A Comprehensive Benchmark Suite for Simulation of Particle Laden Flows Using the Discrete Element Method with Performance Profiles from the Multiphase Flow with Interface eXchanges (MFiX) Code

Peiyuan Liu, Timothy Brown,
William D. Fullmer, Thomas Hauser, and
Christine Hrenya
University of Colorado Boulder

Ray Grout and Hariswaran Sitaraman
National Renewable Energy Laboratory

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Technical Report
NREL/TP-2C00-65637
January 2016

Contract No. DE-AC36-08GO28308

NREL
NATIONAL RENEWABLE ENERGY LABORATORY

A Comprehensive Benchmark Suite for Simulation of Particle Laden Flows Using the Discrete Element Method with Performance Profiles from the Multiphase Flow with Interface eXchanges (MFiX) Code

Peiyuan Liu, Timothy Brown,
William D. Fullmer, Thomas Hauser, and
Christine Hrenya
University of Colorado Boulder

Ray Grout and Hariswaran Sitaraman
National Renewable Energy Laboratory

Prepared under Task No. DOFE.1000

National Renewable Energy Laboratory
15013 Denver West Parkway
Golden, CO 80401
303-275-3000 • www.nrel.gov

Technical Report
NREL/TP-2C00-65637
January 2016

Contract No. DE-AC36-08GO28308

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Available electronically at SciTech Connect http:/www.osti.gov/scitech

Available for a processing fee to U.S. Department of Energy and its contractors, in paper, from:

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
OSTI http://www.osti.gov
Phone: 865.576.8401
Fax: 865.576.5728
Email: reports@osti.gov

Available for sale to the public, in paper, from:

U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Road
Alexandria, VA 22312
NTIS http://www.ntis.gov
Phone: 800.553.6847 or 703.605.6000
Fax: 703.605.6900
Email: orders@ntis.gov

Cover Photos by Dennis Schroeder: (left to right) NREL 26173, NREL 18302, NREL 19758, NREL 29642, NREL 19795.

NREL prints on paper that contains recycled content.

Acknowledgments

This work was supported by the U.S. Department of Energy under Grant No. DE-FE0026298.

This work utilized the JANUS supercomputer, which is supported by the National Science Foundation (award number CNS-0821794) and the University of Colorado Boulder. The Janus supercomputer is a joint effort of the University of Colorado at Boulder, the University of Colorado Denver, and the National Center for Atmospheric Research.

The authors thank Aaron Lattanzi and Dr. Aaron Morris for their insightful discussions about MFiX's heat transfer capabilities.

Nomenclature

Latin

d_p	particle diameter
e	restitution coefficient
f	frequency
$\mathbf{g}$	gravity vector
k_n, k_t	normal and tangential spring constants
m_p	particle mass
N_p	number of particles
np	number of processors
N_x, N_y, N_z	number of CFD cells in the x-, y- and z-directions
T	granular temperature
t	time
V	volume
x, y, z	spatial coordinates

Greek

γ	thermal drag
Δ_x, Δ_y, Δ_z	CFD grid size in x-, y- and z- directions
$\Delta^* = (\Delta_x\Delta_y\Delta_z)^{1/3}/d_p$	non-dim CFD grid size
ε_g	gas (fluid) concentration
ζ_0	zeroth-order collisional cooling rate
η_n, η_t	normal and tangential spring
θ	inclination angle
μ_g	gas viscosity
μ_{pp}	particle-particle coefficient of friction
ρ	density
ϕ	solids concentration
χ	radial distribution function
ω	angular frequency

Subscripts and Superscripts

0	initial
g	gas-phase
s	solids-phase
p	particulate (solids) phase
w	wall

List of Acronyms

BC	boundary condition
CFD	computational fluid dynamics
CPU	central processing unit
DEM	discrete element method
DNS	direct numerical simulation
FB	fluidized bed
HCS	homogeneous cooling system
IC	initial condition
IO	
MFiX	Multiphase Flow with Interface eXchanges
NETL	National Energy Technology Laboratory
NSW	no-slip wall
SIMD	single instruction, multiple data

Executive Summary

This technical report describes activities performed under a project named "MFIX DEM Enhancement for Industry-Relevant Flows" as part of the National Energy Technology Laboratory's Crosscutting Technology Research Program's Transitional Technology Development to Enable Highly Efficient Power Systems with Carbon Management initiative; the report contents also served as a milestone document for Task 2 that project.

Five benchmark problems are developed, discussed, and simulated with the computational fluid dynamics and discrete element method code Multiphase Flow with Interface eXchanges (MFiX). The benchmark problems span dilute and dense regimes and consider statistically homogeneous and inhomogeneous (both clusters and bubbles) particle concentrations and a range of particle and fluid dynamic computational loads. Several variations of the benchmark problems are also discussed to extend the computational phase space to cover granular (particles only), bidisperse, and heat transfer cases. A weak scaling analysis is performed for each benchmark problem, and in most cases the scalability of the code appears reasonable up to ~10^3 cores. Profiling the benchmark problems indicates that the most substantial computational time is being spent on particle-particle force calculations, drag force calculations, and interpolating between discrete particle and continuum fields. Hardware performance analysis was also carried out showing significant Level 2 cache miss ratios and a rather low degree of vectorization. These results provide a baseline to which we can compare future developments of the code as well as a preliminary indicator of where to best focus optimization efforts.

Table of Contents

List of Figures

List of Tables

1 Introduction

This project is funded by the U.S. Department of Energy's National Energy Technology Laboratory's (NETL) Crosscutting Technology Research Program (NETL 2015a), which is intended to bridge the gap between academic research and industrial applications. Keeping with the crosscutting theme, the primary goal of this project seeks to improve the performance of NETL's computational fluid dynamics (CFD) and discrete element method (DEM) code, Multiphase Flow with Interface eXchanges (MFiX) (NETL 2015c)—particularly the DEM aspect of the code.

Short of direct numerical simulation (DNS), which solves for all length scales of particle and fluid motion, CFD-DEM, which employs a larger CFD grid than particle size and uses an interfacial drag closure relation, is the most accurate tool for solving multiphase particulate problems. The accuracy of CFD-DEM is achieved by applying Newton's law and tracking the motion of each individual particle. Obviously, the computational overhead can increase quickly, especially when the particle size is very small or the system is very large, as is the case with most industrial systems. The present capability of CFD-DEM is on the verge of seeing widespread industrial use. Consider the following two cases. One of the largest CFD-DEM studies to date is the recently published work by Capecelatro, Desjardins, and Fox (2015), who studied an unbounded fluidization system of $55 \cdot 10^6$ particles. The unbounded fluidization system mimics the central region of circulating fluidized bed risers. Chew et al. (2012) performed an experimental study on a pilot-scale circulating fluidized bed riser 30 cm in diameter and 18.3 m tall. Two particle sizes were studied: 650 μm and 170 μm. Assuming a solids concentration of 1% in the riser, the pilot-scale system would contain $90 \cdot 10^6$ and $5 \cdot 10^9$ particles for the two diameters. With state-of-the-art academic DEM simulations reaching the level of computational requirements for some pilot-scale systems, it now seems practical to push CFD-DEM capabilities even further so that industrially relevant devices or components may benefit from high-fidelity numerical simulations.

The current project seeks to achieve a speed two times faster than the MFiX 2015-1 CFD-DEM code. To accurately access our progress while the code is being enhanced, the current capabilities of MFiX need to be recorded as a gage, and that is the topic of this document. Five different computational benchmarks, some with a few relatively minor variations, are outlined and discussed in Sections 3–7. The first two benchmark problems—homogeneous cooling system (HCS) and settling—are simple, idealized tests of the code. The latter three benchmark problems—fluidized bed, riser, and square tumbler—are more relevant to industrial problems, though they retain a simple geometry. In this initial phase of the project, it was determined to avoid using geometries requiring cut cells. Additionally, turbulence modeling in the gas phase (e.g., the k-ε model), chemical reactions and cohesion are also excluded from this suite of benchmarks. Polydispersity, granular DEM (no CFD), and heat transfer (in a limited capacity) are all been considered. The five benchmark problems were designed to span a phase space that covers dilute to dense regimes, CFD intensive and DEM intensive, and an order of magnitude in particles per processor.

In addition to recording the simulation time of the benchmark problems, the serial cases are made arbitrarily scalable in size so that a scaling analysis can be performed for each benchmark. In supercomputing, there are two forms of scalability:

- Strong scaling: change in central processing unit (CPU) time with increasing numbers of processors (*np*) for a fixed problem size

- Weak scaling: change in CPU time with increasing *np* for an increasing problem size that scales with *np*.

Ideally we would like to perform a strong scaling analysis because parallelization will minimize the time to obtain the solution for the problems, but this is intractable for the problems at hand. The largest problem sizes should be of the order of 10^6–10^9 particles and may take several hours on several thousand processors. Even assuming the code scales quite poorly, such a simulation may take longer to run in serial than the project duration. Therefore, a weak scaling analysis was performed for each benchmark problem. A supplementary strong scaling analysis was performed on a shortened version of one problem, but this was infeasible for all benchmarks.

