

# Nanoporous Flow Simulations on the Summit Supercomputer

Yidong Xia\*  
Idaho National Laboratory  
Idaho Falls, Idaho, USA

Lixiang Luo  
IBM Center of Excellence at ORNL  
Oak Ridge, Tennessee, USA

Ansel Blumers  
Brown University  
Providence, Rhode Island, USA

Joshua Kane  
Idaho National Laboratory  
Idaho Falls, Idaho, USA

Jan Goral  
University of Utah  
Salt Lake City, Utah, USA

Yu-Hang Tang  
Lawrence Berkeley National Lab  
Berkeley, California, USA

Zhen Li  
Clemson University  
Clemson, South Carolina, USA

Hai Huang  
Idaho National Laboratory  
Idaho Falls, Idaho, USA

Milind Deo  
University of Utah  
Salt Lake City, Utah, USA

## ABSTRACT

Fluid flow behaviors in nanoporous materials can be distinct from those following the continuum physics at higher scales. Numerical simulations can be a complement to laboratory experiments. This work presents a dissipative particle dynamics (DPD) package for GPU-accelerated mesoscale flow simulations in nanoporous materials. In an ideal benchmark that minimizes load imbalance, the package delivered nearly perfect strong- and weak-scaling (with up to 4 billion DPD particles) on up to 1,536 V100 GPUs on Oak Ridge National Laboratory's Summit supercomputer. More remarkably, in a benchmark to measure its usefulness with realistic nanopores in SBA-15 silica, the package exhibited more than 20x speedup over its LAMMPS-based CPU counterpart with the same number nodes (e.g., 384 V100 GPUs vs. 2,688 POWER9 cores). It is also worth highlighting that the NVLink2 Host-to-Device interconnects kept the cost of CPU-GPU memory copy as low as only 10% of GPU activity time per rank, which is 4 times less than their PCIe counterparts.

## KEYWORDS

Summit supercomputer, GPU computing, nanoporous materials, dissipative particle dynamics

### ACM Reference Format:

Yidong Xia, Lixiang Luo, Ansel Blumers, Joshua Kane, Jan Goral, Yu-Hang Tang, Zhen Li, Hai Huang, and Milind Deo. 2019. Nanoporous Flow Simulations on the Summit Supercomputer. In *SC19, Nov 17–22, 2019, Denver, CO*. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 INTRODUCTION

In geosciences, the permeability of natural nanoporous materials such as kerogen-rich tight shale is difficult to characterize, because of their highly non-uniform distribution of pores, because of their

\*Corresponding author: Yidong Xia (yidong.xia@inl.gov). INL/EXP-19-54808

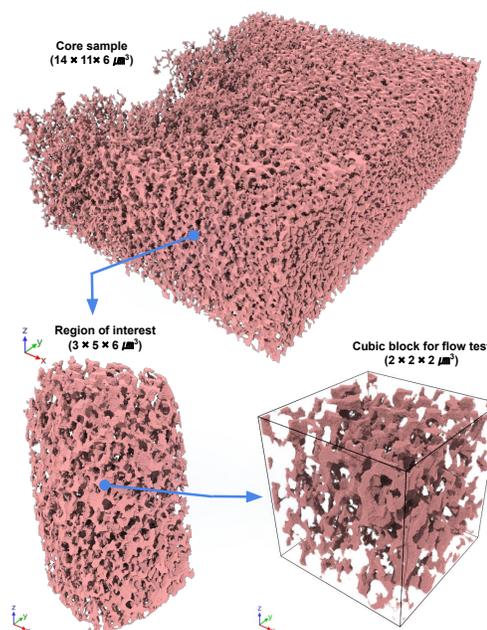
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

Xia '19, November 17–22, 2019, Denver, CO

© 2019 Association for Computing Machinery.

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

heterogeneity of material properties that can impact the fluid inside physically and chemically, and because of the limited capacity of laboratory instrumentation. It is hard to agree on a common approach to study flow and transport phenomena in natural nanoporous materials, as too many coupled variables should be considered.



**Figure 1: Pore skeleton of a SBA-15 silica sample obtained by FIB-SEM (focused ion beam scanning electron microscope) with voxel resolution of  $7.7^3 \text{ nm}^3$  (top). A region of interested (ROI) extracted for fluid-permeability analysis (bottom left). A cubic block extracted from the ROI for DPD flow simulations (bottom right).**

Synthetic mesoporous silica are often used for studying fluid flow and transport at the nanoscale, as they are architected with controlled material properties, pore sizes and distributions. This work uses SBA-15 silica as the nanoporous media to measure water-permeability dependence. Selection of a proper computational domain is illustrated in Figure 1. Over 100 million DPD particles are

required to run for an equivalent number of timesteps for measurement. A simulation code that is able to take full advantage of the latest supercomputers is critical to the success of the study.

