# *MFIX-DEM Enhancement for Industrially-Relevant Flows*

NREL
NATIONAL RENEWABLE ENERGY LABORATORY

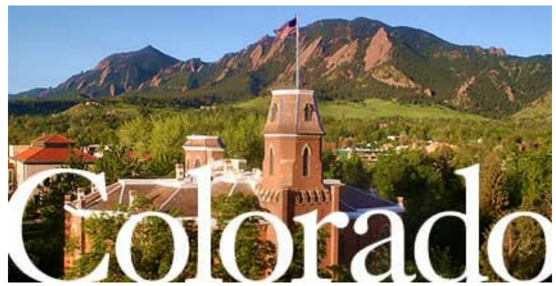


PSRI

***Project Leads:***

***Dr. Ray Cocco (PSRI, co-PI)***
***Dr. Ray Grout (NREL, co-PI)***
***Prof. Thomas Hauser (Univ. CO, co-PI)***
***Prof. Christine Hrenya (Univ. CO, PI)***

# Project Objectives

- Task 1 - Project Management and Planning
- Task 2 - Profiling MFIX DEM
  - Task 2a: Benchmark serial MFIX DEM
  - Task 2b: Benchmark parallelized MFIX DEM
- Task 3 - Determine Optimization Frameworks
- Task 4 - Perform optimization and vectorization of serial DEM
  - Task 4a: Employ optimization techniques
  - Task 4b: Verify enhanced DEM code for numerical correctness
- Task 5 - Optimize and enhance hybrid (MPI + accelerator) DEM
  - Task 5a: Implement hybrid parallelization method (MPI + OpenMP)
  - Task 5b: Use extensive parallel profiling to optimize parallel code
  - Task 5c: Compare enhancements on multiple Xeon/Xeon Phi architectures
- Task 6 - Industrially Relevant Problem

  Task 6a: Survey of PSRI member companies

  Task 6b: Experiments of Interacting Nozzles
- Task 7 - Uncertainty Quantification
  - Task 7a: Test Problem
  - Task 7b: Challenge Problem
  - Task 7c: Industrially Relevant Problem

# Project Team

**University of Colorado**
**Chemical & Biological Engineering**
*DEM modeling of grnanlar and gas-solid flows, MFIX*

**Prof. Christine Hrenya**

**Dr. William Fullmer**

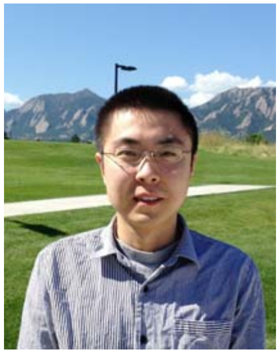

**Dr. Peiyuan Liu**

**University of Colorado**
**Research Computing**
*High-performance computing, CFD*

**Prof. Thomas Hauser**

**Tim Brown**

**NREL**
**Computational Science**
*High-performance computing, CFD*

**Dr. Ray Grout**

**Dr. Hari Sitaraman**

**PSRI**
*Industrial Application and Experiments of Particle Flows*

**Dr. Ray Cocco**

**Dr. Allan Issangya**

# Technical Background

## Continuum or "Two-fluid Model" (TFM)

- Gas = continuum
  (averaged over *many* particles)
- Solids = continuum

$$\frac{\partial}{\partial t}\left(\varepsilon_s \rho_s \mathbf{V}_s\right) + \nabla \cdot \left(\varepsilon_s \rho_s \mathbf{V}_s \mathbf{V}_s\right) = \nabla \cdot \boxed{\boldsymbol{\tau}_s} + \varepsilon_s \rho_s \mathbf{g} + \boxed{\mathbf{F}_{\text{drag}}}$$

## Discrete Element Method (DEM)

- Gas = continuum
- Solids = discrete

$$m_i \frac{d\mathbf{V}_{si}}{dt} = m_i \mathbf{g} + \mathbf{F}_{ci} + \boxed{\mathbf{F}_{drag,i}}$$

Typical CPU times for DEM
Serial processor: $O(10^5$ particles)
Parallel processors: $O(10^8$ particles)

| $d_p$ | $N_p$ in 1 cup |
|---|---|
| 100 μm (sand) | $O(10^8)$ particles |
| 50 μm | $O(10^{10})$ particles |

***Less CPU time***

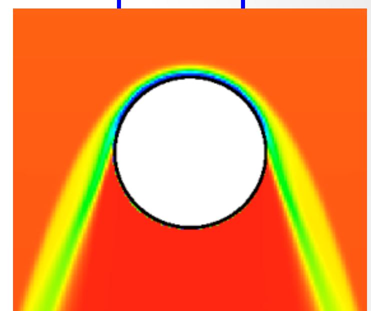

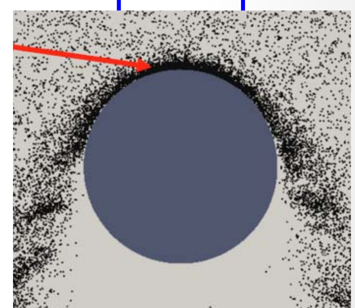

***More detail, Fewer closures***

# Significance of the Results of the Work

- Expanded industrial use of DEM
- Indirect: Improved physics in continuum and hybrid modeling (DEM as benchmark data)
- Direct: Aid in design/optimization of industrial systems

**System-size independent measurements:**
Defluidization of Cohesive/Non-Cohesive Particles
(LaMarche et al., *AIChE J.*, 2015)

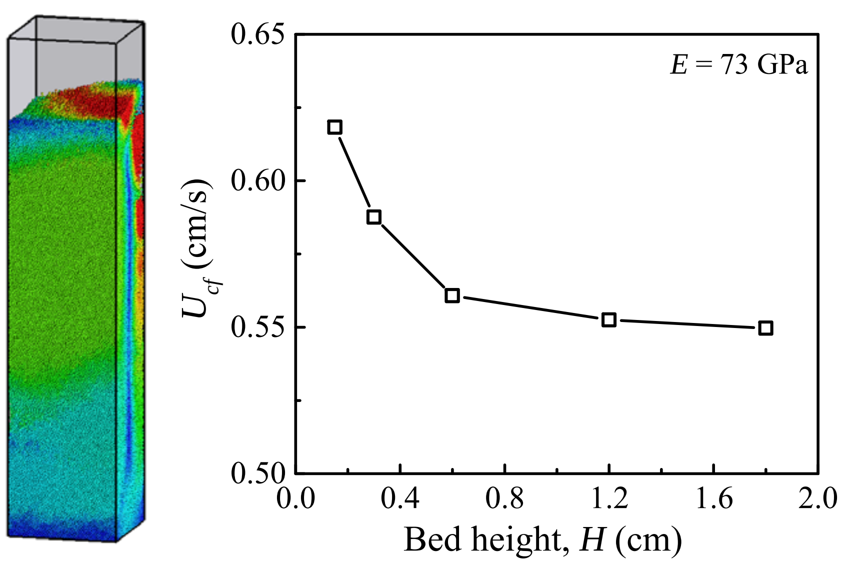


**Fully-developed characteristics:**
Heat transferred to particles falling over heated tubes
(Morris et al., *Solar Energy*, submitted)

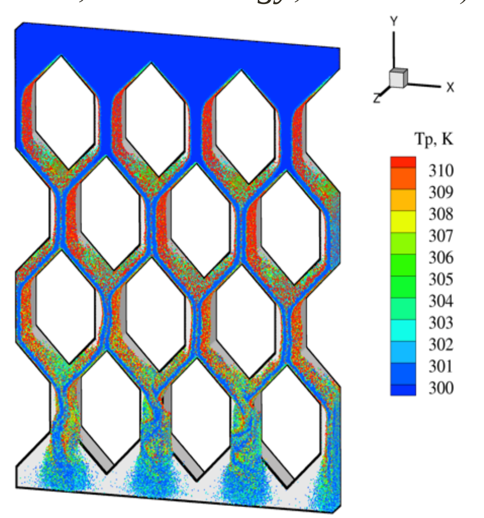

# Survey to Assess Industry Needs

- PSRI initiated
- 18 companies responded
- Interest in DEM applied to
  - Bin/hopper discharge
  - Large valve systems
  - Drug deliver through inhaler device
  - Reacting fluidized bed systems
  - Die filling of non-spherical particles
  - Gas distributers, transfer lines
  - Standpipe flow

