



Exascale Computational Fluid Dynamics in Heterogeneous Systems

David Trebotich

Applied Mathematics and Computational Research
Division,
Lawrence Berkeley National Laboratory,
1 Cyclotron Road, MS 50A-3123,
Berkeley, CA 94720
e-mail: DPTrebotich@lbl.gov

Exascale computing has extended the reach of resolved flow simulations in complex, heterogeneous systems far beyond conventional computational fluid dynamics capabilities. As a result, unprecedented pore and microscale resolution have been achieved in domains that have been traditionally modeled by, and limited to, continuum, effective medium approaches. By making use of computational resources on the new exascale supercomputer, Frontier, at the Oak Ridge Leadership Computing Facility, we performed flow simulations that have pushed the limits of domain-to-resolution ratios by several orders of magnitude for heterogeneous media. Our approach is an incompressible, Navier–Stokes CFD solver based on adaptive, embedded boundary (EB) methods supported by the Chombo software framework for applied partial differential equations (PDEs). The computational workhorse in the CFD application code is an elliptic solver framework in Chombo for pressure-Poisson and viscous, Helmholtz terms that leverages a PETSc-hypre software interface tuned for accelerator-based platforms. We demonstrate scalability of the approach by replicating a unit cylinder packed with microspheres to achieve over 400×10^9 degrees-of-freedom simulated. These simulations model domain lengths of over 20 meters with channel volumes of over 400 cm^3 and containing millions of packed spheres with 20 micron grid resolution, challenging current understanding of what it means to be a representative elementary volume (REV) of the continuum scale in heterogeneous media. We also simulate a range of Reynolds numbers to demonstrate wide applicability and robustness of the approach.

[DOI: 10.1115/1.4064534]

Introduction

Computational fluid dynamics is central to predictive modeling of multiscale, multiphysics problems in heterogeneous systems. In the geologic subsurface, for example, simulation of flow in underground reservoirs can help indicate flow paths to access subsurface energy resources such as enhanced geothermal systems and oil recovery while also identifying locations to sequester other sources like CO_2 . Similarly, in energy storage systems, computational models of flow and related transport processes in these materials can help predict long-term reliability of current lithium ion battery chemistry as well as the discovery of alternative chemistries to increase energy densities, which is critical to long-term electric vehicle and grid storage industries. In addition, many manufacturing processes rely on CFD to help optimize and design devices, shortening product cycles, and reducing resource consumption in energy-intensive manufacturing processes (e.g., desalination, paper drying/manufacturing).

Common to successful modeling of these problems is the ability to predict multiscale flow coupled with multiphysics and chemistry behavior on a macroscopic scale—such as a geologic reservoir in the subsurface or an engineering device for energy storage or manufacturing. However, macromodels do not represent and cannot

capture microscale emergent processes, which can evolve into large-scale dynamics over time. Examples of these types of emergent processes include dissolution of subsurface asperities on the micron scale causing local fracture collapse when CO_2 is injected into an underground reservoir (km scale); or, in the case of energy storage devices, submicron dendrite growth of lithium plating into and through the ceramic material causing short circuiting of the cell during device operation life cycles (cm scale). The importance of resolving CFD at engineering scales is also seen in flows in permeable membranes. For desalination devices, coupled turbulent flow, transport, and reaction processes lead to membrane fouling, seriously limiting the performance of reverse osmosis, nanofiltration, ultrafiltration, microfiltration, and membrane distillation systems. Predictive models linking meso-scale phenomena such as concentration polarization that occur on and within membranes to system-level desalination performance are challenged by turbulent flow regimes in these confined devices, yet are needed to optimize brine treatment for beneficial reuse. As another engineering example, paper manufacturing is water-intensive and can be treated as a bulk flow medium; however, the coupled hydromechanical drying process in the paper press is multiphase and porous, while trying to optimize the balance between dewetting and rewetting of the pulp. All of these examples require high resolution simulation capabilities to capture emergent processes at microscopic scales that can affect macroscopic behavior or performance at the continuum or device scale.

Pore (micro) scale modeling has been applied successfully to heterogeneous media, both geologic and engineered, much in the context of subsurface reacting flow and transport applications where fluid–solid boundaries are resolved to accurately represent the

Contributed by the Fluids Engineering Division of ASME for publication in the JOURNAL OF FLUIDS ENGINEERING. Manuscript received August 15, 2023; final manuscript received January 18, 2024; published online February 9, 2024. Assoc. Editor: Marianne Francois.

