



A scalable, GPU-portable super-droplet method for mixed-phase cold microphysics in ERF

Debojyoti Ghosh¹, Hassan Beydoun², Emily K. de Jong², Coleman J. Kendrick¹, Dana L. McGuffin^{2,a}, David J. Gardner¹, and Katherine A. Lundquist²

¹Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, 7000 East Ave., Livermore, CA 94550, United States

²Atmospheric, Earth, and Energy Division, Lawrence Livermore National Laboratory, 7000 East Ave., Livermore, CA 94550, United States

^anow at: Reflective, United States

Correspondence: Debojyoti Ghosh (ghosh5@llnl.gov)

Abstract. Cloud microphysics controls precipitation, latent heating, and the radiative properties of clouds, and it remains a source of complexity in atmospheric models. The super-droplet method is a Lagrangian, particle-based approach that evolves per-particle attributes and makes no assumptions about the shape of the size distribution. Its fidelity comes at a higher computational cost than that of bulk schemes. We present a scalable, portable implementation of the super-droplet method for cold, mixed-phase microphysics in the Energy Research and Forecasting (ERF) code. It includes ice processes such as freezing, deposition, sublimation, riming, aggregation, and a gradual heat-balance melting that retains the mixed-phase state of each particle. The implementation is portable by construction across NVIDIA, AMD, and Intel graphics processing units through the AMReX framework, and we demonstrate it on NVIDIA and AMD hardware. We verify the method against theory and published benchmarks. We then simulate a three-dimensional deep-convection case based on the GARP Atlantic Tropical Experiment (GATE), where the method resolves the ice-particle morphology and the phase structure of the mature cloud without predefined hydrometeor categories. The method scales efficiently to 512 graphics processing units on the deep-convection case with more than a billion super-droplets, and we report and compare the computational cost of the microphysical processes.

1 Introduction

Cloud microphysics governs the dynamics of cloud and precipitating particles. It controls surface precipitation, the release of latent heat, and the radiative properties of clouds. Atmospheric flow simulations therefore couple a fluid dynamical core to a microphysics model. Numerical weather prediction and cloud-resolving codes such as WRF, MPAS, ICON, and SAM all include such models (Powers et al., 2017; Skamarock et al., 2012; Zängl et al., 2015; Khairoutdinov and Randall, 2003) and often use one of two approaches. Bulk schemes predict a small number of moments of an assumed particle size distribution (Kessler, 1969; Lin et al., 1983; Seifert and Beheng, 2006; Morrison et al., 2009; Thompson et al., 2008). They are inexpensive and robust, which makes them the operational standard. Their accuracy is limited by the assumed distribution shape and by empirical conversions between predefined hydrometeor categories. The P3 scheme (Morrison and Milbrandt, 2015) relaxes the fixed ice



categories by predicting continuous bulk ice properties, though the prediction remains per moment rather than per particle. Bin schemes instead discretize the size distribution into tens of bins and explicitly integrate the equations governing its evolution, condensational growth and the stochastic collection equation (Berry, 1967; Tzivion et al., 1987; Khain et al., 2015). They re-
 25 solve spectral structure that bulk schemes cannot. They are, however, approximately five to twenty times as expensive as bulk schemes (Khain et al., 2015). The most detailed models may be one to two orders of magnitude more expensive (Morrison et al., 2020). They also broaden the size distribution artificially through numerical diffusion in the discretized size space (Khain et al., 2015; Morrison et al., 2020).

The super-droplet method (SDM) is a Lagrangian, particle-based alternative to these Eulerian paradigms. It belongs to the
 30 broader family of super-particle methods, in which one computational particle represents many physical particles. Such methods are long established across the physical sciences: the particle-in-cell method of plasma physics (Hockney and Eastwood, 1988; Birdsall and Langdon, 1991; Dawson, 1983), the direct simulation Monte-Carlo method of rarefied gas dynamics (Bird, 1994), the modeling of liquid sprays (Dukowicz, 1980), and the tracking of entrained sediment in rotorcraft brownout (Govindarajan et al., 2013). Their per-particle work maps naturally onto graphics processing units (GPUs), and particle-in-cell and
 35 direct simulation Monte-Carlo codes have reported strong GPU performance and scaling (Bussmann et al., 2013; Su et al., 2012). Shima et al. (2009) introduced the method for cloud microphysics, in which a finite ensemble of computational particles represents the physical particle population. Each super-droplet carries a multiplicity and a set of attributes, and evolves along flow trajectories. The method makes no assumption about the size-distribution shape. It avoids the size-space numerical diffusion of bin schemes, since the particles are not advected through a fixed size grid. A warm-rain intercomparison shows
 40 that the Lagrangian schemes agree with one another, while the bin schemes broaden the distribution nonphysically (Hill et al., 2023). It also carries per-particle attributes, such as mass and ice habit. Thus, category conversion is not empirical like in the bulk schemes (Morrison et al., 2020; Grabowski et al., 2019).

The SDM has since been developed and applied in several codes. Arabas and Shima (2013) simulated precipitating trade-wind cumuli and compared the results with aircraft measurements. Independent Lagrangian cloud models were developed
 45 by Riechermann et al. (2012) and Andrejczuk et al. (2010). The collision algorithm and its convergence have been studied in detail (Unterstrasser et al., 2017; Dziekan and Pawlowska, 2017; Morrison et al., 2024). Sato et al. (2018) compared the numerical convergence of the particle-based and Eulerian microphysics for a shallow-convection cloud field. Shima et al. (2020) extended the method to mixed-phase clouds, predicting ice particle morphology without predefined categories or mass–dimension relationships. Related Lagrangian ice models include McSnow (Brdar and Seifert, 2018) and the Sölch and Kärcher
 50 (2010) cirrus model. The principal drawback of the SDM is sampling noise. A sufficient number of super-droplets per grid cell is required to control it (Shima et al., 2009; Dziekan and Pawlowska, 2017).

The higher fidelity of the SDM entails an increase in computational cost. The dominant operations are performed per particle: phase changes, stochastic collisional processes, and advection with sedimentation. The collision step is a Monte-Carlo procedure over particle pairs within each grid cell (Shima et al., 2009). Shima et al. (2020) reported that a mixed-phase
 55 method, applied to a deep-convective cloud with 128 super-droplets per grid cell, was approximately thirty times as expensive as a two-moment bulk scheme. Several implementations have addressed this cost on high-performance computing systems.



SCALE-SDM is written in Fortran and parallelized with the Message Passing Interface (MPI). It has been scaled to 36,864 nodes (147,456 MPI processes, 4 per node) of the Fugaku system (Matsushima et al., 2023). The `libcloudph++` library offloads the particle computations to a GPU through the Thrust library (Arabas et al., 2015), and is used by the UWLCM large-eddy simulation model (Dziekan et al., 2019; Dziekan and Zmijewski, 2022). `PySDM` provides a Python implementation with multi-core and GPU backends (Bartman et al., 2022), and `CLEO` adopts the Kokkos performance-portability library (Bayley et al., 2026a). These GPU implementations target the hardware of a single vendor, with the exception of the Kokkos-based `CLEO`. The computational cost of the individual cold-microphysical processes, and the strong scaling of the coupled method on these accelerators, remain little studied.

This paper presents a scalable and portable implementation of the SDM for cold microphysics in the Energy Research and Forecasting (ERF) model. ERF is a regional atmospheric modeling code for mesoscale and microscale flows (Lattanzi et al., 2025; Almgren et al., 2023). It reproduces the dynamics of the Weather Research and Forecasting model, and it provides a range of Eulerian bulk microphysics schemes (Sect. 2). ERF is built on the AMReX framework, which provides a performance-portable programming model (Zhang et al., 2021). The same source compiles for multi-core central processing units and for GPUs from NVIDIA, AMD, and Intel, through OpenMP, CUDA, HIP, and SYCL respectively. ERF therefore combines regional atmospheric dynamics with multi-vendor GPU portability in a single code. The SDM in ERF has been applied to warm-rain microphysics (Beydoun et al., 2026). This paper extends it to the cold, mixed-phase processes and characterizes the computational cost and the scalability of the coupled method.

The primary contributions of this paper are a cold, mixed-phase super-droplet scheme with a gradual melting model, and a scalable, portable implementation of it. The melting model carries a per-super-droplet mixed-phase state and follows the heat-balance melting (Rasmussen and Heymsfield, 1987; Phillips et al., 2007), in contrast to the instantaneous melting of earlier super-droplet ice schemes (Shima et al., 2020). The implementation is based on AMReX data structures and can be built without modification on multiple architectures. We implement the ice microphysical processes: immersion freezing, deposition, sublimation, riming, aggregation, and melting. The ice particles are represented as porous spheroids following Shima et al. (2020). Processes not represented here include collisional raindrop breakup (de Jong et al., 2023), secondary ice production by rime splintering (Hallett and Mossop, 1974; Field et al., 2017), and the shedding of meltwater from large melting particles (Rasmussen and Heymsfield, 1987). We describe the algorithm of each process, and we verify the implementation against theory and published benchmarks. We simulate a three-dimensional deep convection, examining the predicted ice-particle morphology, the phase structure of the mature cloud, and the run-to-run variability of the stochastic method. Finally, we analyze the memory requirements and the computational cost of each microphysical process and assess the strong scaling of the coupled method to 512 GPUs.

The outline of the paper is as follows. Section 2 summarizes the ERF dynamical core. Section 3 describes the super-droplet method for cold microphysics and its implementation in ERF. Section 4 reports the verification and the three-dimensional deep-convection case. Section 5 analyzes the memory requirement, the computational cost of each microphysical process, and the strong scaling of the coupled method. Conclusions are drawn in Sect. 6.



2 ERF dynamical core

The super-droplet method is implemented in the open-source Energy Research and Forecasting (ERF) code (Lattanzi et al., 2025; Almgren et al., 2023; ERF Development Team, 2025), which solves the compressible, non-hydrostatic equations for moist flow on regional-scale domains. It reproduces the dynamics of the Weather Research and Forecasting model (Powers et al., 2017) and is built on the AMReX framework (Zhang et al., 2021). We summarize here the governing equations and their discretization, since the microphysics of Sect. 3 supplies several of their source terms. We describe the fully compressible formulation; an anelastic option is also available (Lattanzi et al., 2025). The prognostic variables are the dry-air density ρ_d , the momentum $\rho_d \mathbf{u}$, and the density-weighted dry potential temperature $\rho_d \theta_d$. They evolve as:

$$\frac{\partial \rho_d}{\partial t} = -\nabla \cdot (\rho_d \mathbf{u}), \quad (1)$$

$$100 \quad \frac{\partial (\rho_d \mathbf{u})}{\partial t} = -\nabla \cdot (\rho_d \mathbf{u} \mathbf{u}) - \frac{1}{1 + q_t} (\nabla p' - \mathbf{B}) - \nabla \cdot \boldsymbol{\tau} + \mathbf{F}_u, \quad (2)$$

$$\frac{\partial (\rho_d \theta_d)}{\partial t} = -\nabla \cdot (\rho_d \mathbf{u} \theta_d) + \nabla \cdot (\rho_d \alpha_\theta \nabla \theta_d) + F_\theta + H_n + H_p, \quad (3)$$

where $\mathbf{u}$ is the velocity, p' is the pressure perturbation from a hydrostatic reference state $p_0(z)$, and q_t is the total water mixing ratio per unit dry-air mass. The term $\mathbf{B} = (\rho - \rho_0) \mathbf{g}$ is the buoyancy, with ρ the total density, $\rho_0(z)$ the reference density, and $\mathbf{g} = (0, 0, -g)$ the gravity vector, so air denser than the reference is forced downward. The factor $1/(1 + q_t)$ accounts for water content. The viscous and subfilter stress is $\boldsymbol{\tau}$, α_θ is the thermal diffusivity, and $\mathbf{F}_u$ and F_θ collect the remaining forcings, including the Coriolis term. The quantities H_n and H_p are the latent heating rates from non-precipitating and precipitating phase changes, and are computed by the microphysics model.

The total water mixing ratio is the sum of the vapor and the condensate species,

$$q_t = q_v + q_c + q_r + q_i + q_s + q_g, \quad (4)$$

110 where q_v is water vapor, q_c and q_r are cloud water and rain, and q_i , q_s , and q_g are cloud ice, snow, and graupel, respectively, each per unit dry-air mass. The super-droplet method models the condensate mixing ratios $(q_c, q_r, q_i, q_s, q_g)$ with the particles. The dynamical core transports the water-vapor mixing ratio q_v , which satisfies

$$\frac{\partial (\rho_d q_v)}{\partial t} = -\nabla \cdot (\rho_d \mathbf{u} q_v) + \nabla \cdot (\rho_d \alpha_q \nabla q_v) + F_v, \quad (5)$$

115 where α_q is the moisture diffusivity and F_v is the vapor source from the microphysics. The super-droplet method evaluates F_v , together with the latent heating H_n and H_p , as described in Sect. 3. The bulk schemes available in ERF advance the set of condensate mixing ratios by similar transport equations with a sedimentation flux for the precipitating species. The system is closed by the equation of state

$$p = P_{00} \left(\frac{R_d \rho_d \theta_d M(q_v)}{P_{00}} \right)^\gamma, \quad M(q_v) = 1 + \frac{R_v}{R_d} q_v, \quad (6)$$

120 where R_d and R_v are the gas constants of dry air and water vapor, c_p is the dry-air specific heat at constant pressure, $\gamma = c_p/(c_p - R_d)$, and $P_{00} = 10^5$ Pa is a reference pressure. The factor $M(q_v)$ introduces the vapor loading, and the dry potential



temperature θ_d is the prognostic thermodynamic variable. The temperature is recovered as $T = \theta_d(p/P_{00})^{R_d/c_p}$. The pressure and density perturbations p' and $\rho - \rho_0$ in Eq. (2) are measured from a hydrostatically balanced reference state $p_0(z)$, $\rho_0(z)$ that satisfies $dp_0/dz = -\rho_0 g$. This reference state sets the buoyancy **B**.

The equations are discretized on a staggered Arakawa C-grid (Arakawa and Lamb, 1977). Scalars (ρ_d , θ_d , and the mixing ratios) are stored at cell centers, while momentum components are stored on the corresponding cell faces. Terrain is represented through a height-following coordinate with map factors for lateral grid stretching. Advective fluxes are computed with centered schemes of second through sixth order, odd-ordered upwind schemes, or weighted essentially non-oscillatory (WENO) schemes (Jiang and Shu, 1996; Balsara and Shu, 2000) of third, fifth, and seventh order. The third-order WENO scheme is the default for microphysics variables, since it limits the spurious overshoots that would otherwise produce negative mixing ratios. The fully compressible equations are advanced with the three-stage Wicker and Skamarock (2002) Runge-Kutta method. The fast acoustic and gravity-wave modes are integrated within each stage on shorter acoustic substeps (Klemp et al., 2007): the horizontal momentum is updated explicitly, while the vertical momentum, density, and potential temperature are advanced implicitly through a tridiagonal solve. The solve is off-centered toward the implicit side, which damps the vertically propagating acoustic modes, and a Rayleigh layer near the model top absorbs the remainder. The moisture scalars are advected once per Runge-Kutta stage; the microphysics, by contrast, is applied once per full time step, after the three stages complete. The dynamics and the microphysics are thus coupled by a first-order (Lie) operator splitting, in which the microphysics advances the state produced by the full dynamical step.

The subfilter stress τ is supplied by a turbulence closure. ERF provides the constant-coefficient Smagorinsky (1963) eddy-viscosity model and the prognostic 1.5-order Deardorff (1980) turbulent kinetic energy model for large-eddy simulation, together with planetary-boundary-layer closures for coarser mesoscale grids. A constant molecular viscosity may be added for direct numerical simulation. Aside from SDM, several bulk microphysical schemes are implemented in ERF; these include the warm-rain Kessler (1969) scheme, the single-moment SAM scheme (Khairoutdinov and Randall, 2003), the WSM6 scheme (Hong and Lim, 2006), and the two-moment Morrison et al. (2009) scheme. These models advance the condensate and return the moisture sources, the fall speed, and the latent heating, and they update $\rho_d \theta_d$ together with the moisture scalars at each step. The boundary conditions depend on the case. ERF supports periodic, inflow, outflow, and externally specified lateral boundaries. It applies a Monin–Obukhov surface layer or a wall at the bottom, and a slip-wall or specified condition at the top, optionally with a Rayleigh-damping layer that absorbs vertically propagating waves.

