

Towards an integrated multiscale simulation of turbulent clouds on PetaScale computers

Lian-Ping Wang¹, Orlando Ayala¹, Hossein Parishani¹, Wojciech W. Grabowski², Andrzej A. Wyszogrodzki², Zbigniew Piotrowski^{2,3}, Guang R. Gao¹, Chandra Kambhamettu¹, Xiaoming Li¹, Louis Rossi¹, Daniel Orozco¹ & Claudio Torres¹

¹University of Delaware, Newark, Delaware, USA

²National Center for Atmospheric Research, Boulder, Colorado, USA

³Institute of Meteorology and Water Management, Warsaw, Poland

E-mail: lwang@udel.edu

Abstract. The development of precipitating warm clouds is affected by several effects of small-scale air turbulence including enhancement of droplet-droplet collision rate by turbulence, entrainment and mixing at the cloud edges, and coupling of mechanical and thermal energies at various scales. Large-scale computation is a viable research tool for quantifying these multiscale processes. Specifically, top-down large-eddy simulations (LES) of shallow convective clouds typically resolve scales of turbulent energy-containing eddies while the effects of turbulent cascade toward viscous dissipation are parameterized. Bottom-up hybrid direct numerical simulations (HDNS) of cloud microphysical processes resolve fully the dissipation-range flow scales but only partially the inertial subrange scales. It is desirable to systematically decrease the grid length in LES and increase the domain size in HDNS so that they can be better integrated to address the full range of scales and their coupling. In this paper, we discuss computational issues and physical modeling questions in expanding the ranges of scales realizable in LES and HDNS, and in bridging LES and HDNS. We review our on-going efforts in transforming our simulation codes towards PetaScale computing, in improving physical representations in LES and HDNS, and in developing better methods to analyze and interpret the simulation results.

1. Introduction

Reliable weather and climate prediction at both local and global scales depends on our understanding of microphysical processes and small-scale cloud dynamics. This is because clouds introduce strong multiscale inhomogeneities and coupling that span spatial scales from sub-centimeter to hundreds of meters and beyond (Baker, 1997; Shaw, 2003; Wang & Grabowski, 2009). Large-scale computation is becoming a viable research tool for probing multiscale systems such as turbulent clouds and weather and climate systems. In a given computer simulation, however, only a limited range of scales is accessible, with a typical ratio of domain size to grid length at ~ 1000 . This has led to different classes of atmospheric models addressing different ranges of scales, from general climate model (GCM) with horizontal grid length of $10^4 \sim 10^5$ m, to numerical weather prediction (NWP) model addressing mesoscale systems with a grid length of $\sim 10^3$ m, to cloud-resolving large-eddy simulations (LES) focusing on interactions of shallow atmospheric clouds and boundary layer dynamics with a grid length of $10 \sim 100$ m. Increasing

computational resources mean that each model can cover a wider range of scales and that a significant overlap in scales between these models can be achieved. This provides opportunities to improve sub-grid scale parameterizations used in the coarser grained models.

There are many challenging sub-grid scale parameterization issues in cloud-resolving LES when liquid droplets or ice particles form and precipitate. For example, what factors determine the conversion rate of cloud droplets to rain drops? How to treat the entrainment and mixing of dry air with cloudy air at the edges of clouds? How does small-scale air turbulence (part of which is not resolved in LES) affect these microphysical processes? Clearly, LES cannot address these questions. Experimental observations are often inadequate in providing accurate data for physical processes occurring at the droplet ($\sim 10 \mu\text{m}$) to centimeter scales (Shaw, 2003; Devenish *et al.*, 2011). In recent years, direct numerical simulations (DNS) resolving the smallest scales of turbulence have been developed to study these open microphysical questions (Vaillancourt *et al.*, 2002; Andrejczuk *et al.*, 2006; Franklin *et al.*, 2007; Ayala *et al.*, 2008a; Wang *et al.*, 2008; Lanotte *et al.*, 2009). Since the domain size or equivalently the flow Reynolds number in DNS is relatively small, there are two general questions: (1) is DNS a good approach for a specific microphysical problem? (2) if so, how to incorporate or integrate results from DNS into cloud-resolving LES?

In this study, we focus on the effect of air turbulence on collision-coalescence of cloud droplets and its impact on warm rain development. We intend to develop a hybrid and integrated computational approach by combining a hybrid direct numerical simulation (HDNS) of cloud microphysics and a cloud-resolving large-eddy simulation (LES) of cloud dynamics. With potential capabilities offered by PetaScale computers, we hope to close the current scale gap between HDNS domain size [$\mathcal{O}(10 \text{ cm})$] and LES grid length [$\mathcal{O}(10 \text{ m})$].