The United States Government retains, and by accepting the article for publication, the publisher acknowledges that the United States Government retains, a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for United States Government purposes.

reactive surface areas that drive coupled processes [1]. Because it is computationally intensive to resolve small characteristic length scales in heterogeneous pore space, a goal of pore scale modeling has been to develop effective parameter upscaling (e.g., permeability, tortuosity), i.e., to model domains with pore scale resolution that are consistent with what it means to be an REV of a macroscopic, continuum scale model [2]. For example, an REV for subsurface materials is generally of the order of a cubic centimeter; with 1 micron resolution, CFD simulations in such systems has required supercomputing capability and resources (e.g., [3,4]). For low permeability materials, an REV can be even smaller, on the order of a cubic millimeter, as nanoscale resolution is required to resolve flow paths in tight pore space (e.g., shale [5]). In both examples, three to four orders of magnitude in domain-to-resolution ratio have pushed HPC CFD capabilities to their limits on CPU-based machines. However, with the advent of exascale computing, more capability has become available on accelerator-based machines to perform CFD simulations for much higher domain-to-resolution ratios than ever before using codes that have undergone the requisite performance portability and engineering (e.g., [6]).

Mathematical Model

For our CFD model, we solve the incompressible Navier–Stokes equations

$$\mathbf{u}_t + (\mathbf{u} \cdot \nabla)\mathbf{u} + \nabla p = \nu \Delta \mathbf{u} \quad (1)$$

where $\mathbf{u} = \mathbf{u}(\mathbf{x}, t)$ is the fluid velocity, $\nabla p, p = p(\mathbf{x}, t)$ is the pressure gradient, and $\nabla \cdot \mathbf{u} = 0$ for incompressibility. The boundary conditions reflect pressure-driven flow through a channel, in the x -direction: inlet, $\mathbf{u} = (\frac{3}{2}\bar{u}(1 - y^2/a^2), 0)$, $\frac{\partial p}{\partial x} = 0$, outlet, $\frac{\partial \mathbf{u}}{\partial x} = 0$, $p = 0$, and solid boundaries, $\mathbf{u} = 0$, $\frac{\partial p}{\partial y} = 0$. Given initial conditions $\mathbf{u}^0 = \mathbf{u}(\mathbf{x}, 0)$, $p^0 = p(\mathbf{x}, 0)$, this system of equations constitutes an initial boundary value problem (IBVP) that can be solved by a number of different numerical approaches. For complex systems such as those found in engineering or natural media, we must be able to solve the equations of incompressible fluid dynamics in arbitrarily complex geometries with sufficient resolution to resolve the dominant viscous boundary layers in large surface area-to-volume ratios represented in heterogeneous media (e.g., [7–9]). To do so with numerical efficiency, it is important to choose a computational approach that supports consistency between geometry generation and discretization. We use the Chombo-based adaptive, embedded boundary method described in Ref. [5].

Algorithmic Approach. Chombo is a set of C++ libraries and FORTRAN macros for solving PDEs in complex geometries using adaptive, finite volume methods. It provides a flexible toolset of algorithmic and software components that can be assembled to simulate a broad range of complex multicomponent physical systems in which PDEs play a central role. The Chombo approach has a coherent algorithmic focus based on the use of adaptive, finite volume discretizations on locally rectangular grids. This approach is meant to permit the use of multiresolution, high performance strategies through the use of block-structured adaptive mesh refinement (AMR) combined with embedded boundary (EB) representations of complex geometries.

The embedded boundary, finite volume method is a numerical technology that provides a stable and accurate volume-of-fluid approach to complex geometry on a structured grid that enables direct numerical simulation (DNS) of geometries obtained from image data, engineering design, or constructive solid geometry (CSG). In this approach, the irregular domain is discretized as a collection of control volumes formed by the intersection of the problem domain with the cubic Cartesian grid cells, as in a “cut-cell” approach. The various operators are approximated by applying the divergence theorem, i.e., using finite volume differences of fluxes on

irregular control volumes, with the fluxes computed using the primary discretized dependent variables which approximate the solution evaluated at the centers of the original Cartesian cells. We use the CFD approach described in Ref. [5] to solve Eq. (1) that combines EB with AMR, or EBAMRINS.