# Sector Served / Size

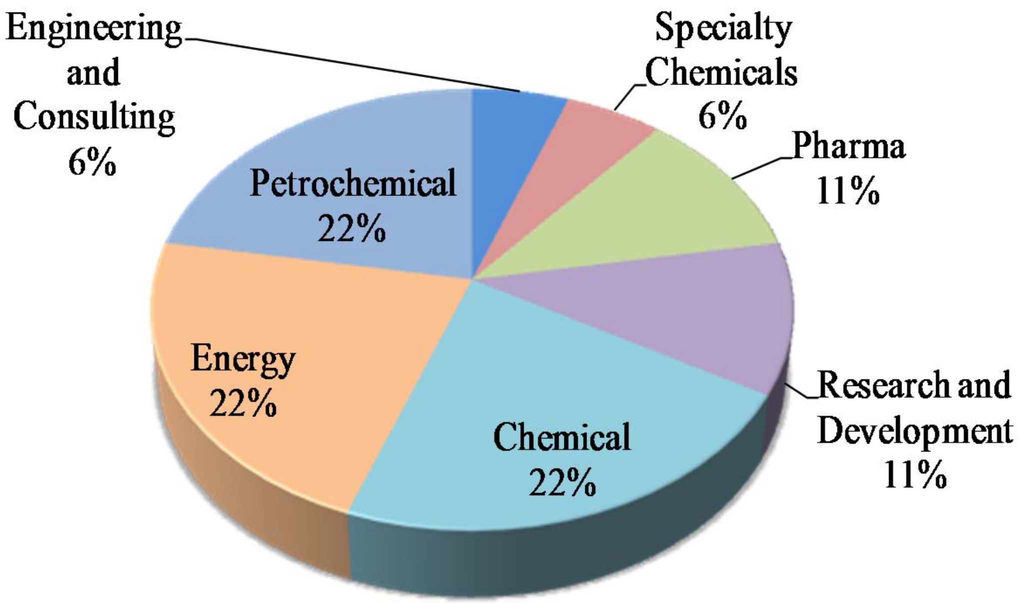


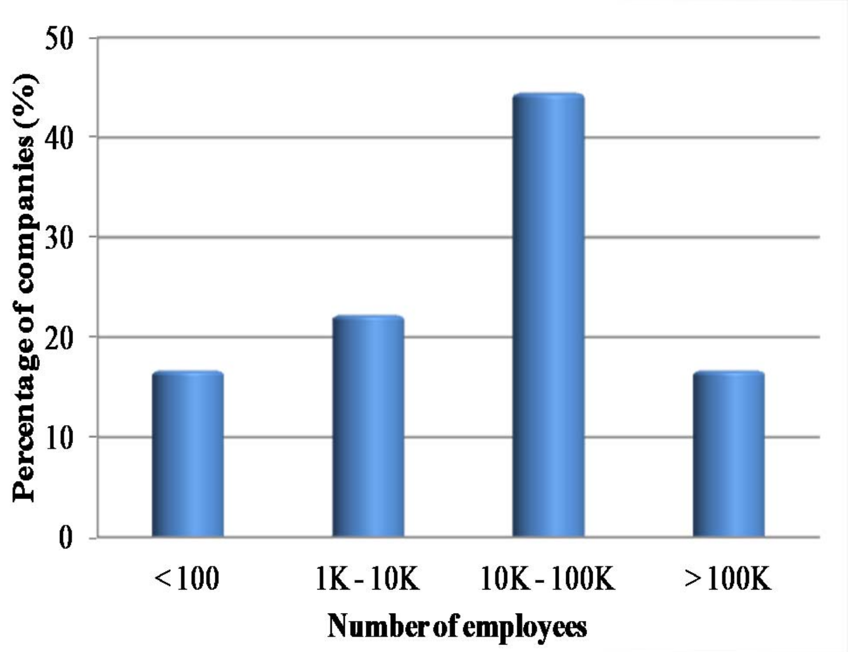

# CFD/DEM Modeling

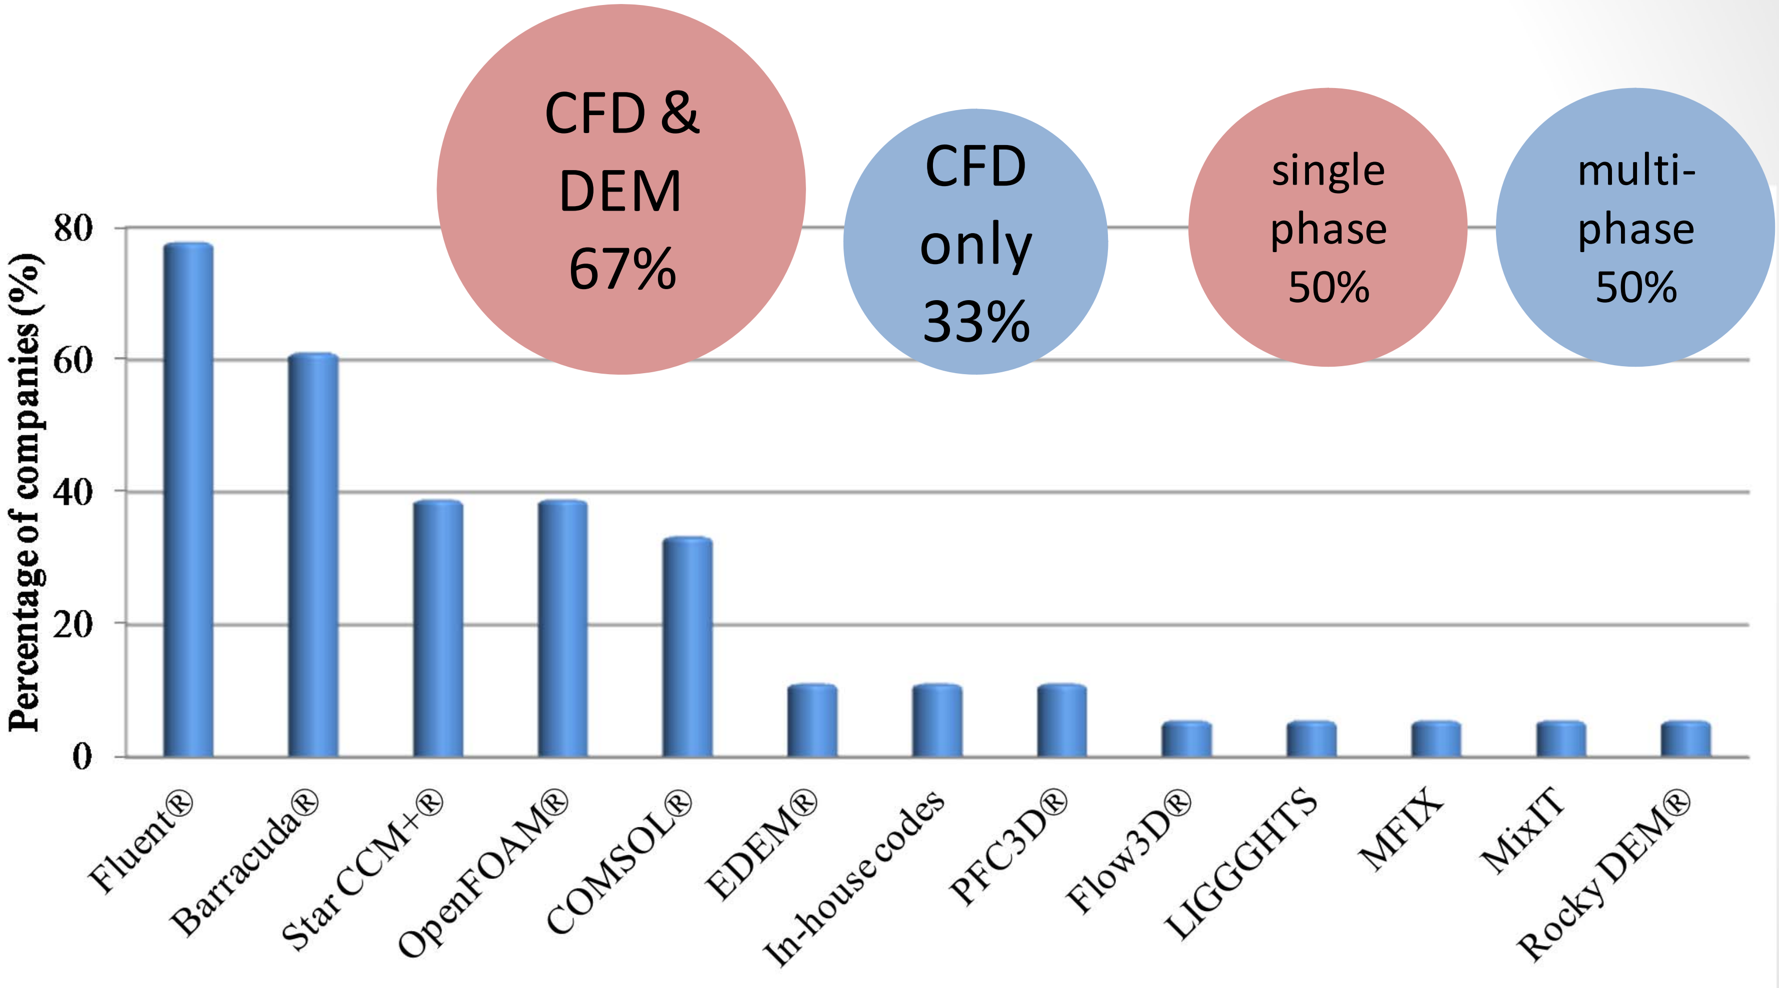

# Criteria for Choosing DEM software

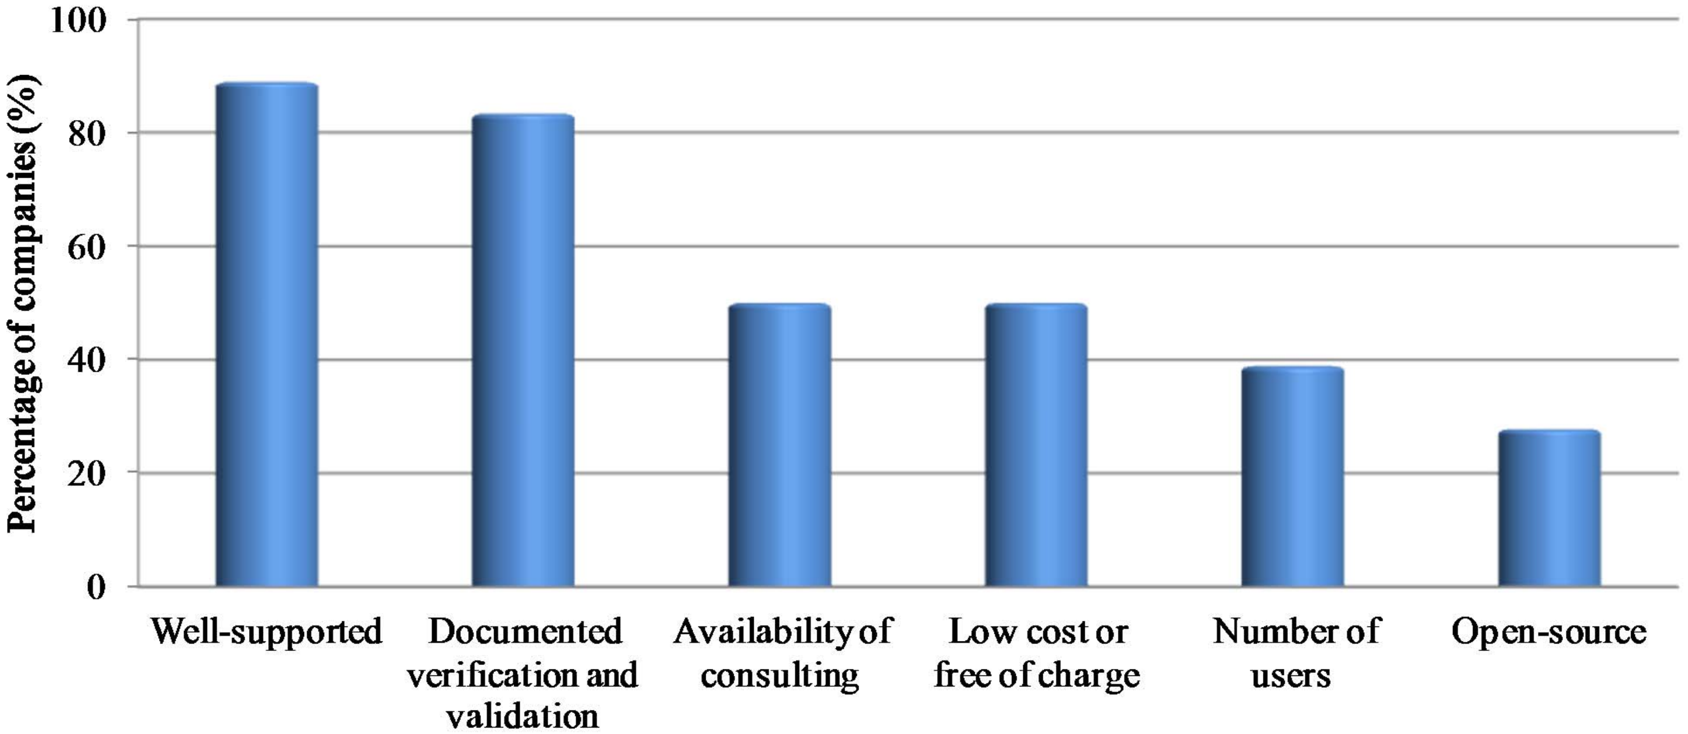

# Usage and Training

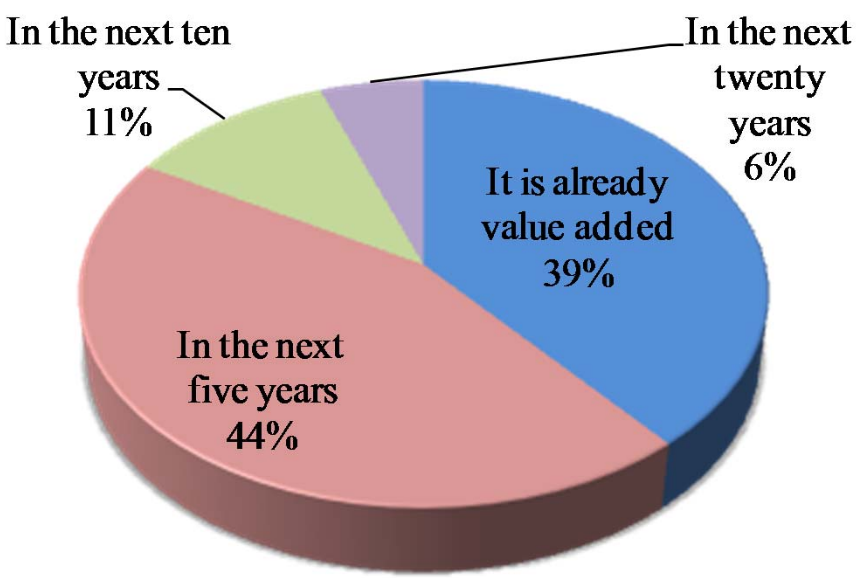


Value added through DEM