3 Super-droplet method for cold microphysics

Microphysical processes are driven by the hydrometeor population: the aerosol particles that serve as cloud condensation and ice nuclei, the cloud droplets and raindrops of the liquid phase, and the cloud ice, snow, and graupel of the ice phase. Cloud droplets are small liquid drops suspended in the flow, whereas raindrops are larger and sediment appreciably. Among the frozen hydrometeors, cloud ice consists of pristine crystals, snow of crystal aggregates, and graupel of heavily rimed particles. We also include mixed liquid–ice particles that are partially melted. Ice is represented as a porous spheroid (Shima et al., 2020).



Therefore, its habit (the crystal shape, here the aspect ratio of the spheroid), the apparent density, and the fall speed of a particle
155 follow from its geometry and mass rather than from prescribed mass–dimension relationships.

The SDM represents this population by a finite ensemble of super-droplets that are computational particles representing many physical particles sharing common attributes (Shima et al., 2009). A super-droplet i represents ξ_i identical physical particles, where ξ_i is the multiplicity. Along with its multiplicity, position $\mathbf{x}_i$, and velocity $\mathbf{v}_i$, each super-droplet carries:

- the equivalent radius R_i of the liquid water, equivalently the water mass $m_{w,i} = \frac{4}{3}\pi\rho_w R_i^3$, with ρ_w the density of water;
- 160 – the soluble and insoluble masses of an arbitrary number of aerosol species, $m_{a,i}^{(1)}, m_{a,i}^{(2)}, \dots$, which serve as cloud condensation and ice nuclei;
- the characteristic freezing temperature $T_{fz,i}$, the threshold below which the super-droplet freezes by immersion;
- the ice mass $m_{i,i}$ and its geometry: the equatorial and polar semi-axes a_i and c_i of the porous spheroid, the accreted rime mass $m_{\text{rime},i}$, and the number of primary crystals, or monomers, $n_{\text{mono},i}$.

165 Multiplicity is the only computational attribute; all other attributes are physical characteristics of the particle. Liquid water and ice may be present simultaneously on the same super-droplet, so its phase is liquid, ice, or mixed. The melt fraction is defined as $\ell = m_w / (m_w + m_i)$, the fraction of its condensate that is liquid, where m_w and m_i are the liquid-water and ice masses of the super-droplet. The aerosol masses determine both the super-droplet’s activation as a droplet and its freezing as an ice particle. Each microphysical step advances the particle population through phase change, then advection with sedimentation, then
170 coalescence, in the order described below. The derived mixing ratios and latent heat then inform the evolution of the dynamical core. Super-droplets are created at initialization and, optionally, new ones enter through injection during the simulation.

3.1 Initialization

Super-droplets are created within one or more regions of the domain, each specified as a rectangular box or an ellipsoidal bubble. A fixed number N_{sd} of super-droplets is placed at random positions in every grid cell that the region intersects,
175 drawing from a target size distribution and a target physical number density n_{phys} . The physical count in a cell, $n_{\text{phys}} \Delta V$ with ΔV the cell volume, is shared among the N_{sd} super-droplets of the cell, indexed $k = 1, \dots, N_{\text{sd}}$, in one of two ways. The constant-multiplicity mode draws the size r_k of each super-droplet from the target distribution by inverting its cumulative distribution, and assigns every super-droplet the same multiplicity,

$$\xi_k = \left\lceil \frac{n_{\text{phys}} \Delta V}{N_{\text{sd}}} \right\rceil, \quad (7)$$

180 so the distribution is represented by the spread of sizes. The sampled-multiplicity mode used in the simulations here instead draws the sizes uniformly in $\ln r$, and weights each super-droplet by the size-distribution density f at its size,

$$\xi_k = \frac{n_{\text{phys}} \Delta V f(r_k)}{\sum_{k'=1}^{N_{\text{sd}}} f(r_{k'})}, \quad (8)$$

so the multiplicities carry the distribution and sum to $n_{\text{phys}} \Delta V$ over the cell. The sampled mode populates the full size range, including the sparsely populated tail, at the cost of a wider spread of multiplicities.



185 Either mode draws the sizes from a target size distribution, chosen independently for each water and aerosol species. The distribution is a constant, an exponential (specified by a mean and a minimum mass), or a log-normal (specified by a geometric mean radius and a geometric standard deviation). The equivalent radius of a new droplet follows from its water, soluble, and insoluble masses,

$$r_{\text{eff}} = \left(\frac{m_w + m_s + (\rho_w / \rho_{\text{ins}}) m_{\text{ins}}}{\frac{4}{3} \pi \rho_w} \right)^{1/3}, \quad (9)$$

190 where m_s and m_{ins} are the soluble and insoluble masses and ρ_{ins} is the density of the insoluble core. Each new super-droplet is also assigned its characteristic freezing temperature T_{fz} , drawn from the ice-nucleating active sites over the insoluble surface of its aerosol (Sect. 3.4). The particles begin unfrozen and at rest, so their ice mass, geometry, and velocities are initialized to 0. The number of super-droplets may be scaled with the local air density, so that the physical concentration decreases with height. Regions with distinct distributions may be superposed to build layered or multi-cloud initial states.

195 3.2 Advection




Figure 1. Fall speed of hydrometeors from the ERF implementation. (a) Terminal velocity versus maximum dimension D at air density $\rho_a = 1.1 \text{ kg m}^{-3}$, against the Böhm (1989) reference values, plotted as open markers matching the color and shape of each habit's curve; the three graupel habits use the Locatelli and Hobbs (1974) mass-dimension densities, and the dashed line is the Stokes law for a 100 kg m^{-3} ice sphere. (b) Fall speed of two melting particles of equal mass (equivalent radius 2 mm), a dense graupel ($\rho_i = 700 \text{ kg m}^{-3}$) and a low-density snow particle ($\rho_i = 100 \text{ kg m}^{-3}$), against the melt fraction ℓ , with dotted lines at each dry-ice speed v_{∞}^{ice} and the shared raindrop speed v_{∞}^{rain} .



Each super-droplet is assumed to move at the local fluid velocity and its terminal fall speed. Its position evolves by

$$\frac{d\mathbf{x}_i}{dt} = \mathbf{v}_i, \quad \mathbf{v}_i = \mathbf{U}(\mathbf{x}_i) - v_{\infty,i} \hat{\mathbf{z}}, \quad (10)$$

where $\mathbf{U}(\mathbf{x}_i)$ is the fluid velocity at the particle, $v_{\infty,i}$ is the terminal velocity of super-droplet i , and $\hat{\mathbf{z}}$ is the upward unit vector. The fluid velocity is interpolated to the particle from the staggered, face-centered velocities of the C-grid. The equation
 200 is integrated with the forward Euler method. The terminal velocity is calculated based on the size and type of the particle. The smallest cloud droplets ($d < 19 \mu\text{m}$) follow the analytic Stokes limit,

$$v_{\infty} = \frac{(\rho_w - \rho) g d^2}{18 \mu} \left(1 + 2.51 \frac{\lambda}{d} \right), \quad (11)$$

where ρ and ρ_w are the densities of air and water, μ is the dynamic viscosity of air, g is the acceleration due to gravity, and d is the diameter. The final factor is the Cunningham slip correction, with λ the mean free path of air.

205 The terminal velocity of all other particles is calculated from the Reynolds number N_{Re} ,

$$v_{\infty} = \frac{\mu N_{\text{Re}}}{\rho d}, \quad (12)$$

which is obtained from the particle mass and shape and the fluid properties through a fall-speed-free number: the Davies or Bond number for liquid water, and the Best number for ice. Liquid water has $d = 2R$, and its Reynolds number follows from the three-regime Beard (1976) parameterization, with the Stokes limit above as the first regime. It is calculated for the other
 210 two regimes as

$$N_{\text{Re}} = \begin{cases} c_{\text{sc}} \exp[f_b(\ln N_{\text{Da}})], & 19 \mu\text{m} \leq d < 1.07 \text{ mm}, \\ N_P^{1/6} \exp\left[f_c\left(\ln(N_{\text{Bo}} N_P^{1/6})\right)\right], & d \geq 1.07 \text{ mm}, \end{cases} \quad (13)$$

where $c_{\text{sc}} = 1 + 2.51 \lambda/d$ is the slip factor, and $N_P = \rho^2 \sigma^3 / [(\rho_w - \rho) g \mu^4]$ is a physical-property group: a dimensionless, size-independent combination of the air density, surface tension, viscosity, and gravity that sets the scale of the large-drop regime. The functions $f_b(x) = \sum_{k=0}^6 b_k x^k$ and $f_c(x) = \sum_{k=0}^5 c_k x^k$ are the Beard (1976) polynomial fits, with

$$\begin{aligned} (b_0, \dots, b_6) &= (-3.18657, 0.992696, -1.53193 \times 10^{-3}, -9.87059 \times 10^{-4}, \\ &\quad -5.78878 \times 10^{-4}, 8.55176 \times 10^{-5}, -3.27815 \times 10^{-6}), \\ (c_0, \dots, c_5) &= (-5.00015, 5.23778, -2.04914, 0.475294, \\ &\quad -5.42819 \times 10^{-2}, 2.38449 \times 10^{-3}). \end{aligned} \quad (14)$$

The Davies and Bond numbers are

$$N_{\text{Da}} = \frac{4\rho(\rho_w - \rho)g d^3}{3\mu^2}, \quad N_{\text{Bo}} = \frac{4(\rho_w - \rho)g d^2}{3\sigma}, \quad (15)$$

respectively, where σ is the surface tension of water. We cap the radius used to calculate the terminal velocity at 4 mm (Gunn and Kinzer, 1949).



220 Ice follows the Böhm (1989) drag relation, in the Shima et al. (2020) porous-spheroid form with $d = 2a$. The Best number is calculated from the particle mass and shape,

$$N_{\text{Be}} = \frac{8mg\rho}{\pi\mu^2 \max(\phi, 1) \max(q^{1/4}, q)}, \quad (16)$$

where m is the particle mass, $\phi = c/a$ is the aspect ratio, and $q = (\rho_p/\rho_{\text{ice}})^{\exp(-\phi)}$ is the area ratio, with ρ_p the apparent particle density and ρ_{ice} the density of solid ice. It is corrected at large values to

$$225 \quad N'_{\text{Be}} = N_{\text{Be}} \frac{1 + (N_{\text{Be}}/N_0)^2}{1 + 1.6(N_{\text{Be}}/N_0)^2}, \quad N_0 = 2.8 \times 10^6. \quad (17)$$

The Reynolds number then follows from the explicit inverse of this drag law,

$$N_{\text{Re}} = \frac{6k}{C_{\text{DP}}} \beta^2 \left[1 + \frac{2\beta e^{-\beta\gamma}}{(2+\beta)(1+\beta)} \right], \quad \beta = \left(1 + \frac{C_{\text{DP}}}{6k} \sqrt{\frac{N'_{\text{Be}}}{C_{\text{DP}}}} \right)^{1/2} - 1, \quad (18)$$

with $\gamma = (C_{\text{DO}} - C_{\text{DP}})/(4C_{\text{DP}})$. The three drag coefficients depend on the aspect ratio ϕ and the area ratio q ,

$$\begin{aligned} k &= \min \left[\max(0.82 + 0.18\phi, 0.85), \min(0.37 + 0.63\phi^{-1/2}, \frac{1.33}{\max(\ln \phi, 0) + 1.19}) \right], \\ C_{\text{DP}} &= \max[1, q(1.46q - 0.46)] \max(0.292k\gamma_\phi, 0.492 - 0.2\phi^{-1/2}), \\ C_{\text{DO}} &= 4.5k^2 \max(\phi, 1), \end{aligned} \quad (19)$$

230 where $\gamma_\phi = \min[1.98, \max(1, 3.76 - 8.41\phi + 9.18\phi^2 - 3.53\phi^3)]$. The inversion is explicit and valid across all flow regimes, so no iteration is needed, and the terminal velocity follows from Eq. (12). The dependence on m , ϕ , and q captures the slower fall of non-spherical, low-density ice. Figure 1(a) compares these fall speeds with the Böhm (1989) reference values across the size range of atmospheric hydrometeors, as a function of the maximum particle dimension $D = 2\max(a, c)$. The three graupel habits reproduce those values, and the small-ice limit recovers the analytic Stokes law.

235 A mixed, partially melted particle falls faster than its dry-ice frame but slower than the equivalent raindrop. We interpolate its terminal velocity between these two limits by the melt fraction ℓ , following Frick et al. (2013),

$$v_\infty = v_\infty^{\text{ice}} + (v_\infty^{\text{rain}} - v_\infty^{\text{ice}}) \Psi(\ell), \quad \Psi(\ell) = 0.246\ell + 0.754\ell^7, \quad (20)$$

where v_∞^{ice} is the Böhm fall speed evaluated for the ice mass alone, and v_∞^{rain} is the Beard (1976) fall speed of a raindrop of the combined mass $m_w + m_i$. The weight $\Psi(\ell)$ is an empirical fit to the Mitra et al. (1990) wind-tunnel data; it rises from 0 for dry ice to 1 once the core has melted. Excluding the liquid mass keeps the mass consistent with the drag area of the spheroid as the core melts. Figure 1(b) shows the blended fall speed for a dense graupel and a low-density snow particle. The two particles have equal mass and melt to the same raindrop, but their dry-ice fall speeds differ. The dense graupel begins near its raindrop speed. As its ice core melts, the core shrinks and its fall speed decreases. Early in melting, this decrease exceeds the blend toward the raindrop, so the mixed particle has a lower fall speed than the dry graupel. The low-density particle begins well below its raindrop speed. The blend toward the faster raindrop then dominates, and its fall speed increases monotonically.

245



3.3 Boundary treatment

A super-droplet that reaches a domain boundary as it advects is treated based on the boundary type. A particle that reaches the bottom surface is considered precipitated. Its velocity is set to 0, and it is deactivated. Unless it is being recycled or saved for diagnostics, it is removed from the simulation. A particle carried above the top boundary is removed in the same way, and a
 250 particle that reaches a lateral wall is again deactivated and removed. A periodic lateral boundary instead returns the particle through the opposite face with its attributes intact. An open inflow or outflow boundary also returns the particle through the opposite face, but resets it to a dry aerosol. Its water and ice are cleared, its multiplicity is restored to the initialization value, and its radius is recomputed from the residual aerosol. The precipitation reaching the ground is computed and saved before the particle removal. The downward super-droplet mass flux through the surface cells is accumulated in time at each step. The
 255 accumulation over a step is

$$\Delta R_s = \frac{\max(0, -F_s^z) \Delta t}{\rho_s}, \quad \Delta M_s = \max(0, -F_s^z) \Delta t \Delta x \Delta y, \quad (21)$$

where F_s^z is the vertical mass flux of species s through the lowest cell face, ρ_s is the density, and $\Delta x \Delta y$ is the horizontal cell area. F_s^z is computed by particle-to-mesh deposition (Sect. 3.6): each super-droplet contributes $\xi_i m_{s,i} [U_z(\mathbf{x}_i) - v_{\infty,i}]$, the product of its represented species mass and the vertical component of its velocity $\mathbf{v}_i$ from Eq. (10), where $U_z(\mathbf{x}_i)$ is the fluid
 260 vertical velocity interpolated to the particle. The contributions are deposited to the cell centers with the trilinear weights and divided by the cell volume. The depth R_s , converted to mm, gives the surface rain accumulation; for the aerosols, the mass M_s (kg) is accumulated on each surface cell.

Two optional mechanisms manage the super-droplet population. These are described for completeness but are not used in the cases presented in this paper. Recycling reuses deactivated particles once their fraction crosses a threshold: each is reset
 265 to a dry aerosol with a resampled aerosol mass and multiplicity, placed at random within a prescribed region, and reactivated. This keeps the particle count approximately constant instead of letting it decay, which suits a statistically steady run with a precipitation. Injection adds new super-droplets at a prescribed rate over a time window and region, representing a continuous source, with fractional per-step counts accumulated until a whole super-droplet is added.