2 Preliminaries

2.1 Computational Environment

A version of MFiX was cloned from NETL's repository to a separate repository shortly after the project start date on September 8, 2015. At that time, MFiX was in between the major stable releases of 2015-1 and 2015-2, which were released on April 10, 2015, and October 27, 2015, respectively. Therefore, we simply dub our baseline code *MFiX 2015-1.5*, which is used for all analyses herein.

The simulations were run on the University of Colorado Boulder's JANUS supercomputer. JANUS consists of 1,368 compute nodes of 12 cores (2.8 GHz) per node for a total of 16,426 individual processors (CU-RC 2015). Each core has 2 GB of random access memory. The nodes are connected with a non-blocking QDR InfiniBand high-performance network. JANUS can achieve 184 trillion floating point operations per second (CU-RC 2015).

MFiX is compiled using the Intel Fortran compiler Version 15.0.2. The optimization is set to Level 3. When compiled in parallel, the Intel Message Passing Interface Version 5.0.3 is used. Allinea MAP, TAU, Intel VTune Amplifier, Scalasca, and Perfsuite are used to profile and obtain performance reports on the benchmark cases.

2.2 Baseline Parameters

In devising the benchmark problems, a certain standard was maintained to keep each problem as similar as possible. In all cases, the interstitial fluid is assumed to be air-like, and the particles are assumed to be glass-like. The material properties of the two phases are provided in Table 1. The drag law of Beetstra, van der Hoef, and Kuipers (2007) is used, which is a best fit to the DNS numerical data.

Similarly, many of the numerical settings are similar for all benchmark problems. The recommended linear solver, stabilized biconjugate gradient (BiCGSTAB) method, is used. The maximum number of iterations is set to 500, and the maximum total residual for convergence is set to $1 \cdot 10^{-3}$. The stall detection algorithm is deactivated, which aids convergence during initialization. Finally, the Superbee flux limiter is used for the discretization of all continuum equations. All other parameters and settings not discussed—e.g., under-relaxation, preconditioners—default to the 2015-1 recommended values (NETL 2015b).

Table 1. Baseline Material Properties for Benchmark Problems

Parameter	Value	Units
Gravity magnitude, $\|\mathbf{g}\|$	980.0	cm/s^2
Gas density, ρ_g	$1.0 \cdot 10^{-3}$	g/cm^3
Gas viscosity, μ_g	$2.0 \cdot 10^{-4}$	g/(cm·s)
Solid density, ρ_s	1.0	g/cm^3
Particle diameter, d_p	100.0	μm
Particle-particle restitution coefficient, e	0.8	-
Particle-particle friction coefficient, μ_{pp}	0.0	-
Particle-wall restitution coefficient, e_w	1.0	-
Particle-wall friction coefficient, μ_{pp}	0.0	-

The DEM solver settings are also standardized. The linear spring dashpot is used to model soft-sphere contacts. The particle-particle and particle-wall normal spring constants, k_n and $k_{n,w}$, respectively, are set to $1 \cdot 10^4$ (dyn/cm). The ratio of tangential to normal spring constants is 2/7 (for both particle-particle and particle-wall), and the ratio of tangential to normal damping coefficients is 1/2 (for both particle-particle and particle-wall). The grid-based neighbor search algorithm is used, allowing the code to internally set the DEM neighbor grid. The maximum number of DEM steps without a neighbor search is set to 25. The DEM is always coupled to the continuum solver except in granular (pure-DEM problems), and GARG_2012 interpolation is used to map between the discrete and continuum variables.

3 Benchmark: Homogeneous Cooling System

3.1 System Description

The HCS was first devised by Haff (1983), the namesake behind Haff's Law, the characteristic decay rate of a granular system. The HCS is the simplest nontrivial particulate system, and it has been widely studied in the physics community because of its simplicity. The HCS considers a uniform distribution of particles in an infinite expanse. There is no net flow in any direction—i.e., $\overline{u}_p = \overline{v}_p = \overline{w}_p = 0$, where $\overline{u}_p$, $\overline{v}_p$ and $\overline{w}_p$ are particle averaged velocities; e.g., $\overline{u}_p = \frac{1}{N_p}\sum_{i=1}^{N_p} u_{p,i}$, in the *x*-, *y*-, and *z*-directions, respectively. However, the particles are uniformly excited with a random velocity such that the state is characterized by a uniform, initial kinetic energy: $KE_0 = \frac{m_p}{2N_p}\sum_{i=1}^{N_p}\left(u_{p,i}^2 + v_{p,i}^2 + w_{p,i}^2\right)$. As long as no internal currents are generated in the HCS, the kinetic energy can be related to the granular temperature, T, a common measure of the fluctuating kinetic energy, via $T = 2KE/(3m_p)$.

Because of the HCS's (statistical) uniformity and the absence of mean flow and boundaries, all spatial gradients vanish and the system simply "cools" in time due to dissipative collisions. Originally a granular problem, here we consider a gas-solid system so that two sources of dissipation arise: collisional and viscous. Under the limiting assumptions of the HCS, the multiphase kinetic-theory-based continuum model of Garzó et al. (2012) reduces to:

$$\frac{dT}{dt} = -\frac{2T}{m_p}\gamma - \zeta_0 T\,, \tag{1}$$

where

$$\begin{aligned}\gamma = {} & 3\pi\mu_g d_p\left(1 + 3\sqrt{\frac{\phi}{2}} + \frac{135}{64}\phi\ln\phi + 11.26\phi(1 - 5.1\phi + 16.57\phi^2 - 21.77\phi^3) + 4.6\phi\chi\right) \\ & + 3\pi\rho_g d_p^2\sqrt{T}\,\frac{0.096 + 0.142\phi^{0.212}}{(1-\phi)^{4.454}}\end{aligned} \tag{2}$$

is the thermal drag model, which was constituted through DNS (Koch and Sangani 1999; Wylie, Koch, and Ladd 2003).

$$\zeta_0 = \frac{8\phi}{\sqrt{\pi}d_p}\left(1 - e^2\right)\chi\left(1 + \frac{3}{16}c\right)\sqrt{T} \tag{3}$$

is the zeroth-order collisional cooling rate, and $\phi = N_p(\pi d_p^3/6)/V_{tot}$ is the mean solids concentration.

$$\chi = \frac{1 - \phi/2}{(1-\phi)^3} \tag{4}$$

Is the radial distribution function—given here by the Carnahan-Starling (1969) approximation, and

$$c = \frac{16(1-e)(1-2e^2)}{81-17e+30e^2(1-e)} \tag{5}$$

is the kurtosis of the particle velocity distribution function. By applying Eqs. (2)–(5), Eq. (1) can be cast in the form:

$$\frac{dT}{dt} = -AT^{3/2} - BT\,, \tag{6}$$

which has the analytical solution (Yin *et al.* 2013):

$$\frac{T}{T_0} = \left[e^{Bt/2} + \frac{A}{B}\sqrt{T_0}\left(e^{Bt/2}-1\right)\right]^{-2} \tag{7}$$

The analytical solution of Eq. (7) is shown in Figure 1 using the condition of this test problem (outline below).

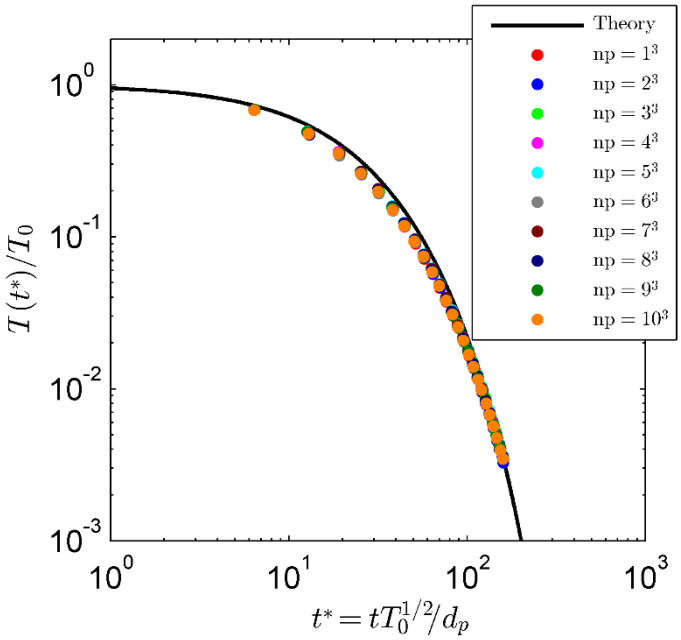

Figure 1. Decay of the initial granular temperature of the HCS test case compared to the theoretical solution of Eq. (7)

Note that the granular temperature is expressed in terms of the difference between the individual particle velocities and *some* local, mean velocity. Defining such a local, mean velocity is difficult; therefore, rather than attempting to do so here, we only expect Eq. (7) to hold in the homogeneous region before mean flows or concentration inhomogeneities develop due to velocity vortex or clustering instabilities, respectively. DNS simulations of a similar two-phase

HCS have verified the accuracy of Eq. (7) when spatial homogeneity is maintained (Yin et al. 2013).

