High-resolution simulation of turbulent collision of cloud droplets

Bogdan Rosa[1], Hossein Parishani[2], Orlando Ayala[2], Lian-Ping Wang[2] and Wojciech W. Grabowski[3]

[1] Institute of Meteorology and Water Management
National Research Institute,
ul. Podlesna 61, 01-673 Warsaw, Poland
bogdan.rosa@imgw.pl,
[2] Department of Mechanical Engineering, 126 Spencer Laboratory,
University of Delaware, Newark, Delaware 19716-3140, USA
lwang@udel.edu,
WWW home page: http://www.me.udel.edu/∼lwang/
[3] Mesoscale and Microscale Meteorology Division,
National Center for Atmospheric Research, PO Box 3000, Boulder,
CO 80307-3000, USA

Abstract. A novel parallel implementation of hybrid DNS (Direct Numerical Simulation) code for simulating collision-coalescence of aerodynamically interacting particles in a turbulent flow has been developed. An important application of this code is to quantify turbulent collision-coalescence rate of cloud droplets, relevant to warm rain formation, under physically realistic conditions. The code enables performing high-resolution DNS of turbulent collisions so the simulation results can be used to begin addressing the question of Reynolds number dependence of pair and collision statistics. The new implementation is based on MPI (Message Passing Interface) library, and thus the code can run on computers with distributed memory. This development enables to conduct hybrid DNS with flow field solved at grid resolutions up to 512^3 while simultaneously track up to several million aerodynamically-interacting droplets. In this paper we discuss key elements of the MPI implementation and present preliminary results from the high resolution simulations. The key conclusion is that, for small cloud droplets, the results on pair statistics and collision kernel appear to reach their saturation values as the flow Reynolds number is increased.

1 Introduction

Turbulent collision-coalescence of cloud droplets is a necessary step for warm rain initiation and development [1]. Rain drops are initiated and further grow in size primarily by colliding with cloud droplets that result from water vapor condensation on cloud condensation nuclei. The small size and small inertial response time of cloud droplets imply that pair statistics (i.e. radial distribution function RDF or droplet relative velocity) relevant to collision-coalescence are

mainly determined by dissipation-range turbulence dynamics, making direct numerical simulation [2, 3] a meaningful approach for this particular application. Due to the computational requirements, the domain size of DNS is typically in the range of 10 cm to 1 m scale only. Such a domain covers the flow dissipation range but an insufficient inertial subrange. As droplet size is increased, some inertial subrange scales of fluid motion can also contribute to the pair statistics. It is thus desirable to systematically increase the range of flow scales covered in DNS and consequently the computational domain size [4] in order to fully simulate the dynamic interactions of droplets and small-scale turbulence. Increase in domain size also implies an increase in the total number of simulated droplets. Furthermore, a long simulation or multiple realizations are often necessary to ensure a small statistical uncertainty of the computed pair statistics.

2 Methodology

Modeling of collision-coalescence of aerodynamically interacting droplets moving in a turbulent flow is a challenging task due to high computational cost and numerical and physical complexities. Dynamic and kinematic statistics of turbulent collision of cloud droplets depend primarily on the small-scale turbulent flow characteristics (e.g., Kolmogorov scales), settling velocity, and particle inertia. Statistical uncertainties of these physical quantities depend on the number of droplets followed in time and space. Assuming the dissipation-range flow is fully resolved, increasing the domain size translates to a higher flow Reynolds number, thus making simulations closer to physical conditions in turbulent clouds. A larger domain requires also tracking of a larger number of droplets, under the condition of a prescribed liquid water content (LWC).

The above considerations motivated us to utilize modern supercomputers with architecture based on distributed memory. The key question is how to take advantages of large computational resources (i.e. CPUs and memory) through efficient parallel implementation. Our hybrid DNS combines a pseudo-spectral simulation of turbulent air flow on a fixed spatial grid (Eulerian representation) with numerical integration of the equation of motion for freely moving droplets (Lagrangian representation). This combination of a large number of degrees of freedom on fixed grid points and a large number of degrees of freedom associated with moving droplets presents a significant challenge to efficient parallel implementation, as the two representations create different data structures that require different data distribution methods on available processors. Another challenge is that the pseudo-spectral method involves three-dimensional Fast Fourier transform (FFT) which requires global (i.e., whole-domain) data access or global data communication. These factors make the code communication-intensive.