Towards this goal, we first need to increase the scale ranges covered in HDNS and LES and improve physical representations of the individual processes (e.g., condensation and turbulent collision-coalescence of cloud droplets, turbulent mixing and entrainment). This requires the development of highly scalable implementation of our simulation codes targeted for PetaScale computers of $\mathcal{O}(100,000)$ processors. We also need to resolve various issues related to coupling the HDNS and LES, with and without scale gap between the two. In this paper, we describe our on-going efforts in these directions and some preliminary results.

2. HDNS of turbulent collision-coalescence of cloud droplets

Direct numerical simulation is a bottom-up approach where turbulent air motion at the dissipation-range scales (mm to cm scales) and a limited range of inertial-subrange scales – currently up to $\mathcal{O}(50 \text{ cm})$ – are resolved, but larger-scale motion is represented by a forcing scheme. In recent years, our group has developed a hybrid direct numerical simulation (HDNS) approach (Wang *et al.*, 2005; Ayala *et al.*, 2007) to simulate the collision rate of cloud droplets in a turbulent air where both the inertia and sedimentation of cloud droplets are considered. Another important aspect of HDNS is to incorporate droplet-droplet local aerodynamic interactions by imbedding analytical Stokes disturbance flows due to droplets in a pseudo-spectral simulation of the background air turbulence. The methodology allows us to study both the geometric collision rate and collision efficiency of cloud droplets in a turbulent carrier flow. The flow dissipation rate, droplet inertia and settling velocity are prescribed to mimic the conditions of atmospheric clouds. To render the many-body interaction problem tractable, the disturbance flow due to a given droplet is truncated and it is assumed that such a truncation has little effect on pair collision statistics (Ayala *et al.*, 2007).

Since the grid length dx in HDNS is determined by the Kolmogorov scale η of the air turbulence, the number of grid points (N or grid resolution) used in each spatial direction determines the domain size L_B and simulated flow Taylor microscale Reynolds number R_λ (Wang

et al., 2009)

$$L_B \approx 2N\eta = 2N \left(\frac{\nu^3}{\varepsilon} \right)^{0.25}, \quad R_\lambda \approx 3N^{2/3}, \quad (1)$$

where ε is the viscous energy dissipation per unit mass per unit time, $\nu = 0.17 \text{ cm}^2/\text{s}^3$ is the air kinematic viscosity. For a typical dissipation rate of $\varepsilon = 200 \text{ cm}^2/\text{s}^3$, the above approximate relations yield a domain size of 0.18 m, 0.36 m, 0.72 m, 1.44 m, and 2.88 m; and R_λ of 76, 121, 192, 305, and 484; for $N = 128, 256, 512, 1024, 2048$, respectively.

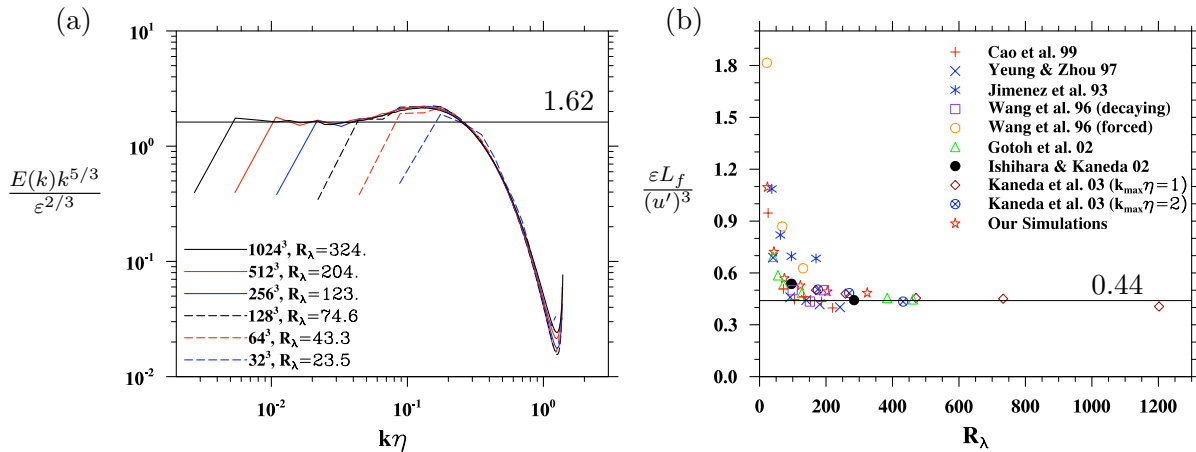


Figure 1. Flow statistics from different grid resolutions: (a) compensated energy spectra; (b) normalized average flow dissipation rate where ε is the average flow dissipation, L_f is the integral length scale associated with longitudinal spatial velocity correlation, and u' is the component r.m.s. fluctuation velocity. The line at 1.62 in (a) indicates the universal scaling for the inertial subrange as observed in experiments (Sreenivasan, 1995; Ishihara *et al.*, 2009). The line at 0.44 in (b) indicates a saturated normalized dissipation rate when the flow Reynolds number is sufficiently high (*i.e.*, $R_\lambda > 200$).