Linear stability analyses of the continuum equations under the HCS conditions show that the critical length scale for instabilities to set increases with decreasing solids concentration. A very dilute system is specified for this case, with $\phi = 0.01$. At this concentration, the granular linear stability analyses predicts that a non-dimensional system size of $L^* = L/d_p = 220$ is required for the onset of the velocity vortex instability (Garzo 2005). Although this transition will occur at a smaller system size for a gas-solid system, its exact value is currently unknown. Therefore, $L^* = 40$ is chosen for the serial case, which should be sufficiently small to avoid instabilities. The initial granular temperature is set to 1,000 cm^2/s^2, which, with the other baseline parameters shown in Table 1, give the other three non-dimensional variables that characterize the system: $Re_{T0} = \rho_s d_p \sqrt{T}/\mu_g = \sqrt{2.5}$, $e = 0.8$, and $\rho_s/\rho_g = 1{,}000$.

The initial particle positions are generated from a random, uniform distribution. Particle velocities are then set from a Gaussian random distribution, scaled to a specified granular temperature, and then normalized to give zero mean flow in each direction. The initial condition of the gas is zero velocity everywhere, and the concentration field is simply interpolated from the DEM data. The system is cubic with periodic boundary conditions (BCs) in each direction. At the specified concentration system size, $N_p = 1{,}222$. The CFD grid is cubic with $\Delta^* = 2$ ($N_x = N_y = 4$, $N_z = 20$).

3.2 Variants

Currently there are no variants of this problem.

3.3 Weak Scaling Results

In the parallel case, the basic HCS system is scaled in three-dimensional form to maintain the cubic domain. The particle initial conditions (ICs) of the serial case are simply mirrored into new domains—i.e., random positions and velocities are only generated for every 1,222 particles, which are simply copied a specified number of times. The spatial coordinates and CFD grid are also multiplied accordingly. As a result, this case has a rather odd scaling, resulting in an *np* of 1^3 (= serial), 2^3 (= 8), 3^3 (= 27), … 10^3 (= 1,000).

It was expected that instabilities would occur at a certain scaling of the serial problem, causing the granular temperature decay to deviate from the analytical solution of Eq. (7). However, all cases decayed quite similarly, as shown in Figure 1, even at *np* = 1,000 which is nearly twice the (granular) estimated critical system size. It is believed that this enhanced stability is a result of the multilevel periodicity of this particular test problem; because of the system periodicity and the way in which the serial system was copied to achieve scaling, each case is essentially the same.

In the serial simulation of this test case, the CPU time was approximately 810 s on JANUS. The scaling of the CPU times as the system increased is displayed in Figure 2. The reported times are CPU times listed at the end of the LOG file, which do not include input/output (I/O) times. However, I/O time comprised a relatively minor portion of the total wall time—e.g., for *np* = 1,000, I/O time/wall time ≈ 8%. The CPU times initially increase linearly, level off, and then

begin to increase rapidly as *np* approaches 1,000. It would be interesting to see how the problem scales out to at least $np = 16^3 = 4{,}096$ in the future if computational resources allow.

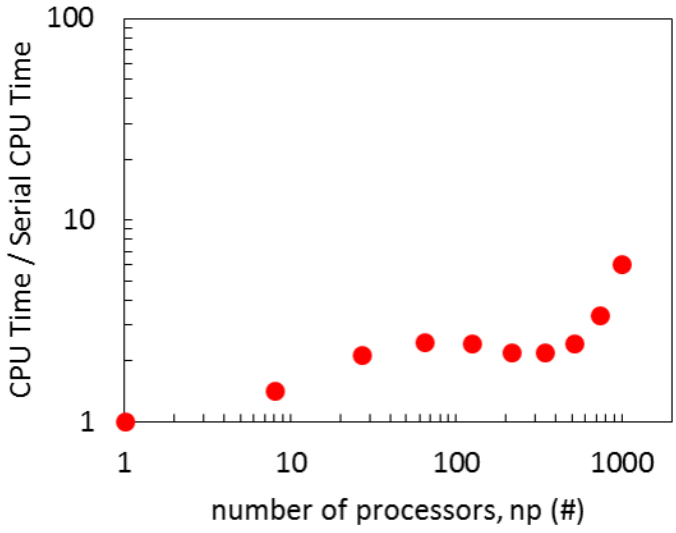

Figure 2. Weak scaling of the HCS test case

3.4 Profiling Results

The profiling results showing the top five most computationally expensive subroutines for the HCS case are provided in Table 2. As the most dilute of the benchmark problems, the interpolation of the mean field variables (from the discrete particle data) takes a significant amount of time. The linear solution of the CFD variables appears as #5 on the list (DGTSV is the tri-diagonal Gaussian solver, part of the BiCGSTAB scheme). It is somewhat concerning that it takes seven times longer to pass discrete/continuous information back and forth than it does for the actual CFD or DEM solutions.

Table 2. Profiling Summary of the Top Five Functions for the HCS Benchmark

Function	% of time
COMP_MEAN_FIELDS0	35.88
FUNTIONS	18.31
DRAG_GS_DES0	6.41
DESGRID_MOD	5.19
DGTSV	4.81

4 Benchmark: Settling

4.1 System Description

The settling case is another simple test case that is designed to run rather quickly in which particles fall to rest in a small box. Initially particles are uniformly, randomly distributed throughout the domain with near-zero initial velocity (T_0 = 0.01 m^2/s^2). The gas velocity (if present) is zero everywhere. Gravity acts in the negative y-direction. In the serial case, the domain is a cubic box with edges of length $L_x = L_y = L_z = 0.15$ cm ($L^* = 15$). A cubic CFD grid of $\Delta^* = 1.5$ ($N_x = N_y = N_z = 10$). No-slip wall's (NSWs) are specified at every BC. The solids concentration is set to $\phi = 0.25$ for N_p = 1,611 uniformly and randomly distributed particles. The simulation runs for 50 ms, long enough for most particles to come to rest. Four snapshots, from the beginning to the end of the serial simulation, are displayed in Figure 3.

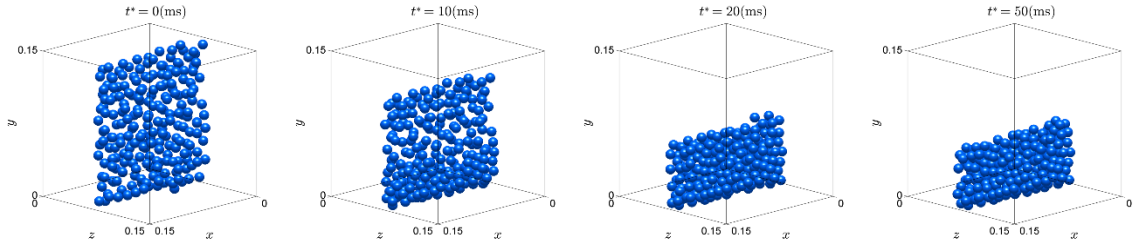

Figure 3. Visualization of the settling case (serial scaling) from the initial state (left) to the final state (right)

4.2 Variants

In addition to the two-fluid case (i.e., coupled CFD-DEM), a granular system is considered (i.e., pure-DEM). The input of the granular variant is nearly identical to the CFD-DEM case except that the CFD and DEM solvers are decoupled, no concentration interpolation is specified, and the gas-phase momentum equations are not solved.

4.3 Weak Scaling Results

To avoid load-balancing issues, the *y*-dimension of the domain remains constant, and it is not decomposed into sub-domains when the settling case is scaled for parallel simulations. (Note that the final condition has all particles in the bottom half and no particles in the top half.) The system is therefore stretched two-dimensionally by successively doubling in the *x*- and *z*-directions. The particle IC's of the serial case are simply mirrored into new sub-domains so that each processor solves for *essentially* the same problem. (Sub-domains near the outer wall are the exception.) The spatial coordinates and CFD grid are also multiplied accordingly.

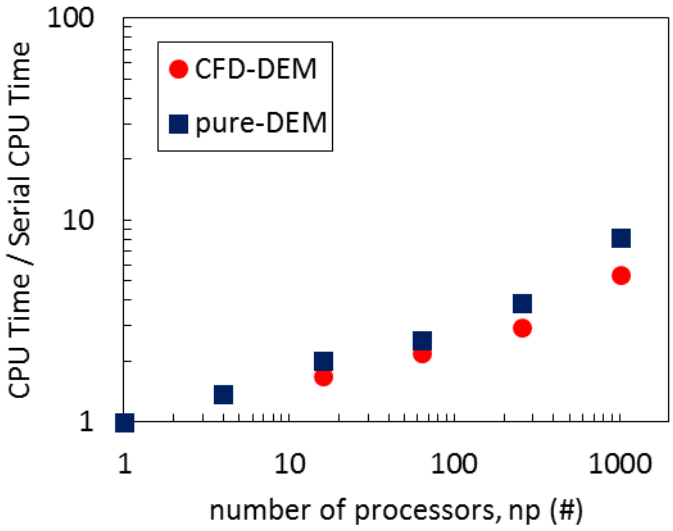

Figure 4. Weak scaling for the granular and CFD-DEM settling problem

The reported simulation times are an average of three runs on JANUS compute nodes. In serial, the CPU times for the CFD-DEM and pure-DEM simulations were 348.5 s and 160.1 s, respectively. The serial times are used to normalize the scaled parallel simulations, as shown in Figure 4. In the DEM simulation, the IO and CPU times were not listed, and therefore the wall time from the screen output was used. To keep an apples-to-apples comparison to the CFD-DEM scaling, the wall time was used here also. The substitution of the wall clock time for the CPU produced a noticeable difference for the settling case because IO time is more substantial for this short simulation—e.g., for np = 1,024, IO time/wall time ≈ 40%.