## 2 APPROACH

This work employs userMESO 2.5 [6] – the latest release to the userMESO framework developed as a GPU extension package to LAMMPS [4] for DPD simulations [1, 5]. New features mainly include implementation of the many-body DPD model [3, 7] and an impenetrable wall boundary model for arbitrary geometries [2]. Notable innovative features of the package are discussed in [5]. In userMESO, all computations and host-device communications are handled by the extension package, while I/O related tasks such as inter-rank communications are attended by LAMMPS. All releases of userMESO (i.e. version 1.0, 2.0 and 2.5) are able to achieve excellent strong- and weak-scaling on ORNL’s Titan supercomputer as well as competitive speedup over its CPU counterpart [1, 5, 6].

## 3 RESULTS

Code readiness on Summit has been made for userMESO 2.5. Like most people did to show scalability using a benchmark with minimal load imbalance, we used a manufactured, strictly uniform pore system, with which userMESO 2.5 delivered nearly perfect strong- and weak-scaling (with up to 4 billion DPD particles) on up to 1,536 V100 GPUs on Summit (6 GPUs and 42 usable Power9 CPU core per node); see Figure 2. To measure how userMESO 2.5 performs

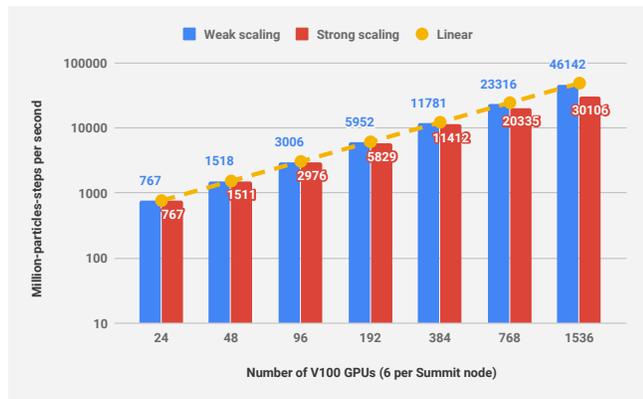


Figure 2: Strong- and weak-scaling performance on Summit.

for non-uniform pores, strong-scaling was conducted with a cubic block SBA-15 silica geometry as shown in Figure 1 (about 70 million DPD particles), exhibiting more than 20x speedup over its CPU counterpart with the same number nodes (e.g., 384 GPUs vs. 2,688 CPU cores); see Figure 3. A maximum load imbalance factor of 2.9 is observed. This indicates that GPU computing, with fewer partitions, effectively reduces load imbalance than its CPU counterpart. Moreover, a breakdown of rank-averaged GPU activity time is displayed in Figure 4, showing a low ratio of time cost of host-device memory copy thanks to the NVLink2 Host-to-Device interconnects. Note that it is 4 times less than its PCIe counterparts on DGX-1 and half of the cost than NVLink1 on SummitDev.

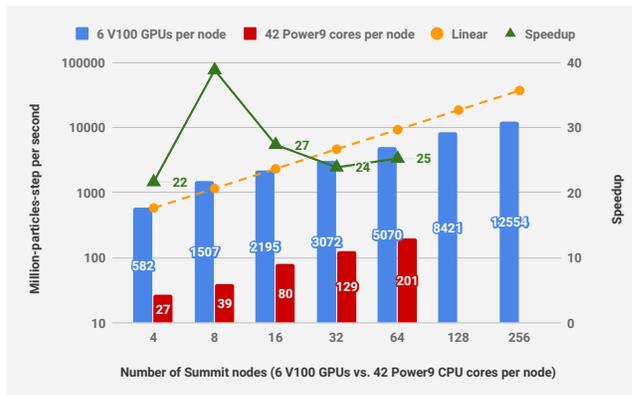


Figure 3: GPUs vs. CPUs with the same nodes on Summit.

