











Modelling the evolution and influence of dust in cosmological simulations that include the cold phase of the interstellar medium

James W. Trayford ¹★, Joop Schaye ², Camila Correa ², Sylvia Ploeckinger ³, Alexander J. Richings ^{4,5}, Evgenii Chaikin ², Matthieu Schaller ^{2,6}, Alejandro Benítez-Llambay ⁷, Carlos Frenk ⁸ and Filip Huško ²

¹*Institute of Cosmology and Gravitation, University of Portsmouth, Dennis Sciama Building, Burnaby Road, Portsmouth PO1 3FX, UK*

²*Leiden Observatory, Leiden University, PO Box 9513, NL-2300 RA Leiden, the Netherlands*

³*Department of Astrophysics, University of Vienna, Türkenschanzstrasse 17, A-1180 Vienna, Austria*

⁴*Centre for Data Science, Artificial Intelligence and Modelling, University of Hull, Cottingham Road, Hull HU6 7RX, UK*

⁵*E. A. Milne Centre for Astrophysics, Department of Physics and Mathematics, University of Hull, Cottingham Road, Hull HU6 7RX, UK*

⁶*Lorentz Institute for Theoretical Physics, Leiden University, PO Box 9506, NL-2300 RA Leiden, the Netherlands*

⁷*Dipartimento di Fisica G. Occhialini, University of Milano-Bicocca, Piazza della Scienza, 3, I-20126 Milano MI, Italy*

⁸*Institute for Computational Cosmology, Department of Physics, Durham University, South Road, Durham DH1 3LE, UK*

Accepted 2025 November 14. Received 2025 October 21; in original form 2025 May 19

ABSTRACT

While marginal in mass terms, dust grains play an outsized role in both the physics and observation of the interstellar medium (ISM). However, explicit modelling of this ISM constituent remains uncommon in large cosmological simulations. In this work, we present a model for the life-cycle of dust in the ISM that couples to the forthcoming COLIBRE galaxy formation model, which explicitly simulates the cold ISM. We follow six distinct grain types: three chemical species, including carbon and two silicate grains, with two size bins each. Our dust model accounts for seeding of grains from stellar ejecta, self-consistent element-by-element metal yields and growth by accretion, grain size transfer (shattering and coagulation) and destruction of dust by thermal sputtering in the ISM. We detail the calibration of this model, particularly the use of a clumping factor, to account for unresolved gas clouds in which dust readily evolves. We present a fiducial run in a 25^3 cMpc^3 cosmological volume that displays good agreement with observations of the cosmic evolution of dust density, as well as the $z = 0$ galaxy dust mass function and dust scaling relations. We highlight known tensions between observational data sets of the dust-to-gas ratio as a function of metallicity depending on which metallicity calibrator is used; our model favours higher normalization metallicity calibrators, which agree with the observations within 0.1 dex for stellar masses $> 10^9 M_{\odot}$. We compare the grain size distribution to observations of local galaxies, and find that our simulation suggests a higher concentration of small grains, associated with more diffuse ISM and the warm-neutral medium, which both play a key role in boosting H_2 content. Putting these results and modelling approaches in context, we set the stage for upcoming insights into the dusty ISM of galaxies using the COLIBRE simulations.

Key words: dust, extinction – galaxies: evolution – galaxies: ISM.

1 INTRODUCTION

While the interaction of stars and multiphase gas have been a primary focus in the field of galactic astrophysics, cosmic dust, a vital piece of the puzzle of how galaxies form, has received less attention. Despite comprising a marginal fraction of the baryonic mass within galaxies, these small, solid-state particulates in the interstellar medium (ISM) play a prominent role in our understanding of the physics of galaxies and their evolution. The influence of dust belies its marginal mass contribution and is felt in a variety of ISM physics; influencing cooling and heating of gas, fostering

conditions for star formation, and various dynamical effects. Dust depletes metals from the gas phase (e.g. E. B. Jenkins 2009; A. De Cia et al. 2016), catalyses molecular hydrogen (H_2) formation (S. Cazaux & A. G. G. M. Tielens 2002), can function as a ‘critical coolant’ in astrophysical gas (including the intracluster medium, the ISM of elliptical galaxies, protostellar clouds e.g. W. G. Mathews & F. Brighenti 2003; L. A. Montier & M. Giard 2004; G. Dopcke et al. 2011), and mediates radiative transfer through absorption and re-emission of stellar radiation (e.g. K. E. Whitaker et al. 2017; I. Barišić et al. 2020; F. Galliano et al. 2021).

In addition to its physical importance, dust shapes much of what we can learn from galaxy observations, due to the reprocessing and modulating of radiation. Dust directly shapes spectral energy distributions (SEDs) through extinction curves and dust-to-gas ratios,

* E-mail: james.trayford@port.ac.uk

which are essential nuisance parameters in deducing key physical properties, such as stellar masses and star formation rates, that dictate galaxy evolution across cosmic time. This is exacerbated by the fact that converting from the *attenuated* emergent SED to the *intrinsic* SED (that may be decomposed into simple stellar populations) is also a complex function of the star-dust geometry in galaxies (e.g. D. Narayanan et al. 2018; J. W. Trayford et al. 2020). Approximately half of the starlight produced in galaxies throughout cosmic time having been processed by dust and re-radiated in the far infrared (FIR) (J. D. Finke, S. Razzaque & C. D. Dermer 2010), and even higher proportions in high-redshift star-forming galaxies, where the processing of UV/optical radiation into infrared emission dominates the integrated light of star-forming galaxies (A. W. Blain et al. 2002). Atacama Large Millimeter/submillimeter Array (ALMA) has revealed significant dust content in distant galaxies ($z > 6$, e.g. M. Palla et al. 2024), underscoring its early role in obscuring and shielding star forming gas. Dust grain populations and abundances may vary widely across high-redshift systems and have differential effects on observables, further highlighting the need for detailed models of dust evolution (e.g. K.-C. Hou et al. 2019; G. L. Granato et al. 2021).

The dust grain size distribution, composition, and spatial configuration in galaxies is critical for understanding how radiation propagates through galaxies, but remain poorly constrained outside of the Milky Way and its satellites due to the complex interplay between dust production/destruction mechanisms and the ISM environment. The majority of dust is thought to comprise carbonaceous (C) grains and silicate-based grains (dominated by O, Si, Mg, Fe). Dust forms in stellar ejecta (e.g. supernovae, AGB stars), grows via accretion of gas-phase metals, and is destroyed by sputtering in hot plasmas, while shattering and coagulation modulate the grain size distribution in turbulent regions (e.g. H. Hirashita 2012; R. S. Asano et al. 2013a). These processes drive a transition in dominant dust growth mechanisms with metallicity: at low metallicity, stellar ejecta dominate, while accretion becomes critical at higher metallicities ($Z \gtrsim 0.1 Z_{\odot}$), explaining observed breaks in dust-to-gas versus metallicity relations (e.g. A. Rémy-Ruyer et al. 2014; K.-C. Hou et al. 2019).

Despite its importance, modelling dust evolution poses significant challenges due to additional computational complexity and profound uncertainties in physical processes like grain growth and destruction, often happening deep within dense, obscured media. Early modelling efforts focused on one-zone chemical evolution models (e.g. E. Dwek 1998), but recent advances have allowed incorporation of dust into (cosmological) hydrodynamical simulations, including dust models using either a minimal grain size distribution (via e.g. a simple *two-size* approximation, H. Hirashita e.g. 2015; K.-C. Hou et al. e.g. 2019; G. L. Granato et al. e.g. 2021), or more extensive size distribution, modelling many grain sizes (e.g. R. McKinnon et al. 2018). The two-size approximation, explored in *smoothed-particle hydrodynamics* (SPH) and *adaptive mesh refinement* (AMR) codes (across idealized galaxy, zoom-in and even cosmological, hydrodynamical simulations), tracks separate populations of small ($r_{\text{grain}} \lesssim 0.03 \mu\text{m}$) and large grains ($r_{\text{grain}} \gtrsim 0.1 \mu\text{m}$) to capture size-dependent effects critical for modelling extinction curves and cooling rates, while forgoing excessive computational cost and the memory footprint associated with tracking a full size distribution per resolution element (H. Hirashita 2015; S. Aoyama et al. 2017; K.-C. Hou et al. 2019). However, spatial inhomogeneity of dust properties, evident in Milky Way extinction curve variations and observations of nearby galaxies (e.g. K. D. Gordon et al. 2003), and the need for multispecies treatments (carbon versus silicate grains, e.g. I. Shivaei

et al. 2020) remain important challenges (G. L. Granato et al. 2021; Y. Dubois et al. 2024).

The dust properties that emerge across galaxies and galaxy populations are a product of galaxy evolution, with dust itself driving ongoing galaxy-scale processes through its physical influence in the ISM. To capture this connection, models should consider dust evolution in a realistic galaxy formation context, and include the interdependence of dust and the ISM wherever relevant. Idealized simulations of Milky Way like discs can provide a detailed and high-cadence test bed for dust evolution, where it is feasible to explore variations and compare to a wealth of observational data. Including modelling for the dust life-cycle in this context provides insight into the build up of dust in our galaxy, and the processes needed to yield the extinction observed in the Milky Way (e.g. R. McKinnon, P. Torrey & M. Vogelsberger 2016; S. Aoyama et al. 2017; K.-C. Hou et al. 2019; G. L. Granato et al. 2021). When also paired with a multiphase model for the ISM, we begin to capture the range of environments relevant for dust evolution self-consistently; from the cold, dense ISM in which dust grows, to the hot, diffuse shocks in which it is destroyed (e.g. C. R. Choban et al. 2022; Y. Dubois et al. 2024). Simpler dust models may also provide insights in this multiphase ISM context. For example A. J. Richings et al. (2022) use an empirical, instantaneous model for dust depletion to show that it is necessary to reproduce FIR emission lines. Overall, however, idealized simulations lack the cosmological context, and thus the co-evolution of dust with metal enrichment and the structural assembly of galaxies over cosmic time.

Cosmological zoom-in simulations provide a means to self-consistently capture dust-gas co-evolution over cosmic time-scales. These can follow how dust properties emerge with the enrichment of the ISM, avoiding the artificial initial conditions of isolated disc models (E. Gjergeric et al. 2018; C. R. Choban et al. 2024; Y. Dubois et al. 2024; C. R. Choban et al. 2025). Zoom-in simulations can afford more computationally and memory intensive representations of the ISM, allowing more sophisticated modelling of physical processes and finer spatial resolution. However, these simulations lack the statistical population sampling of larger volume simulations. Large cosmological volume simulations including dust allow galaxy dust scaling relations and distributions to be tested, providing an exacting test of dust models (e.g. P. Dayal, H. Hirashita & A. Ferrara 2010; K. Bekki 2013; R. McKinnon et al. 2018; Q. Li, D. Narayanan & R. Davé 2019). Previously, volume simulations lacked the physics to explicitly trace the cold ISM, imposing an equation of state on the ISM gas, so cannot explicitly model the cool and cold ($\ll 10^4$ K) ISM that is critically important for dust evolution processes, pushing more of the relevant physics to subgrid modelling. Given these limitations for population scales for simulations, semi-analytic models have provided insight into the potential balance of dust evolutionary processes, and have provided a means to explore the poorly understood parameter space of dust modelling on a population scale (e.g. M. D. Bennassuti et al. 2014; H. Hirashita 2015; G. Popping, A. Puglisi & C. A. Norman 2017; A. P. Vijayan et al. 2019; R. M. Yates et al. 2024).

Despite progress with simulations and insights from semi-analytical models, discrepancies persist between models and observations of high-redshift dust-rich galaxies, suggesting models may be insufficient at early cosmic epochs (e.g. M. Parente 2025). Future efforts must integrate detailed grain physics with multiphase ISM modelling to explain the observed diversity in extinction curves, e.g. the Milky Way's prominent 2175 Å bump versus smoother features in Magellanic Clouds (e.g. S. Noll et al. 2009), and reconcile dust

evolution across metallicity and redshift regimes (e.g. Y. Dubois et al. 2024).

In this work, we present a dust evolution model for incorporation into the upcoming COLIBRE suite of next-generation cosmological hydrodynamical simulations (E. Chaikin et al. 2025; J. Schaye et al. 2025), which include cold gas physics in the multiphase ISM. We track dust via a relatively lightweight grain representation, comprising two grain sizes and three material species (one carbonaceous and two silicate subspecies). In Section 2, we first briefly describe the elements of the COLIBRE galaxy formation model outside the dust model presented here, focussing on those most relevant or connected to our dust implementation. In Section 3 we then describe the components of our dust model itself (in terms of dust composition and life-cycle, with gas–dust and dust–dust interaction processes). We then explore the connection between dust and ISM physics, and salient model parameters mediating this in Section 3.7. We present results of this model in our fiducial test simulation run in Section 4, focussing on dust properties and scaling relations of the low-redshift ($z = 0.1$) population. We discuss our overall findings and present a summary of the paper and main conclusions in Section 5.

We aim to describe our dust model in a standalone context, but we note that due to the interconnected nature of the dust presented in this work with other novel modules constituting the COLIBRE galaxy formation model, and their co-development, this is not completely possible. We describe elements of other works or refer to them where necessary, and use some forward referencing to motivate choices made in the dust model which were made in the context of developing the COLIBRE model.

2 THE COLIBRE MODEL AND SIMULATIONS

The dust evolution model we present below (Section 3) is integrated into the COLIBRE model of galaxy formation physics (J. Schaye et al. 2025), run using the gravity and hydrodynamics solver SWIFT (M. Schaller et al. 2024). In the multiresolution parlance of COLIBRE, we use m6 to refer to a gas (DM) particle mass of $1.84 \times 10^6 M_\odot$ ($2.42 \times 10^6 M_\odot$). We note that COLIBRE employs $4\times$ more dark matter particles than gas, achieving relatively high DM resolution. In this section, we first describe the general COLIBRE model, and then provide technical details of the simulation runs and parameter variations used in this work for context (Section 2.5).

2.1 Energy feedback

The COLIBRE model accounts for the influence of energetic processes in the ISM that transfer energy to the gas phase, termed *feedback*. These processes are contributed by other baryonic phases; particularly stellar and black hole particles.

For a nascent stellar population, the very first feedback processes are the ionizing radiation, stellar winds, and radiation pressure exerted by the most massive stars. These processes are collectively termed early stellar feedback and detailed in A. Benítez-Llambay et al. (in preparation). In short, the stellar age- and metallicity-dependent rates for stellar wind momentum injection and hydrogen ionization are from BPASS (J. J. Eldridge et al. 2017; E. R. Stanway & J. J. Eldridge 2018), assuming a Chabrier (C03) IMF. The BPASS spectra are further processed through a pressure-dependent column density using CHIMES (A. J. Richings, J. Schaye & B. D. Oppenheimer 2014a, b) to calculate the absorbed radiation, and therefore radiation pressure, around the young star particle. The available momentum from stellar winds and radiation pressure is injected stochastically as kinetic energy with a target velocity of

50 km s^{-1} . The hydrogen ionization rate is used to calculate the time-dependent size of H II regions (Strömgren spheres) around young star particles, for which a temperature floor of 10^4 K is imposed.

Once a massive star reaches the end of its life, it may erupt in a violent core-collapse supernova (CCSN). While individual SNe are below the resolution of cosmological simulations, when acting in concert these CCSNe may drive large-scale galactic winds, forcing gas away from star forming regions and out of the galaxy itself. COLIBRE implements this ‘superbubble feedback’, where each feedback event comprises many CCSN to mitigate numerical cooling losses. The CCSN energy budget is computed per time-step, where the number of stars expiring in an SSP is inferred from tabulated lifetimes of stars (L. Portinari, C. Chiosi & A. Bressan 1998) with masses of $\geq 8 M_\odot$ given our C03 IMF. This energy can be released back into surrounding gas following the thermal, stochastic approach of EAGLE (C. Dalla Vecchia & J. Schaye 2012). This approach is refined following E. Chaikin et al. (2023), where thermal and kinetic injection are combined. The kinetic component, constituting 10 per cent of the energy, goes into diametrically opposed, low-velocity (target of $\Delta v = 50 \text{ km s}^{-1}$) kicks for particle pairs. The relatively low energy increment required for these kicks ensures good sampling of feedback events in the kinetic component, and has been shown to induce turbulence in the surrounding media, helping to reduce star formation (E. Chaikin et al. 2023).

For the thermal part, the target heating temperature, $\Delta T_{\text{SN}}(n_{\text{H}})$, is an increasing function of local gas density with resolution-dependent floor and ceiling values ($6.5 < \log_{10} \Delta T_{\text{SN, min}}/\text{K} < 7$ and $7.5 < \log_{10} \Delta T_{\text{SN, max}}/\text{K} < 8$) calibrated from test runs. In the intermediate range between these extrema, $\Delta T_{\text{SN}}(n_{\text{H}}) = 10^{6.5} \text{ K}(n_{\text{H}}/n_{\text{H, pivot}})^{2/3}$. This approach mitigates overly destructive feedback and extensive bubbles in low-density gas. Finally, we adopt the statistically isotropic gas particle selection for energy injection of E. Chaikin et al. (2022), with the thermal feedback also using a variable fraction of energy per single CCSN in units of 10^{51} erg . This is subject to a variable coupling efficiency, f_{E} , to avoid overcooling at high density (R. A. Crain et al. 2015). In COLIBRE this is implemented as a sigmoid in SSP birth pressure, calibrated to reproduce the $z = 0$ galaxy stellar mass function and galaxy sizes (E. Chaikin et al. 2025).

For completeness, COLIBRE also implements SNIa feedback. This is purely thermal, following the CCSN thermal component implementation, with the exception that a fixed $f_{\text{E}} = 1$ is used. Rather than using stellar lifetimes, an exponential *delay time distribution* (DTD) provides the number of SNIa per time-step, calibrated to reproduce the observed cosmic SNIa rate (F. Nobels et al., in preparation). We note that SNIa have marginal impact relative to CCSN.

Alongside stars, COLIBRE accounts for feedback by *supermassive black holes* (SMBHs). SMBH particles are first seeded for friends-of-friends haloes (identified on-the-fly) in a resolution-dependent way, using an SMBH particle of fixed mass ($3 \times 10^4 M_\odot$ at m6), for haloes above a halo friends-of-friends (FoF) mass ($10^{10} M_\odot$ at m6). These particles can then accrete gas, merge and feed energy back to the surrounding media. The COLIBRE model includes two prescriptions for active galactic nucleus (AGN) feedback: purely thermal energy injections and a hybrid model combining thermal and kinetic jet feedback. In this work, we employ the purely thermal feedback prescription. Following similar arguments for effective feedback in CCSN, feedback is quantized to heat surrounding gas by a target temperature increment, ΔT_{AGN} , to mitigate numerical cooling losses. SMBHs build up an energy reservoir until enough is accrued for a single feedback event. The energy for feedback

comes from gas accretion; mass energy is converted to a luminosity, that couples back to the gas with fixed efficiency (5 percent at m6 resolution). In COLIBRE, ΔT_{AGN} is a function of black hole mass, again enforcing floor and ceiling values ($\log_{10} \Delta T_{\text{AGN, min}}/\text{K} = 6.5$ and a resolution-dependent $9.5 < \log_{10} \Delta T_{\text{SN, max}}/\text{K} < 10$) helping to maintain a reasonable sampling and regularity of BH feedback across the mass range of AGN. Feedback is meted out to the gas particle closest to the SMBH at feedback time.

Following S. Ploeckinger & J. Schaye (2020) and S. Ploeckinger et al. (2025), we also account for ambient energetic fields acting on gas particles, particularly the UV/X-ray background (via C.-A. Faucher-Giguère 2020), an interstellar radiation field (ISRF) and an ionizing cosmic ray background. The ISRF normalization and cosmic ray rate, as well as the shielding column density, vary with gas pressure to reflect environments with different levels of nearby star formation and with different coherence lengths.

2.2 Star formation

Our star formation (SF) prescription is detailed in F. S. J. Nobels et al. (2024). Star formation rates are applied assuming a Schmidt law, with an SF efficiency per free-fall time of 1 per cent, which yields good agreement with the Kennicutt–Schmidt law observed locally (e.g. B. B. Ochsendorf et al. 2017; D. Utomo et al. 2018; R. Pokhrel et al. 2021). To be eligible for star formation, gas particles must satisfy a gravitational instability criterion, comparing turbulent and thermal velocity support to the gravitational field within their SPH kernel.

Star formation rates are computed for eligible gas particles, which correspond to a probability of wholesale conversion into stars over a given time increment. Star formation then proceeds stochastically, sampling gas particles that are to be converted to star particles in the following time-step.

2.3 Stellar mass-loss and chemical enrichment

As star particles form in our simulations, their stellar mass-loss and chemical enrichment are modelled following the approach described in C. Correa et al. (in preparation), building on R. P. C. Wiersma et al. (2009). We track the evolution of 14 chemical elements: H, He, C, N, O, Ne, Mg, S, Si, Ca, Fe, Sr, Ba, and Eu. For simplicity, S and Ca are assumed to trace Si at fixed solar abundance ratios. Each star particle represents a simple, coeval stellar population (SSP). We follow the evolution of the relevant enrichment phases: CCSN and asymptotic giant branch (AGB) evolution, based on the mass- and metallicity-dependent stellar lifetimes from L. Portinari et al. (1998). These lifetimes are used at each time-step to determine when an SSP injects mass and metals into the surrounding gas particles.

Stellar enrichment is split into two components: nucleosynthetic (new elements formed during the star’s lifetime), and throughput (pre-existing elements present at birth that are carried out with the ejected material, and assumed to be well mixed within the star). AGB enrichment is attributed to stars with initial masses of 1–8 M_{\odot} . The AGB yields are drawn from a composite of studies, including C. L. Doherty et al. (2014), C. K. Fishlock et al. (2014), A. I. Karakas (2010), A. I. Karakas & M. Lugaro (2016), and G. C. Cinquegrana, M. Joyce & A. I. Karakas (2022). CCSN enrichment comes from stars born with masses of 8–40 M_{\odot} . These stars also undergo significant mass-loss through stellar winds both during and after the main sequence. We adopt CCSN and pre-supernova wind yields from K. Nomoto, C. Kobayashi & N. Tominaga (2013)

and C. Kobayashi et al. (2006), respectively. R. P. C. Wiersma et al. (2009) noted that relative abundance ratios predicted from nucleosynthesis yields may only be accurate within a factor of ≈ 2 , even assuming a fixed IMF. To account for this uncertainty, we apply ‘boost factors’ to the CCSN yields, increasing the yields of C and Mg from massive stars by a factor of 1.5, calibrated to better match observed stellar abundance ratios, especially those relative to Fe from APOGEE (H. Jönsson et al. 2018). Our simulations use the C03 IMF spanning the mass range 0.1–100 M_{\odot} . Stars born with masses greater than 40 M_{\odot} are assumed to collapse directly into black holes, re-accreting all their material and locking it out of the ISM.

For SNIa, we adopt an exponential DTD. The minimum delay time is set at 40 Myr, corresponding to the time required for the first stars in binary systems to evolve into compact remnants capable of accreting from a companion. The DTD time-scale and normalization are calibrated to match the observed cosmic SNIa rate (F. Nobels et al., in preparation). SNIa yields are taken from S.-C. Leung & K. Nomoto (2018). Rare r-process elements are treated similarly, with a DTD approach, but their production is attributed to neutron star mergers, common envelope jet supernovae, and collapsars. These sources use different calibrations and functional forms, along with a stochastic implementation, to reflect the rarity and variability of such enrichment events.

Enrichment time-steps for star particles are designed to balance accurate sampling of stellar evolution within each SSP against computational cost. To capture the early, most active phase of enrichment, we use fine time-steps (≤ 1 Myr) for SSPs younger than 40 Myr. For older populations (> 100 Myr), the time intervals after which mass transfer is performed increase proportionally to age.

Dust seeding is integrated into this enrichment framework, budgeting material from both CCSN and AGB yields to self-consistently form (and deplete, see Section 3.9) dust, as described in Section 3.2, providing a unified picture of the build-up of metals through cosmic time.

2.4 COLIBRE heating and cooling without live dust

A key aspect for this work is the treatment of dust in the cooling module. The fiducial treatment of cooling processes in COLIBRE is fully described in S. Ploeckinger et al. (in preparation) (hereafter, P25). In short, the species abundances of H, He, and free electrons are evolved for each gas particle¹ in non-equilibrium throughout the simulation by the chemical network CHIMES (A. J. Richings et al. 2014a, b), integrated in SWIFT. The species abundances of C, N, O, Ne, Mg, S, Si, Ca, and Fe are pre-tabulated with the standalone version of CHIMES, assuming ionization equilibrium and steady-state chemistry. The total cooling rates are the sum of the cooling rates from the non-equilibrium and the equilibrium species. The equilibrium cooling rates are rescaled to account for the non-equilibrium densities of free electrons from H and He species. Of the physics modelled in this module, dust-related processes are naturally most relevant for this work, and are a key aspect in certain physical conditions.

Given the important role of dust in gas heating and cooling, the model of P25 itself requires a prescription for dust content to

¹While this limited set was chosen to reduce runtime, we note our model is compatible with the with tracking all elements present in the full CHIMES chemical network in non-equilibrium.

compute cooling and heating rates if a live dust model (as described in Section 3) is *not* used. This is handled through an assumed instantaneous dust-to-gas ratio, $\mathcal{DTG}_{\text{hyb}}$, based on the physical state of the gas², which we hereafter refer to as P25 dust. This is the dust prescription used by the cooling model of our *uncoupled* dust runs (i.e. *Uncoupled* and *SeedOnly*), while in other runs we use the model described in Section 3.