3.4 Phase change

270 Phase change transfers mass between the vapor, liquid, and ice held by a super-droplet, and exchanges latent heat with the surrounding air. The ambient temperature, pressure, and saturation are interpolated to the particle from the grid. The three transitions are applied in sequence: freezing and melting between solid and liquid, then condensation and evaporation between liquid and vapor, then deposition and sublimation between solid and vapor. This ordering allows a droplet that has just frozen to grow by vapor deposition within the same step. Two quantities recur in the substeps below: the ventilation factor f_v , which
 275 enhances the vapor and heat exchange for a falling particle, and, for the ice phases, the electrostatic capacitance C of the



spheroid, which generalizes the radius of the spherical case,

$$C = \begin{cases} \frac{\varepsilon c}{\ln[(1+\varepsilon)\phi]}, & \phi > 1, \quad \varepsilon = \sqrt{1-\phi^{-2}}, \\ \frac{\varepsilon a}{\arcsin \varepsilon}, & \phi < 1, \quad \varepsilon = \sqrt{1-\phi^2}, \end{cases} \quad (22)$$

where $\phi = c/a$ is the aspect ratio and ε the eccentricity; the capacitance reduces to the radius a for a sphere ($\phi \rightarrow 1$).

3.4.1 Freezing

280 Liquid freezes by immersion through the Niemand et al. (2012) singular ice-nucleating active site (INAS) description. The active-site surface density is

$$n_s(T) = a_0 \exp(-a_1 T_c + a_2), \quad (23)$$

where $T_c = T - 273.15 \text{ K}$ is the temperature relative to the melting point and $(a_0, a_1, a_2) = (1 \text{ m}^{-2}, 0.517 \text{ K}^{-1}, 8.934)$. The characteristic freezing temperature T_{fz} (Sect. 3.1) follows from inverting the cumulative freezing probability $P =$
 285 $1 - \exp(-A n_s)$ for the insoluble surface area A of its immersed particle, so a larger or more active nucleus freezes at a warmer temperature. Heterogeneous freezing is bounded between -12°C and -36°C in the parameterization, and a super-droplet that exposes no active site (soluble aerosols) freezes homogeneously at -38°C , reported as 235 K by Koop et al. (2000). When the local temperature falls to T_{fz} and the air is supersaturated with respect to water, the liquid converts to a spherical ice core with a single monomer and no rime. The supersaturation requirement is inherited from the condensa-
 290 tion/immersion freezing scheme in Shima et al. (2020), who note that it can be restrictive for the immersion freezing of large supercooled droplets in subsaturated air.

3.4.2 Melting

Ice melts gradually, by a time-dependent heat balance (Rasmussen and Heymsfield, 1987; Phillips et al., 2007), in the form given by Seifert and Beheng (2006). The melt rate of the ice core is

$$295 \quad \frac{dm_i}{dt} = -\frac{4\pi C}{L_f} \left[K f_v (T - T_0) + \frac{L_v D}{R_v} f_v \left(\frac{e_\infty}{T} - \frac{e_{s,w}(T_0)}{T_0} \right) \right], \quad (24)$$

where L_f and L_v are the latent heats of fusion and vaporization, K is the thermal conductivity of air, D is the vapor diffusivity, R_v is the vapor gas constant, $T_0 = 273.15 \text{ K}$ (0°C) is the melting point, e_∞ is the ambient vapor pressure, and $e_{s,w}(T_0)$ is the saturation pressure over water at T_0 ; the air viscosity that enters f_v is taken from Pruppacher and Klett (1997). The two bracketed terms are the sensible heat conducted from the warm air and the latent heat carried by the vapor flux to the wet
 300 surface. As the particle descends through the melting level, the meltwater accumulates on the ice, so the super-droplet becomes mixed. The ice core shrinks while its aspect ratio and apparent density are preserved, and the meltwater is released as rain only when the core has melted completely. A mixed particle that re-enters below-freezing air refreezes. This gradual treatment retains, for each particle, the mixed-phase state through which it passes.



3.4.3 Condensation and evaporation

305 The liquid radius R evolves by the Köhler growth equation

$$R \frac{dR}{dt} = f_v \left[(S - 1) - \frac{a_K}{R} + \frac{b_K}{R^3} \right] (F_k + F_d)^{-1}, \quad (25)$$

where S is the ambient saturation ratio over liquid water, $a_K = 2\sigma/(\rho_w R_v T)$ is the curvature (Kelvin) coefficient, and b_K is the solute (Raoult) coefficient. The coefficient b_K is proportional to the dissolved solute mass, so the activation of cloud condensation nuclei emerges from the solute term and requires no separate activation spectrum. The terms F_k and F_d represent

310 the thermal conduction and the vapor diffusion that limit growth,

$$F_k = \left(\frac{L}{R_v T} - 1 \right) \frac{L}{KT}, \quad F_d = \frac{R_v T}{D e_s}, \quad (26)$$

with L the latent heat of vaporization, K the thermal conductivity of air, R_v the vapor gas constant, D the vapor diffusivity, and e_s the saturation vapor pressure. The ventilation factor f_v enhances the vapor exchange for a falling droplet and follows the Bayley et al. (2026b) fit,

$$315 \quad f_v = \min \left(1 + \frac{x_1 x_2}{x_1 + x_2}, f_{\text{cap}} \right), \quad x_1 = \alpha_1 R^{\beta_1}, \quad x_2 = \alpha_2 R^{\beta_2}, \quad (27)$$

which blends two power laws of the droplet radius R , with fit coefficients α_1, α_2 and exponents β_1, β_2 , and caps the result at f_{cap} . Equation (25) is stiff, since small droplets relax rapidly toward equilibrium. It is integrated by default with a backward-Euler step solved by Newton iteration that allows time steps comparable to the microphysical time step.

3.4.4 Deposition and sublimation

320 An ice particle grows or shrinks by vapor exchange at the rate

$$\frac{dm_i}{dt} = \frac{4\pi C (S_i - 1) f_v}{F_k + F_d}, \quad (28)$$

where S_i is the saturation ratio over ice and C is the capacitance introduced above. The ventilation factor takes the Hall–Pruppacher form (Pruppacher and Klett, 1997),

$$f_v = \begin{cases} 1 + 0.14 \chi^2, & \chi \leq 1, \\ 0.86 + 0.28 \chi, & \chi > 1, \end{cases} \quad \chi = N_{\text{Sc}}^{1/3} N_{\text{Re}}^{1/2}, \quad (29)$$

325 where χ combines the Schmidt and Reynolds numbers N_{Sc} and N_{Re} . The deposited mass is distributed between the two semi-axes according to the Chen and Lamb (1994) inherent growth ratio $\Gamma(T)$, tabulated from 0 to -30°C and retains its terminal value $\Gamma = 1.28$ below. The semi-axes a_i and c_i therefore evolve anisotropically, and the particle develops a plate-like or columnar habit according to the temperature at which it grows.

Figure 2 shows the crystal mass, aspect ratio, semi-axis lengths, and apparent density against the Takahashi et al. (1991)
 330 laboratory data. The ERF implementation reproduces the temperature dependence of the habit diagram. Plates form near



–15°C, where the mass and the a -axis reach their maximum in the dendrite regime. Columns form near –6°C and at the cold end, and the plate-to-column transition occurs near –9°C. The crystal mass at –15°C is approximately 3.3 μg, compared with the Takahashi value of 2.8 μg. Three differences remain: the warm columns are over-elongated, with an aspect ratio near 14 against 7 in the laboratory data; the coldest crystals are too columnar; and the crystal mass at and below –20°C is approximately twice the laboratory value. These are known limitations of the first-generation Chen and Lamb (1994) inherent growth function. Welss et al. (2024) document the same behavior and propose a revised form, while the present implementation uses the first-generation growth function.

3.5 Collisional processes

Collisions are treated by the Shima et al. (2009) Monte-Carlo procedure, which descends from the stochastic-coalescence framework (Gillespie, 1972, 1975). The super-droplets in a grid cell of volume ΔV are grouped at random into $\lfloor N_{sd}/2 \rfloor$ non-overlapping candidate pairs, where N_{sd} is the number of super-droplets in the cell. Examining only these pairs, rather than all $N_{sd}(N_{sd} - 1)/2$ of them, reduces the cost of the collision step from $\mathcal{O}(N_{sd}^2)$ to $\mathcal{O}(N_{sd})$. A candidate pair of super-droplets j and k collides over the microphysical step Δt with probability

$$P_{jk} = \left(\frac{N_{sd}(N_{sd} - 1)/2}{\lfloor N_{sd}/2 \rfloor} \right) \max(\xi_j, \xi_k) \frac{K_{jk} \Delta t}{\Delta V}, \quad (30)$$

where K_{jk} is the collision kernel, and the first term rescales the reduced set of candidate pairs to the full collision rate. The pair undergoes multiple collisions when P_{jk} exceeds 1. The number of events is $\gamma = \lfloor P_{jk} \rfloor + 1$ with probability $P_{jk} - \lfloor P_{jk} \rfloor$, and $\gamma = \lfloor P_{jk} \rfloor$ otherwise, and it is capped at $\lfloor \xi_j/\xi_k \rfloor$ so that no multiplicity becomes negative. The more numerous super-droplet retains its size and loses multiplicity, $\xi'_j = \xi_j - \gamma \xi_k$, while the less numerous one retains its multiplicity and absorbs the collected mass. When the two multiplicities are equal, the pair is split into two equal super-droplets. This update maintains the number of super-droplets nearly constant as the physical particles grow. All three collision types share a geometric kernel,

$$K_{jk} = EA |v_{\infty,j} - v_{\infty,k}|, \quad (31)$$

where A is the collision cross-section, $|v_{\infty,j} - v_{\infty,k}|$ is the difference of the terminal fall speeds, and E is the collection efficiency. The processes differ in the cross-section A , the efficiency E , and the collision outcome, all set by the phases of the two colliding super-droplets.

Two liquid-water particles coalesce through their liquid films. A mixed particle collides through its ice frame and is therefore treated as ice. A liquid particle and an ice or mixed particle consequently rime. The collected water freezes onto the ice below the melting point and is added to the meltwater at and above it. Two particles that both carry ice aggregate.

3.5.1 Coalescence

When two liquid super-droplets collide, they merge. The cross-section in Eq. (31) is $A = \pi(R_j + R_k)^2$, and the efficiency E uses the tabulated Hall (1980) values. The Golovin (1963) kernel, $K_{jk} = b(V_j + V_k)$ with $V = \frac{4}{3}\pi R^3$, which admits an analytic solution, is also available and is used for verification. The merge conserves water and solute: the less numerous partner k retains




Figure 2. Single-crystal vapor deposition over ten minutes at water saturation, from the ERF implementation (red) and the Takahashi et al. (1991) laboratory data (open squares, read from Welss et al. (2024)). (a) Crystal mass, (b) aspect ratio $\phi = c/a$, (c) semi-axis lengths, and (d) apparent density, versus growth temperature.



its multiplicity and grows, while the more numerous partner j loses multiplicity as above,

$$R'_k = (R_k^3 + \gamma R_j^3)^{1/3}, \quad m'_{a,k} = m_{a,k} + \gamma m_{a,j}, \quad T'_{fz,k} = \max(T_{fz,j}, T_{fz,k}), \quad (32)$$

where γ is the number of collision events and m_a denotes each aerosol mass, so the merged particle pools the insoluble nuclei and inherits the warmer freezing temperature. This is the warm-rain growth process, and it produces rain from cloud particles. It also operates between two mixed particles, which collide through their liquid films.

3.5.2 Riming

When liquid water and an ice particle collide, the faster-falling particle is the collector, and the cross-section A in Eq. (31) is the pair cross-section $\pi(r_w + a)(r_w + \max(a, c))$, reduced by the porosity of the ice when the ice collects. The collection efficiency depends on which particle collects,

$$E = \begin{cases} E_{ice}(N_{Re}, N_{Fr}; \phi), & v_{\infty, ice} > v_{\infty, water}, \\ E_{water}(N_{Re}, N_{St}), & v_{\infty, ice} \leq v_{\infty, water}, \end{cases} \quad (33)$$

where $v_{\infty, ice}$ and $v_{\infty, water}$ are the terminal fall speeds of the ice particle and the water, and N_{Re} , N_{Fr} , and N_{St} are the Reynolds, Froude-like impact, and Stokes numbers. The efficiencies follow Beard and Grover (1974); when the ice particle collects, E_{ice} additionally blends the plate-like and columnar Erfani and Mitchell (2017) corrections by the aspect ratio $\phi = c/a$. The collected water of mass m_w depends on the temperature,

$$(\Delta m_i, \Delta m_{rime}, \Delta m_w) = \begin{cases} (m_w, m_w, 0), & T < T_0, \\ (0, 0, m_w), & T \geq T_0. \end{cases} \quad (34)$$

The water freezes as rime below the melting point and adds to both the ice and rime masses, while at or above it the water is retained as meltwater and the ice core is unchanged. The rime is laid down at the Heymsfield and Pflaum (1985) rime density ρ_{rime} , evaluated from the Rasmussen and Heymsfield (1985) impact velocity and bounded to $[100, 910] \text{ kg m}^{-3}$. Supercooled riming grows the spheroid by the rimed volume m_w / ρ_{rime} ,

$$V' = \frac{4}{3} \pi a^2 c + \frac{m_w}{\rho_{rime}}, \quad x'_{ctrl} = \max(x_{ctrl}, r_w (\rho_w / \rho_{rime})^{1/3}), \quad (35)$$

where the controlling semi-axis x_{ctrl} is the equatorial a for $\phi < 0.8$ and for $1 \leq \phi < 1.25$, and the polar c otherwise (Shima et al., 2020); the quasi-spherical band is exchanged because such particles tumble, and the remaining semi-axis is set from the new volume V' . If the collected water exceeds the ice particle in size, $r_w > \max(a, c)$, the product is instead reset to a sphere of the combined ice mass, $a' = c' = [3m'_i / (4\pi\rho_{ice})]^{1/3}$. The monomer count is unchanged, so riming densifies a single crystal into graupel.

3.5.3 Aggregation

When two ice particles collide, they stick and form an aggregate. The cross-section A in Eq. (31) is $(\sqrt{A_1} + \sqrt{A_2})^2$, formed from the projected area A_k of each particle, and the efficiency is fixed at $E = 0.1$ (Shima et al., 2020); aircraft observations of



ice self-collection support a constant efficiency near this value ($E_a = 0.09$) (Field et al., 2006). The two colliders, labeled A and B with the γ copies of A merging into B, add their ice and rime masses and monomer counts,

$$\begin{aligned} m'_i &= \gamma m_{i,A} + m_{i,B}, \\ m'_{\text{rime}} &= \gamma m_{\text{rime},A} + m_{\text{rime},B}, \\ n'_{\text{mono}} &= \gamma n_{\text{mono},A} + n_{\text{mono},B}, \end{aligned} \quad (36)$$

so n_{mono} records the number of crystals in the aggregate. The aggregate ice mass m'_i occupies an apparent volume, which defines the apparent density,

$$V = \frac{4}{3}\pi(\gamma a_A^2 c_A + a_B^2 c_B), \quad \bar{\rho} = \frac{m'_i}{V}. \quad (37)$$

The larger maximum dimension of the two colliders is preserved, so the aggregate is at least as large as its largest component. The remaining semi-axis x is bounded between a compact, volume-filling value $x_{\min}$ and an extended value $x_{\max}$ that stacks the smaller particle onto the larger. The two bounds are blended by the apparent density $\bar{\rho}$,

$$x = (\rho_{\text{ice}} - \rho_{\min}) \left(\frac{\rho_{\text{ice}} - \bar{\rho}}{x_{\min}} + \frac{\bar{\rho} - \rho_{\min}}{x_{\max}} \right)^{-1}, \quad (38)$$

between the solid-ice density ρ_{ice} and a floor $\rho_{\min} = 10 \text{ kg m}^{-3}$. A dense pair maps to the stacked limit, which opens the aggregate and lowers its density, while a pair already near the floor remains compact. Aggregation thereby builds low-density snow from denser pristine ice, with $\rho_{\min}$ setting the open-structure limit. More detailed geometric models exist for rimed and aggregated particles, such as the Seifert et al. (2019) rimed-aggregate geometry and the Jensen and Harrington (2015) riming-driven aspect-ratio evolution, and could replace the spheroid updates used here.

