Towards Enabling Digital Twins Capabilities for a Cloud Chamber

JIAQI YANG*†, Emory University, USA
MOHAMMAD ATIF*, Brookhaven National Laboratory, USA
VANESSA LÓPEZ-MARRERO*, Brookhaven National Laboratory and IACS, Stony Brook University, USA
TAO ZHANG, KWANGMIN YU, MEIFENG LIN, LINGDA LI, FAN YANG, and YANGANG LIU, Brookhaven National Laboratory, USA
ABDULLAH AL MUTI SHARFUDDIN and FOLUSO LADEINDE, Stony Brook University, USA

CCS Concepts: • **Computing methodologies**;

Additional Key Words and Phrases: modeling and simulation, PETSc, MPI, CUDA, neural operator learning

ACM Reference Format:
Jiaqi Yang, Mohammad Atif, Vanessa López-Marrero, Tao Zhang, Kwangmin Yu, Meifeng Lin, Lingda Li, Fan Yang, Yangang Liu, Abdullah Al Muti Sharfuddin, and Foluso Ladeinde. 2023. Towards Enabling Digital Twins Capabilities for a Cloud Chamber. In *SC '23: The International Conference for High Performance Computing, Networking, Storage, and Analysis, Nov 12–17, 2023, Denver, CO.* ACM, New York, NY, USA, 2 pages.

INTRODUCTION: Particle-resolved direct numerical simulation (PR-DNS) models, which resolve not only the smallest turbulent eddies but also track the development and motion of individual particles, are arguably an essential tool for exploring aerosol-cloud-turbulence interactions at the fundamental level [1, 2]. For instance, PR-DNS may complement experimental facilities designed to study key physical processes in a controlled environment and therefore serve as digital twins for such cloud chambers [2, 3]. This requires accelerating the PR-DNS solver with tools from traditional high performance computing (HPC) as well as replacing computationally expensive modules with machine learning models. In this poster we present our ongoing work aimed at enabling the use of the PR-DNS model from [1] to serve as a digital twin.

The PR-DNS model is summarized in Figure 1. A full description of the model and numerical method used to solve the model equations can be found in [1]. To describe the numerical solution method in a nutshell, the systems of ordinary differential equations modeling the development and motion of the particles are solved using the implicit (or backward) Euler scheme. A projection method [4] is used to solve the incompressible Navier-Stokes equations and, as a result, an elliptic equation for pressure must be solved at each time step. The equations were discretized using finite differences. Upon discretization, at each time step, a sequence of systems of linear equations must be solved to compute the fluid velocity $\mathbf{u}$, pressure p, temperature T, and water vapor mixing ratio q_v. Krylov subspace methods from the PETSc library [5] are employed to solve each such linear system.

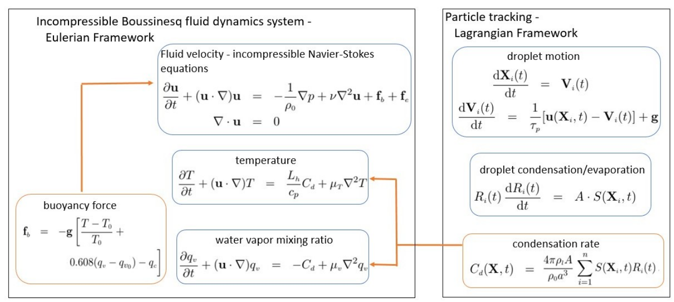

Fig. 1. Particle-Resolved Direct Numerical Simulation (PR-DNS) model. *Left:* Incompressible Boussinesq fluid dynamics system, where $\mathbf{u}$ is the fluid velocity field and p is pressure. The buoyancy force $\mathbf{f}_b$ depends on the temperature T, and vapor mixing ratio q_v. It thus couples the fluid velocity equations with those for T and q_v. *Right:* Equations describing the motion and condensation / evaporation of cloud droplets, where $\mathbf{X}_i(t)$, $\mathbf{V}_i(t)$, and $R_i(t)$ denote, respectively, the position coordinate, velocity, and radius of the i-th droplet. The condensation rate C_d depends on the droplet radii $R_i(t)$ and acts as a source term for the equations for T and q_v in the fluid dynamics model. Hence it couples the fluid dynamics model with the equations for droplet motion and development.