There are several advantages to using adaptive, finite volume methods in the Chombo framework to simulate fluid dynamics in heterogeneous systems. The EB approach is a sharp interface method that yields explicit resolution of fluid–solid interfaces, which must be resolved for large surface area-to-volume ratios in confined pore space of heterogeneous materials. DNS from raw image data, engineering design, or CSG is an established, tractable workflow in Chombo [3,4,7–9]. In addition, AMR is an efficient numerical technique that adds resolution in areas of a domain only where it is needed while leaving the rest of the domain at a coarse resolution requiring less computational resources. In particular, AMR allows computations to be dynamically focused in areas of the domain where the solution is rapidly changing, such as near material and multiphase interfaces, sources, or turbulent flow. Furthermore, the finite volume representation is a natural description of multiphase flows with resolved interfaces, providing a model for sharp moving interfaces between two fluids, including interfacial tension, mass transfer, and wettability [10]. In general, the embedded boundary method coupled with adaptive mesh refinement in a high-performance computational fluid dynamics framework such as the Chombo-based EBAMRINS code provides a powerful high-resolution tool for modeling multiscale, multiphysics problems.

Algebraic Multigrid. Chombo uses multigrid to solve linear systems resulting from elliptic problems in incompressible flow such as pressure-Poisson and viscous, Helmholtz. The EB implementation of Chombo makes use of a PETSc interface for solving elliptic problems in arbitrarily complex geometries by algebraic multigrid (e.g., hypr BoomerAMG) for cases when native Chombo geometric coarsening and multigrid are not solvable [3]. The Chombo-PETSc interface combines the robustness of algebraic multigrid for complex geometries where geometric coarsening is not sufficient with an efficient and accurate finite volume method to resolve microscale/nanoscale flow and transport [3]. The approach is specifically intended to facilitate DNS in geometries obtained from image data of heterogeneous media such as those found in the geologic subsurface or engineering design [7–9]. In these complex geometry scenarios, the effectiveness of the coarsening operation in geometric multigrid is challenged by the presence of certain geometric features: multiply-connected domains, thin pore throats, sharp cusps, long aspect ratios, and semidisconnected cavities. Near a solid boundary in this representation, standard 7-point stencils in 3D do not apply. In fact, our conservative finite volume approach for complex boundaries in Ref. [5] elicits 27-point stencils in some cases. However, the introduction of algebraic multigrid into the structured grid formulation dramatically improves efficiency and robustness of the elliptic solvers, making more tractable solution to flow and advection-diffusion transport equations in realistic pore space obtained from image data or engineering design.

Grid Generation for Complex Geometry. In order to perform grid generation at large scale and with high grid resolution, the gridding and discretization techniques need to be consistent and efficient. The meshing technique used for gridding the geometry must be amenable to surface extraction from engineering designs or experimentally derived image data; and the computational model should be able to directly simulate flows on the meshes obtained from imagery without loss of geometric detail in a fast, accurate, stable, and conservative manner. Complex geometries have more traditionally been treated with conforming, body-fitted grids, especially in the finite element community. However, traditional structured or unstructured grid generation techniques can be time-consuming and involve much user interaction. Furthermore, it is difficult to generate grids that both preserve the fidelity of the

geometry and lead to robust solvers for the PDEs. The EB approach affords this kind of consistency between tractable gridding and efficient discretization.

The Chombo software toolbox contains capabilities for automatic mesh generation of complex geometries obtained from CSG or raw data derived from engineering designs or imagery using an implicit function representation of surfaces on a Cartesian grid. We have successfully applied this approach to a number of problems, including experimental validation of pore scale flows in both engineered and natural materials [4,5,7–9]. This fast and flexible geometry generation technology, coupled with the finite volume approach to discretization of the flow equations, provides a consistent methodology to proceed from geometry data to simulation with error control. In particular, we make use of CSG to develop surrogate models of heterogeneous media in a controllable and scalable manner that tests both performance (by taxing the EB algorithm) and spatial reach of the CFD solver.

