statistics of the stellar mass function as a function of time from the beginning of SF. We highlight the most 'extreme' variations: neglecting the density maximum and virial sink formation criteria in turn, reducing R_{sink} by a factor of 1/4 (to ~ 5 au) without reducing the softening, reducing the minimum density for sink formation ρ_{th} by a factor of 10^{-3} , and increasing R_{sink} and the stellar softening S_{\star} by a factor of 100 (to 1800 au). A variety of other overlapping sink parameter/prescription variations are plotted in grey, listed in full in Appendix A. Our predictions are fairly insensitive to most of these variations, provided they are within reasonable physical limits.

and cosmic ray heating and cooling processes. This cooling module has had various evolutionary updates since Hopkins et al. (2018b) that are not important for our results here [e.g. updating to the Faucher-Giguère (2020) UV background, which is similar to the previously-used Faucher-Giguère et al. (2009) UVB at z=0], but will be detailed in full in an upcoming paper (Hopkins et al. 2021, in preparation). Note that our treatment of hot (> 10^6 K) gas considers the dominant radiative cooling mechanisms (i.e. freefree emission and metal lines) as described in Hopkins et al. (2018b), but neglects heat conduction by thermal electrons by default, which may moderate expansion of wind and SN bubbles.

3.1 Background radiation

In the intermediate-density ($\sim 100-10^4 {\rm cm}^{-3}$) gas that makes up the bulk of the mass of GMCs, the thermal structure is set mainly by the balance of photoelectric heating and molecular or fine-structure cooling (Glover & Clark 2012), necessitating some treatment of this background. When modelling solar circle conditions, we assume an isotropic Draine (1978) background $e_{\rm FUV} = 9 \times 10^{-14} {\rm erg \ cm}^{-3}$ for purposes of photoelectric heating, i.e. 1.7 times the Habing (1968) flux of photons in the range 6–13.6 eV. For each gas cell, we evaluate the optical depth to the FUV background on-the-fly using the TreeCol algorithm (Clark, Glover & Klessen 2012), i.e. summing the optical depths of tree nodes grouped into angular bins during the pass through the gravity tree. We default to a

simple 6-bin angular binning of the sky, and assume an opacity of $\kappa_{FUV}=500\, cm^{-2}Z/Z_{\odot}$.

We also model the background radiation due to galactic dust emission as a black-body spectrum with energy density $0.31\,\mathrm{eV}\,\mathrm{cm}^{-3}$ and an effective temperature of 20 K. When we evolve this radiation component with explicit RHD, we simply implement this radiation energy density and temperature as the initial conditions, and allow both to evolve freely (Section 4.5). Without RHD, it is simply held fixed.

The background radiation components quoted here are as measured in the Solar neighbourhood, and can be re-scaled to appropriate values for other environments.

3.2 Dust cooling and heating

Dust cooling and heating are dominant at high ($\gtrsim 10^6\,\mathrm{cm^{-3}}$) ISM densities (Goldsmith & Langer 1978). The dust heating/cooling term Λ_{dust} is (Meijerink & Spaans 2005)

$$\Lambda_{\text{dust}} = 1.12 \times 10^{-32} \,\text{erg s}^{-1} \,\text{cm}^{3} \,\left(T - T_{\text{dust}}\right) \, T^{1/2} \\ \times \left(1 - 0.8 \,\exp\left(\frac{-75}{T}\right)\right) \left(\frac{f_{\text{d}}}{0.01}\right), \tag{36}$$

where T is the gas temperature in K, $T_{\rm dust}$ is the dust temperature, and $f_{\rm d}$ is the local dust-to-gas ratio, which we take to be $f_{\rm d}=0.01Z/Z_{\odot}$ (i.e. assume a constant dust-to-metals ratio equal to the local value) in simulations which do not explicitly follow dust dynamics (otherwise

2214 M. Y. Grudić et al.

this is the actual local value $\rho_{\rm dust}/\rho_{\rm gas}$). How $T_{\rm dust}$ is determined depends on whether or not we are using an explicit RHD solver.

3.2.1 Simulations with explicit RHD

If explicit RHD is enabled, we co-evolve the gas, dust, and radiation temperature self-consistently as in Hopkins et al. (2020a), including the stellar luminosity in various bands accounting for photon transport, absorption and emission using dust opacity fits from Semenov et al. (2003). Dust cooling is handled by including equation (36) as a radiation source term for the IR band, so that energy lost to dust cooling is transported away by the RHD solver. This automatically handles the trapping of cooling radiation in the optically thick limit (setting e.g. the 'opacity limit' for fragmentation, Rees 1976). We explain our RHD treatment fully in Section 4.5.

3.2.2 Simulations without explicit RHD

If an explicit RHD solver is not enabled, we either make the minimal assumption that $T_{\rm dust}$ is constant, or use a simple, inexpensive RT approximation similar to Guszejnov, Krumholz & Hopkins (2016) and Federrath et al. (2017). This approximation uses the LEBRON radiative transfer algorithm (Hopkins et al. 2018b) to estimate the IR radiation energy density from local sources at the position of a gas cell in the gravity tree pass, summing over contributions from all stars:

$$e_{\mathrm{IR},g} = \sum_{s} \frac{L_s}{4\pi r_{\mathrm{gs}}^2 c},\tag{37}$$

where L_s are the respective bolometric luminosities of the sink particles, r_{gs} are the distances from the gas cell to the sinks. We then solve for $T_{\rm dust}$ assuming local equilibrium between absorption and emission according to a $\beta=1$ opacity law [i.e. $\kappa(\nu) \propto \nu$, e.g. Draine 2006]. $T_{\rm dust}$ is then the solution to the quintic equation

$$(T_{\text{dust}}/2.92 \,\text{K})^5 = \sum_{k} (T_{\text{rad}, k}/\text{K}) (e_k/\text{eV} \,\text{cm}^{-3}),$$
 (38)

where the index k runs over three radiation field components with respective energy densities and effective temperatures: the above component from local sources, which is assumed to have $T_{\rm rad,\,IR} = T_{\rm dust,\,IR}$, the CMB with $e_{\rm CMB} = 0.262\,(1+z)^4\,{\rm eV\,cm^{-3}}$ and $T_{\rm rad,\,CMB} = 2.73\,{\rm K}(1+z)$, and the dust-reprocessed component of the interstellar radiation field (ISRF) with $e_{\rm ISRF} = 0.31\,{\rm eV\,cm^{-3}}$ and $T_{\rm rad,\,ISRF} = {\rm max}\,(20\,{\rm K},\,T_{\rm rad,\,CMB})$, with fiducial values appropriate for solar neighborhood conditions, and adjustable depending upon the simulated environment. Note that this differs slightly from Guszejnov et al. (2016) and Federrath et al. (2017), who adopted the optically thick grey-opacity radiative transfer solution in the static diffusion limit $T \propto e_r^{1/4}$, as was found in Offner et al. (2009).

3.3 Optically thick cooling and the opacity limit

If the trapping of cooling radiation is not being solved self-consistently by our RHD solver, we also adopt a simple prescription for interpolating the cooling rate between the optically thin and thick regimes. If the net absolute heating/cooling rate is $|\Lambda_{Net}|$, then we enforce

$$|\Lambda_{\rm Net}| < \Lambda_{\rm BB}$$
 (39)

$$\Lambda_{\rm BB} \equiv 5.67 \times 10^{-5} \, T^4 \, \left(\frac{\mu}{a \Sigma_{\rm eff}}\right) \frac{1}{1 + \kappa_{\rm eff} \, a \Sigma_{\rm eff}} \, n_{\rm H}^{-1}, \tag{40}$$




Figure 8. Thermal evolution of the densest gas cell in the centre of a collapsing M_{\odot} core as a function of its density ρ_{max} , replicating the test problem in Masunaga et al. (1998) with our default physics using our self-consistent M1 RT solver with gas–dust–radiation coupling (Section 4.5) and our simple optically thick cooling approximation based upon the TreeCol algorithm (equation 40).

where $\mu \approx 2.4$ is the mean molecular weight and $\Sigma_{\rm eff}$ is estimated via the TreeCol algorithm, $\kappa_{\rm eff}$ is the effective opacity detailed in Hopkins et al. (2018b), and a=0.2 is an uncertain geometric factor chosen to reproduce detailed RHD protostellar collapse calculations (Masunaga et al. 1998). This limits the cooling (or heating) rate to the bound for a slab geometry derived in Rafikov (2007). This is still approximate, but is more realistic than an 'effective equation of state' that transitions from isothermal to adiabatic (e.g. Bate et al. 2003; Bate 2009a): we are still always allowing for heating and cooling, with radiation terms that explicitly account for optical depth, which is the physically relevant quantity for radiative cooling. In a single isolated collapsing clump in a spherical geometry, an equation of state may be justified because $\Sigma_{\rm eff} \approx \rho \lambda_{\rm J}$, but this does not hold in general, and particularly not in systems that are optically thick to dust globally ($\Sigma_{\rm eff} \gtrsim 1 {\rm g \, cm}^{-3}$).

3.4 Tests

To test the code's ability to capture the transition from isothermal to adiabatic behaviours as the ISM gets optically thick to cooling radiation (important e.g. for the opacity limit for fragmentation), we simulate collapse of a 1 M_☉, uniform-density Jeans-unstable core with both methods described here: explicit RHD with the M1 solver and the simpler TreeCol-based prescription (equation 40). We initialize the cloud with a mass resolution of $10^{-5}\,M_{\odot}$ (sufficient to marginally resolve the first Larson core, Bate et al. 2003), arranging the cells in a uniform-density glass configuration with density $5.3 \times 10^{-18} \mathrm{g \, cm^{-3}}$, and assume solar metallicity with a dust-to-gas ratio of 0.01, as has been simulated by many other RHD studies (Larson 1969; Masunaga et al. 1998; Vaytet & Haugbølle 2017). In Fig. 8, we plot the thermal evolution of the centre of the core as a function of its density, showing good agreement with previous calculations: in all instances, the transition from isothermal ($T \sim \text{const.}$) to adiabatic $(T \propto \rho^{2/5})$ evolution occurs at $\sim 10^{-13} \mathrm{g \, cm^{-3}}$. Small




Figure 9. Illustration of the mesh-free local injection procedure used for coupling photons and stellar winds with unresolved free expansion to gas cells in STARFORGE (adapted from Hopkins et al. 2018a). Local injection adds momentum, energy, and either mass or photons to pre-existing interacting gas cells (coloured domains with circles representing the mesh-generating points) around the sink particle (red star), in proportion to the solid angle subtended by each cell according to the 'effective' faces constructed around each source from the neighbouring mesh-generating points (thick black lines). We use a weighting scheme that ensures statistical isotropy and exact momentum and energy conservation, and in the case of photons accurately accounts for unresolved extinction between the star and cell centre (Section 4.5).

residual differences between the four calculations at the level seen here are expected, due to varied assumptions about radiation initial and boundary conditions, dust properties and opacities, molecular equation of state, etc.

4 FEEDBACK

We now describe the respective physical model and numerical implementations of each feedback mechanism in turn. For those with well-understood or analytic solutions (winds, SN, and radiation), we will verify that each module performs accurately. Where such solutions are not available (e.g. jets), we will do the next best thing: test for robustness to numerical details and agreement with other codes.

4.1 Generic injection algorithms

4.1.1 Local injection

When coupling feedback in the form of mass, momentum, and energy from stars, a natural approach is analogous to the manner in which the MHD equations are solved: simply determine the fluxes of those respective quantities at the faces of surrounding gas cells, or more generally distribute those quantities in some weighted fashion in the local hydrodynamic stencil – we refer to this technique as 'local injection' (illustrated for a meshless/unstructured mesh code in Fig. 9). This is what is done in virtually all grid codes and is the technique typically used in GIZMO simulations for local radiative feedback, SNe, and stellar winds (Hopkins et al. 2018a, b, 2020a). We adopt this algorithm for feedback coupling (described in full in Hopkins et al. 2018a) where appropriate: for stellar winds when the free-expansion radius is unresolved (Section 4.3) and




Figure 10. Illustration of the cell spawning technique for injecting mass return from sinks (winds, jets, and SNe, Section 4.1.2). We create new gas cells of mass $\Delta M_{\rm w}$ (blue circles) at the sink radius (or at a fraction of the local inter-cell spacing, whichever is smaller), at antipodal positions and velocities to ensure conservation of centre of mass and momentum. The angle θ from the sink angular momentum vector $J_{\rm sink}$ can be controlled to allow arbitrarily-collimated injection. $\Delta M_{\rm w}$ can be set smaller than Δm (the 'normal' mass resolution of ambient gas) to improve resolution in diffuse feedback-driven cavities. Note that this depicts the launching of a pair of gas cells (as for jets and winds), but for SN we launch 24 at once in an angular grid (Section 4.4), designed to ensure exact conservation and isotropy.

for photon injection (Section 4.5). This method ensures machineprecision conservation of mass, momentum, and energy, and ensures that feedback is injected isotropically (or according to the desired angular weighting) even when the local spatial arrangement of cells is anisotropic (unlike e.g. a simple kernel weighting).

4.1.2 Cell spawning

Local injection with Lagrangian methods can run into a major challenge: the resolution is concentrated where gas dens, but feedback structures such as jet cavities, SN remnants, and wind bubbles can be very diffuse. Moreover, if feedback is driving mass away then this problem grows worse over time. And if the inter-cell spacing becomes sufficiently large, it may cease to be a good approximation to instantaneously dump mass, momentum, and energy, because the gas has finite traveltime. Moreover, if we restrict injection to the nearest neighbour cells we cannot inject outflows more collimated than the solid angle subtended by a neighbouring cell (which can be large). So where local injection is not feasible or appropriate, we instead *create* new gas cells, in a procedure we refer to as 'cell spawning' (Fig. 10). A similar technique has been used previously in SPH simulations of stellar winds (Dale & Bonnell 2008) and protostellar outflows (Rohde et al. 2019), and recently in GIZMO to simulate AGN jets (Torrey et al. 2020). We adapt it here to simulate protostellar jets, stellar winds (when resolvable), and SN ejecta.