Due to relatively small droplet inertial response time relative to Kolmogorov time and small droplet size relative to Kolmogorov scale, we argue that the dissipation-range flow dynamics makes the primary contribution to the droplet pair statistics and turbulent collision kernel. However, several secondary aspects must be considered in order to completely account for all influence of air turbulence. Firstly, pair statistics of large cloud droplets (30 to 60 μm in radius) can be affected by a finite range of flow dynamic scales including some inertial subrange scales, it is then desirable to increase the range of flow scales in HDNS so all relevant scales affecting the pair statistics are simulated. Secondly, the probability of locally volume-averaged flow dissipation or intermittency is known to depend on flow Reynolds number (Sreenivasan & Antonia, 1997). Although it is impossible to reproduce the cloud turbulence Reynolds number, it is still desirable to increase the flow Reynolds number in HDNS so the sensitivity on flow Reynolds number can be better understood. For these reasons, we are working to perform HDNS at higher grid resolutions, from the current 256^3 to 512^3 grid resolutions to 1024^3 and eventually 2048^3 flow grids, so a larger computational domain (of 1 m size) or a wider range of flow Reynolds numbers could be realized. As shown in Figure 1, the simulated flows at 512^3 and above include all dissipation-range scales and a portion of the inertial sub-range scales, leading to a universal energy spectrum observed in experiments and a converged normalized average flow dissipation rate. This observation indicates that direct effects of small scale flows and intermittency relevant

to cloud droplets could be systematically accounted for in the simulations, and the indirect effects of flow Reynolds number and inertial-subrange intermittency will be modeled separately.

Table 1. The enhancement factor of collision efficiency by turbulence, obtained from HDNS.

a_1 (μm)	a_2 (μm)	a_2/a_1	Grid resolution / R_λ		
			64^3 / 43.0	128^3 / 72.4	256^3 / 120.
30.0	15.0	0.500	1.123 \pm 0.030	1.127 \pm 0.035	1.122 \pm 0.020
	17.5	0.583	1.182 \pm 0.029	1.164 \pm 0.040	1.137 \pm 0.022
	20.0	0.667	1.202 \pm 0.030	1.137 \pm 0.032	1.319 \pm 0.019
	22.5	0.750	1.246 \pm 0.035	1.267 \pm 0.043	1.220 \pm 0.023
	25.0	0.833	1.377 \pm 0.042	1.346 \pm 0.048	1.382 \pm 0.026
50.0	30.0	0.60	1.119 \pm 0.021	1.086 \pm 0.042	1.222 \pm 0.008
	35.0	0.70	1.121 \pm 0.032	1.063 \pm 0.043	1.088 \pm 0.008
	40.0	0.80	1.098 \pm 0.029	1.103 \pm 0.045	1.104 \pm 0.010
	45.0	0.90	1.263 \pm 0.034	1.116 \pm 0.035	1.163 \pm 0.011

A larger computational domain also implies proportionally larger number of droplets in the domain with a prescribed cloud liquid water content. We have scaled our HDNS codes using MPI (Message Passing Interface) to take advantage of scalable, distributed-memory computers in order to enable these higher resolution HDNS. Results in this direction using one-dimensional domain decomposition are reported in a companion paper in this proceeding (Rosa *et al.*, 2011). Table 1 shows the newly obtained enhancement factors on collision efficiency by turbulence at 256^3 grid and compares them with previous results at lower grid resolutions. Here a_1 and a_2 are the radii of colliding droplets. Overall, it appears that the dependence on R_λ is weak, but results at higher resolutions (512^3 and 1024^3) are needed to confirm this observation.

Incorporation of local aerodynamic interactions in HDNS, however, requires intensive computation and a large memory. For example, performing simulations at 1024^3 grid with a liquid content of 1 g/m^3 (or 10 to 100 M droplets) would require 1 to 10 Terascale floating point operations per time step ($\sim 2 \times 10^{17}$ floating point operations per production run) and ~ 1 Tera bytes of CPU memory. PetaScale computing is necessary. In order to prepare our codes for PetaScale computers, we are implementing MPI based on two-dimensional domain decomposition to replace the one-dimensional domain decomposition currently used in our production code, so that $\mathcal{O}(100,000)$ processors can be utilized. Figure 2 shows the scalability of 3D FFT and DNS flow simulation based on 2D domain decomposition, and results are compared to those from 1D domain decomposition. Three grid resolutions and two computers were used to gather the data. While 1D domain decomposition is better for small processor numbers, the 2D domain decomposition performs better for larger numbers of processors with a sustained scalability. Currently, we are implementing 2D domain decomposition for dynamics of droplets in order to generate a highly scalable HDNS code.