For P25 dust, it is assumed that grains are immediately destroyed at $T > 10^5$ K and do not exist in relatively diffuse gas, represented by a minimum hydrogen column density, ($N_{\text{min}} = 3.1 \times 10^{15} \text{ cm}^{-2}$), such that $\mathcal{DTG}_{\text{hyb}}(T > 10^5 \text{ K}) = 0$ and $\mathcal{DTG}_{\text{hyb}}(N_{\text{H}} \leq N_{\text{min}}) = 0$. This corresponds to a range in volumetric hydrogen number densities at which shielding processes can occur of $n_{\text{H}} > 10^{-8} \text{ cm}^{-3}$. At high column densities ($N_{\text{H}} > 10^{20} \text{ cm}^{-2}$), a constant dust-to-metal ratio is assumed, such that³ $\mathcal{DTG}_{\text{hyb}} \propto Z_{\text{gas}}/Z_{\odot}$. For intermediate column densities ($3.1 \times 10^{15} < N_{\text{H}}/\text{cm}^{-2} < 10^{20}$), an additional taper is introduced to smooth the transition between the metallicity-dependent $\mathcal{DTG}_{\text{hyb}}(N_{\text{H}} = 10^{20} \text{ cm}^{-2}) \propto Z_{\text{gas}}/Z_{\odot}$ and $\mathcal{DTG}_{\text{hyb}}(N_{\text{H}} = 3.1 \times 10^{15}) = 0$, assuming a power law in column density with index $\alpha = 1.4$. This is chosen based on the Kennicutt-Schmidt relation (R. C. Kennicutt 1998), and assuming a scaling between the star formation rate density and dust production. At high density and solar metallicity, the \mathcal{DTG} is normalized to a representative MW value of $\mathcal{DTG}_{\text{MW}} = 6.6 \times 10^{-3}$.

One way dust influences the thermal and chemical state of the gas is through grain–gas collisions. This plays a particularly important role for molecule formation, which can happen readily on grain surfaces. For the H_2 formation rate on dust grains, R_{H_2} , CHIMES adopts the approach of S. Cazaux & A. G. G. M. Tielens (2002) (their equation 18)

$$R_{\text{H}_2} = \frac{1}{2} n_{\text{HI}} v_{\text{HI}} n_{\text{d}} \sigma_{\text{d}} \epsilon_{\text{H}_2} S_{\text{H}}, \quad (1)$$

where v_{HI} and n_{HI} are the thermal velocity and the physical number density of atomic hydrogen,⁴ ϵ_{H_2} is the recombination efficiency of molecular hydrogen, S_{H} is a dimensionless sticking coefficient, and $n_{\text{d}} \sigma_{\text{d}}$ is the number density of dust grains times the grain cross section, i.e. the total cross-section per unit volume. While S_{H} and σ_{d} are temperature-dependent, a constant dust temperature of 10 K is assumed.⁵ It is then the $n_{\text{d}} \sigma_{\text{d}}$ term that depends on the local mass fraction and size distribution of grains, with the other terms depending on the material properties and local gas properties. Dust grains can also contribute to catalysis of ion recombination via available dust surfaces (see e.g. J. C. Weingartner & B. T. Draine 2001), and transfer energy between the dust and gas (A. J. Richings et al. 2014b).

The influence of dust can also be felt indirectly by processing radiative energy via extinction. A particularly important example of this is photoelectric heating by dust grains; absorption of FUV photons by dust produces energetic free electrons that contribute to a heating of the gas component. In CHIMES, the volumetric heating rate, Γ_{PE} , follows that of M. G. Wolfire et al. (2003) (their

equations 19 and 20⁶), and is given by

$$\Gamma_{\text{PE}} = (1.3 \times 10^{-24} \text{ erg cm}^{-3} \text{ s}^{-1}) \epsilon G_0, \quad (2)$$

where G_0 is the integrated incident UV radiation energy density, and ϵ represents the heating efficiency of the gas.⁷ Here, the pre-factor $1.3 \times 10^{-24} \text{ cm}^{-3} \text{ s}^{-1}$ depends on the abundance as well as mass and material properties of grains (see P25). Similarly, incidence of photodissociating and photoionizing photons are reduced through shielding of gas (A. J. Richings et al. 2014a, b). The extinction, A_V/N , also scales with the $n_{\text{d}} \sigma_{\text{d}}$ term to first order.

With the grain cross-section per unit volume the relevant quantity for all of these processes, we can modulate them using the masses and sizes of grains traced through our live dust model. We describe their coupling in Section 3.7. A final dust-sensitive aspect to consider is the role of depletion. The same chemical elements that constitute dust also play an important cooling role in the ISM when existing in the gas-phase, particularly through metal-line cooling. These processes are handled via per-ion cooling rates tabulated for COLIBRE and based upon B. D. Oppenheimer & J. Schaye (2013) (P25 equation 19⁸). These rates are based on gas-phase metals, so it is necessary to assume a per-element depletion alongside the $\mathcal{DTG}_{\text{hyb}}$, in lieu of a live dust model. These assume depletion in the fixed relative proportions of the Milky Way as inferred by E. B. Jenkins (2009). We discuss how we then adjust these rates for the case of live dust in Section 3.7, consistent with the depletions we calculate for our dust model (Section 3.9).

2.5 Simulation runs

Prior to describing our dust model in detail (Section 3, next), we describe the simulations used in this work and our choice of parameter variations. Throughout this work we assume the 3×2 pt cosmology, combining all data and constraints presented with the DES Y3 results (T. M. C. Abbott et al. 2022); $h = 0.681$, $\Omega_{\text{m}} = 0.306$, $\Omega_{\text{b}} = 0.0486$, $\sigma_8 = 0.807$, $n_s = 0.967$.

We use cosmological volume initial conditions for the simulations in this work. This choice allows us to evolve a population of galaxies in a cosmological context, and inspect galaxy scaling relations, which is important to illustrate the behaviour of the model over the cosmological remit of COLIBRE. We focus on a 25^3 cMpc^3 volume, with an initial 376^3 gas particles of $m_{\text{gas}} = 1.8 \times 10^6 M_{\odot}$ and 4×376^3 dark matter particles with $m_{\text{DM}} = 2.4 \times 10^6 M_{\odot}$, representing matched ($4 \times$ higher) baryonic (dark matter) resolution relative to fiducial EAGLE (J. Schaye et al. 2015).

We identify collapsed structures using the FoF algorithm post processed using the publicly available HBT-HERONS subhalo finder described by V. J. Forouhar Moreno et al. (2025) to assign particles to self-bound substructures. This builds on the HBT + code for N -body simulations (J. Han et al. 2018), with critical improvements for the hydrodynamical context, considering the histories and hierarchical formation of subhaloes. In particular, the HBT-HERONS code exhibits more consistent allocation of mass between satellite and central galaxies during close passes, as compared to various other phase-space halo finders (Á. Chandro-Gómez et al. 2025). The subhalo locations and, optionally, the particles are then taken up by

²For the sake of brevity, we forgo a full derivation of these relationships here, but see Section 2.1 (and in particular equation 16) of P25.

³As in P25, we use the $Z_{\odot} = 0.0134$ value of M. Asplund et al. (2009) where appropriate.

⁴We note that dust properties assumed by the cooling module do not make use of the density boost introduced in Section 3.8.

⁵The H_2 formation rate was shown to be approximately constant for $6 \text{ K} \lesssim T_{\text{dust}} \lesssim 50 \text{ K}$ in fig. A1 of A. J. Richings et al. (2014a).

⁶This is an update of what is described in A. J. Richings et al. (2014a), described by A. J. Richings & C.-A. Faucher-Giguère (2018).

⁷Tabulated values, depending on the gas temperature T , the electron number density n_e , and incident UV energy density, G_0 .

⁸Updated from S. Ploekinger & J. Schaye (2020)

Table 1. Simulation variations used in this work. All runs represent a 25^3 Mpc^3 cosmological box, initialized with 376^3 gas particles of mass $m_g = 1.8 \times 10^6 M_\odot$ (and four times as many dark matter particles, with mass $m_{\text{DM}} = 2.4 \times 10^6 M_\odot$). We adopt ‘m6’ as shorthand for this resolution (from the COLIBRE parlance). Simulations are run using the COLIBRE galaxy formation model (Schaye et al., in preparation), using the `swift` code (M. Schaller et al. 2024, see Section 2.5). The columns from left to right indicate: (1) the simulation name, (2) if the dust model is coupled to the cooling (see Section 2.4), (3) if the model includes small grains rather than large, if we include processes of (4) accretion, (5) destruction, (6) or size transfer, (7) the maximum C value, (8) if C has a fixed value, (9) the size representing the small grain bin, (10) the chemical group assumed for silicates, and (11) the diffusion constant (Section 3.6).

(1) Name	(2) Coupled?	(3) Size	(4) Accretion	(5) Destruction	(6) Size transfer	(7) Max. C (a)	(8) Const. C	(9) Small grain r_{grain} [μm]	(10) Silicates	(11) C_d
Fiducial	✓	✓	✓	✓	✓	100.	✗	0.01	Olivine	0.01
FidUncoupled	✗	✓	✓	✓	✓	100.	✗	0.01	Olivine	0.01
SeedOnly	✗	✓	✗	✗	✗	n/a	✗	0.01	Olivine	0.01
NoDestruction	✓	✓	✓	✗	✓	100.	✗	0.01	Olivine	0.01
OneSize	✓	✗	✓	✓	✗	100.	✗	0.01	Olivine	0.01
MaxC10	✓	✓	✓	✓	✓	10.	✗	0.01	Olivine	0.01
constC30	✓	✓	✓	✓	✓	30.	✓	0.01	Olivine	0.01
NoC	✓	✓	✓	✓	✓	1.	✓	0.01	Olivine	0.01
PyroSil	✓	✓	✓	✓	✓	100.	✓	0.01	Pyroxene	0.01
SmallerGrains	✓	✓	✓	✓	✓	100.	✗	0.005	Olivine	0.01
NoDiff	✓	✓	✓	✓	✓	100.	✗	0.01	Olivine	0
LoDiff	✓	✓	✓	✓	✓	100.	✗	0.01	Olivine	0.001
HiDiff	✓	✓	✓	✓	✓	100.	✗	0.01	Olivine	0.1

Note. (a) i.e. the clumping factor

the SOAP tool (R. McGibbon et al. 2025), computing a wide range of aggregated properties using a range of 3D and projected apertures of varying physical size.

Here, we use gravitationally bound particles within a 50 pkpc, 3D aperture about the most-bound particle to define ‘galaxy’ properties, unless stated otherwise. These apertures are defined about the most-bound particle (i.e. closest to the minimum of the local potential well), as defined by the HBT-HERONS code.

The simulations run for analysis in this work and their defining parameters are listed in Table 1. The `Fiducial` run uses the full dust model described in Section 3, and also includes the effects of ‘coupling’ our dust model to gas cooling, heating and phase transition processes, detailed in Section 3.7. Our runs incorporate dust coupling by default (excepting `SeedOnly` where dust fractions are too unphysical), with an uncoupled counterpart to isolate the influence of the dust evolution model on the cooling, e.g. `FidUncoupled`. The `SeedOnly` and `NoDestruction` simulations isolate evolutionary processes. Size transfer effects are tested by `LargeGrainsOnly`. A number of runs are dedicated to exploring the influence of parameters; particularly a subgrid clumping factor, C , (described in Section 3.8; `NoC`, `ConstC30`, and `MaxC10`) and the turbulent diffusion constant, C_d (see Section 3.6; `NoDiff`, `LoDiff`, and `HiDiff`). We also present runs exploring pyroxene silicate chemistry, `PyroSil`, and the representative size of small grains, `SmallerGrains`. The use of coupled cooling physics by default means we see the impact of these dust model changes on the physics and resultant galaxy populations. We will explore the most pertinent variations alongside our primary low-redshift results in Section 4, as well as presenting some additional variations in Appendix A.

3 MODELLING THE FORMATION AND EVOLUTION OF DUST

We now describe the constituent elements of our dust model, including the assumed properties of grains and the evolutionary processes of grain seeding/nucleation, grain growth, grain destruction, and grain size transfer. Particular focus is placed on aspects of the modelling

that differ from previous studies. Our dust model is designed for compatibility with SWIFT SPH simulations (M. Schaller et al. 2016, 2024), particularly the upcoming COLIBRE simulations, as detailed above. All the processes detailed below (seeding, growth, and destruction of grains), are implemented on a particle-by-particle basis, using the properties of a given gas resolution element. Throughout, we characterize the dust content in terms of the mass fractions, either as the dust-to-gas ratio (DTG) or dust-to-metal ratio (DTZ). With the exception of direct destruction by supernovae (SNe) or astration, all our processes depend on gas density. For dense-gas processes (such as grain growth), we consider an *effective* density indicated using ‘ (i.e. ρ' , n'_H), which is related to the physical density by a *clumping factor*, C , as detailed in Section 3.8. The *total* element abundances include both the abundances of metals that remain in the gas-phase, which we refer to as the *gas-phase* component, and the abundance that is *depleted* into dust grains. Throughout, we use the fiducial COLIBRE stellar initial mass function (IMF) of G. Chabrier 2003, with a range $\in [0.1, 100] M_\odot$ (hereafter `C03`).

3.1 Types of dust grains

Interstellar dust grains are diverse, in terms of both their structure and chemical composition. Practically, representing the full diversity of grains is infeasible within a simulation, so a number of simplifying assumptions are necessary to represent dust. Here we discuss and motivate the assumptions that we make.

The chemical composition of grains (see Table 2) is potentially highly complex. For simplicity, we consider two primary species; *silicate* and *graphite* grains. This dichotomous representation is well established (e.g. J. C. Weingartner & B. T. Draine 2001). Carbonaceous grains are generally thought to be either homonuclear (i.e. graphite or diamond) or hydrocarbons. Consequently, the vast majority of their mass comprises carbon. We therefore assume pure carbon grains. This simplifies the budgeting of dust- and gas-phase carbon.

Conversely, silicates are generally heteronuclear, with a variety of different elements contributing significantly to their mass. Strong

Table 2. Composition and nucleation properties of the chemical species of grain adopted in our model. While graphite is homonuclear, the silicates are heteronuclear, comprising *fayalite* (Fe_2SiO_4) and *forsterite* (Mg_2SiO_4) subspecies. For seeding purposes, we assume an equal abundance of molecules in the *fayalite* and *forsterite* subspecies *fayalite* (i.e. an effective molecule FeMgSiO_4 , Section 3.2), yielding a value of $A_G = 171.853$ used in equation (3).

Species	Subspecies	Seed composition	A_G	η_{CCSN}	j	A_j	τ_G (Myr)
Carbonaceous	Graphite	C	12.01	0.15	C	12.01	180
Silicate (Olivine)	Fayalite, Forsterite	FeMgSiO_4	171.85	3.5×10^{-4}	Si Fe Mg	28.09 55.85 24.31	99.3

candidate minerals from stoichiometry are *olivine*, *pyroxene* and intermediate/composite grains (e.g. S. Fogerty et al. 2016). We explore both *olivine* and *pyroxene* mineral series, adopting *olivine* as our fiducial silicate. We then treat Mg and Fe end-members as distinct subspecies tracked independently: *forsterite* (Mg_2SiO_4) and *fayalite* (Fe_2SiO_4) in the case of olivine.⁹ While we limit the silicates to a particular mineral series, the choice to track Mg and Fe end-member silicates separately allows the ratio of these species to freely float given the local chemical abundances, such that they are not artificially limited by a fixed ratio of the two. In particular, the ratio of Mg to Fe species can then reflect variations in α -enrichment, $[\alpha/\text{Fe}]$.

A common simplifying assumption for grain structure is that of spherical grains. This allows a grain’s structure to be characterized by a single value, its radius r_{grain} . While this greatly simplifies the equations for dust evolution, the assumption of sphericity has certain limitations. Some degree of asphericity in cosmic dust is evident in harvested samples (e.g. the STARDUST mission, D. Brownlee 2014), and by the polarization of light in the ISM (e.g. J. Cho & A. Lazarian 2007). In addition, grains may have a degree of porosity, increasing their surface area and collisional and optical properties (H. Hirashita & V. B. Il’in 2022). While the degree of asphericity and porosity across interstellar environments is hard to constrain, the assumption of spherical, aporous grains is extreme in that it minimizes their surface-to-volume ratio, with implications for their interaction with gas or other grains. However, sphericity is degenerate with other physical properties modulating the growth/destruction rates of dust, which are themselves typically empirically determined or calibrated. For our purposes, the assumption of an effective spherical grain is adopted for simplicity.

Given our spherical grain paradigm, grain sizes are characterized solely by their radii, r_{grain} , and are another fundamental property of the dust. As for the chemical composition, representing an arbitrary number of grain sizes is not feasible for our simulations. Instead, we adopt fixed grain size bins. We employ a two-size grain model assuming a $0.1 \mu\text{m}$ and $0.01 \mu\text{m}$ radius for large and small grains, respectively.

In combination, this leads to three distinct chemical bins for dust grains; Carbon, Mg-Silicate and Fe-Silicate, each with two possible sizes *Large* and *Small*. We adopt this six-component dust model in all of our simulations.

We note that polycyclic aromatic hydrocarbon (PAH) grains are not explicitly traced in our modelling. These are extremely small ($< 1 \text{ nm}$), and as such, subjected to different forces (e.g. S. C. Madden et al. 2006; A. G. G. M. Tielens 2008; M. S. Murga et al. 2019), more in the realm of molecular chemistry. The creation mechanism for PAHs is uncertain (whether through grain shattering, condensation in AGB winds or spontaneously through chemical reactions; A. Li 2020), and the process of aromatization itself is driven by local-scale radiation (S.-J. Rau, H. Hirashita & M. Murga 2019), requiring radiative transfer calculations. With these caveats in mind, PAH

presence could be inferred using empirical fractions relative to the carbon depletion in the simulations.

3.2 Seeding grains

With our assumed grain properties characterized, we now first consider how grains are *seeded* before considering physical processes that act upon dust grains and the influence of grains in the ISM. We detail the channels of grain seeding below.

A requirement is that any dust yields should be compatible with the underlying nucleosynthetic yields from stars, such that the abundances of constituent elements of the dust are conserved overall. In the context of the COLIBRE galaxy formation model (see Section 2.5) we build on the composite yields of C. Correa et al. (in preparation), and budget dust-phase metals from overall metal yields and rates of mass return from stars of differing mass and chemical composition. In the C. Correa et al. (in preparation) chemical enrichment scheme, yields are sensitive to the particular enrichment pattern of the enriching star particle, due to their implementation as a linear combination of throughput elements (present at the creation of a star particle) and a net change owing to creation or destruction by stellar evolutionary processes. For a given total metal mass fraction (Z), we assume a fixed solar abundance pattern when displaying these dust yields, based on the M. Asplund et al. (2009) solar values used by C. Correa et al. (in preparation), hereafter C25.

3.2.1 AGB grain production

Although AGB and Super-AGB (SAGB) stars are a relatively transient phase of intermediate mass stars ($2 M_{\odot} < M_{\text{init}} < 8 M_{\odot}$), they are key to the enrichment of the ISM due to their prodigious mass-loss via strong stellar winds. In addition, the circumstellar envelopes around AGB stars are sufficiently cool and dense at distances of 3–10 stellar radii to facilitate the formation of dust in substantial quantities (e.g. A. S. Ferrarotti & H. P. Gail 2006).

Empirical studies of dust associated with (S)AGB have historically been limited to low-metallicity, extragalactic stars, due to the difficulty in obtaining accurate distances locally. *Gaia* parallax measurements should eventually serve to improve this situation (M. Di Criscienzo et al. 2016), probing (S)AGB up to solar metallicity. In the meantime, modelling of AGB evolution, ejecta chemistry and circumstellar formation of different dust grain species have been extended to higher metallicities through enhanced theoretical insights and observations.

We adapt the dust yields of F. Dell’Agli et al. (2017), which include AGB models up to solar metallicity. These are illustrated in their fig. 5. Dust yields are computed using the ATON stellar models (V. M. Canuto & I. Mazzitelli 1991) for AGB stars with zero-age main sequence (ZAMS) stellar masses over the range $1 \lesssim M_{\text{init}}/M_{\odot} \lesssim 8$ and fixed metallicities. We use the yields for three absolute stellar metallicities; $Z_* \in [0.004, 0.008, 0.018]$. These are consistent with the C. Correa et al. (in preparation) nucleosynthetic yields for AGB stars, with the exception of the highest metallicities, which use a slightly higher value of $Z = 0.019$. In this case, dust yields are scaled

⁹Enstatite (MgSiO_3) and ferrosilite (FeSiO_3) for pyroxenes.

up with metallicity from $Z = 0.018$. For consistency with C. Correa et al. (in preparation), we match AGB yields to stars in the initial mass range $1 \leq M_{\text{init}}/M_{\odot} \leq 8$, assuming no dust is yielded below the F. Dell’Agli et al. (2017) range (given yields fall precipitously towards the lowest ZAMS masses).

With the dust mass yields in place, *silicate* and *carbon* grains (see Table 2) are partitioned by the ZAMS stellar mass of AGB progenitors, assuming AGB with $M_{\text{init}}/M_{\odot} \leq 3.5$ produce carbon grains, with more massive stars producing silicates. In the lower mass subset, this is attributed to the formation of carbon enriched outer atmospheres and the formation of carbon grains in their circumstellar envelopes. In higher mass stars, hot bottom burning destroys surface carbon, and favours silicate production in stronger winds (R. Schneider et al. 2014; P. Ventura et al. 2014; M. Di Criscienzo et al. 2016; F. Dell’Agli et al. 2017).

3.2.2 SN grain production

In addition to synthesizing large quantities of metals, the shocks from SNe are thought to facilitate significant dust grain formation. Local SN remnants provide valuable case studies for these processes (e.g. L. Dunne et al. 2009; M. Shahbandeh et al. 2023; D. Milisavljevic et al. 2024). However, the overall dust yields from SNe are particularly uncertain, as strong reverse shocks can also destroy grains (e.g. F. Kirchschrager et al. 2019, 2023, see Section 3.4). At very early times, core-collapse SNe are the primary candidates for dust grain formation (e.g. E. Dwek 1998; A. G. G. M. Tielens 1998). Observations of $z \approx 7$ radio galaxies and quasars with $M_{\text{dust}} > 10^7 M_{\odot}$ have been taken to suggest that SN on average produce $\approx 1 M_{\odot}$ of dust (R. Valiante et al. 2009, 2011; D. Watson et al. 2015), while populations of galaxies exhibiting much lower dust-to-metal ratios have also been observed down to $z \approx 4$ (D. Burgarella et al. 2025).

It is less clear whether type Ia SNe (SNIa) net produce dust (e.g. T. Kozasa et al. 2009; T. Nozawa et al. 2011). Some studies have proposed SNIa as a viable dust production mechanism (D. D. Clayton, S. Amari & E. Zinner 1997; C. Travaglio et al. 1999). However, observations of local SNIa remnants do not generally show evidence of significant net dust formation (e.g. H. L. Gomez et al. 2012), though evidence of dust formation in a SNIa remnant has been observed by L. Wang et al. (2024). T. Nozawa et al. (2011) posits that this lack of evidence may be explained by a lower formation radius and subsequent rapid destruction in SNIa, preventing them from being an appreciable source of ISM dust grains. Given these findings, we omit dust production by SNIa, and consider only CCSN production.

We adapt CCSN dust nucleation rates from S. Zhukovska, H.-P. Gail & M. Tielens (2008) to compute dust yields. These are computed using the overall element yields of C. Correa et al. (in preparation), where the mass in a particular grain species is

$$M_G = \eta_G M_j \frac{A_G}{A_j}, \quad (3)$$

where M_G is the mass returned for a given grain species, M_j is the total mass returned of the ‘bottleneck’ element j , η_G is the species condensation efficiency for CCSN, A_G is the mass of the grain molecule in atomic units, and A_j is the atomic weight of element j . The bottleneck element is the element limiting the number of grain molecules that can be synthesized for a given SN, and for our assumed silicates varies by mass and metallicity of the progenitor star. This is determined by minimizing N_j/i where N_j is the number of elements produced and i the number of atoms in one molecule of the silicate species (S. Zhukovska et al. 2008).

While this approach dictates what fraction of each ejected material is deposited into dust grains, it makes no distinction between subspecies of grain that comprise like elements. Given no *a priori* preference, we adopt a fixed ratio of 0.5 between Fe- and Mg-endmember grains. This yields an effective silicate molecule of FeMgSiO_4 , treated as a single reservoir for grain seeding purposes. Note that while this approach does not allow the ratio of silicate grain types to free-float, as in accretion processes detailed below, the choice of ratio has little influence on our final dust composition, due to the importance of accretion in building the majority of dust mass in our model (see Section 3.3).

3.2.3 AGN grain production

Another possible avenue for grain nucleation is in the tori of AGN and their nearby environment. Partial motivation for this are studies identifying very dusty high-redshift systems, requiring $\approx 1 M_{\odot}$ of dust per CCSN (R. Valiante et al. 2009, 2011; D. Watson et al. 2015), while other studies identify very low levels of dust (e.g. R. J. Bouwens et al. 2012). This could point to a non-stellar channel, such as AGN, being important for nucleating dust. In addition, sightlines to AGN are often better fit with dust extinction and emission models unlike those of the Galactic ISM (e.g. E. Sturm et al. 2005; S. Srinivasan et al. 2017). While stellar sources are still needed to synthesize the metals that constitute grains, some models employ AGN tori as nucleation channels for dust (e.g. A. Sarangi, E. Dwek & D. Kazanas 2019). For simplicity, we neglect grain formation in AGN.