4.4 Strong Scaling Results

A strong scaling analysis was also performed for the CFD-DEM settling problem. The geometry for strong scaling corresponds to the largest system size of the weak scaling analysis—i.e., the settling problem was scaled to np = 1,024, and the simulation time was reduced from 50 ms to 20 ms. This shorter time allows the particles to settle, but it ignores the secondary transient wherein the particles pack into a static bed. The results of the strong scaling analysis are presented in Figure 5, which shows good scalability up to approximately 1,000 cores. Note that in this strong scaling analysis, all cases except the largest case (np = 1,024) have domain decompositions with $N_{x,y,z} > 15$. It is known from previous experience that MFiX scalability begins to deteriorate with spatial decompositions of $N_{x,y,z} \approx 15$; therefore, it is indeterminate whether the departure of perfect scaling shown in Figure 5 is due to the increasing np or the decreasing $N_{x,y,z}/np$.

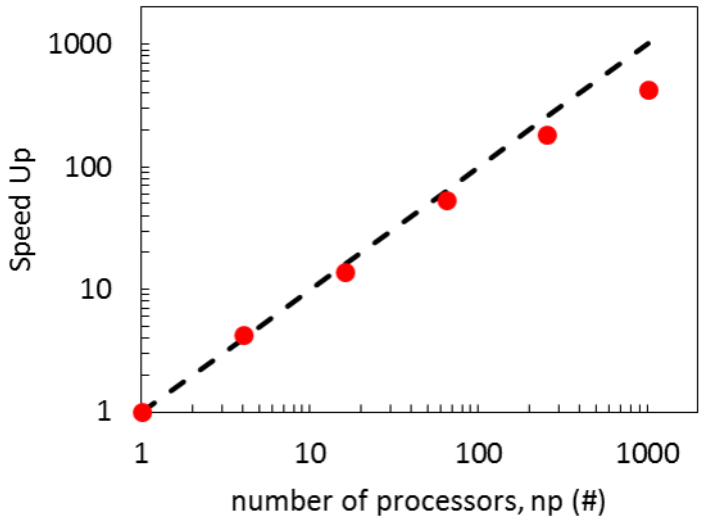

Figure 5. Strong scaling of the case where *np* = 1,024 from the weak scaling analysis of the CFD-DEM settling benchmark

4.5 Profiling Results

The profiling results showing the top five most computationally expensive subroutines for the settling case are provided in Table 3. Not surprisingly, the contact force calculation takes up nearly 50% of the overhead. Although the problem begins as relative sparse with no collisions, the particles are packed with multiple and enduring contacts during the second half of the simulation, as illustrated in Figure 3.

Table 3. Profiling Summary of the Top Five Functions for the CFD-DEM Settling Benchmark

Function	% of time
CALC_FORCE_DEM	49.26
COMP_MEAN_FIELDS0	21.96
DRAG_GS_DES0	10.01
DESGRID_NEIH_BUILD	8.65
CFNEWVALUES	8.39

5 Benchmark: Fluidized Bed

5.1 System Description

The fluidization is a process in which many solid particles confined in a pipe or channel are converted from a static, solid-like state (packed bed) to a dynamic, fluid-like state (fluidized bed) under the drag force introduced by a flow of fluid, frequently a gas. The fluidized bed (FB) is one of the most widely used devices for fluidizing solid materials in chemical, biological, and pharmaceutical industries because of its superb heat and mass transfer capabilities (Rhodes 2008). Multiphase flow research into FBs has become a very active field in recent decades because the understanding of solid behaviors in FBs is crucial for essential engineering processes such as drying, mixing, granulation, coating, and heat transfer (Fan and Zhu 1998).

Four regimes of gas-solid fluidization were established empirically by Geldart (1973) by summarizing an extensive database of FB experiments using air at ambient pressure and temperature. At a given density difference, $\Delta\rho = \rho_s - \rho_g$, and increasing particle diameter, d_p, the following four regimes are encountered:

- Group C: fine, cohesive powders that are difficult to fluidize due to stronger attractive inter-particle forces (i.e., van der Waals force) relative to particle weight
- Group A: smaller particles that will fluidize but experience a noticeable regime of homogeneous bed expansion (non-bubbling fluidization) beyond minimum fluidization velocity, U_{mf}, before bubbling commences. Inter-particle forces may not yet be negligibly small, but they are of relatively minor importance compared to Group C particles.
- Group B: larger, sand-like particles that begin bubbling immediately after U_{mf}
- Group D: very large or dense particles that exhibit relatively poor mixing behavior upon fluidization but readily produce deep spouted beds.

The boundaries for the four regimes determined by Geldart (1973) are provided in Figure 6.

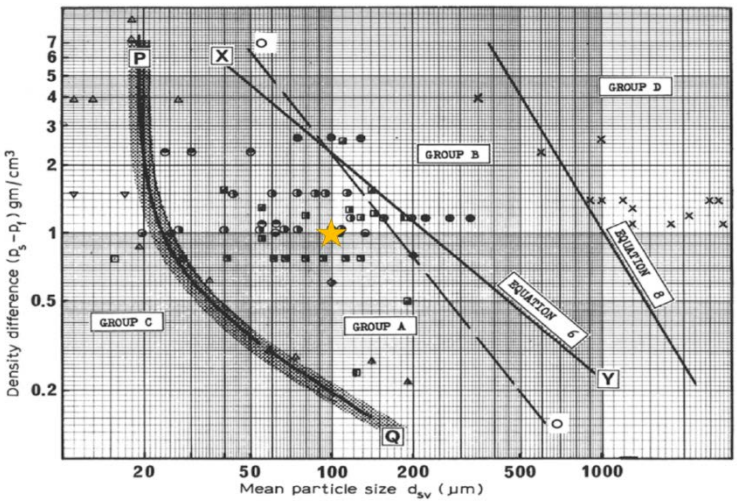

Figure 6. Geldart fluidization diagram showing the approximate location of the benchmark case, as indicated by the golden star. *Image modified from Geldart (1973*)

As one of the benchmark cases in this project, we simulate a rectangular fluidized bed with a square distributor plate for uniform fluidization. The baseline material properties are used (see Table 1), yielding Group A behavior, as indicated by the golden star shown in Figure 6. Abrahamsen and Geldart (1980) give a *very* empirical correlation for the minimum bubbling velocity of Group A particles:

$$U_{mb} = 2.07 \frac{d_p \rho_g^{0.06}}{\mu_g^{0.347}} e^{0.716 F_{fines}}, \tag{8}$$

where F_{fines} is the fraction of particles with $d_p < 45$ μm ($F_{fines} = 0$ in this idealized case). (Note that Eq. (8) is dimensionally inconsistent, and material properties should be specified in International System of Units for U_{mb} in m/s.) According to Eq. (8), homogeneous fluidization should give way to bubbling at approximately $U_{mf} = 0.88$ cm/s. The inlet velocity of this test problem is set at 1.5 cm/s, well above the estimated value, to ensure that the bed is operating in the bubbling regime where gas-solid and solid-solid interactions are equally important.

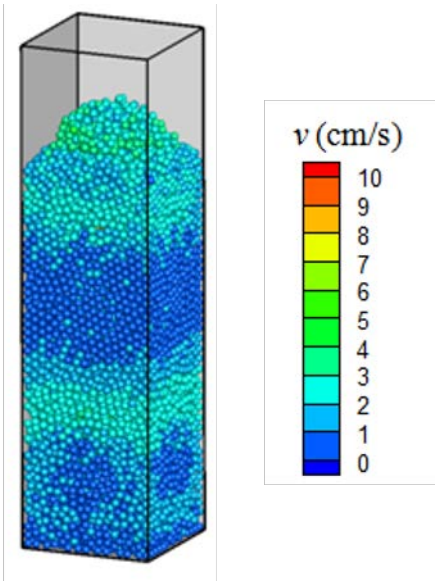

Figure 7. A snapshot of a fluidized bed simulation of 10,000 particles

The square bed is 0.08 cm wide, 0.08 cm deep, and 1.0 cm tall. The distributor plate is modeled with a uniform gas inflow BC at the bottom of the bed. A pressure outflow BC is set at the top, and all four side walls are specified as NSW BCs. The CFD grid is cubic with $\Delta^* = 2$ ($N_x = N_z = 4$, $N_y = 50$). A snapshot of a parallelized simulation containing 10,000 particles is shown in Figure 7.

The original IC consists of 2,500 particles at rest on the bottom of the bed in an ordered packing. The gas flow is also zero throughout the domain. An initial transient period (the time it takes to reach a statistically steady state from the original IC) of 1,000 ms is simulated off-line, i.e., not part of the benchmark test. The temporal evolution of the system is shown in Figure 8. The state of the system at t_0 = 1,000 ms is used as the restart IC for subsequent benchmark tests. The ensuing simulations are run for either 50 ms or 200 ms for the short and long benchmark tests, respectively. As shown in Figure 8, the characteristic bubbling frequency is close to the shorter simulation time of 50 ms. Therefore, the longer runs will provide better statistics for weak scaling analyses but at the cost of increased CPU time.

