

High-Resolution Simulation of Pore-Scale Reactive Transport Processes Associated with Carbon Sequestration

David Trebotich, Mark F. Adams, Sergi Molins, and Carl I. Steefel | Lawrence Berkeley National Laboratory
 Chaopeng Shen | Penn State University

New investigative tools, combined with experiments and computational methods, are being developed to build a next-generation understanding of molecular-to-pore-scale processes in fluid-rock systems. A new numerical simulation modeling capability, known as Chombo-Crunch, resolves flow and transport processes in geometric features obtained from image data of realistic pore space at unprecedented scale and resolution.

When CO₂ is injected into the Earth's subsurface, the system is forced far from equilibrium, both chemically and mechanically. The result is a nonlinear dynamical regime in which emergent behavior develops at the pore and greater scales. These nonlinear interactions among multiphase flow, solute transport, mineral dissolution, and precipitation have been addressed in a cursory fashion at the larger porous-continuum scale, where grains and pores aren't resolved, but they've largely been neglected at the pore scale, where the chemical and physical environments undergo strong variations locally in space and time. Interest in the pore scale stems from the need to resolve explicitly (in contrast to the porous-continuum scale) the key fluid-fluid (H₂O and CO₂) and fluid-solid interfaces in order to arrive at a mechanistic understanding of how to control CO₂ injection in the subsurface. By carefully understanding processes at the pore scale, the overall scientific goal of this work is to bring such knowledge to bear on the macroscopic scale of a reservoir, which is the relevant scale for carbon sequestration.

This article reports a recent advancement that combines an embedded boundary (EB) algorithm with algebraic multigrid (AMG) method to produce an efficient, robust, and scalable approach for modeling flow and transport in very complex 3D

geometries obtained from image data. Primarily, we identify the observed convergence problem of geometric multigrid (GMG) methods in the presence of very complex geometries and demonstrate that the introduction of AMG into our elliptic solvers leads to a robust solution. We also demonstrate the new combined approach's performance gains, with scalability to more than 100,000 processor cores. We apply the method to complex pore-scale domains obtained from both constructive solid geometry as well as from image data of experiments that model processes associated with carbon sequestration.

Background

Several classes of methods model subsurface flow and transport in porous media, but most typically address flow and transport in aquifers over distances of 1 to 10 km. In these models, the pore space's microscale geometry is treated as an "effective medium," and flow and transport processes are replaced with averaged variables, balance equations, and bulk properties. For example, fluid momentum balance is governed by Darcy's law as a function of a pressure gradient, elevation gradient, and a medium-dependent permeability parameter. However, processes that emanate at microscopic length scales (dissolution-diffusion-convection, for example) can evolve into larger-scale dynamics over time,¹ underlining the need to consider

the resolution of the pore-scale geometry and microscale processes.²

The geometry of subsurface pore space is often characterized by long, thin channels meandering between multiply connected domains in highly heterogeneous microscopic environments. Other characteristics of these media are low porosity, very high specific interfacial area, heterogeneous preferred flow paths (or wormholing), sharp cusps, and semidisconnected cavities. Consequently, in direct numerical simulation of pore-scale systems, the fluid incompressibility constraint and both viscous and diffusive forces result in the need for elliptic solvers in very complex geometries. However, realistic geometries pose several challenges to numerical schemes and solvers in terms of stability, accuracy, flexibility, memory, and CPU demands.

An advantage of direct numerical simulation at the pore scale is that it can be used to perform validation studies of porous media experiments. This requires direct simulation in realistic geometries obtained from image data, as in a geologic core sample or packed bed column. For the “image to simulation” process to be successful, the meshing technique used for gridding the geometry must not only be amenable to surface extraction techniques for experimentally-derived image data but also consistent with an efficient and convergent discretization method for the equations of motion.

Complex geometries have more traditionally been treated with conforming, body-fitted grids. However, as noted, porous media are typically highly irregular and dense, especially in the geologic subsurface. Under such circumstances, unstructured grid generation techniques are time-consuming and involve much user interaction. Furthermore, it’s difficult to generate grids that both preserve the geometry’s fidelity and lead to robust solvers. In addition, if the pore-space boundary changes due to, say, precipitation or dissolution in the geologic case, then a body-fitted gridding process, which is already complicated, also becomes time-dependent.

Finite-volume methods have been shown to be computationally efficient as well as accurate and stable for modeling microscale flow and transport in irregular geometries. In particular, the EB method is a finite-volume approach to gridding complex geometry that’s specifically intended to enable direct simulation from image data. 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 on the irregular control volumes, with the fluxes computed by using primary discretized dependent variables that approximate the solution evaluated at the centers of the original Cartesian cells. Away from the boundary, the finite-volume method reduces to a standard finite-difference approximation. This approach has been used as the basis for second-order accurate methods for elliptic, parabolic, and hyperbolic partial differential equations in two and three dimensions,^{3,4} shown to scale efficiently for elliptic problems,⁵ and extended to flow and transport in complex microscale geometries⁶ as well as time-dependent fixed and free boundary problems.⁷ Coupled with adaptive mesh refinement (AMR), the EB approach provides a powerful high-resolution tool for modeling multiphysics problems.