Results

We present results from simulations of the Chombo-based EBAMRINS CFD solver on the exascale supercomputer, Frontier, at Oak Ridge Leadership Computing Facility (OLCF) as a part of the Department of Energy (DOE) Exascale Computing Project (ECP). We demonstrate the spatial reach (and performance) of the capability by extending historical scaling data for the embedded boundary CFD solver initially reported in [11] and with further scaling in [3,5]. The scaling test is based on simulation of flow in a cylinder packed tightly with spheres where the geometry is replicated as shown in Fig. 1. This replication process generates a weak scaling test for computational performance as each subsequent replicated geometry is simulated with twice the number of processors as the previous geometry. The replication unit is a cylinder that is 0.5 cm long with 0.5 cm diameter and packed with approximately 750 spheres, each with 250 micron radius. For visualization purposes, we show the first 3 geometries in the replication with 1:1, 2:1, and 4:1 aspect ratios, respectively, in Fig. 1 as in Ref. [5]. The full range of replicated geometries is listed in Table 1. We perform 10 timesteps of the flow algorithm for each geometry and average the time spent in each time-step to obtain the performance metric.

Scaling results of the EBAMRINS CFD solver on Frontier for the replicated geometries in Table 1 are shown in Fig. 2. Each plot is a weak scaling study itself (theoretically flat), where the initial datapoint on a given plot starts at a different node count (indicated in the legend) for a fixed degree-of-freedom (x -axis). Degrees of freedom are calculated by the number of grid cells (number of boxes $\times 32^3$ cells) multiplied by the number of flow variables (6, \mathbf{u} , and ∇p) where the number of boxes ranges from 512 to 2,097,152 along the x -axis, and each box contains 32^3 cells. For a given degree-of-freedom on the x -axis, the datapoints along a vertical for fixed x depict strong scaling performance, i.e., factor of 2 reduction in time as the number of nodes are doubled. In Fig. 2(a), we show several rounds of the scaling test based on two phases of early access to the machine. The orange data are from the second round and compare with the blue data from the first round in the left plot, but with an additional datapoint that reaches over 20 m in cylinder length with over 3×10^6 packed spheres, all-in-all accounting for over 400×10^9 degrees-of-freedom simulated and over 400 cm^3 of pore space. The data from the orange plot are also shown in tabular form in Table 1 to display the details of the replication (weak) scaling. The primary difference between the two phases of Frontier access in the left plot is the default software stack on Frontier from the module-based system, and the versions of PETSc and hypre configured with this software stack. The left plot also demonstrates strong scaling in the vertical direction. In Fig. 2(b), we duplicate the orange data on the left as a baseline to depict performance improvements in subsequent performance engineering cycles in the light green and purple datasets. We attribute the performance improvement in these latter data to faster performance of the PETSc-hypre solvers built in




Fig. 1 Packed cylinder geometry replication. (a) 1:1 aspect ratio — 0.5 cm long, 0.5 cm diameter cylinder packed with 750 spheres ($256 \times 256 \times 256$ grid cells). (b) 2:1 aspect ratio — 1 cm long, 0.5 cm diameter cylinder packed with 1500 spheres ($512 \times 256 \times 256$ grid cells), resolution $\Delta x \approx 19.5 \mu\text{m}$. (c) 4:1 aspect ratio — 2 cm long, 0.5 cm diameter cylinder packed with 3000 spheres, resolution $\Delta x \approx 19.5 \mu\text{m}$ ($1024 \times 256 \times 256$ grid cells). (d) Velocity data. Grid resolution: $\Delta x \approx 19.5 \mu\text{m}$. Sphere radii = 250 μm . Larger aspect ratios, up to 4096:1, are listed in Table 1.

Table 1 Replication (weak) scaling of the Chombo-based CFD code on Frontier

Aspect	32^3 boxes	spheres	D.O.F.	Nodes	MPI ranks	Length (cm)	Time (s)
1:1	512	750	1.01×10^8	1	8	0.5	41.30
2:1	1024	1500	2.01×10^8	2	16	1	47.72
4:1	2048	3000	4.03×10^8	4	32	2	53.02
8:1	4096	6000	8.05×10^8	8	64	4	54.87
16:1	8192	12,000	1.61×10^9	16	128	8	56.47
32:1	16,384	24,000	3.22×10^9	32	256	16	56.58
64:1	32,768	48,000	6.44×10^9	64	512	32	57.60
128:1	65,536	96,000	1.29×10^{10}	128	1024	64	58.40
256:1	131,072	192,000	2.58×10^{10}	256	2048	128	58.37
512:1	262,144	384,000	5.15×10^{10}	512	4096	256	64.25
1024:1	524,288	768,000	1.03×10^{11}	1024	8192	512	71.16
2048:1	1,048,576	1,536,000	2.06×10^{11}	2048	16,384	1024	86.97
4096:1	2,097,152	3,072,000	4.12×10^{11}	4096	32,768	2048	119.88

These performance numbers and problem specifications correspond to the orange data in left and right plots of Fig. 2.