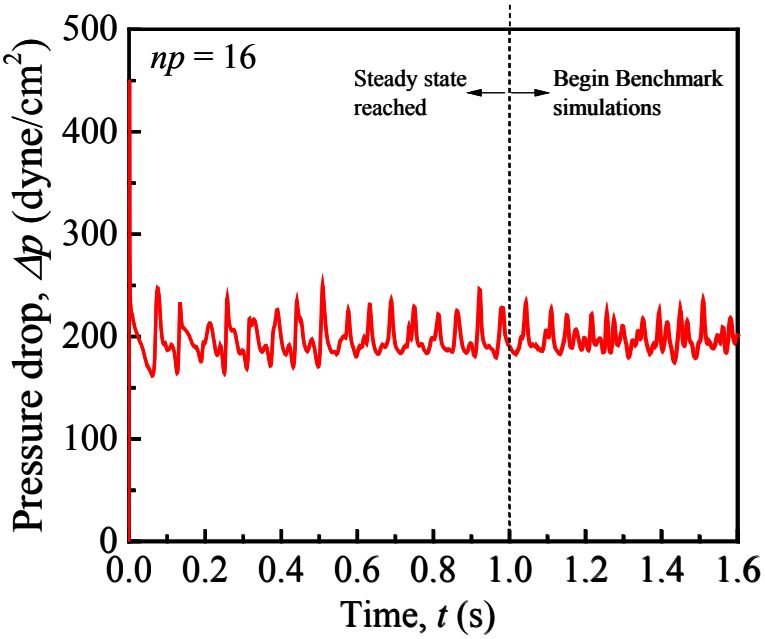

Figure 8. Temporal evolution of the fluidized bed benchmark problem

5.2 Variants

Currently there are no variants of this problem. However, heat transfer is slated to be added to the fluidized bed problem in the near future. The heat transfer variant considered specified a discontinuous IC in which hot particles (T_l = 300 K) were superimposed over cold particles (T_l = 300 K). The gas IC and inlet temperature were set to T_l. However, we were unable to achieve a converged solution using the 2015-1.5 version. The initial non-convergence was determined to be due to the linear solution of the gas-phase temperature equation, which was intractably stiff due to the mismatch between the CFD and DEM time steps. Essentially, as the DES time steps progress without CFD iteration, the particles are allowed to convect too much energy into the gas phase due to its static temperature.[1] In an attempt to circumvent this deficiency, a particle phase energy source is used to heat the solids temperature during the DES time steps. The changes that must be reflected in the 2015-1 subroutine des_thermo_conv.f are:

- After line 27 insert: use geometry.
- Replace or modify line 58 to: Tg = Tg + DES_ENERGY_SOURCE(IJK) / (ROP_G(IJK) * C_PG(IJK) * VOL(IJK)).

After making the above changes to the source code, a *np* = 4 scaled simulation of the fluidized bed including heat transfer was successful. However, scaling this variant has not yet been performed. The corrections listed above should be reviewed by the NETL staff and patched into future MFiX releases, or a more thorough solution to this issue should be implemented.

[1] Personal communications with authors and A. Lattanzi, W.D. Fullmer, and P. Liu, from the University of Colorado at Boulder, on December 9, 2015.

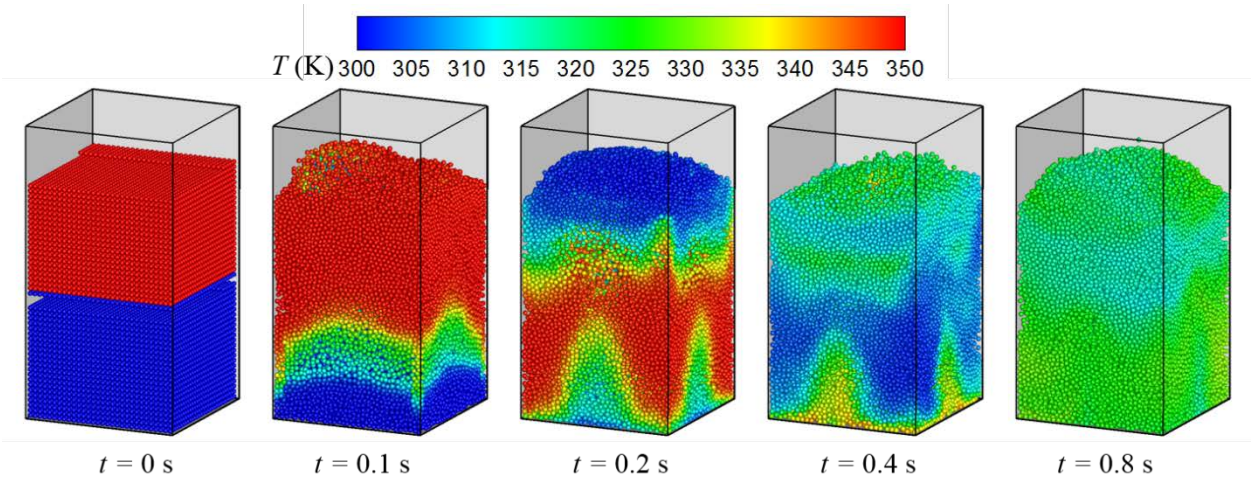

Figure 9. Progression of the *np* = 4 bubbling bed with heat transfer

5.3 Weak Scaling Results

The FB benchmark case is scaled in the *x*- and *z*- dimensions while keeping the serial conditions in the *y*-dimension. The serial CPU time was 335.5 s for the benchmark test with a simulated time of 50 ms. The serial times are used to normalize the scaled parallel simulations, as shown in Figure 10.

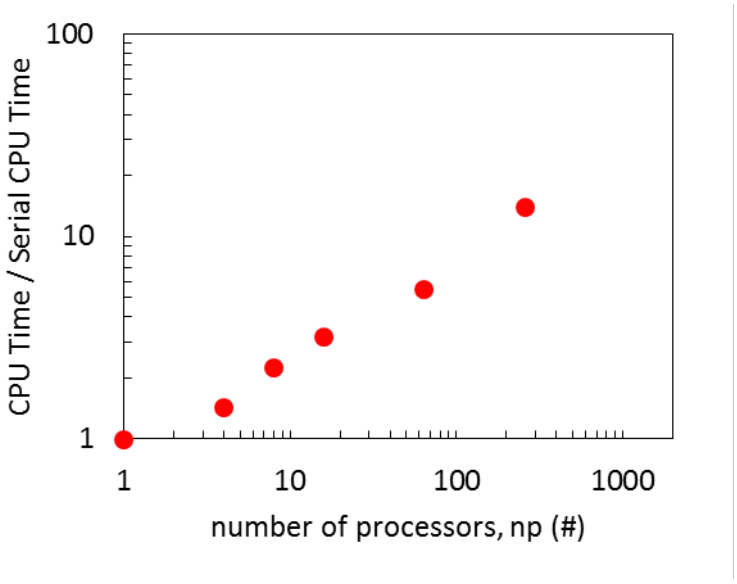

Figure 10. Weak scaling results of the FB benchmark problem

5.4 Profiling Results

A summary of the profiling results is given in Table 4. Somewhat surprisingly for this case, which has many and enduring contacts, the contact force calculation does not take the most time. Calculating the drag force takes a surprisingly large 7.2% of the computation time. This points to an area of potential improvement. Even with the most recent DNS-based drag law, discrepancies

between the correlation and the data can be as large as 4% relative error (Tang et al. 2015). This begs the question: do we really need double-precision calculation of a correlation that implicitly contains 4% relative error?

Table 4. Profiling Summary of the Top Five Functions for the FB Benchmark

Function	% of time
COMP_MEAN_FIELDS0	23.13
DRAG_GS_DES0	12.17
CALC_FORCE_DEM	10.31
DRAG_GS	7.20
CFNEWVALUES	6.20

6 Benchmark: Riser

6.1 System Description

Riser flows occur when a fluidized bed runs at higher superficial gas velocities than the bubbling bed. A bed of particles are moved in bulk by the gas flow in the riser. Due to the increased gas-solid drag, riser flows are highly unsteady, with large fluctuations of particle concentrations and formations of particle clusters that may settle under gravity (Agrawal et al. 2001). Understanding riser flows is crucial to designing a circulating fluidized bed for heat transfer or reacting flows (Capecelatro, Pepiot, and Desjardins 2014; Capecelatro, Pepiot, and Desjardins 2015). Similar to the bubbling fluidized bed, gas-solid and solid-solid interactions are both important in riser flows.

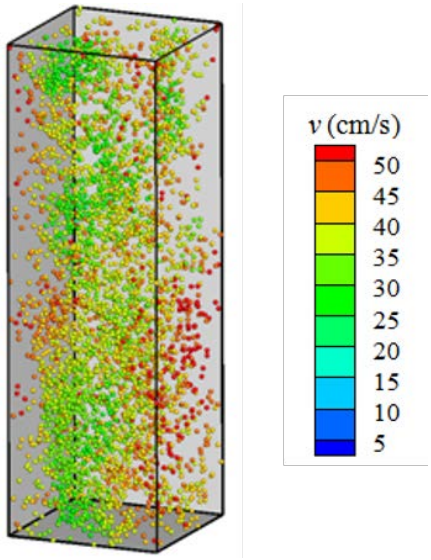

Figure 11. A snapshot of a rise flow simulation of 4,000 particles

In the benchmark case for riser flows, we simulate a rectangular bed periodic for solid phase in all three dimensions. For the gas phase, cyclic BC with zero pressure drop was applied in the two lateral dimensions. To drive the gas flow, a cyclic BC with specified pressure drop is set in the vertical direction. The benchmark simulations start when the gas flow reaches the fully-developed regime and the gas velocity gradient vanishes along the vertical direction. The solid volume fraction and bed dimensions can be adjusted to control the system scale for weak/strong analysis. A snapshot of a parallelized simulation containing 4,000 particles is shown in Figure 11.