The EB method makes use of matrix-free GMG methods to solve the elliptic Poisson and positive definite Helmholtz equations that result from a semi-implicit predictor-corrector formulation for incompressible viscous flow and transport (scalar advection-diffusion-reaction [ADR]).^{6,8} We’ve observed that generic GMG performs poorly for very complex flow domains, and, in some cases, it doesn’t even converge. An important reason for GMG’s poor performance is that the geometrically coarsened grid doesn’t represent the physical boundary in a porous medium; technically, the operator has a null space in the pressure-Poisson case. As an alternative, AMG methods^{9,10} construct coarse grid spaces algebraically and, more importantly, create coarse grid operators using an algebraic Galerkin process, which is very robust with respect to complex geometry. The AMG approach has also been applied to Navier-Stokes equations in complex geometries and demonstrated for simple problems in two dimensions.¹¹

Technical Approach

Our technical approach is a predictor-corrector type based on a Godunov projection method.

Equations

First, we consider the incompressible Navier-Stokes (INS) equations with constant density and the scalar transport equations

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} = -\nabla p + \nu \Delta \mathbf{u} \quad (1)$$

$$\nabla \times \mathbf{u} = 0 \quad (2)$$

$$\frac{\partial c}{\partial t} + (\mathbf{u} \cdot \nabla) c = \nabla \cdot D \nabla c + R, \quad (3)$$

where \mathbf{u} is fluid velocity, ∇p is pressure gradient, ν is kinematic viscosity, c is a scalar density, D is the scalar's diffusivity, and R is a reaction source term.

We've developed a semi-implicit predictor-corrector projection method¹² based on a previous¹³ and more recent work^{6,7} to solve the equations of motion. This method previously used GMG-based elliptic solvers to model reactive transport at the pore scale, as in a separate study.⁸ Here, we discuss the details of these new elliptic solvers based on the hybrid EB-AMG approach needed by the predictor-corrector method to solve flow and transport in very complex geometries.

Two types of elliptic problems must be solved within this predictor-corrector framework. The first is the Poisson equation resulting from the pressure projection formulation:

$$\Delta \phi = \nabla \times \mathbf{u}^{*,n+1} \text{ on } \Omega \quad (4)$$

$$\mathbf{n} \times \nabla \phi = 0 \text{ on } \partial \Omega. \quad (5)$$

The second is a positive definite Helmholtz equation resulting from the implicit treatment of the viscous term:

$$(\alpha I - \beta \Delta) \mathbf{u}^{*,n+1} = f(\mathbf{u}^n, \mathbf{u}^{n+1/2}) \text{ on } \Omega \quad (6)$$

$$\mathbf{u} = 0 \text{ on } \partial \Omega. \quad (7)$$

Similar to Equation 6, a positive definite Helmholtz equation results from scalar diffusion in transport problems. For ease of discussion, we address the Helmholtz problem resulting from the velocity equation and note that the approach also applies to transport, although the relaxation parameter is different (species diffusivity). In each of these cases, an H^1 -elliptic problem must be solved.

Algebraic Multigrid

Multigrid is well known as an optimal solution method for H^1 -elliptic problems on structured meshes with a serial complexity of $O(n)$ (with n degrees of freedom) and $O(\log n)$ parallel computational depth.^{9,14–16} Multigrid methods are motivated by the observation that a low-resolution discretization of an operator can capture error modes or components that are expensive to

compute directly on a highly resolved discretization. For example, an inexpensive iterative method such as Gauss-Seidel is effective at reducing high-frequency or high-energy error via geometrical smoothing. More generally, any poorly locally determined solution component has the potential to be resolved with coarser representation. This process can be applied recursively with a series of coarse grids, thereby requiring that each grid resolve only the components of the error that it can solve efficiently. These coarse grids have fewer grid points—typically, about a factor of two in each dimension—such that the total amount of work in multigrid iterations can be expressed as a geometric sum that converges to a small factor of the work on the finest mesh.

Multigrid communicates between grids through prolongation and restriction. The columns of prolongation operator, P , are the discrete representation of the coarse grid functions; we use P to map corrections to the solution from the coarse grid to the fine grid. Residuals are mapped from the fine grid to the coarse grid with the restriction operator, R , which is often equal to P^T . The coarse grid matrix can be formed in one of two ways: either algebraically to form Galerkin (or variational) coarse grids ($A_{\text{coarse}} \leftarrow RA_{\text{fine}}P$) or by creating a new discrete problem on each coarse grid (if an explicit coarse grid is available), thereby letting the application construct the operator. Figure 1 shows the standard multigrid V -cycle with smoother $x \leftarrow S(A, x_0, b)$.