3.6 Interface with the dynamical core

The super-droplets exchange water and heat with the Eulerian state of Sect. 2 through six moisture mixing ratios: vapor q_v , cloud water q_c , rain q_r , cloud ice q_i , snow q_s , and graupel q_g . The vapor is the only prognostic field, and it is advanced by the dynamical core. The five condensate mixing ratios are calculated from the super-droplet state at each step. Their advection, sedimentation, and phase-change sources in the transport equations of Sect. 2 are represented by the motion and growth of the particles rather than evolved on the grid. The local vapor mixing ratio at each particle location is read from the prognostic vapor density divided by the dry-air density, and it drives the phase-change calculations of Sect. 3.4. The vapor consumed or released by condensation, evaporation, deposition, and sublimation is returned to the vapor field as the source F_v in Eq. (5). Algorithm 1 summarizes the coupling over one microphysics step.

The condensate mixing ratios are computed by accumulating the super-droplet masses onto the mesh. Each super-droplet contributes its represented mass $\xi_i m_i$ to the neighboring cell centers through a cloud-in-cell, trilinear weighting. The accumulated mass is then sorted into categories and divided by the air density to yield the mixing ratios. The sorting is where the frozen water is classified. Liquid is split by size: a droplet of radius below $40 \mu\text{m}$ contributes to cloud water q_c , and a larger



Algorithm 1 Coupling of the super-droplet method to the dynamical core over one microphysics step.

- 1: Advance the dynamical core (Sect. 2), updating ρ_d , $\rho_d \mathbf{u}$, $\rho_d \theta_d$, and the vapor density $\rho_d q_v$.
 - 2: Copy the conserved state to the microphysics grid: ρ_d , θ_d , T , p , and q_v .
 - 3: Interpolate the ambient fields needed by the particle processes (ρ_d , p , T , q_v , and the flow velocity $\mathbf{u}$) to each super-droplet position $\mathbf{x}_i$.
 - 4: Inject new super-droplets.
 - 5: Phase change: freezing and melting, then condensation and evaporation, then deposition and sublimation.
 - 6: Advect and sediment the super-droplets.
 - 7: Apply the boundary treatment (Sect. 3.3): deactivate super-droplets that reach the surface or a domain boundary, accumulate the surface precipitation, and recycle or free the deactivated particles.
 - 8: Coalesce the super-droplets.
 - 9: Deposit $\xi_i m_i$ onto the mesh, classify it into q_c , q_r , q_i , q_s , and q_g , and accumulate the latent heating into θ_d .
 - 10: Copy back to the conserved state: the moisture densities $\rho_d q_v$, $\rho_d q_c$, $\dots$, $\rho_d q_g$ and $\rho_d \theta_d$.
-

one to rain q_r . Frozen water is classified by the rime-mass fraction $f = m_{\text{rime}}/m_i$ and the monomer number n_{mono} , with a threshold $f^* = 0.3$:

- 420 – cloud ice q_i : pristine crystals with $f < f^*$ and $n_{\text{mono}} = 1$;
 – snow q_s : unrimed or lightly rimed aggregates with $f < f^*$ and $n_{\text{mono}} > 1$;
 – graupel q_g : heavily rimed particles with $f \geq f^*$.

These categories are diagnostic. They allow the super-droplet method to report the same six mixing ratios as the bulk schemes of Sect. 2. Here, the categories emerge from the per-particle attributes, rather than from prognostic variables with prescribed
 425 conversion rates between them.

Finally, the latent heat of the phase changes is returned to the potential temperature. Each transition contributes an increment $\Delta\theta_d = (\theta_d/T)(L/c_p)\Delta q$, with T the temperature and c_p the specific heat, where the latent heat L and the mass change Δq depend on the transition:

- condensation and evaporation, between liquid and vapor: $L = L_v$, and Δq is the change in vapor;
 430 – deposition and sublimation, between ice and vapor: $L = L_s$, and Δq is the change in vapor;
 – freezing and melting, between liquid and ice: $L = L_f$, and Δq is the change in liquid.

Here L_v , L_f , and $L_s = L_v + L_f$ are the latent heats of vaporization, fusion, and sublimation. These increments form the heating $H_n + H_p$ in Eq. (3), and the density-weighted potential temperature $\rho_d \theta_d$ is updated accordingly.

3.7 Implementation in ERF

435 ERF is built on the AMReX framework (Zhang et al., 2021), which supplies performance-portable data structures and parallel loops that compile for several architectures from a single source. Field variables are stored in distributed grid data containers: the domain is partitioned into cuboidal boxes that are spread across the distributed-memory ranks, and each rank owns the data on its boxes. A computation is expressed by passing a function to a parallel loop that applies it to every grid cell, or to



every particle. The loop runs across threads on a multi-core processor and as a kernel over the same elements on a graphics
 440 processor. This applies to hardware from NVIDIA, AMD, and Intel, through their respective programming models. The same
 super-droplet source therefore runs on all of these architectures without modification.

The super-droplets are stored in a particle container, the particle counterpart of the grid data container, distributed across the
 same ranks as the mesh. A container stores each attribute in its own contiguous array spanning all the particles, rather than each
 particle storing a contiguous record of its own attributes. A parallel loop over particles then reads or writes one attribute with
 445 unit stride, which vectorizes on a processor and yields coalesced memory access across the threads of a graphics processor.
 Each super-droplet stores the attributes described in Sect. 3 together with additional quantities that the time integration needs:
 the terminal velocity, a unique identifier, and an active flag. The active flag distinguishes super-droplets active in the simulation
 from those that have exited the domain. The kernels skip the inactive particles, which are then removed or reused as new
 injected super-droplets. The memory requirement and computational cost of this representation are analyzed in Sect. 5.

450 4 Results

We first verify the `ERF` implementation of the cold microphysics, examining each microphysical process in isolation. We then
 apply the method to three-dimensional deep convection, which exercises the coupled system in a dynamically active flow. A
 final ensemble quantifies the run-to-run variability of the stochastic method.

4.1 Computational platforms

455 The simulations are run on the Tuolumne (Livermore Computing, Lawrence Livermore National Laboratory, 2024) and Ma-
 trix (Livermore Computing, Lawrence Livermore National Laboratory, 2025) systems at Lawrence Livermore National Labo-
 ratory. Tuolumne is an HPE Cray system whose nodes each contain a 96-core AMD EPYC processor and four AMD MI300A
 accelerated processing units that share 512 GB of unified memory. Matrix is a Dell system whose nodes each contain two
 56-core Intel Xeon Platinum 8480+ processors and four NVIDIA H100 GPUs with 80 GB of memory each. The simulations
 460 run on the GPUs of both systems. On Matrix, the code is compiled with the CUDA toolkit version 12.9 and the GNU host
 compiler version 12.1. On Tuolumne, it is compiled with the ROCm toolchain and the Clang compiler version 19.

4.2 Verification

We verify the fundamental cold-microphysics processes against laboratory data, published benchmarks, and independent solu-
 tions of the governing rate equations. The warm-rain processes were verified previously by Beydoun et al. (2026). Each process
 465 is isolated in a configuration in which the environment is prescribed and held fixed. The measured response of the particle pop-
 ulation then depends only on the process under examination. Most of the processes share a common column configuration,
 which we describe next. Ice aggregation uses a well-mixed box, described with that case.

The shared configuration is a one-dimensional vertical column. The domain spans a height specific to the process. It is 4 cells
 wide in x and y , the minimum permitted by the spatial stencils in `ERF`, with periodic lateral boundaries. The thermodynamic




Figure 3. Melting layer. (a) Fraction of mass melted versus distance below the 0 °C level for the six Mitra et al. (1990) Table 1 cases: the ERF column (solid), an independent integration of the melt rate law (dashed), and the melt depths reported by Mitra et al. (1990) (open squares). (b) Fall speed versus melt fraction for the 10 mm, 100%, 0.006 K m⁻¹ case, against the Frick et al. (2013) blend with a shrinking ice core (solid) and with fixed endpoint speeds (dashed).

environment is prescribed as a function of height and is held fixed. The temperature follows a linear profile with a constant lapse rate, or is uniform. The pressure is in hydrostatic balance. The relative humidity over water or ice is prescribed, and it may be uniform or vary with height. The latent heat and moisture released by the microphysics do not feed back to the environment. The profiles therefore remain stationary while the particles evolve within them. The mean flow is either 0 or a specified vertical updraft. A particle population enters the column at a steady rate through one end. Particles enter at the top for the falling tests and at the bottom for the rising test. The injected size is either monodisperse or drawn from an exponential mass distribution. Each particle is transported by its terminal fall speed and the prescribed flow, and it undergoes the single process under test at the local environment. The injection, the environment, and the transport are stationary, so the column reaches a steady state. The diagnostic is the steady vertical profile of a per-particle quantity. We compare it with an independent solution of the same rate equation, or with a published profile.

The falling and rising tests compare the ERF column with an independent integration of the same governing rate equation. This reference is a forward-Euler quadrature of the rate law along the fall path, marching in height with a small step through the prescribed environment and using material coefficients comparable to those of the microphysics. It shares the physics of the ERF kernels but not their particle solver, so agreement isolates the numerics of the coupled column from the rate law itself.




Figure 4. (a) Sublimation column: ice mass fraction remaining m/m_0 versus height for three subsaturations over ice at -15°C ; markers are the ERF column, lines the diffusion ODE. (b) Immersion freezing: frozen fraction of the rising droplets versus height, with the corresponding air temperature on the right ordinate, against the Niemand et al. (2012) active-site curve. (c) Riming column: ice mass ratio m/m_0 versus height for three cloud water contents at -10°C ; markers are the ERF column, lines the continuous-collection ODE. (d) Ice aggregation box: ice mass-density distribution $g(\ln R)$ at four times, for four super-droplet counts N_{sd} ; the layout follows Brdar and Seifert (2018), Fig. 5.




Figure 5. Ice and snow column (Brdar and Seifert, 2018), Fig. 7. (a) Prescribed environment: temperature and relative humidity over water. (b) Number density, (c) number flux, (d) mass density, (e) mass flux, and (f) super-droplet count, each a steady vertical profile for four process levels. S is sedimentation, SD adds deposition, and SDA adds aggregation at collection efficiency E_a . Dotted lines show the water-saturated layer between 1 km and 2 km.



4.2.1 Melting layer

Figure 3 shows the melting-layer test. Monodisperse low-density ice flakes are injected at the top of an 800m column. The flakes are seeded at the low apparent density of dry snow. The temperature is 0°C at the top of the column and increases downward at the case lapse rate. Six cases reproduce the six combinations of flake diameter (5mm and 10mm), relative humidity (90, 95, and 100%), and lapse rate (0.006K m^{-1} and 0.01K m^{-1} , tabulated as 0.6°C and 1.0°C per 100m in Mitra et al. (1990, Table 1)). Deposition is disabled for the mixed particles, so the flakes only melt. We compare the column with the Mitra et al. (1990) results and with an independent integration of the melt-rate law of Eq. (24). Figure 3(a) shows the fraction of mass melted against the distance below the 0°C level. The ERF column follows the independent solution for every case. The ordering of the cases is physical. A smaller flake melts over a shorter distance. A lower relative humidity melts deeper, by approximately 120m for the 90% case relative to the saturated case. A steeper lapse rate melts over a shorter distance. The two subsaturated cases begin melting below the 0°C level, near 90m and 150m for the 95 and 90% cases. The flake surface remains subcooled by sublimation until the flake reaches warmer air. Figure 3(b) shows the fall speed against the melt fraction for one case. It reproduces the Frick et al. (2013) blend with the shrinking ice core. The completion depths are approximately 20m to 40m deeper than the depths reported by Mitra et al. (1990). This offset reflects two modeling choices that act in opposite directions. Our implementation retains the flake's seeded density and shrinks it by mass alone, which yields a large capacitance and a rapid melt. The flakes considered by Mitra et al. (1990) collapse toward a drop as it melts, which slows the melt. We consider a spherical flake that is heavier than the oblate flake of the same maximum dimension in Mitra et al. (1990); this deepens the melt. The two effects act in opposite directions and largely compensate, so the residual difference is small.

4.2.2 Sublimation column

The sublimation test of Fig. 4(a) injects monodisperse solid-ice spheres of diameter 0.35mm at the top of a 1km column. The column is isothermal at -15°C and subsaturated over ice. Three cases set the relative humidity over ice to 75, 82, and 90%. Each sphere falls at its terminal speed and loses mass by vapor diffusion, and it shrinks at a fixed apparent density. We compare the column with an independent integration of the same diffusion law. The figure shows the ice mass fraction remaining against height. The ERF column follows the independent solution across the full column for all three subsaturations. The mass fraction that reaches the ground matches the independent solution to within 0.01.

4.2.3 Immersion freezing

Immersion freezing is verified in a one-dimensional updraft column, with the result in Fig. 4(b). Each droplet carries a dust nucleus and is assigned a freezing temperature from the Niemand et al. (2012) ice-nucleating active-site density, which distributes the individual thresholds between about -12°C and -36°C . A droplet freezes when the local temperature drops below its threshold. The column carries droplets upward in a steady 2ms^{-1} updraft through a water-saturated environment that cools from -8°C at the base to below -38°C aloft. The droplets enter all-liquid at the base and freeze as they rise into colder air. The column therefore couples the freezing to the vertical transport. The figure shows the frozen fraction against height, with



the corresponding air temperature on the right ordinate. The column follows the Niemand et al. (2012) curve, matching it to within 0.02 at every temperature.

4.2.4 Riming column

The riming test, shown in Fig. 4(c), involves a solid ice sphere of diameter 2 mm injected at the top of a 1 km column. The column is water-saturated and isothermal at -10°C . It is seeded with a uniform supercooled cloud of monodisperse $10\mu\text{m}$ droplets. Three cases set the cloud liquid water content (LWC) to $4 \times 10^{-4}\text{kg m}^{-3}$, $7 \times 10^{-4}\text{kg m}^{-3}$, and $1.0 \times 10^{-3}\text{kg m}^{-3}$. Phase change is disabled, so the ice grows only by collecting cloud droplets. We compare the column with an independent integration of the continuous-collection equation, using the same collection kernel and the rime density measured from the run. The figure shows the ice mass ratio against height. The ERF column follows the independent solution for all three water contents. The ice reaches the ground at a mass ratio of 1.35, 1.68, and 2.06, and it matches the independent solution to within 1%. The rime accretes at a density near 540kg m^{-3} in all three cases.

4.2.5 Ice aggregation box

Ice aggregation is verified in a well-mixed box with no flow, and Fig. 4(d) plots the resulting spectrum. The box is seeded with an exponential ice mass distribution at an ice water content of 10^{-3}kg m^{-3} . Aggregation proceeds by differential sedimentation, driven by the fall-speed difference between colliding particles. The reference is the McSnow model (Brdar and Seifert, 2018), which uses the same box configuration. The figure shows the ice mass-density distribution at four times, for four super-droplet counts. The distribution broadens from a sharp peak near 10^{-4}m toward the millimeter scale, which matches the shape of the Brdar and Seifert (2018) result. The four super-droplet counts coincide at each time, and this confirms Monte-Carlo convergence. ERF aggregates approximately 50 times more slowly than the reference. The continuous-collection kernel scales as the product of the collection efficiency and the geometric cross-section swept by the falling particle, and the two codes differ in both. Both codes employ $(\sqrt{A_1} + \sqrt{A_2})^2$ and differ in the area entered into it and in the efficiency. ERF applies a constant aggregation efficiency of 0.1 (Shima et al., 2020; Field et al., 2006) and evaluates A_k as the circumscribed ellipse of the spheroid reduced by its porosity. McSnow multiplies a Böhm collision efficiency, which approaches 1 for millimeter-sized ice, by a temperature-dependent sticking efficiency, and adopts the larger bounding circle on the maximum dimension, $\frac{\pi}{4}(D_1 + D_2)^2$ (Brdar and Seifert, 2018). The smaller ERF efficiency and area reduce its kernel and slow the growth of the aggregate spectrum by the observed factor, without altering its converged shape. A slower aggregation kernel plausibly biases the snow formation low in the deep-convection case of Sect. 4.3; a sensitivity study of the aggregation kernel is noted as future work in the conclusions.