METHOD: We discuss two concurrent efforts – (1) the traditional HPC effort enabling the use of GPUs and (2) employing machine learning to accelerate computational bottlenecks.

HPC Efforts: In this poster, we focus on the performance of linear system solvers when solving for pressure, as an initial code profiling revealed it as the most time consuming among the various linear systems being solved for. In addition, studies to assess the parallel performance and scalability for the various other components of the PR-DNS model are being conducted. CUDA ports are also being developed for particle tracking and thermodynamical field's solvers. A summary of results from our scalability experiments on NERSC Perlmutter appears in Figure 2, whereas Figure 3 compares various preconditioners for solving linear systems, on the BNL Institutional Cluster (BNL-IC).

Machine Learning: The second effort concerns emerging machine learning methods. One goal is to understand the performance of the Fourier Neural Operator (FNO) method [6] to define fast and accurate surrogates for various components of the PR-DNS model.

*Authors contributed equally to this research.
†Work performed while a Guest Research Associate at Brookhaven National Lab, as a participant in the Sustainable Research Pathways (SRP) Program.

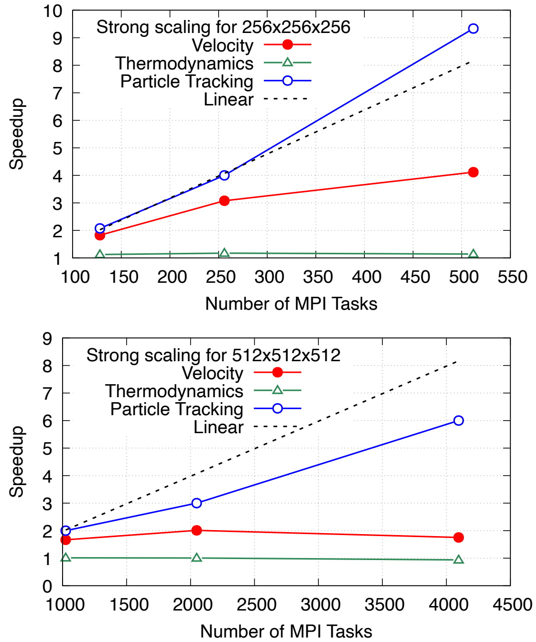

Fig. 2. Strong scaling behavior of PR-DNS model components, on NERSC Perlmutter. Top: 256^3 grid points; Bottom: 512^3 grid points. Particle tracing solver scales well up to 32 nodes (4096 MPI tasks). Thermodynamical field's solver shows poor strong scaling – CUDA ports of compute-intensive kernels under development to address this. Solver for velocity field shows sub-linear scaling – PETSc's CUDA implementations for Krylov subspace solvers can help, as well as CUDA ports of compute-intensive PR-DNS kernels.