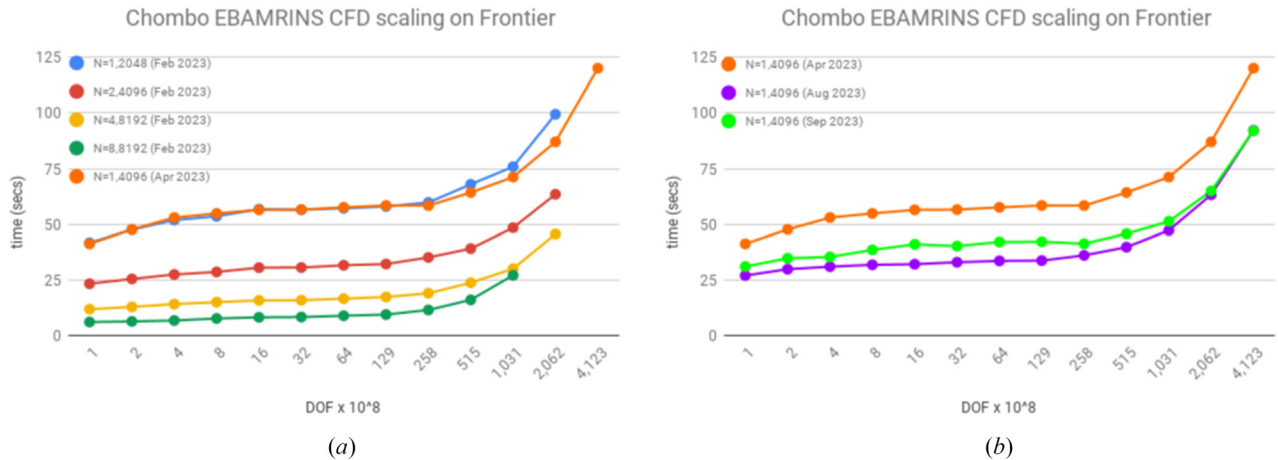


Fig. 2 Scaling of the Chombo-based EBAMRINS CFD code on OLCF Frontier. (a) Weak scaling is shown in each data line; strong scaling is shown in the collection of datapoints along a vertical direction for fixed degrees-of-freedom at an x location in the left figure. Orange data on the left is a later extension of the blue data. (b) Orange data line is repeated in right plot as baseline for a later scaling study with an improved software stack and performance engineering. Degrees of freedom are calculated by the number of grid cells (number of boxes $\times 32^3$ cells) multiplied by the number of flow variables (6). The number of boxes ranges from 512 to 2,097,152.

the GNU software stack environment compared to the Cray environment.

To also demonstrate applicability and relevance of this high-resolution CFD capability, we simulate a range of Reynolds numbers that include DNS of applied turbulent flows at engineering scales. Figure 3 represents a desalination device comprised of a spacer geometry and a permeable membrane as described in [12]. We show both laminar ($Re = 6.437$, Fig. 3(a)) and turbulent ($Re = 643.7$, Fig. 3(b)) flow conditions in the left column for this confined device-scale simulation. We also demonstrate transport polarization in the right column by solving an advection-diffusion equation for a nonreacting scalar component with the same solvers used for the flow equations. The flow boundary condition described in Ref. [12] is applied at the permeable membrane for these simulations.

Discussion

These new CFD scaling results extend previously reported results for lower, full machine concurrencies on CPU-based machines (e.g., [3,5]) to full machine capacity on Frontier GPUs. Simulation of over 3,000,000 microspheres in a 20 m cylinder with 0.5 cm diameter with over 400,000,000,000 degrees-of-freedom and over 400 cm³ of volume is one of the most computationally intensive simulations to date (cf. nuclear reaction simulations in Ref. [6] and subsequent news articles at Argonne that discuss Frontier results). Our previous

results on CPU-based machines reached over 100,000,000,000 degrees-of-freedom. We can now simulate a problem that is 4 times bigger with the same resolution than the previous capability in [3,5]. Furthermore, since the maximum problem size was simulated with half the machine (4096 nodes), there is room for additional capability on Frontier to simulate a problem 2 times the maximum problem size. The current limitation is not the grid generation for the next datapoint in the scaling test but, rather, the packing algorithm which specifies the geometry for packing over 6×10^6 spheres in a 40 m long cylinder. A simulation of this size would approach 1 trillion degrees-of-freedom.