Although multigrid was originally developed with explicit coarse grids (geometric multigrid), the method only requires a prolongation operator if $R = P^T$. Over the past decade, AMG methods have been developed for unstructured problems that generalize the construction of the prolongation operator. AMG methods, by the broadest definition, construct the coarse grid operators with a Galerkin process. More specifically, AMG methods refer to multigrid methods that construct the coarse grid spaces and operators from the matrix alone. Many AMG methods have been developed, such as smoothed aggregation¹⁷ and “classic” algebraic multigrid.¹⁰

In this work, we apply a standard smoothed aggregation method in the Portable Extensible Toolkit for Scientific computing (PETSc) for pressure correction in Equation 4 and the positive definite Helmholtz in Equation 6. (The operators that result from the EB formulation used here aren't symmetric and are discussed in the following section.)

```

function  $M\ GV(A_i, b_i)$ 
  if there is a coarser grid  $i + 1$ 
     $x_i \leftarrow S^{v_1}(A_i, 0, b_i)$            -  $v_1$  iterations of the (pre)smoother
     $r_i \leftarrow b_i - Ax_i$ 
     $r_{i+1} \leftarrow R_{i+1}(r_i)$            - restriction of residual to coarse grid
     $e_{i+1} \leftarrow M\ GV(R_{i+1} A_i P_{i+1}, r_{i+1})$  - the recursive application
     $x_i \leftarrow x_i + P_{i+1}(e_{i+1})$        - prolongation of coarse grid correction
     $x_i \leftarrow S^{v_2}(A_i, x_i, b_i)$      -  $v_2$  iterations of the (post)smoother
  else
     $x_i \leftarrow A_i^{-1} b_i$            - exact solve of coarsest grid
  return  $x_i$ 

```

Figure 1. Multigrid $V(v_1, v_2)$ -cycle algorithm.




Figure 2. Embedded boundary (EB) approach.⁶ (a) Example of an irregular geometry on a Cartesian grid, (b) close-up view of embedded boundaries “cutting” regular cells, and (c) single irregular cut cell showing boundary fluxes (right). Shaded area represents volume of cells excluded from domain, and dots represent cell centers. X's represent centroids.

Structured Grid Method

Our solver stencils are based on a finite-volume discretization of elliptic problems in the presence of complex geometries.⁶ We use an EB approach whereby the problem domain is intersected with a rectangular Cartesian grid resulting in irregular, or cut, cells near the boundary, as in Figure 2. A discrete form of the divergence theorem can then be applied in the irregular cells to form a conservative approximation to the divergence of a flux:

$$D(\vec{F})_v = \frac{1}{h\kappa_v} \left(\left(\sum_{d=1}^D \alpha_{i+\frac{1}{2}^d} \widetilde{F}_{i+\frac{1}{2}^d}^d - \alpha_{i-\frac{1}{2}^d} \widetilde{F}_{i-\frac{1}{2}^d}^d \right) + \alpha_v^B F_v^B \right), \quad (8)$$

where α is the area fraction of a cell face, h is the mesh spacing, and κ is the cell's volume fraction; $\widetilde{F}_{i+\frac{1}{2}^d}$ indicates that the flux has been interpolated to the face centroid using linear (2D) or bilinear (3D) interpolation of face-centered fluxes. For