Several different parallelization techniques could be considered. For the grid-based flow simulation, the most logical method is domain decomposition. For particle tracking, two different strategies can be considered: the first assigns to each processor a fixed subset of particles and their movements are handled by integrating their equations of motion; alternatively, each processor can treat

a fixed subset of particle pairs and particle-particle interaction forces. Some quantitative evaluation of the efficiencies of these two approaches is given in [5].

Here we present our MPI implementation based on 1D domain decomposition. This approach distributes evenly tasks associated with the computation of the turbulent flow field. The standard pseudo-spectral method is used to integrate the incompressible Navier-Stokes equations in a cubic domain with periodic boundary conditions. The Navier-Stokes equations are transformed into the spectral space where time advancement is performed. Parallel implementation of the 3D FFT based on 1D domain decomposition was developed in [6].

Several implementation issues related to the computation of particle trajectories have to be resolved. The first is the interpolation of the fluid velocity at the location of each particle, from the velocities at the fixed grid points. The second is the proper implementation of periodic boundary conditions for moving particles. The third is sending and receiving particles between neighboring subdomains when they cross the subdomain boundaries. Another is the detection of collision events. These aspects and their parallel implementations were carefully discussed in [7]. When local droplet-droplet aerodynamic interactions are considered [2], disturbance flow around each droplet has to be considered when advancing the location of other droplets within a given distance. This requires a significant data communication between neighboring domains [8].

3 Simulation results

In this section, we present results obtained from our MPI code, including characteristics of the background turbulent flow field, and kinematic and dynamic statistics related to collision-coalescence of both aerodynamically-interacting and noninteracting droplets.

3.1 Background turbulent flow

First we present results of the simulated turbulent flows at different grid resolutions. Figure 1a shows Taylor microscale Reynolds number as a function of the grid resolution used. The Reynolds number is defined as $R_\lambda = \sqrt{15}u'^2\tau_k/\nu$ where u' is r.m.s. fluctuation velocity in a given direction, τ_k is Kolmogorov time and ν is fluid viscosity. We have utilized two different large-scale forcing schemes, namely, a stochastic forcing [9] and a deterministic forcing [10]. The purpose was to examine whether the collision and pair statistics are affected by the nature of large-scale forcing scheme. Using the deterministic forcing scheme we reached $R_\lambda = 205.2$ at grid size 512^3. Isihara et al. [14] and Donzis et al. [15] performed simulations with even larger Taylor microscale Reynolds numbers but their simulations were restricted to the single phase flow only. Figure 1b shows compensated energy spectra of the turbulent flows, using the stochastic forcing scheme, at several different mesh resolutions starting from 32^3 up to 1024^3. This plot demonstrates that the dissipation-range spectra completely overlap when $R_\lambda > 100$. The simulated flows at 512^3 and above also show that a portion of

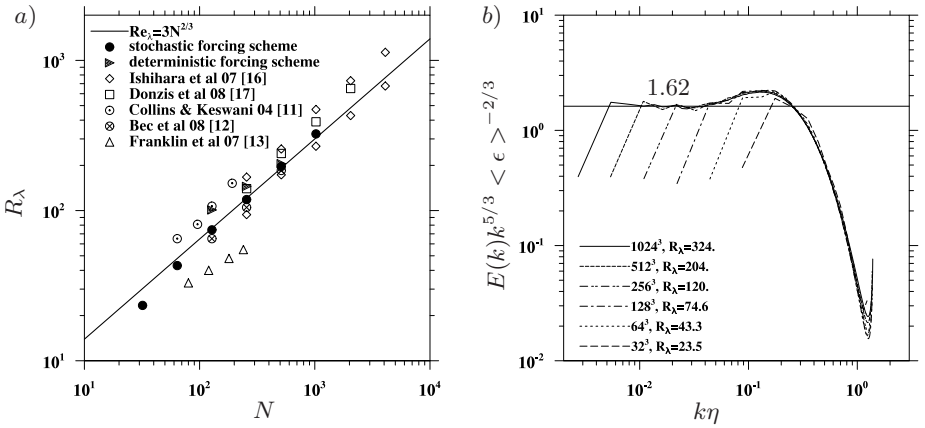

Fig. 1. *a*) Simulated flow Taylor microscale Reynolds number, R_λ, as a function of grid resolution and *b*) compensated energy spectra of the turbulent flows at different R_λ, using a stochastic large-scale forcing scheme. The line in *a*) indicates $R_\lambda = 3N^{2/3}$ [3].

the inertial sub-range scales are correctly represented, since the universal scaling (the line at 1.62) for the inertial subrange observed in experiments [16, 17] is reproduced. Clearly the range of simulated flow scales increases with the flow Reynolds number (or grid resolution).