HDNS provides kinematic pair statistics of cloud droplets such as radial distribution function and radial relative velocity, which can be used to quantify the turbulent collision rate and collision efficiency (Ayala *et al.*, 2008a; Wang *et al.*, 2008) and to guide the development of an analytical parameterization of the turbulent collision kernel (Ayala *et al.*, 2008b). The higher resolution simulation data will be used to develop an improved parameterizations of turbulent collision kernel. Previously, based on low-resolution HDNS simulations, we have shown that air turbulence can enhance the collision kernel by a factor of two to three on average if the flow dissipation rate is significant (Wang & Grabowski, 2009). Having data at various domain sizes also provides the possibility to explore the effect of flow intermittency.

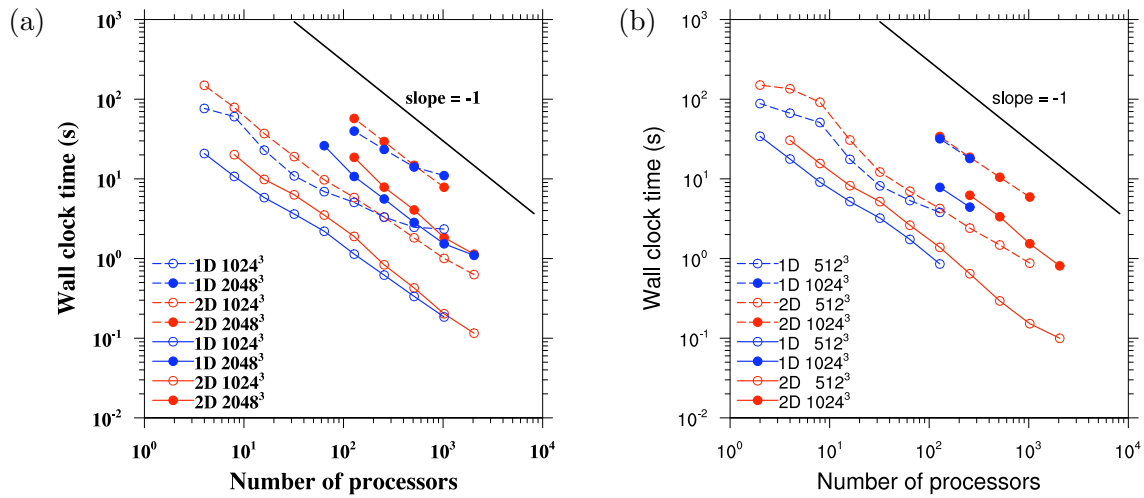


Figure 2. Measured wall clock time as a function of the number of processors used: (a) one forward and one backward 3D FFT, (b) DNS flow simulation per time step. Red symbols show results from 2D domain decomposition, and Blue for 1D domain decomposition. Timing data on Bluefire at NCAR are shown with solid lines, and data on Chimera at UD shown as dash lines. Finally, two grid resolutions are shown with open and filled circles, respectively.

3. Cloud-resolving LES

Atmospheric large-eddy simulation is a top-down approach where the relevant scales of motion are only resolved down to the scale of turbulent energy-containing eddies and the effects of turbulent cascade toward viscous dissipation are parameterized. In atmospheric applications (e.g., atmospheric turbulent boundary layer studies) this typically implies model gridlength in the range of 10 to 100 m. Such an approach is also used in studies concerning shallow atmospheric clouds, such as tropical shallow cumulus or subtropical stratocumulus which are closely tied to the boundary layer dynamics. LES with gridlengths in the range of 10 to 100 m results in realistic simulation of bulk cloud features (such as the cloud field depth, cloud fraction, vertical cloud transports, *etc.* (Siebesma *et al.*, 2003; Stevens *et al.*, 2005). It is unclear, however, if such gridlengths are sufficient to adequately capture the coupling between small-scale cloud dynamics and cloud microphysics. Even smaller gridlengths are needed to resolve the dynamics of the cloud-environment interface instabilities that drive entrainment of environmental cloud-free air and provide kinetic energy input for the cloud turbulence (Grabowski & Clark, 1993). Arguably, gridlengths of the order of 1 m would provide an unprecedented view on the dynamics and turbulence within a small cumulus cloud. As far as cloud microphysics is concerned, one also needs to include a size-resolving representation of cloud droplets and drizzle/rain drops (i.e., the so-called bin microphysics; cf. Grabowski *et al.* (2011) and references therein). Bin microphysics adds a significant computational expense as it requires of the order of a hundred model variables that need to be transported in the physical space.