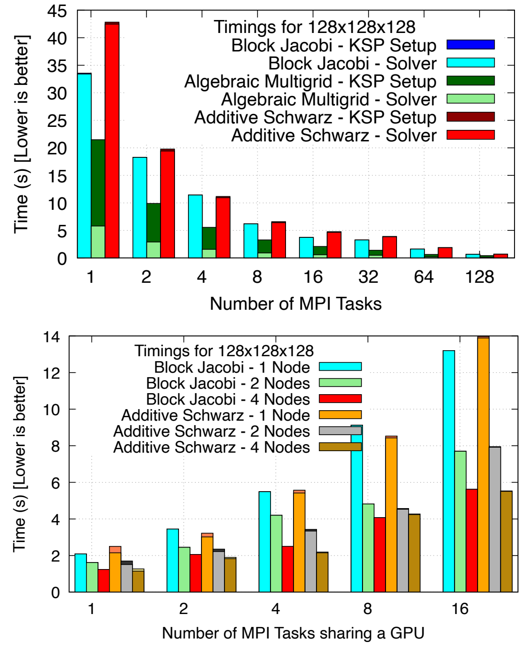

Fig. 3. Impact of preconditioners for GMRES, when solving linear system for pressure (conducted on BNL-IC). Top: CPU only. Algebraic multigrid yields least compute (setup+solver) time, even with higher setup time. Bottom: GPU enabled solution via PETSc, with NVIDIA cuSPARSE library. Best performance when each MPI task has its own GPU.

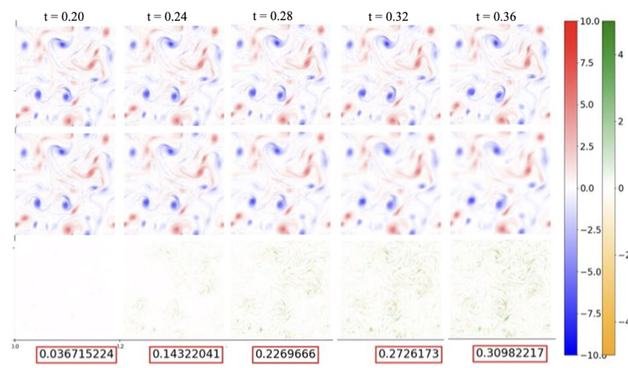

Fig. 4. Data set: 5000 (4000 train, 1000 test) simulations of 2D turbulence with different initial conditions. Image shows vorticity field of one sample: expected (top), FNO's prediction (center), difference (bottom). The numbers in the red box indicate the relative L2 error.

We discuss accumulation of errors with time for vorticity field of 2D turbulence. Refer to Figure 4 for further details.

OUTLOOK: This poster supplements the ongoing efforts to prepare our PR-DNS for utilizing the leadership class exascale supercomputers. Concurrent investigations are also underway with focus on optimizing node-level performance, identifying scalability bottlenecks, and porting compute-intensive kernels to GPUs using CUDA. We also plan to incorporate portable GPU implementations via OpenMP's target offload model, as well as aerosol interactions into the PR-DNS model. ML models for thermodynamical fields will also be investigated in the future.

ACKNOWLEDGMENTS

This work is supported by the LDRD program at Brookhaven National Laboratory, which is sponsored by the US Department of Energy, Office of Science, under Contract Number DE-SC0012704. We also acknowledge administrative support from the Sustainable Research Pathways (SRP) Program, support from a SBU–BNL Seed Grant, and conference travel support from Emory University for JY. This research used resources of the NERSC Center, a User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231 using award NERSC DDR-ERCAP0025875.

REFERENCES

[1] Zheng Gao *et al.* Investigation of turbulent entrainment-mixing processes with a new particle-resolved direct numerical simulation model. *J. Geophys. Res. Atmos.*, 123(4), 2018.

[2] Lulin Xue *et al.* Progress and challenges in modeling dynamics–microphysics interactions: From the Pi chamber to monsoon convection. *Bull Am Meteorol Soc*, 103(5), 2022.

[3] K. Chang *et al.* A laboratory facility to study gas–aerosol–cloud interactions in a turbulent environment: The π chamber. *Bull Am Meteorol Soc*, 97(12), 2016.

[4] David L. Brown *et al.* Accurate projection methods for the incompressible Navier-Stokes equations. *J. Comput. Phys.*, 168(2), 2001.

[5] Satish Balay *et al.* PETSc Web page, 2023. https://petsc.org/.

[6] Z. Li *et al.* Fourier neural operator for parametric partial differential equations, 2021. https://arxiv.org/abs/2010.08895.