3.2.4 Size distribution of seeded grains

In addition to seeding the mass in each chemical species of grain, the initial size distribution of the seeded grains must also be set. We assume that seed grains are dominated by larger size ($a \approx 0.1 \mu\text{m}$). This reflects studies of AGB winds (M. A. T. Groenewegen 1997; Y. Yasuda & T. Kozasa 2012; R. S. Asano et al. 2013b), and is consistent with the inference of $r_{\text{grain}} \gtrsim 0.01 \mu\text{m}$ in core-collapse SNa (e.g. T. Nozawa et al. 2007). We initiate the size distribution with 90 per cent of the mass in large grains and the remaining 10 per cent in small grains. This is an order-of-magnitude estimate, chosen to be between the values seen in AGB winds (e.g. R. S. Asano et al. 2013a) and the observed Milky-Way (MW) mass distribution (e.g. J. S. Mathis, W. Rumpel & K. H. Nordisiek 1977). We note that the size distribution of seed grains, like their absolute masses, have little effect on the dust content of low-redshift galaxies that we present for our model, as this is set by the growth of grains and its equilibration with destructive processes in the ISM (see Section 4.2). However at higher redshift, or especially rapidly forming galaxies, the yield contribution can be more significant, with dust produced through yields before evolutionary processes can equilibrate.¹⁰

3.2.5 Comparing dust yields from seeding channels

Given the considerable uncertainty and variation in the predictions for dust seeding by AGB and CCSNe, it is worthwhile to consider how our curated yields (S. Zhukovska et al. 2008; F. Dell’Agli et al. 2017) may differ from other models. For comparison we consider the dust yields predicted by E. Dwek (1998), which have been employed

¹⁰In Section 4.3, we show that in our model the total contribution from direct stellar yields is subdominant to growth processes for $z \lesssim 11$ (Section 4.3, Fig. 7).

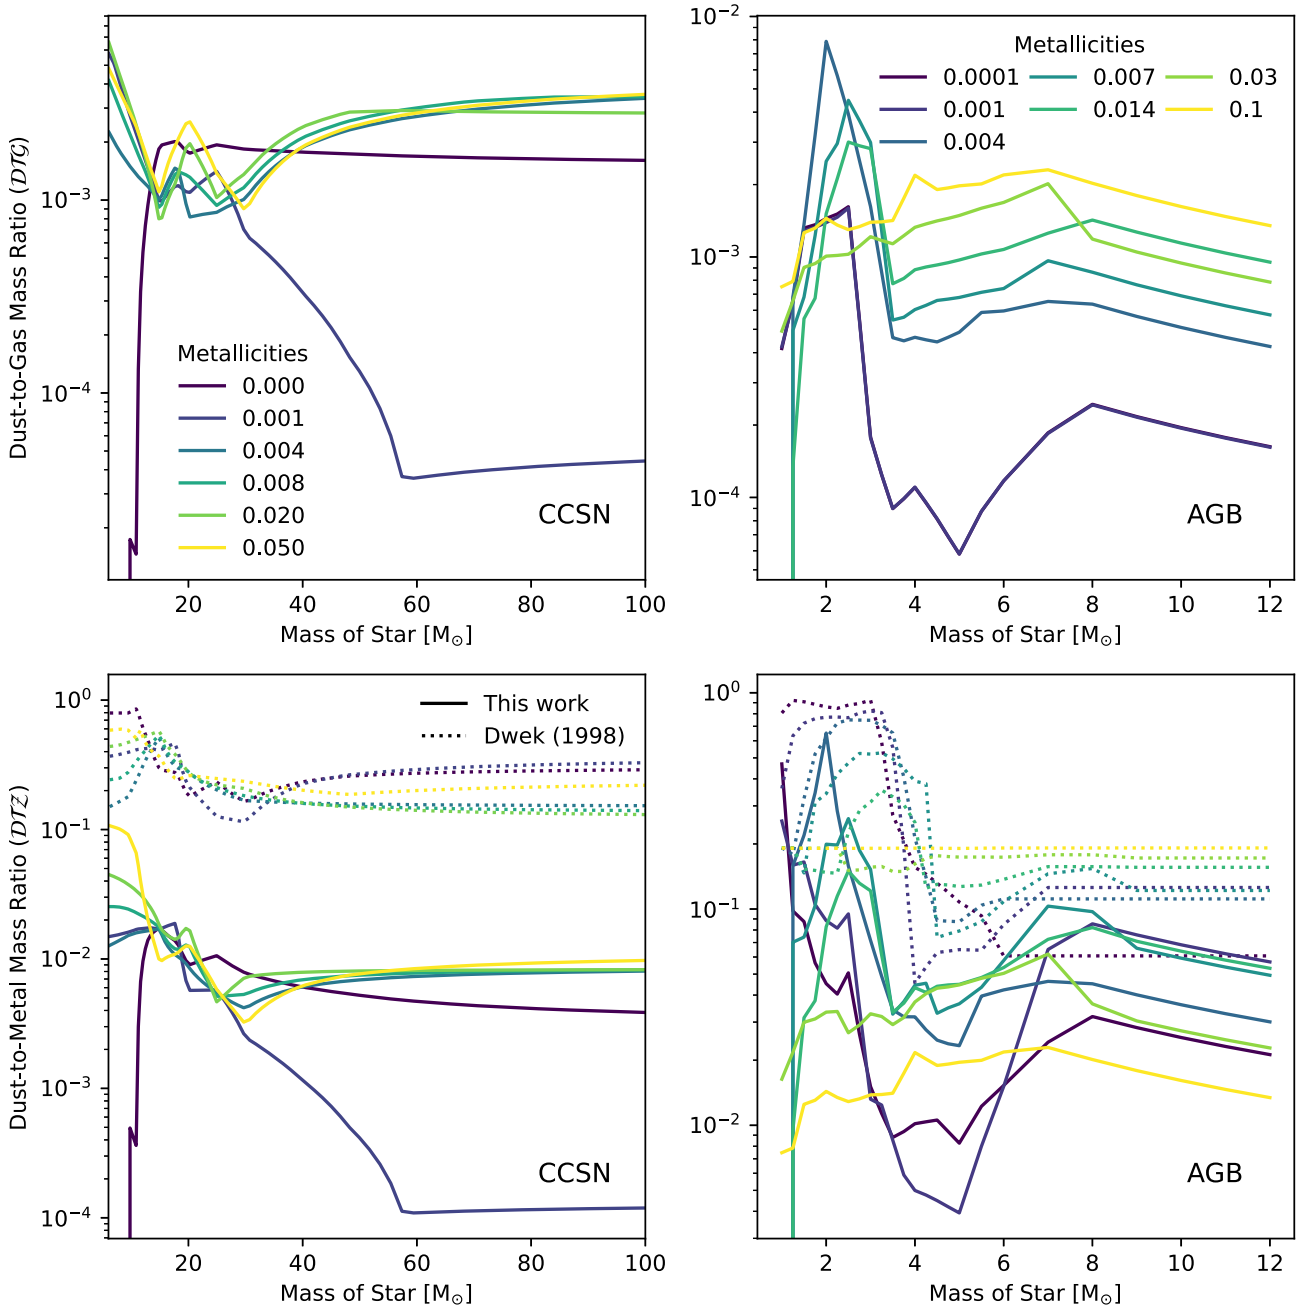


Figure 1. Stellar dust yields used in this work, as a function of zero-age main sequence (ZAMS) stellar mass. The top row shows the dust-to-gas ratio in the ejecta computed for CCSN (left) and AGB (right) channels using our adopted yields (*solid lines*, Section 3.2) for different absolute stellar metallicities (line colours). The bottom row shows the same, except that the y-axis is now the dust-to-metal ratio in the returned material, and the grain yields of E. Dwek (1998) are provided for comparison (*dashed lines*). We note that the two columns use different y-axis ranges. Yields are budgeted from the overall metal yields presented by C. Correa et al. (in preparation). For the contributions to the yields from elements passing through a star (i.e. present at birth), we assume the solar abundance pattern of C25, scaled by the stellar metallicity relative to solar (0.0129). A striking feature of this plot is that the yields and dust-to-metal ratios of E. Dwek (1998) tend to be higher than ours (F. Dell’Agli et al. 2017; S. Zhukovska et al. 2008) by about an order of magnitude. This difference is discussed further in the text.

by various galaxy formation models (e.g. R. McKinnon et al. 2016, 2017; R. Davé et al. 2019).

In Fig. 1 we show the dust yields (top row) and returned dust-to-metal ratios (bottom row) for CCSN (left column) and AGB (right column) as functions of their ZAMS mass, M_{init} . A striking feature is that the yields advocated by E. Dwek (1998) are higher than those we use by about an order of magnitude. In addition, the dust-to-metal

ratio of returned material in the E. Dwek (1998) model is of the same order as that seen in the local ISM (≈ 0.3 , A. K. Inoue 2003). This difference is profound as it implies different roles for evolutionary processes in reaching the dust content of galaxies like the Milky Way. Our model requires dust masses to grow by a factor ~ 10 in the ISM, while the E. Dwek (1998) yields require either no mass evolution, or an equilibrium of destruction and growth mechanisms. In particular,

our lower level of seeding requires modelling of the cold, dense ISM conducive to efficient grain growth (see Section 3.3), as opposed to a level of seeding requiring only maintenance. We note that some studies use yields of order 10 per cent but still require growth to reproduce observables (e.g. S. Lower et al. 2024; C. R. Choban et al. 2025).

While the E. Dwek (1998) yields are empirically motivated, with some arbitrary choices with regards to grain nucleation properties, some more recent observations of post-shock region of supernovae support this high-yield paradigm (e.g. M. Shahbandeh et al. 2023). The S. Zhukovska et al. (2008) yields are intended to be decoupled from subsequent molecular cloud growth, and are consistent with the limited measurements of pre-solar grains in the Solar system, as well as the lower end of supernovae observations. Ultimately, these differences showcase the uncertainty in the direct yields of grains and differential role of growth in dust models, explored further in the following section.

3.3 Grain growth

Once grains have been nucleated in the ejecta of stars, we can consider their evolution. One aspect of this is the growth of grains through the process of *accretion*. A time-scale argument comparing our dust injection into the ISM by stellar sources with dust destruction by sputtering suggests the need for significant growth of dust within the ISM to explain the observed dust content of galaxies (A. K. Inoue 2003; S. Zhukovska et al. 2008; B. T. Draine 2009; A. Pipino et al. 2010; R. Valiante et al. 2011; R. S. Asano et al. 2013a).

For accretion, we adapt the accretion time-scale equations of H. Hirashita & N. V. Voshchinnikov (2013), using

$$\tau_{\text{acc}} = \tau_{\text{G}} \left(\frac{r_{\text{g}}}{0.1 \mu\text{m}} \right) \left(\frac{\epsilon_{j,\odot}}{\epsilon_j} \right) \left(\frac{10 \text{ cm}^{-3}}{n'_{\text{H}}} \right) \left(\frac{10 \text{ K}}{T} \right)^{0.5} \left(\frac{0.3}{S_{\text{acc}}} \right), \quad (4)$$

where n'_{H} is the effective local number density of hydrogen, T is the local gas temperature, S_{acc} is the sticking probability for accretion, τ_{G} is the accretion time-scale normalization for a given grain species, and ϵ_j and $\epsilon_{j,\odot}$ represent the local and solar abundance of the bottleneck abundance in the gas-phase relative to hydrogen, respectively. For our purposes, we assume a fixed S_{acc} value of 0.3, with τ_{G} for silicate and carbon grains taken from H. Hirashita & N. V. Voshchinnikov (2013) and listed in Table 2. The fractional change in the mass of a given dust species due to accretion is computed over a time-step of Δt for each gas particle, as $\exp(\Delta t/\tau_{\text{acc}})$.

The main variation from H. Hirashita & N. V. Voshchinnikov (2013) is the $\epsilon_{j,\odot}/\epsilon_j$ factor, replacing a Z_{\odot}/Z factor. In the case of solar abundance patterns (assuming the R. P. C. Wiersma et al. 2009 abundances), these terms are equivalent, but considering individual element abundances allows us to go beyond this assumption in the ISM, exploiting the individual traced elements. Another advantage is that as the local abundance of a depleted element approaches zero, τ_{acc} approaches infinity. This allows us to transfer mass in depleted elements from the gas phase to the dust phase consistently, in the fixed proportions dictated by their chemical compounds (Table 2). Equation (4) is applied in the context of our dust model, to the two overabundant dust species (silicate and carbon). For carbon grains, this is straightforward; with only one, homonuclear, subspecies of grain the bottleneck abundance is always C, i.e. $\epsilon_j = \epsilon_{\text{C}}$. For silicates, comprising O, Mg, Si, and Fe, this is found by maximizing the $\epsilon_{j,\odot}/\epsilon_j$ term,

$$\frac{\epsilon_{j,\odot}}{\epsilon_j} = \max \left(\frac{\epsilon_{\text{O},\odot}}{\epsilon_{\text{O}}}, \frac{\epsilon_{\text{Si},\odot}}{\epsilon_{\text{Si}}}, \frac{\epsilon_{\text{Mg+Fe},\odot}}{\epsilon_{\text{Mg+Fe}}} \right), \quad (5)$$

such that j represents the bottleneck abundance, where $\text{Mg} + \text{Fe}$ is the composite abundance of the two silicate subspecies endmember elements Mg and Fe, i.e.

$$\frac{\epsilon_{\text{Mg+Fe},\odot}}{\epsilon_{\text{Mg+Fe}}} = \frac{\epsilon_{\text{Mg},\odot} + \epsilon_{\text{Fe},\odot}}{\epsilon_{\text{Mg}} + \epsilon_{\text{Fe}}}. \quad (6)$$

This is reflective of the fact that total accretion in the silicate species is freely divisible between Mg and Fe endmembers, such that the total abundance of Mg and Fe can be considered a single reservoir of gas-phase metals available for accretion. With O and Si always present in silicates with a fixed abundance ratio, these elements can be treated as independent reservoirs.

This accretion for the overall silicate species is computed and the total accreted mass is divided between endmember species proportionally to the relative abundance ratio of gas-phase endmember elements in the ISM.

Accretion proceeds most efficiently in cold, dense environments where dust constituent elements are available in the gas-phase. Accounting for how dust is distributed, the majority of accretion is posited to take place in the cold, neutral medium (e.g. E. B. Jenkins 2009; H. Hirashita 2012). This represents local densities of $n_{\text{H}} \gtrsim 10 \text{ cm}^{-3}$ (H. Hirashita 2000), where n_{H} is the hydrogen number density where $n_{\text{H}} = m_{\text{p}}^{-1} X_{\text{g}} \rho$ for proton mass m_{p} and hydrogen mass fraction, X_{g} . Efficient accretion in these contexts is supported by the enhanced depletion observed in dense clouds (B. D. Savage & K. R. Sembach 1996). For this reason, to properly account for accretion it is important that dense, molecular media are represented. S. Zhukovska et al. (2008) attribute smaller dust nucleation rates compared to E. Dwek (1998) from stellar sources to a more explicit treatment of dust growth in molecular clouds. A conflation of stellar nucleation and accretion could be one explanation for the significantly higher dust yields displayed for E. Dwek (1998) in Fig. 1, though higher dust yields now observed in individual CCSNe may support the higher level grain seeding paradigm (M. Shahbandeh et al. 2023) comparable to E. Dwek (1998).

Dust yields that implicitly include some molecular cloud grain growth can be useful for simulations that do not attempt to model such dense gas structures (R. McKinnon et al. 2017; R. Davé et al. 2019; A. P. Vijayan et al. 2019). However, our model has been developed for galaxy formation simulations that include multiphase gas and that can attempt (at sufficient resolution) to distinguish molecular gas clouds from the diffuse ISM, such that using yields that implicitly include an accreted component could effectively ‘double count’ accretion. Despite the relatively high densities reached in our simulations, we are still far from resolving the dense cores of molecular clouds where these grains are thought to accrete efficiently. As a result we apply a clumping factor, C , to boost the densities input into the accretion rate. Our fiducial scheme uses a variable boost factor that rises monotonically from $C = 1$ to 100 between $-1 < \log_{10} n_{\text{H}}/\text{cm}^{-3} < 2$ and is clipped to the boundary values on either side. This is found to reproduce observed dust-to-metal ratios within gas at ISM densities ($n_{\text{H}} \gtrsim 0.1 \text{ cm}^{-3}$). This provides a means to calibrate across regimes in resolution. We investigate the clumping factor further in Section 3.8. For our fiducial clumping with representative properties of the cold, neutral medium ($T = 100 \text{ K}$, $n_{\text{H}} = 10 \text{ cm}^{-3}$) at $0.1 Z_{\odot}$ a small carbon (silicate) grain in an environment where 5 per cent of the bottleneck element is available for depletion in the gas-phase gives $\tau_{\text{acc}} = 26.4 \text{ Myr}$ ($\tau_{\text{acc}} = 14.6 \text{ Myr}$), of comparable order to the ‘enhanced’ accretion model found necessary by C. R. Choban et al. (2025).

3.4 Grain destruction

Just as grains may grow and accrete material in the ISM, grains may also be shrunken and destroyed through their interaction with local media. We employ a few distinct mechanisms,¹¹ detailed below. For computed destruction time-scales, the fractional mass change of a given grain is computed for each particle over a time-step of length Δt as $\exp(-\Delta t/\tau)$.

3.4.1 Thermal sputtering

Thermal sputtering is the erosion of dust grains through high-velocity interactions with gas particles in hot gas. We use the thermal sputtering prescription of J. C. Tsai & W. G. Mathews (1995), which uses the grain shrinkage rate

$$\frac{dr_{\text{grain}}}{dt} = -3.2 \times 10^{-18} \text{ cm s}^{-1} \left(\frac{n_{\text{H}}}{1 \text{ cm}^{-3}} \right) \left[1 + \left(\frac{T}{T_0} \right)^{-2.5} \right]^{-1}, \quad (7)$$

where the time-scale for sputtering the dust mass for uniform spherical grains can then be computed as

$$\tau_{\text{sp}} = 0.85 \text{ Myr} \left(\frac{r_{\text{grain}}}{0.1 \mu\text{m}} \right) \left(\frac{n_{\text{H}}}{1 \text{ cm}^{-3}} \right)^{-1} \left[1 + \left(\frac{T}{T_0} \right)^{-2.5} \right], \quad (8)$$

where T_0 is the temperature normalization, set to 2×10^6 K. We note that the sputtering time-scale is highly sensitive to the temperature, such that dust in gas heated to temperatures $T > T_0$ will be destroyed very quickly at ISM densities. In hot gas at the lowest CGM densities the time-scale can become long again. Note that because thermal sputtering occurs in the diffuse-gas phase, we use the unmodulated density, n_{H} , for sputtering (as opposed to n'_{H} , modulated by a clumping factor, C).

3.4.2 Astration

Another process to consider is the destruction of grains upon their reincorporation into stellar atmospheres, also known as *astration*. In our model, the dust mass that is present in a gas particle as it is being converted to a stellar population is transferred back to the *gas-phase* element mass fraction array stored in each particle, in the proportions dictated by each grain compound. These gas-phase elemental mass fractions then come to represent the *total* elemental mass fractions for the resultant star particle.

3.4.3 SN shocks

The effect of SN shocks on grains may be considered a complement to the thermal sputtering we implement to account for resolved heating processes. The shocks of individual SNe are not resolved in galaxy-scale simulations. Instead, in COLIBRE bulk feedback and star formation events are modelled stochastically (Sections 2.1 and 2.2). For consistency, we implement a stochastic approach linked directly to the feedback module – dust is totally destroyed in the particles stochastically selected to receive SN feedback.¹² This is handled in the same way as astration – on being hit, mass in dust is redistributed into the gas-phase of the constituent elements, and dust abundances are set to zero, all within the struck particle. The practical,

¹¹We opt for a minimal-necessary model. For our grain population, such as photo-evaporation of grains, are neglected as by theory minor contributors to grain destruction (e.g. A. Nanni et al. 2024).

¹²For either thermal or kinetic implementations of SN feedback.

stochastic implementation of this approach paired with the complex varying properties of the gas receiving feedback make it difficult to compute a time-scale from first principles. To estimate the time-scale for dust destruction by SNe in our simulations, we compare the average feedback-hit particles in the last 100 Myr (before $z = 0$) with the $z = 0$ dust properties in our fiducial model (see Section 2.5), the ratio of which provides a cosmic time-scale of around 420 Myr.

3.5 Grain size evolution

Other processes have been identified that conserve the mass in dust, but modify the grain size distribution. Despite not affecting the grain mass directly, this has a knock-on effect on the evolution of grains, with the accretion and sputtering time-scales being sensitive to grain size, with small grains exhibiting shorter time-scales due to their higher surface to volume ratios.

3.5.1 Shattering

Mass can transfer from the large to small grain component via *shattering*, where energetic grain–grain collisions break their constituents into many smaller grains. For shattering, we follow the formula of S. Aoyama et al. (2017), adapted by G. L. Granato et al. (2021),

$$\tau_{\text{sh}} = 54.1 \text{ Myr} \left(\frac{DTG}{0.01} \right)^{-1} \left(\frac{a_L}{0.1 \mu\text{m}} \right) \times \begin{cases} \left(\frac{n_{\text{H}}}{1 \text{ cm}^{-3}} \right)^{-1} & \text{for } \frac{n_{\text{H}}}{1 \text{ cm}^{-3}} < 1, \\ \left(\frac{n_{\text{H}}}{1 \text{ cm}^{-3}} \right)^{-\frac{1}{3}} & \text{for } \frac{n_{\text{H}}}{1 \text{ cm}^{-3}} \geq 1, \end{cases} \quad (9)$$

where the time-scale normalization value assumes a material density of 3 g cm^{-3} (intermediate between typical silicate and carbon grains), and a_L is the radius of large grains ($0.1 \mu\text{m}$ in our model). The G. L. Granato et al. (2021) prescription differs from previous applications in that the shattering does not shut off abruptly above a gas density of $n_{\text{H}} = 1 \text{ cm}^{-3}$, but instead increases more gradually above this density. Note, as in sputtering, we are using the unmodulated density, n_{H} (as opposed to n'_{H}), as shattering is a diffuse-gas process.

3.5.2 Coagulation

The collisions of grains can conversely lead to mass transfer from small to large sizes, as the constituents of lower velocity grain–grain collisions can *coagulate* to form a single larger grain. To represent these processes, we again follow the prescription of S. Aoyama et al. (2017), modified to,

$$\tau_{\text{co}} = \tau_{\text{co},0} \left(\frac{a_s}{0.1 \mu\text{m}} \right) \left(\frac{DTG}{0.01} \right)^{-1} \left(\frac{v_{\text{co}}}{0.1 \text{ km s}^{-1}} \right)^{-1} \times \left(\frac{n'_{\text{H}}}{10 \text{ cm}^{-3}} \right)^{-1}, \quad (10)$$

where the time-scale normalization, $\tau_{\text{co},0}$ is taken to be 27.1 Myr and incorporates an assumed material density of 3 g cm^{-3} for grains, and a_s is the radius of small grains ($0.01 \mu\text{m}$). We also assume a fixed coagulation velocity of $v_{\text{co}} = 0.2 \text{ km s}^{-1}$. In the original prescription of S. Aoyama et al. (2017) this process is assumed to take place in molecular clouds, with a fixed density of 10^3 cm^{-3} . We adapt the S. Aoyama et al. (2017) equation to depend on an effective density, n'_{H} , i.e. the local simulated gas density boosted by a clumping factor to represent unresolved structure (see Section 3.8), as with the other collisional processes assumed here. Equation (10) makes use of the higher densities afforded by the higher resolution and cold gas

physics of our simulations, and does not depend on a subresolution criterion for molecular gas content. This allows us to maintain a self-consistent approach across the dust evolution processes. This is particularly important when dust is coupled to cooling and molecule formation processes (Section 3.7).

Despite our higher resolution and cold gas physics, we cannot resolve the interiors of molecular clouds, as addressed in Section 3.3. The coagulation rates are therefore also subject to a variable clumping factor, explained further in Section 3.8. In addition to the clumping factor, we note that the coagulation time-scale $\tau_{\text{co},0}$ is relatively uncertain, so calibration of $\tau_{\text{co},0}$ to match second-order dust properties, such as the size distribution of grains, is a reasonable approach. In particular, a lower $\tau_{\text{co},0}$ could be argued for, particularly due to the assumption of spherical grains; more prolate/oblate grains could increase grain–grain collision rates, and increase the growth of grains at fixed density. We will consider the option of boosting this coagulation rate in future work.