2216 *M. Y. Grudić et al.*

Cell spawning can be viewed as the inverse of the gas cell deletion operation that occurs during sink accretion: a new cell is created at a certain position with a certain mass, velocity, and internal energy, and the volume partition in its vicinity is re-computed to accommodate it in the next density iteration. We take the distance between the centre-of-mass mesh-generating point of the spanned cell and the sink to be

$$R_{\text{spawn}} = \min \left(R_{\text{sink}}, \Delta x_s / 2 \right), \tag{41}$$

where Δx_s is the average inter-cell spacing in the vicinity of the sink. We prescribe the initial radial direction and velocity according to the desired angular pattern of the feedback mechanism being realized (Sections 4.2–4.4). We assign purely radial initial velocities, but non-radial velocities could potentially be used to model angular momentum transport (Federrath et al. 2010). Spawned cells are assigned an initial temperature of 10^4 K, and a very small, random initial magnetic field scaled such that the initial plasma β is 10^6 .

Spawning is allowed to occur when the sink's internal respective feedback reservoir (Fig. 4) contains enough mass to produce at least $N_{\rm spawn} \times \Delta M_{\rm w}$ cells, where the number of cells spawned at a time $N_{\rm spawn}$ and spawned cell mass resolution ΔM_w are specified for the respective feedback channel. Note that cells do not necessarily have to be spawned with a mass resolution equal to the nominal average 'ambient' gas cell mass Δm : rather $\Delta M_{\rm w}$ can be chosen to be smaller to achieve better time resolution of feedback and to improve spatial resolution within diffuse feedback cavities. For jets and winds, we have generally found the choice $\Delta M_w = 0.1 \Delta m$ to be a good compromise between computational cost and resolution. For SNe, we simply take $\Delta M_w = \Delta m$. With these choices, we note the spatial resolution in diffuse feedback bubbles will generally be fairly coarse (~ 1 pc) for a typical $\Delta m \approx 10^{-3}$ M_{\odot}, possibly making it challenging to resolve channels and leakage of hot gas (e.g. Rogers & Pittard 2013).

Care must be taken when handling MHD interactions between cells of greatly differing masses and sizes in MFM, particularly for new cells which change the local volume partition and can perturb $\nabla \cdot \boldsymbol{B}$ (interacting with out $\nabla \cdot \boldsymbol{B}$ cleaning algorithm). Spawned cells with $\Delta M_{\rm w} < \Delta m/2$ default to a lower-order but morerobust reconstruction in the Riemann problem. We also limit the magnitude of the oriented effective face area $A_{gg'}$ between cells (all interacting cells, wind or not) to the lesser of their geometric areas, $A_{\rm max} = \min\left(\pi h_g^2, \pi h_{g'}^2\right)$. A spawned cell g is merged into a normal cell g if they are hydrodynamically interacting neighbours, they are moving towards each other, and $|\boldsymbol{v}_g - \boldsymbol{v}_{g'}| < \min\left(c_{s,g}, c_{s,g'}\right)$. We have found that this is a good indicator that the spawned cell has merged with the surrounding ISM, and hence its additional resolution is no longer required.

Lastly, to ensure physical and convergent results, it is important for any feedback algorithm to ensure conservation of momentum and centre of mass. We achieve this by always spawning cells in multiples of 2, such that each cell has an antipodal counterpart in the opposite direction, giving machine-precision conservation.

4.2 Jets

4.2.1 Physics prescription

Collimated, bipolar protostellar outflows (jets) have special importance as a feedback mechanism. Many works have demonstrated their potential importance both in setting the IMF and SFE, because they are the only channel can be both prompt and omnipresent, immediate

ately regulating the growth of individual stars without requiring e.g. massive stars with dynamically relevant wind or radiative fluxes to be present (Matzner & McKee 1999a; Krumholz et al. 2019). We find that they are likely to be an important ingredient for the IMF in Paper 2, consistent with previous studies (Krumholz, Klein & McKee 2012; Myers et al. 2014; Federrath et al. 2014; Li et al. 2018; Cunningham et al. 2018). The details of the MHD jet launching mechanism (e.g. whether it is better-described by the canonical 'X-wind' or 'D-wind' models, Pudritz & Norman 1983; Shu et al. 1994) are the subject of active research, and depend upon a variety of complex microphysics operating at sub-au scales that are not practical to resolve in our simulations (although the simulations may provide important context for subsequent 'zoom-in' studies of individual stars). Thus, we model jets according to a simple phenomenological prescription following Cunningham et al. (2011), parametrizing the jets' properties in three parameters: the fraction $f_{\rm w}$ of mass accreted by the envelope-discstar system that is diverted to the jet (see Fig. 4), the fraction f_K of the Keplerian velocity at the protostellar radius R_{\star} at which jets are launched, such that

$$v_{\rm jet} = f_{\rm K} \sqrt{\frac{GM_{\star}}{R_{\star}}},\tag{42}$$

and the collimation angle θ_0 , such that the angular distribution of injected wind momentum is given by (Matzner & McKee 1999b):

$$\xi(\theta, \theta_0) = \left(\ln\left(\frac{2}{\theta_0}\right)\sin^2\theta + \theta_0^2\right)^{-1},\tag{43}$$

where θ is the angle with respect to the angular momentum axis of the sink J_s of the sink. This concentrates injection in a narrow cone of angular size $\approx \theta_0$ about the angular momentum axis of the star (see Fig. 11). Our default choices are $f_{\rm w}=f_{\rm K}=0.3$ and $\theta_0=$ 0.01, following Cunningham et al. (2011) and other authors. These choices put the momentum loading $f_w f_K$ roughly in the middle of the observed range (see Cunningham et al. 2011, section 2.4; Federrath et al. 2014, section 3.5; and references therein). The values of f_w and f_K do matter: in Paper 2, we find that the product $f_W f_K$ affects the IMF peak. The specific value of θ_0 is likely to be unimportant because in any realistic turbulent accretion scenario J_s will generally tend to precess during accretion over an angular region much larger than θ_0 (Rosen & Krumholz 2020), and even without precession jet cavities will expand in the perpendicular direction, opening up an ever-increasing solid angle (Arce & Sargent 2006; Offner et al. 2011). In our tests, we will show that our results are insensitive to variations in θ_0 of at least a factor of 10, which is consistent with prior hydrodynamic outflow simulations carried out by Offner & Arce (2014).

4.2.2 Numerical methods

We always couple jets via cell spawning (Section 4.1.2), waiting until sufficient mass is available in the jet reservoir to spawn two cells of mass ΔM_w . The angular direction with respect to the sink angular momentum \mathbf{J}_s is sampled randomly for the first cell from equation (43), and the second cell is pointed in the opposite direction, conserving momentum and centre of mass. Our prescription ignores both the angular momentum and magnetic flux content of the jet material. Although outflows can be the dominant mechanism of angular momentum transport within the disc (i.e. on scales smaller than the sink radius), the material in the disc already had to have very low angular momentum to get to the base of the jet, so it is unlikely that this angular momentum has important effects on larger




Figure 11. Examples of the protostellar jet module (Section 4.2) in action in SF simulations. Left: Idealized laminar rotating core collapse problem forming a single star, run at high $(10^{-5}\,M_\odot)$ resolution. As the star accretes from a disc, jets clear out high-velocity diffuse cavities along the poles, entraining material away from the core. Right: Bipolar outflows (highlighted in orange) permeate a $2\times10^4\,M_\odot$ GMC run at $10^{-3}\,M_\odot$ resolution (typical of STARFORGE runs), with the largest penetrating out to $\sim10\,\mathrm{pc}$ scales before merging with the ISM. This map colours by 1D line-of-sight velocity dispersion (purple is $\sim0.1\,\mathrm{km\,s^{-1}}$, orange is $\sim10\,\mathrm{km\,s^{-1}}$ and modulates the lightness to encode surface density information (lighter is denser).

scales once transported back out in the jet (however it is considered by other jet feedback models, such as by Federrath et al. 2014 and Rohde et al. 2019). The effects of the magnetic field in the jet are less readily dismissed; however, this would be nontrivial model in the present numerical framework due to the problem of determining what initial **B** should be assigned to the spawned cells while observing conservation laws.

We plot some examples of the effects of the jet module in simulations in Fig. 11. We generally observe realistic-looking structures in the simulations, with broad bipolar cavities penetrated by a narrow jet, surrounded by a disc (if angular momentum support is important) or a pseudo-disc (of infalling material funnelled by the bipolar cavities). $v_{\rm jet}$ tends to increase as the star accretes (because R_{\star} varies only weakly in equation 42), so the jet tends to catch up with itself, piling up and cooling in a plume-like region reminiscent of Herbig–Haro objects (e.g. Bally 2016).

Lacking a test problem for our jet module that admits an analytic or universally agreed-upon solution, we resort to heuristic methods to validate it: checking for numerical convergence, checking for robustness to uncertain or arbitrary numerical parameters, and finally comparing with another published solution from a different code.

4.2.3 Resolution tests

Here we consider the effects of numerical resolution upon a simulation of the M2e3 GMC model introduced in Paper 0 with gravity, MHD, the cooling module without explicit RT or protostellar radiation (Section 3), and the jet module enabled. The initial condition is a spherical, uniform-density GMC with mass $2\times 10^3\,M_\odot$, radius 3 pc, and virial parameter 2 (for an initial turbulent Mach number of \sim 9). The GMC is surrounded by a diffuse medium with 1/1000 the

density filling a 30 pc box, and the initial magnetic field is uniform throughout the box with a strength of 2.3 μ G. The normal mass resolution Δm varies from 0.1 to 10^{-4} M $_{\odot}$, and the jet mass resolution $\Delta M_{\rm w}$ varies from 0.01 to 10^{-5} M $_{\odot}$. The sink formation density threshold, sink radius, and sink softening scale are varied according to the mass resolution (section 2.5), scaling $\rho_{\rm th} \propto \Delta m^2$ over a factor of 10^6 between 3×10^{-17} and 3×10^{-11} g cm $^{-3}$, and scaling $R_{\rm sink} = S_{\star} \propto \Delta m$ from 1800 to 1.8 au. Hence, there is no purely-numerical, resolution-related quantity that is held fixed in our resolution study, so the possibility of inferring 'false' convergence is ruled out. 12

In the top left-hand panel of Fig. 12, we plot the evolution of the SFE for the 10 simulations in our resolution study. Star formation tends to start sooner at lower resolution, opposite to what is seen in the collapse of Jeans-mass clumps (Fig. 2), possibly owing to increased numerical dissipation of turbulence at low resolution. The star formation history ceases to appear sensitive to resolution below $\approx 0.01 \, \mathrm{M}_{\odot}$. This corresponds to the resolution criterion we derived in Paper 0, that the sonic mass $M_{\mathrm{sonic}} \approx M_0 M^{-4} \approx 0.3 \, \mathrm{M}_{\odot}$ be resolved in at least ≈ 30 gas cells.

We examine the resolution dependence of the stellar mass spectrum at fixed total stellar mass in panels 2-6 of Fig. 12: the

¹²Scaling all purely numerical, dimensional sink-related quantities with resolution is important for resolution studies in SF simulations. Otherwise, it is possible that the constant value of e.g. the density threshold or sink radius imprints a characteristic Jeans mass or length, leading one to falsely infer convergence. In near-isothermal problems with Lagrangian codes, this entails scaling $\rho_{\rm th} \propto \Delta m^2$ and $R_{\rm sink} \propto \Delta m$. For AMR codes, one must scale $\rho_{\rm th} \propto \Delta x_{\rm min}^{-1/2}$ and $R_{\rm sink} \propto \Delta x_{\rm min}$, where $\Delta x_{\rm min}$ is the spatial resolution at the finest refinement level (and assume a Truelove et al. (1997) Jeans refinement scheme). AMR resolution studies may also need to scale the based-level grid resolution to resolve turbulent fragmentation (Haugbølle et al. 2018), i.e. fixing the number of refinement levels.




Figure 12. Resolution study of a simulation of a $2000 \, \rm M_{\odot}$ GMC with initial radius 3 pc, with MHD, gravity, cooling physics, and protostellar jets (Section 4.2). We vary the mass resolution Δm from 10^{-4} –0.1 $\rm M_{\odot}$ (darker is finer), scaling sink accretion and softening radii $\propto \Delta m$ from 2-2000 au (equation 22) and scaling the sink density threshold $\propto \Delta m^{-2}$ (equation 19), so that there is no characteristic numerical scale that could cause false convergence. We plot the resulting star formation history and the mass-weighted median, mean, and median stellar masses M_{50} , $M_{\rm mean}$, and $M_{\rm med}$, and black dashed lines correspond to the value for a Kroupa (2002) IMF. SFE and IMF statistics cease to scale systematically with resolution past a certain threshold.

number of stars N_{\star} , the mass-weighted median stellar mass M_{50} , and the median, mean, and maximum stellar masses. All IMF statistics, whether mass- or number-weighted, eventually cease to change systematically with resolution. As in the isothermal case, the resolution threshold for M_{50} and $M_{\rm max}$ to stabilize is $\approx 0.01\,{\rm M}_{\odot}$. Number-weighted statistics require somewhat higher resolution, with $\approx 10^{-3}\,{\rm M}_{\odot}$ being the marginal value for accurately predicting the mean stellar mass and $4\times 10^{-4}\,{\rm M}_{\odot}$ for the median. We also plot the effects of resolution upon statistics taken over different stellar mass cuts in Appendix B, finding that the resolution required depends upon the mass cut (consistent with a simple resolution-dependent low-mass incompleteness effect, with incompletness starting below $< 100\,{\rm A}\,m$)

By design, there is no purely numerical dimensional quantity that could imprint a characteristic stellar mass here, so it is likely that the predicted IMF is shaped largely by the physical processes modelled in the simulation, i.e. there may exist a well-defined, physical IMF that emerges from the combined physics of cooling, MHD, gravity, stellar dynamics, and protostellar outflows, and this IMF resembles the observed one. We explore this IMF prediction across a wide parameter space in Paper 2. Assuming that other feedback processes do not demand further resolution requirements, this experiment gives some idea of the mass resolution needed to predict e.g. the mean stellar mass in STARFORGE simulations. Note that some incompleteness in the IMF may persist to higher resolution - to obtain a complete IMF, we may require a resolution of $\sim 10^{-5} \,\mathrm{M}_{\odot}$ to fully resolve the collapse of clumps at the opacity limit, forming the smallest brown dwarfs (Bate et al. 2003). But brown dwarfs contain only a small fraction of the total number and mass in stars and are not expected to

exert significant feedback. Hence, many major questions involving cluster formation, feedback, and the physics underlying the *typical* mass of stars can be addressed at much lower resolution. We adopt $10^{-3}\,\rm M_\odot$ as our standard resolution for these purposes, but note that the coincidence of our results here with the Paper 0 sonic mass resolution criterion $\Delta m \lesssim 0.03 M_{\rm sonic} \sim 0.03 M_0 M^{-4} \approx 0.01 M_\odot (T/10\,\rm K)^2 \left(\Sigma/100 M_\odot \,\rm pc^{-2}\right)^{-1}$ suggests that this is the more-general convergence criterion. This resolution criterion can be more demanding for e.g. high-surface density GMCs found in the Galactic centre (Oka et al. 2001; Longmore et al. 2012) but not necessarily because such clouds can also be warmer.