In this study, a multiscale massively parallel anelastic finite-difference model EULAG (Smolarkiewicz & Margolin, 1997; Grabowski & Smolarkiewicz, 2002; Prusa *et al.*, 2008) is applied as the LES model. Recently bin warm-rain microphysics was implemented in the EULAG code (Wyszogrodzki *et al.*, 2011). As an initial test, we applied the bin EULAG model to the problem of above-the-cloud-base (or in-cloud) activation of cloud droplets (see discussion in Slawinska *et al.* (2011)) in shallow convective clouds observed during the Barbados Oceanographic and Meteorological Experiment (BOMEX; Holland & Rasmusson, 1973) and used in the model in-

tercomparison study described in Siebesma *et al.* (2003). In the BOMEX case, the 1.5-km-deep trade-wind convection layer overlays 0.5-km-deep mixed layer near the ocean surface and is covered by 0.5-km-deep trade-wind inversion. The cloud cover is about 10% and quasi-steady conditions are maintained by the prescribed large-scale subsidence, large-scale moisture advection, surface heat fluxes, and radiative cooling. The model is run for 6 hours as in Siebesma *et al.* (2003) and results from the last hour are used in the analysis. The simulations consider CCN characteristics corresponding to the pristine aerosol from Grabowski *et al.* (2011).




Figure 3. Snapshot showing isosurfaces of cloud water mixing ratio (light shaded) at 0.05 g kg^{-1} and isosurfaces of the activation tendency at 1 (mg s)^{-1} (patchy dark shaded areas).

Figure 3 presents a snapshot example of the 3D distribution of a cloud field at time of 6 hrs. The figure identifies areas where the activation of cloud droplets takes place by rendering cloudy volumes with a condensed water mixing ratio larger than 0.05 g kg^{-1} , and the cloud droplet concentration tendency due to activation larger than 1 (mg s)^{-1} . As anticipated, the maximum activation occurs at the cloud base with the peak at the height of 640 m. Additional activation (the in-cloud activation) takes place in patchy areas that extend across the entire cloud depth. The spatial pattern of these areas changes in space and time as individual clouds evolve.

The in-cloud activation has important effect on the cloud droplets spectra, an aspect important for the warm-rain processes. In particular, bimodal droplet spectra (i.e., spectra characterized by two peaks) can be found in the vicinity of volumes with significant in-cloud activation. This is illustrated in Fig. 4 that shows cloud droplet spectra at selected heights (680, 1040, 1340, and 1540 m) above cloudy points where the maximum activation occurs at each height. Bimodality of the droplet spectra is apparent in each panel, with the secondary peak of the distribution to the left of the primary spectra peak at larger sizes. As expected, the primary peak shifts towards larger sizes as the height increases.

We are also developing an approach to include turbulent collision kernels into LES model prior to closing the gap between HDNS computational domain and LES model gridbox. To account for possible effect of flow intermittency, a PDF-based method is being developed to bridge the scale gap, namely, to address the question of how to incorporate the parameterization derived from HDNS into LES models with bin microphysics. Our basic idea is to employ the Kolmogorov refined similarity theory (Kolmogorov, 1962) to describe the distribution of HDNS-domain-averaged dissipation rates within a LES grid volume, and apply this distribution to calculate the average collision kernel in the LES grid volume. Preliminary results show that