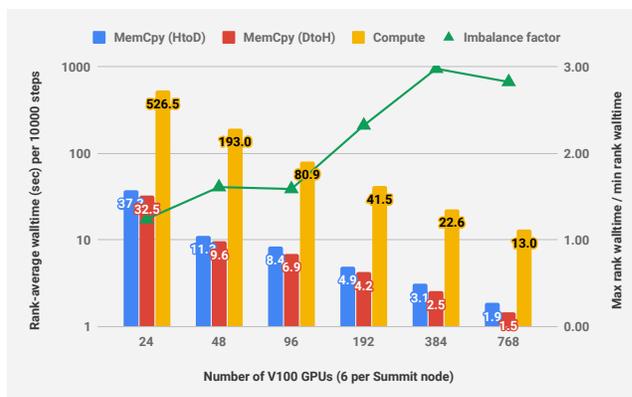


Figure 4: Breakdown of GPU activity time on Summit.

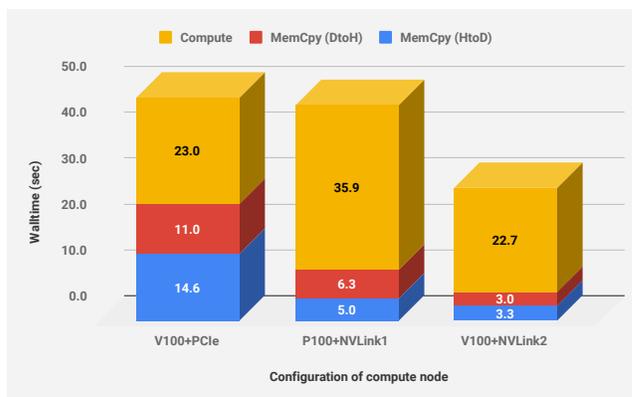


Figure 5: Single-node breakdown of GPU activity time.

Above all, our study has demonstrated the success of userMESO 2.5 and DPD methods on Summit. A broader implication is that other particle-based methods such as smoothed particle hydrodynamics (SPH) and discrete element method (DEM) may also benefit tremendously from the dense GPU configuration and the overall GPU-computing power provided by Summit.

## ACKNOWLEDGMENTS

The software development in this work is supported through the Idaho National Laboratory (INL) Laboratory Directed Research & Development (LDRD) Program under the U.S. Department of Energy Idaho Operations Office Contract DE-AC07-05ID14517.

The weak- and strong-scaling benchmarks were primarily performed at Oak Ridge Leadership Computing Facility (OLCF) through the OLCF Director's Discretion Program under project GEO133, which is supported by the Office of Science of the U.S. Department of Energy under Contract DE-AC05-00OR22725.

The FIB-SEM and image processing of synthetic nano-porous silica was supported as part of the EFRC-MUSE, an Energy Frontier Research Center funded by the U.S. Department of Energy, Office of Science, Basic Energy Sciences under Award No. DE-SC0019285.

## REFERENCES

- [1] A. L. Blumers, Y-H Tang, Z. Li, X. Li, and G. E. Karniadakis. 2017. GPU-accelerated red blood cells simulations with transport dissipative particle dynamics. *Computer Physics Communications* 217 (2017), 171–179.
- [2] Z. Li, X. Bian, Y-H Tang, and G.E. Karniadakis. 2018. A dissipative particle dynamics method for arbitrarily complex geometries. *J. Comput. Phys.* 355 (2018), 534–547.
- [3] Z. Li, G-H Hu, Z-L Wang, Y-B Ma, and Z-W Zhou. 2013. Three dimensional flow structures in a moving droplet on substrate: A dissipative particle dynamics study. *Physics of Fluids* 25, 7 (2013), 072103.
- [4] S. Plimpton. 1995. Fast parallel algorithms for short-range molecular dynamics. *J. Comput. Phys.* 117, 1 (1995), 1–19.
- [5] Y. Tang and G. Karniadakis. 2014. Accelerating dissipative particle dynamics simulations on GPUs: Algorithms, numerics and applications. *Computer Physics Communications* 185, 11 (2014), 2809–2822.
- [6] Y. Xia, A. Blumers, Z. Li, L. Luo, Y-H Tang, J. Kane, H. Huang, M. Andrew, M. Deo, and J. Goral. 2019. A GPU-accelerated package for simulation of flow in nanoporous source rocks with many-body dissipative particle dynamics. *arXiv preprint arXiv:1903.10134* 1, 1 (2019).
- [7] Y. Xia, J. Goral, H. Huang, I. Miskovic, P. Meakin, and M. Deo. 2017. Many-body dissipative particle dynamics modeling of fluid flow in fine-grained nanoporous shales. *Physics of Fluids* 29, 5 (2017), 056601.