We take this opportunity to comment on the computational cost and scaling of these simulations with feedback. Our fiducial resolution ($10^{-3} \,\mathrm{M}_{\odot}$, 2 × 10^{6} cells) run cost roughly 10 000 corehours run on the Frontera supercomputer at the Texas Advanced Computing Center equipped with 2.7 GHz Intel Xeon 'Cascade Lake' processors (56 cores per node), and the simulations in the following section at the same resolution all had comparable cost. The largest simulation shown in Fig. 12 (mass resolution $10^{-4}\,\mathrm{M}_\odot$, 2 × 10^7 gas cells) required roughly 100 000 core-hours, running for 17 wall-clock days on 4 Frontera nodes. The largest STARFORGE simulation with jet feedback run so far ($M_{\rm GMC}=$ $2 \times 10^5 \,\mathrm{M}_{\odot}$ with 2×10^8 cells, Paper 2) required roughly 4.8M core-hours in 70 wall-clock days on the Stampede-2 machine at the Texas Advanced Computing Center with 2.1 Ghz Intel Xeon 'Skylake' processors (24 cores per node). Note that these are numbers for simulations without explicit RHD; we anticipate that our full RHD STARFORGE simulations currently in progress will be a factor of \sim 5-10 more expensive than their non-RHD




Figure 13. Effects of various numerical and physical variations upon the star formation history, kinematics, and IMF evolution of the same $2000 \,\mathrm{M}_\odot$ GMC simulated in Fig. 12, at $10^{-3} \,\mathrm{M}_\odot$ resolution. Quantities are as described in Fig. 12, plus the evolution of the 3D mass-weighted RMS Mach number M (bottom left). Reference: Baseline settings, with cooling, MHD, and protostellar jets enabled. No jets: Jet module disabled. No jets; $f_{\rm acc} = 0.5$: Jet module disabled, and sinks delete half the mass that they accrete. Protostellar heating: Including the approximate protostellar heating prescription described in Section 4.1.2, with physical and artificially large (\times 10) luminosities. Isotropic jets: spawning jet cells in random (versus highly collimated, equation 43) directions. $\theta_0 = X$: Changing the jet collimation angle θ_0 (equation 43) from the standard value of 0.01. No MHD: Setting the magnetic field to $0.t_{\rm acc} \times 10$: Scaling the accretion smoothing time-scale (33) \times 10. No jets if <X: Jets are disabled for stars with mass < X. $\Delta M_{\rm w} = \Delta m$: Setting the jet cell mass resolution to the nominal mass resolution Δm , as opposed to the standard $\Delta m/10$. Ang. mom return: Returning accreted angular momentum to surrounding gas as in Hubber et al. (2013).

counterparts, mainly due to their more stringent time-step constraints.

4.2.4 Effect of physics variations and numerical details

In Fig. 13, we explore the effects of 13 other variations on this set-up (both numerical and physical) upon the same SFH and IMF statistics, as well as the GMC kinematics. Neglecting jet feedback altogether results in much higher terminal SFE, even if we delete half the accreted mass from the simulation, a simple prescription used by previous works to model jet feedback, deleting either on-the-fly or in post-processing (Padoan, Haugbølle & Nordlund 2012; Federrath & Klessen 2012; Haugbølle et al. 2018). Models that do not explicitly treat feedback also seriously underestimate the level of turbulence in the GMC, and predict a much more top-heavy IMF (e.g. greater values of M_{50}). Haugbølle et al. (2018) found that deleting half the accreted mass gave a good fit to the observed IMF, but simulated a somewhat different GMC set-up, and it can easily be reasoned on dimensional grounds that the accretion efficiency one needs to emulate the effect of jets could generally be problem-dependent. Overall, runs with jet feedback behave dramatically different to the baseline run with feedback: both the total stellar mass formed and average individual stellar masses are roughly an order of magnitude smaller, because the momentum content of the jets disrupts both local protostellar accretion flows and eventually the cloud itself (see Paper 2).


All results are fairly insensitive to variations in the collimation angle θ_0 from 0.01 to 1, the jet mass resolution $\Delta M_{\rm w}$ from 0.1 to $1\Delta m$, and the accretion smoothing time-scale $t_{\rm acc}$ from 1 to $10 \times \Delta m c_{\rm s}^3/G$ (equation 33). Results are also insensitive to whether we allow jets

from stars $<0.1\,M_{\odot}$, whether we include protostellar heating, and whether we return accreted angular momentum to the surrounding gas as in Hubber et al. (2013) (which would influence the stellar angular momenta and jet directions in turn). Results only differ significantly for variations that are ruled out observationally and/or unphysical: neglecting magnetic fields (resulting in a much more bottom-heavy IMF, Guszejnov et al. 2018), assuming jets are emitted isotropically (reducing the typical stellar masses and increasing the overall SFE), artificially increasing the coupled protostellar heating luminosity by a factor of 10 (which made the IMF noticeably more top-heavy, as in Krumholz et al. 2012), and disabling jets for all stars $<1\,M_{\odot}$ (which also made the IMF more top-heavy).

4.2.5 Comparison with AMR simulations

To conclude our testing of the jet module, we test its results against a code that differs from ours in all regards except for the physical equations solved and the physical assumptions underlying our feedback models. Our objective here is to verify that the results of simulations with jets are robust to such details.

We have reproduced the set-up described in Cunningham et al. (2018), who simulated low-mass cluster formation in a periodic box of side length 0.65 pc containing $185\,\mathrm{M}_\odot$ with the ORION2 AMR constrained-transport MHD code with radiative transfer and protostellar jets (Cunningham et al. 2011; Li et al. 2012). We re-run their driven turbulence simulation with an initial mass-to-flux ratio $\mu=2.17$, initially driving turbulence at $M\sim6.6$ for $\sim1\,\mathrm{Myr}$ and then switching on gravity. We did not use exactly the same turbulent driving pattern, so our results should only be compared statistically, hence we ran an ensemble of simulations from






Figure 14. Comparison between the results of the driven $\mu = 2.17$ simulation in Cunningham et al. (2018) with MHD and protostellar heating and jets, and three different STARFORGE replications at standard $10^{-3} \, \mathrm{M}_{\odot}$ resolution, using our simple approximate RT treatment of dust heating and our protostellar jet module (modified so that $v_{\mathrm{jet}} = \min\left(\sqrt{GM_{\star}/R_{\star}}, 60 \, \mathrm{km \, s^{-1}}\right)$, as in Cunningham et al. 2018). Top: Star formation efficiency versus time. Bottom: Stellar mass function at the final SFE of ~ 8 per cent, comparing the stacked mass function from the three STARFORGE runs with Cunningham et al. (2018), showing the respective median stellar masses as vertical lines.

three different initial turbulence realizations. We adopted the same modified prescription for the jet speed as Cunningham et al. (2018), $v_{\rm jet} = {\rm min}\left(\sqrt{GM_\star/R_\star}, 60\,{\rm km\,s^{-1}}\right)$, and the two codes' respective protostellar evolution modules setting R_\star both follow Offner et al. (2009). Cunningham et al. (2018) account for protostellar radiation with a flux-limited diffusion solver, while we use the inexpensive tree-based approximation described in Section 3. Our simulations adopt a mass resolution of $10^{-3}\,{\rm M}_\odot$ and cost roughly 200 core-hours each when run on a single Frontera CLX node.

We plot the resulting star formation history and IMF at 8 per cent SFE in Fig. 14. The specific shape of the SF history does appear to be a function of the initial turbulent driving, but we had two seeds that matched that of Cunningham et al. (2018) for an appreciable fraction of the SF history, both modulo a small time difference due to the different initial turbulent states. Therefore predictions regarding the regulation of SF due to protostellar feedback appear similar for the two codes.

The final IMFs are also in fair agreement: the median stellar mass predicted by both codes is $\sim 0.15 \, \mathrm{M}_\odot$ (shown as vertical lines). The only statistically significant difference is between the predicted numbers of $0.01-0.03 \, \mathrm{M}_\odot$ brown dwarfs, with STARFORGE runs finding only $\sim 1/4$ as many. This may be due any number of details in the cooling and RT modules that differ (with more higher temperatures or more radiative heating suppressing brown dwarfs, Bate 2009b; Offner et al. 2009), or a mere resolution effect (as we expect some numerical IMF incompleteness at masses $\lesssim 30 \, \Delta m$).

In summary, the respective implementations of STARFORGE and ORION2 find very similar results for the regulation of SF and the

stellar mass range of the IMF in low-mass star cluster formation, despite these two codes' detailed numerical implementations differing in every regard. We scale a set-up similar to this to GMCs as much as $>1000\times$ more massive in Paper 2, exploring the broader implications of protostellar feedback in massive GMCs.

4.3 Stellar winds

4.3.1 Physics prescription

We allow main-sequence stars more massive than $2\,M_\odot$ to inject stellar winds. Stellar mass loss rates are subject to considerable theoretical and observational uncertainties, with various unresolved discrepancies between theory and observations (Smith 2014), so we default to a simple phenomenological prescription. Wind-emitting stars feed their wind reservoir from the stellar mass at a base rate of

$$\frac{\dot{M}_{\text{wind}}}{M_{\odot} \text{yr}^{-1}} = \min \left(10^{-6} L_{\text{MS}}^{1.5}, 10^{-7.7} L_{\text{MS}}^{2.9} \right) Z_{\star}^{0.7}, \tag{44}$$

where the main-sequence luminosity $L_{\rm MS}$ and the metallicity Z_{\star} are in solar units. This is a fit to the envelope of the 'de Jager/3' and 'weak wind problem' scalings given in Smith (2014), hence it is a conservative model accounting for the widely acknowledged overestimation of \dot{M} by theoretical line-driven stellar wind models [i.e. it is generally weaker than widely used models such as Vink, de Koter & Lamers (2001)]. The velocity of the winds is

$$v_{\rm wind} = \sqrt{\frac{2GM_{\star}}{R_{\star}}} \times \begin{cases} 0.7 & T_{\rm eff} < 12500 \,\mathrm{K} \\ 1.3 & 12500 \,\mathrm{K} < T_{\rm eff} < 21\,000 \,\mathrm{K} \\ 2.6 & T_{\rm eff} \ge 21\,000 \,\mathrm{K} \end{cases} , \quad (45)$$

following Lamers, Snow & Lindholm (1995).

Much of the energy and momentum in stellar winds from a stellar population originates in Wolf–Rayet stars. We use a simple model for the Wolf–Rayet phase for $> 20\,\mathrm{M}_\odot$ stars, multiplying \dot{M}_wind by a factor of 10 at the end of its lifetime. The time spent in the Wolf–Rayet phase is given by

$$t_{\rm WR} = 1.5 \,\mathrm{Myr} \,\mathrm{min} \left(1, \frac{M_{\star}/\mathrm{M}_{\odot} - 20}{80}\right) \left(\frac{Z_{\star}}{\mathrm{Z}_{\odot}}\right)^{0.5},\tag{46}$$

an approximate fit to results from Meynet & Maeder (2005).

4.3.2 Numerical methods

Our numerical method for coupling winds uses either local injection or cell spawning, where appropriate. In the regime where the wind's free-expansion radius $R_{\rm free} = \sqrt{\dot{M}_{\rm wind}}/v_{\rm wind}\rho$ is much less than the size of a wind cell $\Delta x_{\rm w} = (\Delta\,M_{\rm w}/\rho)^{1/3}$, a spawned cell will generally stop within a single cell length, collide with the ISM, and thermalize its kinetic energy, so it is more efficient and accurate to instantaneously inject that mass, momentum, and energy isotropically into the neighbouring cells. But when the free-expansion radius is well resolved, the simulations resolve the traveltime of the wind before it merges with the ISM, so cell spawning is more appropriate. Hence we switch between the two modules adaptively, based on whether $R_{\rm free}$ is resolved by at least 1 wind cell length. As with jets we spawn cells 2 at a time, but with an isotropic angular distribution instead of collimated.

4.3.3 Tests

We test this module on the problem of the self-similar expansion of a wind bubble propagating into a uniform medium with an




Figure 15. Self-similar expansion of a stellar wind bubble with an adiabatic interior and radiative outer shell, as simulated with our stellar wind module (Section 4.3) for a star with $\dot{M}_{\rm wind}=10^{-5}{\rm M}_{\odot}$ and $v_{\rm wind}=3000~{\rm km\,s^{-1}}$ in a 16 pc box containing $2\times10^4{\rm\,M}_{\odot}$, using different numerical methods (local injection (Section 4.1.1), cell spawning (Section 4.1.2), and a hybrid method that switches between them adaptively (Section 4.3). Top: Comparison of the evolution of the radius of the swept-up shell $R_{\rm shell}$ (corresponding to the number of swept-up gas cells $N_{\rm shell}$) to the known similarity solution, $R_{\rm shell}=0.763\left(\frac{L_{\rm wind}}{\rho}\right)^{1/5}t^{3/5}$ (Weaver et al. 1977). Bottom: Numerical velocity anisotropy of the bubble (= 0 in the similarity solution), generally \lesssim 10 per cent except in very early phases when the bubble is not well-resolved (contains just a few cells). The switch between different coupling methods at 0.02 Myr is apparent when the 'Hybrid' curve deviates from the 'Local Injection' curve.

adiabatic interior and radiative exterior with negligible exterior pressure, which has the analytic solution $R_{\text{shell}} = 0.763 \left(\frac{L_{\text{wind}}}{a}\right)^{1/5} t^{3/5}$ (Weaver et al. 1977). We place a star with $\dot{M}_{\rm wind} = 10^{-5} M_{\odot}$ and $v_{\rm wind} = 3000 \; {\rm km \, s^{-1}}$ in a 16 pc box containing 2 \times 10⁴ M_{\odot}, with $\Delta m = 0.01 \,\mathrm{M}_{\odot}$ and $\Delta M_{\mathrm{w}} = 10^{-3} \,\mathrm{M}_{\odot}$, with initial temperature 10K. In Fig. 15, we plot the bubble expansion and find that agreement with the similarity solution is good whether we use pure local injection, pure cell spawning, or our hybrid method.¹³ We also examine the velocity anisotropy $\sigma_{\rm v,\,max}/\sigma_{\rm v,\,min}$ - 1 where $\sigma_{\rm v,\,max}$ and $\sigma_{\rm v,\,min}$ are the maximum and minimum gas velocity dispersions along the principal axes of the gas momentum distribution, which is 0 in the exact solution. This is typically $\lesssim 10$ per cent, except in the very early phase of the run with pure cell spawning (because the free expansion radius is not yet well resolved, so shot noise from individual injection steps is still apparent). For the hybrid method, the transition between methods occurs smoothly, with no clear spurious numerical artefacts.