In Table 1, we list the flow statistics and indicate the computational domain size if a physical dissipation rate of 400 $\mathrm{cm^2 s^{-3}}$ is assumed in a turbulent cloud. Also listed is the number of droplets in the computational domain for a prescribed LWC at 1 $\mathrm{g/m^3}$, assuming that the system is bidisperse with half of the particles are 20 μm in radius and the other half 30 μm. These cloud droplet sizes are most relevant to warm rain initiation, as condensation and gravitational collision are both slow in growing these droplets and additional mechanism such as effects of turbulence could become critical. A larger number of droplets will reduce statistical uncertainty. In our previous OpenMP implementation, the total number of droplets in one single simulation might be limited by the amount of available memory. The MPI code can handle a larger number of particles ($\sim 10^7$) without such a technical restriction.

3.2 Parallel performance

Parallel scalability of the new MPI implementation is evaluated using two modern supercomputers, *i.e.*, Lynx and Chimera. Lynx is a single cabinet Cray XT5m machine installed at National Center for Atmospheric Research (USA). The computer has seventy-six nodes, each with twelve processors, on two hex-core AMD 2.2 GHz Opteron chips. Chimera at the University of Delaware contains 3,168 cores which are grouped in 66 nodes. Each node has 4 CPUs AMD Opteron 6164HE 12-core running at 1.7GHz.

In the first test, the influence of compiler options on multi-CPU performance of our MPI code is investigated. The runs were conducted on Lynx employing two popular compilers, PGI and Intel. For each series of runs, all components

Table 1. Implications of increasing DNS grid resolutions. N is domain size (number of grid points in one direction), $< \epsilon >$ is the average energy dissipation rate in DNS units and u' is r.m.s. of fluid velocity

N	R_λ	$< \epsilon >$	Domain size (cm) ($400\ \mathrm{cm}^2/\mathrm{s}^3$)	u'	Number of droplets LWC = $1\mathrm{g/m}^3$
Stochastic forcing scheme					
32	23.5	3646	6.0	7.08	1.0×10^3
64	43.3	3529	11.9	9.61	8.0×10^3
128	74.6	3589	23.9	12.61	6.6×10^4
256	120.	3690	48.1	16.18	5.4×10^5
512	204.	3900	97.5	20.84	4.5×10^6
1024	324.	3777	193.0	26.29	3.5×10^7
Deterministic forcing scheme					
256	144.9	0.2011	26.32	17.56	2.6×10^5
512	205.2	0.2146	45.02	20.90	1.2×10^6

of the code were compiled with the highest optimization level. The simulations employed a grid resolution of 512^3 and 5 million of non-interacting particles. The total number of processors was fixed to 64, but the number of processes per node changed between runs. Figure 2a shows the average wall clock time per time step from each run. The wall clock times for the flow solver and for tracking particles are shown separately. We conclude that the total wallclock time with the Intel compiler is shorter by 6 do 12 % than that with the PGI compiler. The difference also depends on the distribution of processes between nodes and is mainly related to particle tracking. Additionally, Fig. 2a shows that the parallel efficiency drops (by more than 50%) as the number of processes per node is increased, due to memory bus saturation in this multi-core system.

In the second test, the speedup factors of the MPI implementation on two different machines are examined. The same testing problem is used, namely, 512^3 flow grid and 5 million particles. The runs on Lynx were performed for two different numbers of processes per node, *i.e.*, 12 which uses the least number of nodes, and 1 which corresponds to the best performance. The run with 128 processors on Lynx was done with 2 processes per node, to insufficient number of nodes. The runs employing all 12 processors per node were limited to 32 processes due to the memory constraint. Figure 2b displays results from both Lynx and Chimera. Also shown are the scalability data of the previous OpenMP implementation conducted on the IBM Power 575 cluster (4064 POWER6 processors running at 4.7 GHz) at NCAR.

3.3 Kinematic and dynamic statistics for droplets

Figure 3 shows single-droplet r.m.s. fluctuation velocities in horizontal *a*) and vertical *b*) direction, respectively. Thick lines represent the theoretical prediction from [18, 19]. For vertical direction the theoretical prediction have been developed following the procedure from [18] for the horizontal direction.