4.2.6 Ice and snow column

The final case, in Fig. 5, couples sedimentation, deposition and aggregation in a 5 km column (Brdar and Seifert, 2018). Ice enters at the top with an exponential mass distribution and a mean radius near $130\mu\text{m}$. The environment has three zones in



height. The air is supersaturated over ice above 2 km, and the crystals grow by deposition. It is saturated over water between 1 km and 2 km, and subsaturated over ice below 1 km, where the crystals sublimate. We build the process up in stages: S with sedimentation alone, SD with deposition added, and SDA with aggregation added at collection efficiencies of 0.1 and 1. Figure 5(a) shows the prescribed environment, and Figs. 5(b)–5(f) show the steady profiles. Sedimentation alone (S) lets the injected ice fall at a size-sorted speed, and the number density is uniform with height. Adding deposition allows the crystals to grow and fall faster, so the number density decreases downward, and the mass density attains its maximum near the water-saturated layer. The crystals sublimate below 1 km, which produces a number-density spike and then removes the ice before it reaches the ground. Adding aggregation at the low efficiency results in profiles similar to the deposition-only case. Raising the efficiency to 1 reduces the number density by about a factor of 3 and lowers the number flux, while the mass flux is conserved. These profiles reproduce the structure of the results presented in Brdar and Seifert (2018). Results from the two GPU platforms we use agree to within 0.1%.

4.3 Three-dimensional deep convection




Figure 6. Schematic of the deep-convection (GATE) case: a $100 \times 100 \text{ km} \times 21 \text{ km}$ domain ($200 \times 200 \times 230$ cells, 500 m horizontal spacing and a stretched vertical grid) with periodic lateral boundaries, a Rayleigh-damping layer in the top 2 km, and a Monin–Obukhov surface layer. The column is initialized from the observed GATE sounding and driven by the prescribed large-scale forcing (advective tendencies, subsidence, and a geostrophic wind through the Coriolis term at 15°N).

We simulate a three-dimensional deep-convection case that exercises the coupled cold microphysics in a realistic tropical environment. It is based on the Global Atmospheric Research Program (GARP) Atlantic Tropical Experiment (GATE), a standard benchmark for tropical cloud-resolving modeling (Grabowski et al., 1996, 1998). The domain spans $100 \times 100 \text{ km}$ in the horizontal and 21 km in the vertical, discretized with $200 \times 200 \times 230$ cells. The horizontal spacing is 500 m. The vertical spacing is 50 m below 1.25 km, widens to 100 m between approximately 5 km and 18 km, and stretches to 162 m at the model top; the lower troposphere is therefore the best resolved. The lateral boundaries are periodic, a Rayleigh layer in the top 2 km damps vertically propagating waves, and subgrid turbulence is closed by a Smagorinsky model (Smagorinsky,




Figure 7. Initial conditions for the deep-convection case. (a) Aerosol number size distributions for the two sulfate modes and the dust mode. (b) Environmental sounding: temperature (bottom axis) and relative humidity over water and over ice (top axis) against height; the dotted line marks the -38°C homogeneous-freezing level.

1963). The domain is initialized from the observed GATE sounding and driven by the prescribed large-scale forcing of the campaign. Horizontal-advective tendencies of potential temperature and water vapor, a large-scale subsidence profile, and a geostrophic wind acting through the Coriolis term at the GATE latitude of 15°N maintain a convectively unstable mean state. Surface fluxes of heat, moisture, and momentum are supplied separately through a Monin–Obukhov surface layer (Monin and Obukhov, 1954). Convection develops from small random perturbations of temperature and water vapor, of amplitude 0.1 K and 0.025 g kg^{-1} , applied below 1.6 km in the initial state. The time step is 0.5 s , and the simulations advance to $t = 36 \text{ h}$. Figure 6 summarizes the configuration.

We simulate two configurations, with and without dust. Both configurations carry two soluble ammonium-bisulfate (NH_4HSO_4) modes, log-normal in radius, with geometric-mean radii of 30 nm and 140 nm , geometric standard deviations of 1.28 and 1.75 , and number concentrations of $9 \times 10^7 \text{ m}^{-3}$ and $1.5 \times 10^7 \text{ m}^{-3}$ (Fig. 7(a)). They carry no ice-nucleating surface since they are soluble. The with-dust configuration additionally includes an ice-nucleating aerosol: dust sampled from a log-normal distribution with a geometric-mean radius of $2.5 \mu\text{m}$, a geometric standard deviation of 1.70 , and a number concentration of 10^6 m^{-3} , reported as 1 cm^{-3} by Ryder et al. (2018). We truncate this distribution to radii between $0.5 \mu\text{m}$ and $15 \mu\text{m}$. Each dust particle carries a thin NH_4HSO_4 coating so that it activates as a cloud droplet. These number concentrations are surface values ($z = 0$): the super-droplet number density is scaled with the local air density, so the physical concentration decreases with height.

The two soluble modes are each represented by 64 super-droplets per cell. Thus, the simulation without dust involves 128 super-droplets per cell. The with-dust configuration adds a further 16 super-droplets per cell for the dust, raising the count




Figure 8. Three-dimensional hydrometeor isosurfaces for the deep-convection case at $t = 12$ h, 25 h, and 36 h (rows), without dust (left column) and with dust (right column), each species drawn at its own mixing-ratio level: cloud 0.10 g kg^{-1} , rain 0.20 g kg^{-1} , cloud ice 0.02 g kg^{-1} ($z > 13.5 \text{ km}$), 0.10 g kg^{-1} ($z \leq 13.5 \text{ km}$), snow 0.03 g kg^{-1} , and graupel 0.03 g kg^{-1} . The vertical axis is truncated to 17 km.




Figure 9. Horizontally averaged hydrometeor profiles for the deep-convection case at $t = 12, 25$ and 36 h (rows), without dust (left) and with dust (right).




Figure 10. Condensate water paths of the deep-convection case against time, without dust (solid) and with dust (dashed): (a) the liquid categories and (b) the frozen categories. The water path is the condensate mass per unit area, so each curve is the horizontal mean of the column integral.

to 144 super-droplets per cell. These counts follow Shima et al. (2020), who reported ensemble convergence near 128 super-droplets per cell for a comparable deep-convection case. The simulations involve 1.2×10^9 super-droplets in total without dust and 1.3×10^9 with dust. These simulations are performed on 64 nodes of Tuolumne (256 GPUs). The cost of each timestep is approximately 0.39s of wall time without dust and 0.46s of wall time with dust. These correspond to 1.3h and 1.1h of simulated time per hour of wall time, respectively. Consequently, simulating 36h (259,200 timesteps) requires approximately 28h and 33h of total wall time. Figure 8 shows the solution as three-dimensional isosurfaces of the five hydrometeor categories.

The field at 12h is a population of scattered warm convective cells that carry cloud and rain. A thin cloud-ice layer near 15km is also present. A relative-humidity maximum causes the initial super-droplets to condense a small amount of water, and the ambient temperature there ($\sim 204\text{K}$, or $\sim -69^\circ\text{C}$) causes homogeneous freezing. The cloud and rain form deeper mixed-phase cells by 25h, with graupel and snow embedded in the ice and rain, and it remains an actively evolving mixed-phase field at 36h.

Figure 9 shows the horizontally averaged hydrometeor profiles at $t = 12\text{h}$, 25h , and 36h . We observe the warm cloud-and-rain layer that reaches 6km at 12h, with the convectively produced frozen categories still absent. The plume glaciates by 25h. Cloud ice is observed from 4km to 16km, snow from 4km to 10km, and graupel from 4km to 7km. The warm-rain column persists beneath. Its mean increases by approximately a factor of 2 from 12h to 25h, and changes little thereafter. Its maximum at 36h is approximately 0.05g kg^{-1} , at an altitude of 4km, while the cloud-ice column remains deep. Figure 10 follows the condensate water paths, the condensate mass per unit area, through the complete simulation. The frozen categories attain their maximum between 20h and 26h and decrease thereafter. The vertically integrated frozen condensate at 36h is 58% of its value




Figure 11. Predicted ice-particle morphology of the almost-pure-ice super-droplets in the with-dust case, at $t = 25$ h and 36 h (rows). Columns are, from left to right, the fall speed, the mass normalized by the solid-ice sphere of equal maximum dimension, the aspect ratio, and the apparent density, each against maximum dimension D . Marker opacity denotes the population mass density, and the solid line segments are the Locatelli and Hobbs (1974) relations. The dotted lines denote solid-ice sphere properties: its Stokes fall speed in the first column, and unit normalized mass, unit aspect ratio and the solid-ice density in the remaining three columns.

at 25 h without dust and 42% with dust. Cloud ice accounts for most of this loss in absolute terms. However, snow decreases most steeply, by approximately 86% in both configurations. Dust advances the onset of graupel formation, since heterogeneous freezing acts at temperatures well above the homogeneous-freezing threshold. Graupel first appears in trace amounts at 11 h with dust and grows steadily from 13 h, whereas without dust it appears only at 16.5 h.

The super-droplet method evolves the geometry and material attributes of every ice particle. Figure 11 shows the predicted morphology of the almost-pure-ice particles (pure-ice super-droplets and mixed-phase super-droplets with melt fraction below 0.01) in the with-dust case at two glaciated times, $t = 25$ h and 36 h. The fall speed, normalized mass (fraction of the mass of a solid-ice sphere of equal dimension), aspect ratio, and apparent density are shown as a function of the maximum dimension D . Each is colored by the diagnostic (ice/snow/graupe) category. The almost-pure-ice population numbers 3.8 million super-droplets at 25 h and 3.0 million super-droplets at 36 h. The fall speed increases from 10^{-3} m s^{-1} to approximately 10 m s^{-1} as D increases. At the largest sizes, the particles transition from the viscous Stokes regime, marked by the dotted line, to a




Figure 12. Predicted morphology of the mixed-phase super-droplets (melt fraction 0.01 to 0.99) of the deep-convection case with dust, at 25 h (top) and 36 h (bottom). Columns are the fall speed, the mass normalized by a solid-ice sphere, the aspect ratio, and the apparent density, colored by category; the frozen-particle references of Fig. 11 are repeated. The three geometric columns use the ice-skeleton maximum dimension $D = 2 \max(a, c)$ weighted by ice mass, whereas the fall-speed column uses the effective size $D^* = \max(D, D_{eq})$, with D_{eq} the raindrop-equivalent diameter of the total condensate, and is weighted by total condensate mass, since the ice core shrinks as the particle melts.

higher-Reynolds-number drag regime. The three categories separate in the expected order: the dense graupel falls fastest, cloud ice is slowest, and the snow aggregates are intermediate. The normalized mass is below 1 for nearly all particles and decreases with D . Thus, the ice is porous, and its filling fraction falls as deposition and aggregation extend the particle faster than they add mass. The cloud ice carries the lowest normalized mass and the rimed categories the highest. The aspect ratio is skewed toward columns ($\phi > 1$), especially for the cloud ice. This is consistent with the inherent growth ratio, which exceeds 1 from -20°C and remains at its terminal value of 1.28 below -30°C , over the range in which this ice grows. The apparent density follows the normalized mass: it is near the solid-ice value for the smallest crystals and declines toward the large aggregates. Pristine ice therefore begins near-solid and grows progressively more porous. These trends agree qualitatively with the deep-convection simulation of Shima et al. (2020): the porosity increases with size, the apparent density declines from near-solid, the fall speed saturates in the expected category order, and the upper-level cloud ice is columnar. The mass-size and velocity-size distributions emerge from the per-particle attributes, so each category clusters near its own Locatelli and Hobbs (1974) relation,




Figure 13. Melting particles in the deep-convection case with dust. (a) The blend weight Ψ_{obs} of every melting super-droplet against melt fraction ℓ at 36 h, with the Frick et al. (2013) weight $\Psi(\ell)$ overlaid; particles whose ice and raindrop endpoints converge are drawn separately. (b) Liquid mass fraction of the condensate mass against altitude at 25 h and 36 h, taken over the super-droplets that carry ice (solid) and over all condensate (dashed), with the horizontally averaged temperature on the right ordinate. The horizontal dashed line marks the 0°C level, the lower shaded band the melting layer, and the upper band the supercooled mixed-phase layer; both edges are taken from the curves at whichever time reaches furthest.

and the spread records the range of habit, riming, and aggregation history within the population. The velocity comparisons are qualitative, because the Locatelli and Hobbs (1974) relations were measured near the surface. The fall speed instead scales with the inverse square root of the air density, and is larger in the low-density air at higher altitudes. The configuration without dust produces a morphologically similar ice population.

Figure 12 presents the morphology of the mixed-phase super-droplets with melt fraction between 0.01 and 0.99 for the with-dust case at 25 h and 36 h, and shows how melting reshapes the particles. The population is approximately 110,000 super-droplets at 25 h and 60,000 super-droplets at 36 h. This is a small fraction of the almost-pure-ice, because a particle crosses the melting range rapidly. Melting graupel dominates the population, followed by cloud ice and a smaller number of snow aggregates. We observe key differences between these particles and the almost-pure-ice ones of Fig. 11. Overall, a much smaller number of small particles are present in this population; larger particles sediment faster and reach lower (warmer) altitudes. The mass-weighted median maximum dimension at 25 h is 1.2 mm, compared with 0.4 mm for the almost-pure-ice. The normalized mass is substantially higher, and its mass-weighted median at 25 h is 0.57, compared with 0.17 for the almost-pure-ice. This is because the meltwater fills the pores of the ice skeleton, and the rounding particle approaches the solid-ice



sphere of its own maximum dimension. The apparent density measures the ice mass against the volume of the ice skeleton rather than of that sphere. It is already near the solid-ice value in the frozen population, and therefore differs little between the two populations, at 820 kg m^{-3} and 920 kg m^{-3} , respectively. Their aspect ratio decreases toward 1, as the surface tension of the meltwater rounds the particle. The mass-weighted median at 25 h is 1.3, compared with 2.3 for the almost-pure-ice, and the mass fraction with $\phi > 1$ is 0.79, compared with 0.97. The geometry therefore evolves from the porous, anisotropic crystal toward the compact, spherical raindrop.

The fall speed of these melting super-droplets follows the Frick et al. (2013) blend of the ice-core and raindrop limits, evaluated for each particle at its local air state. Figure 13(a) shows that the blend weight Ψ_{obs} agrees with the analytic $\Psi(\ell)$ at 36 h, to a median absolute error of 1.4×10^{-3} . The scatter is confined to the ill-conditioned tail, where the ice-core and raindrop speeds converge for the largest melting graupel. The mixed-phase particles are also confined to a well-defined melting layer. Figure 13(b) shows the mass-weighted liquid fraction $\ell = m_w / (m_i + m_w)$ against altitude over the super-droplets that carry ice, whether as cloud ice, snow, or graupel, and over all super-droplets with condensate. These separate the melting layer from the supercooled mixed-phase layer above it. The liquid fraction of the mixed-phase particles shows the melting layer of the deep-convection case. These particles reach the 0°C altitude, near 4.2 km, with a melt fraction near 0. The mean rises steadily as they descend into warmer air. The partially melted particles are located between 1.3 km and 4.4 km at 25 h, and between 2.7 km and 4.2 km at 36 h. The melting layer is therefore deeper at 25 h, extending ~ 3 km below the freezing altitude, because the larger graupel present at that time carries ice further down before melting completely. Melting completes within ~ 1.5 km of the freezing altitude at 36 h. The liquid fraction over all condensate remains high above the freezing altitude, because the condensate budget there is dominated by supercooled liquid cloud and rain rather than by the ice. It decreases below 0.5 at 4.4 km, or -1.4°C , and is 0.1 at 6.4 km, or -13°C , at 25 h. The corresponding altitudes at 36 h are 5.1 km, or -4.7°C , and 8.4 km, or -26°C . In this supercooled mixed-phase layer, ice and supercooled liquid coexist without the ice melting. The layer therefore deepens as the system matures, and the supercooled liquid reaches approximately 2 km higher at 36 h than at 25 h.