4.4 Supernovae

4.4.1 Physics prescription

We assume that all stars more massive than $8\,M_\odot$ go SN at the end of their lifetime, with the lifetime given by equation (34) (from $\approx 40\,Myr$ for $8\,M_\odot$ to $\approx 3\,Myr$ at $100\,M_\odot$). When flagged as a SN, the star ceases all other forms of feedback, and rapidly expels its mass isotropically with velocity

$$v_{\rm SN} = \sqrt{\frac{2E_{\rm SN}}{M_{\rm ejecta}}} = 3200 \,\rm km \, s^{-1} \left(\frac{E_{\rm SN}}{10^{51} \rm erg}\right)^{1/2} \left(\frac{M_{\rm ejecta}}{10 M_{\odot}}\right)^{-1/2}, \tag{47}$$

where we assume $E_{SN}=10^{51} {\rm erg}$ by default. We assume the entire star is destroyed, but we can in principle allow for finite-mass relic compact objects by reserving a certain final mass. We assume that the SN ejecta have IMF-averaged yields according to Nomoto et al. (2006), with mass fractions (He, C, N, O, Ne, Mg, Si, S, Ca, Fe) = (3.87, 0.133, 0.0479 MAX[Z/Z_{\odot} , 1.65], 1.17, 0.30, 0.0987, 0.0933, 0.0397, 0.00458, 0.0741).

4.4.2 Numerical methods

SNe are realized numerically by the same cell spawning strategy as winds, except that 1) the spawned cells have the standard mass resolution $\Delta M_{\rm w} = \Delta m$ and 2) cells are spawned in shells of $N_{\rm spawn} = 24$ cells at once until the progenitor mass is exhausted. Mass is transferred from the star to the wind reservoir at a rate of

$$\dot{M}_{\rm SN} = \frac{N_{\rm spawn} v_{\rm SN} \Delta m}{R_{\rm sink}},\tag{48}$$

which in our default simulations is $\approx 1\,M_\odot\,yr^{-1}$. Hence, a typical progenitor will actually take several years to eject all its mass, but this does not affect the solution on scales $\gtrsim 0.01\,pc$, where we actually resolve the dynamics. We impose this finite duration because a very massive star could, in a single time-step, spawn $\sim\!10^5$ new cells, which would make operations like load-balancing and controlling $\nabla\cdot \textbf{\textit{B}}$ computationally challenging. Gas cells within a given shell are arranged in a regular angular grid pattern following Bruls, Vollmöller & Schüssler (1999) to avoid pathological cell arrangements and ensure statistical isotropy, and the orientation of the grid is randomized between each shell to reduce grid alignment effects.

4.4.3 Tests

In Fig. 16, we test this algorithm by detonating a $10\,\mathrm{M}_\odot$ progenitor star in the manner described here. The star is initially placed in a 16pc box containing $2\times10^4\,\mathrm{M}_\odot$ in gas (with $0.01\,\mathrm{M}_\odot$ mass resolution in the box and the ejecta), and we follow the evolution of the remnant from the free expansion through Sedov–Taylor through snowplow phases. The radius of the swept-up shell matches the expected similarity solutions in the different phases, and the terminal momentum boost factor from PdV work in the Sedov–Taylor phase is \sim 6, consistent with more-detailed SN simulation studies (Martizzi, Faucher-Giguère & Quataert 2015; Walch & Naab 2015; Haid et al. 2016; Gentry et al. 2017; Hopkins et al. 2018a), given the ambient density $n\sim200\,\mathrm{cm}^{-3}$ and metallicity ($Z=Z_\odot$). Numerical momentum anisotropy is always small, peaking at ~1 per cent

¹³We have also found negligible differences between the different wind methods in full star cluster formation simulations including winds as the only feedback mechanism.




Figure 16. Evolution of the SN remnant from a $10\,\mathrm{M}_\odot$ progenitor in a $16\,\mathrm{pc}$ box containing $2\times10^4\,\mathrm{M}_\odot$ ($n\sim200\,\mathrm{cm}^{-3}$, $Z=Z_\odot$), with cooling physics enabled (Section 3) and with the SN modelled via cell-spawning as described in Section 4.4, SNR evolution starting at the free-expansion phase. Top: Radius of the swept-up shell of dense gas as a function of time, which interpolates between the initial free-expansion ($\propto t$), Sedov–Taylor ($\propto t^{2/5}$), and pressure-driven snowplow phases ($\propto t^{2/7}$) as shown. Middle: radial momentum in units of the initial $p_{\rm ejecta} = M_{\rm ejecta} v_{\rm ejecta}$, initially conserved in the free-expansion phase but boosted by a factor of \sim 6 due to PdV work performed in the Sedov–Taylor phase. Bottom: Numerical velocity anisotropy, which is generally $\lesssim 1$ per cent.

during the Sedov-Taylor phase. We visualize the final morphology of the SN remnant in Fig. 17.

4.5 Radiation

Following Hopkins et al. (2020a), STARFORGE simulations with the radiative transfer module enabled follow the emission, transport, and absorption of photons in five different bands in wavelength λ :

- (i) Hydrogen ionizing ($\lambda < 912$ Å): Ionizing photons emitted by stars and responsible for the dynamics of H II regions, which are widely theorized to be the most important feedback effect from massive stars in typical Galactic conditions on global GMC scales (McKee, van Buren & Lazareff 1984; Krumholz & Matzner 2009; Dale, Ercolano & Bonnell 2012; Geen, Soler & Hennebelle 2017; Kim et al. 2018; Grudić et al. 2019; Olivier et al. 2021).
- (ii) Far-UV/photoelectric (912 Å < λ < 1550 Å): Responsible for heating the ISM via the photoelectric effect on dust grains, and likely an important component of the thermal balance of the cold and warm neutral media in the outer parts of galaxies (Wolfire et al. 1995; Ostriker, McKee & Leroy 2010).
- (iii) Near-UV (1550 Å $< \lambda < 3600$ Å): Contains most of the photon energy and momentum emitted by a young stellar population, and hence is the most important term for direct stellar radiation pressure (Fall et al. 2010; Murray et al. 2010; Raskutti et al. 2016; Kim et al. 2018; Hopkins & Grudić 2019).
- (iv) Optical/near-IR (3600 Å $< \lambda < 3 \mu m$): Contains most of the light from old stellar populations, and carries a non-negligible




Figure 17. Morphology of the SN remnant at the end of the SN test in Section 4.4, when the SN remnant has entered the pressure-driven snowplow phase.

fraction of photon momentum that can potentially couple on larger scales in a GMC due to reduced dust opacity compared to NUV.

(v) *Mid/far-IR* (mainly $\lambda > 3 \mu m$): Radiation absorbed and reradiated by dust, which is the primary cooling mechanism in the densest gas and can dominate the radiation pressure near massive protostars or in ULIRGs (Krumholz et al. 2009; Kuiper et al. 2011; Davis et al. 2014; Rosen et al. 2016; Tsang & Milosavljević 2018). This is treated specially, as a component of the radiation field having a blackbody SED with a local effective temperature $T_{\rm rad}$, which is evolved self-consistently.

We can optionally further 'fine grain' these bands into narrower bins: for example, splitting ionizing and FUV photons into photoelectric, Lyman-Werner, H ionizing, He-ionizing, He-secondary-ionizing, soft X-ray, etc., as described in Hopkins et al. (2020a), but this generally produces second-order effects. Note that, unlike most RHD SF simulations, we *independently* and explicitly evolve the dust temperature $T_{\rm dust}$, radiation temperature $T_{\rm rad}$ (of the IR band), ¹⁴ and gas temperature $T_{\rm gas}$.

All sinks/stars are treated as potential sources for all bands above. In our default simulations, we calculate the emitted flux in each band by treating each sink/star as a blackbody with effective temperature $T_{\rm eff,\star}\approx 5780~{\rm K}~(L_\star/{\rm L}_\odot)^{1/4}~(R_\star/{\rm R}_\odot)^{-1/2}$ with L_\star and R_\star given by our stellar evolution module (Section 2.5.5), integrated over the relevant wavelengths. We ignore 'primary' gas emission at other wavelengths, as this is generally negligible in the problems of interest. Secondary gas/dust (re)-emission is treated as follows: recombination emission from absorbed ionizing photons are re-emitted into the optical/near-IR (OIR) band; the absorbed radiation energy in other bands (where

¹⁴Note that the IR radiation temperature $T_{\rm rad,\,IR}$ parametrizes the spectral shape (or equivalently wavelength of the IR SED peak or mean energy per photon). We therefore allow the IR radiation energy density $u_{\gamma,\,\rm IR}$ and spectral shape or $T_{\rm rad,\,IR}$ to evolve independently, rather than imposing the blackbody assumption $u_{\rm gamma,\,IR} \propto T_{\rm rad,\,IR}^4$ (which is only valid in the infinite optical-depth, tight-coupling limit).

the opacity is dust-dominated) is re-emitted by dust in the IR band with the evolved dust temperature $T_{\rm dust}$.

The limitation of this band-integrated treatment of radiation is that line emission and absorption are not followed explicitly. As such, it does not capture molecular line cooling explicitly (which we treat using approximate formulae, see Section 3). It is also not applicable to dust-free conditions where multiply-scattered Ly α photons are dynamically important (e.g. Smith, Bromm & Loeb 2017).

Absorption and scattering cross-sections/opacities and coupling to our thermo-chemistry (gas heating/cooling) routines largely follow Hopkins et al. (2020a). For ionizing bands, this employs standard photoionizing absorption cross sections which scale with the neutral H and neutral or partially ionized fractions for He, and absorbed ionizing photons directly couple to our detailed photoionization heating rates (see Hopkins et al. 2018b; note these are always calculated, our RHD simulations simply include local sources in addition to the UVB and ISRF). The opacities in FUV, NUV, OIR are given by the grey expressions: $(\kappa_{FUV}, \kappa_{NUV}, \kappa_{OIR}) = (0.2 +$ $2000 f'_d$, $1800 f'_d$, $180 f'_d$) cm² g⁻¹ where $f'_d = f_d/0.01$ is the local dust-to-gas ratio relative to solar, defined as in Section 3. The FUV intensity directly enters the photo-electric heating rate in our thermochemistry calculation (Hopkins et al. 2018b; appendix B); absorbed radiation in FUV+NUV+OIR + IR bands contributes to determine the dust temperature T_{dust} which interacts via dust-gas collisions (Section 3). For the IR band, we calculate the opacity as a function of ionized fraction, local dust-to-gas ratio, dust temperature, and radiation temperature, as $\kappa_{\rm IR} = \kappa_{\rm gas} + \kappa_{\rm dust}^0 f_d'$ where $\kappa_{\rm gas} \approx 0.35 \, x_e \, {\rm cm}^2 \, {\rm g}^{-1}$ from Thompson scattering with x_e the free electron fraction, and $\kappa_{\rm dust}^0(T_{\rm dust}, T_{\rm rad})$ calculated from the tables of Semenov et al. (2003). Specifically, we take the 'standard' model with the 'porous 5-layered sphere' composition in Semenov et al. (2003) and for each dust temperature T_{dust} (which gives a different dust composition) we explicitly calculate the Rosseland-mean opacity for each radiation temperature $T_{\rm rad}$ assuming a blackbody-like spectral shape. We provide detailed fits in Appendix C. We assume a sublimation temperature of $T_{\text{dust}}^{\text{sub}} = 1500 \text{ K}$. Finally, because our RHD methods account for both absorption and scattering we must define the albedo A. For simplicity we assume $A_{ion} = 0$ (pure absorption) for ionizing bands; $A_{\text{FUV, NUV, OIR}} = 1/2$, i.e. equal absorption and scattering opacities which is roughly appropriate for dust grains in FUV through OIR bands (see e.g. Weingartner & Draine 2001); and for IR we assume the Thompson portion of the opacity is pure-scattering while for the dust albedo we can interpolate reasonably accurately between the short-wavelength (A \sim 1/2) and long-wavelength (Rayleigh scattering, $A \to 1$) regimes by taking $A_{\rm IR} \approx (\tilde{T}_r/2 + 1)/(\tilde{T}_r + 1)$ with $\tilde{T}_r \equiv (T_{\rm rad, IR}/725 \,\mathrm{K})^2$.

Photon momentum (radiation pressure) is always transferred appropriately to the gas + dust when radiation is absorbed or scattered (from any band).

4.5.1 Photon injection

Photons from sinks must be injected into the simulation domain before they are propagated by the RT solver. We do this via local injection: constructing effective oriented faces A_{sg} between the sink particle and overlapping gas cells, and injecting photons conservatively with a weighting given by the solid angle subtended by the face (e.g. Fig. 9). A full description of the original algorithm for photons is given in Hopkins & Grudić (2019) Appendix A, but we make a small extension here.

In the original algorithm, an extinction factor $f_{\rm abs} = \exp(-r_{sg}/\lambda_{\rm mfp})$ (for photon mean free path $\lambda_{\rm mfp} = (\kappa \rho)^{-1}$) was applied

to the injected photon energy and momentum of a cell, and the appropriate absorbed photon momentum was imparted. This models sub-resolution extinction, which is crucial for capturing radiation pressure effects when λ_{mfp} is unresolved – which is very often the case in SF problems at practical resolutions (Krumholz 2018; Hopkins & Grudić 2019). Here we also take the photon energy absorbed on unresolved scales and 'downgrade' (re-emit) it to the appropriate band as defined above. The downgraded photons are then injected into their respective bands in addition to the photons originally in that band in the stellar SED. Hence, in practice a star in a highly optically thick accretion flow will usually end up injecting most of its luminosity to the mid/far-IR band, because λ_{mfp} is not resolved.