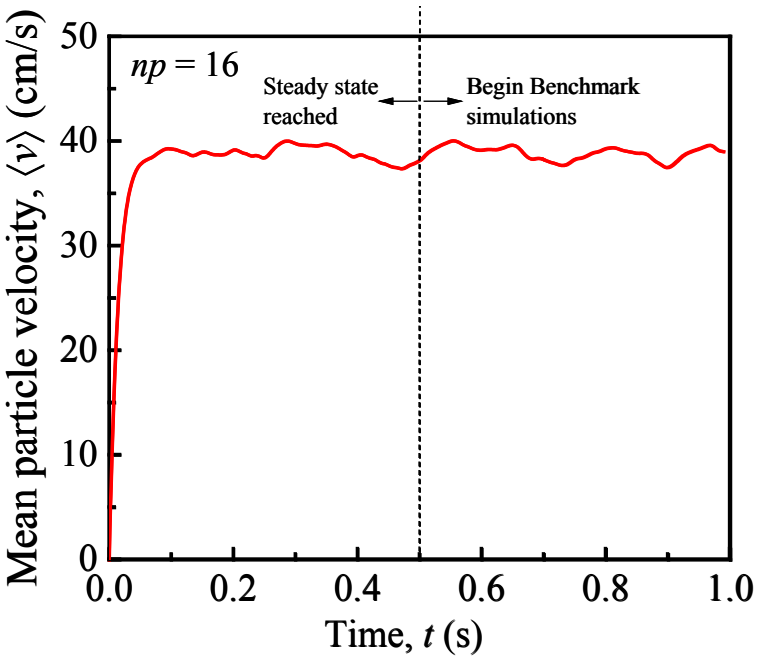

Figure 12. Temporal evolution of the riser benchmark

The original IC consists of 250 particles randomly distributed throughout the entire domain with a solid volume fraction of 2%. The CFD grid is cubic with $\Delta^* = 2$ ($N_x = N_z = 4$, $N_y = 50$). An initial transient period of 500 ms was simulated off-line to reach a statistically steady state from the original IC. The transient from the true initial state to a quasi-steady state is displayed in Figure 12. The state of the system at $t_0 = 500$ ms is used as the restart IC for subsequent benchmark tests. The ensuing simulations are run for either 50 ms or 200 ms for the short and long benchmark tests, respectively.

6.2 Variants

Currently there are no variants of this problem.

6.3 Weak Scaling Results

The riser flow benchmark case is scaled in the *x*- and *z*- dimensions while keeping the serial conditions in the *y*-dimension. The serial CPU time was 176.2 s for the benchmark test with a simulated time of 50 ms. The serial times are used to normalize the scaled parallel simulations, as shown in Figure 13.

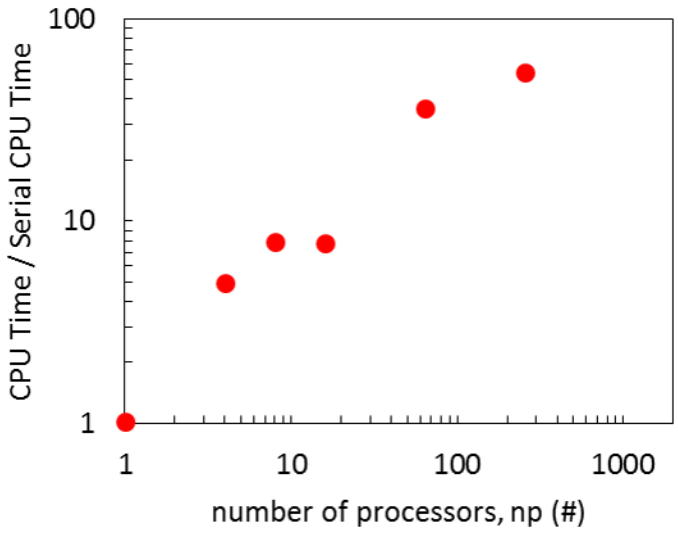

Figure 13. Weak scaling results for the riser benchmark problem

6.4 Profiling Results

A summary of the profiling results for the riser benchmark is provided in Table 5. Even though particles are clustering, as evidenced in Figure 11, the computational times remain CFD-dominated. Subroutines related to transferring discrete and continuum data control the simulation time. The linear solver subroutine has pushed its way into the top five, whereas the DEM force calculation subroutine has dropped out.

Table 5. Profiling Summary of the Top Five Functions for the Riser Benchmark

Function	% of time
FUNCTIONS	21.9
COMP_MEAN_FIELDS0	15.0
LEQSOL	12.76
DGTSV	12.25
DRAG_GS_DES0	4.54

7 Benchmark: Square Tumbler

7.1 System Description

The previous two systems were vertically oriented and induced flow via a carrier fluid. Another type of industrially relevant device is a tumbler or rotating drum, which is perpendicular to gravity and induces particulate flow by causing a continual cascade. Tumblers and dryers can be used for drying, spray coating, granulation, and milling (Yang, Zou, and Yu 2003) and are familiar in the food-processing and pharmaceutical industries (Jacob 2007).

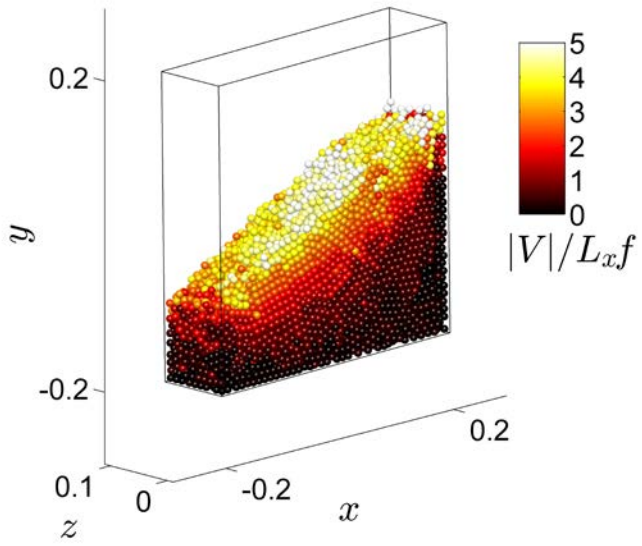

Figure 14. A snapshot of the (serial) square tumbler at the end of initialization

Typically, real devices are cylindrical or canonical in shape and contain internal components such as blades, augers, or sprayers (Jacob 2007). The problem is simplified here using a simple hollow tumbler of a square cross section. The serial version, which resembles a thin slice along the axis, is illustrated in Figure 14. Although a square tumbler does not need a cut-cell geometry, it significantly complicates the boundary motion. Given a constant rotation frequency, the velocity of the outer wall of a cylinder is also constant. For a square tumbler, the wall speed is a function of its location along the perimeter—i.e., the distance from the axis of rotation to the wall is a function of its location on the square. To overcome this difficulty without adding complicated wall BCs, the gravity vector is simply rotated counterclockwise: $g_x = |\mathbf{g}| \sin(\theta), g_y = |\mathbf{g}| \cos(\theta)$ and $g_z = 0$. Although this trick is not physically accurate, it does faithfully reproduce the general dynamics of the intended problem. The time-dependent change to the body force is reflected in the subroutines: bodyforce_mod.f, set_constants.f, and des_time_march.f. The angle of inclination is given by a constant angular frequency: $\theta_t = \omega$, where ω is a constant, or $\theta(t) = \theta_0 + \omega(t - t_0)$. We specify $\theta_0 = 0$ at $t_0 = 0$, and the angular frequency is $\omega = \pi$ (rad/s); or equivalently the period is $T = 2$ (s), or the frequency is $f = 1/2$ Hz = 30 (rpm).

At the true IC, particles are uniformly, randomly distributed throughout the domain with near-zero initial velocity. The gas velocity (if present) is zero everywhere. Initially gravity acts in the negative y-direction. The particles fall to the bottom of the tumbler, similar to the settling

problem albeit in a different domain. This initial transient produces a noticeably rapid increase and decrease in mean particle kinetic energy, as evidenced by Figure 15. As the tumbler begins to "turn" (recall that it is really gravity that is turning), the particles begin to cascade. However, unlike the round tumbler, which cascades quasi-uniformly, the dynamics of the square tumbler are punctuated by larger cascades at a frequency of 4*f*. The larger cascades occur after the box turns from the diamond position back to the square position, which is labeled and sketched in Figure 15. Although the simulation becomes quasi-steady after approximately a quarter turn, the initial transient is run off-line for half a rotation, which is then used as the IC for subsequent tests. The instantaneous dynamics of the serial problem at half a rotation are shown in Figure 14.