For most cases, the droplet r.m.s velocities approach the corresponding fluid r.m.s. velocity as the droplet radius is reduced. The horizontal droplet r.m.s

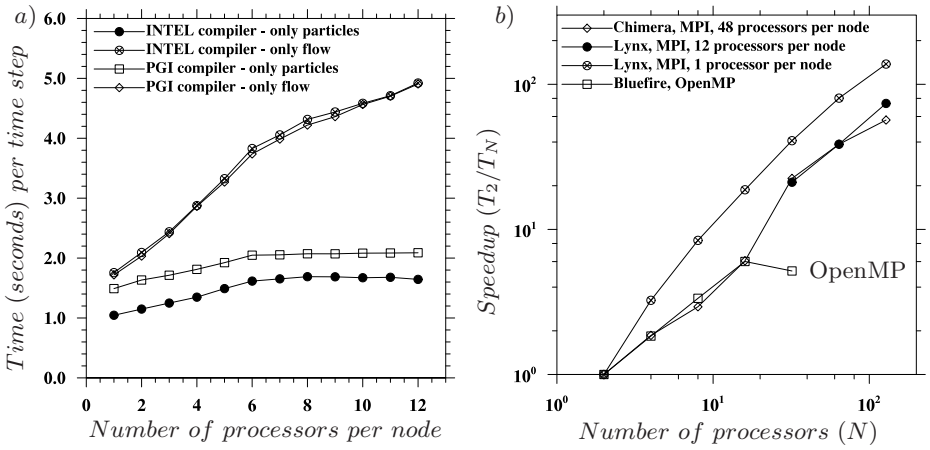

Fig. 2. Scalability data of the MPI code for simulations at 512^3 grid size with 5 million (non-interacting) particles. Panel *a*): wall clock time obtained with two different compilers on Lynx as a function of the number of processors per node (the total number of processes is fixed to 64). Panel *b*): speedup factors as a function of the number of processors. Also speedup data of previous OpenMP implementation performed on NCAR's Bluefire are plotted.

velocity drops more quickly with droplet size than the vertical velocity due to the sedimentation [19]. This difference becomes more evident at larger flow Reynolds numbers. Here the key point is that the single-droplet r.m.s. fluctuation velocity increases with flow Reynolds number or simulation domain size as the fluid r.m.s. velocity increases when more larger scales are included. The usual Reynolds number scaling is $u' \approx 0.5 v_k R_\lambda^{0.5}$ so the single-droplet r.m.s. velocity increases monotonically with R_λ.

Droplet-droplet pair statistics, such as radial distribution function (RDF) [20] and radial relative velocity are kinematic parameters directely proportional collision rate. The monodisperse pair statistics of nearly touching particles are shown in Figs. 4a and 4b, along with results at other flow Reynolds numbers. For larger droplets, the pair statistics increase with the flow Reynolds number. However, for any given droplet size, there is a tendency of saturation, namely, the pair statistics eventually become insensitive to flow Reynolds number. This saturation for smaller droplets is reached at smaller flow Reynolds number, since the range of flow scales affecting the pair statistics is more limited for smaller droplets due to their smaller Stokes number. For example, saturation is observed to occur at $R_\lambda \approx 200$ for 30 μm droplets. The gradual saturation of pair and collision statistics with R_λ is further demonstrated in Fig. 5, where the statistics are plotted as a function of R_λ. The dynamic collision kernel and RDF of the mondisperse system both show evidence of saturation. This observed saturation justifies the hybrid DNS approach using flow Reynolds numbers significantly less than those in real clouds.

Finally, we show pair statistics for aerodynamically interacting (AI) droplets [2, 21] in Fig. 6. The results are normalized by corresponding values without