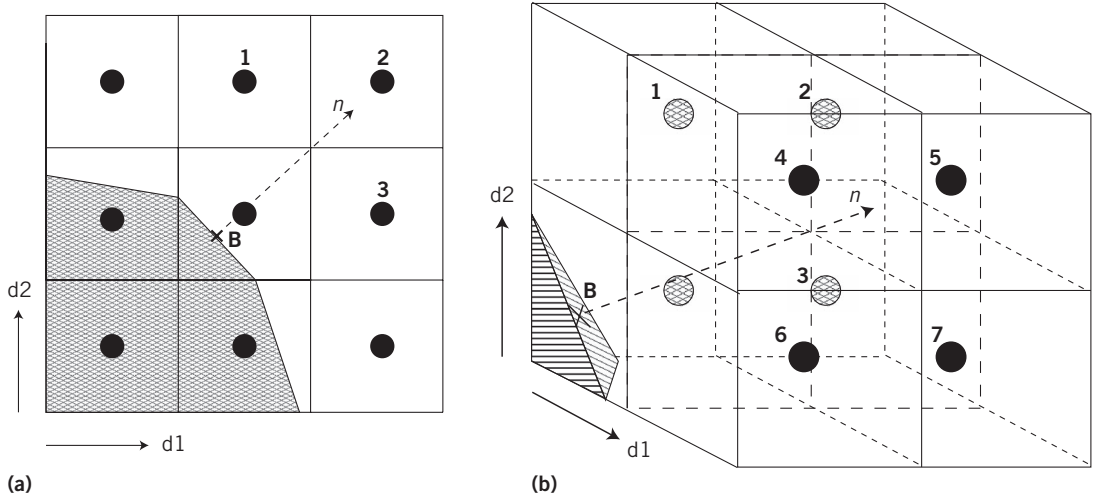


Figure 3. Least squares stencil to obtain flux for Dirichlet boundary condition on embedded boundary in (a) 2D and (b) 3D.⁶

example, given the cell edge with outward normal \hat{e}^1 , with centroid \vec{x} , the 2D linearly interpolated flux in the d ($d \neq 1$) direction is defined by

$$\begin{aligned} \tilde{F}_{i+\frac{1}{2},j}^d &= \eta F_{i+\frac{1}{2},j} + (1-\eta) F_{i+\frac{1}{2},j+\frac{1}{2}} \\ \eta &= 1 - \frac{|\vec{x} \cdot \hat{e}^d|}{h_d} \\ \pm &= \begin{cases} +\vec{x} \cdot \hat{e}^d > 0 \\ -\vec{x} \cdot \hat{e}^d \leq 0. \end{cases} \end{aligned} \quad (9)$$

We can write the 3D bilinear interpolation of the flux for the face with normal \hat{e}^1 as follows:

$$\begin{aligned} \tilde{F}_{i+\frac{1}{2},j,k}^d &= \omega F_{i+\frac{1}{2},j,k}^d + (1-\omega) F_{i+\frac{1}{2},j+\frac{1}{2},k+\frac{1}{2}}^d \\ \omega &= 1 - \frac{|\vec{x} \cdot \hat{e}^{d'}|}{h_{d'}} \\ \pm &= \begin{cases} +\vec{x} \cdot \hat{e}^{d'} > 0 \\ -\vec{x} \cdot \hat{e}^{d'} \leq 0, \end{cases} \end{aligned} \quad (10)$$

where $d' \neq d$, $d' \neq 1$.

The divergence theorem's conservative discretization provides a flux-based formula for elliptic operations in irregular cells. In the context of Poisson's equation, the operator is the Laplacian and the flux is simply the gradient of a scalar $\vec{F} = \nabla \varphi$ with Neumann boundary conditions $F^B = \hat{n} \cdot \nabla \varphi = 0$ at the EB. However, in the context of Helmholtz, the irregular boundary is a no-slip boundary for the velocity, requiring an elliptic

operator with Dirichlet boundary conditions at the EB. In this case, the flux at the EB, $F^B = \hat{n} \cdot \nabla \varphi$, must be constructed while maintaining the method's global accuracy. (For the scalar transport Helmholtz equation, the boundary condition is Neumann.)

We construct a Dirichlet flux at boundaries by solving the least squares linear system for $\nabla \varphi$:

$$\begin{aligned} A \cdot \nabla \varphi &= \delta \varphi \\ A &= (\delta \vec{x}_1, \delta \vec{x}_2, \dots, \delta \vec{x}_p)^T \\ \delta \varphi &= (\delta \varphi_1, \delta \varphi_2, \dots, \delta \varphi_p)^T \\ \delta \vec{x}_m &= \vec{x}_m - \vec{x}_B \\ \delta \varphi_m &= \varphi_m - \varphi_B. \end{aligned}$$

The stencil of points ($m = 1, 2, \dots, p$), excluding the irregular cell that contains the EB, is determined by the list of cells along a monotone path from the irregular root cell given some radius of influence. In 2D, the minimum stencil size for uniqueness is $p = 3$; in 3D, it's $p = 7$. Figure 3 shows exemplary stencils. For very complex geometries, a radius of influence of 3 ensures robustness.

We note that our finite-volume approach to discretization of elliptic operators reduces to a standard finite-difference approximation in regular cells away from boundaries (5-point Laplacian stencil in 2D; 7-point Laplacian stencil in 3D). However, porous media domains require a large number of cut-cell discretizations due to tortuous,

microscopic pore space. Therefore, the required stencils are much more complicated in the irregular cut cells due to interpolation of face and boundary fluxes, thus supporting our claims of complexity and asymmetry compared to standard finite-difference operators.

Results

Our flow and conservative transport solver, known as Chombo-Crunch, is based on adaptive EB methods developed in the Chombo software framework (<http://chombo.lbl.gov>). Chombo generates the computational grid, be it from constructive solid geometry techniques (synthetic) or image data surface reconstruction (realistic), processes the geometries, decomposes the domain into sets of boxes on individual processor cores, and forms the operators (matrix and right-hand side).