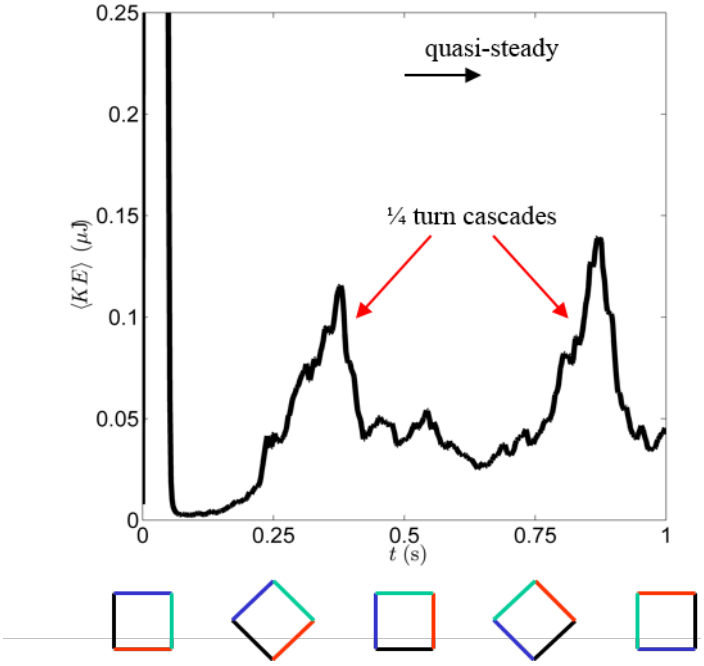

Figure 15. Temporal evolution of the monodispersed square tumbler

In the serial case, the domain resembles a slice with edges of length $L_x = L_y = 40$ cm and $L_z = 0.10$ cm. A cubic CFD grid of $\Delta^* = 2$ is applied ($N_x = N_y = 20$, $N_z = 5$). NSWs are specified at every BC. The solids concentration is set at $\phi = 0.30$ for $N_p = 9{,}165$ particles—the larger particle/processor count of the benchmarks. The simulations run for an additional quarter turn (0.25 s) after the initial transient.

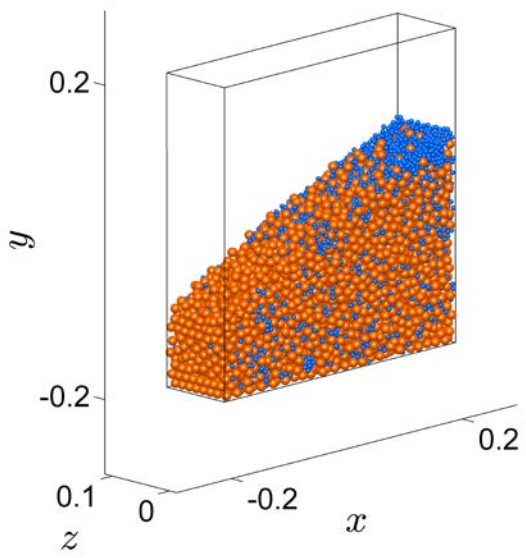

Figure 16. A snapshot of the (serial) square tumbler with a bidisperse 50/50 (number) mixture of larger particles (orange) and smaller particles (blue) with a diameter ratio of two

7.2 Variants

The square tumbler is also used as a test bed for polydispersity. A bidisperse mixture of 50/50 concentration replaces the previous monodisperse particles. The particle diameter ratio is specified as $d_{p1}/d_{p2} = 2$, and the concentration remains the same as it is in the monodisperse case. An algebraic equation can then be used to find $d_{p1} = 121.2$ μm and $d_{p2} = 60.6$ μm. Because 9,165 does not split evenly, the extra particle goes to d_{p2}. The serial case is displayed in Figure 16 after initialization.

7.3 Weak Scaling Results

The serial problem is scaled in one dimension by replicating the serial problem into *np* sub-domains. The true IC and CFD grid of the serial problem is used in each sub-domain. Due to the initial transient period and the end walls, some slight discrepancies may exist in the computational load of each processor. However, visual observations show that the system remains relatively uniform even at large *np*, i.e., long tumblers.

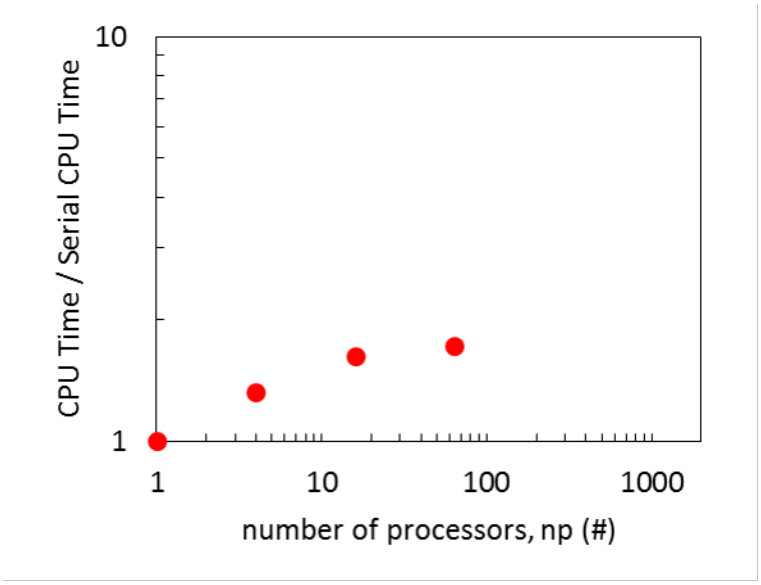

Figure 17. Weak scaling results for the monodisperse square tumbler problem

The weak scaling results for the monodisperse case are presented in Figure 17. The serial simulation took 2.8 h of real time. Unfortunately, the largest cases, *np* = 256 and *np* = 1,024, experienced an initial non-convergence issue that did not appear in the smaller cases. Future studies will test different ICs—i.e., longer initial transient periods—to see if this can overcome the difficulty. Otherwise it may be necessary to decrease the minimum allowable time step, but such an adjustment could offset the ideal scaling needed for a consistent analysis.

7.4 Profiling Results

Not surprisingly, the profile of the tumbler benchmark in Table 6 resembles that of the settling benchmark. The DEM contact force calculation tops the list, owing to many enduring contacts on the bottom of the tumbler. A subroutine related to neighbors (for contact detection) also makes the top five, but it is not the same subroutine as in the settling benchmark.

Table 6. Profiling summary of the Top Five Functions for the CFD-DEM Settling Benchmark

Function	% of time
CALC_FORCE_DEM	31.30
CALC_DRAG_DES	20.50
COMP_MEAN_FIELDS	13.30
NEIGHBOUR	6.60
PARTICLES_IN_CELL	5.40

8 Performance Assessment

Several metrics were employed to assess the performance of the runs:

- Loop metric: This is the percentage of time the benchmark problems spend in vectorized loops, scalar loops, and outside of loops.
- Floating point operations per CPU cycle: This is the average percentage of time the CPU spends doing floating point operations per cycle.
- Floating point vectorization: This is the percentage of how vectorizable the floating point operations are.
- Level 2 and Level 3 cache miss ratio: This is the percentage of time the data was unavailable in the respective caches (memory).

The results for the five benchmark problems have been compiled in Table 7. Note that these initial performance metrics were averaged over five runs of the single-process (serial) benchmarks.

Table 7. Comparison of the Loop Metrics for the Five Benchmark Problems

Loop Metrics (%)	HCS	Settling	FB	Riser	Tumbler
Vectorized	6.0	7.2	2.0	6.3	5.1
Scalar	61.2	55.1	58.0	53.0	57.8
Outside	32.8	37.7	40.0	40.7	37.1

Processors are supporting wider single instruction, multiple data (SIMD) operations. The current x86 processors support at least 128-bit SIMD registers. Future processors and co-processors are increasing this width to 512 bits. To take advantage of these hardware improvements, codes must be vectorized. In the five benchmark problems, we see a low percentage of vectorized loops compared to the scalar loop percentage. Future work will focus on increasing the percentage of vectorized loops within MFiX. Increased vectorization can be accomplished by:

1. Aligning data arrays on vectorizable byte boundaries by using compiler directives
2. Providing subroutines with information about the vectorized arrays
3. Using OpenMP SIMD directives on the loops.

The floating point and cache miss metrics are reported in Table 8 for the five benchmark problems. The floating point operations per cycle for the five benchmark problems are relatively low, yielding to longer processing times and time to solution. The wide percentage spread indicates the computational divergence among the problems. The low floating point vectorization percentage that indicates the present code is not able to benefit from significant performance gains that are capable on modern SIMD processors. The memory (Level 2 and Level 3) cache miss ratio indicates the locality of data to the processors. The Level 2 ratio is relatively low; an optimized program will often see this value around the range of 90%. The low Level 3 ratio indicates that the benchmark problems are data extensive and stress the memory bandwidth of the system.

