

Lawrence Berkeley National Laboratory

LBL Publications

Title

Approaches for the Simulation of Coupled Processes in Evolving Fractured Porous Media Enabled by Exascale Computing

Permalink

<https://escholarship.org/uc/item/13z5z602>

Journal

Computing in Science & Engineering, 26(2)

ISSN

1521-9615

Authors

Molins, Sergi
Trebotich, David
Steefel, Carl I

Publication Date

2024-04-01

DOI

10.1109/mcse.2024.3403983

Peer reviewed

Theme Article: Scientific Impact of the Exascale Computing Project (ECP)

Approaches for the simulation of coupled processes in evolving fractured porous media enabled by exascale computing

Sergi Molins, *Lawrence Berkeley National Laboratory, Berkeley, CA, 94720, USA*

David Trebotich, *Lawrence Berkeley National Laboratory, Berkeley, CA, 94720, USA*

Carl I. Steefel, *Lawrence Berkeley National Laboratory, Berkeley, CA, 94720, USA*

Abstract—Models have historically represented fractured porous media with continuum descriptions that characterize the media using bulk parameters. The impact of small-scale features is not captured in these models, although they may be controlling the performance of subsurface applications. Pore-scale models can simulate processes in small-scale features by representing the pore space geometry explicitly but are computationally expensive for large domains. The alternative multiscale approach entails the combination of pore-scale and continuum-scale descriptions in a single framework. We use Chombo-Crunch, a computational capability that discretizes complex geometries with an adaptive, embedded boundary method to contrast these two approaches. Chombo-Crunch takes advantage of recent computational performance and memory bandwidth improvements resulting from the emergence of exascale computing resources. These combined improvements enable the efficient simulation of reactive transport in fractured media with a high degree of fidelity and the ability to capture the control small-scale processes exert on the overall medium evolution.

The performance of reservoir-scale subsurface storage applications like CO₂ sequestration often hinges on the behavior of small features that make up porous media such as cracks or fractures. For example, thermal-hydrological-geochemical-mechanical (THCM) process coupling in these small-scale features control the leakage of fluids from the subsurface reservoirs where they are injected. Geochemical processes in particular play a significant role in the evolution of fractures through dissolution-driven widening, fracture sealing due to precipitation, and facilitating the formation of fractures during subcritical crack growth. In order to be used for prediction, numerical models must be able to capture the effect of these small-scale processes while reaching the large scales relevant for the application at hand.^{1,2}

Fractured porous media are characterized by their physical and mineralogical heterogeneity at spatial scales from nanometers to meters and beyond. Fluid flow, transport of solutes and geochemical reactions within these media may also operate on a broad range of temporal scales, from seconds to year. A large contrast is often observed between fluid residence times in regions of enhanced permeability such as the fractures themselves or macropores and less permeable media such as the rock matrix. In such systems, a separation of scales is possible such that the different compartments may be considered explicitly in conceptual models. The overall challenge is to implement these conceptual models that couple complex multiphysics occurring at multiple scales, from the micro to the kilometer scale, in a high-resolution simulator. This requires an approach that is flexible enough to account for the different time scales while also performing with computational efficiency to resolve non-linear computations in arbitrarily complex fractured porous media

XXXX-XXX © 2024 IEEE
Digital Object Identifier 10.1109/XXX.0000.0000000

Scientific Impact of the Exascale Computing Project (ECP)

geometries.

Here we describe modeling approaches to simulate flow, transport and reactions in evolving, fractured media enabled by exascale computing capabilities. Pore-scale models separate the pore space and the solid phase, explicitly representing the interfacial processes. While the pore-scale approach is limited by the ability to resolve fine scale heterogeneities in mineral composition, high performance computing makes it tractable to perform simulations of evolving fractured media.

Modeling approaches that combine a pore-scale description and a Darcy-scale description make it possible to relax resolution requirements in the regions of the medium by describing them as porous media.³ This makes them more computationally tractable and can be seen as a compromise between resolution and domain size. These multiscale approaches can be broadly classified into 2 groups. The first group comprises approaches that combine the different scale characterizations in a single equation. The second group comprises approaches that use an equation for each scale characterization considered.