4.5.2 Photon transport

GIZMO employs modular RHD solvers, so in principle we can adopt and compare various methods for photon transport. But in our default explicit-RHD simulations we adopt the first-moment or M1 (Levermore 1984) method, which has the advantages of being computationally efficient (well-adapted to hierarchical time-stepping and multiphysics simulations), manifestly momentum and energy conserving in finite-volume form, able interpolate between opticallythick and optically thin limits, and well-tested in simulations of star cluster formation (Geen et al. 2015; Gavagnin et al. 2017; Geen et al. 2017; He et al. 2019). In particular, for questions involving radiation pressure forces on gas, shadowing in an inhomogeneous medium, and the transition between optically thin-thick regimes, it is (by construction) able to capture phenomena which cannot appear in the 0th-order flux-limited-diffusion (FLD) method (see references above and e.g. Davis et al. 2014; Rosdahl et al. 2015; Zhang & Davis 2017; Kannan et al. 2019). For each band i we explicitly evolve the first two moments of the intensity equation in the usual mixedframe approximation keeping all terms to $O(v^2/c^2)$, with all terms appropriately integrated over the relevant bands (Mihalas & Mihalas 1984; Lowrie, Morel & Hittinger 1999). This gives

$$\frac{\partial e_r^i}{\partial t} + \nabla \cdot \boldsymbol{f}_r^i = \left(\dot{e}_{\text{em}}^i - \dot{e}_{\text{abs}}^i\right) + (\psi_a^i - \psi_s^i) \boldsymbol{u} \cdot \boldsymbol{g}_r^i, \tag{49}$$

$$\frac{1}{\tilde{c}^2} \frac{\partial f_r^i}{\partial t} + \nabla \cdot \mathbb{P}_r^i = -(\psi_a^i + \psi_s^i) g_r^i + \frac{u}{c^2} \left(\dot{e}_{\text{em}}^i - \dot{e}_{\text{abs}}^i \right), \tag{50}$$

where \boldsymbol{u} is the local gas/dust velocity, $\dot{e}_{abs} \equiv \tilde{c}^2 \psi_a \, e_r$ and \dot{e}_{em} are the volumetric absorption and emission rates, $\boldsymbol{g}_r \equiv \boldsymbol{f}_r - \boldsymbol{u} \cdot (e_r \, \mathbb{I} + \mathbb{P}_r)$ with e_r and \boldsymbol{f}_r the radiation energy and flux densities, $\psi_{a,s} \equiv \rho \, \kappa_{a,s}/\tilde{c}$ are the absorption+scattering coefficients, \tilde{c} and c the reduced (RSOL) and true speed of light, and $\mathbb{P}_r \equiv e_r \, \mathbb{D}_r$ the radiation pressure tensor (Eddington tensor \mathbb{D}_r). These are discretized and inter-cell fluxes are integrated in the same finite-volume (conservative) form as the MHD equations (Hopkins et al. 2020a).

Note that equations (49)–(50) include terms up to formal $O(v^2/c^2)$ such as the $\boldsymbol{u} \cdot (e_r \mathbb{I} + \mathbb{P}_r)$ term differentiating \boldsymbol{g}_r and \boldsymbol{f}_r or the $(\psi_a - \psi_s) \boldsymbol{u} \cdot g_r$ term representing the work done by radiation pressure which are often dropped in SF simulations where typical speeds $v \ll c$. However, as many authors have pointed out, these terms actually dominate the behaviour in the infinite-optical-depth, tight-coupling or photon-trapped limit, where their actual order scales closer as $O(v/v_{\rm rad, \, diff})$ (with $v_{\rm rad, \, diff}$ the effective bulk speed

¹⁵We adopt the common M1 closure for \mathbb{P}_r , taking $\mathbb{D}_r \to (1/2)[(1-\chi)\mathbb{I} + (3\chi - 1)\hat{f}_r\hat{f}_r]$, $\chi \equiv (3+4\xi^2)/(5+2[4-3\xi^2]^{1/2})$, and $\xi \equiv |f_r|/(\tilde{c}e_r)$, which interpolates between thin and thick regimes (Levermore 1984).

of photon diffusion). Without this terms, the RHD equations in the trapped limit will give unphysical solutions (photons will not properly be advected and arbitrarily large radiation-pressure forces can arise). The terms in $(\dot{e}_{abs}^i - \dot{e}_{em}^i) u/c^2$, on the other hand, represent true relativistic beaming, and are negligible for our problems of interest.

In the gas/dust momentum equation we add the terms $\sum_i (\psi_a^i + \psi_s^i) \, \mathbf{g}_r^i + (\dot{e}_{abs}^i - \dot{e}_{em}^i) \, \mathbf{u}/c^2$, the former representing the normal radiation pressure acceleration and the latter accounting for beaming. In the gas/dust total energy equation we add the terms $\sum_i (\dot{e}_{abs}^i - \dot{e}_{em}^i) + (\psi_s^i - \psi_a^i) \, \mathbf{u} \cdot \mathbf{g}_r^i$, the former representing the energy of absorption + emission (handled with our thermochemistry as described above) and the latter work terms. The gas temperature is then evolved according to the thermodynamics modules in Section 3, coupled also to the dust temperature by way of dust-gas collisions (equation 36). The dust temperature is set assuming grain absorption-emission equilibrium, giving $T_{\rm dust}^4 = (\langle Q_{abs} \rangle/\langle Q_{em} \rangle) \, e_r^{\rm tot} \, c/(4 \, \sigma_T)$ where $e_r^{\rm tot} = \sum_i e_r^i$, and $Q_{abs, \, em}$ are the appropriate absorption and emission efficiencies. ¹⁶

As noted above, the IR radiation field is treated as a blackbody shape with local effective temperature $T_{\rm rad,\,IR}$ and total energy integrated over the cell domain $E_{\rm gamma,\,IR}$, which evolves as new photons are emitted or when radiation is exchanged between cells of different $T_{\rm rad}$. In emission, sinks emit with $T_{\rm rad,\,IR}^{\rm em} = T_{\rm eff}$ and dust has $T_{\rm rad,\,IR}^{\rm em} = T_{\rm dust}$; given a total emitted radiation energy $\Delta E_{\rm em}$ in the cell timestep, the effective $T_{\rm rad,\,IR}$ is then updated to guarantee both radiation energy and photon number conservation, giving $T_{\rm rad,\,IR}(t+\Delta t) = [E_{\gamma,\,\rm IR}(t) + \Delta E_{\rm em}]/[E_{\gamma,\,\rm IR}(t)/T_{\rm rad,\,IR}(t) + \Delta E_{\rm em}/T_{\rm rad,\,IR}^{\rm em}]$. The same update scheme is used when cells exchange radiation energy.

Finally, we follow common practice and adopt a reduced speed of light (RSOL) with $\tilde{c} < c$ to enable larger time-steps. In general, \tilde{c} should be larger than the bulk speeds of radiative diffusion or ionizing front expansion to capture the dynamics; Geen et al. (2015) argue this is satisfied for $\tilde{c} \gtrsim 30~{\rm km\,s^{-1}}$ in our problems of interest, and we verify this below.

4.5.3 Tests

First, we remind the reader of the test in Section 3, which demonstrates the accuracy of our IR RHD + thermochemistry models evolving T_{dust} , T_{rad} , and T_{gas} appropriately in the collapse of a Jeansunstable core. To test other aspects of our RHD models, we next repeat the test setup of an O star in a box described in the previous sections (Sections 4.3–4.4) for radiation. First, we examine the ability of the photon injection, transport, and absorption schemes to capture radiation pressure in two regimes: where the photon mean free path is well-resolved ($\Delta x \ll \lambda_{\rm mfp}$) and totally unresolved ($\Delta x > \lambda_{\rm mfp}$). We inject radiation in a single band with opacity scaled so that the cell opacity $\tau_{\text{cell}} = \Delta x / \lambda_{\text{mfp}}$ ranges from 0.015 to 1.5, and the global optical depth $\tau_{\rm box} = L_{\rm box}/\lambda_{\rm mfp}$ ranges from 1.9 to 190. Because the box is always optically-thick and thermal pressure is negligible, in all cases we expect the expanding shell solution to approach the momentum-conserving similarity solution $R_{\text{shell}} = \left(\frac{L_{\star}}{6\pi\rho_{0}c}\right)^{1/4}t^{1/2}$ at late times with radial momentum approaching the total emitted photon momentum $L_{\star}t/c$, and the solution should be spherically symmetric.

Fig. 18 shows that the dynamics of a radiation pressure-driven shell are captured accurately: all three solutions eventually approach the




Figure 18. Radiation pressure-only test with an O star in a homogeneous box, accounting only for radiation pressure feedback with a range of opacities, varying the global and cell optical depths $\tau_{\rm box}$ and $\tau_{\rm cell}$ and injecting and transporting photons as described in Section 4.5. Top: position of the spherical shell swept up by the radiatively-driven bubble. All solutions approach the analytic similarity solution for momentum-driven bubbles $\propto t^{1/2}$, with the time lag determined by the absorption time-scale as expected. Middle: Radial gas momentum in units of the total emitted photon momentum Lt/c. The correct momentum is always coupled even when $\tau_{\rm cell} > 1$, by accounting for unresolved local extinction in the injection phase (Hopkins & Grudić 2019). Bottom: Numerical velocity anisotropy, which is < 1 per cent in all but the earliest (i.e. worst-resolved, few-cell) phases of the expansion.

similarity solution (with the least optically-thick run having a time delay $\lambda_{mfp}/\tilde{c}\approx 0.3$ Myr, owing to the finite light traveltime before absorption). The correct radial momentum is imparted whether or not λ_{mfp} is well-resolved (due to the face-integrated injection method, detailed in Hopkins & Grudić 2019), and numerical anisotropy falls rapidly below ~ 1 per cent once the bubble becomes well-resolved.

We repeat this experiment with our ionizing radiation band enabled, without radiation pressure, to test the ability of the code to follow the dynamics of expanding H II regions (and also to determine an appropriate value for \tilde{c} to accomplish this). We compare with the approximate Hosokawa & Inutsuka (2006) solution for the position of the ionization front:

$$R_{\rm HII} = R_{\rm St} \left(1 + \frac{7}{4} \sqrt{\frac{4}{3}} \frac{c_i t}{R_{\rm St}} \right)^{4/7},$$
 (51)

where $c_i \approx 11 \; \mathrm{km} \, \mathrm{s}^{-1}$ is the isothermal sound speed in ionized gas and the R_{Sr} is the Strömgren radius:

$$R_{\rm St} = \left(\frac{3\dot{N}_{\rm LyC}m_{\rm p}^2}{\alpha_{\rm B}\rho_0^2}\right)^{1/3},\tag{52}$$

where $\dot{N}_{\rm LyC}$ is the emission rate of H ionizing photons and $\alpha_{\rm B}$ is the case-B H II recombination coefficient (invoking the 'on-the-spot' approximation, Osterbrock 1989). We fix the ratio $\dot{N}_{\rm LyC}/\alpha_{\rm B}$ so that $R_{\rm St}$ is initially resolved in only two cell lengths, $R_{\rm St} = 2(\Delta m/\rho_0)^{1/3}$, and survey \tilde{c} values between 5 and 300 km s⁻¹. We also run a version with $\tilde{c} = 30$ km s⁻¹ with the gas fixed in place, to compare with the

 $^{^{16}}$ For IR–IR band interactions, we assume $Q_{\rm abs} \approx Q_{\rm em}$, though this could lead to small differences in $T_{\rm dust}$.




Figure 19. Evolution of an H II region surrounding a single O star in an initially static, homogeneous box, with only ionizing radiation feedback and no photon momentum. Top: Evolution of the ionization front, plotting the interval over which the ionization fraction falls from 99 per cent to 1 per cent, comparing results for a range of reduced speed of light values to the analytic solution (Hosokawa & Inutsuka 2006). The 'No Hydro' solution freezes the gas in place, simulating a static Strömgren sphere. Bottom: Numerical velocity anisotropy as a function of time (0 in the exact solution).

static Strömgren sphere solution and isolate artefacts of the RSOL approximation.

We plot the expansion of the ionization front for the different runs in Fig. 19. The 'No Hydro' frozen solution relaxes to the static Strömgren sphere solution after a time $\sim t_{\rm St} = R_{\rm St}/\tilde{c}$ as expected, and remains statistically (but not exactly) static thereafter. The solutions with $\tilde{c} > 30~{\rm km\,s^{-1}}$ relax to the Strömgren solution similarly, but then start to expand after a time $\sim R_{\rm St}/c_{\rm i}$ and agree well with the equation (51) solution. But for $\tilde{c} = 5~{\rm km\,s^{-1}}$, $t_{\rm St}$ is longer than the physical sound crossing time of the bubble, so the bubble expansion is delayed artificially, confirming the finding of Geen et al. (2015) that $\tilde{c} \sim 30~{\rm km\,s^{-1}}$ is roughly the marginal RSOL value for following the dynamics of H II regions accurately. Numerical anisotropy is again small (a few per cent) once the bubble actually expands and becomes well resolved.

5 DISCUSSION

We have presented, demonstrated, and tested the methods used for STARFORGE simulations, and we refer the reader to Paper 2 for the preliminary science results of the STARFORGE project. We now discuss some further applications of the methods presented here and enumerate several caveats, limitations, and possible extensions to our set-up.

5.1 Applications

The particular suite of physics and numerical methods developed here is optimized and intended for GMC and star cluster formation simulations, but the methods described in this work are potentially suitable for wider applications in astrophysical simulations involving stars and feedback.

5.1.1 Dedicated feedback simulations

The methods we have presented for coupling feedback from individual stars do not necessarily need to be combined with star formation simulations – in principle our feedback implementation is suitable for any problem involving stellar winds, jets, radiation, or SNe from individual stars. Notably, our methods can be used to capture multi-scale flows in complicated geometries, such as the evolution of a SN remnant from the free-expansion phase at sub-AU scales onward in an inhomogeneous ISM, or following interacting binary stellar winds from the scale of the binary separation all the way to interaction with the ISM. These geometries are historically challenging for AMR methods owing to high-velocity, non-gridaligned motion.

5.1.2 Local and global galaxy simulations

Stratified and/or shearing-box simulations have been used to simulate the evolution of a patch of the ISM within a galaxy at a resolution that is generally higher than what is attainable in global galaxy simulations (e.g. Hennebelle & Iffrig 2014; Walch et al. 2015; Martizzi et al. 2016; Kim & Ostriker 2017). These can be used to follow the formation and dispersal of GMCs self-consistently. All the algorithms presented here translate directly to this type of setup – the only differences are the initial conditions, boundary conditions, and additional inertial forces (all currently implemented in GIZMO). Resolved SF simulations could also be performed in a galactic context via a 'zoom-in' re-simulation of a GMC, taking a coarsely resolved GMC in a simulated galaxy (e.g. Guszejnov et al. 2020a) and up-sampling it to higher resolution (Rey-Raposo, Dobbs & Duarte-Cabral 2015).