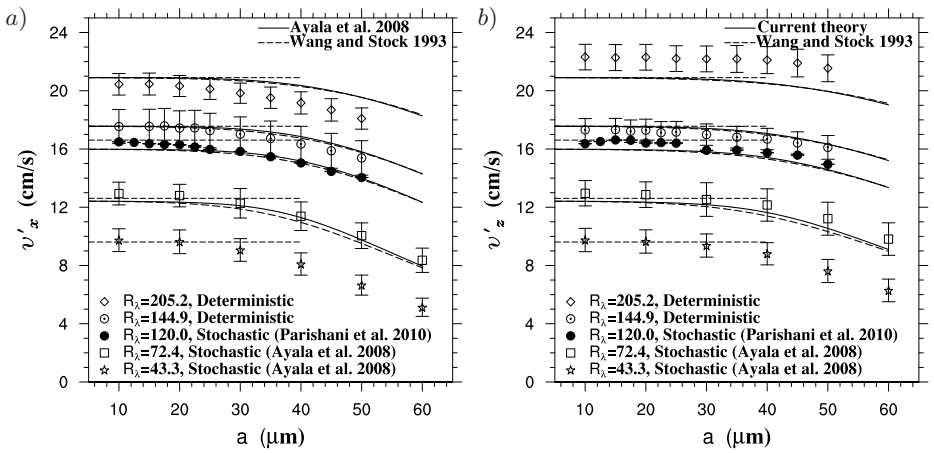

Fig. 3. Single-droplet r.m.s. fluctuation velocity in a) horizontal and b) vertical direction. Dashed thin lines mark r.m.s. fluctuation velocity of the fluid. Thicker lines represent the theoretical prediction from [18, 19].

aerodynamic interactions (no-AI). Here we simulate bidisperse system where one set of droplets is 50 μm in radius. The results show that the aerodynamic interaction has a strong influence on the kinematic statistics for the pairs of very different sizes, leading to a large reduction of $< w_{r,12} >$ and a significant increase in g_{12}. In this case, the disturbance flow of the larger droplet significantly alters the trajectory of the smaller droplet. When the pair are similar in size, the impact of aerodynamic interaction is relatively weak. While it is difficult to infer the precise effect of R_λ due to large statistical uncertainty, we may conclude that the flow Reynolds number effect is relatively weak for the cases considered here.

4 Summary and Conclusions

In this paper, we have reviewed a hybrid DNS approach designed to study turbulent collision-coalescence of cloud droplets. A general question is whether such an approach is justified for this application problem. As a first step to address this question, we have implemented MPI in order to increase the range of flow scales that can be simulated. We reviewed briefly various MPI implementation issues associated with the turbulent flow simulation and with tracking the motion and detecting pairwise interaction of droplets.

The MPI code is then used to simulate the flow and dynamics of droplets at higher flow Reynolds numbers, using up to 1024^3 grid for the flow only simulations and up to 512^3 grid for flows laden with droplets. The results show that the inertial subrange of turbulence can be correctly reproduced at high grid resolutions. As the flow Reynolds number or the computational domain size is increased, the range of flow scales is also increased, leading to increased single-droplet r.m.s. fluctuation velocities. However, we show that the pair statistics and the dynamic collision kernel will reach their saturated values, at least to

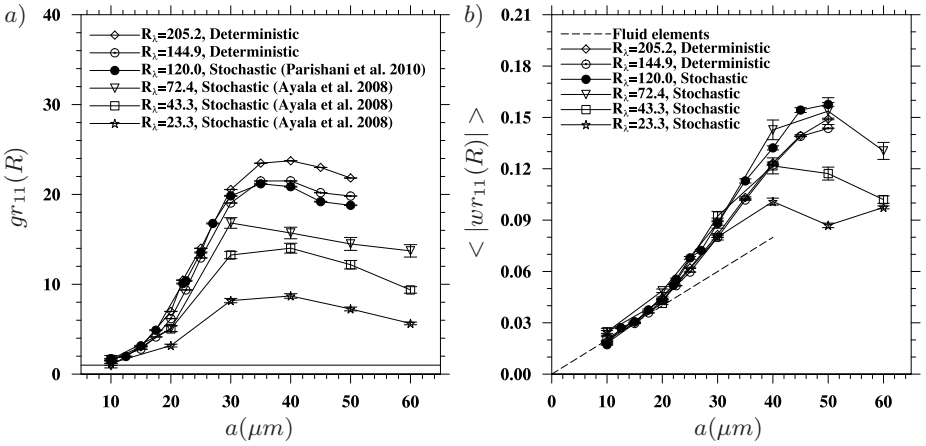

Fig. 4. *a*) Radial distribution function and *b*) relative velocity at contact $r = R$ for monodisperse pairs as a function of droplet size for different values of R_λ.