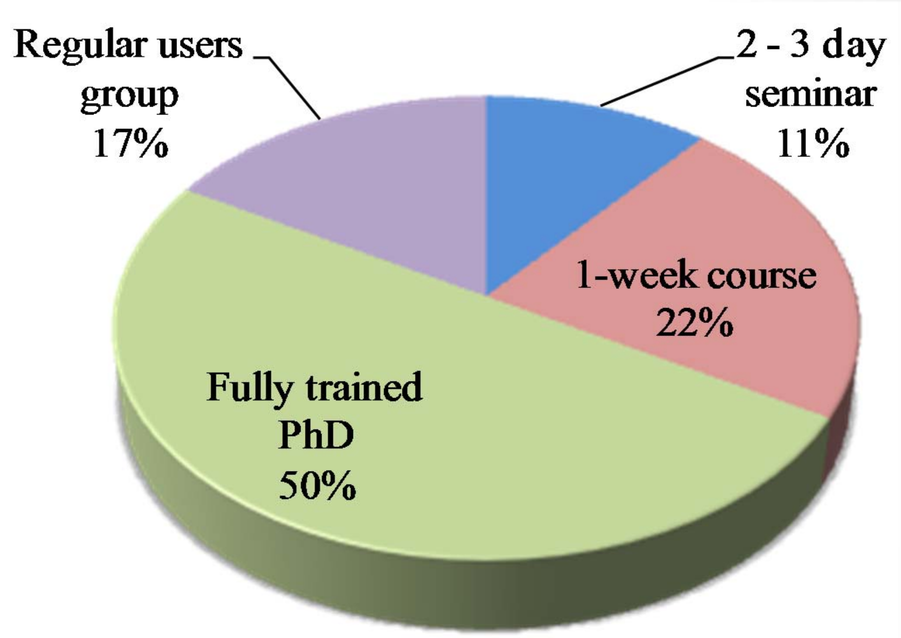


Workforce training

# DEM Physics Enhancements

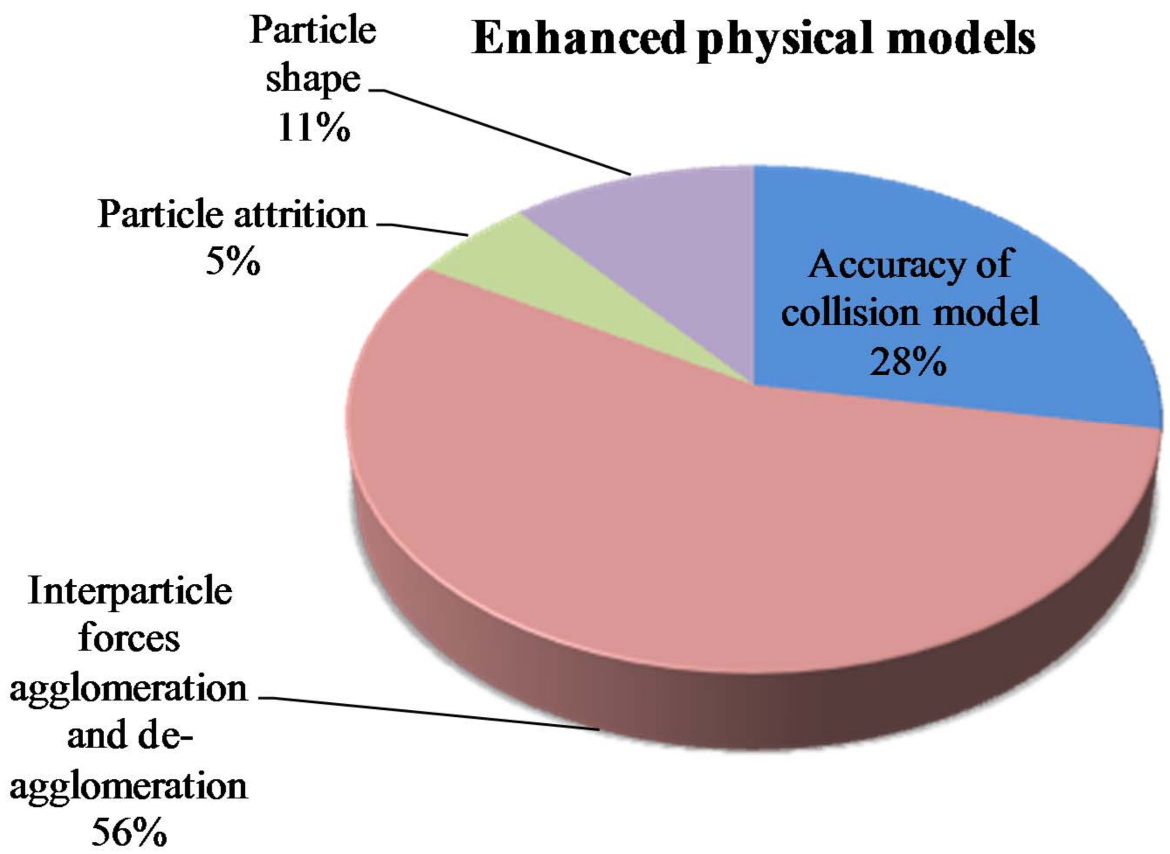

# Computational Improvements

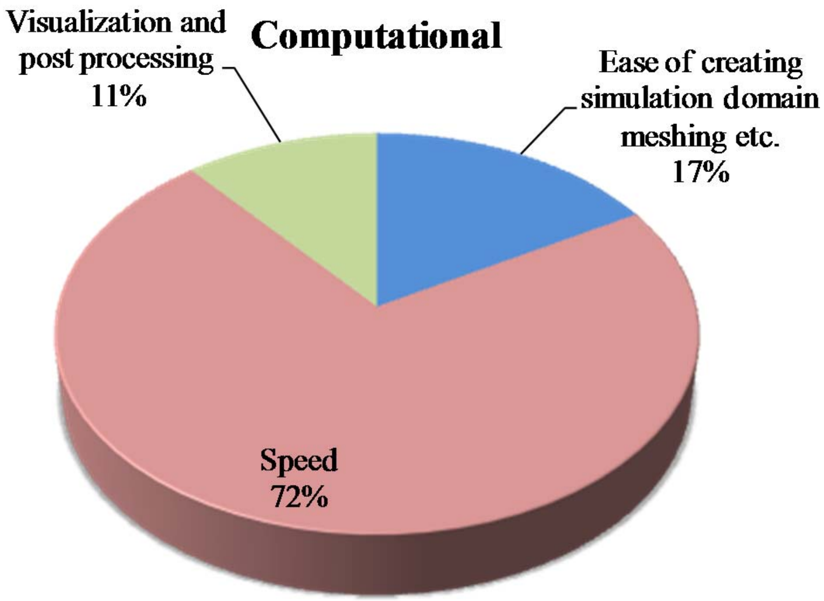


Improvements for more widespread use

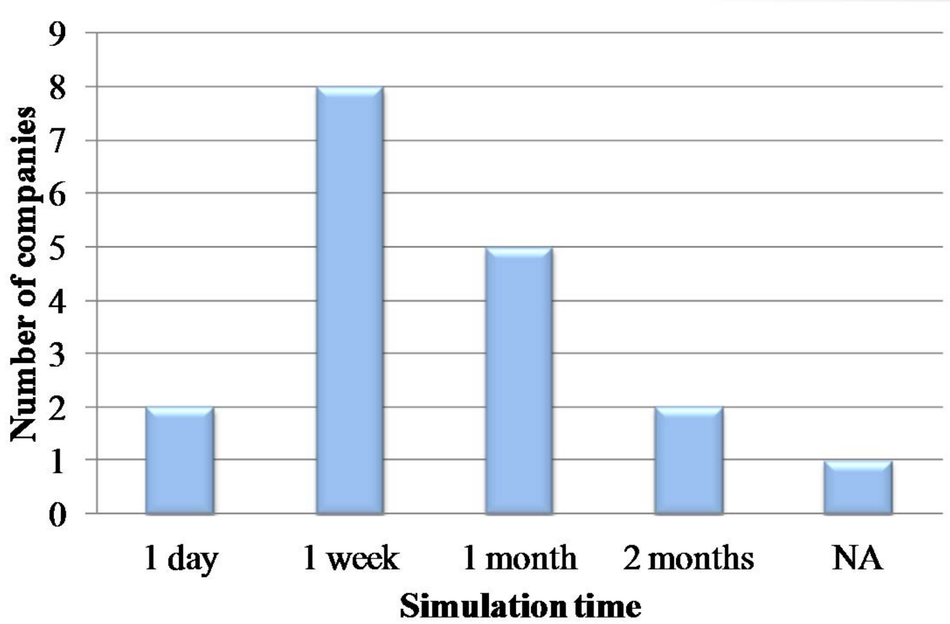


How long should DEM simulations run to add value?