The total stellar masses we form in the largest simulations in Paper 2 ($\gtrsim\!10^4\,M_\odot$) is within an order of magnitude of the total stellar mass in the faintest known dwarf galaxies (Simon 2019; Wheeler et al. 2019), so it may even be possible to perform a global cosmological galactic zoom-in simulation of an ultra-faint dwarf (UFD) with individually-resolved star formation. State-of-theart simulations like Wheeler et al. (2019) simulate UFD formation at a mass resolution of $\sim\!30\,M_\odot$, so much higher resolution would be required, but this could potentially be achieved by using a T97-like refinement criterion to reach the required resolution in the dense gas, while the more diffuse gas not engaged in star formation could be kept at coarser resolution.

5.1.3 Stellar zoom-in simulations

Our sink particle method creates an open boundary condition for gas to flow into a stellar system, but we do not follow physics on scales smaller than $R_{\rm sink}$, for lack of resolution. Meanwhile, detailed simulations of individual protostellar systems can capture processes on smaller scales that we cannot, but lack the broader context of star cluster formation that should inform the accretion history of the system, as well as environmental effects like ionizing radiation and close encounters (Concha-Ramírez, Vaher & Portegies Zwart 2019a; Concha-Ramírez et al. 2019b). Simulations like those here, however, can be used to inform the initial conditions for simulating the formation of individual star systems at much higher ($<10^{-6}\,{\rm M}_\odot$) resolution, sufficient to resolve the structure of the

2226 M. Y. Grudić et al.

inner envelope and disc, and follow the dynamics of dust and nonideal MHD with a suitable adaptive refinement scheme (e.g. Tomida, Okuzumi & Machida 2015; Mocz et al. 2017; Angles-Alcazar et al. 2020).

5.2 Caveats and room for improvement

5.2.1 Gravity and full N-body dynamics

In Section 2.3, we introduced a treatment of stellar dynamics using an integrator that gives superior accuracy to the usual secondorder integrators used in multiphysics simulations (Fig. 1), allowing us to preserve the properties of binaries over the GMC lifetime. However, the accuracy and efficiency of our algorithms in pure Nbody applications still pales in comparison to dedicated N-body codes (e.g. Aarseth 2003; Wang et al. 2015). Standard N-body treatments do not generally require a minimum softening length, using a variety of techniques to optimize binary motion and close encounters (e.g. regularization). The gravitational force is also generally exact to machine precision in pure N-body applications, or approximate to a specified very fine, dynamically controlled tolerance (McMillan & Aarseth 1993). The approximate tree-force is not necessarily an issue, because as discussed in Section 2.3 the error budget in SF simulations is dominated by errors and uncertainties in RMHD algorithms and feedback, but it is not presently clear what physics relevant to SF may be missed when softening is introduced (but we do not find qualitatively-different results with softenings as small as 1.8 AU, Section 4.2.3). A regularization scheme would improve the efficiency and accuracy of binary integration, potentially allowing larger timesteps and optimizing the simulations. However, such methods are non-trivial to couple to multi-physics simulations (see however recent successes with the AMUSE framework, Wall et al. 2020).

5.2.2 Radiative transfer

Although we have validated our implementation of the M1 radiative transfer method in various simple problems relevant to SF (Section 4.5), it is by no means clear that M1 can capture all important radiation phenomena in SF. As a moments method, it does not capture the collisionless nature of photons (colliding streams will shock, not pass through each other), so the radiation field streaming within a cluster will not generally be particularly accurate (but should still be reasonable if a single source is dominant).

Unfortunately the idiosyncrasies of different RT methods often only reveal themselves in complex, nonlinear problems with non-trivial geometries (such as SF), where exact solutions are unknown (as opposed to the simple problems considered here). This motivates an empirical approach to studying the behaviours of different RT methods in SF, i.e. comparing their results in the full SF problem, and determining which best reproduces observations. This will be important to do but is beyond the scope this work.

5.2.3 Resolving disc evolution

Our use of the sink particle method (which identifies each sink with an individual star) effectively ignores any possibility of fragmentation on scales $< R_{\rm sink}$ (~ 20 au in a typical STARFORGE application), and does not model the detailed accretion flow on to the protostar on scales smaller than the sink radius. As discussed in Section 2.5.2, the rate at which mass arrives at the protostar need not be the same as the rate at which it enters the sink boundary, and this difference

in accretion rate would ultimately influence the evolution of feedback rates from the protostar, and the surrounding environment in turn.

If fragmentation occurs frequently on these smaller scales as may happen in gravitationally unstable discs (Kratter et al. 2010; Kratter & Lodato 2016), then our predicted IMF will be incomplete for a given set of physical assumptions. Our resolution study (Section 4.2.3) does not show any hint of IMF incompleteness as \sim au scales start to become resolved, but our simulation assumes ideal MHD and hence may overestimate magnetic braking (Li, Krasnopolsky & Shang 2011). This likely exaggerates accretion on to the central star and reduces its disc mass, suppressing any possible fragmentation. The impact of disc properties and evolution on the IMF should be investigated further with higher resolution simulations accounting for non-ideal MHD and radiative transfer (e.g. Wurster et al. 2019).

5.2.4 Sub-grid accretion and feedback modelling

As previously discussed in Section 2.5.2, another caveat of not following sub-au physics is that the rate at which mass arrives at the protostar, and is launched in an outflow, must be assumed. We show that our results are insensitive to our assumed accretion rate at least the factor of 10 level in protostellar jet simulations (4.2.3), but if the flow is angular momentum-supported at unresolved scales then accretion may proceed slower still, regulated by the rate of angular momentum transport. If accretion proceeds much more slowly, then the rate of accretion-powered protostellar radiation and outflows will be reduced in turn.

Another potential issue is our assumptions about the power and collimation of protostellar outflows. We have assumed a simple parametrized model following Cunningham et al. (2011), with parameters chosen to roughly match observations, but in reality these parameters may exhibit systematic scalings according to e.g. stellar type and accretion rate, non-ideal MHD processes and the dust grain distribution (Pudritz & Ray 2019), and the magnetic field geometry (Gerrard, Federrath & Kuruwita 2019). Because protostellar outflows can have such powerful effects upon SF, efforts should be made to constrain sub-grid prescriptions.

5.2.5 Cooling and chemistry

Our treatment of cooling and chemistry assumes equilibrium abundances, i.e. we do not explicitly follow the formation and destruction of the various molecular species that can serve as coolants or useful observational tracers. In principle this could affect the dynamics of the simulations, if the molecular cooling rate was severely overor underestimated, but in practice the cold, dense initial conditions we simulate can safely be assumed to be fully molecular, and even if not the cooling rate is actually fairly insensitive to the specific species into which e.g. C and O are locked (Glover & Clark 2012). The presence of molecules is a consequence of gas collapse, not a prerequisite (Orr et al. 2018). Rather, the main utility of selfconsistent chemistry in simulations is to enable the simulations to predict molecular emission self-consistently, e.g. to help determine which physical processes or which regions are being probed by different lines. In future work will explore simulations adopting detailed molecular networks such as CHIMES (Richings, Schaye & Oppenheimer 2014a, b), which has been implemented in GIZMO (Richings et al. in preparation).

6 CONCLUSION

We have presented and demonstrated the methods of the STAR-FORGE, combining the physics of MHD, gas self-gravity, stellar dynamics, thermodynamics, and all major dynamically important (proto-) stellar feedback mechanisms into a detailed numerical model of star formation. We have shown that the respective techniques for each mechanism give satisfactory results in test problems with known solutions. We also discovered a remarkable degree of robustness in the sink particle prescription (Section 2.5.8), and found good agreement when comparing with results in similar problems from a code that implements the same physics with completely different numerical methods (Section 4.2.3).

We found stable numerical results for the IMF down to a completeness limit of $\sim 0.1 M_{\odot}$ at the modest (by SF simulation standards) mass resolution of $\approx \! 10^{-3} \, M_{\odot}$ (Section 4.2.3, Fig. 12), and in Paper 2 we scale this setup up to GMCs as massive as $2 \times 10^5 \, M_{\odot}$, mapping out exploring the effects of protostellar jets upon the IMF at this scale for the first time. In subsequent works, we will present the full results of radiation MHD simulations of those same massive GMC models, with all feedback modules described in this work acting in concert

We anticipate that STARFORGE will be a useful theoretical laboratory for disentangling the many physical mechanisms at work in GMCs. By starting with a realistic picture and switching different physics on and off in controlled experiments, it can help distil the essential elements of a working theory of star formation. It can also be used to calibrate sub-resolution prescriptions for effects such as stellar kinematics and protostellar feedback for use in lower-resolution star cluster and galaxy formation simulations, increasing the predictive power of such simulations in the densest gas and star clusters, where the details of prescriptions become important (e.g. Hopkins et al. 2013b; Grudić & Hopkins 2019; Li et al. 2020). It should also be useful as interpretive tool for observations, mapping out the effects of different physics upon the relations between observed gas tracer properties and star formation.

An important goal of this project is to reduce the dependence of SF simulations upon sub-grid prescriptions, which must make highly uncertain assumptions about how individual stars form. Our setup helps to accomplish this, but it only peels back one layer. To simulate feedback and the emergence of the IMF, we must make assumptions about how various sub-au physics (accretion, stellar winds, jet launching, protostellar and stellar evolution/death) proceed, and we list various ways in which incorrect assumptions about these processes could affect our results (Section 5.2). Hence, it is crucial to continue to advance our understanding of the processes governing the formation and internal evolution of individual stars and star systems.

ACKNOWLEDGEMENTS

We warmly thank the theoretical and observational star formation communities for the innumerable enlightening exchanges that informed and motivated this work over the past several years. We are especially grateful to fellow starsmith Anna Rosen and to the referee Chris Matzner, whose careful readings helped improve the manuscript. MYG is supported by a CIERA Postdoctoral Fellowship. DG is supported by the Harlan J. Smith McDonald Observatory Postdoctoral Fellowship. Support for PFH was provided by NSF Collaborative Research Grants 1715847 and 1911233, NSF CAREER grant 1455342, and NASA grants 80NSSC18K0562 an JPL 1589742. SSRO acknowledges support by NSF CAREER

Award AST-1748571, NSF grant AST-1812747, NASA ATP grant 80NSSC20K0507 and a Cottrell Scholar Award from the Research Corporation for Science Advancement. CAFG is supported by NSF through grants AST-1715216, and CAREER award AST-1652522; by NASA through grant 17-ATP17-0067; and by a Cottrell Scholar Award from the Research Corporation for Science Advancement. This work used computational resources provided by XSEDE allocation AST-190018, the Frontera allocation AST-20019, and additional resources provided by the University of Texas at Austin and the Texas Advanced Computing Center (TACC; http://www.tacc.utexas.edu).

DATA AVAILABILITY

The data supporting the plots within this article and the initial conditions used for the numerical tests are available by request to the corresponding authors. A public version of the GIZMO code is available at http://www.tapir.caltech.edu/~phopkins/Site/GIZMO.html.

REFERENCES

Aarseth S. J., 2003, Gravitational N-Body Simulations. Cambridge University Press, Cambridge UK

Angles-Alcazar D. et al., 2020, preprint (arXiv:2008.12303)

Arce H. G., Sargent A. I., 2006, ApJ, 646, 1070

Bally J., 2016, Ann. Rev. Astron. Astrophys., 54, 491

Barnes J., Hut P., 1986, Nature, 324, 446

Bate M. R., 2009a, MNRAS, 392, 590

Bate M. R., 2009b, MNRAS, 392, 1363

Bate M. R., Burkert A., 1997, MNRAS, 288, 1060

Bate M. R., Bonnell I. A., Price N. M., 1995, MNRAS, 277, 362

Bate M. R., Bonnell I. A., Bromm V., 2003, MNRAS, 339, 577

Bleuler A., Teyssier R., 2014, MNRAS, 445, 4015

Bonnell I. A., Bate M. R., 2005, MNRAS, 362, 915

Bonnell I. A., Bate M. R., Vine S. G., 2003, MNRAS, 343, 413

Boss A. P., Bodenheimer P., 1979, ApJ, 234, 289

Bruls J. H. M. J., Vollmöller P., Schüssler M., 1999, A&A, 348, 233

Burkert A., Bodenheimer P., 1993, MNRAS, 264, 798

Chandrasekhar S., 1951, R. Soc. London Proc. Ser. A, 210, 26

Clark P. C., Glover S. C. O., Klessen R. S., 2012, MNRAS, 420, 745

Colín P., Vázquez-Semadeni E., Gómez G. C., 2013, MNRAS, 435, 1701

Colman T., Teyssier R., 2020, MNRAS, 492, 4727

Concha-Ramírez F., Vaher E., Portegies Zwart S., 2019a, MNRAS, 482, 732
 Concha-Ramírez F., Wilhelm M. J. C., Portegies Zwart S., Haworth T. J.,
 2019b MNRAS, 490, 5678

Cunningham A. J., Klein R. I., Krumholz M. R., McKee C. F., 2011, ApJ, 740, 107

Cunningham A. J., Krumholz M. R., McKee C. F., Klein R. I., 2018, MNRAS, 476, 771

Dale J. E., 2015, New Astron. Rev., 68, 1

Dale J. E., Bonnell I. A., 2008, MNRAS, 391, 2

Dale J. E., Ercolano B., Bonnell I. A., 2012, MNRAS, 424, 377

Davis S. W., Jiang Y.-F., Stone J. M., Murray N., 2014, ApJ, 796, 107

Dedner A., Kemm F., Kröner D., Munz C. D., Schnitzer T., Wesenberg M., 2002, J. Comput. Phys., 175, 645

Deng H., Mayer L., Meru F., 2017, ApJ, 847, 43

Deng H., Mayer L., Latter H., Hopkins P. F., Bai X.-N., 2019a, ApJS, 241, 26

Deng H., Reinhardt C., Benitez F., Mayer L., Stadel J., Barr A. C., 2019b, ApJ, 870, 127