Table 8. Comparison of CPU Statistics for the Five Benchmark Problems

Loop Metrics (%)	HCS	Settling	FB	Riser	Tumbler
Floating point operations/cycle	13	21	21	14	28
Floating point vectorization	1	8	7	2	7
Level 2 cache miss ratio	35	28	22	15	17
Level 3 cache miss ratio	0	15	14	0	2

9 Conclusions

This technical report documents an initial benchmarking and profiling summary for NETL's multiphase flow CFD-DEM code, MFiX. The performance of MFiX 2015-1.5 (developmental version cloned on September 8, 2015) is assessed and documented, and it will serve as a baseline for future improvements to the code aimed at speeding up the process. To establish the baseline performance, we developed five benchmark problems:

- HCS: The HCS is the most dilute of the five problems with a solids concentration of only 1%. The system is fully periodic, and the initial granular temperature (a measure of kinetic energy) simply decays with time because of inelastic particle collisions. Because of the double-periodicity of the initial condition, instabilities do not develop as the system size increases (favorable for weak scaling). The system is CFD intensive with a particle-to-grid ratio of $N_p/N_xN_yN_z = 0.153$.

- Settling: The settling case considers a random distribution of particles in a box with NSWs. The particles fall to the bottom of the domain under gravity and reach a static, packed state. The turnaround time for the simulations is relatively fast. The system is DEM intensive with a particle-to-grid ratio of $N_p/N_xN_yN_z = 1.61$. Pure-DEM (granular) simulations of the settling system are also studied.

- FB: The FB benchmark represents a simplified version of a typical industrial bubbling fluidized bed. The system is classified as Geldart Group A, and it is operated in the bubbling regime. Air is injected uniformly at the bottom of the domain, the side walls are no-slip, and a pressure outflow is set at the exit. An initial transient period of 1,000 ms is carried out off-line, and the benchmark tests are specified as restarts (type: restart_2). The problem is well balanced with a particle-to-grid ratio of $N_p/N_xN_yN_z = 1.125$. The FB will also be used in the future to study heat transfer; presently, MFiX 2015-1.5 must be modified to include heat transfer, as discussed in Section 5.

- Riser: The riser problem simulates conditions representative of the centerline of an industrial circulating fluidized bed. It is slightly more dense than the HCS with a concentration of 1%. The system is also fully periodic, but the vertical direction includes a pressure drop, which equals the weight of the bed, and constant gas mass flux is specified. An initial transient of 500 ms is simulated before the restart benchmarks begin—a sufficient amount of time to develop inhomogeneous particle spatial distribution, i.e., clustering. Although the presence of clustering presents some challenges for a weak scaling analysis (e.g., it is unlikely to have an equal computational load per processor), it is a physically relevant phenomena encountered in virtually any practical system. The riser has a particle-to-grid ratio of $N_p/N_xN_yN_z = 0.313$.

- Square tumbler: In contrast to the FB and riser, the tumbler problem is horizontal (perpendicular to gravity) and induces particle flow by (mimicking) boundary motion. To use a rectangular grid, the tumbler cross section is square, which presents an additional challenge: the wall speed is not a constant (due to differences in distance between the location of the wall and the axis of rotation). Therefore, the rotation of the tumbler is faked by simply rotating the gravity vector, which fatefully reproduces the cascading

dynamics in a static domain. Like the settling case, the tumbler is fairly DEM dominate with a particle-to-grid ratio of $N_p/N_xN_yN_z$ = 4.58. The tumbler is also used to study polydispersity. A bidisperse mixture has been considered, which maintains the same solids concentration of the monodisperse case and a species diameter ratio of two.

All five benchmark problems are made spatially scalable to perform a weak scaling analysis. Every scaled case is designed to have the same grid and particle load per processor as in the serial case, although this is likely violated (to varying degrees) for the industrially relevant benchmarks. The weak scaling results were generally favorable up to np = 1,000. However, several benchmarks indicate that a significant divergence in scalability may be encountered beyond 1,000 processors. To achieve the project's final goal of an $N_p = 10^8$ simulation at this level of parallelization, each processor would be required to carry a load of 10^5 N_p/np. Such a large processor load may be too slow to gain support for industrial research and development (consider that the case of $Np \sim 10^4$ of the square tumbler has a real ratio of CPU time to simulation time of nearly $4 \cdot 10^5$). Therefore, improvements to either, if not both, the serial speed of larger N_p/np problems or the scalability of intermediately sized N_p/np problems appears necessary.

The benchmark cases were also profiled with Perfsuite and Intel VTune. Three themes were observed: subroutines related to the DEM force calculation, the interpolations between the discrete data and continuum data, and the interfacial drag force calculation were consistently in the top five. These initial tests indicate that significant gains can be achieved by targeting DEM-specific and CFD-DEM bridging algorithms. The performance assessments of the benchmark cases have shown that the code is underutilizing the floating point performance and vectorization operations available within current processors. By addressing this low performance, significant gains can be achieved through vectorization and by optimizing the floating point operations to target present and future processors.

References

Abrahamsen, A.R., and D. Geldart. 1980. "Behavior of Gas-Fluidized Beds of Fine Powders .1. Homogeneous Expansion." *Powder Technology* 26(1): 35–46.

Agrawal, K, Loezos, P.N., Syamlal, M, and Sundareson, S. 2001. "The Role of Meso-scale Structures in Rapid Gas-Solid Flows." *Journal of Fluid Mechanics* 445: 151–185.

Beetstra, R., M.A. van der Hoef, and J.A.M. Kuipers. 2007. "Drag Force of Intermediate Reynolds Number Flow Past Mono- and Bidisperse Arrays of Spheres." *Aiche Journal* 53(2): 489–501.

Capecelatro, J., O. Desjardins, and R.O. Fox. 2015. "On Fluid-Particle Dynamics in Fully Developed Cluster-Induced Turbulence." *Journal of Fluid Mechanics* 780: 578–635.

Capecelatro, J., P. Pepiot, and O. Desjardins. 2014. "Numerical Characterization and Modeling of Particle Clustering in Wall-Bounded Vertical Risers." *Chemical Engineering Journal* 245: 295–310.

Capecelatro, J., P. Pepiot, and O. Desjardins. 2015. "Numerical Investigation and Modeling of Reacting Gas-Solid Flows in the Presence of Clusters." *Chemical Engineering Science* 122: 403–415.

Carnahan, N.F., and K.E. Starling. 1969. "Equation of State for Nonattracting Rigid Spheres." *Journal of Chemical Physics* 51(2): 635–636.

Chew, J. W., R. Hays, J.G. Findlay, T.M. Knowlton, S.B.R. Karri, R.A. Cocco, and C.M. Hrenya. 2012. "Cluster Characteristics of Geldart Group B Particles in a Pilot-Scale CFB Riser. I. Monodisperse Systems." *Chemical Engineering Science* 68(1): 72–81.

CU-RC. 2015. "JANUS Supercomputer." from https://www.rc.colorado.edu/resources/compute/janus.

Fan, L.-S., and C. Zhu. 1998. *Principles of Gas-Solid Flows*. Cambridge: Cambridge University Press.

Garzo, V. 2005. "Instabilities in a Free Granular Fluid Described by the Enskog Equation." *Physical Review E* 72(2).

Garzó, V., S. Tenneti, S. Subramaniam, and C.M. Hrenya. 2012. "Enskog Kinetic Theory for Monodisperse Gas–Solid Flows." *Journal of Fluid Mechanics* 712: 129–168.

Geldart, D. 1973. "Types of Gas Fluidization." *Powder Technology* 7(5): 285–292.

Haff, P. K. 1983. "Grain Flow as a Fluid-Mechanical Phenomenon." *Journal of Fluid Mechanics* 134(Sep): 401–430.

Jacob, M. 2007. "Granulation Equipment." *Granulation* 11: 417–476.

Koch, D.L., and A.S. Sangani. 1999. "Particle Pressure and Marginal Stability Limits for a Homogeneous Monodiperse Gas-Fluidized Bed: Kinetic Theory and Numerical Simulations." *Journal of Fluid Mechanics* 400(1): 229–263.

NETL. 2015a. "Crosscutting Technology Research." from http://www.netl.doe.gov/research/coal/crosscutting.

NETL. 2015b. *MFIX 2015-2 Users Guide*.

NETL. 2015c. "MFS NETL Multiphase Flow Science." from https://mfix.netl.doe.gov/.

Rhodes, M.J. 2008. *Introduction to particle technology*. Chichester, England: Wiley.

Tang, Y., E.A.J.F. Peters, J.A.M. Kuipers, S.H.L. Kriebitzsch, and M.A. van der Hoef. 2015. "A New Drag Correlation from Fully Resolved Simulations of Flow Past Monodisperse Static Arrays of Spheres." *AIChE Journal* 61(2): 688–698.

Wylie, J.J., D.L. Koch, and A.J. Ladd. 2003. "Rheology of Suspensions with High Particle Inertia and Moderate Fluid Inertia." *Journal of Fluid Mechanics* 480: 95–118.
Yang, R.Y., R.P. Zou, and A.B. Yu. 2003. "Microdynamic Analysis of Particle Flow in a Horizontal Rotating Drum." *Powder Technology* 130(1–3): 138–146.
Yin, X., J.R. Zenk, P.P. Mitrano, and C.M. Hrenya. 2013. "Impact of Collisional Versus Viscous Dissipation on Flow Instabilities in Gas–Solid Systems." *Journal of Fluid Mechanics* 727: R2.

Appendix: Electronic Attachments

The input requirements for all benchmark problems are currently being stored on the JANUS supercomputer in the directory /projects/mfix/benchmarks.