3.6 Grain transport and diffusion

In addition to the processes that govern how dust forms and evolves in each gas particle, we also model how dust mixes through the gas. This mixing occurs across multiple scales; from large-scale gas flows to the fine, fractal structure of turbulent motions. At scales we can resolve, dust mixing is naturally handled by the hydrodynamics: gas motions carry dust out of enriched regions and redistribute it throughout galaxies. For smaller, unresolved scales, we use an explicit diffusion model, that follows the same implementation to the one used for chemical elements (C. Correa et al., in preparation). Our approach of dust mixing adapts the diffusion equation, $dX_i/dt = \rho^{-1}\nabla(D\nabla X_i)$, into the SPH framework (J. J. Monaghan, H. E. Huppert & M. G. Worster 2005). Here, X_i is the dust fraction of gas particle i and D is the diffusion coefficient. The coefficient D depends on the local gas velocity shear, following a Kolmogorov turbulence model (J. Smagorinsky 1963), to represent turbulent mixing. Its normalization is set by a free parameter, C_d . The value of C_d varies significantly in prior studies from values $C_d \approx 0.003$, representing a minimal subgrid diffusion level that limits extreme abundances in forming stars (I. Escala et al. 2018), to $C_d \sim 0.1$ as suggested by Kolmogorov theory directly (e.g. J. Smagorinsky 1963; S. Shen, J. Wadsley & G. Stinson 2010; K.-Y. Su et al. 2018).

Diffusion of individual elements and dust are treated independently. For simplicity, we assume no relative boost in dust diffusion rates with respect to the diffusion of metals. We adopt the fiducial value of metal diffusion coefficient $C_d = 0.01$ (C. Correa et al., in preparation), with the exception of variations explicitly varying C_d , finding that a higher relative dust diffusion rate has only a small influence on dust evolution.

Turbulent diffusion helps to mitigate sampling noise in abundances of dust brought about by the discrete enrichment events. Particularly at low dust abundances, this allows gas that has not been directly enriched by star particles to accrue some dust and begin dust evolution.

We do not account for other kinematic effects of dust on gas, for instance through radiation pressure or the effect of gas-dust drag. Dust drag can be an important process at high densities ($n_{\text{H}} \gg 10^3 \text{ cm}^{-3}$), particularly those seen at circumstellar scales and protoplanetary discs (S. J. Weidenschilling 1977; T. Birnstiel, M. Fang & A. Johansen 2016). Dust drag may also be important internally for dense, molecular clouds (e.g. P. Padoan et al. 2006). However, for the density scales that we resolve ($n_{\text{H}} \lesssim 10^3 \text{ cm}^{-3}$), the drag time-scale considered here is long enough ($\tau_{\text{drag}} \gg 1 \text{ Myr}$) that

the stopping times of grains are long with respect to the gas dynamical times, and we may assume dust motions trace the gas motions well (B. T. Draine 2011; H. Hirashita & N. V. Voshchinnikov 2013).

3.7 Coupling to COLIBRE hybrid cooling and heating

The evolutionary life-cycle model of dust described above emphasizes the influence of stars and gas in the ISM. Given this model, we can in turn account for the influence of local dust properties on various heating and cooling processes, such that dust may in turn influence gas physics generally and even star formation. In Section 2.4 we detail the relevant processes influenced by live dust grains. The relevant quantity for cooling processes linked to grains is the local dust cross-section per unit volume, $n_d\sigma_d$. The hybrid cooling module assumes a constant dust mix internally, through an ‘effective size’ translating to an A_V/N_{H} value (see Section 2.4). Having fixed attenuation curves parametrized solely by their normalization in the V -band, A_V , means it is less complex and computationally expensive to compute heating and cooling rates. With mass-weighted average grain sizes of the MW ($a \sim 0.1 \mu\text{m}$, see e.g. B. T. Draine 2003) consistent with the effective size assumed for our OneSize model (Table 1, taken as our large-grain size in Fiducial and elsewhere) it is straightforward to simply scale extinctions and collisional cooling processes particle-by-particle using our modelled dust-to-gas ratio $\mathcal{DTG}_{\text{model}}$. This is done via a rate scaling factor s_{dust} where $s_{\text{dust}} = \mathcal{DTG}_{\text{model}}/\mathcal{DTG}_{\text{MW}}$, where we take $\mathcal{DTG}_{\text{MW}} = 6.6 \times 10^{-3}$ (following the ISM grain model in Cloudy, M. Chatzikos et al. 2023).

By contrast, our Fiducial dust model uses two grain size bins to capture local variations in the size distribution. At a fixed dust mass and material density, dust-associated heating and cooling rates scale with the surface-to-volume ratio, which scales in inverse proportion with the grain radius ($\propto a^{-1}$). Accounting for the two sizes gives a new expression for s_{dust} :

$$s_{\text{dust}} = \frac{a_{\text{MW}} \left(\frac{1}{a_{\text{L}}} D_{\text{L}} + \frac{1}{a_{\text{S}}} D_{\text{S}} \right)}{\mathcal{DTG}_{\text{MW}}}, \quad (11)$$

where D_{L} and D_{S} are the fraction of a particle’s total mass in small and large grains, respectively, and a_{MW} is the effective size of dust grains in the Milky Way.

We note that for collisional processes (e.g. H_2 formation) the s_{dust} scaling simply reflects the increased surface area per unit mass contributed by small grains. For heating and cooling processes, however, this scaling may strictly depend upon an additional extinction efficiency, $\langle Q_{\text{ext}} \rangle$, i.e. the ratio of the extinguishing cross section (the sum of scattering and absorption cross sections) to the physical cross section of grains (e.g. J. C. Weingartner & B. T. Draine 2001). This can vary with wavelength and differs depending on the material and size properties of grains. While V -band attenuation A_V is used as a standard normalization for extinction curves, the processes relevant to the dust coupling described here (e.g. self-shielding, photoelectric heating) actually depend on the FUV, ionizing or Lyman-Werner photons. While extinction efficiencies of large grains, $\langle Q_{\text{ext},L} \rangle$, dominate at optical wavelengths (e.g. B. T. Draine & A. Li 2001), in the UV the small-grain extinction efficiency, $\langle Q_{\text{ext},S} \rangle$, becomes comparable. For simplicity, we neglect the influence of $\langle Q_{\text{ext}} \rangle$, assuming $\langle Q_{\text{ext},S} \rangle / \langle Q_{\text{ext},L} \rangle \approx 1$ and marginal differences between material species. We look into detailed extinction and attenuation properties of grain in more detail in upcoming studies post-processing COLIBRE galaxies using radiative transfer, and justify our assumptions around extinction efficiency further in Appendix B.

The live dust model may also impact radiative cooling and heating processes more indirectly, through the different depletion of gas-phase elements, affecting gas-phase cooling channels (such as metal-line cooling and heating). The element-by-element cooling channels for gas-phase metals in the hybrid cooling scheme accounts for the assumed P25 dust depletion, via a MW-like relative depletion pattern (E. B. Jenkins 2009) that scales with Z_{gas} (see Section 3.9 for a comparison of equilibrium and live dust depletions). The fractions of metal elements in the gas- and dust-phase of enriched gas are handled implicitly, scaling the tabulated rates and leading to different element-by-element cooling and heating between the coupled and uncoupled cases for gas of a fixed thermodynamic state and metallicity, owing to the differences between our live dust model and the P25 dust.

In Fig. 2 we show differences between dust, cooling and gas properties, comparing our live dust with the instantaneous P25 dust, both within the `Fiducial` run, and when incorporating the coupled physics effects when comparing to `FidUncoupled` runs. In the top row we see the influence of the dust properties input into the heating and cooling modules alone, by comparing the P25 dust to the live dust model within `Fiducial`. This shows how, for like physical conditions, the influence of dust on cooling processes differs across $n_{\text{H}}-T$ space when live dust is included. The shading in this figure shows absolute pixel differences normalized by the maximum positive and negative difference pixel in each panel. This nuanced representation allows us to show where bulk properties differ in phase space. The overall distribution of dust depletion fractions in phase space can be seen more directly in upcoming Section 4.2 with overall phase-space distributions for COLIBRE presented in J. Schaye et al. (2025).

In Fig. 2, we see that for relatively cool ($\log_{10} T/\text{K} \lesssim 3.5$) and moderately dense ($\log_{10} n_{\text{H}}/\text{cm}^{-3} > -2$) gas, the P25 model tends to yield lower DTG ratios (panel a), though with small, exceptional regions (around $\log_{10} n_{\text{H}}/\text{cm}^{-3} \approx -0.5$, $\log_{10} T/\text{K} \approx 2.5$ and $\log_{10} n_{\text{H}}/\text{cm}^{-3} \gtrsim 2$). For hotter or more diffuse gas, the live dust model dominates, particularly due to the propagation of dust outside of the ISM and into the CGM and other phases that is assumed to be completely destroyed in the P25 model. This residual dust that avoids destruction by sputtering and direct annihilation in shocks is seen within H II regions (horizontal sequence at higher density and $\log_{10} T/\text{K} \approx 4$) and in the diffuse hot gas, where despite the high temperatures, dust destruction time-scales can remain long enough due to the lower densities. We highlight the few cells in hot, denser gas showing evidence of live dust (around $\log_{10} T/\text{K} \approx 7$, $\log_{10} n_{\text{H}}/\text{cm}^{-3} \approx 0$). These few particles are those caught in the simulation output *between* the feedback and dust evolution steps (where this dust is rapidly destroyed), due to the order of operations in COLIBRE. The total dust content of the `Fiducial` volume when using P25 or live dust is very similar (within 0.01 dex or ≈ 2 per cent), though this level of consistency is coincidental, varying more in other volumes and environments. Recomputing the instantaneous P25 elemental cooling and heating rates of `Fiducial` using P25 dust (panels b and c, respectively) yields complex differences, particularly higher metal-line cooling and heating rates in H II regions (due to more gas-phase metals) and heating in the densest ISM. However, in overall cosmic terms, the difference in cooling and heating rates per unit density ($\Lambda_{\text{cool}}/n_{\text{H}}$, $\Lambda_{\text{heat}}/n_{\text{H}}$) are marginal ($\lesssim 0.2$ per cent).

The bottom row of Fig. 2 compares the live dust of `Fiducial` to the P25 dust of the `FidUncoupled` run, including the physical differences and run–run variations. To illustrate this, we first show the different gas distributions (panel d). Along with the variations in hot, diffuse gas owing to differing and stochastic feedback histories of each run, there are also differences in the ISM; most notably the existence of an excess of colder gas in the `FidUncoupled`

run at $\log_{10} T/\text{K} \lesssim 100$. Taking this into account, we see that the local grain cross sections are generally higher in the live dust model, and by 0.53 dex ($\approx 3\times$) overall, excepting this cold sequence and a region of the warm-neutral medium (WNM), at $\log_{10} n_{\text{H}}/\text{cm}^{-3} \approx -1$, $\log_{10} T/\text{K} \approx 3.25$. This is attributable to our variable size distributions and abundance of small grains, particularly outside the coldest, densest gas. Another significant difference to highlight is that H_2 extends to slightly warmer temperatures and lower densities in the coupled runs relative to the uncoupled ones (panel f). We attribute this to lower cooling rates in the dense ISM, paired with more small grains at intermediate density, boosting the local grain cross-section for H_2 formation. Ultimately, the distribution of H_2 in $\rho-T$ space leads to higher cosmic H_2 mass densities (by 0.17 dex), boosting galactic H_2 mass fractions. Altogether, this emphasizes the complex differences that emerge from using a live dust model,¹³ where dust evolution is influencing gas thermodynamics throughout cosmic time.

3.8 Clumping factors

As with gas physics in the ISM, dust evolution is a highly multiscale process, ranging from dense molecular clouds to the diffuse CGM. In particular, the dense cores of molecular clouds provide sites for the growth of grains through frequent collisions in relatively cold gas (C. W. Ormel et al. 2009). The core densities of molecular clouds typically range from $n_{\text{H}} = 10^4 \text{ cm}^{-3}$ up to protostellar densities. The spatial scales associated with the densest molecular gas is below what we resolve in our cosmological volume simulations. Inability to represent these denser structures skews the densities experienced by gas to lower values and inhibits the accelerated dust evolution inferred for observed molecular clouds. To represent the presence and structure of unresolved collapsed clouds, we experiment with a ‘clumping factor’, C . This factor can boost input densities such that the effective gas densities (effective hydrogen number density), ρ' (n'_{H}), used by dense-gas processes (such as accretion and coagulation) are related to the physical densities in our simulations, as

$$\rho' = C\rho, \quad (12)$$

or

$$n'_{\text{H}} = Cn_{\text{H}}, \quad (13)$$

where ρ (or n_{H}) is the input physical gas density (or hydrogen number density) and $C > 1$. With the C factor tied to the resolution and small-scale physics limitations of the simulations, we treat C as a free parameter of the model, which must be calibrated. We experiment with both a constant clumping factor as well as a density-dependent clumping factor, $C(\rho_{\text{phys}})$. In the density dependent case

$$C(n_{\text{H}}) = \begin{cases} 1, & \text{if } n_{\text{H}} \leq n_{\text{H}, \text{min}} \\ \left(\frac{n_{\text{H}}}{n_{\text{H}, \text{min}}}\right)^m, & \text{if } n_{\text{H}, \text{min}} < n_{\text{H}} \leq n_{\text{H}, \text{max}} \\ C_{\text{max}}, & n_{\text{H}} > n_{\text{H}, \text{min}} \end{cases}, \quad (14)$$

where $n_{\text{H}, \text{min}}$ and $n_{\text{H}, \text{max}}$ are the minimum and maximum densities of the clumping transition, C_{max} is the maximum clumping value (see Table 1), and $m = \log_{10}(C_{\text{max}})/\log_{10}(n_{\text{H}, \text{max}}/n_{\text{H}, \text{min}})$. For a gradual, transition we take $n_{\text{H}, \text{min}} = 0.1 \text{ cm}^{-3}$ and $n_{\text{H}, \text{max}} = 100 \text{ cm}^{-3}$, as well as a fiducial $C_{\text{max}} = 100$ (yielding $m = 2/3$).

The runs varying the subgrid clumping factor, C (see Table 1), are plotted together in Fig. 3, showing properties as a function of n_{H} . Here, the top panel indicates C for each variation (coloured lines, left

¹³Alongside any stochastic run-to-run variation (J. Borrow et al. 2023).

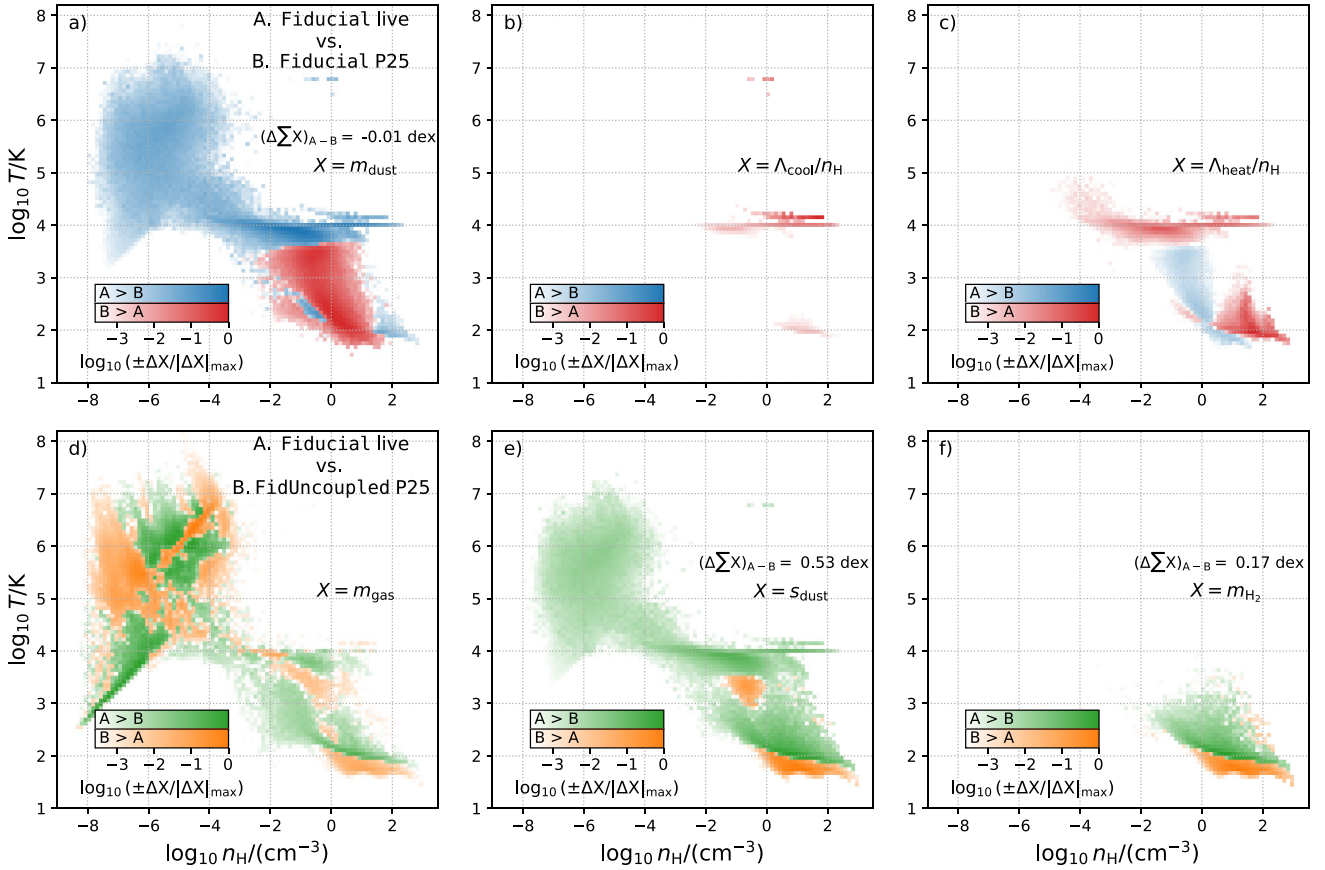


Figure 2. Differential density–temperature ($n_{\text{H}}\text{-}T$ or ‘phase’) diagrams for particles at $z = 0$, comparing intra- and intersimulation gas and dust properties. The *top row* compares properties within the Fiducial run; comparing our live dust model, to values associated with the instantaneous, P25 dust (computed by interpolating the P25 tables). Comparing within the same simulation is intended to isolate differences, given an otherwise identical state of the gas. From left to right we compare total dust content, and elemental cooling and heating rates given these different depletions. The *bottom row* compares between the Fiducial and FidUncoupled runs, considering live and P25 dust for each, respectively. From left to right, this compares the gas mass distribution, the local grain collisional cross-section, and the H_2 mass distribution. For the top row, the absolute difference cell-by-cell (for 100×100 cells) between the live (A) versus P25 dust (B) properties within Fiducial are taken, and shaded blue or red depending on whether there is an excess in the former or latter case, respectively, revealing where significant differences exist. For the bottom row comparing Fiducial live (A) versus FidUncoupled P25 (B) dust, we instead use orange and green. Where significant, the difference in the quantity summed over all particles is written, where positive values indicate an excess in the Fiducial live dust case. In particular we see higher grain cross-sections (by 0.53 dex) and H_2 masses (by 0.17 dex) in the Fiducial run using live dust, when compared to the FidUncoupled using P25 dust.

axis), with H species fractions for the Fiducial run (grey lines, right axis). The bottom panel then shows DTZ (solid lines, left axis) and the ratio of mass in small versus large grains. To indicate how different runs transition from low DTZ at lower densities to the saturated regime in denser gas, we compute the density above which half of the dust in the entire cosmological volume is enclosed, $\rho_{0.5, \text{dust}}$. These transition densities are indicated in the top panel of Fig. 3 using downward arrow markers. We can also see how varying C affects gas properties; the density above which more than half of the gas is in the molecular phase, $\rho_{0.5, \text{H}_2}$, is model-dependent. For each run, $\rho_{0.5, \text{H}_2}$ is indicated using vertical line marks.

For the noC and maxC10 runs, the clumping factor is lower than or equal to that of the Fiducial run at all densities. We see that for maxC10, the transition to higher DTZ is pushed to higher densities by ≈ 1 dex, while grain sizes are systematically larger for $\log_{10}(n_{\text{H}}/\text{cm}^{-3}) \lesssim 1.5$ (bottom panel). We note that the maxC10 model does not reach the saturated regime, with an approximately constant DTZ ratio of ≈ 40 per cent at high ISM densities, as seen in Fiducial. For noC, we see a more extreme

shift in the DTZ transition to higher densities. Interestingly, for $\log_{10}(n_{\text{H}}/\text{cm}^{-3}) \gtrsim 1.5$ the noC model yields smaller dust grains (dashed lines in the bottom panel, right axis). These trends in the grain size distribution are indicative of the lasting imprint of the grain sizes assumed for seed grains in low density gas where dust evolution proceeds slowly, i.e. reflective of seed grains having 90 per cent in large grains by mass (Section 3.2.4), followed by efficient shattering at intermediate densities and coagulation in dense gas, balanced with the higher rates of gas-phase accretion per unit mass for smaller grains.

In terms of $\rho_{0.5, \text{H}_2}$ (top panel), we see that relative to Fiducial, maxC10 is pushed to densities ≈ 0.6 dex higher, while noC is pushed a further ≈ 0.6 dex higher, indicative of the lower rate of molecule formation in less dusty ISM gas for a given density. The differences in $\rho_{0.5, \text{dust}}$ are less intuitive; the maxC10 and noC values are both ≈ 0.2 dex higher than fiducial. The lack of difference between the two variations could be a balance of a more concentrated DTZ with a higher proportion of dust contributed by grain seeding and pushed outside the ISM.

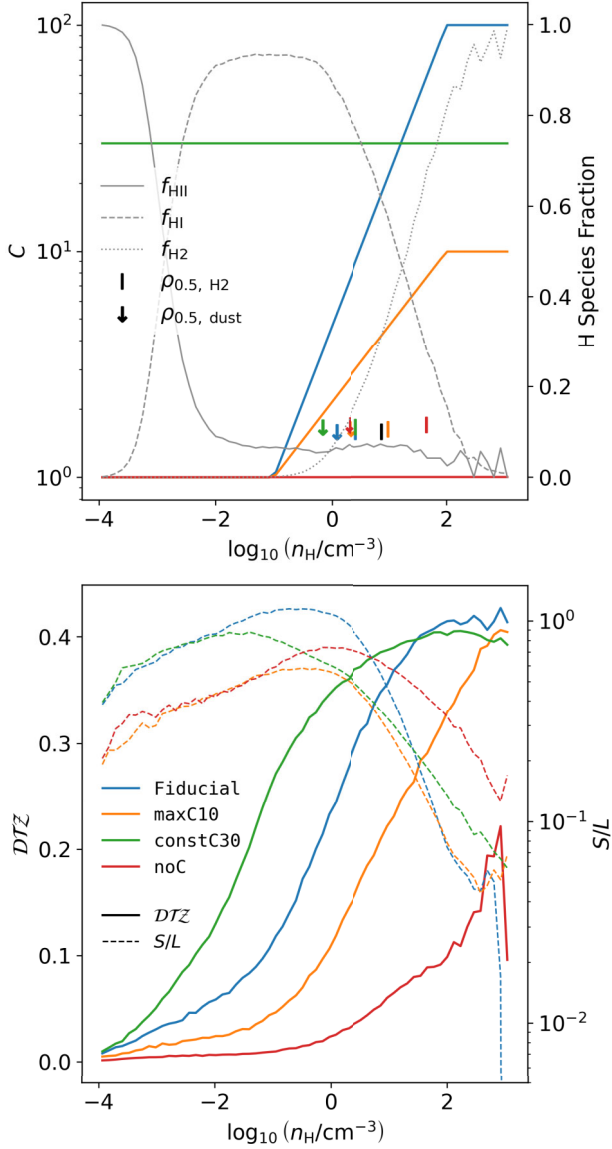


Figure 3. Density-dependent properties of dust models assuming differing subgrid clumping factors. The top panel shows the clumping factor C as a function of n_{H} for the *Fiducial* run, alongside a number of variations (thick coloured lines, left y-axis). For reference the hydrogen species fractions of the *FidUncoupled* run are co-plotted. The dust (H_2) transitional densities (densities enclosing half the cosmic dust mass), are plotted as downward arrow marks (vertical line marks). Markers are given slight vertical offsets for distinguishability. The bottom panel shows properties of dust grains as a function of gas density, displaying the total DTZ (solid lines, left y-axis) and small-to-large grain mass ratio (S/L , dotted lines, right y-axis) in each bin. We see strong variations in the dust properties associated with different treatments of C , but also convergence at the highest densities for all but the *noC* run.

The *constC30* model has a lower clumping factor than *fiducial* for $\log_{10}(n_{\text{H}}/\text{cm}^{-3}) \gtrsim 1$, but stronger clumping at lower densities. In the bottom panel we see how dust extends to significantly lower values in the constant clumping factor case, due to boosted growth at lower densities. The ratio of mass in the small grain bin to that in the large grain bin, S/L , is lower (i.e. larger grains) at intermediate densities, but falls more slowly towards high density as rates of




Figure 4. Elemental depletion and total metal depletion of dust constituents for our model. We compare the *Fiducial* run (\times markers), using graphite and olivine grain chemical species, to the grain chemistry variation *SilPyro* ($+$ markers), with graphite and pyroxene chemistry. For comparison, we plot the Milky Way ISM values as *empty circles* (E. B. Jenkins 2009, corrected for our assumed $Z_{\odot} = 0.0134$), as well as values for a factor 2 reduction in δ_{C} (U. J. Sofia et al. 2011; V. S. Parvathi et al. 2012). C is depleted into the homonuclear graphite grains, while O, Mg, Si, and Fe are depleted into the heteronuclear silicates. For visibility, the $\delta_{\text{Fe, MW}}$ is plotted as an upper limit, given the value of -2.2 . We also plot the total metal depletion δ_{Z} . We see that the *Fiducial* run shows best MW-depletion agreement compared to *PyroSil*, with the exception of Si, which is comparable.

coagulation rise less dramatically without an increasing C value. We see that the *Fiducial* run yields approximately the same $\rho_{0.5, \text{H}_2}$ as the constant C model, but a lower $\rho_{0.5, \text{dust}}$.