Overall, scaling efficiency of 34.5% is excellent for such a wide range of concurrencies, representing nearly 4 orders of magnitude in node count and problem size. Performance at high concurrencies resembles the familiar and expected “hockey stick” with the upswing in time. As a result, if the last datapoint, for example, is not included, then the efficiency will increase to 47.5%, and so on. Previous scaling on CPU-based machines like NERSC Edison (as in Refs. [3,5]), and thereafter on NERSC Cori KNL, shows similar performance behavior at the highest concurrencies, for up to 131,072 cores and 524,288 cores, respectively. The extent of the orange plot line in Fig. 2 and the bottom line in Table 1 represent extreme scaling for aspect ratios and degrees-of-freedom that could not be simulated previously. The significant improvement in performance of the orange dataset (which in and of itself is good) in the right plot of Fig. 2 coincides with the change to the GNU

software stack on Frontier from the Cray software stack in the earlier scaling studies on the left.

Furthermore, the GPU-based architecture of the Frontier node provides more capability on a node-to-node basis, allowing for a strong scaling assessment for the first time with this historical scaling study. We observe approximately $2\times$ reduction in time when the number of nodes is doubled, as seen in the datapoints along a vertical in Fig. 2(a). Our load balancing and domain decomposition sweet spot on CPU-based machines has been one box of 32^3 grid cells per MPI rank. With Frontier GPUs, the memory bandwidth is increased such that we are able to load balance 64 boxes per GPU per MPI rank (cf. 1 box per MPI rank on CPUs). Frontier nodes contain 8 AMD MI250X GPUs.

The impact of this new CFD capability on the fluid dynamics community is significant. For subsurface and other porous media flows, domains that are larger than the conventional scale of an REV can now be simulated. For example, in Ref. [5], the CFD simulation of flooding in shale from image data was performed on a less than 100 micron cubic block with 50 nanometer resolution using the CPU-based machine Cori at NERSC. With this new capability, a domain closer to 1 millimeter cubic block could be simulated with the same resolution. This domain-to-resolution ratio is at the higher end of the REV scale and was previously unattainable for low porosity, low permeability, heterogeneous materials. In addition, the high-resolution turbulence simulation is capable of resolving viscous boundary layers in a device-scale confined flow, which is