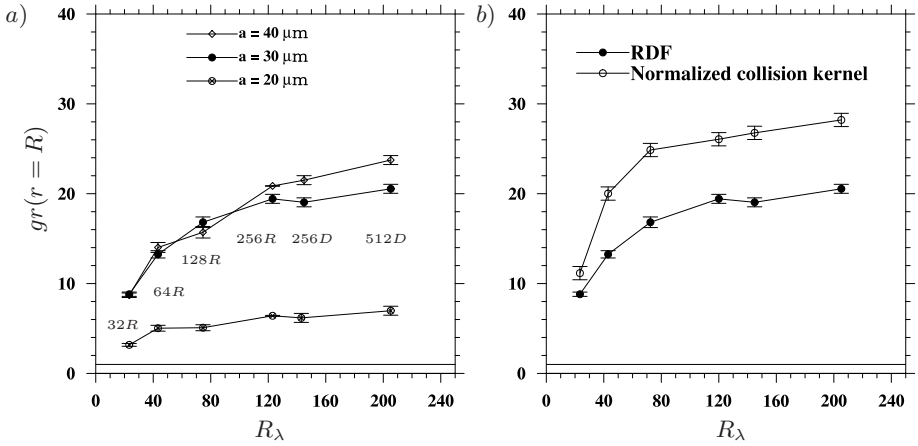

Fig. 5. *a*) Monodisperse RDF as a function of R_λ for three droplet sizes, *b*) dynamic collision kernel of the mondisperse system (30μm) for different R_λ normalized by $2\pi R v_{GRAV}$ where v_{GRAV} is terminal velocity of the particles in the stagnant flow

the leading order in R_λ, if all relevant scales of fluid motion are included in the flow simulation. This supports a fundamental assumption in the hybrid DNS, namely, hybrid DNS at much lower flow Reynolds numbers compared to those in real clouds can be used to quantity turbulent collision-coalescence of cloud droplets. The R_λ dependence of pair and collision statistics found in previous low-resolution simulations is a result of inadequate flow scale separation.

The MPI code is currently based on one-dimensional spatial domain decomposition. We are currently extending the MPI implementation using two-dimensional spatial domain decomposition, so that a much larger number of distributed-memory processors can be used, to further increase the flow Reynolds number or computational domain size. The ultimate goal is that all relevant

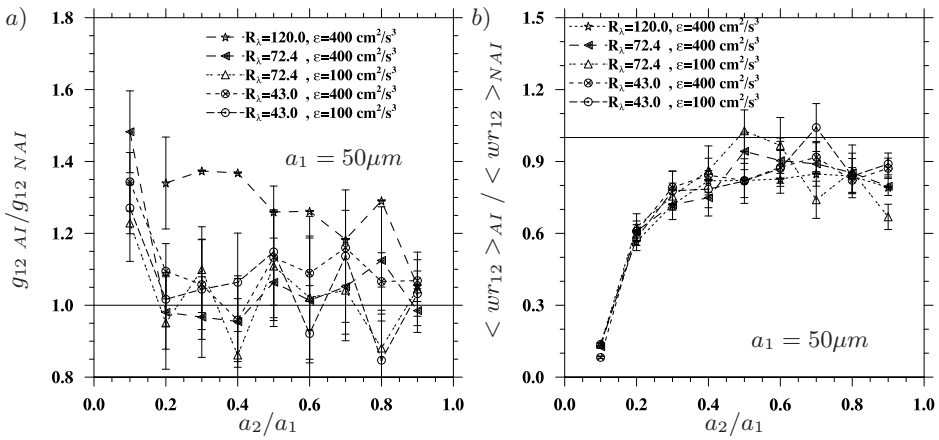

Fig. 6. Ratio of *a*) the radial distribution function and *b*) the radial relative velocity with AI to that corresponding to geometric collisions

scales of small-scale turbulence can be simulated so the pair statistics and pairwise interactions of cloud droplets of radii 100 μm or less can be fully studied.

Acknowledgements

This work was supported by the National Science Foundation (NSF) under grants ATM-0527140, ATM-0730766, OCI-0904534, and CRI-0958512. Computing resources are provided by National Center for Atmospheric Research (NCAR CISL-35751010, CISL-35751014, and CISL-35751015).

References

[1] L-P. Wang, Y. Xue, 0. Ayala, and W.W. Grabowski. Effects of stochastic coalescence and air turbulence on the size distribution of cloud droplets. *Atmospheric Research*, 82:416–432, 2006.