3.9 Depletion

In our development of the dust model, we prioritized total dust content (in terms of absolute masses, DTG and DTZ ratios), but the depletion of individual elements provides more detailed information about the composition of the dust. For element X, depletion is the ratio between gas-phase and total (gas + dust) abundance δ_{X} , i.e.

$$\delta_{\text{X}} = \frac{[\text{X}/\text{H}]_{\text{gas}} - [\text{X}/\text{H}]_{\text{total}}}{[\text{X}/\text{H}]_{\text{total}}}, \quad (15)$$

such that more negative values are indicative of larger fractions of X being depleted onto dust grains. Strong depletion of elements is key to limiting grain growth in the ISM of galaxies, and can contribute to a saturation in the DTZ ratio and thus a near-constant DTZ ratio measured in massive galaxies (e.g. T. Zafar & D. Watson 2013; A. Rémy-Ruyer et al. 2014; P. De Vis et al. 2019). While a secondary consideration to the total dust content, elemental depletion helps constrain our model.

We illustrate depletion patterns in Fig. 4 in order to compare how different elements bottleneck grain formation in different models, and how these compare to what we see in the dense ISM of the Milky Way (from E. B. Jenkins 2009, assuming the strong depletion case, $F_{\star} = 1$, corrected to our solar abundance as in S. Ploeckinger & J. Schaye 2020). We aggregate gas of $n_{\text{H}} > 50 \text{ cm}^{-3}$, to sample the depletion-limited regime as seen in the bottom panel of Fig. 3.

In this regime, the chosen grain chemistry and chemical enrichment set the dust content. We compare our `Fiducial` run to our silicate variant, `SilPyro`, which uses the *pyroxene* chemical group for silicate grains, as opposed to our fiducial *olivine*. We see that for total metal depletion (δ_Z , rightmost points) is slightly lower than the E. B. Jenkins (2009) observations for both models, but best fit by the grain chemistry choices of our `Fiducial` run.

3.9.1 Silicate grain depletion

In the case of heteronuclear silicate grains, the depletion pattern of constituent elements (O, Mg Si, and Fe) depends on the grain chemistry, and its interplay with the yields and gas-phase abundance of those elements in the ISM. As such, depletion patterns can provide motivation for our choice of effective grain molecule. We consider both the *pyroxene* and our fiducial *olivine* groups as candidate silicate molecules, with the form $(\text{Mg,Fe})_2\text{Si}_2\text{O}_6$ and $(\text{Mg,Fe})_2\text{SiO}_4$, respectively.

Comparing the δ_X values of Fig. 4 for silicate constituents shows that the *olivine* chemistry of the `Fiducial` model is closer to the observed MW depletions (E. B. Jenkins 2009) for all elements, excepting Si. We see that Si is overdepleted for *pyroxene* (`SilPyro`) grains. It appears this has the effect of limiting the depletion of the pyroxene Fe and Mg endmembers,¹⁴ leaving their depletion levels far below observed MW values (i.e. around 50 per cent in the gas-phase, as opposed to just 5 per cent).

While O is observed to be the least depleted silicate constituent, it can still contribute significantly to the silicate mass due to its higher abundance compared to the other constituent elements. We find that using higher *olivine* and *pyroxene* silicates both slightly underdepletes O, with the olivine molecule faring slightly better. This is intriguing; olivines represent the most oxygen-rich viable silicate candidates, so a pure olivine model still under-depleting O could be indicative of a potential tension between the metallicities and depletions, perhaps via the metallicity calibration examined further in Appendix C. This difficulty in explaining O depletions has been noted previously (e.g. D. C. B. Whittet 2010).

The two models differ qualitatively for Si; olivine silicate chemistry underdepletes Si, while it is overdepleted for pyroxenes. This is perhaps indicative of Si being the bottleneck element in the more silicon-rich pyroxene grains of the model, and explains the strong underdepletion of Mg and Fe in the pyroxene models.

It is notable that the MW Fe depletion is considerably stronger than we achieve in our models, with only around 0.6 per cent left in the gas-phase compared with 3 per cent for our `Fiducial` model. However, this reflects only a small change in the dust mass, which is the dominant cooling channel for colder gas (fig. 9 of S. Ploeckinger & J. Schaye 2020).

We note that laboratory studies comparing to observations in the X-ray can provide evidence for certain (olivine or pyroxene) silicates or some relative mix between the two (see e.g. S. T. Zeegers et al. 2019; D. Rogantini et al. 2020; I. Psaradaki et al. 2023). While a single silicate group is clearly a simplification, and tracking more grain chemistry species (e.g. including metal oxides, silicon carbide grains), may help to balance the depletion of elements, this would complicate our model¹⁵ and contribute to its memory overhead. We therefore use olivine as a viable, effective molecule that we find to

¹⁴*Ferrosilite* and *enstatite*, respectively.

¹⁵For example, having multiple grain types containing the same elements requires some decision as to how this depletion is split.

reasonably reproduce depletion through its fixed ratio of constituent elements is therefore our preference.

3.9.2 Graphite grain depletion

In the case of graphite, depletion is much simpler as it is a mononuclear grain. However, we found that when depletion of C is uninhibited, it depletes too readily, such that close to 100 per cent of carbon is locked in dust grains in the ISM. As gas-phase molecular C can be an important ISM coolant, this is both undesirable and unphysical.

In observed molecular clouds, the majority of gas-phase carbon is in CO (A. Fuente et al. 2019). The triple bond in CO is the strongest molecular bond and CO only solidifies at temperatures $T \gtrsim 70$ K. We take these properties together to mean that carbon in the CO reservoir, which our simulations do not resolve and which we do not track in our coupled simulations, is not readily available for depletion onto solid grains. We therefore limit carbon depletion by making some fraction of the gas-phase element unavailable for depletion. We choose a maximum dust depletion factor of 3 (66 per cent of carbon in dust), following the empirical study of local molecular clouds of A. Fuente et al. (2019). We note that this limit is only relevant for carbon, and oxygen depletion is instead bottle-necked by the availability of silicate constituent elements.

A similar need to limit accretion of carbonaceous dust, and appeal to a CO reservoir, is employed by C. R. Choban et al. (2022) for their simulations implemented in FIRE-2. C. R. Choban et al. (2022) also highlight that C depletion is particularly poorly constrained by observations, owing to the limited environments probed and scarcity of sightlines in E. B. Jenkins (2009), with the suggestion that the observed C depletion may be too high (U. J. Sofia et al. 2011; V. S. Parvathi et al. 2012). We find that our δ_C is indeed too low compared to E. B. Jenkins (2009), but incorporating this correction factor brings δ_C values into good agreement with our model. For simplicity, and given the limited resolution of our simulations, we use our empirical treatment which avoids tracking CO formation directly. However, we note that employing a CHIMES chemical network that includes C and O would allow us to track this self-consistently to limit the pure C accretion.

3.10 Resolution

While we focus on the m6 resolution in this work, with gas (DM) particles of $1.84 \times 10^6 M_\odot$ ($2.42 \times 10^6 M_\odot$), the dust model presented here is intended to be used across a range of resolutions, and as a single component of the multi-resolution COLIBRE simulation suite. In particular, we defer plots showing the convergence properties of dust (particularly the M_\star - M_{dust} relation and grain size ratio as a function of M_\star) to the presentation of the COLIBRE simulations as a whole by J. Schaye et al. (2025) in context with the convergence of other aspects of galaxies. With COLIBRE we target ‘*weak convergence*’ (e.g. J. Schaye et al. 2015), aiming to achieve a consistent reproduction of target data sets through the choice and calibration of key parameters. This calibration process happens in concert with the other modules that constitute COLIBRE, with feedback calibration detailed in E. Chaikin et al. (2025).

Elements of the modelling were designed to support a range of resolutions, for example by the implementation of a clumping factor (Section 3.8) intended to represent unresolved structures like molecular clouds and their dense cores at intermediate and low resolutions, as opposed to an explicit subgrid model of molecular clouds in simulation. The dust properties in COLIBRE are generally




Figure 5. Density–temperature ($n_{\text{H}}-T$ or ‘phase’) diagrams for binned gas particles in our simulations, with stepped shading of cells to illustrate comparative dust properties. *Left panel* shows the total \log_{10} dust-to-metal (DTZ) ratio for each $n_{\text{H}}-T$ bin in the Fiducial run. The *middle panel* shows the difference in $\log_{10} DTZ$ of the Fiducial run relative to SeedOnly, with the *right panel* comparing to NoDestruction in place of SeedOnly. We see that in our Fiducial run, DTZ is higher in relatively dense, cool ($T < 10^4\text{K}$) gas, boosted strongly above the DTZ in SeedOnly. Meanwhile, the Fiducial run shows a strong reduction in dust in relatively hot, diffuse gas relative to NoDestruction.

converged for galaxies resolved by $\gtrsim 1000$ star particles, or $\gtrsim 2 \times 10^9 M_{\odot}$ at m6 resolution. At the highest masses, $\gtrsim 10^{10.7} M_{\odot}$, convergence deteriorates somewhat owing to the role of AGN feedback calibration. Dust convergence properties are discussed in more detail in section 7.2.5 of J. Schaye et al. (2025).

4 RESULTS

Here, we present results from our simulations (Table 1), and compare to observations. In Section 4.1 we first present and discuss some resolved dust maps for some of our Fiducial run galaxies. In Section 4.2, we present phase diagrams shaded according to DTZ , isolating the role of different evolutionary processes in the proliferation of dust. We then consider the evolution of cosmic dust density with redshift in Section 4.3, comparing our Fiducial run to observations, alongside other variation runs isolating different effects. In Sections 4.4–4.7, we then compare galaxy dust relations from our cosmological volume simulations at $z = 0$ with low-redshift observational data, focussing on our Fiducial model.

4.1 Resolved properties

First, we briefly present a gallery of randomly oriented (i.e. box-projected) dust maps for a handful of galaxies taken from our Fiducial model in Fig. 6. These are the six highest stellar mass galaxies, defined within a 50 kpc spherical aperture, that also have > 10 per cent of their baryonic mass constituted by gas. For each galaxy, we map the highest dust surface density (Σ_{d}) pixels with colour shading, using a kernel density projection of the SPH gas particles, via the `swiftgalaxy`¹⁶ tool (K. A. Oman 2025). These are plotted alongside neutral and molecular hydrogen contours to show the relatively cold-dense gas regions. Colour shadings illustrate various properties of the dust population; the local DTZ ratio, the small-to-large grain ratio, the ratio of silicate to graphite/carbon grains and the ratio of forsterite (Mg-endmember) to fayalite (Fe-endmember) silicates.

The main qualitative feature we aim to get across here is the heterogeneity of these populations and their spatial configuration; we see diverse dust morphologies between galaxies, with distinct

discs, as well as an irregular and compact morphologies. We see some common features; for example, large-grain dominated dense, molecular regions, and small-grain dominated diffuse regions as well as higher depletion and more silicate-heavy central regions of galaxies. There are also distinct differences, for example overall DTZ variations, or dust chemistry differences. For example we see the clumpy and asymmetric property distributions of the first column galaxy, compared to the radial gradients and disc morphology of the last column galaxy.

4.2 Isolating evolutionary processes

In order to show the influence of different dust evolution processes, we compare runs where various processes are isolated. In particular, the SeedOnly run has no ISM dust evolution, only injecting grains from stellar channels. The NoDestruction run then turns on accretion and size transfer processes, leaving processes converting dust mass back into gas-phase metals off. These runs are compared to Fiducial in Fig. 5, where we show gas phase diagrams (T as a function of n_{H}), the left panel shading by total DTZ ratio, and centre and right panels shading by the difference in DTZ with respect to the Fiducial run of SeedOnly and NoDestruction, respectively.

We see an increasing gradient in the Fiducial run DTZ towards denser, cooler gas, reaching values of $DTZ \gtrsim 10$ per cent for $T \lesssim 10^4\text{K}$, $n_{\text{H}} \gtrsim 1\text{cm}^{-3}$. We see that this gradient is driven by accretion in the ISM; the DTZ is over an order of magnitude higher in the Fiducial run relative to SeedOnly (centre panel). We see that in the cool-dense gas, destruction processes have little influence over this, with the NoDestruction run showing similar DTZ to Fiducial (right panel).

By contrast, in the plume of hot, diffuse gas we see in Fig. 5 (i.e. $T > 10^4\text{K}$, $n_{\text{H}} < 10^{-2.5}\text{cm}^{-3}$) we see $DTZ \lesssim 0.01$, and of order the level directly injected from stellar channels (centre panel Fig. 5, comparing with Fig. 1). However, we see that dust destruction plays an important role here; including destruction effects reduces DTZ by more than an order of magnitude relative to runs with only accretion and size-transfer processes (Fiducial versus NoDestruction comparison, right panel). This shows dust destruction processes mitigate dust growth through accretion in the ISM, leading to higher DTZ in hot, diffuse gas.

¹⁶github.com/SWIFTSIM/swiftgalaxy

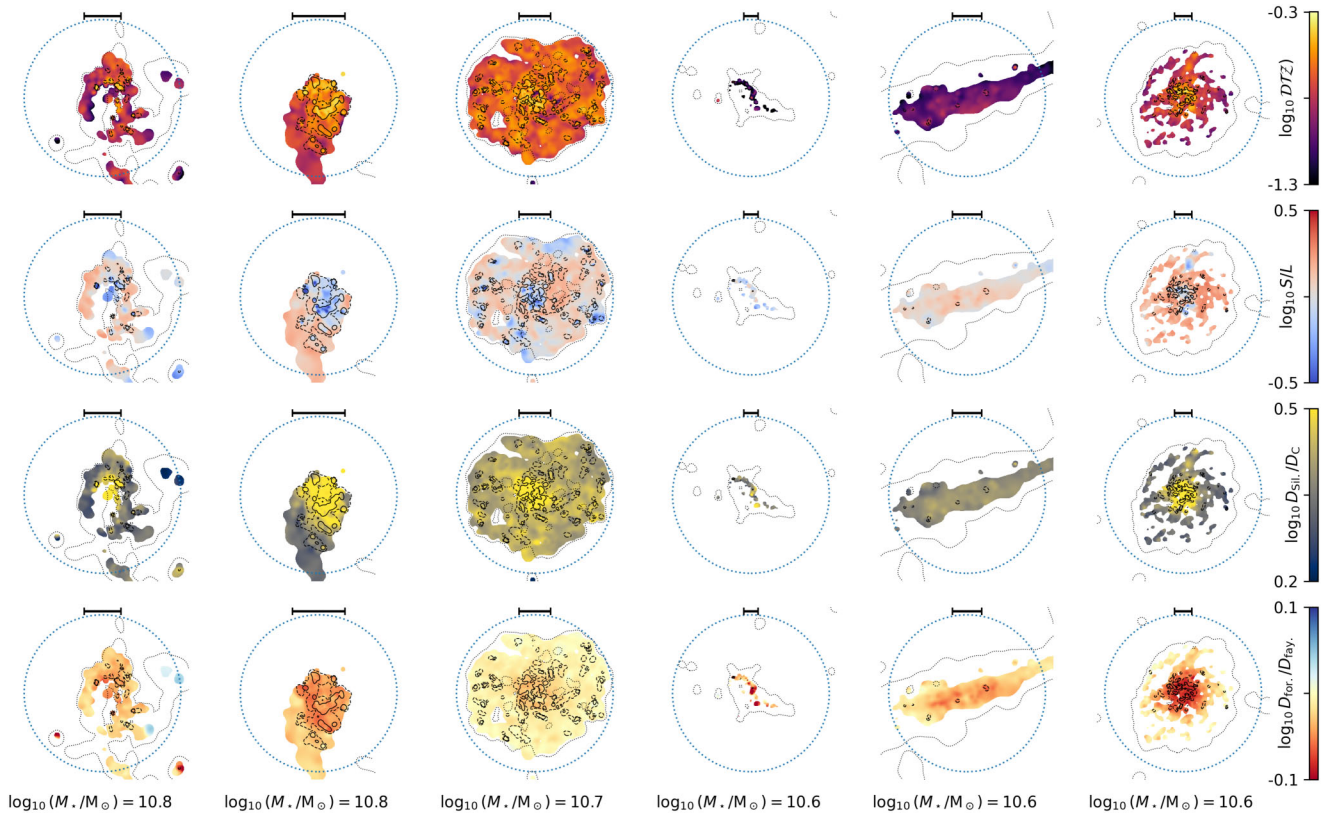


Figure 6. A gallery of the six most massive galaxies with more than 10 per cent of their baryons in gas from our Fiducial run, projected in box coordinates (i.e. randomly oriented), illustrating the heterogeneity of dust morphologies and the variation in the grain population across their diverse ISM. Each column shows a different galaxy, scaled to 1.5 times the 3D gas half-mass radius ($R_{0.5, \text{gas}}$), indicated by the dotted circles, where the constant physical scale of 5 kpc is indicated above each projection using black whiskers. H_2 surface densities are also indicated in each image using contours to map gas species (*dot* and *dashed* enclose 50 per cent of hydrogen in H I and H_2 respectively, with *solid* 90 per cent H_2). The dust distribution above the 60th percentile in dust density is shaded by different grain property ratios by mass in each row. From top to bottom, these show the DTZ ratio, the small-to-large grain ratio, the ratio of silicate to graphite grains, and the ratio of *forsterite* (Mg-endmember) to *fayalite* (Fe-endmember) silicates. Significant intra- and intergalaxy variations can have important implications for the processing and re-emission of radiation.

4.3 Cosmic dust density evolution

For insight into evolutionary processes further to Section 4.2, we investigate how the total dust content evolves with redshift in our simulations. Fig. 7 shows the cosmic dust density, ρ_{dust} , evolution (total dust mass density per co-moving volume, as a function of cosmic lookback time). Observational data (grey markers, M. Fukugita & P. J. E. Peebles 2004; C. Vlahakis et al. 2005; R. A. Beeston et al. 2018; S. P. Driver et al. 2018; L. Dunne et al. 2018; C. Péroux & J. C. Howk 2020; F. Pozzi et al. 2020; A. Traina et al. 2024; S. Berta et al. 2025; Y.-K. Chiang, R. Makiya & B. Ménard 2025) and alternative dust models (grey lines, S. Aoyama et al. 2017; G. Popping et al. 2017; Q. Li et al. 2019; A. P. Vijayan et al. 2019), are plotted here for comparison. Following Fig. 5, *SeedOnly* yields ρ_{dust} values lower than all the data and models, by ≈ 0.5 dex at its closest point ($z \approx 0$), showing the necessity of ISM accretion in our model to reproduce cosmic dust densities. By contrast, *NoDestruction* produces ρ_{dust} much higher than typically observed, and notably above the upper limit cosmic infrared background estimates of Y.-K. Chiang et al. (2025) $z \lesssim 2$ and the cosmic mass budgeting of M. Fukugita & P. J. E. Peebles (2004). The *Fiducial* model shows ρ_{dust} intermediate between the two, and is consistent with the dust given the variance between the different observations. The combination of accretive and destructive processes in *Fiducial* also induces a turnover in the evolution of ρ_{dust} that is not seen in

the other two models. Finally, we see that neglecting small grains, size-related effects and transfer processes (*OneSize*), leads to a ≈ 50 per cent reduction in ρ_{dust} relative to *Fiducial*.

The variety of observational data sets represent a number of different approaches to inferring ρ_{dust} through cosmic time. Galaxy survey approaches (e.g. R. A. Beeston et al. 2018; S. P. Driver et al. 2018; L. Dunne et al. 2018; F. Pozzi et al. 2020; A. Traina et al. 2024; S. Berta et al. 2025; C. Vlahakis et al. 2005), summing values from FIR detected galaxies or integrating the galaxy dust mass function, require some form of completeness corrections for faint galaxies (though see S. Duivenvoorden et al. 2020). Another approach is integrating the total cosmic infrared background (Y.-K. Chiang et al. 2025), which avoids the completeness correction. These FIR approaches can be sensitive to temperature if the shape of the SED is not well constrained, but this can be mitigated by sampling across the Rayleigh-Jeans tail (e.g. A. Traina et al. 2024; Y.-K. Chiang et al. 2025). Another important observational assumption is the wavelength-dependent dust opacity, $\kappa(\lambda)$. Commonly this is taken to be $\kappa(850 \mu\text{m}) = 0.77 \text{ cm}^2 \text{ g}^{-1}$, though estimates can vary by an order of magnitude (C. J. R. Clark et al. 2019). Alternative probes, such as measuring depletion in damped $\text{Ly}\alpha$ absorbers (C. Péroux & J. C. Howk 2020) or cosmic budgetary arguments (M. Fukugita & P. J. E. Peebles 2004) provide independent estimates. Some studies infer significant dust in the CGM of galaxies through




Figure 7. Co-moving cosmic dust mass density ($\rho_{c,dust}$) as a function of lookback time, compared to observationally derived values. Comparing the Fiducial simulation (blue) with the SeedOnly (orange), NoDestruction (green), and OneSize (red) runs. Lines show both the cooling coupled (solid) and uncoupled (dashed) versions of each run (excepting SeedOnly which is always uncoupled). Filled circles include only dust in galaxies with stellar masses $\log_{10}(M_*/M_\odot) > 9$ in 50 pkpc apertures. Observations span a range of redshifts and approaches (SED fitting, elemental depletions, attenuations; C. Vlahakis, L. Dunne & S. Eales 2005; R. A. Beeston et al. 2018; S. P. Driver et al. 2018; L. Dunne et al. 2018; C. Péroux & J. C. Howk 2020; F. Pozzi et al. 2020). The total $\rho_{c,dust}$ of the Fiducial run lies between the SeedOnly and NoDestruction runs, while being comparable to the various data sets. However, when considering just dust in galaxies (coloured data points) the Fiducial and NoDestruction runs are comparable, emphasizing how the excess dust in NoDestruction is outside galaxies (see Fig. 7), while the small impact on Fiducial implies that in the ISM, dust fraction are set by the balance between accretion and metal depletion.

indirect observation and budgetary arguments for metal production ranging from more or comparable dust content in the CGM relative to the galactic ISM (e.g. B. Ménard et al. 2010; M. Fukugita 2011; J. E. G. Peek, B. Ménard & L. Corrales 2015; J. Meinke et al. 2023) to ~ 10 per cent in the ISM (e.g. A. McCormick et al. 2018; M. Romano et al. 2024). To facilitate a comparison with such estimates, we additionally compute ρ_{dust} within galaxies only, by applying a 50 pkpc exclusive spherical aperture about the most-bound particle of each subhalo (the fiducial aperture scale for COLIBRE calibration) for galaxies of $M_* > 10^9 M_\odot$, at $z \in [0, 1, 2, 3]$ for the coupled runs (coloured circles). We see that with this approach, the Fiducial and NoDestruction runs are much closer together, both agreeing well with the observed data, with the OneSize lower (showing a similar offset below the total ρ_{dust} as with Fiducial). This shows the large increase in CGM dust in the NoDestruction run, with ≈ 90 per cent of dust outside galaxies, by mass. The lack of destruction in SeedOnly also leads to a similar fraction of dust existing outside galaxies, such that aperture values all lay below the plotting range (downward arrows). In contrast, the Fiducial run has $\lesssim 30$ per cent of the dust mass in the CGM. NoDestruction (SeedOnly) overproduces (under-produces) CGM dust compared to all observations, while the Fiducial run is consistent with observations finding lower levels of CGM dust (e.g. A. McCormick et al. 2018; M. Romano et al. 2024).

For these runs we also show uncoupled equivalents of runs (i.e. where dust model has no influence on other gas physics, see Section 3.7), indicated by dashed lines, and excepting the SeedOnly run which is only run in an uncoupled context. In other cases we see that the uncoupled runs yield a factor ~ 2 lower ρ_{dust} values by $z = 1$, as the self-consistent dust-treatment leads to greater grain growth by

accretion. In addition, for the SeedOnly run we also show an orange dotted curve illustrating ρ_{dust} boosted by an order of magnitude, and empty orange circles showing the same boost for only dust within the $M_* > 10^9 M_\odot$ sample. This is intended as a rough approximation of dust yield levels comparable to e.g. E. Dwek (1998). We see that a high-yield paradigm appears better at reproducing the observations in lieu of evolutionary processes. However, we see that only counting dust within $M_* > 10^9 M_\odot$ galaxies is still significantly low relative to the data so would require dust growth. This shows that adopting higher dust yields (discussed in Section 3.2) could be a means to reproduce observations regulated with milder evolutionary processes (e.g. a lower clumping factor, Section 3.8), as found by e.g. C. R. Choban et al. (2025).

4.4 Galaxy dust-mass function

In Fig. 8 we plot the Fiducial galaxy-dust mass function (GDMF), the cosmic number density of galaxies per unit logarithmic galaxy dust mass bin width (dex^{-1}), as a function of M_{dust} , corrected to our assumed cosmology. To indicate the influence of resolution, we plot the mass function of galaxies with stellar masses $\log_{10} M_*/M_\odot > 9$ (i.e. resolved by $\gtrsim 500$ star particles), as well as for a broader selection, $\log_{10} M_*/M_\odot > 8$, using thick and thin blue lines, respectively. The narrower selection is plotted down to galaxy masses where it is > 85 per cent complete.