# Computational Project Goals

- Improve the computational performance of MFIX-DEM
  - Create a set of benchmarks to guide optimization work
  - Profile benchmarks
  - Improve serial performance of select subroutines by a factor of 2
    - Vectorization
    - Optimization of memory access
  - Improve parallel performance of MFIX-DEM
    - Hybrid OpenMP + MPI parallelization
- Demonstrate performance enhancements through DEM simulation with $O(10^8)$ particles
- Explore algorithmic changes

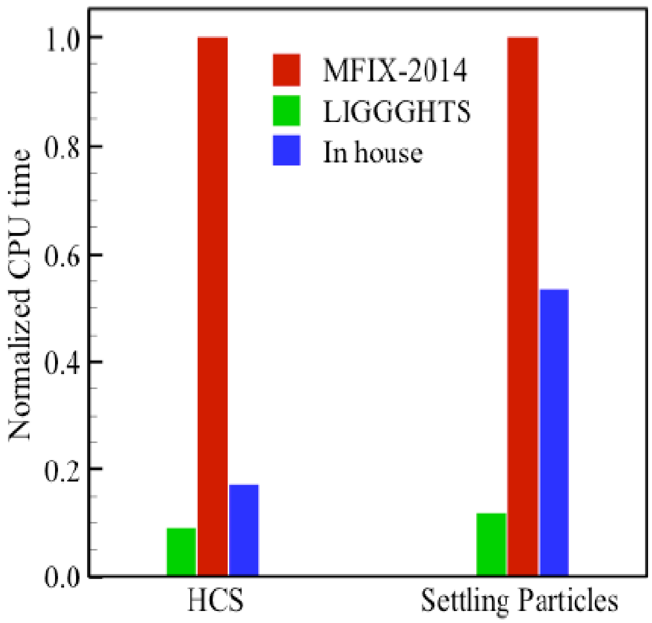


Preliminary results: comparison of CPU time normalized by MFIX CPU time based on serial DEM simulations of homogeneous cooling systems (HCS) and settling particles.

# Computer Architecture 101

## Stampede Supercomputer

Stampede has 6,400 nodes [3]
56 GB/s FDR Infiniband interconnect

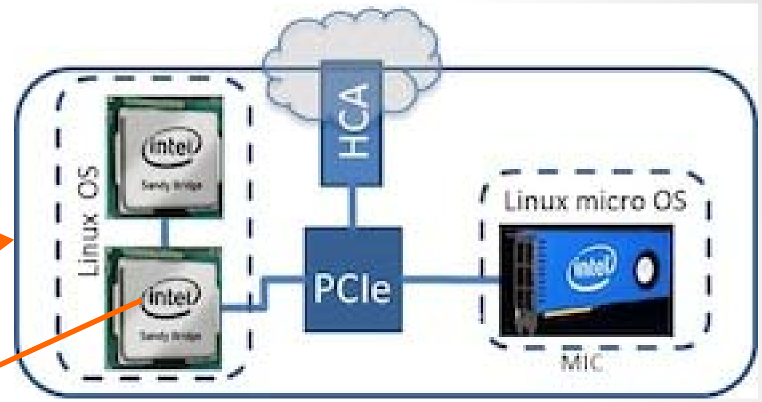


Stampede compute node [3]:

- 2 Sockets per Node → 2 Xeon E5 processors
- 1 Xeon Phi coprocessor
- 32 GB Memory
- 250 GB Disk


www.scan.co.uk

Socket:

- 2.7 GHz
- 8 Cores
- 8 DP FP operations per clock cycle
- 64 GB L1 Cache/core
- Vector width: 4 double precision items

# Parallelism at All Levels

- Parallelism across multiple nodes or processors
- Parallelism across threads
- Parallelism across instructions
- Parallelism on data – SIMD (Single Instruction Multiple Data)

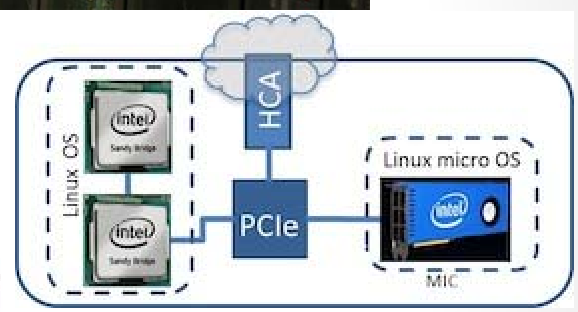



www.scan.co.uk

# Floating Point Performance

$$P = n_{\mathrm{core}} * F * S * \nu$$

- Example: Intel Xeon E5 on Stampede
  - Number of cores: 8 $n_{\mathrm{core}}$
  - FP instructions per cycle: 2 (1 Multiply and 1 add) $F$
  - FP operations / instruction (SIMD): 4 (dp) / 8 (sp) $S$
  - Clock speed: 2.7 GHZ $\nu$

$$P = 173\ GF/s\ (dp) \quad \mathrm{or} \quad 346\ GF/s\ (sp)$$

  - But: P= 5.4 GF/s (dp) for serial, non-SIMD code

# Benchmark #1: Settling

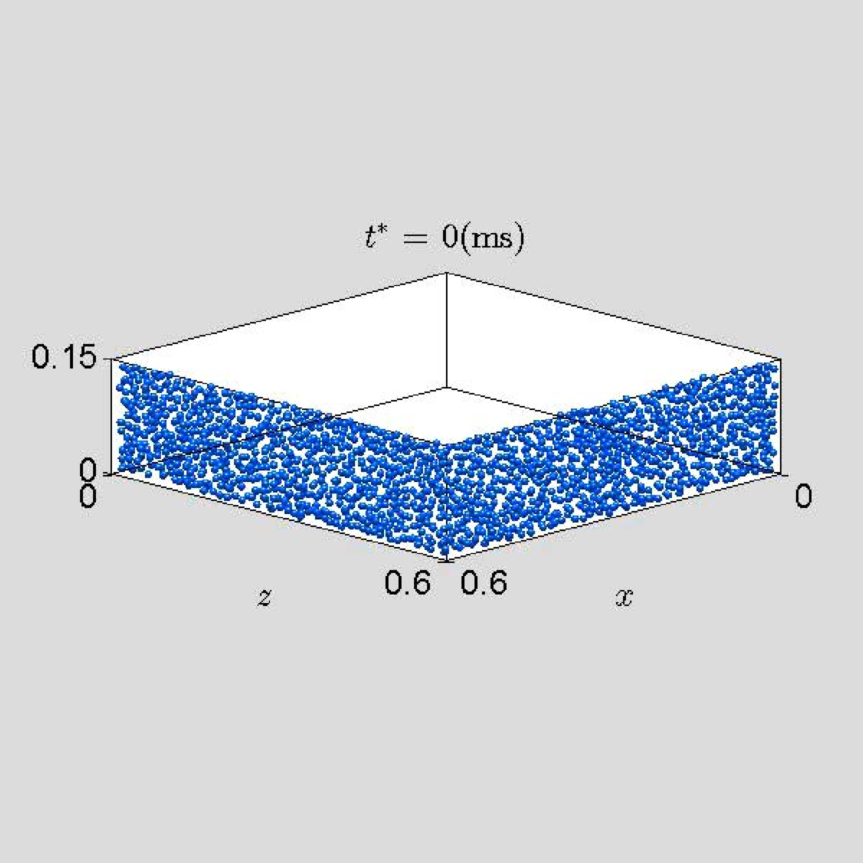


(Only showing $2d_p$ layer near two walls for clarity)

- Uniform, randomly distributed particles fall to rest in enclosed container under gravity
- Initial random generation of particle location with zero speed
- No slip for all walls
- Quick turn-around time
  - Simulation time = 50 ms
  - CPU time < 6 mins (serial)
- Scaled in “2-D” to avoid load balancing issues
  - x-scale = $np^{1/2}$
  - y-scale = none
  - z-scale = $np^{1/2}$
- Variants: also used to test pure-DEM

# Performance results #1: Settling

Weak Scaling

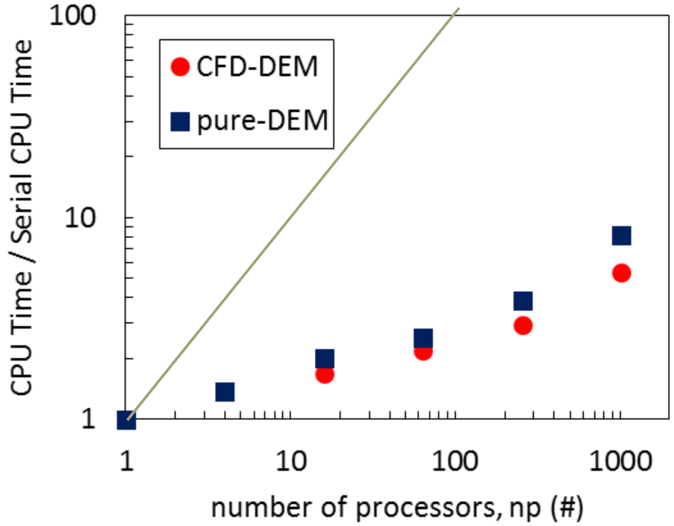


Strong Scaling (of shorter *np* = 1024 case)

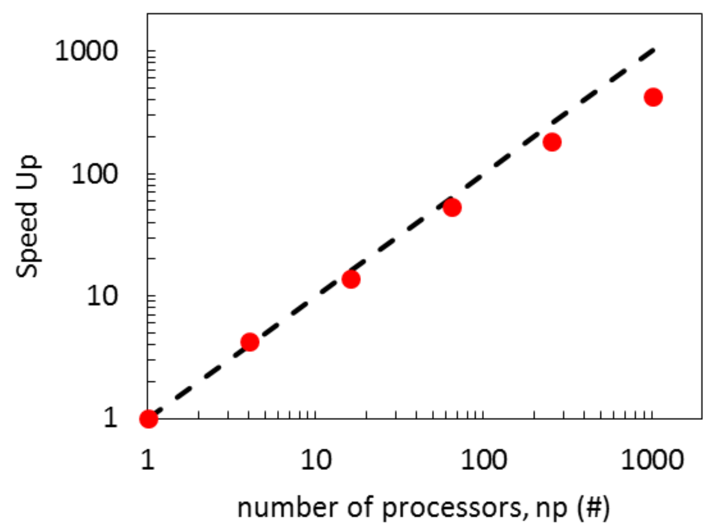


| **Weak scaling analysis for settling case** | | | |
|---|---|---|---|
| $N_p$/np (both) | $N_x N_y N_z$/np (CFD-DEM only) | Serial CPU time (CFD-DEM) | Serial CPU time (pure-DEM) |
| 1611 | 1000 ($10^3$) | 349 (s) | 160 (s) |