A key feature of this framework is the integration of the PETSc scalable equation solver library into the Chombo libraries to facilitate access to many other linear and nonlinear solvers besides GMG.¹⁸ We used a smoothed aggregation AMG solver native to PETSc v3.4 (`-pc type gamg -pc gamg sym graph true`) written by one of the authors (Mark F. Adams). We also compared the gamg solver's performance with hypre (`-pc type hypre -pc hypre type boomeramg`).¹⁹ We used the GMRES Krylov method because EB generates nonsymmetric matrices, preconditioned with one AMG V(1,1)-cycle. The convergence criterion for the norm of the residual error of the elliptic operator was a relative tolerance of 10^{-12} .

Scalability and Parallel Efficiency

To test the scalability of Chombo-Crunch, we performed a weak scaling test for flow through a cylinder packed with spheres. We obtained the geometry via constructive solid geometry techniques that use an implicit function representation on a grid. We performed 10 time steps of the flow code and took the average time per time step in seconds for comparison. We used a sweet spot for domain decomposition and load balancing of one box per processor core, where one box was 32^3 cells. We conducted these tests on the NERSC Cray XC30 system, Edison, for up to 131,072 cores. This weak scaling test is a replication scaling test in which we took a fixed number of spheres (or near fixed, due to random placement of spheres) per 32^3 patch of the domain and assigned one patch per processor core for three different aspect ratios of flow in a cylinder: a 1-to-1 (750 packed spheres) run on

512-4,096-32,768 cores, a 2-to-1 (1,500 packed spheres) run on 1,024-8,192-65,536 cores, and a 4-to-1 (3,000 spheres) run on 2,048-16,384-131,072 cores.

Due to the random placement of the spheres in the packing, we were only able to guarantee true replication (number of spheres per box scaled with number of processors) within one set of runs (512-4,096-32,768) and not necessarily between the three sets (512-1,024-2,048). However, theoretically, the scaling should be very close to replicated data. Figure 4 shows all the data points on one plot, with the three official sets of weak scaling depicted by shape. A slight dip at $N = 4,096$ was likely due to a lower number of irregular cut cells from refinement of the geometry. We observed an increased slope in the weak scaling curve at the two highest concurrencies, but overall, the time only increased 50 percent from the lowest concurrency (512) to the highest (131,072). Overall, we saw about 67 percent efficiency from 43 to 10,923 sockets, which we consider to be excellent for a vast range of concurrencies and considering the complex flow physics and geometry modeled here. We obtained nearly equivalent scaling numbers when we used hypre instead of gamg. Furthermore, this benchmark scaling problem clearly demonstrates that we achieve two to three time steps per minute when using a domain decomposition and load balancing sweet spot of 32^3 patches per core. From experience, we find that one time step per minute is satisfactory for large-scale computations of fluid dynamics in very complex geometries.

The simulation data in Figure 5 demonstrates the flow and reactive transport capability at high resolution ($<5 \mu\text{m}$). The three flow components demonstrate tortuosity in the packed cylinder, with hot colors indicating fast flow paths and cold colors slow flow paths or even recirculation. The reactive component (calcium released in the dissolution of calcite) reveals that rates are a function of flow conditions. In zones with slow flow or near stagnant conditions, dissolution rates are transport controlled—that is, dissolution is faster than transport (mostly by diffusion) of reactants.⁸

Heterogeneous Image Data

We demonstrated the robustness of the EB-AMG algorithm for the simulation of flow in a complex pore-scale geometry obtained from image data of a flow-through capillary tube experiment packed with mineral grains (calcite). The cylinder had a




Figure 4. Weak scaling test results, with average time spent per time step in seconds. Three sets of weak scaling datasets—512-4,096-32,768 squares, 1,024-8,192-65,536 triangles, and 2,048-16,384-131,072 circles—are shown in a single plot due to replication of the domain while maintaining one 32^3 box of cells per core.




Figure 5. Flow and reactive transport in a packed cylinder. From top left, counterclockwise, (a–c) three components of velocity in a 2-cm long cylinder packed with 3,000 spheres computed on 131,072 processor cores on Edison at NERSC. Inflow velocity is 0.01 cm/sec. (d) Transient reactive transport in a 1-cm long cylinder packed with 1,500 spheres run on 65,536 processor cores.

length of 0.71235 cm in z and a diameter of 0.0524 cm; resolution was $512 \times 512 \times 6144$ ($1.19 \mu\text{m}$). The mean inlet velocity was 0.035 cm/sec.