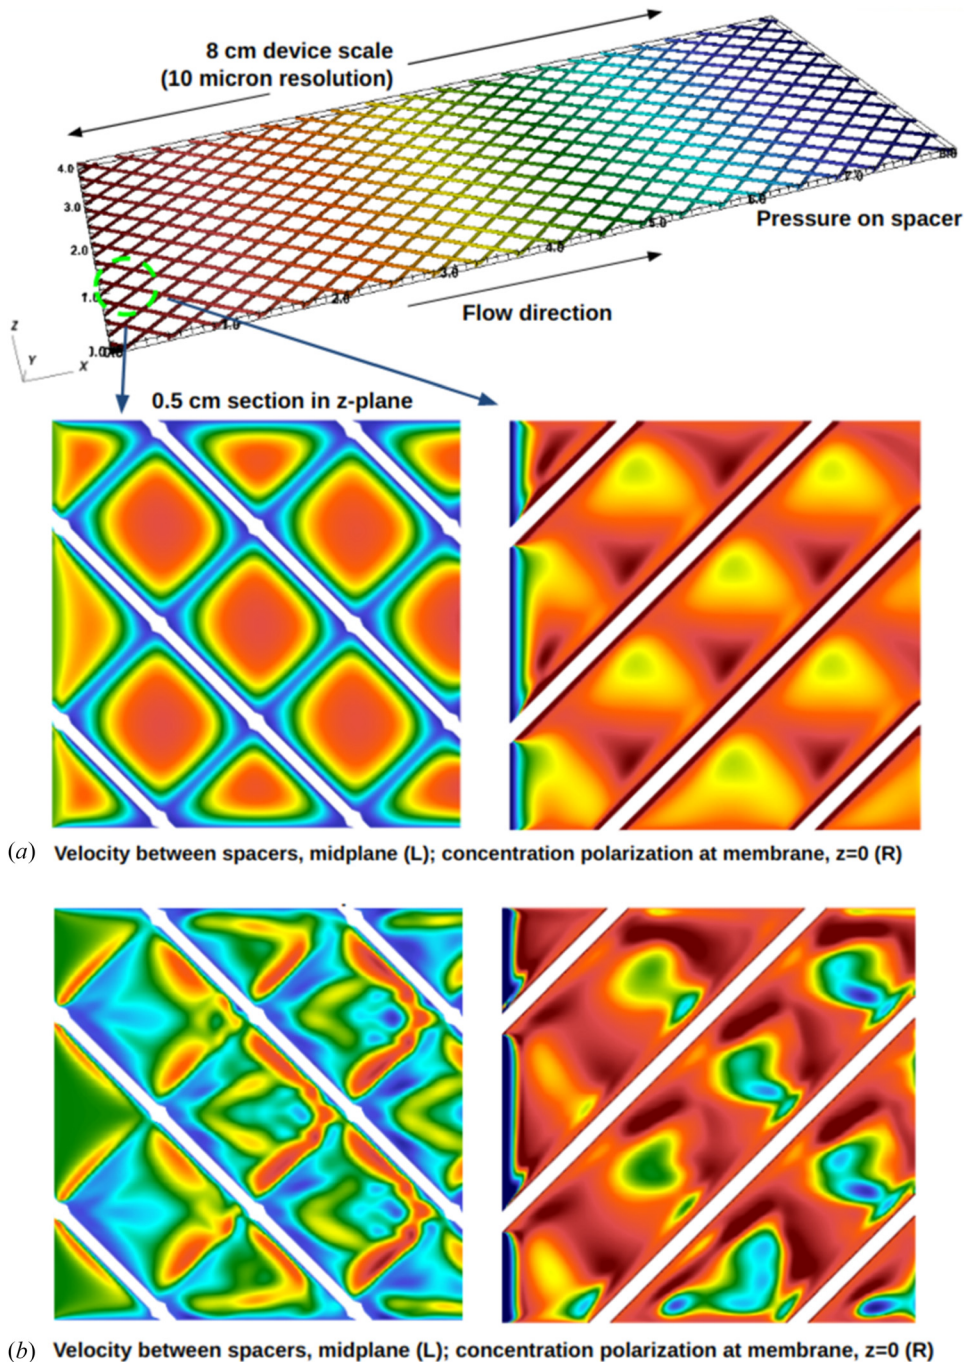


Fig. 3 CFD simulation in a permeable membrane of a desalination device [12]. (a) Laminar flow ($Re \approx 6$, top) and (b) turbulent flow ($Re \approx 600$, bottom) in left column, with corresponding concentration polarization in right column. Geometry generated by Sergi Molins, LBNL.

necessary to capture vortex shedding and other localized effects of turbulent flow that can affect engineering performance. As this was only a partial machine simulation, the capability can be applied to even higher Reynolds number flows in complex geometry environments that require higher resolution near solid boundaries.

Conclusions

The Frontier supercomputer at OLCF provides significant improvement to memory bandwidth due to the GPU-based architecture, providing much more application code capability compared to previous CPU-based machines. This added capability allows for simulation of larger domains with higher resolution, which was not possible before. We demonstrate extreme scaling results for CFD simulations in heterogeneous geometries with unprecedented domain-to-resolution ratios, challenging conventional understanding of the reach of brute force pore scale resolution into the meter scale and greater domain regime. The CFD simulation of over 400×10^9 degrees-of-freedom presented here is likely the largest CFD simulation in complex geometry to date. We use sphere-packed cylinders that serve as surrogates to represent heterogeneous complexity in porous media materials, both engineered and geologic. These constructive solid geometries allow for control of porosity-to-permeability ratios, which can be tested, verified, and validated for realistic materials, while also providing a portable benchmark for performance on different computational platforms. With the increased capability, we are also able to demonstrate strong scaling for the first time in this scaling test framework at high concurrencies. And because these packed geometries are reflective of porous media, they also have application to practical problems such as proppant transport and packing in subsurface fractures. The capability extends to a wide range of Reynolds numbers, from creeping, Stokes flow to turbulence, as well as differing application domains, especially confined flow systems with large aspect ratios.