The latter group is broadly termed hybrid multi-scale, and has the key advantage that by splitting the domain into two or more subdomains, it can take into consideration the characteristic time scales of each sub-problem. This can be used to efficiently solve the coupled problem over large domains and long periods while capturing small-scale processes.

In this work, we conceptualize the fractured media domain as composed of a fracture opening and the rock matrix that surrounds it. We present the equations that describe the flow, transport and reactive processes. We start with a pore-scale description, where the fluid in the fracture opening is the continuum medium the equations are written for and the interface with the rock matrix is a solid boundary. We follow with a hybrid multiscale description where the fracture matrix, rather than a solid, is a porous continuum characterized by a heterogeneously distributed porosity and mineral surface area. The terms pore-scale and Darcy-scale are not associated with any specific spatial scale but rather imply a conceptual framework that uses different process descriptions for natural fractured porous media. In other words, they can be applied to the range of spatial scales that characterize these media, although the example presented focuses on a small (2.5-cm by 0.95-cm) experimental fractured core.

PORE-SCALE APPROACH

In the pore-scale approach, the fluid phase is the continuum and the rock-matrix interface is the external

boundary of the domain. Fluid flow is solved within this pore-scale sub-domain with the incompressible Navier-Stokes equations

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

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

where ρ is the fluid density, ν is the kinematic viscosity, ρ is the fluid pressure and \mathbf{u} is the fluid velocity.

The advection-diffusion equation is used to simulate the transport of solutes, or dissolved species, in the same sub-domain

$$\frac{\partial \psi_i}{\partial t} + \nabla \cdot \mathbf{u} \psi_i = \nabla \cdot D_i \nabla \psi_i \quad (i = 1, \dots, N_c) \quad (3)$$

where D_i is the diffusion coefficient of component i and ψ_i is its total concentration. The total concentration of a component is the sum of the individual species concentrations that contain this component (4)

$$\psi_i = \rho \left(c_i + \sum_{j=1}^{N_x} \xi_{ij} m_j \right) \quad (i = 1, \dots, N_c) \quad (4)$$

where N_c is the number of components, N_x is the number of aqueous complexation reactions, c_i and m_j are the concentration of primary species i and secondary species j , and ξ_{ij} is the stoichiometric coefficient of primary species i in reaction j . The law of mass action describes thermodynamic equilibrium between primary and secondary species

$$m_j = (K_j \gamma_j)^{-1} \prod_{i=1}^{N_c} (\gamma_i c_i)^{\xi_{ij}} \quad (j = 1, \dots, N_x) \quad (5)$$

where γ_i and γ_j are the activity coefficients of primary and secondary species, and K_j is the equilibrium constant of reaction j . The rate of mineral dissolution-precipitation is calculated at the rock matrix interface with a transition state theory-type rate law

$$r_k = k_k \prod (\gamma_i c_i)^{n_i} \left(1 - \frac{IAP}{K_k} \right) \quad (6)$$

where k_k is the rate constant, n_i the reaction order exponent, and K_k the equilibrium constant. The mineral reaction can be described as a surface reaction and thus can be expressed as a boundary condition:

$$-\rho D_i \nabla c_i = \sum_{k=1}^{N_m} \chi_k \xi_{ik} r_k \quad (i = 1, \dots, N_c) \quad (7)$$

where N_m is the number of mineral reactions, ξ_{ik} is the stoichiometric coefficient of species i in the




FIGURE 1. Computational and conceptual representation of a 2.52-cm long 0.95-cm wide fractured core,⁴ showing the fracture surfaces that separate the fracture opening (shown in the foreground cutout) and the fracture matrix (shown in light gray) that extends away the fracture surfaces and is a porous medium. In the limit, one can assume that this porous medium has little-to-no-porosity, in which case the overall representation is equivalent to a solid-pore pore-scale model. The aperture of the fracture is variable with the fracture closing in some areas, shown in white.

mineral reaction k , and χ_k is the fraction of the surface composed of mineral k such that $\sum_{k=1}^{N_m} \chi_k = 1$.