The top of Figure 6 shows the discrete EB representation of crushed calcite packed in a cylinder with two magnifications below; the middle panel is the steady-state velocity field, and the

bottom panel is a transient snapshot of a passive scalar being advected and diffused in the domain. We ran these simulations on 49,152 processor cores at NERSC. In terms of performance for this problem, the code averaged one time step per minute, which we found to be an adequate benchmark for achieving validatable scientific results.²⁰

Chombo-Crunch is an efficient and scalable framework for direct numerical simulation of pore-scale flows and reactive transport from experimental image data. It uses a new approach for solving elliptic partial differential equations (PDEs) in very complex geometries based on AMG and an EB finite-volume method. The EB-AMG approach integrated into the Chombo-Crunch code framework can facilitate modeling pore-scale flow and transport in geometries obtained from image data of heterogeneous media, as in the geologic subsurface or energy materials. The advantage of AMG is evident for a range of geometries, from a thin single pore throat to the more complex pore space derived from experiments. Scaling results show not only converged solutions over previous GMG results, which didn't converge in some cases, but also improved performance at higher concurrencies.⁵ A more appropriate load-balancing algorithm that includes hyperbolic solver performance in addition to elliptic will potentially improve the parallel speedup for scaling tests. We also plan on investigating further some of the dips in the weak scaling plot, as well as performing a comparison of the `gamg` solver with an AMG solver in `hypr`.

Predictive simulation of the flow field at unprecedented resolution serves as a basis for simulating reactive transport processes associated with carbon sequestration.^{2,20} This new capability is useful for future reactive transport studies regarding the scale dependence of parameters, interactions between processes spanning multiple scales, and distillation of apparent physical laws from fine-scale simulations. It can also be used to investigate diffusion control of chemical reaction rates at solid-fluid interfaces. ■

Acknowledgments

This work was funded by the Office of Advanced Scientific Computing Research, in part by the Center for Nanoscale Control of Geologic CO_2 , an Energy Frontier Research Center funded by the Office of Basic Energy