4.4 Run-to-run variability

The super-droplet method is stochastic, and there are two sources of physical stochasticity: the initial particle attributes drawn by sampling of the aerosol distribution, and the Monte-Carlo collisions. When dust is present, the INAS-sampled freezing temperature adds another stochastic factor. The multi-threaded GPU implementation introduces an additional stochasticity due to the non-deterministic assignment of particles to GPU threads. This manifests through atomic operations that particle-to-mesh depositions use; however, this is of the order of round-off errors. We analyze the robustness of our simulations to these variations. An ensemble with sufficient realizations of the entire 36 h evolution is too expensive, so we quantify the run-to-run spread in two shorter windows. Each window is a ten-member ensemble of the deep-convection case with dust on 256 GPUs. The first window begins from $t = 0$ and measures the sensitivity to the initial Monte-Carlo sampling of the aerosol and collisional processes that form rain. Each realization of this window is given a distinct random seed, so it draws an independent initial droplet sampling and an independent collision stream. The convection is still shallow and warm at 2 h. A boundary-layer cloud has formed and rain has begun, but the deep updrafts and the ice phase have not yet appeared. Figure 14(a) shows




Figure 14. Run-to-run variability of the horizontally averaged hydrometeor profiles for two ten-member ensembles of the deep-convection case with dust; each of the ten realizations is drawn as a thin line per species. (a) Started from $t = 0$, at $t = 2$ h, when the convection is still shallow and warm. (b) Started from $t = 23$ h (the collisional peak), at $t = 24$ h.

the cloud and rain profiles of the ten realizations. The cloud spread is small, with a median coefficient of variation of 0.18, while the rain spread is considerably larger, at 0.54. The rain is still forming at this time, and its mixing ratio is more than an order of magnitude below that of the cloud, so a small absolute difference between the realizations is a large relative one. The second window restarts from the state at 23 h, when the collisional processes are most active, and measures the sensitivity of the fully developed cold microphysics to the collisions. It advances 1 h. Figure 14(b) shows the hydrometeor profiles of the ten realizations, which nearly coincide. The median coefficient of variation is between 0.003 and 0.052 for the five species, and it does not exceed 0.308 at any altitude. Overall, the small differences between the realizations show that the simulations are physically robust.

5 Computational cost and scalability

The super-droplet method is computationally more expensive than an Eulerian bulk scheme, because it carries and updates many particles per cell rather than a few fields per cell. This section analyzes that cost. We measure and analyze the wall



times and strong scaling of the three-dimensional deep-convection simulation of Sect. 4.3. We then estimate the memory requirements for this case.

5.1 Computational cost




Figure 15. Strong scaling of the deep-convection case with dust. (a) Per-step wall time of the total, the super-droplet method, and the dynamical core, against the ideal $1/N$ line. (b) The corresponding parallel efficiency, relative to the 8-GPU point of each machine. Solid lines with filled markers are Tuolumne (MI300A, 8 to 512 GPUs); dashed lines with open markers are Matrix (H100, 8 to 40 GPUs).

We measure the strong scaling of the coupled method on the deep-convection case with dust. The mesh is decomposed horizontally across MPI processes, in x and y only, so each subdomain spans the full depth. Each MPI process is assigned a bundle of grid columns. The maximum number of MPI processes is constrained by the requirement that each column be at least 4 cells wide, which is the size of the spatial discretization stencils. This caps the 200×200 horizontal grid at $\lfloor 200/4 \rfloor^2 = 2500$ MPI processes. Each super-droplet belongs to the process that owns the grid cell it is in and is reassigned as advection carries it across a subdomain boundary. The scaling studies begin from the state at 23h, with 1.22×10^9 active super-droplets, when the collisional processes are most active. The domain is partitioned evenly among the accelerators, and we advance 250 timesteps. We use 8 to 512 GPUs (2 to 128 nodes) of Tuolumne and 8 to 40 GPUs (2 to 10 nodes) of Matrix, since the latter is a smaller resource. One MPI process is assigned one GPU. We measure the cost of each operation, averaged over the accelerators.

Figure 15(a) shows the per-step total wall time and its contributions from the super-droplet method and the dynamical core on Tuolumne and Matrix. The super-droplet method scales near-ideally on Tuolumne. The wall time reduces from 6.2s to 0.15s per step between 8 and 512 GPUs. Its parallel efficiency, shown in Fig. 15(b), is 84% and 63% at 256 and 512 GPUs, respectively. The lower efficiency at the large GPU counts is due to fixed per-step overheads. We observe near-ideal scaling on Matrix as well, from 6.1s to 1.3s per step, a $4.6\times$ speed-up with $5\times$ more GPUs. The per-step cost is lower than that of the MI300A at the same GPU counts. The dynamical core scales through 48 GPUs and then saturates near 0.17s per step. Its




Figure 16. Composition of the super-droplet cost, the fraction spent in each process, on Tuolumne (MI300A, 8 to 512 GPUs) and Matrix (H100, 8 to 40 GPUs).




Figure 17. Per-count breakdown of the collision cost on Tuolumne (MI300A): (a) the parallel efficiency of the three stages and of their total, normalized to the 8-GPU point, and (b) each stage as a fraction of the collision cost, labeled with its wall time per step.




Figure 18. (a) The mesh-particle coupling stencil, the two-dimensional equivalent of the eight-point three-dimensional stencil: an ambient field ϕ is interpolated from the cell centers $\phi_{i,j}$ to the super-droplet and its mass $\xi_i m_i$ deposited back, both with the trilinear weights $w_{i,j}$, each the fractional volume of the diagonally opposite sub-cell (shaded). (b) Scaling of the coupling operations against the ideal $1/N$ line; solid lines with filled markers are Tuolumne, dashed lines with open markers are Matrix (H100).




Figure 19. The compute-plus-communication time split into the particle and dynamical-core computation and the three kinds of communication, as a fraction of the total, on Matrix (H100, left) and Tuolumne (MI300A, right).



subdomains contain too few cells to occupy the accelerators, and its memory-bound stencil operations are no longer efficient. The super-droplet method dominates the cost at the smaller number of GPUs, at 18 times the dynamical core on 8 GPUs. The cost of the SDM is comparable to that of the dynamical core at ≥ 256 GPUs, which demonstrates the necessity of GPUs to make the SDM computationally tractable.

Figure 16 shows the cost of each SDM process as a fraction of the method’s per-step cost. On the MI300A, the composition is nearly independent of the GPU count through 256 GPUs, because every process scales at the same near-ideal rate. The super-droplet operations share a common work complexity but differ in how readily they thread. The implementation is parallel at two levels: the domain is decomposed across MPI processes, and the work within each process is then multithreaded in a shared address space. Advection, sedimentation, and phase changes are local to each super-droplet. Advection and sedimentation contribute 18%. Phase changes account for 80% of the cost, with condensation and evaporation being the most expensive, because they integrate a stiff condensation equation for each super-droplet with an iterative backward-Euler solver. Freezing, melting, deposition, and sublimation are explicit single-step updates. The implicit solver adds arithmetic operations per particle with no synchronization. The collisional processes are 2% on the MI300A. A remaining (“other”) component grows to 15% at 384 and 512 GPUs: the per-step diagnostics, which count the collisions through a global reduction, together with the boundary handling and the kernel-launch latency. These are fixed per step and do not scale with the particle work. This composition depends on the architecture, and it differs on Matrix from that on Tuolumne. The H100 is faster on the arithmetic operations (the per-particle advection and sedimentation, the explicit phase-change substeps, and the dynamical core) but slower on the branch-heavy iterative kernels. These are the backward-Euler condensation solver and, most of all, the per-particle collision update, which draws a random number and evaluates an extended sequence of conditional branches for each candidate pair. The collision share therefore rises from 2% of the super-droplet cost on the MI300A to nearly 20% on the H100. The two accelerators reach a comparable time to solution through different balances of the same operations.

The wall times and scaling for the three stages of the collisional processes are shown in Fig. 17. They couple the super-droplets within a cell. The particles are first sorted into their collision cells by a counting sort, which is parallel over the particles but relies on atomic counters and a prefix sum. The candidate pairs are then formed within each cell, parallel over the grid cells, which are far fewer than the particles. Each pair is finally tested for a collision and, if one occurs, the attributes of its two members are updated. The pairs are disjoint, so this last stage is parallel over the pairs with no conflicts. The total work is similar to other per-particle processes, but the parallelism of the pairing is coarser and carries more synchronization. The pair update is the largest, and it scales near-ideally. The pairing scales super-ideally, as its per-cell loop shrinks with the particles per cell. The binning is the exception. Its counting sort and within-cell shuffle stop improving below a few tenths of a millisecond per step, limited by a per-tile host-to-device transfer in the shuffle. As a result, the binning share climbs from 15 to 23% across the ladder while staying small in absolute terms.

Two operations couple the super-droplets to the Eulerian fields, both shown in their two-dimensional equivalent in Fig. 18(a); in three dimensions each stencil spans the eight surrounding cell centers. Mesh-to-particle interpolation carries an ambient field ϕ from its cell values $\phi_{i,j}$ to the particle as $\phi(\mathbf{x}_i) = \sum w_{i,j} \phi_{i,j}$, where the trilinear weight $w_{i,j}$ of a node is the fractional volume of the diagonally opposite sub-cell. Particle-to-mesh deposition spreads the represented condensate mass $\xi_i m_i$ of the



735 super-droplet to the same cell centers with the same weights $w_{i,j}$. Summing these contributions over all super-droplets and
 dividing by the air density yields the Eulerian mixing ratios of Sect. 3.6. Both operations thread over the particles. The inter-
 740 polation only reads the mesh and is embarrassingly parallel. The deposition accumulates its writes through atomic operations,
 and it is completed by a reduction that sums the contributions in the ghost cells across the neighboring processes. Figure 18(b)
 shows the scaling of the coupling operations. The particle-to-mesh reduction that forms the mixing ratios is the largest coupling
 cost, between a third and a half of the per-step time of the super-droplet method. It is evaluated repeatedly to refresh these fields
 between the phase-change substeps. This time is measured within the phase-change processes of Fig. 16, so it accounts for
 much of the phase-change share reported there. The source-term deposition, which spreads the microphysical vapor and heat
 sources onto the grid, is an order of magnitude smaller. The mesh-to-particle interpolation is performed inside the per-particle
 745 kernels rather than as a separate pass, so it has no separate timer. We take the particle advection, which interpolates the ambient
 fields to each super-droplet, as its proxy. All three scale near-ideally, at 70 to 85% of the ideal over the full range, because the
 interpolation and the deposition are local to each particle and its neighboring cells.

Figure 19 shows the fractional wall times of the compute and communication components of a time step. The particle
 computation is the largest share at the smaller counts, and the communication grows from 3 to 18% of the total between 8 and
 512 GPUs on Tuolumne as the per-process work shrinks. We observe a similar increase on Matrix; however, the maximum
 750 number of GPUs is small enough that the communication costs are not significant. Three kinds of communication contribute:
 the mesh halo exchange of the dynamical core, the ghost-cell reduction of the super-droplet deposition, and the redistribution
 of super-droplets among the processes. The redistribution scales the worst and is the most variable, because its message count
 grows with the process count and its size is proportional to the number of particles. Overall, the per-particle work of the
 super-droplet method is well-suited to a GPU. The processes are independent across super-droplets and involve many more
 755 floating-point operations relative to data movement. The per-attribute contiguous layout maintains coalesced memory access.
 A simulation exposes one independent thread per super-droplet, far more parallel work than the single thread per cell of a
 bulk scheme. GPU acceleration therefore benefits the particle work more than the memory-bound Eulerian core, and it is what
 makes particle-based microphysics tractable for large-scale simulations.

5.2 Memory requirement

760 Each super-droplet carries 18 double-precision attributes (144 bytes): the physical and kinematic attributes of Sect. 3, the ter-
 minal velocity and unique identifier of Sect. 3.7, and two cached quantities, the condensate mass and the ambient temperature,
 that are interpolated once per step and reused across the phase-change substeps (Sect. 3.4). The position and the AMReX-
 internal particle identifier, which is distinct from the unique identifier stored as an attribute, add 32 bytes, and an active flag
 adds 4 bytes, for a total of 180 bytes per super-droplet. The particle representation is therefore memory-intensive relative to an
 765 Eulerian scheme, which stores a small, fixed number of fields per cell. The GATE deep convection simulation is initialized with
 1.32×10^9 super-droplets that need approximately 240 GB. The requirement is largest at this initial state and decreases as the
 storm develops: super-droplets that precipitate out of the domain are deactivated and removed from memory (Sect. 3.3). The
 active count falls from 1.32×10^9 at the start to 1.09×10^9 at 36 h. The strong-scaling study of Sect. 5.1 restarts from the state at



23 h, near the collisional peak, where the simulation contains 1.22×10^9 active super-droplets that need approximately 220 GB of particle storage. Distributed over the accelerators, this is 27 GB per device on 8 GPUs and falls to 0.4 GB on 512 GPUs. The transient working memory is larger, since the particle redistribution and the input/output staging allocate communication buffers and copies on top of the resident storage. Across the strong-scaling study, the peak allocation is 1.2 to 1.3 times the resident particle storage over 8 to 128 GPUs. The simulation needs 33 GB per device at the lower end of that range. These memory requirements fit within the device memory: the AMD MI300A provides 128 GB of unified memory (Advanced Micro Devices, Inc., 2024) and the NVIDIA H100 provides 80 GB of high-bandwidth memory per device (NVIDIA Corporation, 2023). The particle data size exceeds the mesh data size by more than an order of magnitude, so the super-droplets, rather than the Eulerian fields, set the number of devices needed for a simulation. The requirement grows linearly with the domain size and with the number of super-droplets per cell.

6 Conclusions

We implement the super-droplet method for cold, mixed-phase microphysics in the Energy Research and Forecasting (ERF) code. The implementation resolves immersion freezing, vapor deposition and sublimation, riming, aggregation, and melting, and it represents ice particles as porous spheroids whose morphology evolves with their growth history. Melting is gradual: each super-droplet carries a mixed-phase state as it descends through the melting level. The particle data structures and compute kernels are portable across NVIDIA and AMD graphics processing units (GPUs), through the AMReX framework on which the implementation is built. We verify each process against theoretical solutions and published benchmarks, including the depth of the melting layer, the fall speed of melting particles, and the rates of sublimation, freezing, and riming. Applied to three-dimensional deep convection (GATE), the method reproduces the structure of a deep convective cloud without predefined hydrometeor categories or mass–dimension relationships. It predicts the ice-particle morphology directly from the per-particle masses. It also recovers the phase structure of the cloud: an ice layer aloft, a supercooled mixed-phase layer through the updraft, and warm rain below the freezing level. Two ten-member ensembles are used to measure the run-to-run variability of the stochastic method. The hydrometeor profiles of the mature cloud nearly coincide across the realizations, while the nascent rain of the developing cloud varies more strongly, because its onset is rate-limited by the stochastic collisions.

This increased fidelity of the super-droplet method is accompanied by a higher computational cost. The method is more expensive than an Eulerian bulk scheme, both in computations and in memory requirements. The memory requirement is dominated by the particle data. However, its structure is well-suited for parallelizing on GPUs. The GATE deep-convection case involved more than a billion super-droplets, and it fits within the memory of a modest number of GPUs. The particle operations scale near-ideally, retaining 84% and 63% of the ideal at 256 and 512 GPUs, respectively. Phase change dominates the cost, specifically the condensation and evaporation of liquid water, which integrate a stiff growth equation for every particle with an implicit solver. The stochastic collisions, by contrast, are relatively inexpensive, because each cell contains few particles and every stage of the collision is local to the cell. However, this depends on the GPU itself, with some key differences in relative cost between the NVIDIA H100 and AMD MI300A. Overall, the particle-based microphysics is not the bottleneck with a



sufficient number of GPUs. Its cost is then comparable to that of the Eulerian dynamical core. This number (for example, 256 for the GATE simulation here) is accessible on most high-performance computing platforms.

This paper focused on the individual verification of the processes as well as a demonstration of this method with a benchmark atmospheric flow simulation. These results are promising, but several extensions and verifications are planned for the future. The model omits collisional raindrop breakup, secondary ice production, and the shedding of meltwater from large melting particles; these processes will be added. We will quantify the sensitivity of the coupled solution to the number of super-droplets per cell. We will assess the various empirical models used in the constituent processes (for example, ice aggregation efficiency, ice deposition model) against other, more recent approaches in the literature. We will simulate other benchmarks, compare our results with measurements and other published models, and participate in model inter-comparison studies.

Code and data availability. The super-droplet method described here is implemented and available in ERF (<https://github.com/erf-model/ERF>), distributed under a BSD 3-Clause license with a US Department of Energy clause. The exact version used for every simulation reported here is archived at <https://doi.org/10.5281/zenodo.22050730> (Ghosh et al., 2026). It contains the ERF source tree, with the AMReX and RRTMGP submodules in full, the input files of every case reported here, and the postprocessing scripts used to analyze the solutions and draw the figures. The simulation outputs are not archived, so a figure is reproduced by running the archived input file for its case and then running that case's scripts.