We plot a variety of observed GDMFs from the literature. For values published in the last 10 yr, we indicate each bin value as individual data points. For older mass functions, we plot their best-fitting single Schechter functions (P. Schechter 1976), as collated by C. J. R. Clark et al. (2015). All Schechter functions are extended




Figure 8. The galaxy dust mass function (GDMF). We plot values for our *Fiducial* simulation (blue, solid lines), indicating galaxies with stellar mass $\log_{10} M_*/M_\odot > 9$ (thick blue line) for bins where the selection is ≥ 85 per cent complete, as well as a broader selection with $\log_{10} M_*/M_\odot > 8$ (thin blue line). We plot observational data from a number of literature sources. Recent literature mass functions are plotted using individual data points for each M_{dust} bin (C. J. R. Clark et al. 2015; R. A. Beeston et al. 2018; F. Pozzi et al. 2020), while those that are older are presented as single-Schechter fits (patterned, translucent lines C. Vlahakis et al. 2005; L. Dunne et al. 2011; M. S. Clemens et al. 2013), as collated by C. J. R. Clark et al. (2015). Where Schechter fits are extrapolated beyond the fitting range, a thinner line-style is used. We see a general consistency with observational data.

down to $\log_{10} M_{\text{dust}}/M_\odot = 5$, but we indicate values extrapolating the fitting range using a thinner line style.

We first note differences between the observational GDMF values. A variety of low-mass slopes and normalizations are exhibited, from close to flat (L. Dunne et al. 2011, $\alpha = -1.01$) to negative (C. Vlahakis et al. 2005, $\alpha = -1.67$), and even positive (C. J. R. Clark et al. 2015). This can be attributed to the smaller observational volumes over which low-mass galaxies can be detected, inducing larger uncertainty in the low-mass GDMF. As a result, the Schechter functions are fit down to different limits and need to be extrapolated down to $\log_{10} M_{\text{dust}}/M_\odot = 5$, where they vary by ≈ 1 dex. The number density of galaxies around the knee of the mass function (where the contribution to the cosmic mass density peaks) shows two distinct normalizations, with R. A. Beeston et al. (2018); L. Dunne et al. (2011); F. Pozzi et al. (2020) and C. Vlahakis et al. (2005) within ≈ 0.1 dex of $\log_{10} \Phi / (\text{cMpc}^{-3} \text{dex}^{-1}) = -2.5$ and M. S. Clemens et al. (2013) and C. J. R. Clark et al. (2015) at $\log_{10} \Phi / (\text{cMpc}^{-3} \text{dex}^{-1}) \approx -2.1$. An important consideration for comparing the GDMF is the evolution observed in the late universe, with cosmic dust densities reducing by a factor ≈ 5 over the past 3–5 Gyr for some observations (L. Dunne et al. 2011; F. Pozzi et al. 2020; R. A. Beeston et al. 2024, see fig. 7). Here, we compare dust mass functions up to $z = 0.25$.

Comparing to these observations, we see that our *Fiducial* simulation GDMF exhibits general consistency with the range of observations. At lower dust masses ($5.5 \lesssim \log_{10} M_{\text{dust}}/M_\odot \lesssim 7.25$) the *Fiducial* GDMF traces the R. A. Beeston et al. (2018) GDMF well, which extends well-constrained masses significantly below the limits of prior works ($\log_{10} M_{\text{dust}}/M_\odot \sim 4.5$) owing to the large

number of galaxies in the optically-selected GAMA (S. P. Driver et al. 2011) galaxy sample. At higher masses ($\log_{10} M_{\text{dust}}/M_\odot \gtrsim 7.5$) we see that the knee of the GDMF extends to masses marginally higher than those bracketed by observations, and agrees better with the high-knee set of observations (M. S. Clemens et al. 2013; C. J. R. Clark et al. 2015). While the *Fiducial* model GDMF agrees well with observational sets in the high- and low-mass regime, it does not follow one observed GDMF consistently. The higher knee of the GDMF we predict appears consistent with the slightly higher $\rho_{\text{c,dust}}$ we predict at low redshift.

4.5 Galaxy dust-mass scaling relations

To further evaluate the relationship between dust masses and the other phases of the ISM in our modelling, we plot galactic dust masses (M_{dust}) as a function of various salient properties in Fig. 9, comparing to the observed scaling relations from the *DustPedia* catalogue, as presented by S. Bianchi et al. (2018) and P. De Vis et al. (2019). On the observational side, these are derived by applying the CIGALE Bayesian SED-fitting code to galaxies in the *DustPedia* catalogue. In particular, we compare to the galaxy stellar masses (M_*), the ongoing galaxy star formation rates (SFRs), the neutral (i.e. non-ionized) gas masses ($M_{\text{neut}} = M_{\text{HI}} + M_{\text{H}_2}$) and the H_2 masses (M_{H_2}).

For the M_*-M_{dust} relation, we see a generally concave increasing relation, superlinear for $\log_{10} M_*/M_\odot < 8$, close to linear for $8 < \log_{10} M_*/M_\odot < 10$, and exhibiting a turnover for $\log_{10} M_*/M_\odot > 10$. This relation is relatively tight, with a scatter of $\lesssim 0.5$ dex at $\log_{10} M_*/M_\odot = 9$. At higher M_* , the turnover is induced by an increasing scatter to low M_{dust} in $\log_{10} M_*/M_\odot = 10$ galaxies, increasing to ~ 2 dex. Physically, this scatter to lower M_{dust} owes to increasing passive fractions at high M_* , as evidenced by the relatively tight relation displayed by the $M_{\text{dust}}-\text{SFR}$ relation (top right, discussed below), tightening with M_* . Comparing to the observations of S. Bianchi et al. (2018), we see that the simulated galaxies are consistent with the data, tracing well the upper ridge line amongst the data points. These observations similarly exhibit a marked increase in scatter at higher masses ($M_*/M_\odot > 10$), though they appear to show a larger number of galaxies scattered to low M_{dust} extending to low-intermediate stellar masses ($\log_{10} M_*/M_\odot > 8.5$).¹⁷ We also compare to the multisurvey compilation of I. De Looze et al. (2020), which exhibits less scatter at intermediate masses, agreeing very well with the *Fiducial* galaxies in the $8.5 \lesssim \log_{10} M_*/M_\odot \lesssim 10$ range. This compilation also exhibits increased scatter at high M_* , though significantly less than seen in the S. Bianchi et al. (2018) data points. This owes to the differing selection functions of the surveys, in particular the heavy sampling of Virgo and Fornax clusters given the median sample distance of 20 Mpc, and the stronger representation of early type galaxies in *DustPedia* scattering to lower dust-to-stellar ratios (J. I. Davies et al. 2019). The *Fiducial* simulation appears intermediate between these two samples.

The median $M_{\text{dust}}-\text{SFR}$ relation of Fig. 9 (top right) is close to linear in the range $-2 < \log_{10} \text{SFR}/(M_\odot \text{yr}^{-1}) < 0.5$, with increasing scatter towards lower SFR values, $\lesssim 0.5$ dex at $1 M_\odot \text{yr}^{-1}$, and reaching ~ 2 dex for $0.001 M_\odot \text{yr}^{-1}$. This relation follows the S. Bianchi et al. (2018) data, albeit with a slightly steeper slope; on the lower side at $-3 < \log_{10} \text{SFR}/(M_\odot \text{yr}^{-1}) < -2$, and on the higher

¹⁷We note that the COLIBRE hybrid AGN feedback model, calibrated in larger volumes than presented here, enhances the high-mass scatter in this relation by 0.2–0.3 dex for galaxies of $\log_{10} M_*/M_\odot > 10$ in a 100^3Mpc^3 volume at m6 resolution.

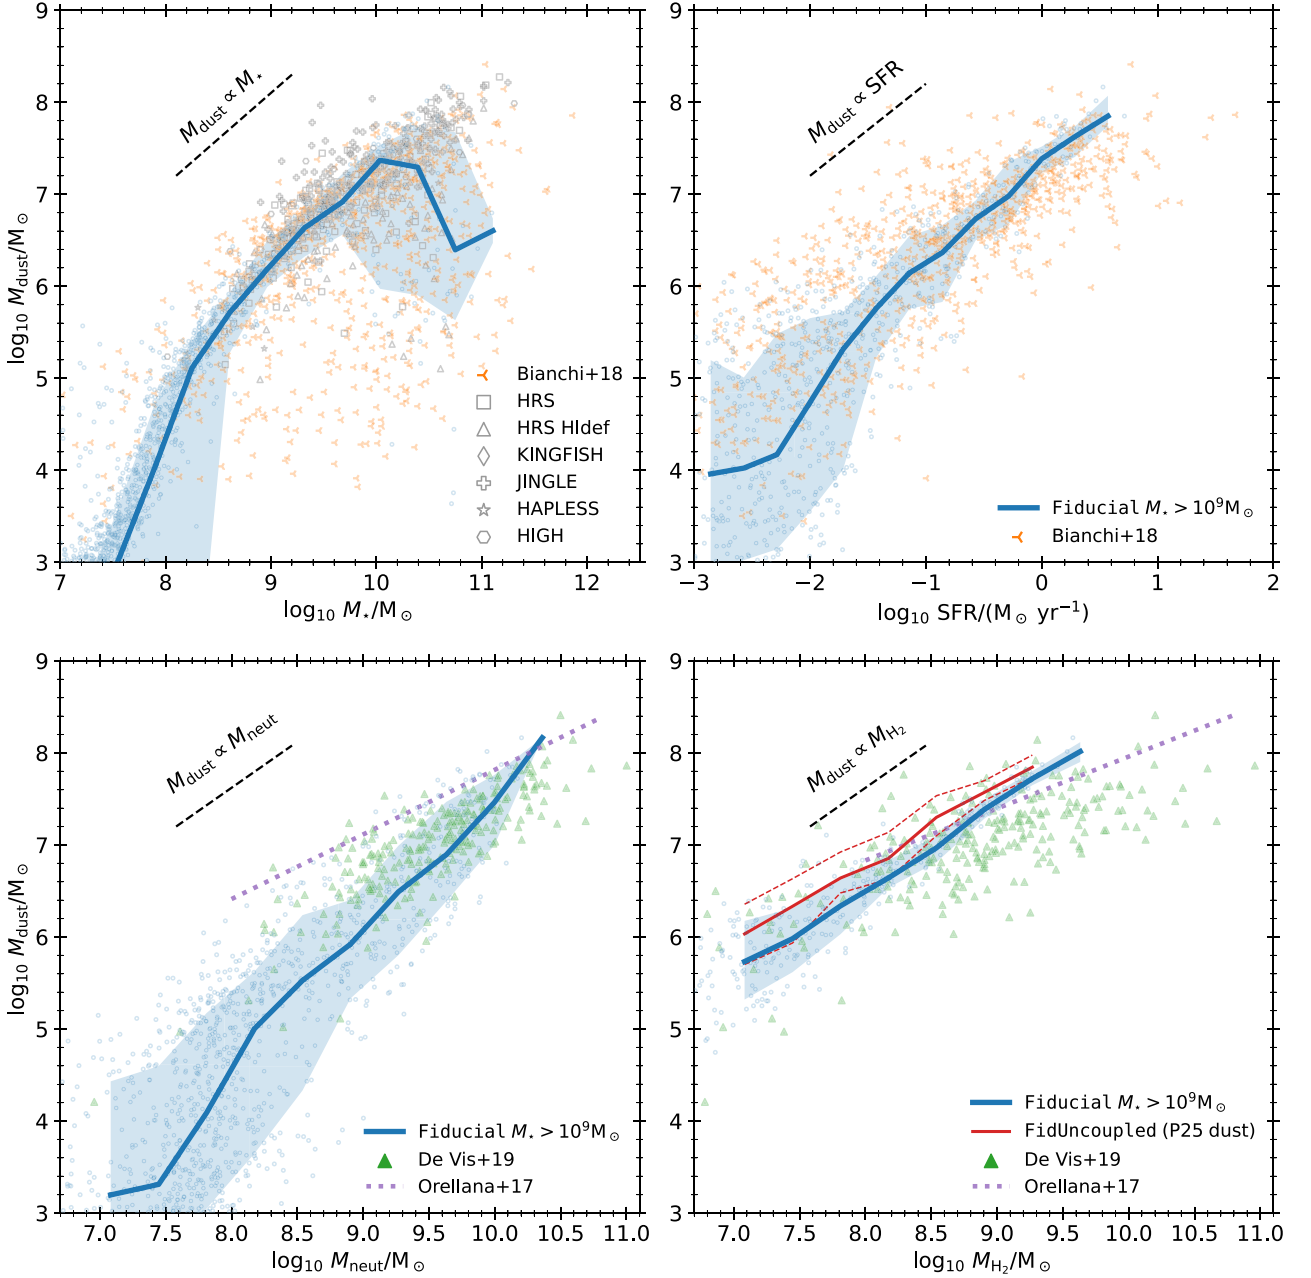


Figure 9. Galaxy dust scaling relations, showing the dust mass (M_{dust}) as a function of other baryonic (star and gas) properties of galaxies. In each panel, we show the median relation for galaxies taken from the `Fiducial` simulation (solid blue line), with shading indicating the 16th-84th percentile range (translucent blue). Individual galaxies are also plotted as low opacity blue data points. The *top left* panel shows the $M_{*}-M_{\text{dust}}$ relation, with the *top right* panel shows the $\text{SFR}-M_{\text{dust}}$ relation, comparing both to *DustPedia* data presented S. Bianchi et al. (2018). We also compare $M_{*}-M_{\text{dust}}$ specifically to the survey compilation of I. De Looze et al. (2020) (indicated for each contributing survey using grey markers). The *bottom left* panel shows the $M_{\text{neut}}-M_{\text{dust}}$ relation, while the *bottom right* panel shows $M_{\text{H}_2}-M_{\text{dust}}$, comparing to data using the masses of hydrogen phases in *DustPedia* galaxies, as presented by P. De Vis et al. (2019), as well as a power-law fit to the Planck-2MRS data (G. Orellana et al. 2017). To show if the tight $M_{\text{H}_2}-M_{\text{dust}}$ relation is driven by the dust coupling, we additionally plot medians in the bottom-right panel and scatter from the `FidUncoupled` runs, aggregating the P25 dust properties with SOAP.

side by $\log_{10} \text{SFR}/(M_{\odot} \text{ yr}^{-1}) \approx 0$. The observations show a similar trend for the scatter, with slightly larger amplitude. Additional scatter in the observations could be attributed to the larger uncertainties in measuring dim FIR sources corresponding to the lowest dust masses.

Comparing to gas masses in the bottom row, the $M_{\text{neut}}-M_{\text{dust}}$ relation (bottom left) for median `Fiducial` galaxies exhibits a steadily increasing, slightly superlinear slope for $\log_{10}(M_{\text{neut}}/M_{\odot}) > 8.25$, with a steeper slope below this, owing to a transition to lower dust

densities at $\log_{10}(M_{\text{neut}}/M_{\odot}) \approx 8$. This is likely due to a physical transition in galaxies at an ISM density threshold, where the \mathcal{DTZ} relation transitions from being set by the seeding, to being set by a balance between accretion and the availability of gas-phase elements for depletion (see following section).

Comparing to the data, we see a general consistency with the cloud of observed data points from P. De Vis et al. (2019). The median relation appears steeper than observed, with an excess of dust

at the highest neutral masses [$\log_{10}(M_{\text{neut}}/M_{\odot}) \sim 10.25$] compared to P. De Vis et al. (2019). Comparing to the line representing the data from G. Orellana et al. (2017), a power-law fit to the values they derive for their 2MRS + Planck selection, we see a higher normalization than P. De Vis et al. (2019), but actually agreeing well with Fiducial galaxies at $\log_{10}(M_{\text{neut}}/M_{\odot}) \approx 10.25$. However G. Orellana et al. (2017) also shows a shallower, sublinear relation than Fiducial, such that the simulated galaxies diverge increasingly towards lower masses. The higher M_{dust} values obtained for the *Planck* observations is consistent with the picture we see for the GDMF (Fig. 8), where the M. S. Clemens et al. (2013) GDMF shows a higher normalization of the knee and a steeper low-mass slope compared to other observations, which agrees well with Fiducial at high M_{dust} , but is above the simulation and other data at lower M_{dust} . The Fiducial scatter appears comparable to that of P. De Vis et al. (2019), narrowing towards high M_{dust} and M_{neut} , attributable to the limited sampling of massive galaxies in these 25^3 cMpc^3 volumes.

The bottom right panel then plots dust mass against molecular-phase hydrogen mass ($M_{\text{dust}}-M_{\text{H}_2}$). Here the median Fiducial relation is distinctly sublinear, where an increase in dust mass corresponds a greater proportional change in M_{H_2} . This is perhaps unsurprising; alongside the regular mass trends, dust catalyses molecule formation, due to the crucial role of grain surfaces as formation sites for molecules and shielding gas from radiation. This relationship is notably tighter than that of $M_{\text{dust}}-M_{\text{neut}}$, with a scatter of 0.2–0.3 dex in M_{dust} for $\log_{10}(M_{\text{H}_2}/M_{\odot}) > 8.25$. This could imply that the physical relationship between M_{dust} and M_{H_2} is regulating the scatter. To test this, we also plot the fiducial run with our dust model de-coupled from the cooling (using the hybrid model P25 dust, see Section 2.4); *FidUncoupled*, where medians and 16–84th percentile range shown via red solid and dashed lines, respectively. Here we see that the relation is offset high by up to ≈ 0.3 dex for the *FidUncoupled* run when aggregating the implied P25 dust relative to *Fiducial* galaxies, with a comparable level of scatter. This suggests that the dust-model coupling boosts the H_2 mass at a fixed M_{dust} . We again see steeper relationships than are observed, by P. De Vis et al. (2019) and G. Orellana et al. (2017). This again leads to an excess of M_{dust} in *Fiducial* galaxies at high M_{H_2} relative to the P. De Vis et al. (2019), and one that is more pronounced than seen in the M_{neut} comparison (up to 0.5 dex above the data for $\log_{10}(M_{\text{dust}}/M_{\odot}) \approx 8$). The G. Orellana et al. (2017) fit to Planck + 2MRS also shows a familiar offset toward higher M_{dust} , and again agrees better with *Fiducial* at the highest masses, and *FidUncoupled* at lower masses.

4.6 Dust-to-gas ratio

To assess how the proportions of ISM material depleted into the solid dust phase in our simulated galaxies compare to observation, we present the galaxy dust-to-gas (DTG) mass ratios in Fig. 10.

For DTG , plotting values as a function of gas-phase metallicity can reveal how dust content varies with the ongoing enrichment of the ISM. Observationally, this metallicity is typically measured in the gas-phase via oxygen emission lines, taken as a proxy for total metallicity, $12 + \log_{10}(\text{O}/\text{H})$. For our simulated galaxies, we compute this per subhalo for our simulated galaxies in 50 pkpc apertures about the most-bound particles, measuring the gas-phase oxygen fractions in relatively cold, dense gas ($\log_{10} T/\text{K} < 4.5$, $\log_{10} n_{\text{H}}/\text{cm}^{-3} > -1$) gas, to approximate the regions probed observationally. Here, we limit the galaxy selection to actively star-forming galaxies (with specific star formation rates $\text{SFR}/M_{\star} > 0.01 \text{ Gyr}^{-1}$), with $12 + \log_{10}(\text{O}/\text{H})$ abundances aggregated linearly over the subhalo




Figure 10. The dust-to-gas (DTG) ratio of relatively *active* ($\text{sSFR} = \text{SFR}/M_{\star} > 0.01 \text{ Gyr}^{-1}$) galaxies as a function of their gas-phase oxygen abundances, measured in cool ($\log_{10} T/\text{K} < 4.5$), dense ($\log_{10} n_{\text{H}}/\text{cm}^{-3} > -1$) gas, where O/H values are aggregated as the total linear ratio of oxygen to hydrogen in selected, in-aperture gas. These are compared to a compilation of observations (F. Galliano, E. Dwek & P. Chanial 2008; M. Galamez et al. 2011; P. De Vis et al. 2019; P. Cigan et al. 2021; C. Konstantopoulou et al. 2023). *Blue lines* show median values while *blue data points* show individual galaxies. The *solid line* shows the median for galaxies with $M_{\star} > 10^9 M_{\odot}$ and $\pm 1\sigma$ ranges indicated by the shaded region and *blue squares* showing individual galaxies. The *dotted line* shows the median relation for all processed galaxies, with *blue dots* indicating galaxies outside the high mass selection. We also use (KDE-smoothed) contours (with 0.46 dex spacing) to indicate the histogram of P. De Vis et al. (2019) data using an alternative metallicity calibration (T04, C. A. Tremonti et al. 2004), as opposed to their fiducial one (PG16R, L. S. Pilyugin & E. K. Grebel 2016). Metallicity calibrations are investigated further in Appendix C.

material (such that O/H is the ratio of total oxygen to hydrogen across all the cool-dense, in-aperture gas). Another important consideration for the comparison in this case is that with $12 + \log_{10}(\text{O}/\text{H})$ as the independent variable, there can be considerable blending in particle resolution across the x -axis, particularly due to low-mass galaxies scattered to high metallicity bins. By default, we apply a stellar mass cut of $\log_{10}(M_{\star}/M_{\odot}) > 9$ (corresponding to galaxies resolved by $\gtrsim 500$ star particles) to our primary median relation (solid line), while also plotting a median relation using a lower $\log_{10}(M_{\star}/M_{\odot}) > 8$ cut ($\gtrsim 50$ star particles) to show the influence of less well resolved systems.

We see that the median DTG increases with metallicity as seen in the observations, following a slightly steeper relationship than the point cloud of observations. In this comparison, we also observe a noticeable offset between the data and the simulated DTG values at fixed $12 + \log_{10}(\text{O}/\text{H})$, with the simulations $\approx 0.4 - 0.2$ dex lower DTG for a fixed metallicity. This is perhaps surprising, given that dust masses appear to be in good agreement or even offset above observed relations in previous plots (knee of the GDMF in Fig. 8, $\log_{10} M_{\star}/M_{\odot} \approx 7.8$). Given the uncertainty in the absolute calibration of observationally-inferred metallicity values in the gas-phase (P. De Vis et al. 2019), we also make use of the wider *Dustpedia* catalogues provided by P. De Vis et al. (2019), tabulating numerous

metallicity values per galaxy, via a variety of metallicity probes and calibrations from the literature. To show the extent of this variation, we also use contours to plot the P. De Vis et al. (2019) O3N2 calibration of C. A. Tremonti et al. (2004), which is in good agreement with the simulated data at high metallicity. We expand on this calibration issue in Appendix C, also showing similar levels of agreement with the C. A. Tremonti et al. (2004) mass–metallicity relation. We note that the gas selection and aggregation of total O/H for galaxies can also help explain a portion of these differences. In lieu of a more observational, forward modelled measurement of O/H for our galaxies (e.g. fig. 2 of D. Nelson et al. 2018), we stick to a simple linear aggregation.

Regarding the scatter in the relation between \mathcal{DTG} and metallicity, we first see that the intrinsic scatter in our simulated galaxies appears significantly smaller than that of the cloud of observed data points, with a $\pm 1\sigma$ range below 0.4 dex across the metallicity range for galaxies of $M_\star > 10^9 M_\odot$. Considering these data sets individually, we see that a lot of this global scatter comes from systematic offsets between data sets, though the majority of observed sets individually exhibit scatter that is larger than predicted.¹⁸ An exception is the relation of C. Konstantopoulou et al. (2023), which is significantly tighter. It is possible that a lack of scatter could be due to the limited resolution of the simulations, given we are unable to directly model ISM structure below the resolution scale, that could provide an additional source of variance between galaxies. However, variation in scatter between data sets, given intrinsic uncertainties associated with metallicity calibrations and dust measurements, could contribute significantly to the larger observed scatter.

Comparing the relation for better resolved ($\gtrsim 500$ star particles) galaxies with a broader selection using a lower resolution threshold ($\gtrsim 50$ star particles), we see that for $(12 + \log_{10}(\text{O}/\text{H})) > 8.4$ the broader selection is marginally lower ($\lesssim 0.1$ dex) owing to the inclusion of less-resolved galaxies. We begin to see the broad selection \mathcal{DTG} s diverge to lower values at $12 + \log_{10}(\text{O}/\text{H}) \sim 8.25$, falling super-linearly to $12 + \log_{10}(\text{O}/\text{H}) \lesssim 8.1$, and approximately linearly again below that. We note that this transition feature appears to be resolution-dependent, shifting to lower metallicities at higher resolution (Vijayan et al., in preparation). For the narrower selection $12 + \log_{10}(\text{O}/\text{H}) \lesssim 8.2$ we run out of galaxies, owing to a lack of low-metallicity massive objects.