The velocity of the interface between the fracture and the matrix (v_n^Γ) can be described by

$$v_n^\Gamma = \sum_k^{N_m} \chi_k V_k r_k \quad (8)$$

where V_k is the molar volume of the k -th mineral. As the geometry evolves, parameters associated with this geometry such as the value of the reactive surface area also change. Equation (8) modifies the geometry of the domain, thus affecting the boundary conditions both for reactive transport (4), as well as for flow:

$$\mathbf{u} = v_n^\Gamma \quad (9)$$

MULTISCALE APPROACH

In the multiscale approach, processes are simulated both in the fracture opening and the porous rock matrix. It is assumed here that flow is negligible in the porous matrix, and therefore only solved in the pore-scale domain with (1, 2). However, the coupling approach used for transport processes described below could be extended for the solution of flow as well if Darcy's law were used for flow in the porous matrix.

Solute transport is solved in the pore-scale sub-domain according to (3) and in the Darcy-scale domain according to the following equation

$$\frac{\partial \theta \psi_i}{\partial t} = \nabla \cdot (\theta \tau D_i \nabla \psi_i) + R_i \quad (i = 1, \dots, N_c) \quad (10)$$

It is important to highlight the differences between (3) and (10). The first difference is that due to the assumption of negligible flow, there is no advection in (10), and the second difference is in the use of bulk parameters and process contributions. The porosity (θ) and tortuosity (τ) of the medium are used in the Darcy-scale description, and so is the source-sink term from the mineral reactions (R_i). Because mineral surfaces are not explicitly captured in the porous sub-domain, mineral dissolution-precipitation is captured as a bulk process

$$R_i = \sum_{k=1}^{N_m} \xi_{ik} A_k r_k \quad (i = 1, \dots, N_c) \quad (11)$$

A Darcy-scale description of mineral dissolution-precipitation captures the changes of the porosity of the porous matrix as the volume fractions of the minerals change

$$\Delta \theta = \sum_{k=1}^{N_m} V_k R_k \quad (12)$$

The evolution of the domain is purely driven by changes in porosity in the matrix, capturing the transport limitations first within the pore space (3) and then through the porous matrix (10). To do so, the

model must capture the continuity of fluxes and concentrations at the interface (Γ) between pore-scale and Darcy-scale sub-domains. This is accomplished with the following conditions

$$C_{i,pore}|^{\Gamma} = C_{i,Darcy}|^{\Gamma} \quad (13)$$

$$D_i \nabla C_{i,pore}|^{\Gamma} = \theta \tau D_i \nabla C_{i,Darcy}|^{\Gamma} \quad (14)$$

In the hybrid multiscale approach, the concentrations at the interface are thus explicitly calculated as unknowns of the problem in the solution approach as detailed in the following section.

SOLUTION AND NUMERICAL APPROACHES

The time scales associated with each of the processes—continuum, Darcy and pore scale—can be significantly different. For carbonate dissolution, under the range of spatial dimensions and inlet fluid velocities in the systems of interest, the rate of the (mineral-fluid) boundary displacement is much slower than that of reactive transport which in turn is much slower than that of flow. We take advantage of this fact to solve these processes separately when possible in order to resolve fine-scale dynamics (e.g. flow in complex fracture geometries) while still reaching the large scales due to mineral morphology changes caused by dissolution or precipitation reactions.

The solution of the coupled pore-scale problem entails the sequential solution of flow, transport, mineral-dissolution precipitation and interface evolution using subcycling. Interface evolution is the slowest process and uses the largest time step Δt . The boundary geometry is updated based on a reaction rate that does not change over the duration of this time step, which is limited by a Courant-Friedrichs-Lewy (CFL) criterion applied to the boundary velocity. The solutions of flow and reactive transport are obtained via sub-time stepping within this Δt . The approach is to solve pore-scale flow in the fracture to obtain a steady-state advective velocity using a transient solver for (1,2). The steady-state assumption is justified due to the much shorter time scales associated with pore-scale flow, with the constraints for time stepping being the viscous time scale in very low Reynolds number, or Stokes, flow regime. Once the steady-state velocity is obtained, the transient reactive transport problem is solved in the fracture, with time stepping constrained by the CFL criterion applied to the fluid velocity. Mineral dissolution-precipitation is solved as part of the reactive transport problem, i.e., the rates calculated at the

interface contribute to the mass balance of the solutes being transported. These rates are then used to determine the boundary displacement that determines the interface evolution, where the new interface location marks the start of the next time step of the algorithm.