Author contributions. DG implemented the cold-microphysics super-droplet method in ERF, set up and ran the verification tests, ran and analyzed the deep-convection simulations, and wrote the manuscript. HB contributed to the formulation of the ice microphysics, design and setup of the deep-convection simulation, the interpretation of the results, and review/editing of the paper. EKdJ contributed to the super-droplet implementation, analysis of the results, and review/editing of the paper. CJK contributed to the software implementation. DLM contributed to the formulation of the SDM implementation. DJG contributed to the software testing and project administration. KAL supervised the work, administered the project, and acquired funding.

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. We thank Ann Almgren and Aaron Lattanzi of Lawrence Berkeley National Laboratory for their help with the implementation and for their stewardship of the ERF model. We thank Andrew Myers of Lawrence Berkeley National Laboratory for help with the particle containers and related code in AMReX. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344, LLNL-JRNL-2014569. The United States Government retains, and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes. The authors used Claude (Anthropic), a generative artificial intelligence tool, during this work. It assisted with the de-



velopment and testing of the software, with the scripts that analyze the simulations and draw the figures, and with the editing and formatting of the manuscript text. The authors reviewed, revised, and verified all such content against the source code and the simulation data, and they take full responsibility for the content of this publication.

Financial support. This work was supported by Laboratory Directed Research and Development project 24-SI-001 at Lawrence Livermore
835 National Laboratory.



References

- Advanced Micro Devices, Inc.: AMD Instinct MI300A Accelerator, <https://www.amd.com/en/products/accelerators/instinct/mi300/mi300a.html>, product page; accessed June 2026, 2024.
- Almgren, A., Lattanzi, A., Haque, R., Jha, P., Kosovic, B., Mirocha, J., Perry, B., Quon, E., Sanders, M., Wiersema, D., Willcox, D., Yuan, X.,
 840 and Zhang, W.: ERF: Energy Research and Forecasting, *Journal of Open Source Software*, 8, 5202, <https://doi.org/10.21105/joss.05202>, 2023.
- Andrejczuk, M., Grabowski, W. W., Reisner, J., and Gadian, A.: Cloud-aerosol interactions for boundary layer stratocumulus in the Lagrangian Cloud Model, *Journal of Geophysical Research: Atmospheres*, 115, <https://doi.org/10.1029/2010jd014248>, 2010.
- Arabas, S. and Shima, S.-i.: Large-Eddy Simulations of Trade Wind Cumuli Using Particle-Based Microphysics with Monte Carlo Coalescence, *Journal of the Atmospheric Sciences*, 70, 2768–2777, <https://doi.org/10.1175/jas-d-12-0295.1>, 2013.
 845
- Arabas, S., Jaruga, A., Pawlowska, H., and Grabowski, W. W.: libcloudph++ 1.0: a single-moment bulk, double-moment bulk, and particle-based warm-rain microphysics library in C++, *Geoscientific Model Development*, 8, 1677–1707, <https://doi.org/10.5194/gmd-8-1677-2015>, 2015.
- Arakawa, A. and Lamb, V. R.: Computational design of the basic dynamical processes of the UCLA general circulation model, in: *Methods*
 850 *in Computational Physics: Advances in Research and Applications*, vol. 17, pp. 173–265, Elsevier, 1977.
- Balsara, D. S. and Shu, C.-W.: Monotonicity Preserving Weighted Essentially Non-oscillatory Schemes with Increasingly High Order of Accuracy, *Journal of Computational Physics*, 160, 405–452, <https://doi.org/10.1006/jcph.2000.6443>, 2000.
- Bartman, P., Bulenok, O., Górski, K., Jaruga, A., Łazarski, G., Olesik, M. A., Piasecki, B., Singer, C. E., Talar, A., and Arabas, S.: PySDM v1: particle-based cloud modeling package for warm-rain microphysics and aqueous chemistry, *Journal of Open Source Software*, 7, 3219,
 855 <https://doi.org/10.21105/joss.03219>, 2022.
- Bayley, C. J. A., Kölling, T., Naumann, A. K., Vogel, R., and Stevens, B.: CLEO: the fundamental design of a new superdroplet model for high computational performance, *Geoscientific Model Development*, 19, 6099–6120, <https://doi.org/10.5194/gmd-19-6099-2026>, 2026a.
- Bayley, C. J. A., Naumann, A. K., Poydenot, F., Vogel, R., Stevens, B., and Shima, S.-I.: CLEO: the numerical methods of a new superdroplet model including a droplet breakup algorithm, *Geoscientific Model Development*, 19, 6121–6142, <https://doi.org/10.5194/gmd-19-6121-2026>, 2026b.
 860
- Beard, K. V.: Terminal Velocity and Shape of Cloud and Precipitation Drops Aloft, *Journal of the Atmospheric Sciences*, 33, 851–864, [https://doi.org/10.1175/1520-0469\(1976\)033<0851:tvasoc>2.0.co;2](https://doi.org/10.1175/1520-0469(1976)033<0851:tvasoc>2.0.co;2), 1976.
- Beard, K. V. and Grover, S. N.: Numerical Collision Efficiencies for Small Raindrops Colliding with Micron Size Particles, *Journal of the Atmospheric Sciences*, 31, 543–550, [https://doi.org/10.1175/1520-0469\(1974\)031<0543:ncefsr>2.0.co;2](https://doi.org/10.1175/1520-0469(1974)031<0543:ncefsr>2.0.co;2), 1974.
- Berry, E. X.: Cloud Droplet Growth by Collection, *Journal of the Atmospheric Sciences*, 24, 688–701, [https://doi.org/10.1175/1520-0469\(1967\)024<0688:cdgbc>2.0.co;2](https://doi.org/10.1175/1520-0469(1967)024<0688:cdgbc>2.0.co;2), 1967.
 865
- Beydoun, H., Ghosh, D., De Jong, E. K., McGuffin, D., Kendrick, C., Lee, J., Gardner, D. J., and Lundquist, K. A.: Super Droplet Cloud Microphysics in the Energy Research and Forecasting Model, *ESS Open Archive [preprint]*, <https://doi.org/10.22541/essoar.176858434.43447149/v1>, 2026.
- Bird, G. A.: *Molecular Gas Dynamics and the Direct Simulation of Gas Flows*, Clarendon Press, Oxford, 1994.
- Birdsall, C. K. and Langdon, A. B.: *Plasma Physics via Computer Simulation*, Adam Hilger, Bristol, 1991.



- Böhm, H. P.: A General Equation for the Terminal Fall Speed of Solid Hydrometeors, *Journal of the Atmospheric Sciences*, 46, 2419–2427, [https://doi.org/10.1175/1520-0469\(1989\)046<2419:ageftt>2.0.co;2](https://doi.org/10.1175/1520-0469(1989)046<2419:ageftt>2.0.co;2), 1989.
- 875 Brdar, S. and Seifert, A.: McSnow: A Monte-Carlo Particle Model for Riming and Aggregation of Ice Particles in a Multidimensional Microphysical Phase Space, *Journal of Advances in Modeling Earth Systems*, 10, 187–206, <https://doi.org/10.1002/2017ms001167>, 2018.
- Busmann, M., Burau, H., Cowan, T. E., Debus, A., Huebl, A., Juckeland, G., Kluge, T., Nagel, W. E., Pausch, R., Schmitt, F., Schramm, U., Schuchart, J., and Widera, R.: Radiative Signatures of the Relativistic Kelvin-Helmholtz Instability, in: *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis (SC '13)*, pp. 1–12, ACM, <https://doi.org/10.1145/2503210.2504564>, 2013.
- 880 Chen, J.-P. and Lamb, D.: The Theoretical Basis for the Parameterization of Ice Crystal Habits: Growth by Vapor Deposition, *Journal of the Atmospheric Sciences*, 51, 1206–1222, [https://doi.org/10.1175/1520-0469\(1994\)051<1206:ttbftp>2.0.co;2](https://doi.org/10.1175/1520-0469(1994)051<1206:ttbftp>2.0.co;2), 1994.
- Dawson, J. M.: Particle simulation of plasmas, *Reviews of Modern Physics*, 55, 403–447, <https://doi.org/10.1103/RevModPhys.55.403>, 1983.
- de Jong, E., Mackay, J. B., Bulenok, O., Jaruga, A., and Arabas, S.: Breakups are complicated: an efficient representation of collisional breakup in the superdroplet method, *Geoscientific Model Development*, 16, 4193–4211, <https://doi.org/10.5194/gmd-16-4193-2023>, 2023.
- 885 Deardorff, J. W.: Stratocumulus-capped mixed layers derived from a three-dimensional model, *Boundary-Layer Meteorology*, 18, 495–527, <https://doi.org/10.1007/bf00119502>, 1980.
- Dukowicz, J. K.: A particle-fluid numerical model for liquid sprays, *Journal of Computational Physics*, 35, 229–253, [https://doi.org/10.1016/0021-9991\(80\)90087-X](https://doi.org/10.1016/0021-9991(80)90087-X), 1980.
- Dziekan, P. and Pawlowska, H.: Stochastic coalescence in Lagrangian cloud microphysics, *Atmospheric Chemistry and Physics*, 17, 13 509–13 520, <https://doi.org/10.5194/acp-17-13509-2017>, 2017.
- 890 Dziekan, P. and Zmijewski, P.: University of Warsaw Lagrangian Cloud Model (UWLCM) 2.0: adaptation of a mixed Eulerian–Lagrangian numerical model for heterogeneous computing clusters, *Geoscientific Model Development*, 15, 4489–4501, <https://doi.org/10.5194/gmd-15-4489-2022>, 2022.
- Dziekan, P., Waruszewski, M., and Pawlowska, H.: University of Warsaw Lagrangian Cloud Model (UWLCM) 1.0: a modern large-eddy simulation tool for warm cloud modeling with Lagrangian microphysics, *Geoscientific Model Development*, 12, 2587–2606, <https://doi.org/10.5194/gmd-12-2587-2019>, 2019.
- 895 ERF Development Team: ERF: Energy Research and Forecasting, <https://github.com/erf-model/ERF>, open-source software, 2025.
- Erfani, E. and Mitchell, D. L.: Growth of ice particle mass and projected area during riming, *Atmospheric Chemistry and Physics*, 17, 1241–1257, <https://doi.org/10.5194/acp-17-1241-2017>, 2017.
- 900 Field, P. R., Heymsfield, A. J., and Bansemer, A.: A Test of Ice Self-Collection Kernels Using Aircraft Data, *Journal of the Atmospheric Sciences*, 63, 651–666, <https://doi.org/10.1175/JAS3653.1>, 2006.
- Field, P. R., Lawson, R. P., Brown, P. R. A., Lloyd, G., Westbrook, C., Moisseev, D., Miltenberger, A., Nenes, A., Blyth, A., Choulaton, T., Connolly, P., Buehl, J., Crosier, J., Cui, Z., Dearden, C., DeMott, P., Flossmann, A., Heymsfield, A., Huang, Y., Kalesse, H., Kanji, Z. A., Korolev, A., Kirchgaessner, A., Lasher-Trapp, S., Leisner, T., McFarquhar, G., Phillips, V., Stith, J., and Sullivan, S.: Secondary Ice Production: Current State of the Science and Recommendations for the Future, *Meteorological Monographs*, 58, 7.1–7.20, <https://doi.org/10.1175/AMSMONOGRAPH-D-16-0014.1>, 2017.
- 905 Frick, C., Seifert, A., and Wernli, H.: A bulk parametrization of melting snowflakes with explicit liquid water fraction for the COSMO model, *Geoscientific Model Development*, 6, 1925–1939, <https://doi.org/10.5194/gmd-6-1925-2013>, 2013.



- Ghosh, D., Beydoun, H., de Jong, E. K., Kendrick, C. J., McGuffin, D. L., Gardner, D. J., and Lundquist, K. A.: ERF-SDM: a scalable,
 910 GPU-portable super-droplet method for mixed-phase cold microphysics, Zenodo [code], <https://doi.org/10.5281/zenodo.22050730>, 2026.
- Gillespie, D. T.: The Stochastic Coalescence Model for Cloud Droplet Growth, *Journal of the Atmospheric Sciences*, 29, 1496–1510,
[https://doi.org/10.1175/1520-0469\(1972\)029<1496:TSCMFC>2.0.CO;2](https://doi.org/10.1175/1520-0469(1972)029<1496:TSCMFC>2.0.CO;2), 1972.
- Gillespie, D. T.: An Exact Method for Numerically Simulating the Stochastic Coalescence Process in a Cloud, *Journal of the Atmospheric
 Sciences*, 32, 1977–1989, [https://doi.org/10.1175/1520-0469\(1975\)032<1977:AEMFNS>2.0.CO;2](https://doi.org/10.1175/1520-0469(1975)032<1977:AEMFNS>2.0.CO;2), 1975.
- 915 Golovin, A. M.: The solution of the coagulation equation for cloud droplets in a rising air current, *Izvestiya Akademii Nauk SSSR, Seriya
 Geofizicheskaya*, pp. 482–487, 1963.
- Govindarajan, B., Leishman, J. G., and Gumerov, N. A.: Particle-Clustering Algorithms for the Prediction of Brownout Dust Clouds, *AIAA
 Journal*, 51, 1080–1094, <https://doi.org/10.2514/1.J051907>, 2013.
- Grabowski, W. W., Wu, X., and Moncrieff, M. W.: Cloud-Resolving Modeling of Tropical Cloud Systems during Phase III of
 920 GATE. Part I: Two-Dimensional Experiments, *Journal of the Atmospheric Sciences*, 53, 3684–3709, [https://doi.org/10.1175/1520-0469\(1996\)053<3684:crmotc>2.0.co;2](https://doi.org/10.1175/1520-0469(1996)053<3684:crmotc>2.0.co;2), 1996.
- Grabowski, W. W., Wu, X., Moncrieff, M. W., and Hall, W. D.: Cloud-Resolving Modeling of Cloud Systems during Phase III
 of GATE. Part II: Effects of Resolution and the Third Spatial Dimension, *Journal of the Atmospheric Sciences*, 55, 3264–3282,
[https://doi.org/10.1175/1520-0469\(1998\)055<3264:crmocs>2.0.co;2](https://doi.org/10.1175/1520-0469(1998)055<3264:crmocs>2.0.co;2), 1998.
- 925 Grabowski, W. W., Morrison, H., Shima, S.-I., Abade, G. C., Dziekan, P., and Pawlowska, H.: Modeling of Cloud Microphysics: Can We Do
 Better?, *Bulletin of the American Meteorological Society*, 100, 655–672, <https://doi.org/10.1175/bams-d-18-0005.1>, 2019.
- Gunn, R. and Kinzer, G. D.: The Terminal Velocity of Fall for Water Droplets in Stagnant Air, *Journal of Meteorology*, 6, 243–248,
[https://doi.org/10.1175/1520-0469\(1949\)006<0243:TTVOFF>2.0.CO;2](https://doi.org/10.1175/1520-0469(1949)006<0243:TTVOFF>2.0.CO;2), 1949.
- Hall, W. D.: A Detailed Microphysical Model Within a Two-Dimensional Dynamic Framework: Model Description and Preliminary Results,
 930 *Journal of the Atmospheric Sciences*, 37, 2486–2507, [https://doi.org/10.1175/1520-0469\(1980\)037<2486:admmwa>2.0.co;2](https://doi.org/10.1175/1520-0469(1980)037<2486:admmwa>2.0.co;2), 1980.
- Hallett, J. and Mossop, S. C.: Production of secondary ice particles during the riming process, *Nature*, 249, 26–28,
<https://doi.org/10.1038/249026a0>, 1974.
- Heymsfield, A. J. and Pflaum, J. C.: A Quantitative Assessment of the Accuracy of Techniques for Calculating Graupel Growth, *Journal of
 the Atmospheric Sciences*, 42, 2264–2274, [https://doi.org/10.1175/1520-0469\(1985\)042<2264:aqaoa>2.0.co;2](https://doi.org/10.1175/1520-0469(1985)042<2264:aqaoa>2.0.co;2), 1985.
- 935 Hill, A. A., Lebo, Z. J., Andrejczuk, M., Arabas, S., Dziekan, P., Field, P., Gettelman, A., Hoffmann, F., Pawlowska, H., Onishi,
 R., and Vié, B.: Toward a Numerical Benchmark for Warm Rain Processes, *Journal of the Atmospheric Sciences*, 80, 1329–1359,
<https://doi.org/10.1175/JAS-D-21-0275.1>, 2023.
- Hockney, R. W. and Eastwood, J. W.: *Computer Simulation Using Particles*, Adam Hilger, Bristol, 1988.
- Hong, S.-Y. and Lim, J.-O. J.: The WRF Single-Moment 6-Class Microphysics Scheme (WSM6), *Journal of the Korean Meteorological
 940 Society*, 42, 129–151, 2006.
- Jensen, A. A. and Harrington, J. Y.: Modeling Ice Crystal Aspect Ratio Evolution during Riming: A Single-Particle Growth Model, *Journal
 of the Atmospheric Sciences*, 72, 2569–2590, <https://doi.org/10.1175/JAS-D-14-0297.1>, 2015.
- Jiang, G.-S. and Shu, C.-W.: Efficient implementation of weighted ENO schemes, *Journal of Computational Physics*, 126, 202–228,
<https://doi.org/10.1006/jcph.1996.0130>, 1996.
- 945 Kessler, E.: On the Distribution and Continuity of Water Substance in Atmospheric Circulations, *Meteorological Monographs*, 10, 1–84,
https://doi.org/10.1007/978-1-935704-36-2_1, 1969.