- Poor weak scaling up to np = ~1000 (~85 nodes on CU-Boulders supercomputer)
- CFD-DEM and pure-DEM scale very similarly
  (note: $N_{x,y,z}$ = 10 and gas is nearly static not a CFD-intensive problem)

# Performance results #1: Settling

| **Benchmark Statistics (Percentages)** | |
|---|---|
| FP operations/cycle | 21.5 |
| FP vectorization | 8 |
| Level 2 cache miss ratio | 28 |
| Level 3 cache miss ratio | 15 |

| **Loop Metrics (Percentages)** | |
|---|---|
| Vectorized | 7.2 |
| Scalar | 55.1 |
| Outside | 37.7 |

**Profiling Summary (Top 5 functions)**

| Function | Percentage 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 |

Intel(R) Xeon(R) CPU X5660 @ 2.80GHz
24 Processors per node
24 GB RAM per node

# Benchmark #2: Fluidized Bed

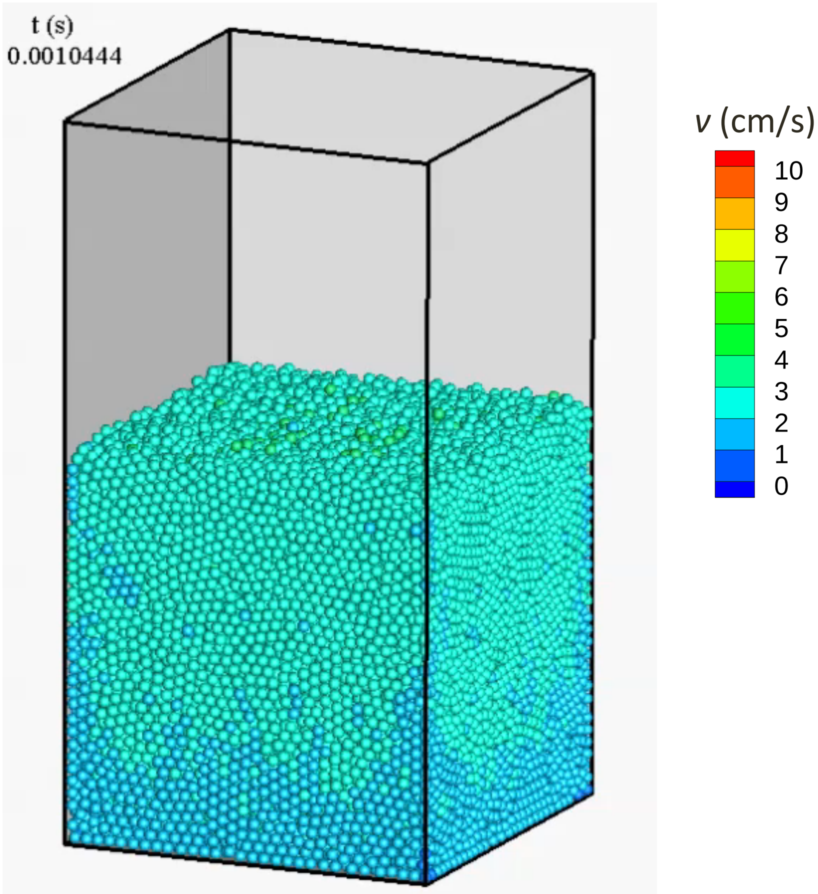


- Rectangular bed with uniform gas inlet and no-slip side walls
- Coexistence of dilute (bubble) and dense (emulsion) phases of solids
- Importance of lasting contact and friction force
- Experimental data available from NETL Challenge Problem: Small Scale Problem I
- Wide industrial applications
- Variants: can test heat transfer

# Performance results #2: Fluidized Bed

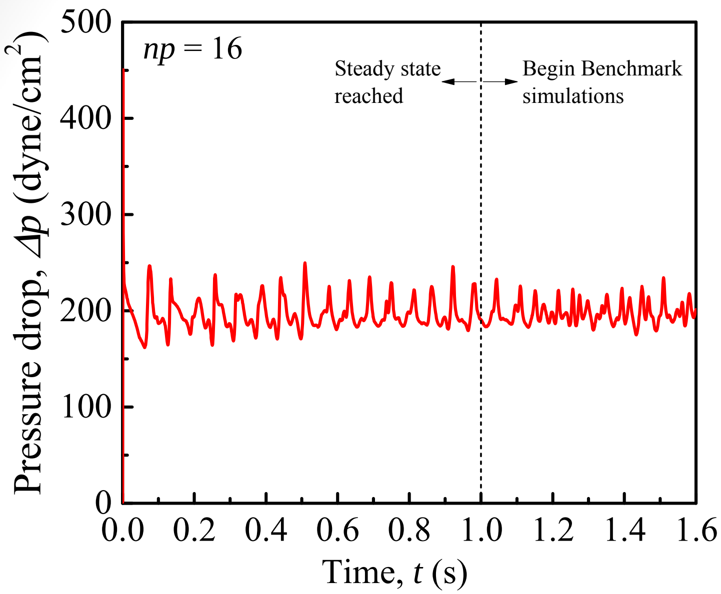


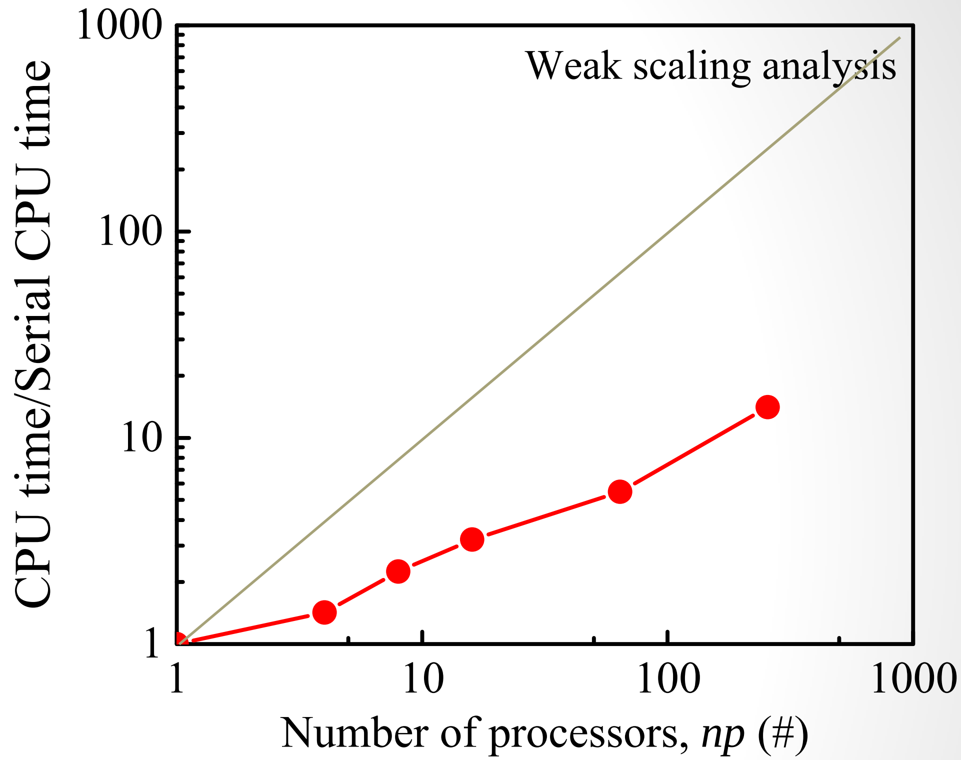


| Weak scaling analysis for fluidized bed | |
|---|---|
| Number of particles, $N_p$ | 2500 ~ 2560000 |
| Number of procs, np | 1 ~ 1024 |
| Number of particles per procs | 2500 |
| Number of CFD cell per procs | 800 |
| Simulated time, $t_s$ (s) | 0.05 |
| Serial CPU time, $t_{cpu, serial}$, (h) | 0.1 |

Benchmark simulations begins after fluidized bed reached steady state, reflected the time evolution of bed pressure drop

# Performance results #2: Fluidized Bed

| Benchmark Statistics (Percentages) | |
|---|---|
| FP operations/cycle | 21 |
| FP vectorization | 7 |
| Level 2 cache miss ratio | 22 |
| Level 3 cache miss ratio | 14 |

| Loop Metrics (Percentages) | |
|---|---|
| Vectorized | 2 |
| Scalar | 58 |
| Outside | 40 |

| Profiling Summary (Top 5 functions) | |
|---|---|
| Function | Percentage of time |
| COMP_MEAN_FIELDS0 | 23.13 |
| DRAG_GS_DES0 | 12.17 |
| CALC_FORCE_DEM | 10.31 |
| DRAG_GS | 7.2 |
| CFNEWVALUES | 6.2 |

Intel(R) Xeon(R) CPU X5660 @ 2.80GHz
24 Processors per node
24 GB RAM per node

# Benchmark #3: Riser Flow

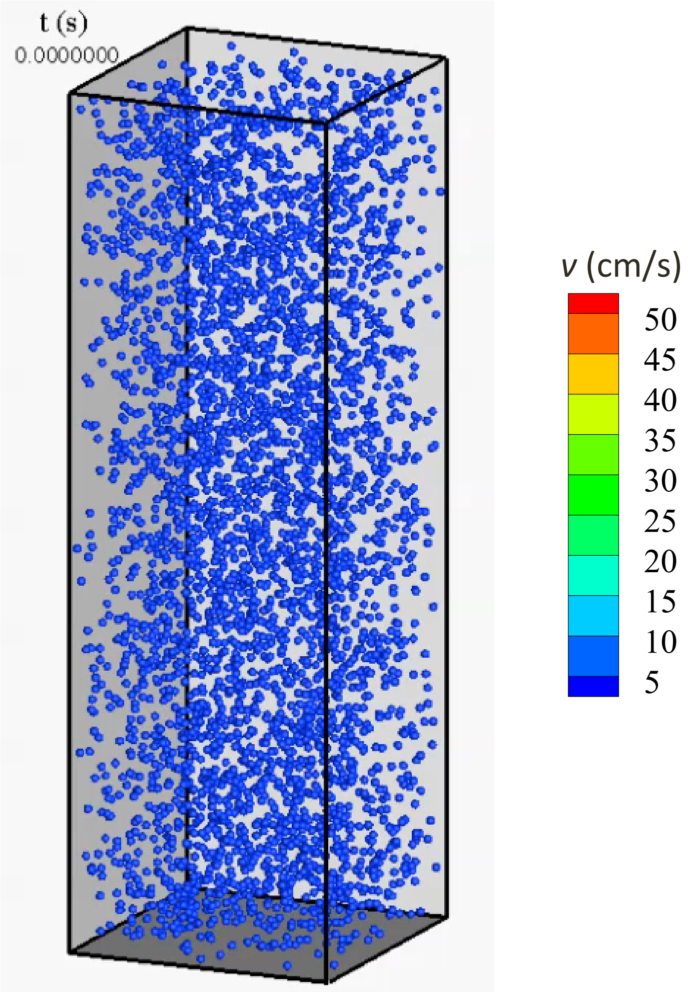


scaled to np = 16

- Rectangular periodic domain
- Constant gas flux and solid concentration maintained
- Solid-solid interactions dominated by brief binary collisions
- Industrial relevance: mimicking fully-developed flow in freeboard of circulating fluidized bed (CFB)