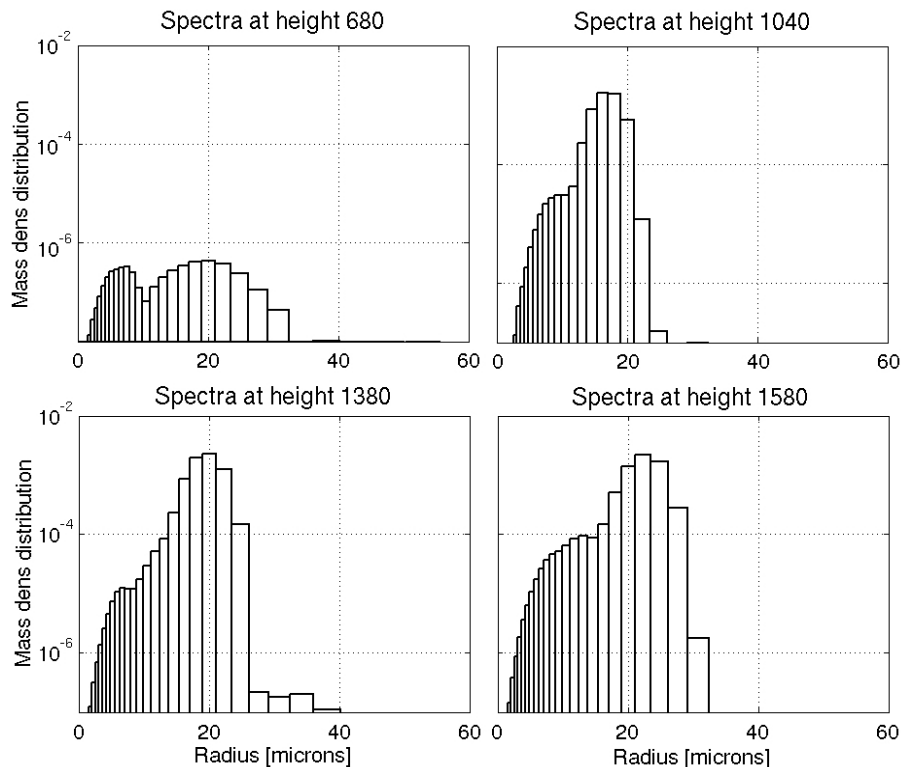


Figure 4. Droplets spectra near regions with elevated activation at four selected heights. Spectra are computed at the final time step, within the cloud above the point with maximum in-cloud activation at each level.

the effect of intermittency on the average kernel in the LES gridbox is negligible, although the local effect within the HDNS volume may be significant. We are currently integrating turbulent collision-coalescence parameterization into the bin EULAG model.

4. PetaScale implementation and data interpretation issues

There are several implementation and data interpretation issues for HDNS and LES. Efforts to address these are described below, which are being coordinated to yield a more scalable and efficient simulation codes, better physical interpretations, and more reliable results.

First, treating the local aerodynamic interaction of cloud droplets in HDNS requires an efficient solver for the large, sparse linear system induced by Stokes flow interactions (Ayala *et al.*, 2007). In our original work, the method of solution was a block Jacobi method. Using this algorithm, the droplet-droplet interaction calculations accounted for approximately 80% of the entire computational cost. Though the system is sparse, solving it is especially challenging because Stokes flow interactions decay slowly, like $1/r$. In fact, the sparsity is imposed by truncation at 50 droplet radii, not by rapid decay. To accelerate the droplet interaction part of HDNS, we have introduced the Generalized Minimal Residual (GMRES) method (Yousef, 2003) as the iterative solver and determined that it is twice as efficient as the block Jacobi method in total CPU time. Our current emphasis is to identify an effective preconditioner to accelerate the GMRES solver. The most direct option is to use an additive restricted Schwarz preconditioner (Cai & Sarkis, 1999). Through experimentation, we find that the Schwarz preconditioner considerably reduces the the number of GMRES iterations

required for the disturbance flows to converge to a given accuracy. Unfortunately, the added overhead associated with applying the preconditioner exceeds the benefit from the reduction in GMRES iterations. Our experiments were performed on a system with a specialized high speed interconnect so that network latency was very low. In systems where latency is more of an issue, the Schwarz preconditioner may still offer distinct advantages. We continue to search for better preconditioning strategies in order to achieve a more efficient GMRES solver with preconditioning.

Second, we are exploring the possibility of tapping the computational resource of graphical processing units (GPU) to perform flow simulation and 3D FFT. The GPU is designed as a highly parallel computing platform for computation-intensive applications such as the HDNS simulation. Our work focuses on optimizing one of the computation kernels of the simulation program, i.e., 3D FFT for GPU. The optimization of 3D FFT for GPU is conducted in two phases. In the first phase, we look at FFT problems whose data can all fit into the relative small GPU memory. Therefore, the CPU-GPU communication is tentatively excluded from the consideration, and the main optimization task is to find the fastest computing strategy. We propose a Cooley-Tukey algorithm based multi-dimensional FFT framework that can represent all reasonable implementations of 2D/3D FFT on GPU and consider GPU architectural features in search for the most efficient implementation. Overall, we achieve a peak of 256 GFlops for 3D FFT on a NVIDIA GTX280 GPU card (Gu *et al.*, 2010). In the second phase, we focus on larger FFT problems that cannot fit into GPU memory, i.e., whose data must be stored in CPU memory, and thereby the communication between CPU and GPU must be taken into consideration. In this study, we propose several FFT decomposition algorithms to increase the data locality and consequently improve the communication efficiency. Furthermore, based on the FFT representation framework, we co-optimize communication and computation to balance their performance in a FFT implementation. Overall, we achieve 15 GFlops for 3D FFT on a single NVIDIA C2070 GPU card (Gu *et al.*, 2011).