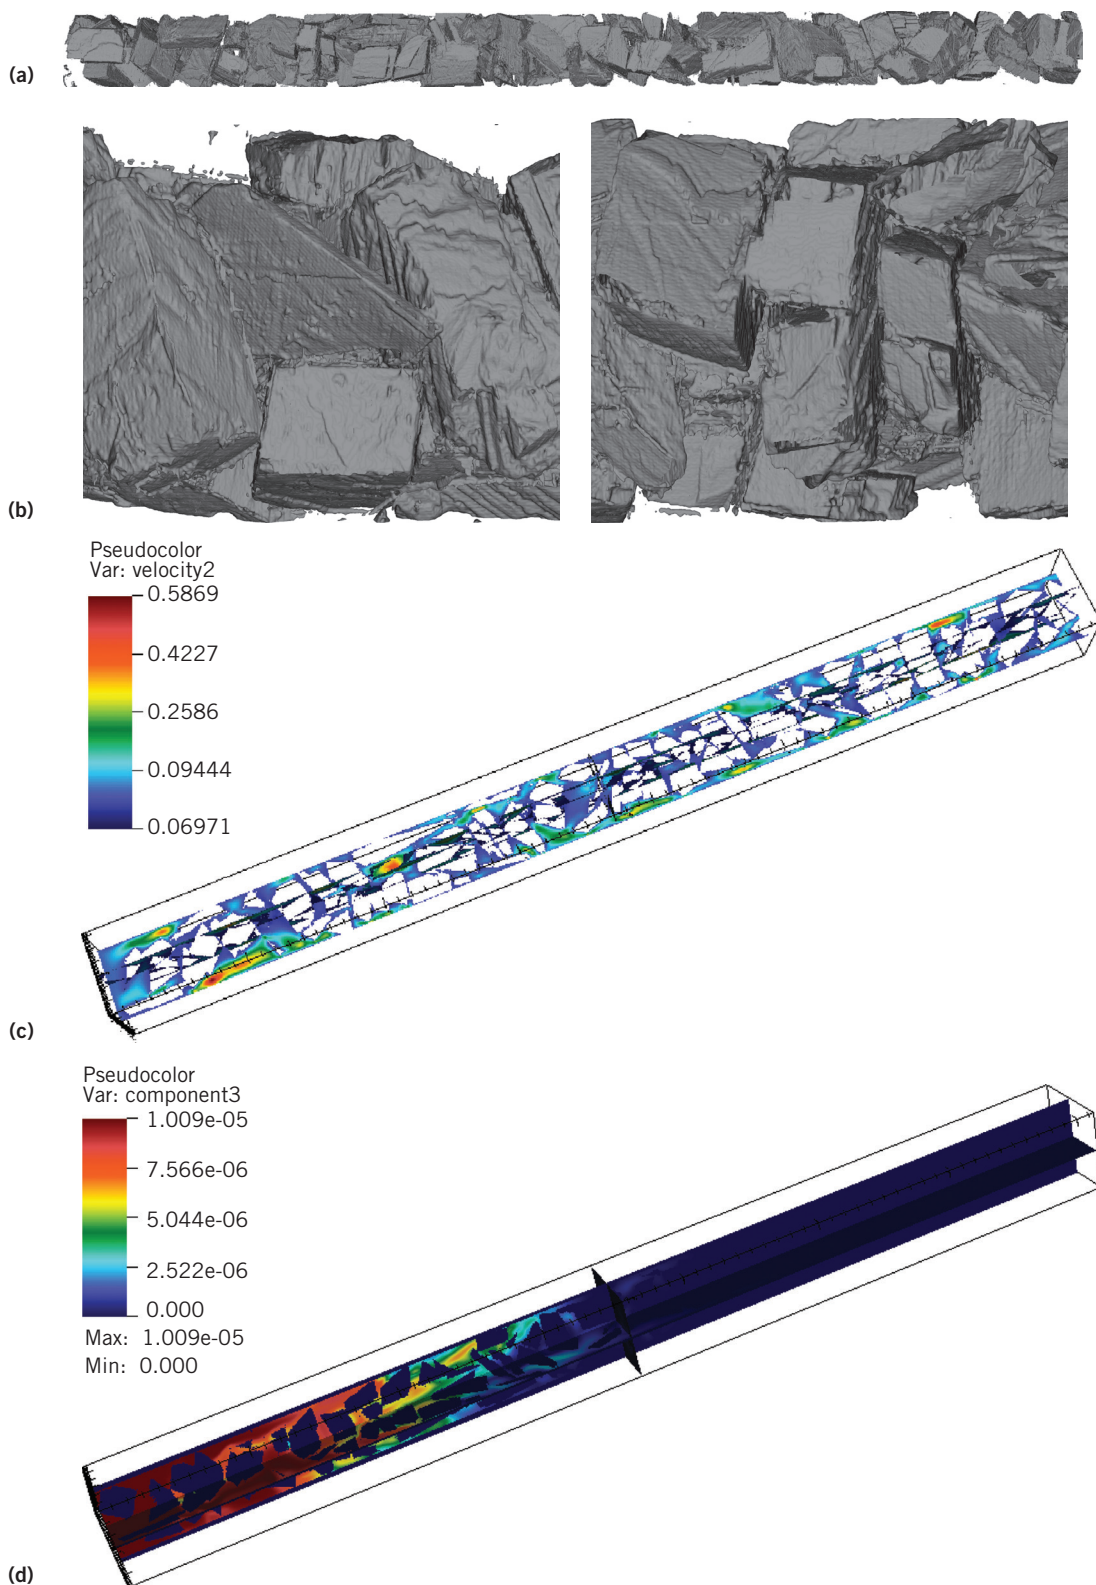


Figure 6. Flow simulation in a geometry obtained from images of a capillary tube experiment. From top to bottom, full computational domain, two magnifications of domain, steady-state velocity, and the nonreactive component at 3 seconds. Experimental image data courtesy of J. Ajo-Franklin and L. Yang, Lawrence Berkeley National Laboratory.

Sciences; it also used resources of the National Energy Research Scientific Computing Center, all supported by the Office of Science of the US Department of Energy under contract DE-AC02-05CH11231.