# Performance results #3: Riser Flow

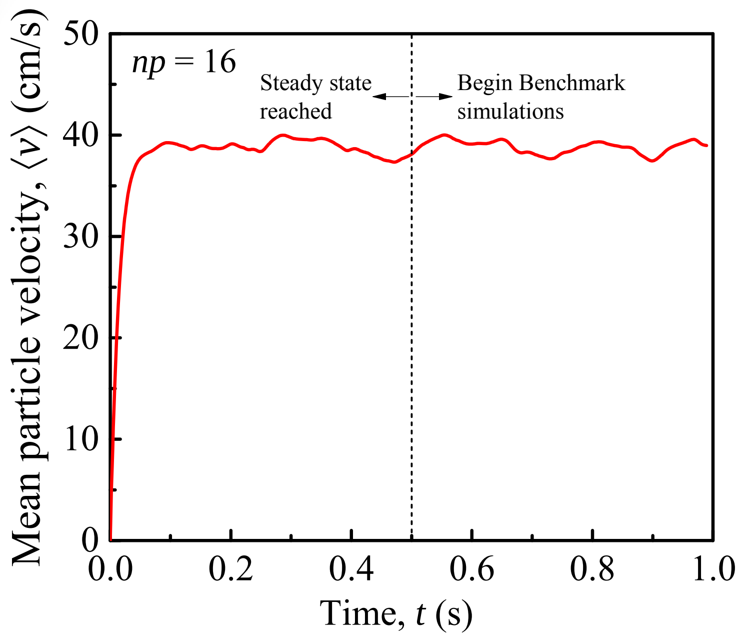


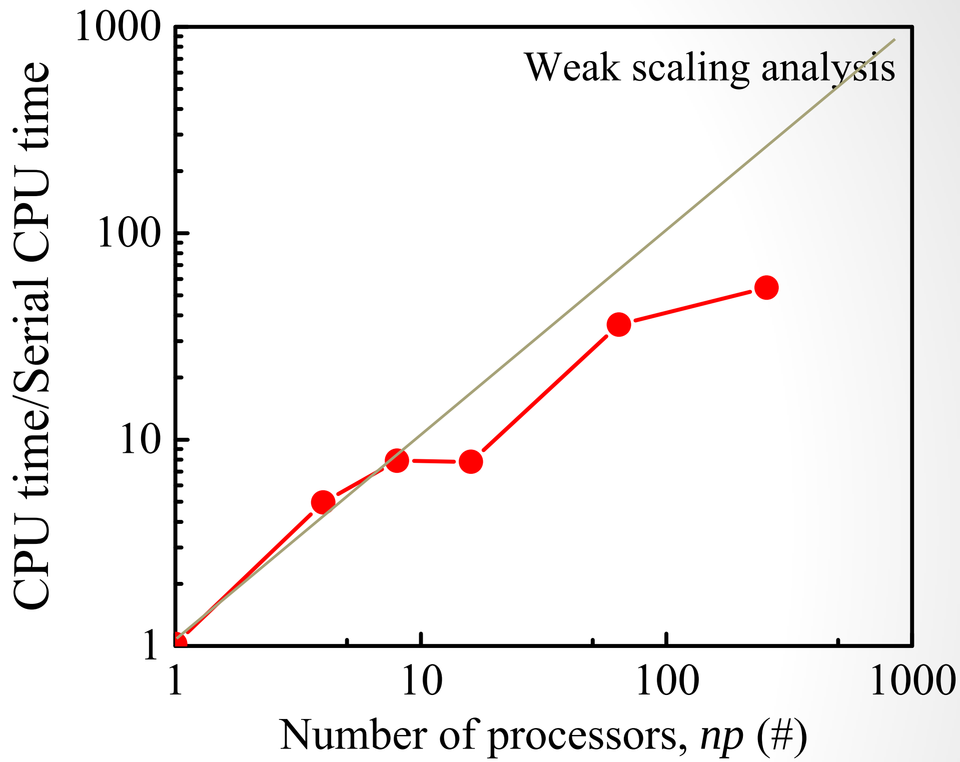


| Weak scaling analysis for riser | |
|---|---|
| Number of particles, $N_p$ | 250 ~ 256000 |
| Number of procs, np | 1 ~ 1024 |
| Number of particles per procs | 250 |
| Number of CFD cell per procs | 800 |
| Simulated time, $t_s$ (s) | 0.05 |
| Serial CPU time, $t_{cpu, serial}$, (h) | 0.05 |

Steady of riser flow reached after the average particle speed levels off, then benchmark cases start

# Performance results #3: Riser Flow

| Benchmark Statistics (Percentages) | |
|---|---|
| FP operations/cycle | 14 |
| FP vectorization | 2 |
| Level 2 cache miss ratio | 15 |
| Level 3 cache miss ratio | 0 |

| Loop Metrics (Percentages) | |
|---|---|
| Vectorized | 6.3 |
| Scalar | 53.0 |
| Outside | 40.7 |

| Profiling Summary (Top 5 functions) | |
|---|---|
| Function | Percentage of time |
| FUNCTIONS | 21.9 |
| COMP_MEAN_FIELDS0 | 15.0 |
| LEQSOL | 12.76 |
| DGTSV | 12.25 |
| DRAG_GS_DES0 | 4.54 |

Intel(R) Xeon(R) CPU X5660 @ 2.80GHz
24 Processors per node
24 GB RAM per node

# Benchmark #4: Square tumbler

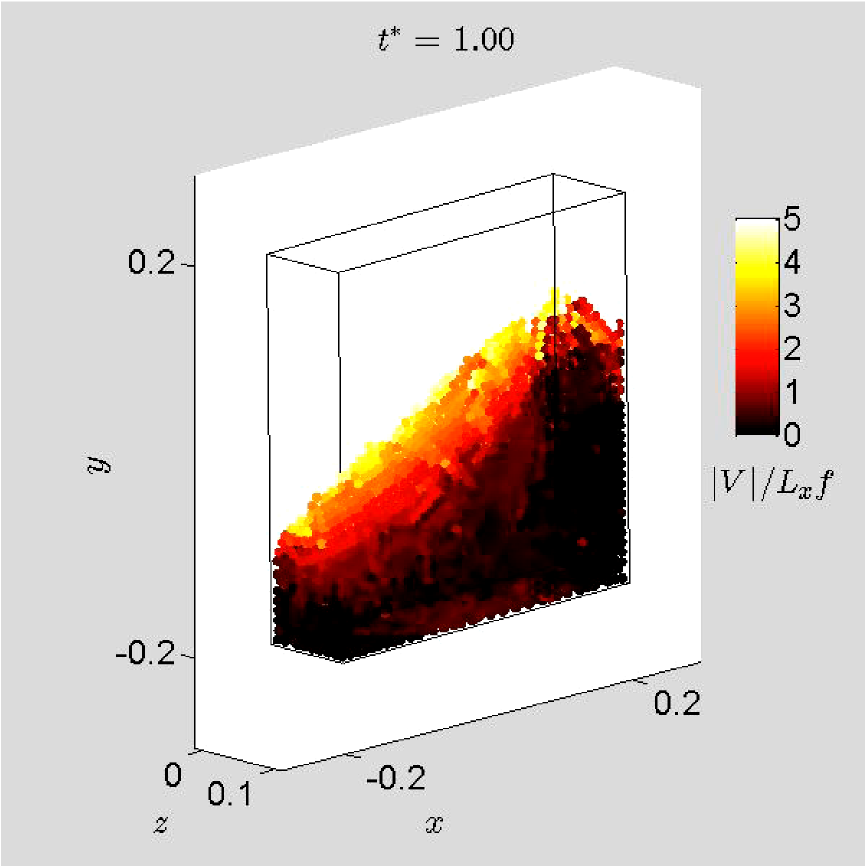


serial case (np = 1)

- Enclosed pipe of square cross-section
- Scaled in “1D”
  - *x*-scale = none
  - *y*-scale = none
  - *z*-scale = np
- Total concentration of 30% → particles packed in ~½ domain
- Avoid moving walls: gravity vector rotated through *xy*-plane
- Speed: $\omega = \pi$ (rad/s) = 30 (rpm)
- Variants: bi-disperse mixture
  - $d_{p1}/dp_2 = 2$
  - 50/50 mixture
  - Same total concentration as monodisperse case