In addition, profiling tools are being used to identify bottleneck in our codes and to improve load balancing. Overall, our codes are both computation and communication intensive. Better algorithms and better runtime optimization are needed to achieve a good overall parallel efficiency. We are also developing event-driven visualization tools to better probe physics of turbulence-droplet and droplet-droplet interactions.

5. Summary

In this paper, we outlined an integrated computational approach to address the important multiscale problem in atmospheric sciences: the development of warm rain in shallow convective clouds covering a scale range extending from 1 km (the scale of individual shallow convective clouds) down to $10\ \mu\text{m}$ (the scale of droplet sizes). The approach combines a top-down LES (bin EULAG) of cloud dynamics with a bottom-up hybrid DNS of cloud microphysics. Although we limit here the hybrid DNS to turbulent collision-coalescence due to the important roles of small-scale cloud turbulence, other aspects such as condensational growth and entrainment and mixing could also be included in the HDNS. In order to bridge the scale gap between LES and HDNS, our codes must be made highly scalable to take advantage of imminent PetaScale computational resources. For this, we are implementing and testing multi-dimensional domain decompositions. Preliminary results indicate that our codes can be run on PetaScale computers of $O(100,000)$ cores. Other related efforts to address computation and communication bottlenecks of our codes, such as incorporation of GMRES for solving large linear system and profiling tools, are also being undertaken. The computational approach will provide rigorous tools to address open questions in turbulent clouds such as parameterization of turbulent collision-coalescence, in-cloud activation, and entrainment and mixing on the characteristics of cloud droplets within

turbulent clouds and on precipitation formation.

6. Acknowledgments

This research has been supported by the National Science Foundation through grants OCI-0904534, OCI-0904449, and ATM-0730766 and by the National Center for Atmospheric Research (NCAR). NCAR is sponsored by the National Science Foundation. Computing resources are provided by National Center for Atmospheric Research through CISL-35751010, CISL-35751014 and CISL-35751015.

References

- ANDREJCZUK, M., GRABOWSKI, W. W., MALINOWSKI, S. P. & SMOLARKIEWICZ, P. K. 2006. Numerical simulation of cloud-clear air interfacial mixing: effects on cloud microphysics. *J. Atmos. Sci.* **63**, 3204–3225.
- AYALA, O., GRABOWSKI, W. W. & WANG, L.-P. 2007 A hybrid approach for simulating turbulent collisions of hydrodynamically-interacting particles. *J. Comp. Phys.* **225**, 51–73.
- AYALA O., ROSA B., WANG L.-P. & GRABOWSKI W. W. 2008a Effects of Turbulence on the Geometric Collision Rate of Sedimenting Droplets: Part 1. Results from direct numerical simulation. *New J. Phys.* **10**, 075015.
- AYALA, O., ROSA, B. & WANG, L.-P. 2008b Effects of Turbulence on the Geometric Collision Rate of Sedimenting Droplets: Part 2. Theory and Parameterization. *New J. Phys.* **10**, 075016.
- BAKER, M. B. 1997 Cloud microphysics and climate. *Science* **276**, 1072–1078.
- CAI, X.-C. & SARKIS, M. 1999 A Restricted Additive Schwarz Preconditioner for General Sparse Linear Systems. *SIAM J. on Sci. Computing* **21**, 792–797.
- DEVENISH, B. J., BARTELLO, P., BRENGUIER, J.-L., COLLINS, L. R., GRABOWSKI, W. W., IJZERMANS, R. H. A., MALINOWSKI, S. P., REEKS, M. W., VASSILICOS, J. C., WANG, L. P., WARHAFT, Z. 2011 Droplet growth in warm turbulent clouds. *Q. J. R. Meteorol. Soc.*, in review.
- FRANKLIN, C. N., VAILLANCOURT, P. A. & YAU, M. K. 2007 Statistics and parameterizations of the effect of turbulence on the geometric collision kernel of cloud droplets. *Journal of the Atmospheric Sciences* **64**, 938–954.
- GRABOWSKI, W. W. & CLARK, T. L. 1993 Cloud-environment interface instability, Part II: Extension to three spatial dimensions. *J. Atmos. Sci.* **50**, 555–573.
- GRABOWSKI, W. W. & SMOLARKIEWICZ, P. K. 2002 A multiscale anelastic model for meteorological research. *Mon. Weather Rev.* **130**, 939–956.
- GRABOWSKI, W. W., ANDREJCZUK, M., WANG, L.-P. 2011 Droplet growth in a bin warm-rain scheme with Twomey CCN activation. *Atmos. Res.* **99**, 290–301.
- GU, L., LI, X. & SIEGEL, J. 2010 An empirically tuned 2D and 3D FFT library on CUDA GPU. *Proceedings of the 24th ACM International Conference on Supercomputing ICS '10*, 305–314.
- GU, L., SIEGEL, J. & LI, X. 2011 Using GPUs to compute large out-of-card FFTs. *Proceedings of the international conference on Supercomputing ICS '11*, 255–264.
- HOLLAND, J. Z. & RASMUSSEN, E. M. 1973 Measurements of the atmospheric mass, energy, and momentum budgets over a 500-kilometer square of tropical ocean. *Mon. Wea. Rev.* **101**, 44–55.
- ISHIHARA, T., GOTOH, T. & KANEDA, Y. 2009 Study of high Reynolds number isotropic turbulence by direct numerical simulation. *Annu. Rev. Fluid Mech.* **41**, 165–180.