Computational fluid dynamics simulations applied brute force at these extreme scales produce vast amounts of data, which can comprise ground truth science in the absence of experimental data. These new results demonstrate further spatial reach of pore scale resolution on domains that are much larger than the conventional understanding of an REV in heterogeneous, porous media. This accomplishment will lead to more accurate direct upscaling of continuum parameters such as permeability in Darcy flow and reaction rates in transport. Furthermore, the data created by this new capability will lead to new opportunities in AI/ML to perform dynamic data-informed multiscale, multiphysics prediction.

Acknowledgment

The author acknowledges the PETSc and hypre software teams for performance portability of their software on Frontier. The author also acknowledges Terry Ligocki for development of computational

geometry software to construct packed sphere geometries and Sergi Molins for geometry generation and problem specification of the permeable membrane simulation. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

Funding Data

- U.S. Department of Energy (DOE) Exascale Computing Project (contract No. DE-AC02-05CH11231; Funder ID: 10.13039/100000015).

Data Availability Statement

The datasets generated and supporting the findings of this article are obtainable from the corresponding author upon reasonable request.

References

- [1] Molins, S., Trebotich, D., Steefel, C. I., and Shen, C., 2012, "An Investigation of the Effect of Pore Scale Flow on Average Geochemical Reaction Rates Using Direct Numerical Simulation," *Water Resour. Res.*, **48**(3), p. W03527.
- [2] Bear, J., 1988, *Dynamics of Fluids in Porous Media*, trade paperback ed., Dover Publications, Dover Civil and Mechanical Engineering, New York.
- [3] Trebotich, D., Adams, M. F., Steefel, C. I., Molins, S., and Shen, C., 2014, "High Resolution Simulation of Pore Scale Reactive Transport Processes Associated With Carbon Sequestration," *Comput. Sci. Eng.*, **16**(6), pp. 22–31.
- [4] Molins, S., Trebotich, D., Yang, L., Ajo-Franklin, J. B., Ligocki, T. J., Shen, C., and Steefel, C. I., 2014, "Pore-Scale Controls on Calcite Dissolution Rates From Flow-Through Laboratory and Numerical Experiments," *Environ. Sci. Technol.*, **48**(13), pp. 7453–7460.
- [5] Trebotich, D., and Graves, D., 2015, "An Adaptive Finite Volume Method for the Incompressible Navier–Stokes Equations in Complex Geometries," *Commun. Appl. Math. Comput. Sci.*, **10**(1), pp. 43–82.
- [6] Fischer, P., Kerkemeier, S., Min, M., Lan, Y.-H., Phillips, M., Rathnayake, T., Merzari, E., et al., 2022, "NekRS, a GPU-Accelerated Spectral Element Navier–Stokes Solver," *Parallel Comput.*, **114**, p. 102982.
- [7] Kneafsey, T., Trebotich, D., and Ligocki, T., 2023, *Direct Numerical Simulation of Flowthrough Nanoscale Shale Pores in a Mesoscale Sample*, Springer International Publishing, Cham, p. 87.
- [8] Trebotich, D., and Ligocki, T., 2023, *High Resolution Simulation of Fluid Flow in Press Felts Used in Paper Manufacturing*, Springer International Publishing, Cham, p. 132.
- [9] Molins, S., and Trebotich, D., 2023, *Pore-Scale Controls on Calcite Dissolution Using Direct Numerical Simulations*, Springer International Publishing, Cham, p. 135.
- [10] Miller, G. H., and Trebotich, D., 2012, "An Embedded Boundary Method for the Navier–Stokes Equations on a Time-Dependent Domain," *Comm. App. Math. Comp. Sci.*, **7**(1), pp. 1–31.
- [11] Trebotich, D., Straalen, B. V., Graves, D. T., and Colella, P., 2008, "Performance of Embedded Boundary Methods for CFD With Complex Geometry," *J. Phys.: Conf. Ser.*, **125**, p. 012083.
- [12] Van Wagner, E. M., Sagle, A. C., Sharma, M. M., and Freeman, B. D., 2009, "Effect of Crossflow Testing Conditions, Including Feed pH and Continuous Feed Filtration, on Commercial Reverse Osmosis Membrane Performance," *J. Membr. Sci.*, **345**(1–2), pp. 97–109.