# Performance results #4: Square tumbler

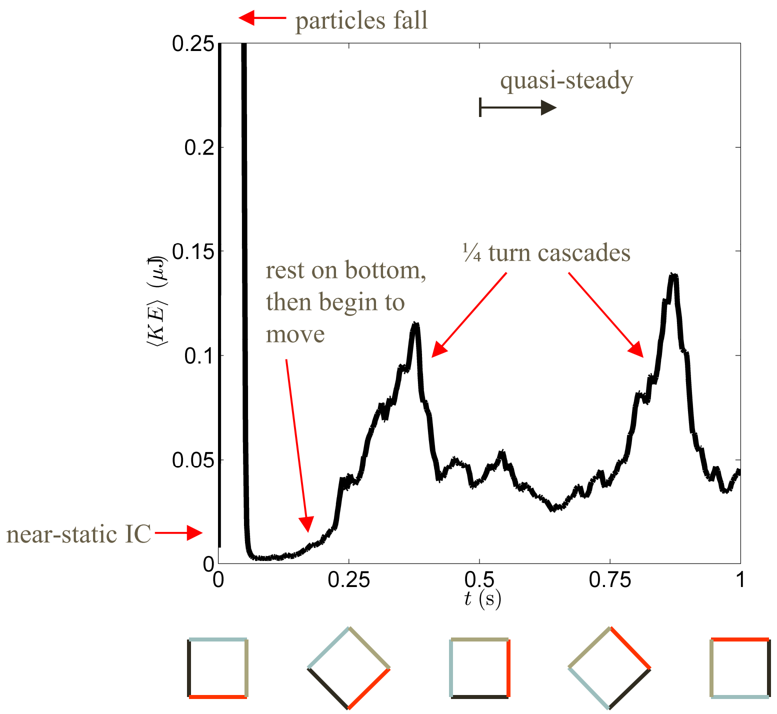


Initial non-convergence issues above np = 64 on restart

Monodispersed

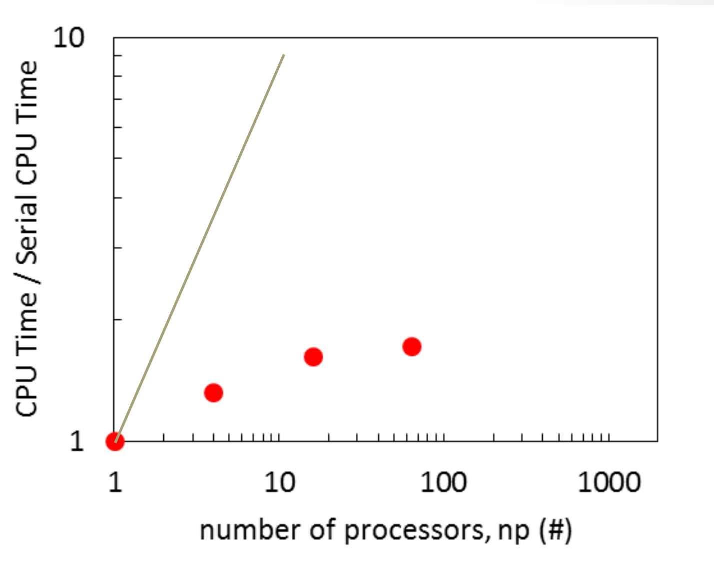


| Weak scaling analysis for tumbler case | |
|---|---|
| $N_p$/np (both) | 9165 |
| $N_x N_y N_z$ /np (CFD-DEM only) | 2000 ($20^2 \cdot 5$) |
| Serial CPU time (CFD-DEM) | 2.8 (h) |

# Performance results #4: Square tumbler

| Benchmark Statistics (Percentages) | |
|---|---|
| FP operations/cycle | 28 |
| FP vectorization | 7 |
| Level 2 cache miss ratio | 17 |
| Level 3 cache miss ratio | 2 |

| Loop Metrics (Percentages) | |
|---|---|
| Vectorized | 5.1 |
| Scalar | 57.8 |
| Outside | 37.1 |

**Profiling Summary (Top 5 functions)**

| Function | Percentage of time |
|---|---|
| CALC_FORCE_DEM | 31.3 |
| CALC_DRAG_DES | 20.5 |
| COMP_MEAN_FIELDS | 13.3 |
| NEIGHBOUR | 6.6 |
| PARTICLES_IN_CELL | 5.4 |

# Benchmark #5: Homogeneous Cooling System

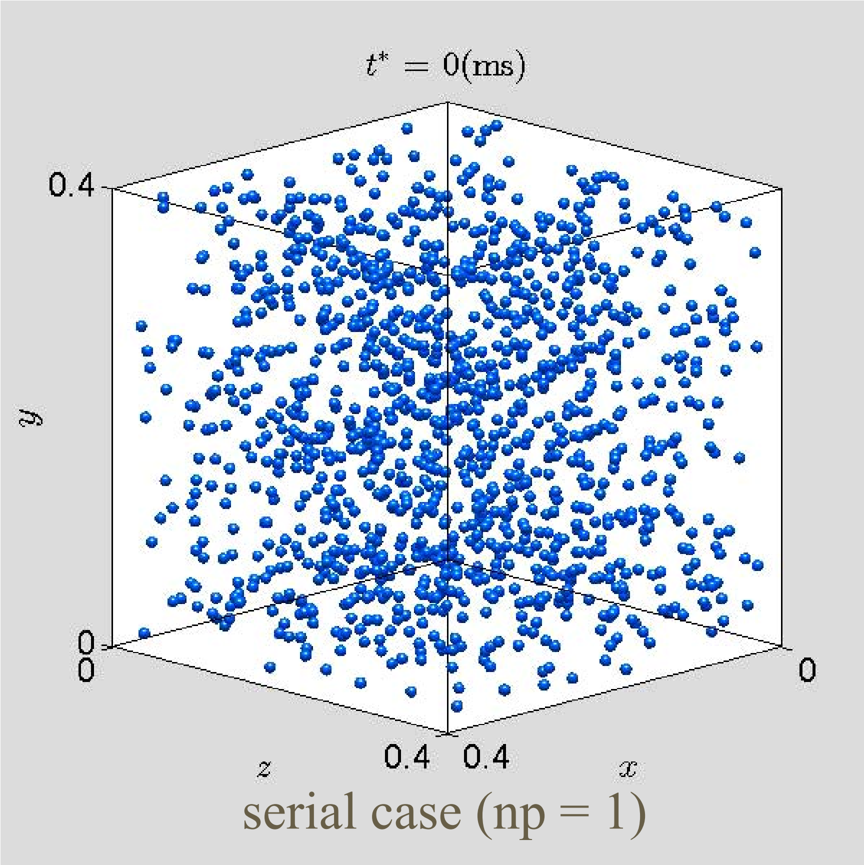

serial case (np = 1)

- Very common problem in physics community
- Traditionally used to study onset of instabilities – here we stay in the homogeneous regime ($\phi = 0.01$)
- Particles uniformly, randomly distributed
- Particle have normal, random velocity but zero net velocity in each direction (i.e. $T_0$)
- Fully periodic
- Granular temperature, $T$, decays in time through elastic collisions and interfacial drag

# Performance results #5: Homogeneous Cooling System

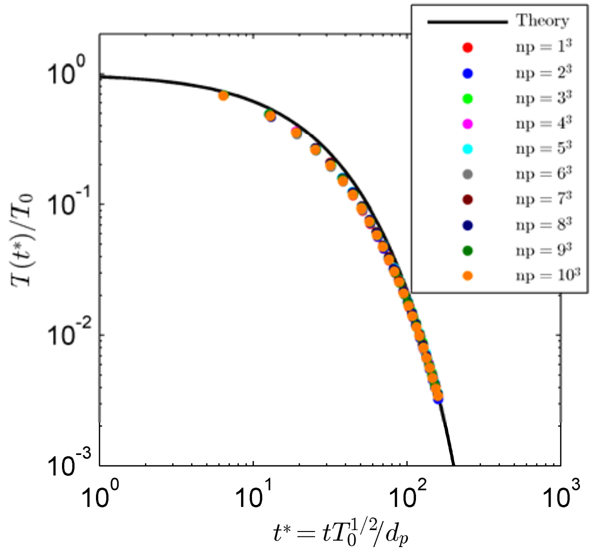


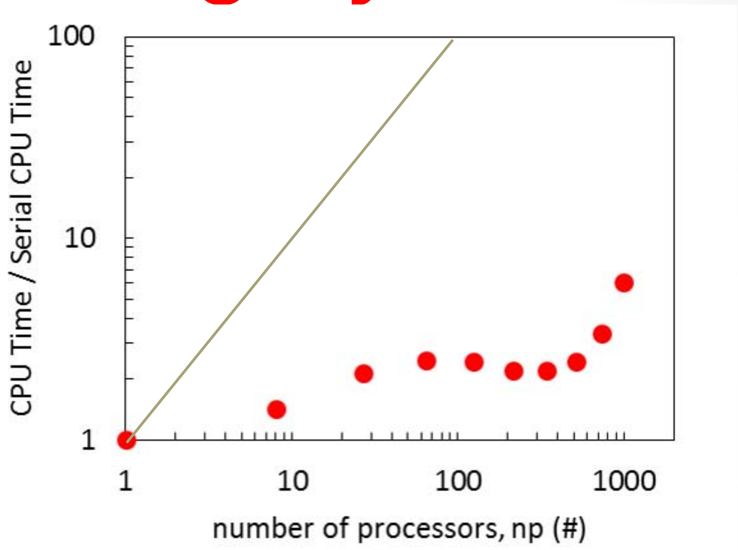


| Weak scaling analysis for HCS case | |
|---|---|
| $N_p$/np (both) | 1222 |
| $N_x N_y N_z$/np (CFD-DEM only) | 8000 ($20^3$) |
| Serial CPU time (CFD-DEM) | 812 (s) |

- Remains homogeneous even at large *L* due to $L_{serial}$-periodicity

# Scaling Summary

Strong Scaling of Settling

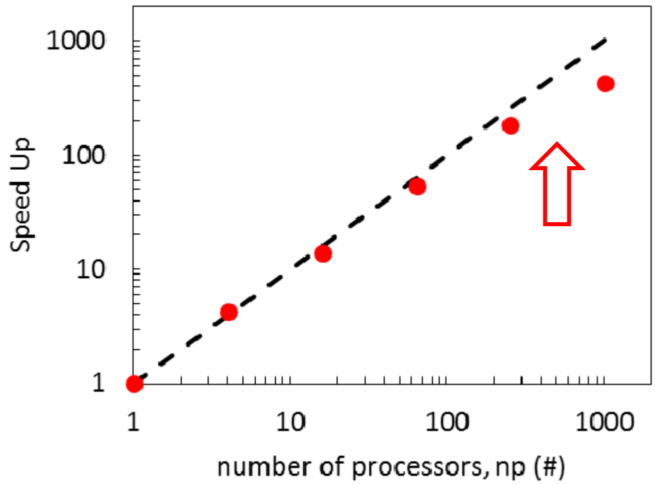


Current Capability

- $10^3$ $N_p/core$: Good speed for quick turn-around
- $10^4$ $N_p/core$: Slow, but likely acceptable
- $10^5$ $N_p/core$: Too slow for R&D turn-around

Weak Scaling of HCS

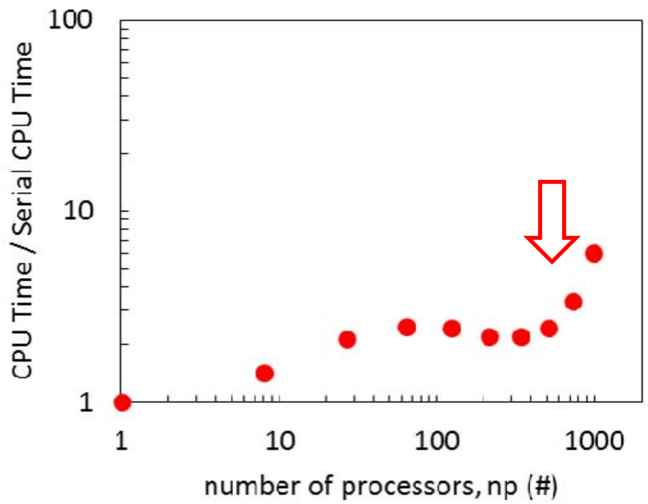


Goal: $N_p = 10^8$ for realistic simulations

- $10^3$ $N_p/core \cdot 10^5$ *cores*: Unrealistic industrial capability
- $10^4$ $N_p/core \cdot 10^4$ *cores*: Sweet spot, but need to improve

  serial speed and scalability for *cores* >$10^3$
- $10^5$ $N_p/core \cdot 10^3$ *core*: MFIX too slow