[2] O. Ayala, W-W. Grabowski, and L-P. Wang. A hybrid approach for simulating turbulent collisions of hydrodynamically-interacting particles. *JCP*, 225.

[3] L-P. Wang, B. Rosa, H. Gao, G.W. He, and G.D. Jin. Turbulent collision of inertial particles: point-particle based, hybrid simulations and beyond. *Int. J. Multiphase Flow*, 35:854–867, 2009.

[4] 0. Ayala, B. Rosa, L-P. Wang, and W.W. Grabowski. Effects of turbulence on the geometric collision rate of sedimenting droplets: Part 1. results from direct numerical simulation. *New J. Physics*, 10(075015), 2008.

[5] S. Plimpton. Fast parallel algorithms for short-range molecular-dynamics. *Journal of Computational Physics*, 117(1):1–19, 1995.

[6] P. Dmitruk, L. P. Wang, W. H. Matthaeus, R. Zhang, and D. Seckel. Scalable parallel FFT for spectral simulations on a Beowulf cluster. *Parallel Computing*, 27(14):1921–1936, 2001.

[7] B. Rosa and L-P. Wang. Parallel implementation of particle tracking and collision in a turbulent flow. *Lecture Notes in Computer Science*, pages 388–397, 2010.

[8] H. Parishani, B. Rosa, W-W Grabowski, and L-P. Wang. Towards high-resolution simulation of turbulent collision of cloud droplets. *Proceedings of the 7th International Conference on Multiphase Flow, Tampa, Fl, USA*, 2010.

[9] L. P. Wang and M. R. Maxey. Settling velocity and concentration distribution of heavy-particles in homogeneous isotropic turbulence. *Journal of Fluid Mechanics*, 256:27–68, 1993.

[10] L-P. Wang and B. Rosa. A spurious evolution of turbulence originated from round-off error in pseudo-spectral simulation. *Computer and Fluids*, 38:1943–49, 2009.

[11] L. R. Collins and A. Keswani. Reynolds number scaling of particle clustering in turbulent aerosols. *New Journal of Physics*, 6(119), 2004.

[12] J. Bec, L. Biferale, G. Bofftta, A. Celani, M. Cencini, A. Lanotte, S. Musacchio, and F. Toschi. Acceleration statistics of heavy particles in turbulence. *Journal of Fluid Mechanics*, 550:349–358, 2006.

[13] C. N. Franklin, P. A. Vaillancourt, and M. K. Yau. Statistics and parameterizations of the effect of turbulence on the geometric collision kernel of cloud droplets. *Journal of the Atmospheric Sciences*, 64(3):938–954, 2007.

[14] T. Ishihara, Y. Kaneda, M. Yokokawa, K. Itakura, and A. Uno. Small-scale statistics in high-resolution direct numerical simulation of turbulence: Reynolds number dependence of one-point velocity gradient statistics. *J. Fluid Mech.*, pages 335–366, 2007.

[15] D.A. Donzis, P.K. Yeung, and K.R. Sreenivasan. Dissipation and enstrophy in isotropic turbulence: resolution effects and scaling in direct numerical simulations. *Phys. Fluids*, 20(045108), 2008.

[16] T. Ishihara, T. Gotoh, and Kaneda. Study of high-reynolds number isotropic turbulence by direct numerical simulation. *Annual Review of Fluid Mechanics*, 41:165–180, 2009.

[17] K.R. Sreenivasan. On the universality of the Kolmogorov constant. *Physics of Fluids*, 7:2778–84, 1995.

[18] O. Ayala, B. Rosa, and L. P. Wang. Effects of turbulence on the geometric collision rate of sedimenting droplets. Part 2. Theory and parameterization. *New Journal of Physics*, 10, 2008.

[19] L-P. Wang and D.E. Stock. Dispersion of heavy particles by turbulent motion. *J. Atmos. Sci.*, 50(13), 1993.

[20] Zhou Yong Wang L-P., Wexler A. S. Statistical mechanical description and modelling of turbulent collision of inertial particles. *J. Fluid Mech.*, pages 117–153, 2000.

[21] L-P. Wang, 0. Ayala, S.E. Kasprzak, and W.W. Grabowski. Theoretical formulation of collision rate and collision efficiency of hydrodynamically-interacting cloud droplets in turbulent atmosphere of one-point velocity gradient statistics. *J. Atmos. Sci.*, 62:2433–2450, 2005.