References

1. G.S.H. Pau et al., "High Resolution Simulation and Characterization of Density-Driven Flow in CO₂ Storage in Saline Aquifers," *Advances in Water Resources*, vol. 33, no. 4, 2010, pp. 443–455.
2. C.I. Steefel, S. Molins, and D. Trebotich, "Pore Scale Processes Associated with Subsurface CO₂ Injection and Sequestration," *Geochemistry of Geologic CO₂ Sequestration*, D.J. DePaolo et al., eds., Mineralogical Soc. Am. Geochemical Soc., 2013, pp. 259–303.
3. P. McCorquodale, P. Colella, and H. Johansen, "A Cartesian Grid Embedded Boundary Method for the Heat Equation on Irregular Domains," *J. Computational Physics*, vol. 173, no. 2, 2001, pp. 620–635.
4. P. Colella et al., "A Cartesian Grid Embedded Boundary Method for Hyperbolic Conservation Laws," *J. Computational Physics*, vol. 211, no. 1, 2006, pp. 347–366.
5. D. Trebotich et al., "Performance of Embedded Boundary Methods for CFD with Complex Geometry," *J. Physics: Conf. Series*, vol. 125, no. 1, 2008, p. 012083.
6. D. Trebotich and D.T. Graves, "An Adaptive Finite Volume Method for Incompressible Flow in Complex Geometries," under editorial revision with *Comm. Applied Mathematics and Computational Science*, 2014.
7. G.H. Miller and D. Trebotich, "An Embedded Boundary Method for the Navier-Stokes Equations on a Time-Dependent Domain," *Comm. Applied Mathematics and Computational Science*, vol. 7, no. 1, 2012, pp. 1–31.
8. S. Molins et al., "An Investigation of the Effect of Pore Scale Flow on Average Geochemical Reaction Rates Using Direct Numerical Simulation," *Water Resources Research*, vol. 48, no. 3, 2012; doi:10.1029/2011WR011404.
9. U. Trottenberg et al., *Multigrid*, Academic Press, 2001.
10. A. Brandt, S.F. McCormick, and J.W. Ruge, "Algebraic Multigrid (AMG) for Sparse Matrix Equations," *Sparsity and Its Applications*, D.J. Evans, ed., Cambridge Univ. Press, 1984.
11. M. Griebel, T. Neunhoffer, and H. Regler, "Algebraic Multigrid Methods for the Solution of the Navier-Stokes Equations in Complicated Geometries," *Int'l J. Numerical Methods in Fluids*, vol. 26, no. 3, 1998, pp. 281–301.
12. A.J. Chorin, "Numerical Solutions of the Navier-Stokes Equations," *Mathematics of Computation*, vol. 22, 1968, pp. 745–762.
13. D. Trebotich and P. Colella, "A Projection Method for Incompressible Viscous Flow on Moving Quadrilateral Grids," *J. Computational Physics*, vol. 166, 2001, pp. 191–217.
14. A. Brandt, "Multi-Level Adaptive Solutions to Boundary Value Problems," *Mathematics of Computation*, vol. 31, 1977, pp. 333–390.
15. W. Hackbusch, "Multigrid Methods and Applications," *Computational Mathematics*, vol. 4, Springer-Verlag, 1985.
16. W.L. Briggs, *A Multigrid Tutorial*, SIAM, 1987.
17. P. Vanvek, J. Mandel, and M. Brezina, "Algebraic Multigrid by Smoothed Aggregation for Second and Fourth Order Elliptic Problems," *Proc. 7th Copper Mountain Conf. Multigrid Methods*, 1995, pp. 179–196.
18. S. Balay et al., "PETSc Users' Manual," tech. report ANL-95/11, v3.3, Argonne Nat'l Laboratory, 2012.
19. V.E. Henson and U.M. Yang, "BoomerAMG: A Parallel Algebraic Multigrid Solver and Preconditioner," *Applied Numerical Mathematics*, vol. 41, no. 1, 2002, pp. 155–177.
20. S. Molins et al., "Pore-Scale Controls on Calcite Dissolution Rates from Flow-Through Laboratory and Numerical Experiments," *Environmental Science and Technology*, vol. 48, no. 13, 2014, pp. 7453–7460.

David Trebotich is a computational scientist in the Applied Numerical Algorithms Group of the Computational Research Division at Lawrence Berkeley National Laboratory. He's the original developer of the Chombo-based flow and transport solver and the combined Chombo-Crunch code. Trebotich has a PhD in computational science from the University of California, Berkeley. Contact him at dptrebotich@lbl.gov.

Mark F. Adams is a computational mathematician in the Applied Numerical Algorithms Group of the Computational Research Division at Lawrence Berkeley National Laboratory. He's the original developer of the gamg algebraic multigrid code that's native in PETSc. Adams has a PhD in civil engineering from the University of California, Berkeley. Contact him at dptrebotich@lbl.gov.

Sergi Molins is a computational geoscientist in the Earth Sciences Division at Lawrence Berkeley

National Laboratory. He was instrumental in interfacing the CrunchFlow geochemistry module with the Chombo flow and transport solver. Molins has a PhD in computational science from the University of California, Berkeley, and in civil engineering from the University of British Columbia. Contact him at smolins@lbl.gov.

Carl I. Steefel is a geoscientist in the Earth Sciences Division at Lawrence Berkeley National Laboratory. He's the original author of CrunchFlow. Steefel has a PhD in geochemistry from Yale University. Contact him at smolins@lbl.gov.

Chaopeng Shen is an assistant professor of civil engineering at Penn State University. His research interests are large-scale computational hydrology, land surface processes, hydrologic scaling issues, and high-performance subsurface reactive transport modeling. Shen has a PhD in civil and environmental engineering from Michigan State University. Contact him at cshen@engr.psu.edu.



Selected articles and columns from IEEE Computer Society publications are also available for free at <http://ComputingNow.computer.org>.



Call for Articles

IEEE Software seeks practical, readable articles that will appeal to experts and nonexperts alike.

The magazine aims to deliver reliable information to software developers and managers to help them stay on top of rapid technology change.

Submissions must be original and no more than 4,700 words, including 200 words for each table and figure.

**IEEE
Software**

Author guidelines: www.computer.org/software/author.htm
Further details: software@computer.org
www.computer.org/software