# Performance Summary

- To enable industrial flows
  - $10^4$ to $10^5$ particles per core
  - Performance of MFIX needs to improve
- Barriers to better performance
  - Low vectorization of inner loops
  - High number of cache misses
  - Poor weak scaling

# MFIX-DEM Improvements

- Current status
- General floating point optimization
  - Drag calculation
- Algorithmic changes to improve memory performance and aid vectorization
  - Sorting particles
    - Spatially
    - By type of particle
  - Memory alignment improvements

# Floating Point Optimization - Drag

Expressions with divisions and repeated operations in Koch Hill drag

```
F_0 = (1.0D0-w) * (1.0D0 + 3.0D0*dsqrt(phis/2.0D0) + &
        135.0D0/64.0D0*phis*LOG(phis) + 17.14D0*phis) / &
        (1.0D0 + 0.681D0*phis - 8.48D0*phis*phis + &
        8.16D0*phis**3) + w*10.0D0*phis/(1.0D0-phis)**3
```

- Optimize using
    - Pre-calculated variables for dsqrt and log
    - Use contiguous memory

```
phivals(1)=phis
phivals(2)=phis*phis
phivals(3)=1.d0/phis
phivals(4)=dsqrt(phis*0.5d0)
phivals(5)=LOG(phis)
one_m_phis_3_inv = (1.d0-phis)**(-3.d0)
```

# Spatial reordering of particles

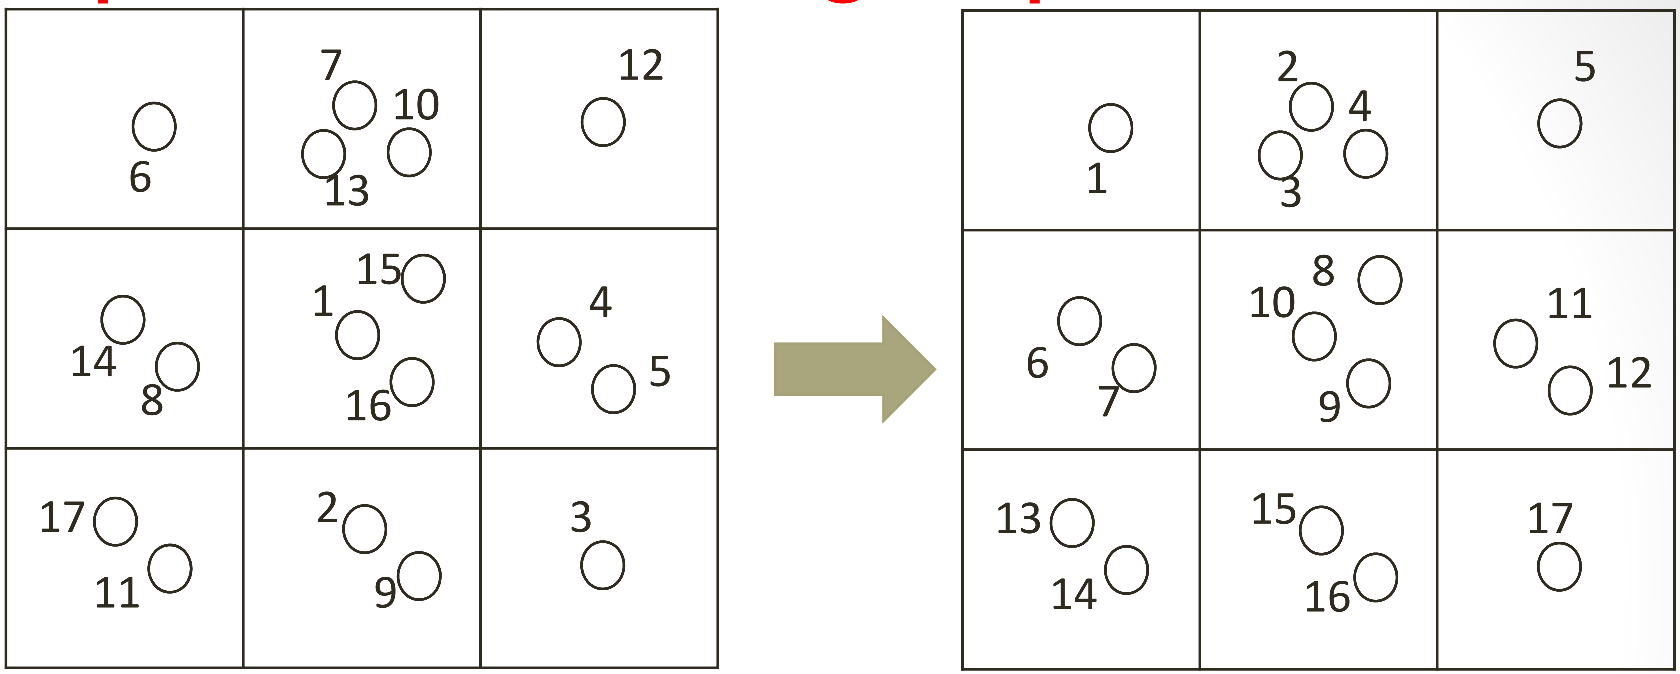


- This will ensure spatial locality in memory while
  - Finding neighbors
  - Calculating inter particle forces
  - Drag calculations
  - Extrapolating mean fields.
- This reordering need not be done every des time step
  - Done only once at the beginning of des time march

# Rearrangement based on state

Particle array

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| n | n | ghst | X | n | n | ghst | X | n | n |

Sorted Particle array

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| n | n | n | n | n | n | ghst | X | ghst | X |

Np_normal

<u>General particle loop structure</u>

Do i=1,MAX_PIP

If not a normal particle cycle
If …
If ..

Do operations
A(i) = B(i) + C(i)

Removal of this condition for most particles that are normal can aid in vectorization

# Sorting neighbor arrays

Neighbor information
(human readable)

particles

| 1 | 2 | 3 | 4 | 5 |
|---|---|---|---|---|
| 2 | 3 | 4 | 5 | - |
| 3 | 5 | 5 | | |
| 4 | | | | |

↓ neighbors

MFiX data structure

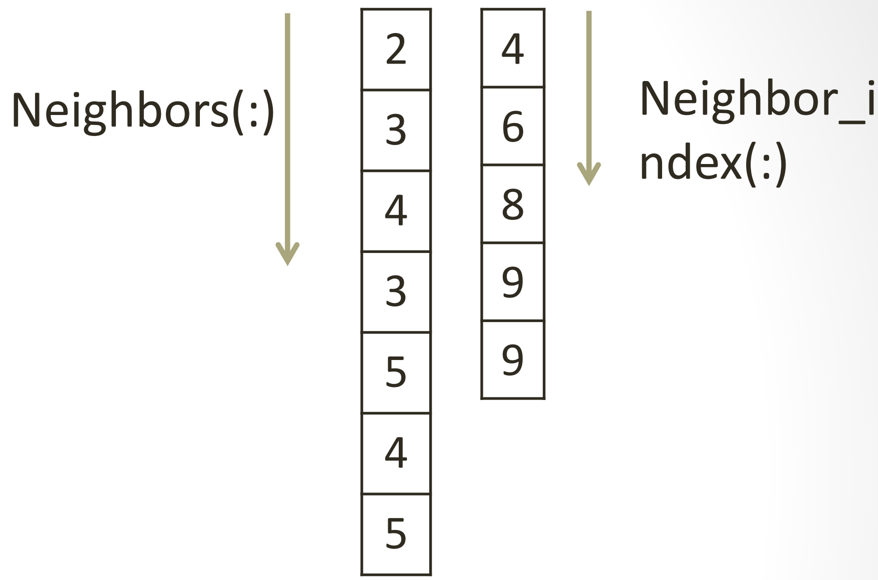


- When sorting arrays based on state
  - Particles are exchanged (ie say part. 1 becomes part. 4)
  - When sorting neighbor arrays
    - Need to populate new ids
    - Need to adjust neighbor_index
    - Need to offset sets of indices

## Square tumbler case ( ~ 9000 particles)

### Original code

| | |
|---|---|
| comp_mean_fields0 | 83.337 |
| __libm_pow_e7 | 54.358 |
| drag_gs_des0 | 45.799 |
| calc_force_dem | 39.25 |
| drag_bvk | 38.008 |
| cfnewvalues | 34.01 |
| drag_weightfactor | 27.209 |
| des_drag_gp | 26.959 |
| desgrid_neigh_build | 25.129 |
| drag_interpolation | 22.889 |
| calc_dem_force_with_wall_stl | 18.63 |
| particles_in_cell | 18.57 |
| dgtsv | 8.53 |
| fluid_at | 7.68 |
| is_normal | 5.11 |
| is_nonexistent | 4.62 |
| leq_iksweep | 4.61 |
| set_interpolation_stencil | 4.41 |
| is_on_mype_wobnd | 4.4 |
| funijk | 4.28 |
| check_cell_movement | 3.59 |
| cfrelvel | 3.4 |

Total wall clock time = 531 sec

### Optimized code

| | |
|---|---|
| comp_mean_fields0 | 76.708 |
| __libm_pow_e7 | 52.549 |
| calc_force_dem | 36.109 |
| drag_gs_des0 | 33.509 |
| cfnewvalues | 33.059 |
| des_drag_gp | 32.13 |
| drag_interpolation | 23.669 |
| drag_bvk | 23.469 |
| drag_weightfactor | 22.28 |
| desgrid_neigh_build | 20.199 |
| particles_in_cell | 18.43 |
| calc_dem_force_with_wall_stl | 11.5 |
| fluid_at | 8.62 |
| dgtsv | 8.19 |
| leq_iksweep | 4.82 |
| is_normal | 4.71 |
| set_interpolation_stencil | 4.36 |
| is_on_mype_wobnd | 4.28 |
| is_normal | 4.24 |
| funijk | 3.91 |
| check_cell_movement | 3.83 |
| cfrelvel | 3.26 |

Total wall clock time = 486 sec

# Next Steps

- Removing if conditions in loops for vectorization
- Sorting particles by state
  - Sort non-existent particles also and move it to the end
- Looping over cells instead of particles and computing inter-particle forces

# Other Contributions to MFIX-DEM

- Initial bug fixes for the Fortran dependency generator
- A new Fortran dependency generator
- autoconf macro to query the alignment size from the compiler.

# Questions ?