Deng H., Ogilvie G. I., Mayer L., 2021, MNRAS, 500, 4248

Draine B. T., 1978, ApJS, 36, 595

Draine B. T., 2006, ApJ, 636, 1114

Draine B. T., Salpeter E. E., 1979, ApJ, 231, 77

Duffell P. C., MacFadyen A. I., 2011, ApJS, 197, 15

- Duffell P. C., D'Orazio D., Derdzinski A., Haiman Z., MacFadyen A., Rosen A. L., Zrake J., 2020, ApJ, 901, 25
- Evans Neal J. I., Heiderman A., Vutisalchavakul N., 2014, ApJ, 782, 114
- Evrard A. E., 1988, MNRAS, 235, 911
- Fall S. M., Krumholz M. R., Matzner C. D., 2010, ApJ, 710, L142
- Faucher-Giguère C.-A., 2020, MNRAS, 493, 1614
- Faucher-Giguère C.-A., Lidz A., Zaldarriaga M., Hernquist L., 2009, ApJ, 703, 1416
- Federrath C., 2015, MNRAS, 450, 4035
- Federrath C., Klessen R. S., 2012, ApJ, 761, 156
- Federrath C., Banerjee R., Clark P. C., Klessen R. S., 2010, ApJ, 713, 269
- Federrath C., Schrön M., Banerjee R., Klessen R. S., 2014, ApJ, 790, 128
- Federrath C., Krumholz M., Hopkins P. F., 2017, in Journal of Physics Conference Series. IOP Publishing, Bristol, UK, p. 012007
- Few C. G., Dobbs C., Pettitt A., Konstandin L., 2016, MNRAS, 460, 4382
- Gavagnin E., Bleuler A., Rosdahl J., Teyssier R., 2017, MNRAS, 472, 4155
- Geen S., Hennebelle P., Tremblin P., Rosdahl J., 2015, MNRAS, 454, 4484
- Geen S., Soler J. D., Hennebelle P., 2017, MNRAS, 471, 4844
- Gentry E. S., Krumholz M. R., Dekel A., Madau P., 2017, MNRAS, 465, 2471
- Gerrard I. A., Federrath C., Kuruwita R., 2019, MNRAS, 485, 5532
- Girichidis P. et al., 2020, Space Sci. Rev., 216, 68
- Glover S. C. O., Clark P. C., 2012, MNRAS, 421, 9
- Goldsmith P. F., Langer W. D., 1978, ApJ, 222, 881
- Gong H., Ostriker E. C., 2013, ApJS, 204, 8
- Grudić M. Y., 2020, preprint (arXiv:2010.13792)
- Grudić M. Y., Hopkins P. F., 2019, MNRAS, 488, 2970
- Grudić M. Y., Hopkins P. F., 2020, MNRAS, 495, 4306
- Grudić M. Y., Hopkins P. F., Faucher-Giguère C.-A., Quataert E., Murray N., Kereš D., 2018a, MNRAS, 475, 3511
- Grudić M. Y., Guszejnov D., Hopkins P. F., Lamberts A., Boylan-Kolchin M., Murray N., Schmitz D., 2018b, MNRAS, 481, 688
- Grudić M. Y., Hopkins P. F., Lee E. J., Murray N., Faucher-Giguère C.-A., Johnson L. C., 2019, MNRAS, 488, 1501
- Grudić M. Y., Kruijssen J. M. D., Faucher-Giguère C.-A., Hopkins P. F., Ma X., Quataert E., Boylan-Kolchin M., 2020, preprint (arXiv:2008.04453)
- Guszejnov D., Krumholz M. R., Hopkins P. F., 2016, MNRAS, 458, 673
- Guszejnov D., Hopkins P. F., Grudić M. Y., Krumholz M. R., Federrath C., 2018, MNRAS, 480, 182
- Guszejnov D., Grudić M. Y., Offner S. S. R., Boylan-Kolchin M., Faucher-Gigère C.-A., Wetzel A., Benincasa S. M., Loebman S., 2020a, MNRAS, 492, 488
- Guszejnov D., Grudić M. Y., Hopkins P. F., Offner S. S. R., Faucher-Giguère C.-A., 2020b, MNRAS, 496, 5072 (Paper 0)
- Guszejnov D., Grudić M. Y., Hopkins P. F., Offner S. S. R., Faucher-Giguère C.-A., 2021, MNRAS, 502, 3646 (Paper 2)
- Habing H. J., 1968, Bull. Astron. Inst. Netherlands, 19, 421
- Haid S., Walch S., Naab T., Seifried D., Mackey J., Gatto A., 2016, MNRAS, 460, 2962
- Haugbølle T., Padoan P., Nordlund Å., 2018, ApJ, 854, 35
- Heggie D. C., 1975, MNRAS, 173, 729
- He C.-C., Ricotti M., Geen S., 2019, MNRAS, 489, 1880
- Hennebelle P., Chabrier G., 2008, ApJ, 684, 395
- Hennebelle P., Chabrier G., 2013, ApJ, 770, 150
- Hennebelle P., Iffrig O., 2014, A&A, 570, A81
- Hopkins P. F. et al., 2018a, MNRAS, 477, 1578
- Hopkins P. F. et al., 2018b, MNRAS, 480, 800
- Hopkins P. F. et al., 2020b, MNRAS, 492, 3465
- Hopkins P. F., 2012, MNRAS, 423, 2016
- Hopkins P. F., 2014, ApJ, 797, 59
- Hopkins P. F., 2015, MNRAS, 450, 53 (H15)
- Hopkins P. F., 2016, MNRAS, 462, 576
- Hopkins P. F., 2017, MNRAS, 466, 3387
- Hopkins P. F., Grudić M. Y., 2019, MNRAS, 483, 4187
- Hopkins P. F., Lee H., 2016, MNRAS, 456, 4174
- Hopkins P. F., Raives M. J., 2016, MNRAS, 455, 51 (HR16)
- Hopkins P. F., Quataert E., Murray N., 2012, MNRAS, 421, 3488

- Hopkins P. F., Narayanan D., Murray N., 2013a, MNRAS, 432, 2647
- Hopkins P. F., Narayanan D., Murray N., Quataert E., 2013b, MNRAS, 433, 69
- Hopkins P. F., Grudić M. Y., Wetzel A., Kereš D., Faucher-Giguère C.-A., Ma X., Murray N., Butcher N., 2020a, MNRAS, 491, 3702
- Hosokawa T., Inutsuka S.-I., 2006, ApJ, 648, L131
- Howard C. S., Pudritz R. E., Harris W. E., 2017, MNRAS, 470, 3346
- Hubber D. A., Goodwin S. P., Whitworth A. P., 2006, A&A, 450, 881
- Hubber D. A., Walch S., Whitworth A. P., 2013, MNRAS, 430, 3261
- Hubber D. A., Rosotti G. P., Booth R. A., 2018, MNRAS, 473, 1603
- Hu C.-Y., Chiang C.-T., 2020, ApJ, 900, 29
- Jones M. O., Bate M. R., 2018, MNRAS, 480, 2562
- Kannan R., Vogelsberger M., Marinacci F., McKinnon R., Pakmor R., Springel V., 2019, MNRAS, 485, 117
- Kim C.-G., Ostriker E. C., 2017, ApJ, 846, 133
- Kim J.-G., Kim W.-T., Ostriker E. C., Skinner M. A., 2017, ApJ, 851, 93
- Kim J.-G., Kim W.-T., Ostriker E. C., 2018, ApJ, 859, 68
- Klessen R. S., Burkert A., 2000, ApJS, 128, 287
- Kratter K., Lodato G., 2016, ARA&A, 54, 271
- Kratter K. M., Matzner C. D., Krumholz M. R., Klein R. I., 2010, ApJ, 708, 1585
- Krause M. G. H. et al., 2020, Space Sci. Rev., 216, 64
- Kroupa P., 2002, Science, 295, 82
- Krumholz M. R., 2014, Phys. Rep., 539, 49
- Krumholz M. R., 2018, MNRAS, 480, 3468
- Krumholz M. R., Federrath C., 2019, Front. Astron. Space Sci., 6, 7
- Krumholz M. R., Matzner C. D., 2009, ApJ, 703, 1352
- Krumholz M. R., McKee C. F., Klein R. I., 2004, ApJ, 611, 399
- Krumholz M. R., McKee C. F., Klein R. I., 2005, Nature, 438, 332
- Krumholz M. R., McKee C. F., Kielii K. I., 2003, Nature, 438, 332 Krumholz M. R., Stone J. M., Gardiner T. A., 2007, ApJ, 671, 518
- Krumholz M. R., Klein R. I., McKee C. F., Offner S. S. R., Cunningham A. J., 2009, Science, 323, 754
- Krumholz M. R., Klein R. I., McKee C. F., 2012, ApJ, 754, 71
- Krumholz M. R., McKee C. F., Bland -Hawthorn J., 2019, ARA&A, 57, 227
- Kuiper R., Klahr H., Beuther H., Henning T., 2011, ApJ, 732, 20
- Lahén N., Naab T., Johansson P. H., Elmegreen B., Hu C.-Y., Walch S., 2019, ApJ, 879, L18
- Lahén N., Naab T., Johansson P. H., Elmegreen B., Hu C.-Y., Walch S., 2020, ApJ, 904, 71
- Lamers H. J. G. L. M., Snow T. P., Lindholm D. M., 1995, ApJ, 455, 269
- Larson R. B., 1969, MNRAS, 145, 271
- Larson R. B., 1981, MNRAS, 194, 809
- Larson R. B., 2005, MNRAS, 359, 211
- Lee H., Hopkins P. F., Squire J., 2017, MNRAS, 469, 3532
- Lee A. T., Offner S. S. R., Kratter K. M., Smullen R. A., Li P. S., 2019, ApJ, 887, 232
- Lee Y.-N., Hennebelle P., 2018, A&A, 611, A88
- Levermore C. D., 1984, J. Quant. Spec. Radiat. Transf., 31, 149
- Li P. S., Martin D. F., Klein R. I., McKee C. F., 2012, ApJ, 745, 139
- Li P. S., Klein R. I., McKee C. F., 2018, MNRAS, 473, 4220
- Li H., Vogelsberger M., Marinacci F., Gnedin O. Y., 2019, MNRAS, 487, 364
 Li H., Vogelsberger M., Marinacci F., Sales L. V., Torrey P., 2020, MNRAS, 499, 5862
- Li Z.-Y., Krasnopolsky R., Shang H., 2011, ApJ, 738, 180
- Longmore S. N. et al., 2012, ApJ, 746, 117
- Lowrie R. B., Morel J. E., Hittinger J. A., 1999, ApJ, 521, 432
- Lupi A., Volonteri M., Silk J., 2017, MNRAS, 470, 1673
- Mac Low M.-M., Klessen R. S., 2004, Rev. Mod. Phys., 76, 125
- Makino J., Aarseth S. J., 1992, PASJ, 44, 141
- Martizzi D., Faucher-Giguère C.-A., Quataert E., 2015, MNRAS, 450, 504 Martizzi D., Fielding D., Faucher-Giguère C.-A., Quataert E., 2016, MNRAS,
- 459, 2311
- Masunaga H., Miyama S. M., Inutsuka S.-I., 1998, ApJ, 495, 346
- Matzner C. D., McKee C. F., 1999a, in Nakamoto T., ed., Star Formation 1999. Nobeyama Radio Observatory, Nobeyama, Japan, p. 353
- Matzner C. D., McKee C. F., 1999b, ApJ, 526, L109
- Matzner C. D., McKee C. F., 2000, ApJ, 545, 364
- McKee C. F., Ostriker E. C., 2007, ARA&A, 45, 565

- McKee C. F., Williams J. P., 1997, ApJ, 476, 144
- McKee C. F., van Buren D., Lazareff B., 1984, ApJ, 278, L115
- McMillan S. L. W., Aarseth S. J., 1993, ApJ, 414, 200
- Meijerink R., Spaans M., 2005, A&A, 436, 397
- Meynet G., Maeder A., 2005, A&A, 429, 581
- Mihalas D., Mihalas B. W., 1984, Foundations of Radiation Hydrodynamics. Oxford University Press, New York, p. 731
- Mocz P., Burkhart B., Hernquist L., McKee C. F., Springel V., 2017, ApJ, 838 40
- Moseley E. R., Squire J., Hopkins P. F., 2019, MNRAS, 489, 325
- Mouschovias T. C., Spitzer L. J., 1976, ApJ, 210, 326
- Muñoz D. J., Miranda R., Lai D., 2019, ApJ, 871, 84
- Murray N., Rahman M., 2010, ApJ, 709, 424
- Murray N., Quataert E., Thompson T. A., 2010, ApJ, 709, 191
- Murray D., Goyal S., Chang P., 2018, MNRAS, 475, 1023
- Myers A. T., McKee C. F., Cunningham A. J., Klein R. I., Krumholz M. R., 2013, ApJ, 766, 97
- Myers A. T., Klein R. I., Krumholz M. R., McKee C. F., 2014, MNRAS, 439, 3420
- Nakamura F., Li Z.-Y., 2007, ApJ, 662, 395
- Nakano T., Hasegawa T., Morino J.-I., Yamashita T., 2000, ApJ, 534, 976
- Nomoto K., Tominaga N., Umeda H., Kobayashi C., Maeda K., 2006, Nucl. Phys. A, 777, 424
- Offner S. S. R., Arce H. G., 2014, ApJ, 784, 61
- Offner S. S. R., Chaban J., 2017, ApJ, 847, 104
- Offner S. S. R., Liu Y., 2018, Nature Astron., 2, 896
- Offner S. S. R., Klein R. I., McKee C. F., Krumholz M. R., 2009, ApJ, 703, 131
- Offner S. S. R., Lee E. J., Goodman A. A., Arce H., 2011, ApJ, 743, 91
- Oka T., Hasegawa T., Sato F., Tsuboi M., Miyazaki A., Sugimoto M., 2001, ApJ, 562, 348
- Olivier G. M., Lopez L. A., Rosen A. L., Nayak O., Reiter M., Krumholz M. R., Bolatto A. D., 2021, ApJ, 908, 68
- Orr M. E. et al., 2018, MNRAS, 478, 3653
- Osterbrock D. E., 1989, Astrophysics of Gaseous Nebulae and Active Galactic Nuclei. University Science Books, Melville, NY, USA
- Ostriker E. C., McKee C. F., Leroy A. K., 2010, ApJ, 721, 975
- Padoan P., Nordlund A., Jones B. J. T., 1997, MNRAS, 288, 145
- Padoan P., Nordlund Å., 2002, ApJ, 576, 870
- Padoan P., Nordlund A., 2011, ApJ, 741, L22
- Padoan P., Haugbølle T., Nordlund Å., 2012, ApJ, 759, L27
- Padoan P., Haugbølle T., Nordlund Å., Frimann S., 2017, ApJ, 840, 48
- Padoan P., Pan L., Juvela M., Haugbølle T., Nordlund Å., 2020, ApJ, 900, 82
- Panuelos J., Wadsley J., Kevlahan N., 2020, J. Comput. Phys., 414, 109454
 Pontzen A., Rey M. P., Cadiou C., Agertz O., Teyssier R., Read J., Orkney M. D. A., 2021, MNRAS, 501, 1755
- Portegies Zwart S. F., Makino J., McMillan S. L. W., Hut P., 1999, A&A, 348, 117
- Powell K. G., Roe P. L., Linde T. J., Gombosi T. I., De Zeeuw D. L., 1999, J. Comput. Phys., 154, 284
- Power C., Navarro J. F., Jenkins A., Frenk C. S., White S. D. M., Springel V., Stadel J., Quinn T., 2003, MNRAS, 338, 14
- Price D. J., Bate M. R., 2007, MNRAS, 377, 77
- Price D. J., Monaghan J. J., 2007, MNRAS, 374, 1347
- Pudritz R. E., Norman C. A., 1983, ApJ, 274, 677
- Pudritz R. E., Ray T. P., 2019, Front. Astron. Space Sci., 6, 54
- Rafikov R. R., 2007, ApJ, 662, 642
- Raskutti S., Ostriker E. C., Skinner M. A., 2016, ApJ, 829, 130
- Rees M. J., 1976, MNRAS, 176, 483
- Rennehan D., Babul A., Hopkins P. F., Davé R., Moa B., 2019, MNRAS, 483, 3810
- Rey-Raposo R., Dobbs C., Duarte-Cabral A., 2015, MNRAS, 446, L46
- Richings A. J., Schaye J., Oppenheimer B. D., 2014a, MNRAS, 440, 3349
- Richings A. J., Schaye J., Oppenheimer B. D., 2014b, MNRAS, 442, 2780
- Robertson B. E., Kravtsov A. V., Gnedin N. Y., Abel T., Rudd D. H., 2010, MNRAS, 401, 2463
- Rogers H., Pittard J. M., 2013, MNRAS, 431, 1337

