




# A Parallel-efficient GPU Package for Multiphase Flow in Realistic Nano-pore Networks

Yidong Xia, PhD<sup>1</sup> | Ansel Blumers<sup>1,2</sup> | Zhen Li, PhD<sup>3</sup> | Lixiang Luo, PhD<sup>4</sup> | Jan Goral<sup>5</sup> | Matthew Andrew, PhD<sup>6</sup> | Joshua Kane, PhD<sup>2</sup> | Yu-Hang Tang, PhD<sup>7</sup>

<sup>1</sup> Materials Science & Engineering Department, Idaho National Laboratory, Idaho Falls, ID; <sup>2</sup> Department of Physics, Brown University, Providence, RI; <sup>3</sup> Division of Applied Mathematics, Brown University, Providence, RI; <sup>4</sup> IBM Center of Excellence at ORNL, Oak Ridge, TN;

<sup>5</sup> Department of Chemical Engineering, University of Utah, Salt Lake City, UT; <sup>6</sup> Carl Zeiss X-Ray Microscopy, Pleasanton, CA; <sup>7</sup> Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, CA




## Objectives and technical challenges

- Challenges of fluid flow simulations in oil and gas shale rocks: multiscale pore systems ranging from a few nanometers to a few micrometers; complex fluid-fluid/solid interactions; **significant computational costs**.
- Computational model**: robust and viable **mesoscale fluid models** are required for multiphase flow in nano- to micro-scale pore systems in shale with pore geometries of nanometer-scale resolutions.
- Solution**: **efficient use of top-tier GPGPUs to accelerate** simulations of heterogeneous porous systems with  $10^8 \sim 10^9$  particles and  $10^7 \sim 10^8$  timesteps for predicting material properties such as permeability.



## Fluid flow in oil/gas shale rocks

- Most of the shale oil and gas resources contained in kerogen-hosted pores.
- Low porosity & permeability; multiscale ( $10^{-9} \sim 10^{-6}$  m) in pore sizes.
- Strong thermal fluctuation and diffusion at fluid-fluid and fluid-solid interfaces.
- Continuum-based fluid models are no longer valid in nano- to micro-pores.




A **mesoscale fluid model** is required for modeling **multiphase flow** in shale – **many-body dissipative particle dynamics (mDPD)**

(Reference: Y. Xia et al. (2017), *Physics of Fluids*, 29(5), 056601)




Fluid-fluid and fluid-solid interfaces in a **100-nanometer**-wide pore are stable



Thermodynamic fluctuation and diffusion in a **2-nanometer**-wide pore are strong

From 3D digital imaging to simulations of shale oil/gas recovery



## Mathematical model: many-body dissipative particle dynamics (mDPD)

$$\mathbf{f}_i^{\text{int}} = \sum_{j \neq i} \mathbf{F}_{ij} = \sum_{j \neq i} \left( \mathbf{F}_{ij}^{\text{C}} + \mathbf{F}_{ij}^{\text{D}} + \mathbf{F}_{ij}^{\text{R}} \right) \quad \mathbf{F}_{ij}^{\text{C}} = \left\{ A_{ij} w_c(r_{ij}) + B_{ij} [\bar{\rho}_i + \bar{\rho}_j] w_d(r_{ij}) \right\} \hat{\mathbf{r}}_{ij}$$

Conservative Force      Dissipative Force      Random Force

- A coarse-grained molecular dynamics model. Motion of particles governed by the Newton's second law.
- The conservative force which depends on **local particle densities** consists of a **long-range attraction** and **short-range repulsion**. (Reference: Warren, P. B. (2003). *Physics Review E*, 68(6), 066702)

## Solution: GPU-accelerated particle simulator package *USER*MESO 2.5

- USER*MESO 2.5 is designed for **flow simulations in 3D arbitrarily complex nano-/micro-porous media**:

- ✓ Implementation of mDPD model for multiphase flow.
- ✓ No-slip boundary condition for fluid flow in arbitrarily complex geometries.
- ✓ Voronoi-inspired particle partitioning for load balancing (under R&D).