- KOLMOGOROV, A. N. 1962. A refinement of previous hypotheses concerning the local structure of turbulence in a viscous incompressible fluid at high Reynolds number. *J. Fluid Mech.* **13**, 82–85.
- LANOTTE, A.S., SEMINARA, A., TOSCHI, F. 2009 Cloud Droplet Growth by Condensation in Homogeneous Isotropic Turbulence. *J. Atmos. Sci.* **66**, 1685–1697.
- PRUSA, J. M., SMOLARKIEWICZ, P. K., WYSZOGRODZKI, A. A. 2008 EULAG, a computational model for multiscale flows. *Comp. Fluids* **37**, 1193–1207.
- ROSA, B., PARISHANI, H., AYALA, O., WANG, L.-P. & GRABOWSKI, W.W. 2011 Kinematic and dynamic pair collision statistics of sedimenting inertial particles relevant to warm rain initiation. *The 13th European Turbulence Conference*, Sept. 12–15, 2011, Warsaw, Poland.
- SHAW, R. A. 2003 Particle-turbulence interactions in atmospheric clouds. *Ann. Rev. Fluid Mech.* **35**, 183–227.
- SIEBESMA, A. P. & COAUTHORS 2003 A large eddy simulation intercomparison study of shallow cumulus convection. *J. Atmos. Sci.* **60**, 1201–1219.
- SLAWINSKA, J., GRABOWSKI, W. W., PAWLOWSKA, H. & MORRISON, H. 2011 Droplet activation and mixing in large-eddy simulation of a shallow cumulus field, *J. Atmos. Sci.*, submitted.
- SMOLARKIEWICZ, P. K. & MARGOLIN, L. G. 1997 On forward-in-time differencing for fluids: An Eulerian/semi-Lagrangian nonhydrostatic model for stratified flows. *Atmos.-Ocean Special* **35**, 127–152.
- SREENIVASAN, K. R. 1995 On the universality of the Kolmogorov constant. *Phys. Fluids* **7**, 27782784.
- SREENIVASAN, K. R. & ANTONIA, R.A. 1997 The Phenomenology of small-scale turbulence. *Annu. Rev. Fluid Mech.* **29**, 435472.
- STEVENS, B. & COAUTHORS 2005 Evaluation of large-eddy simulations via observations of nocturnal marine stratocumulus. *Mon. Wea. Rev.* **133**, 1443–1462.
- VAILLANCOURT, P. A., YAU, M. K., BARTELLO, P., GRABOWSKI W. W. 2002 Microscopic approach to cloud droplet growth by condensation. Part II: Turbulence, clustering, and condensational growth. *J. Atmos. Sci.* **59**, 3421–3435.
- WANG, L.-P., AYALA, O., KASPRZAK & S.E., GRABOWSKI, W. W. 2005 Theoretical formulation of collision rate and collision efficiency of hydrodynamically-interacting cloud droplets in turbulent atmosphere. *J. Atmos. Sci.* **62**, 2433–2450.
- WANG, L.-P., AYALA, O., ROSA, B. & GRABOWSKI, W.W. 2008 Turbulent collision efficiency of heavy particle relevant to cloud droplets. *New J. Phys.* **10**, 075013.
- WANG, L.-P., ROSA, B., GAO, H., HE, G. W. & JIN, G.-D. 2009 Turbulent collision of inertial particles: Point-particle based, hybrid simulations and beyond. *Int. J. Multiphase Flow* **35**, 854–867.
- WANG, L.-P. & GRABOWSKI, W.W. 2009 The role of air turbulence in warm rain initiation. *Atmos. Sci. Lett.*, **10**, 1–8.
- WYSZOGRODZKI, A. A., GRABOWSKI, W. W. & WANG, L.-P. 2011 Activation of cloud droplets in bin-microphysics simulation of shallow convection. *Acta Geophysica (EULAG topical issue)*, submitted.
- YOUSEF, S. 2003 *Iterative Methods for Sparse Linear Systems*, second edition. SIAM.