4.7 Dust grain sizes

Another property we can probe with our dust evolution model is the grain size distribution. In Fig. 11, we plot the logarithmic mass ratio of small grains relative to large grains for galaxies, $\log_{10} S/L$. We first focus on ratios computed for all gas within the 50 pkpc exclusive aperture defining each galaxy (blue lines, shading, and points). This provides a metric for the grain size distribution in galaxies, where positive values are small-grain mass dominated and negative values are large-grain mass dominated. We see that for our Fiducial model, median galaxy dust masses are dominated by large (0.1 μm) grains for $\log_{10}(M_\star/M_\odot) > 8.5$. In general, we see a weakly decreasing trend for our median values, such that the average galaxy goes from rough parity between the mass in small and large grains, to most massive galaxies having ≈ 60 per cent of their dust mass in large grains. This trend is significantly smaller than the scatter, however, which is of order 0.4–0.6 dex, and exhibiting diverse extremes; extending between close to an order of magnitude




Figure 11. Ratio of depleted mass in small to large grains, as a function of galaxy stellar mass, measured within exclusive 50 pkpc apertures. We use the ratio of sizes in our 2 grain-size bins, $a_L = 0.1 \mu\text{m}$ and $a_S = 0.01 \mu\text{m}$, to indicate the size distribution and compare it to the observationally inferred data of M. Relaño et al. (2020), inferred through fitting large silicate grain (BG), very small grain and PAH components to the IR SEDs of local galaxies. The lines show median $\log_{10} S/L$ values. For the Fiducial run we show mass weighted values among all bound aperture gas (solid, with shaded 16–84th percentiles), as well as medians computed within an imposed moderate-density ($\log_{10} n_{\text{H}}/\text{cm}^{-3} > 1$, dotted) and high-density ($\log_{10} n_{\text{H}}/\text{cm}^{-3} > 1$, dot-dashed) gas selection. We also show H_2 -mass weighted S/L medians, across all aperture gas (dashed).

more mass in large grains, to around 50 per cent more mass in small grains. This diversity in the size distribution is profound, and will necessarily impact the extinction properties of grains (which are particularly relevant for post-processing observables from simulated galaxies), as well as the rates of dust-related cooling and heating processes at a fixed \mathcal{DTG} ratio.

Obtaining extragalactic dust sizes observationally is a challenging endeavour. These are very difficult to decode from the attenuation curves of external galaxies; the confounding influence of star-dust geometries (J. Fischera, M. A. Dopita & R. S. Sutherland 2003; D. Narayanan et al. 2018; J. W. Trayford et al. 2020) alongside the actual extinction properties (shaped by the underlying population of grains). While certain grain populations may induce distinctive features in attenuation (e.g. the 2175 Å bump, S. Noll et al. 2009), the overall mass-weighted grain distributions are more elusive. By also appealing to the shape of the FIR SED, and making assumptions about the heating of differently sized grains, we may be able to infer sizes indirectly. We make use of values derived by M. Relaño et al. (2020, 2022) for nearby galaxies. These are again presented as a single $\log_{10} S/L$ value per galaxy. M. Relaño et al. (2020) obtain their estimates through fitting three fixed-shape components to the FIR SEDs of galaxies; big (silicate) grains (BG), very small grains (VSG), and PAHs. These use the classical dust templates of F. X. Desert, F. Boulanger & J. L. Puget (1990), and their defined size ranges of each component; with BGs of $a_{\text{BG}} > 0.015 \mu\text{m}$, VSGs of $0.0012 \mu\text{m} < a_{\text{VSG}} < -0.015 \mu\text{m}$ and PAHs with $a_{\text{PAH}} < 0.0012 \mu\text{m}$.

We see that while some observed galaxies are inferred to have a comparable mass in small grains relative to large, as exhibited by the all-gas median values (blue) in our Fiducial simulation

¹⁸The P. De Vis et al. (2019) data set can be seen isolated in Appendix C.

($\log(S/L) \sim -0.1$), most observed galaxies are assigned significantly lower small-to-large grain mass fractions ($\log(S/L) \lesssim -0.5$). These inferred values do come with large error bars, typically ≈ 0.5 dex. The data has a similarly weak trend with M_* , exhibiting a large scatter in small-to-large grain ratio at fixed mass.

To investigate this apparent inconsistency further, we can look at the small-to-large grain ratio within a more stringent gas selection. With S/L ratios showing strong trends in gas density and phase (see e.g. Fig. 5), we overplot medians for H_2 -weighted average S/L values (red), as well as after imposing lower limits on the hydrogen number density, $n_H > 10 \text{ cm}^{-3}$ and $n_H > 100 \text{ cm}^{-3}$. We observe that considering H_2 -weighted gas does generally push S/L to lower values, but only by ≈ 0.1 dex for $\log_{10}(M_*/M_\odot) > 9$. Considering high-density gas selections has a stronger impact, with $n_H > 10 \text{ cm}^{-3}$ gas 0.3–0.5 dex lower than the all-gas ratios, and $n_H > 100 \text{ cm}^{-3}$ with an offset of > 1 dex below all-gas, consistent with the lower sequence of observed values. Along with the gas selection, we note that again resolution appears to play a role here, with the higher m5 (lower m7) resolution yields higher (lower) S/L ratios respectively (Vijayan et al., in preparation).

This shows the influence of dense gas processes, particularly coagulation, in setting the grain sizes, and is suggestive that our fiducial coagulation time-scales may not be short enough to reproduce observational size distributions. There are caveats with this assessment; our two discrete bin sizes do not correspond ideally to the dust templates used to fit observations, there are general challenges with obtaining these size distributions observationally and the selection effects associated with these observations. Dust in denser gas, particularly around young stellar populations, will tend to process significantly more radiation. However, black- or grey-body fits to IR SEDs are taken as a general measure of overall dust content. We await future results with the full COLIBRE model to explore size distributions further, particularly making use of radiative transfer post-processing, where we could reproduce the observational inference of sizes from FIR SEDs,¹⁹ and explore the importance of dense versus diffuse media. In the meantime, we consider an additional boost factor to the coagulation that can be applied to our lower resolution models in future work.

5 DISCUSSION AND CONCLUSIONS

We have presented a model for the life-cycle of dust and its interaction with the physics of the ISM, implemented using the SWIFT SPH code (M. Schaller et al. 2024), as part of a network of physics modules constituting the COLIBRE suite of galaxy formation simulations (J. Schaye et al. 2025). This model uses an established two-size (0.01 and 0.1 μm) grain population paradigm (H. Hirashita 2015), as well as three chemical species of dust (carbonaceous grain, as well as Mg- and Fe-endmember silicate species), providing a model lightweight enough to be carried particle-by-particle in large SPH cosmological volume simulations. This dust is seeded from stellar channels, budgeting for consistency with overall metal enrichment, and can accrue mass from the gas-phase ISM (accretion), change its size distribution (shattering and coagulation), and ultimately be destroyed (sputtering, direct shock annihilation, or astration). During this life-cycle, dust properties are coupled to the cooling and heating processes, accounting for the role of dust in e.g. absorbing and re-emitting ionizing radiation, providing surfaces for molecule

¹⁹Also considering how to use our two size bins to parametrize a continuous size distribution for radiative modelling.

formation and depleting metals from the gas-phase that would otherwise contribute to cooling and heating processes. We presented results from $(25 \text{ cMpc})^3$ moderate resolution ($m_{\text{gas}} \approx m_{\text{DM}} \approx 10^6 M_\odot$) simulations showing the effect of parameter and modelling choices on the distribution, depletion, and evolution of dust, as well as results comparing galaxy dust distributions and scaling relations at redshift $z = 0$ to observations.

Generally, we find that our model produces agreement with a variety of observed properties. Our `Fiducial` run exhibits a cosmic dust density ($\rho_{\text{c,dust}}$) evolution in good agreement with a wide range of observations across different redshift windows and observational probes (Fig. 7). We put particular stock in the recent Y.-K. Chiang et al. (2025) results, which integrate the cosmic infrared background to probe all dust emission. We regard this as more representative of the $\rho_{\text{c,dust}}$ value that we compute for the simulations, where we sum all the dust in the simulated volume. This shows a new level of agreement across the redshift range for cosmological models of dust evolution, compared to various literature models we consider. While it is arguable that the predicted $\rho_{\text{c,dust}}$ evolution is flatter than observed over the range $0 < z < 2$, on the high side compared to low redshift observations and vice-versa, this is consistent with volume effects in a relatively small 25^3 cMpc^3 box that lacks the massive haloes that assemble vigorously at high redshift but are largely evacuated of ISM at low redshift.

The $z = 0$ GDMF (Fig. 8) also shows broad concordance with observations, favouring the set of observations that find more galaxies at the highest dust masses (e.g. M. S. Clemens et al. 2013; C. J. R. Clark et al. 2015), consistent with being on the high side of low- z $\rho_{\text{c,dust}}$ observations. The simultaneous low-mass slope preference for lower normalization GDMFs (e.g. R. A. Beeston et al. 2018), however, may point to a different shape of the GDMF overall. Dust scaling relations (Fig. 9) convey a similar story, showing generally good agreement, agreeing better with higher normalized observations, particularly compared to atomic and molecular gas masses (e.g. G. Orellana et al. 2017). We note that the inclusion of small grains in our coupled dust model helps to improve the general agreement relative to the P25 dust model, increasing the cosmic H_2 densities (Fig. 2) and H_2 abundance at fixed dust mass (bottom right, Fig. 9).

A curious tension appears when considering the DTG - Z relation (Fig. 10); despite the aforementioned dust comparisons suggesting similar or higher dust content in `Fiducial` compared to observations at $z = 0$, the DTG - Z relation is generally low compared to some of the observational data. Making use of the multiple metallicity calibrations of P. De Vis et al. (2019), however, we find that a calibration yielding higher metallicities (e.g. C. A. Tremonti et al. 2004) confers excellent agreement with our simulations (explored further in Appendix C). This combination of dust content (constrained well by FIR observations), depletion constraints (largely saturated in the ISM) and the DTG ratio may suggest a preference for the higher gas-phase metallicity calibration.

A more elusive tension is seen for grain sizes (Fig. 11), where we see a systematic overabundance of small grains relative to the majority of M. Relaño et al. (2020, 2022) observations. It is worth noting that inferring these observed grain sizes is difficult, requiring fitting multicomponent fits to the FIR SED, with many potential systematic effects. However, we do find improved agreement using more exclusive selections of cold, dense gas where coagulation can be effective. It is unclear to what extent the observationally inferred values are more representative of dust in dense regions, and we look to investigate this connection by re-deriving size distributions using virtual FIR SEDs in future work. Other works in isolated galaxy simulations have found they can reproduce these more large-grain

size dominated distributions, with MW-like net extinction curves²⁰ at the same time, with assumptions on how these simple two-size approximations constrain a broader size distribution (e.g. K.-C. Hou et al. 2019; G. L. Granato et al. 2021; Y. Dubois et al. 2024). Still, in our coupled cooling model these smaller grains also help increase the cosmic density of molecular gas, and its correlation with dust masses. We note that the grain size distributions are resolution-dependent, with increased (reduced) small-to-large grain size ratios for higher (lower) resolution. We will explore applying a coagulation boost to lower resolution runs in upcoming work as a means to reduce the small-to-large grain ratios.

Fig. 6 shows the intra- and intergalaxy variations in dust grain populations, that may have important implications for the observable properties, largely mediated by dust. Observationally, these effects may seem intractable from complex star-dust geometries, which can, by themselves, yield significant diversity in attenuation properties (e.g. D. Narayanan et al. 2018; J. W. Trayford et al. 2020). However we can account for the grain population and geometric effects self-consistently using radiative transfer forward modelling, tracing the propagation of starlight through the dusty medium and accounting for absorption, scattering and re-emission. This was done for e.g. the EAGLE simulations using the SKIRT code (P. Camps & M. Baes 2015, 2020) assuming a fixed dust mix (P. Camps et al. 2016, 2018; J. W. Trayford et al. 2017), but when coupled with a live dust model such as ours, it can offer broad insight into the properties shaping attenuation, and its potential effects on galaxy property inference. We will explore forward modelling with dust radiative transfer in upcoming work.

Detailed dust evolution within galaxy formation models that incorporate a cold phase, with molecular gas and dust physics, are emerging, but dust typically only has a *passive* role (e.g. for FIRE-2 by C. R. Choban et al. 2022, where radiation pressure, shielding, heating/cooling, and prescriptions for molecule formation are unaware of the modelled dust). Among dust evolution models in general, it is still uncommon to have dust *coupling* to gas physics through the self-consistent influence on heating and cooling processes associated with grains (though see e.g. M. Vogelsberger et al. 2019; O. Osman, K. Bekki & L. Cortese 2020). G. L. Granato et al. (2021) present such a model within cosmological zoom-in simulations, using the MUPPI framework, which implements semi-analytic models for multiphase ISM within GADGET SPH particles (G. Murante et al. 2010, 2015; M. Valentini et al. 2017, 2019, 2023), and noting the important role of this coupling. This model shows the coevolution of dust and gas over cosmic time within simulations of MW-like galaxies, exhibiting universally lower small-to-large grain size ratios, in closer agreement with observational inference (M. Relaño et al. 2020, 2022) than we find here. Cosmological simulations adapting this dust model are presented in M. Parente et al. (2022) and C. Ragone-Figueroa et al. (2024) at a baryonic mass resolution $m_g = 7 \times 10^6 M_\odot$. While the $z = 0$ population of M. Parente et al. (2022) presents remarkable concordance with the galaxy dust mass functions and consistency with *DTG-Z* observations, this is in the context where the mass–metallicity relation is too low and the cosmic star formation rate density and dust masses are underpredicted at high redshift, attributed to the direct application of the G. L. Granato et al. (2021) model to cosmological simulations. C. Ragone-Figueroa et al. (2024) calibrate this model to improve the cosmic evolution for dust and gas properties in follow-up simulations, similar to those we

present in our higher resolution runs with an explicit cold phase, with more results imminent from the larger and higher-fidelity volumes of the COLIBRE suite.

Overall, we present a model generally consistent with a range of observations at low redshift, and able to reproduce the evolution of dust content well across cosmic time. While lightweight, the power of this model is in its self-consistent interaction with the panoply of sophisticated physics modules constituting the large-scale upcoming COLIBRE simulations, while setting the stage for novel insight into a profusion of contemporary dust-mediated observations through a forward modelling approach.

ACKNOWLEDGEMENTS

We are thankful to our referee whose detailed suggestions have helped to improve a number of parts of the manuscript significantly. We gratefully acknowledge everyone who has contributed to the COLIBRE project for their input. This work used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The equipment was funded by BEIS capital funding via STFC capital grants ST/K00042X/1, ST/P002293/1, ST/R002371/1, and ST/S002502/1, Durham University and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure. JT acknowledges support of his STFC *Early Stage Research and Development grant* (ST/X004651/1). We are grateful to Maarten Baes, Andrea Gebek, and Nick Andreadis for helpful discussions in the latter phase of this work. This work made generous use of the pipeline tools developed for the COLIBRE project, representing many years of work; including the *swiftsimio* (J. Borrow & A. Borrisov 2020) and SOAP²¹ tools. SP acknowledges support from the Austrian Science Fund (FWF) through project V 982-N. ABL acknowledges support by the Italian Ministry for Universities (MUR) program ‘Dipartimenti Di Eccellenza 2023–2027’ within the Centro Bicocca di Cosmologia Quantitativa (BiCoQ), and support by UNIMIB’s Fondo Di Ateneo Quota Competitiva (project 2024-ATEQC-0050). CSF acknowledges support from ERC Advanced Investigator Grant, DMIDAS (GA 786910). This project has received funding from the Netherlands Organization for Scientific Research (NWO) through research programme Athena 184.034.002.

DATA AVAILABILITY

The data presented in this article can be made available to individuals upon reasonable request. The SWIFT and CHIMES codes are already available at www.swiftsim.com and richings.bitbucket.io/chimes/home.html, respectively. The COLIBRE code and data are scheduled for release some time after the initial COLIBRE papers.

REFERENCES

- Abbott T. M. C. et al., 2022, *Phys. Rev. D*, 105, 023520
 Aoyama S., Hou K.-C., Shimizu I., Hirashita H., Todoroki K., Choi J.-H., Nagamine K., 2017, *MNRAS*, 466, 105
 Asano R. S., Takeuchi T. T., Hirashita H., Inoue A. K., 2013a, *Earth Planets Space*, 65, 213
 Asano R. S., Takeuchi T. T., Hirashita H., Nozawa T., 2013b, *MNRAS*, 432, 637
 Asplund M., Grevesse N., Sauval A. J., Scott P., 2009, *ARA&A*, 47, 481

²⁰We defer analysis of extinction curves to future work, concerned with the post-processed extinction and attenuation properties of galaxies.

²¹github.com/openjournals/joss-reviews/issues/7851

- Barišić I. et al., 2020, *ApJ*, 903, 146
- Beeston R. A. et al., 2018, *MNRAS*, 479, 1077
- Beeston R. A., Gomez H. L., Dunne L., Maddox S., Eales S. A., Smith M. W. L., 2024, *MNRAS*, 535, 3162
- Bekki K., 2013, *MNRAS*, 432, 2298
- Bennassuti M. D., Schneider R., Valiante R., Salvadori S., 2014, *MNRAS*, 445, 3039
- Berta S. et al., 2025, *A&A*, 696, A193
- Bianchi S. et al., 2018, *A&A*, 620, A112
- Birnstiel T., Fang M., Johansen A., 2016, *Space Sci. Rev.*, 205, 41
- Blain A. W., Smail I., Ivison R. J., Kneib J. P., Frayer D. T., 2002, *Phys. Rep.*, 369, 111
- Borrow J., Borrisov A., 2020, *J. Open Source Softw.*, 5, 2430
- Borrow J., Schaller M., Bahé Y. M., Schaye J., Ludlow A. D., Ploekinger S., Nobels F. S. J., Altamura E., 2023, *MNRAS*, 526, 2441
- Bouwens R. J. et al., 2012, *ApJ*, 754, 83
- Brownlee D., 2014, *Annu. Rev. Earth Planet. Sci.*, 42, 179
- Burgarella D. et al., 2025, *A&A*, 699, A336
- Camps P., Baes M., 2015, *Astron. Comput.*, 9, 20
- Camps P., Baes M., 2020, *Astron. Comput.*, 31, 100381
- Camps P., Trayford J. W., Baes M., Theuns T., Schaller M., Schaye J., 2016, *MNRAS*, 462, 1057
- Camps P. et al., 2018, *ApJS*, 234, 20
- Canuto V. M., Mazzitelli I., 1991, *ApJ*, 370, 295
- Cazaux S., Tielens A. G. G. M., 2002, *ApJ*, 575, L29
- Chabrier G., 2003, *PASP*, 115, 763
- Chaikin E., Schaye J., Schaller M., Bahé Y. M., Nobels F. S. J., Ploekinger S., 2022, *MNRAS*, 514, 249
- Chaikin E., Schaye J., Schaller M., Benítez-Llambay A., Nobels F. S. J., Ploekinger S., 2023, *MNRAS*, 523, 3709
- Chaikin E. et al., 2025, preprint (arXiv:2509.04067)
- Chandro-Gómez Á. et al., 2025, *MNRAS*, 539, 776
- Chatzikos M. et al., 2023, *Rev. Mex. Astron. Astrofis.*, 59, 327
- Chiang Y.-K., Makiya R., Ménard B., 2025, *ApJ* 992 65
- Cho J., Lazarian A., 2007, *ApJ*, 669, 1085
- Choban C. R., Kereš D., Hopkins P. F., Sandstrom K. M., Hayward C. C., Faucher-Giguère C.-A., 2022, *MNRAS*, 514, 4506
- Choban C. R., Kereš D., Sandstrom K. M., Hopkins P. F., Hayward C. C., Faucher-Giguère C.-A., 2024, *MNRAS*, 529, 2356
- Choban C. R., Salim S., Kereš D., Hayward C. C., Sandstrom K. M., 2025, *MNRAS*, 537, 1518
- Cigan P. et al., 2021, *AJ*, 162, 83
- Cinquegrana G. C., Joyce M., Karakas A. I., 2022, *ApJ*, 939, 50
- Clark C. J. R. et al., 2015, *MNRAS*, 452, 397
- Clark C. J. R. et al., 2019, *MNRAS*, 489, 5256
- Clayton D. D., Amari S., Zinner E., 1997, *Ap&SS*, 251, 355
- Clemens M. S. et al., 2013, *MNRAS*, 433, 695
- Crain R. A. et al., 2015, *MNRAS*, 450, 1937
- Dalla Vecchia C., Schaye J., 2012, *MNRAS*, 426, 140
- Davé R., Anglés-Alcázar D., Narayanan D., Li Q., Rafieferantsoa M. H., Appleby S., 2019, *MNRAS*, 486, 2827
- Davies J. I. et al., 2019, *A&A*, 626, A63
- Dayal P., Hirashita H., Ferrara A., 2010, *MNRAS*, 403, 620
- De Cia A., Ledoux C., Mattsson L., Petitjean P., Srianand R., Gavignaud I., Jenkins E. B., 2016, *A&A*, 596, A97
- De Looze I. et al., 2020, *MNRAS*, 496, 3668
- De Vis P. et al., 2019, *A&A*, 623, A5
- Dell'Agli F., García-Hernández D. A., Schneider R., Ventura P., La Franca F., Valiante R., Marini E., Di Criscienzo M., 2017, *MNRAS*, 467, 4431
- Desert F. X., Boulanger F., Puget J. L., 1990, *A&A*, 237, 215
- Di Criscienzo M. et al., 2016, *MNRAS*, 462, 395
- Doherty C. L., Gil-Pons P., Lau H. H. B., Lattanzio J. C., Siess L., 2014, *MNRAS*, 437, 195
- Dopcke G., Glover S. C. O., Clark P. C., Klessen R. S., 2011, *ApJ*, 729, L3
- Draine B. T., 2003, *ARA&A*, 41, 241
- Draine B. T., 2009, in Henning T., Grün E., Steinacker J., eds, ASP Conf. Ser. Vol. 453, Cosmic Dust – Near And Far (Interstellar Dust Models and Evolutionary Implications). Astron. Soc. Pac., San Francisco, p. 453
- Draine B. T., 2011, *ApJ*, 732, 100
- Draine B. T., Lee H. M., 1984, *ApJ*, 285, 89
- Draine B. T., Li A., 2001, *ApJ*, 551, 807
- Driver S. P. et al., 2011, *MNRAS*, 413, 971
- Driver S. P. et al., 2018, *MNRAS*, 475, 2891
- Dubois Y. et al., 2024, *A&A*, 687, A240
- Duivendoorn S. et al., 2020, *MNRAS*, 491, 1355
- Dunne L. et al., 2009, *MNRAS*, 394, 1307
- Dunne L. et al., 2011, *MNRAS*, 417, 1510
- Dunne L. et al., 2018, *MNRAS*, 479, 1221
- Dwek E., 1998, *ApJ*, 501, 643
- Eldridge J. J., Stanway E. R., Xiao L., McClelland L. A. S., Taylor G., Ng M., Greis S. M. L., Bray J. C., 2017, *Publ. Astron. Soc. Aust.*, 34, e058
- Escala I. et al., 2018, *MNRAS*, 474, 2194
- Faucher-Giguère C.-A., 2020, *MNRAS*, 493, 1614
- Ferrarotti A. S., Gail H. P., 2006, *A&A*, 447, 553
- Finke J. D., Razzaque S., Dermer C. D., 2010, *ApJ*, 712, 238
- Fischera J., Dopita M. A., Sutherland R. S., 2003, *ApJ*, 599, L21
- Fishlock C. K., Karakas A. I., Lugaro M., Yong D., 2014, *ApJ*, 797, 44
- Fogerty S., Forrest W., Watson D. M., Sargent B. A., Koch I., 2016, *ApJ*, 830, 71
- Forouhar Moreno V. J., Helly J., McGibbon R., Schaye J., Schaller M., Han J., Kugel R., Bahe Y., 2025, *MNRAS*, 543, 1339
- Fuente A. et al., 2019, *A&A*, 624, A105
- Fukugita M., 2011, preprint (arXiv:1103.4191)
- Fukugita M., Peebles P. J. E., 2004, *ApJ*, 616, 643
- Galametz M., Madden S. C., Galliano F., Hony S., Bendo G. J., Sauvage M., 2011, *A&A*, 532, A56
- Galliano F., Dwek E., Charnial P., 2008, *ApJ*, 672, 214
- Galliano F. et al., 2021, *A&A*, 649, A18
- Gjergo E., Granato G. L., Murante G., Ragone-Figueroa C., Tornatore L., Borgani S., 2018, *MNRAS*, 479, 2588
- Gomez H. L. et al., 2012, *MNRAS*, 420, 3557
- Gordon K. D., Clayton G. C., Misselt K. A., Landolt A. U., Wolff M. J., 2003, *ApJ*, 594, 279
- Granato G. L. et al., 2021, *MNRAS*, 503, 511
- Groenewegen M. A. T., 1997, *A&A*, 317, 503
- Han J., Cole S., Frenk C. S., Benitez-Llambay A., Helly J., 2018, *MNRAS*, 474, 604
- Hirashita H., 2000, *PASJ*, 52, 585
- Hirashita H., 2012, *MNRAS*, 422, 1263
- Hirashita H., 2015, *MNRAS*, 447, 2937
- Hirashita H., Il'ín V. B., 2022, *MNRAS*, 509, 5771
- Hirashita H., Voshchinnikov N. V., 2013, *MNRAS*, 437, 1636
- Hou K.-C., Aoyama S., Hirashita H., Nagamine K., Shimizu I., 2019, *MNRAS*, 485, 1727
- Inoue A. K., 2003, *PASJ*, 55, 901
- Jenkins E. B., 2009, *ApJ*, 700, 1299
- Jönsson H. et al., 2018, *AJ*, 156, 126
- Karakas A. I., 2010, *MNRAS*, 403, 1413
- Karakas A. I., Lugaro M., 2016, *ApJ*, 825, 26
- Kennicutt R. C. Jr, 1998, *ARA&A*, 36, 189
- Kewley L. J., Ellison S. L., 2008, *ApJ*, 681, 1183
- Kirchschlager F., Schmidt F. D., Barlow M. J., Fogerty E. L., Bevan A., Priestley F. D., 2019, *MNRAS*, 489, 4465
- Kirchschlager F., Schmidt F. D., Barlow M. J., De Looze I., Sartorio N. S., 2023, *MNRAS*, 520, 5042
- Kobayashi C., Umeda H., Nomoto K., Tominaga N., Ohkubo T., 2006, *ApJ*, 653, 1145
- Konstantopoulou C. et al., 2023, *A&A*, 674, C1
- Kozasa T., Nozawa T., Tominaga N., Umeda H., Maeda K., Nomoto K., 2009, in Henning T., Grün E., Steinacker J., eds, ASP Conf. Ser. Vol. 414, Cosmic Dust – Near and Far. Astron. Soc. Pac., San Francisco, p. 43
- Leung S.-C., Nomoto K., 2018, *ApJ*, 861, 143
- Li A., 2020, *Nat. Astron.*, 4, 339
- Li Q., Narayanan D., Davé R., 2019, *MNRAS*, 490, 1425
- Lower S., Narayanan D., Hu C.-Y., Privon G. C., 2024, *ApJ*, 965, 123