- Rohde P. F., Walch S., Seifried D., Whitworth A. P., Clarke S. D., Hubber D. A., 2019, MNRAS, 483, 2563
- Rosdahl J., Schaye J., Teyssier R., Agertz O., 2015, MNRAS, 451, 34
- Rosen A. L., Krumholz M. R., 2020, AJ, 160, 78
- Rosen A. L., Krumholz M. R., McKee C. F., Klein R. I., 2016, MNRAS, 463, 2553
- Saitoh T. R., Makino J., 2009, ApJ, 697, L99
- Semenov D., Henning T., Helling C., Ilgner M., Sedlmayr E., 2003, A&A, 410, 611
- Shakura N. I., Sunyaev R. A., 1973, A&A, 500, 33
- Shi Y., Grudić M. Y., Hopkins P. F., 2020, preprint (arXiv:2008.12290)
- Shu F. H., 1977, ApJ, 214, 488
- Shu F., Najita J., Ostriker E., Wilkin F., Ruden S., Lizano S., 1994, ApJ, 429, 781
- Simon J. D., 2019, ARA&A, 57, 375
- Smith N., 2014, ARA&A, 52, 487
- Smith A., Bromm V., Loeb A., 2017, MNRAS, 464, 2963
- Sormani M. C., Treß R. G., Klessen R. S., Glover S. C. O., 2017, MNRAS, 466, 407
- Springel V., 2005, MNRAS, 364, 1105 (S05)
- Springel V., 2010, MNRAS, 401, 791
- Steigman G., Strittmatter P. A., Williams R. E., 1975, ApJ, 198, 575
- Su K.-Y. et al., 2018, MNRAS, 480, 1666
- Su K.-Y., Hopkins P. F., Hayward C. C., Faucher-Giguère C.-A., Kereš D., Ma X., Robles V. H., 2017, MNRAS, 471, 144
- Tomida K., Okuzumi S., Machida M. N., 2015, ApJ, 801, 117
- Torrey P. et al., 2020, MNRAS, 497, 5292
- Tout C. A., Pols O. R., Eggleton P. P., Han Z., 1996, MNRAS, 281, 257
- Truelove J. K., Klein R. I., McKee C. F., Holliman II J. H., Howell L. H., Greenough J. A., 1997, ApJ, 489, L179 (T97)
- Tsang B. T. H., Milosavljević M., 2018, MNRAS, 478, 4142
- Vaidya B., Mignone A., Bodo G., Massaglia S., 2015, A&A, 580, A110
- Vaytet N., Haugbølle T., 2017, A&A, 598, A116
- Vink J. S., de Koter A., Lamers H. J. G. L. M., 2001, A&A, 369, 574
- Vogelsberger M., White S. D. M., Helmi A., Springel V., 2008, MNRAS, 385, 236
- von Hoerner S., 1960, ZAp, 50, 184
- Walch S. et al., 2015, MNRAS, 454, 238
- Walch S., Naab T., 2015, MNRAS, 451, 2757
- Wall J. E., Mac Low M.-M., McMillan S. L. W., Klessen R. S., Portegies Zwart S., Pellegrino A., 2020, ApJ, 904, 192
- Wang P., Li Z.-Y., Abel T., Nakamura F., 2010, ApJ, 709, 27
- Wang L., Spurzem R., Aarseth S., Nitadori K., Berczik P., Kouwenhoven M. B. N., Naab T., 2015, MNRAS, 450, 4070
- Weaver R., McCray R., Castor J., Shapiro P., Moore R., 1977, ApJ, 218, 377 Weingartner J. C., Draine B. T., 2001, ApJ, 548, 296
- Wheeler C. et al., 2019, MNRAS, 490, 4447
- Whitworth A. P., Boffin H. M. J., Francis N., 1998, MNRAS, 299, 554
- Wolfire M. G., Hollenbach D., McKee C. F., Tielens A. G. G. M., Bakes E. L. O., 1995, ApJ, 443, 152
- Wurster J., Bate M. R., Price D. J., 2019, MNRAS, 489, 1719
- Zel'dovich Y. B., 1970, A&A, 5, 84
- Zhang D., Davis S. W., 2017, ApJ, 839, 54
- Zhu Q., Li Y., 2016, ApJ, 831, 52

APPENDIX A: LIST OF SINK PARTICLE TESTS

- In Section 2.5.8 we re-run the **M2e4** GMC simulation from Paper 0 with a large space of parameters and prescriptions for our sink particle algorithm, with results shown in Fig. 7. These tests include the following:
- (i) Fiducial parameters and prescriptions, simply re-running the simulation in Paper 0 with exactly the same code version as the other tests, with the methods described in Section 2.
- (ii) Simultaneously increasing $R_{\rm sink}$ and S_{\star} by $\times 10$ and $\times 100$ (from 18 au to 180–1800 au). These tests generally formed the

upper envelope of the predicted IMF mass scale statistics because accretion is made much easier and IMF incompleteness is introduced by accreting independently-collapsing cores within the sink radius. However, these effects are small.

- (iii) Increasing $R_{\rm sink}$ and $\times 10$ and $\times 100$ without rescaling S_{\star} . These had similar results to tests in which we scaled S_{\star} as well.
 - (iv) Decreasing S_{\star} from 18 to 1.8 au. This had negligible effects.
- (v) Changing the sink formation density threshold ρ_{th} by a factor of 10^{-3} and 10^3 . The $\rho_{th} \times 10^{-3}$ test also neglected the thermal term in the virial criterion (equation 21), which would otherwise have effectively imposed a density threshold of its own. This has the effect of rescaling the threshold number of Jeans wavelengths per cell at the sink formation threshold from 0.5 to 1.58 and 0.15, respectively [i.e. strongly violating versus satisfying the Truelove et al. (1997) criterion]. All results of the $\times 10^3$ run were difficult to distinguish from the fiducial run, possibly because following collapse far beyond ρ_J is unlikely to reveal any new fragments (Section 2.4). The $\rho_{th} \times 10^{-3}$ run formed a slightly larger number of sinks than the fiducial run, because the other sink formation criteria are not perfect predictors of runaway collapse when looking at gas of modest density, but effects were still quite small.
- (vi) Rescaling $R_{\rm sink}$ from 18 to 9 au and 4.5 au, respectively, while keeping the softening radius fixed at 18 au. The $R_{\rm sink}=9$ au run had no clear systematic difference from the fiducial run. The 4.5 au run (highlighted in Fig. 7) was the largest outlier of our survey, forming stars with a slight delay with respect to the modal SFE, and producing a noticeably larger (×2) number of sinks, affecting the median and mean sink masses in turn. Mass-weighted statistics such as M_{50} and $M_{\rm max}$ were affected more modestly, but were also systematically lower than the modal solution. Note that this is not a particularly reasonable prescription: it forces sink particles to have a volume 1/64 the volume of a gas cell at the resolution limit, resulting in a gross mismatch between $R_{\rm sink}$, S_{\star} , and the gas resolution scale at the time of sink formation, making it very difficult to satisfy all accretion criteria.
- (vii) A minimal accretion prescription requiring only $r < R_{\rm sink}$. This had negligible effects.
- (viii) A simple accretion prescription requiring boundedness (equation 23) and $r < R_{\text{sink}}$. This had negligible effects.
- (ix) Rescaling $R_{\rm sink}$ by a factor of 1/4 while *also* rescaling S_{\star} by the same factor and $\rho_{\rm th}$ by a factor of 64 so that the cell spacing at sink formation matched $R_{\rm sink}$. This was much closer to the reference run than reducing $R_{\rm sink}$ alone.
- (x) Running our fiducial sink formation prescription in conjunction with the simpler sink accretion prescription given in Bate et al. (1995) (except neglecting the correction terms to the hydro force). This amounts to neglecting thermal and magnetic pressure in the boundedness calculation (equation 23), and not requiring that gas cells physically fit inside the sink volume. This had negligible effects because there is considerable redundancy between the different checks.
- (xi) Ignoring the $\nabla \cdot \boldsymbol{v}$, virial, density maximum, tidal, and infall sink formation criteria in turn. These all had negligible effects except for neglecting the density maximum and virial criteria, which were at the upper envelope of number of sinks formed.
- (xii) Including a version of the Hubber et al. (2013) angular momentum return prescription, exerting a net torque $\tau = \frac{J_s}{t_{acc}}$ upon the surrounding gas to transfer angular momentum from the sink back to the gas (taking t_{acc} to be 500 yr, likely much faster than the actual angular momentum transfer time-scale in a protoplanetary disc). This had negligible effects, but may be more pronounced in problems

where protostellar discs are well-resolved and angular momentum support is important.

(xiii) Enforcing the additional criterion of 'collapse along all three axes', a stricter version of the $\nabla \cdot \boldsymbol{v}$ criterion. We check that all three eigenvalues of the symmetric component of $\nabla \boldsymbol{v}$ are negative (as opposed to merely their sum $\nabla \cdot \boldsymbol{v}$), similar to prescriptions used in Federrath et al. (2010) and Gong & Ostriker (2013). This had negligible effects.

APPENDIX B: RESOLUTION DEPENDENCE OF IMF STATISTICS WITH VARIOUS MASS CUTS

In Section 4.2, we perform a resolution study of a $2000\,M_\odot$ GMC with mass resolution ranging from 0.1 to $10^{-4}\,M_\odot$ with cooling, MHD, and protostellar jet physics enabled, and found that the predicted IMF statistics stabilized at sufficient resolution (Fig. 12). In Figs B1 and B2, we remake the relevant panels from Fig. 12 while cutting stars $<\!0.1\,M_\odot$ and $<\!1\,M_\odot$, respectively, to determine the resolution requirements for statistics computed on different mass ranges of the IMF. Cutting at $<\!0.1\,M_\odot$ (Fig. B1), a mass resolution of $\approx\!2\times10^{-3}\,M_\odot$ ap-




Figure B1. Effect of numerical resolution upon various IMF statistics in a GMC simulation with cooling, MHD, and jets, as in Fig. 12, but computed after cutting masses $<0.1 \, M_{\odot}$.




Figure B2. Effect of numerical resolution upon various IMF statistics in a GMC simulation with cooling, MHD, and jets, as in Fig. 12, but computed after cutting masses $<1\,M_{\odot}$.

pears marginally sufficient to predict the mean stellar mass, and $10^{-3} \,\mathrm{M}_{\odot}$ is marginally sufficient to predict the median. Cutting at $1 \, \mathrm{M}_{\odot}$ (Fig. B2), $0.01 \, \mathrm{M}_{\odot}$ is sufficient for all three statistics. This suggests that the effect of numerical resolution is simply to impose a lower completeness limit on the predicted IMF, without seriously affecting larger masses (to a point). Rigorous comparisons with the observed IMF should ideally take both observational and numerical incompleteness functions into consideration.

APPENDIX C: DUST OPACITY FITS

For the opacities used in our RHD treatment of the IR band (Section 4.5), we fit to results from Semenov et al. (2003) for the 'porous five-layered sphere' composition as a function of both dust temperature $T_{\rm dust}$, which determines the dust composition, and radiation temperature $T_{\rm rad}$, which affects the opacity seen by the radiation. We assume a dust sublimation temperature of 1500 K,

above which we assume dust to be absent and the opacity to be zero. Otherwise, if $T_{\text{dust}} < 1500 \,\text{K}$, we use the fit

$$\kappa_{\text{dust,IR}} = f_{\text{d}} \exp(0.57 \max(x - 7, 0)) \times \exp(c_1 + c_2 x + c_3 x^2 + c_4 x^3 + c_4 x^4),$$
 (C1)

where $x = 4 \log_{10} (T_{\text{rad}}/\text{K}) - 8$, f_{d} is the local dust-to-gas ratio and the coefficients c vary with the dust temperature range as

$$c = \begin{cases} (0.728, 0.751, -0.0722, -0.0116, 0.00249) & T_{\rm dust} < 160 \, {\rm K}, \\ (0.166, 0.701, -0.0423, -0.0113, 0.00213) & 160 \, {\rm K} \le T_{\rm dust} < 275 \, {\rm K}, \\ (0.0358, 0.684, -0.0379, -0.0113, 0.00213) & 275 \, {\rm K} \le T_{\rm dust} < 425 \, {\rm K}, \\ (-0.766, 0.571, -0.0123, -0.0104, 0.00198) & 425 {\rm K} \le T_{\rm dust} < 680 \, {\rm K}, \\ (-2.24, 0.812, 0.0801, 0.00862, -0.00272) & 680 {\rm K} \le T_{\rm dust} < 1500 \, {\rm K}. \end{cases}$$

This paper has been typeset from a TFX/IATFX file prepared by the author.