- USER*MESO 2.5 is a **hybrid GPU/CPU code designed to align with LAMMPS code structure**:


- ✓ Programmed using CUDA C/C++ with MPI and OpenMP.
- ✓ As customizable as LAMMPS.

- USER*MESO 2.5 is an **extension to the original *user*MESO** package<sup>[1,2]</sup> which include key features:

- ✓ An atomics-free warp-synchronous neighbor list construction algorithm.
- ✓ A 2-level particle reordering scheme, which aligns with the cell list lattice boundaries for generating strictly monotonic neighbor list.
- ✓ Overlapping pairwise force evaluation with halo exchange using CUDA streams for hiding communication and kernel launch latency.
- ✓ Radix sort with GPU stream support.



References: [1] Y-H. Tang et al. (2014), *Computer Physics Communications*, 185 (11), 2809-2822  
[2] A. Blumers et al. (2017), *Computer Physics Communications*, 217, 171-179

## Benchmarks




### Benchmark-problem:

- USER*MESO 2.5 is benchmarked for a body force-driven flow in a quasi-2D domain with uniform pore distributions in the domain and periodic domain boundaries in all the three directions.
- The pore network is represented by the space (**narrowest gap = 4 nm**) between the circles (**14 nm in diameter**) which symbolize the wall of the porous channel. A no-slip boundary condition is prescribed at the wall surfaces, which prevents fluid particles from penetrating the walls.



**Benchmark #1:** Benchmark tests of pore-network flow show **nearly perfect strong- and weak-scalings** up to 512 k20x GPUs (one GPU per node) on TITAN, ORNL. For strong-scaling, the system size is kept at 2.1 million particles. For weak-scaling, the system size is fixed at 1 million particles per node. (MPS is short for Million-Particles-Steps.)

**Benchmark #2:** Our code benefits greatly from more powerful and newer generations of top-tier GPUs, as demonstrated in a **single-GPU performance test**. The high-speed Host-to-Device interconnect, NVLink, can significantly boost performance, as the case on ORNL SummitDev and Summit. We also compared with a CPU node featuring two Intel Xeon E5-2695.



**Benchmark #3:** The **speedup brought by NVLink is evident** from inspecting the **breakdown of walltime** of a single-GPU simulation **on GPU related tasks**. On DGX-1 (V100) with no NVLink, transfers take up 53% of GPU related walltime. On SummitDev with Host-to-Device NVLink1, transfers only take 24%. However, the blazing-fast V100 compensates for the slow transfer speed. When both V100 and NVLink2 are available, walltime on GPU related tasks is further shortened.

**NVLink 1** with P100+Power8: 160 Gbyte/s total data rate (out+in)  
**NVLink 2** with V100+Power9: 300 Gbyte/s total data rate (out+in)

### Computing resources

**Tesla k20x:** TITAN, ORNL, AMD Opteron 6274  
**V100:** Galena (DGX-1), INL, Intel Xeon E5-2698 v4  
**V100 + NVLink2:** IBM internal test node, IBM Power9

**Titan Xp:** Desktop workstation, Intel i7-8700K  
**P100 + NVLink1:** SummitDev, ORNL, IBM Power8  
**Intel Xeon E5-2695:** Falcon, INL, CPU node

### Author contribution

**Yidong Xia** - principal investigator  
**Zhen Li** - reference CPU source codes  
**Jan Goral & Matthew Andrew** - 3D digital rock imaging

**Ansel Blumers & Yu-Hang Tang** - *USER*MESO 2.5 source code  
**Lixiang Luo** - technical support of Titan & SummitDev HPC  
**Joshua Kane** - technical support of imaging data conversion

### Acknowledgement

This work is supported through the Idaho National Laboratory (INL) Laboratory Directed Research & Development (LDRD) Program under Department of Energy Idaho Operations Office Contract DE-AC07-05ID14517. This work used resources of the Oak Ridge Leadership Computing Facility at Oak Ridge National Laboratory, which is supported by the Office of Science of the Department of Energy under Contract DE-AC05-00OR22725.