- Madden S. C., Galliano F., Jones A. P., Sauvage M., 2006, *A&A*, 446, 877
- Mathews W. G., Brighenti F., 2003, *ApJ*, 590, L5
- Mathis J. S., Rumpl W., Nordsieck K. H., 1977, *ApJ*, 217, 425
- McCormick A. et al., 2018, *MNRAS*, 477, 699
- McKinnon R., Torrey P., Vogelsberger M., 2016, *MNRAS*, 457, 3775
- McKinnon R., Torrey P., Vogelsberger M., Hayward C. C., Marinacci F., 2017, *MNRAS*, 468, 1505
- McKinnon R., Vogelsberger M., Torrey P., Marinacci F., Kannan R., 2018, *MNRAS*, 478, 2851
- McGibbon R., Helly John C., Schaye J., Schaller M., Vandenbroucke B., 2025, *J. Open Source Softw.*, 10, 8252
- Meinke J., Cohen S., Moore J., Böckmann K., Mauskopf P., Scannapieco E., 2023, *ApJ*, 954, 119
- Ménard B., Scranton R., Fukugita M., Richards G., 2010, *MNRAS*, 405, 1025
- Milislavjevic D. et al., 2024, *ApJ*, 965, L27
- Monaghan J. J., Huppert H. E., Worster M. G., 2005, *J. Comput. Phys.*, 206, 684
- Montier L. A., Giard M., 2004, *A&A*, 417, 401
- Murante G., Monaco P., Giovalli M., Borgani S., Diaferio A., 2010, *MNRAS*, 405, 1491
- Murante G., Monaco P., Borgani S., Tornatore L., Dolag K., Goz D., 2015, *MNRAS*, 447, 178
- Murga M. S., Wiebe D. S., Sivkova E. E., Akimkin V. V., 2019, *MNRAS*, 488, 965
- Nanni A., Cristallo S., Donevski D., Michałowski M. J., Romano M., Sawant P., 2024, *A&A*, 684, A163
- Narayanan D., Conroy C., Davé R., Johnson B. D., Popping G., 2018, *ApJ*, 869, 70
- Nelson D. et al., 2018, *MNRAS*, 475, 624
- Nobels F. S. J., Schaye J., Schaller M., Ploekinger S., Chaikin E., Richings A. J., 2024, *MNRAS*, 532, 3299
- Noll S. et al., 2009, *A&A*, 499, 69
- Nomoto K., Kobayashi C., Tominaga N., 2013, *ARA&A*, 51, 457
- Nozawa T., Kozasa T., Habe A., Dwek E., Umeda H., Tominaga N., Maeda K., Nomoto K., 2007, *ApJ*, 666, 955
- Nozawa T., Maeda K., Kozasa T., Tanaka M., Nomoto K., Umeda H., 2011, *ApJ*, 736, 45
- Ochsendorf B. B., Meixner M., Roman-Duval J., Rahman M., Evans N. J. II, 2017, *ApJ*, 841, 109
- Oman K., 2025, *J. Open Source Softw.*, 10, 9278
- Oppenheimer B. D., Schaye J., 2013, *MNRAS*, 434, 1043
- Orellana G. et al., 2017, *A&A*, 602, A68
- Ormel C. W., Paszun D., Dominik C., Tielens A. G. G. M., 2009, *A&A*, 502, 845
- Osman O., Bekki K., Cortese L., 2020, *MNRAS*, 497, 2002
- Padoan P., Cambrésy L., Juvela M., Kritsuk A., Langer W. D., Norman M. L., 2006, *ApJ*, 649, 807
- Palla M. et al., 2024, *MNRAS*, 528, 2407
- Parente M., 2025, preprint (arXiv:2504.10585)
- Parente M., Ragone-Figueroa C., Granato G. L., Borgani S., Murante G., Valentini M., Bressan A., Lapi A., 2022, *MNRAS*, 515, 2053
- Parvathi V. S., Sofia U. J., Murthy J., Babu B. R. S., 2012, *ApJ*, 760, 36
- Peek J. E. G., Ménard B., Corrales L., 2015, *ApJ*, 813, 7
- Péroux C., Howk J. C., 2020, *ARA&A*, 58, 363
- Pilyugin L. S., Grebel E. K., 2016, *MNRAS*, 457, 3678
- Pipino A., Fan X. L., Matteucci F., Calura F., Silva L., Granato G., Maiolino R., 2010, *A&A*, 525, A61
- Ploekinger S., Schaye J., 2020, *MNRAS*, 497, 4857
- Ploekinger S., Richings A. J., Schaye J., Trayford J. W., Schaller M., Chaikin E., 2025, *MNRAS*, 543, 891
- Pokhrel R. et al., 2021, *ApJ*, 912, L19
- Popping G., Puglisi A., Norman C. A., 2017, *MNRAS*, 472, 2315
- Portinari L., Chiosi C., Bressan A., 1998, *A&A*, 334, 505
- Pozzi F., Calura F., Zamorani G., Delvecchio I., Gruppioni C., Santini P., 2020, *MNRAS*, 491, 5073
- Psaradaki I. et al., 2023, *A&A*, 670, A30
- Ragone-Figueroa C., Granato G. L., Parente M., Murante G., Valentini M., Borgani S., Maio U., 2024, *A&A*, 691, A200
- Rau S.-J., Hirashita H., Murga M., 2019, *MNRAS*, 489, 5218
- Relaño M., Lisenfeld U., Hou K. C., De Looze I., Vílchez J. M., Kennicutt R. C., 2020, *A&A*, 636, A18
- Relaño M. et al., 2022, *MNRAS*, 515, 5306
- Rémy-Ruyer A. et al., 2014, *A&A*, 563, A31
- Richings A. J., Faucher-Giguère C.-A., 2018, *MNRAS*, 478, 3100
- Richings A. J., Schaye J., Oppenheimer B. D., 2014a, *MNRAS*, 440, 3349
- Richings A. J., Schaye J., Oppenheimer B. D., 2014b, *MNRAS*, 442, 2780
- Richings A. J., Faucher-Giguère C.-A., Gurvich A. B., Schaye J., Hayward C. C., 2022, *MNRAS*, 517, 1557
- Rogantini D., Costantini E., Zeegers S. T., Mehdipour M., Psaradaki I., Raassen A. J. J., de Vries C. P., Waters L. B. F. M., 2020, *A&A*, 641, A149
- Romano M. et al., 2024, *A&A*, 683, L9
- Sarangi A., Dwek E., Kazanas D., 2019, *ApJ*, 885, 126
- Savage B. D., Sembach K. R., 1996, *ARA&A*, 34, 279
- Schaller M., Gonnet P., Chalk A. B. G., Draper P. W., 2016, Proc. Platform for Advanced Scientific Computing Conference on – PASC’16. Association for Computing Machinery, New York, p. 2
- Schaller M. et al., 2024, *MNRAS*, 530, 2378
- Schaye J. et al., 2015, *MNRAS*, 446, 521
- Schaye J. et al., 2025, preprint (arXiv:2508.21126)
- Schechter P., 1976, *ApJ*, 203, 297
- Schneider R., Valiante R., Ventura P., dell’ Agli F., Di Criscienzo M., Hirashita H., Kemper F., 2014, *MNRAS*, 442, 1440
- Shahbandeh M. et al., 2023, *MNRAS*, 523, 6048
- Shen S., Wadsley J., Stinson G., 2010, *MNRAS*, 407, 1581
- Shivaei I., Darvish B., Sattari Z., Chartab N., Mobasher B., Scoville N., Rieke G., 2020, *ApJ*, 903, L28
- Smagorinsky J., 1963, *Mon. Weather Rev.*, 91, 99
- Sofia U. J., Parvathi V. S., Babu B. R. S., Murthy J., 2011, *AJ*, 141, 22
- Srinivasan S. et al., 2017, *Planet. Space Sci.*, 149, 56
- Stanway E. R., Eldridge J. J., 2018, *MNRAS*, 479, 75
- Sturm E. et al., 2005, *ApJ*, 629, L21
- Su K.-Y. et al., 2018, *MNRAS*, 480, 1666
- Tielens A. G. G. M., 1998, *ApJ*, 499, 267
- Tielens A. G. G. M., 2008, *ARA&A*, 46, 289
- Traina A. et al., 2024, *A&A*, 690, A84
- Travaglio C., Gallino R., Amari S., Zinner E., Woosley S., Lewis R. S., 1999, *ApJ*, 510, 325
- Trayford J. W. et al., 2017, *MNRAS*, 470, 771
- Trayford J. W., Lagos C. d. P., Robotham A. S. G., Obreschkow D., 2020, *MNRAS*, 491, 3937
- Tremonti C. A. et al., 2004, *ApJ*, 613, 898
- Tsai J. C., Mathews W. G., 1995, *ApJ*, 448, 84
- Utomo D. et al., 2018, *ApJ*, 861, L18
- Valentini M., Murante G., Borgani S., Monaco P., Bressan A., Beck A. M., 2017, *MNRAS*, 470, 3167
- Valentini M., Borgani S., Bressan A., Murante G., Tornatore L., Monaco P., 2019, *MNRAS*, 485, 1384
- Valentini M. et al., 2023, *MNRAS*, 518, 1128
- Valiante R., Schneider R., Bianchi S., Andersen A. C., 2009, *MNRAS*, 397, 1661
- Valiante R., Schneider R., Salvadori S., Bianchi S., 2011, *MNRAS*, 416, 1916
- Ventura P., Dell’ Agli F., Schneider R., Di Criscienzo M., Rossi C., La Franca F., Gallerani S., Valiante R., 2014, *MNRAS*, 439, 977
- Vijayan A. P., Clay S. J., Thomas P. A., Yates R. M., Wilkins S. M., Henriques B. M., 2019, *MNRAS*, 489, 4072
- Vlahakis C., Dunne L., Eales S., 2005, *MNRAS*, 364, 1253
- Vogelsberger M., McKinnon R., O’Neil S., Marinacci F., Torrey P., Kannan R., 2019, *MNRAS*, 487, 4870
- Wang L. et al., 2024, *Nat. Astron.*, 8, 504
- Watson D., Christensen L., Knudsen K. K., Richard J., Gallazzi A., Michałowski M. J., 2015, *Nature*, 519, 327
- Weidenschilling S. J., 1977, *MNRAS*, 180, 57
- Weingartner J. C., Draine B. T., 2001, *ApJ*, 548, 296

- Whitaker K. E., Pope A., Cyburski R., Casey C. M., Popping G., Yun M. S., 2017, *ApJ*, 850, 208
- Whittet D. C. B., 2010, *ApJ*, 710, 1009
- Wiersma R. P. C., Schaye J., Theuns T., Dalla Vecchia C., Tornatore L., 2009, *MNRAS*, 399, 574
- Wolfire M. G., McKee C. F., Hollenbach D., Tielens A. G. G. M., 2003, *ApJ*, 587, 278
- Yasuda Y., Kozasa T., 2012, *ApJ*, 745, 159
- Yates R. M., Hendriks D., Vijayan A. P., Izzard R. G., Thomas P. A., Das P., 2024, *MNRAS*, 527, 6292
- Zafar T., Watson D., 2013, *A&A*, 560, A26
- Zeegers S. T., Costantini E., Rogantini D., de Vries C. P., Mutschke H., Mohr P., de Groot F., Tielens A. G. G. M., 2019, *A&A*, 627, A16
- Zhukovska S., Gail H.-P., Trieloff M., 2008, *A&A*, 479, 453

APPENDIX A: ADDITIONAL PARAMETER VARIATIONS

Here, we expand upon the dust model parameter variations presented in the main text, particularly those that are not directly calibrated, using the additional runs of Table 1. Comparing these variations to our Fiducial run demonstrates how these parameters influence the model.

A1 Smaller small grain sizes

The choice of small-grain size influences all evolutionary process rates, as well as rates of dust-mediated cooling and surface nucleation. To show the effect of this choice, we compare to a run using a small-grain size of $0.005 \mu\text{m}$; a value used in a number of dust evolution models (e.g. H. Hirashita 2015; K.-C. Hou et al. 2019; G. L. Granato et al. 2021). We plot the DTZ and small-to-large grain mass ratios in logarithmic bins of density in Fig. A1.

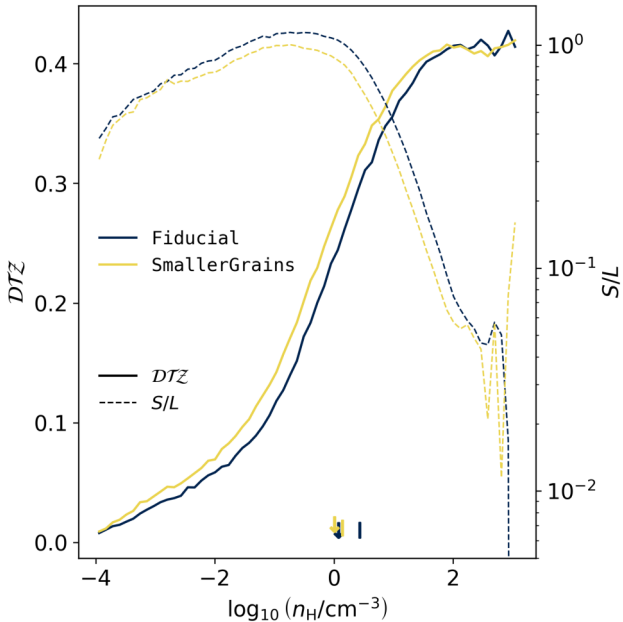


Figure A1. Properties of dust grains in logarithmic bins of gas density (as lower panel of Fig. 3), comparing the Fiducial run with the SmallerGrains run. Solid lines show the dust-to-metal (DTZ) ratios (left y-axis), with dashed-lines indicating the mass in small to large grains (right y-axis). Density transitions for molecular gas ($\rho_{0.5, \text{H}_2}$) and dust ($\rho_{0.5, \text{dust}}$). Using a small grain radius of $0.005 \mu\text{m}$ (relative to the fiducial $0.01 \mu\text{m}$) slightly increases the DTZ , and increases S/L . The H_2 - H_2 transition density is reduced by ≈ 0.3 dex for the SmallerGrains run.

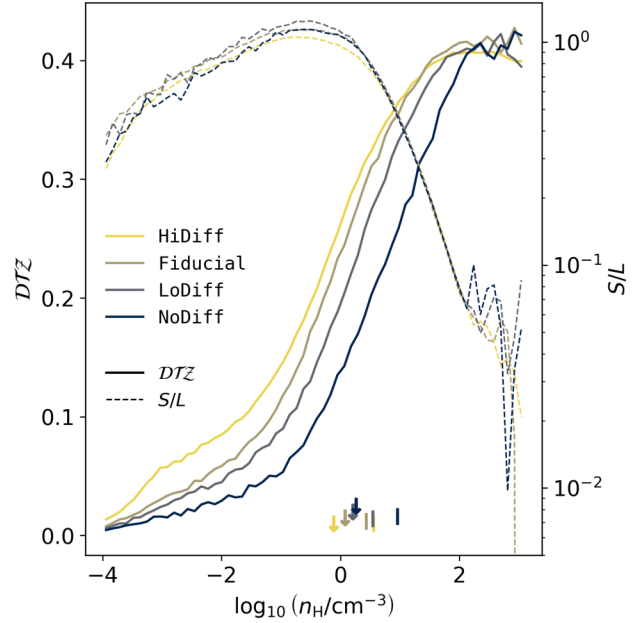


Figure A2. Properties of dust grains in logarithmic bins of gas density – following Fig. A1, now comparing runs with varying levels of turbulent diffusion. Lines from dark blue to light yellow indicate increasing diffusion, from no diffusion, $0.1 \times$ fiducial, fiducial to $10 \times$ fiducial. We see that stronger diffusion leads to a lower transition density from low to high DTZ , with only a marginal change to the S/L ratio. The transition density for dust (H_2) decreases by 0.3 dex (0.4 dex) between the no and high diffusion coefficient cases.

We see that smaller grains shift the depletion-limited DTZ regime to lower densities, and boost the S/L grain ratio by ≈ 0.1 dex relative to Fiducial in $-2 \lesssim \log_{10} n_{\text{H}}/\text{cm}^{-3} \lesssim 2$. While the change in these DTZ -density relations is relatively small, we see that this acts in concert with the smaller grain size to produce a more marked shift in the H_2 transition (vertical line mark) to lower densities; a higher number of grains with higher surface-to-volume ratios at intermediate densities speeds up H_2 nucleation.

A2 Turbulent diffusion

Another process that can influence the evolution of the dust, and its associated influence on gas physics, is turbulent diffusion. As described in Section 3.6, the diffusion of dust follows that of gas-phase elements described in C. Correa et al. (in preparation). The amount of diffusion is modulated by a diffusion constant, C_d , where elements or grains are diffused more readily for higher values. We compare a number of variations with different C_d values, shown in Fig. A2.

These runs are shaded from dark blue to light yellow, in order of lowest to highest diffusion coefficient, where the effective C_d values are 0, 0.001, 0.01, and 0.1 respectively. We see that increased diffusion leads to a lower $\rho_{0.5, \text{dust}}$, pushing the saturated DTZ regime to lower densities. This is relatively intuitive; as growth by accretion is important in our model and is limited by the depletion in dense gas, diffusion of dust grains from dense gas into lower density gas with lower depletion allows this dust to grow further.

At ISM densities ($n_{\text{H}} \gtrsim 1 \text{cm}^{-3}$) we see the largest difference between the NoDiff and LoDiff runs. This shows the important role even low-level of diffusion can have in simulating dust growth;

diffusion allows transport of dust into gas particles that have not been directly enriched by stellar sources. As accretion relies on the dust–gas interaction, without these low seeding levels of grains, dust cannot grow to deplete elements at the levels we observe in galaxies. This demonstrates that while the dust mass in our simulated galaxies is much greater than directly seeded (see e.g. Fig. 5), the production and distribution of seed grains is crucial to manifest the dust content of galaxies. We also note that the HI-H_2 transition density $\rho_{0.5, \text{H}_2}$ differs most for the `NoDiff` run; with dust only residing in gas particles that have been directly enriched, the important role of dust in nucleating H molecules is lacking in neighbouring gas.

We see that the differences between the `Fiducial` and `HiDiff` appears greater at low density. We attribute this to the direct effect of diffusion. While at the lowest densities gas cannot accrete efficiently, the increased flux of dust into these particles boosts their \mathcal{DTZ} levels.

APPENDIX B: EXTINCTION EFFICIENCIES AND DUST COUPLING

In Section 3.7 we detail how we ignore the influence of extinction efficiency, $\langle Q_{\text{ext}} \rangle$, in deriving our dust size-dependent scaling factor s_{dust} , for dust-dependent cooling rates. We justify this by assuming $\langle Q_{\text{ext},S} \rangle / \langle Q_{\text{ext},L} \rangle \approx 1$.

In Fig. B1 we plot the $\langle Q_{\text{ext},S} \rangle / \langle Q_{\text{ext},L} \rangle$ ratio for silicate and carbonaceous grains from the B. T. Draine & H. M. Lee (1984) model as a function of the wavelength (λ) of incident light. We see that while this ratio is not constant, and can fall to values $\langle Q_{\text{ext},S} \rangle / \langle Q_{\text{ext},L} \rangle \ll 1$ for optical wavelengths ($\lambda \approx 0.4 \mu\text{m}$), it is closer to 1 for the Lyman-Werner and ionizing radiation bands ($\lambda \gtrsim 0.1 \mu\text{m}$) the extinction processes we consider important for




Figure B1. Ratio of extinction efficiencies for large and small grains as a function of the wavelength of incident light. The curves are derived for graphite and silicates taking the extinction ratio of large and small grains derived for a log-normal size distribution centred on the grain radius (0.1 and 0.01 μm , respectively) and with a standard deviation of 0.75 dex, using grain optical and calorimetric properties of B. T. Draine & H. M. Lee (1984). Relevant photon energy ranges are hatched as indicated. We see that for Lyman–Werner and ionizing photons, the extinction efficiency ratio is centred around $\langle Q_{\text{ext},S} \rangle / \langle Q_{\text{ext},L} \rangle \approx 1$ for both chemical species, suggesting we may ignore Q_{ext} for our purposes.

dust physics, varying ≈ 20 per cent about unity. These differences are small relative to the particle–particle differences in grain masses and species distributions, which justifies our simplifying assumption that extinction optical depths can be taken to scale with the surface-to-volume ratio of grains alone, and our use of the s_{dust} factor to scale dust rates (see equation 11).

APPENDIX C: DUST AND METALLICITY CALIBRATIONS

A curious tension in the development of these models has been the diverging conclusions on the abundances of dust in our simulated galaxies that can be drawn through comparison to different data sets. While comparisons like the cosmic dust density, galaxy dust mass functions, and dust scaling relations (see Figs 7, 8, and 9) may point to a slight surfeit of dust, \mathcal{DTG} and \mathcal{DTZ} values seem to imply a deficit of dust in the `Fiducial` run compared to observation (see e.g. Fig. 10).

One way to reconcile these results could be via the gas-phase metallicity calibrations used for the observations; while absolute dust masses can be measured in a relatively reliable way through black- or grey-body fits to the infrared emission of galaxies, absolute gas-phase metallicity calibrations can be uncertain, with ≈ 0.5 dex differences between average measured Z_{gas} for a given M_* (see the distinct branches in the literature gas-phase mass–metallicity relations, e.g. L. J. Kewley & S. L. Ellison 2008).

In Fig. C1 we make use of the *Dustpedia* data presented in P. De Vis et al. (2019), which usefully provides a number of different




Figure C1. The gas-phase oxygen abundances as a function of stellar mass of galaxies. We plot the `Fiducial` run results, where oxygen abundances are measured in cold, dense gas ($\log_{10} T/\text{K} < 4.5$, $\log_{10} n_{\text{H}}/\text{cm}^{-3} > -1$) and within a 50 kpc aperture about the subhalo potential minima. *Lines* show median values while *data points* show individual galaxies. *Shaded regions* indicate 16–84th percentile ranges. We compare the `Fiducial` run with the *Dustpedia* results of P. De Vis et al. (2019) for different observational calibrations; the [O II]-exclusive calibration of L. S. Pilyugin & E. K. Grebel (2016, PS16S) and that of C. A. Tremonti et al. (2004, T04). We see the simulation data overlaps best with T04, particularly at intermediate metallicities.




Figure C2. The dust-to-gas (DTG) ratio of galaxies as a function of their gas-phase oxygen abundances, measured in cold, dense gas, with the same galaxy selection and O/H aggregation as described in Fig. 10. Lines show median values while squares show individual galaxies. The solid line shows the median for galaxies with $M_* > 10^9 M_\odot$ and 16–84th percentile ranges indicated by the shaded region. The dotted line shows the median relation for all processed galaxies. Top panel shows comparison to the P. De Vis et al. (2019) data calibrated using the [O II]-exclusive calibration of L. S. Pilyugin & E. K. Grebel (2016, PS16S). Bottom panel compares instead of the P. De Vis et al. (2019) data calibrated using the C. A. Tremonti et al. (2004, T04) approach. We see that the Fiducial galaxies better reproduce the T04 calibration in this instance.

inferred metallicity values (in the form of oxygen abundance) based on differing indices and calibrations. In particular, we plot the [O II]-exclusive calibration of L. S. Pilyugin & E. K. Grebel (2016) (taken as the default in P. De Vis et al. 2019) and the calibration of C. A. Tremonti et al. (2004) (T04). We see a ≈ 0.5 dex offset between the two median relations, with the T04 calibration comparing best to our simulation for $\log_{10}(M_*/M_\odot) \lesssim 10.25$.

In Fig. C2, we then plot the DTG values of Fiducial simulation galaxies as a function of their gas-phase oxygen abundances, against the P. De Vis et al. (2019) data for these two calibrations. We see that the simulated data agrees much better with the DTG values when using the C. A. Tremonti et al. (2004)-calibrated metallicities. This demonstrates how our modelling shows a preference for certain metallicity calibrations.

Given that dust in the ISM of our simulated galaxies is set largely by the balance between accretion and the (strong) depletion of metals (Fig. 5), with little headroom to deplete more given viable dust grain chemistries (Fig. 4 and associated discussion), our modelling suggests dust masses are a feasible means to help break the degeneracy in absolute metallicity calibrations.

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