- Khain, A. P., Beheng, K. D., Heymsfield, A., Korolev, A., Krichak, S. O., Levin, Z., Pinsky, M., Phillips, V., Prabhakaran, T., Teller, A., van den Heever, S. C., and Yano, J.-I.: Representation of microphysical processes in cloud-resolving models: Spectral (bin) microphysics versus bulk parameterization, *Reviews of Geophysics*, 53, 247–322, <https://doi.org/10.1002/2014rg000468>, 2015.
- 950 Khairoutdinov, M. F. and Randall, D. A.: Cloud Resolving Modeling of the ARM Summer 1997 IOP: Model Formulation, Results, Uncertainties, and Sensitivities, *Journal of the Atmospheric Sciences*, 60, 607–625, [https://doi.org/10.1175/1520-0469\(2003\)060<0607:crmota>2.0.co;2](https://doi.org/10.1175/1520-0469(2003)060<0607:crmota>2.0.co;2), 2003.
- Klemp, J. B., Skamarock, W. C., and Dudhia, J.: Conservative Split-Explicit Time Integration Methods for the Compressible Nonhydrostatic Equations, *Monthly Weather Review*, 135, 2897–2913, <https://doi.org/10.1175/mwr3440.1>, 2007.
- 955 Koop, T., Luo, B., Tsias, A., and Peter, T.: Water activity as the determinant for homogeneous ice nucleation in aqueous solutions, *Nature*, 406, 611–614, <https://doi.org/10.1038/35020537>, 2000.
- Lattanzi, A., Almgren, A., Quon, E., Natarajan, M., Kosovic, B., Mirocha, J., Perry, B., Wiersema, D., Willcox, D., Yuan, X., and Zhang, W.: ERF: Energy Research and Forecasting Model, *Journal of Advances in Modeling Earth Systems*, 17, <https://doi.org/10.1029/2024ms004884>, 2025.
- 960 Lin, Y.-L., Farley, R. D., and Orville, H. D.: Bulk Parameterization of the Snow Field in a Cloud Model, *Journal of Climate and Applied Meteorology*, 22, 1065–1092, [https://doi.org/10.1175/1520-0450\(1983\)022<1065:bpotsf>2.0.co;2](https://doi.org/10.1175/1520-0450(1983)022<1065:bpotsf>2.0.co;2), 1983.
- Livermore Computing, Lawrence Livermore National Laboratory: Tuolumne, <https://hpc.llnl.gov/hardware/compute-platforms/tuolumne>, 2024.
- Livermore Computing, Lawrence Livermore National Laboratory: Matrix, <https://hpc.llnl.gov/hardware/compute-platforms/matrix>, 2025.
- 965 Locatelli, J. D. and Hobbs, P. V.: Fall Speeds and Masses of Solid Precipitation Particles, *Journal of Geophysical Research*, 79, 2185–2197, <https://doi.org/10.1029/JC079i015p02185>, 1974.
- Matsushima, T., Nishizawa, S., and Shima, S.-i.: Overcoming computational challenges to realize meter- to submeter-scale resolution in cloud simulations using the super-droplet method, *Geoscientific Model Development*, 16, 6211–6245, <https://doi.org/10.5194/gmd-16-6211-2023>, 2023.
- 970 Mitra, S. K., Vohl, O., Ahr, M., and Pruppacher, H. R.: A Wind Tunnel and Theoretical Study of the Melting Behavior of Atmospheric Ice Particles. IV: Experiment and Theory for Snow Flakes, *Journal of the Atmospheric Sciences*, 47, 584–591, [https://doi.org/10.1175/1520-0469\(1990\)047<0584:awtats>2.0.co;2](https://doi.org/10.1175/1520-0469(1990)047<0584:awtats>2.0.co;2), 1990.
- Monin, A. S. and Obukhov, A. M.: Basic laws of turbulent mixing in the surface layer of the atmosphere, *Contributions of the Geophysical Institute of the Academy of Sciences of the USSR*, 24, 163–187, 1954.
- 975 Morrison, H. and Milbrandt, J. A.: Parameterization of Cloud Microphysics Based on the Prediction of Bulk Ice Particle Properties. Part I: Scheme Description and Idealized Tests, *Journal of the Atmospheric Sciences*, 72, 287–311, <https://doi.org/10.1175/JAS-D-14-0065.1>, 2015.
- Morrison, H., Thompson, G., and Tatarskii, V.: Impact of Cloud Microphysics on the Development of Trailing Stratiform Precipitation in a Simulated Squall Line: Comparison of One- and Two-Moment Schemes, *Monthly Weather Review*, 137, 991–1007, <https://doi.org/10.1175/2008mwr2556.1>, 2009.
- 980 Morrison, H., van Lier-Walqui, M., Fridlind, A. M., Grabowski, W. W., Harrington, J. Y., Hoose, C., Korolev, A., Kumjian, M. R., Milbrandt, J. A., Pawlowska, H., Posselt, D. J., Prat, O. P., Reimel, K. J., Shima, S.-I., van Diedenhoven, B., and Xue, L.: Confronting the Challenge of Modeling Cloud and Precipitation Microphysics, *Journal of Advances in Modeling Earth Systems*, 12, e2019MS001689, <https://doi.org/10.1029/2019ms001689>, 2020.



- 985 Morrison, H., Chandrakar, K. K., Shima, S.-I., Dziekan, P., and Grabowski, W. W.: Impacts of Stochastic Coalescence Variability on Warm Rain Initiation Using Lagrangian Microphysics in Box and Large-Eddy Simulations, *Journal of the Atmospheric Sciences*, 81, 1067–1093, <https://doi.org/10.1175/jas-d-23-0132.1>, 2024.
- Niemand, M., Möhler, O., Vogel, B., Vogel, H., Hoose, C., Connolly, P., Klein, H., Bingemer, H., DeMott, P., Skrotzki, J., and Leisner, T.: A Particle-Surface-Area-Based Parameterization of Immersion Freezing on Desert Dust Particles, *Journal of the Atmospheric Sciences*, 69, 3077–3092, <https://doi.org/10.1175/jas-d-11-0249.1>, 2012.
- 990 NVIDIA Corporation: NVIDIA H100 Tensor Core GPU, <https://www.nvidia.com/en-us/data-center/h100/>, product page; accessed June 2026, 2023.
- Phillips, V. T. J., Pokrovsky, A., and Khain, A.: The Influence of Time-Dependent Melting on the Dynamics and Precipitation Production in Maritime and Continental Storm Clouds, *Journal of the Atmospheric Sciences*, 64, 338–359, <https://doi.org/10.1175/jas3832.1>, 2007.
- 995 Powers, J. G., Klemp, J. B., Skamarock, W. C., Davis, C. A., Dudhia, J., Gill, D. O., Coen, J. L., Gochis, D. J., Ahmadov, R., Peckham, S. E., Grell, G. A., Michalakes, J., Trahan, S., Benjamin, S. G., Alexander, C. R., Dimego, G. J., Wang, W., Schwartz, C. S., Romine, G. S., Liu, Z., Snyder, C., Chen, F., Barlage, M. J., Yu, W., and Duda, M. G.: The Weather Research and Forecasting Model: Overview, System Efforts, and Future Directions, *Bulletin of the American Meteorological Society*, 98, 1717–1737, <https://doi.org/10.1175/bams-d-15-00308.1>, 2017.
- 1000 Pruppacher, H. R. and Klett, J. D.: *Microphysics of Clouds and Precipitation*, vol. 18 of *Atmospheric and Oceanographic Sciences Library*, Kluwer Academic Publishers, Dordrecht, 2nd edn., 1997.
- Rasmussen, R. M. and Heymsfield, A. J.: A Generalized Form for Impact Velocities Used to Determine Graupel Accretional Densities, *Journal of the Atmospheric Sciences*, 42, 2275–2279, [https://doi.org/10.1175/1520-0469\(1985\)042<2275:agffiv>2.0.co;2](https://doi.org/10.1175/1520-0469(1985)042<2275:agffiv>2.0.co;2), 1985.
- Rasmussen, R. M. and Heymsfield, A. J.: Melting and Shedding of Graupel and Hail. Part I: Model Physics, *Journal of the Atmospheric Sciences*, 44, 2754–2763, [https://doi.org/10.1175/1520-0469\(1987\)044<2754:masoga>2.0.co;2](https://doi.org/10.1175/1520-0469(1987)044<2754:masoga>2.0.co;2), 1987.
- 1005 Riechelmann, T., Noh, Y., and Raasch, S.: A new method for large-eddy simulations of clouds with Lagrangian droplets including the effects of turbulent collision, *New Journal of Physics*, 14, 065 008, <https://doi.org/10.1088/1367-2630/14/6/065008>, 2012.
- Ryder, C. L., Marengo, F., Brooke, J. K., Estelles, V., Cotton, R., Formenti, P., McQuaid, J. B., Price, H. C., Liu, D., Ausset, P., Rosenberg, P. D., Taylor, J. W., Choularton, T., Bower, K., Coe, H., Gallagher, M., Crosier, J., Lloyd, G., Highwood, E. J., and Murray, B. J.: Coarse-mode mineral dust size distributions, composition and optical properties from AER-D aircraft measurements over the tropical eastern Atlantic, *Atmospheric Chemistry and Physics*, 18, 17 225–17 257, <https://doi.org/10.5194/acp-18-17225-2018>, 2018.
- 1010 Sato, Y., Shima, S.-i., and Tomita, H.: Numerical Convergence of Shallow Convection Cloud Field Simulations: Comparison Between Double-Moment Eulerian and Particle-Based Lagrangian Microphysics Coupled to the Same Dynamical Core, *Journal of Advances in Modeling Earth Systems*, 10, 1495–1512, <https://doi.org/10.1029/2018ms001285>, 2018.
- 1015 Seifert, A. and Beheng, K. D.: A two-moment cloud microphysics parameterization for mixed-phase clouds. Part 1: Model description, *Meteorology and Atmospheric Physics*, 92, 45–66, <https://doi.org/10.1007/s00703-005-0112-4>, 2006.
- Seifert, A., Leinonen, J., Siewert, C., and Kneifel, S.: The Geometry of Rimed Aggregate Snowflakes: A Modeling Study, *Journal of Advances in Modeling Earth Systems*, 11, 712–731, <https://doi.org/10.1029/2018MS001519>, 2019.
- Shima, S., Kusano, K., Kawano, A., Sugiyama, T., and Kawahara, S.: The super-droplet method for the numerical simulation of clouds and precipitation: a particle-based and probabilistic microphysics model coupled with a non-hydrostatic model, *Quarterly Journal of the Royal Meteorological Society*, 135, 1307–1320, <https://doi.org/10.1002/qj.441>, 2009.
- 1020



- Shima, S.-i., Sato, Y., Hashimoto, A., and Misumi, R.: Predicting the morphology of ice particles in deep convection using the super-droplet method: development and evaluation of SCALE-SDM 0.2.5-2.2.0, -2.2.1, and -2.2.2, *Geoscientific Model Development*, 13, 4107–4157, <https://doi.org/10.5194/gmd-13-4107-2020>, 2020.
- 1025 Skamarock, W. C., Klemp, J. B., Duda, M. G., Fowler, L. D., Park, S.-H., and Ringler, T. D.: A Multiscale Nonhydrostatic Atmospheric Model Using Centroidal Voronoi Tessellations and C-Grid Staggering, *Monthly Weather Review*, 140, 3090–3105, <https://doi.org/10.1175/mwr-d-11-00215.1>, 2012.
- Smagorinsky, J.: General Circulation Experiments with the Primitive Equations: I. The Basic Experiment, *Monthly Weather Review*, 91, 99–164, [https://doi.org/10.1175/1520-0493\(1963\)091<0099:gcewtp>2.3.co;2](https://doi.org/10.1175/1520-0493(1963)091<0099:gcewtp>2.3.co;2), 1963.
- 1030 Su, C.-C., Smith, M. R., Kuo, F.-A., Wu, J.-S., Hsieh, C.-W., and Tseng, K.-C.: Large-scale simulations on multiple Graphics Processing Units (GPUs) for the direct simulation Monte Carlo method, *Journal of Computational Physics*, 231, 7932–7958, <https://doi.org/10.1016/j.jcp.2012.07.038>, 2012.
- Sölch, I. and Kärcher, B.: A large-eddy model for cirrus clouds with explicit aerosol and ice microphysics and Lagrangian ice particle tracking, *Quarterly Journal of the Royal Meteorological Society*, 136, 2074–2093, <https://doi.org/10.1002/qj.689>, 2010.
- 1035 Takahashi, T., Endoh, T., Wakahama, G., and Fukuta, N.: Vapor Diffusional Growth of Free-Falling Snow Crystals between -3 and -23°C , *Journal of the Meteorological Society of Japan. Ser. II*, 69, 15–30, https://doi.org/10.2151/jmsj1965.69.1_15, 1991.
- Thompson, G., Field, P. R., Rasmussen, R. M., and Hall, W. D.: Explicit Forecasts of Winter Precipitation Using an Improved Bulk Microphysics Scheme. Part II: Implementation of a New Snow Parameterization, *Monthly Weather Review*, 136, 5095–5115, <https://doi.org/10.1175/2008mwr2387.1>, 2008.
- 1040 Tzivion, S., Feingold, G., and Levin, Z.: An Efficient Numerical Solution to the Stochastic Collection Equation, *Journal of the Atmospheric Sciences*, 44, 3139–3149, [https://doi.org/10.1175/1520-0469\(1987\)044<3139:aenst>2.0.co;2](https://doi.org/10.1175/1520-0469(1987)044<3139:aenst>2.0.co;2), 1987.
- Unterstrasser, S., Hoffmann, F., and Lerch, M.: Collection/aggregation algorithms in Lagrangian cloud microphysical models: rigorous evaluation in box model simulations, *Geoscientific Model Development*, 10, 1521–1548, <https://doi.org/10.5194/gmd-10-1521-2017>, 2017.
- Welss, J.-N., Siewert, C., and Seifert, A.: Explicit Habit-Prediction in the Lagrangian Super-Particle Ice Microphysics Model McSnow, *Journal of Advances in Modeling Earth Systems*, 16, e2023MS003 805, <https://doi.org/10.1029/2023MS003805>, 2024.
- 1045 Wicker, L. J. and Skamarock, W. C.: Time-Splitting Methods for Elastic Models Using Forward Time Schemes, *Monthly Weather Review*, 130, 2088–2097, [https://doi.org/10.1175/1520-0493\(2002\)130<2088:tsmfem>2.0.co;2](https://doi.org/10.1175/1520-0493(2002)130<2088:tsmfem>2.0.co;2), 2002.
- Zhang, W., Myers, A., Gott, K., Almgren, A., and Bell, J.: AMReX: Block-structured adaptive mesh refinement for multiphysics applications, *The International Journal of High Performance Computing Applications*, 35, 508–526, <https://doi.org/10.1177/10943420211022811>, 2021.
- 1050 Zängl, G., Reinert, D., Rípodas, P., and Baldauf, M.: The ICON (ICOsahedral Non-hydrostatic) modelling framework of DWD and MPI-M: Description of the non-hydrostatic dynamical core, *Quarterly Journal of the Royal Meteorological Society*, 141, 563–579, <https://doi.org/10.1002/qj.2378>, 2015.