The solution of the coupled multiscale problem entails adding the solution of Darcy-scale processes to the pore-scale problem. Similar to the solution approach described above, pore-scale flow is solved in the fracture. For the reactive transport problem, the solution for each time step is obtained by solving sequentially the pore-scale problem and the Darcy-scale problem, updating the boundary conditions applied to the interface between the two sub-domains, and iterating the solution as needed until convergence for these boundary conditions. An iterative method is employed to obtain the solution of the coupled problem, i.e., (3-4)+(6-10)+(13-14). This "flux matching" approach entails solving each sub-problem subject to boundary conditions determined from the solution of the complementary sub-problem. Specifically, Darcy-scale concentrations at the interface are used as Dirichlet boundary conditions for the pore-scale problem, (13). And then, upon solution of the pore-scale sub-problem, concentration gradients are evaluated at the interface and used as Neumann boundary conditions for the Darcy-scale sub-problem, (14). Convergence of the flux matching procedure is used to determine whether the solution of the iterative problem has been achieved. The rates of mineral dissolution-precipitation and the resulting changes of porosity in the porous matrix are solved as part of the Darcy-scale reactive transport problem. Because of this porous-medium treatment, this component of the model is much less computationally demanding than the pore-scale counterpart that requires capturing the geometry explicitly.

The model is implemented in Chombo-Crunch, an exascale simulation capability based on the adaptive, embedded boundary (EB) method for flow and conservative transport⁵ and the geochemistry module of CrunchFlow.⁶ Chombo-Crunch, originally a pore-scale simulator for reactive transport in non-evolving porous media,^{7,8} has become a flexible framework for the simulations in evolving geometries and multiscale domains.^{9,10} This framework allows for subcycling in time and the sequential, iterative solution of the pore-scale/Darcy-scale coupled problem. It also enforces the appropriate boundary conditions on the shared embedded boundary that represents the interface between the two sub-domains. Adaptive mesh refinement (AMR) is used around the interface to match the resolution from the coarser Darcy-scale domain to the pore scale resolution. This approach is computationally

efficient in that it allows for a coarser mesh to be used for the Darcy-scale, continuum domain, where it is appropriate, while fine resolution is applied at the pore scale where it is necessary. The advantage of this adaptive, EB approach is the ability to explicitly resolve fluxes at the pore scale while using AMR only near the boundary to match concentration values and fluxes.

FRACTURED CORE SIMULATIONS

We simulate the fractured core experiment performed by Ajo-Franklin and co-authors.⁴ The sample used for the experiment was a core (0.95 cm in diameter and 2.52 cm long) cut from a piece of dolomite from the Duperow Formation, a regionally extensive reservoir and seal in the Williston Basin. The fractured core was subject to constant flow of CO₂-acidified brine for 113 h under high pressure (pore pressure 96.5 bar) at room temperature. The influent solution was in equilibrium with CO₂ at 35 bar at 25°C, yielding a total carbon concentration of 1.2 mol/kgw. X-ray powder diffraction was used to determine the mineralogical composition of the core, comprising dolomite (86.9 wt %), calcite (9.5 wt %), plagioclase (1.8 wt %), smectite clay (1.8 wt %), and quartz (<1 wt %), and X-ray computed microtomography (or xCMT) images of the fractured core were collected at different stages of the experiment. The experimental results indicated that the increase of fracture aperture was localized with the fracture plane, that a porous altered layer had developed due to rapid calcite reactions, and that the fracture surface had become rougher.^{2,4}

We present two simulations. In the first, a purely pore-scale description of the problem is used where dissolution alters the surface of the fracture, which subsequently becomes rougher and recedes. In the second, the altered layer that develops around the fracture surface due the faster dissolution of calcite is captured with a Darcy-scale model, which adds to the pore-scale model used for the fracture opening, in a hybrid multiscale simulation. In both cases, segmentation by thresholding of the grayscale xCMT images was used to generate the geometry of the fracture surface. In the pore-scale simulation, the initial mineral composition of the surface is also obtained from the grayscale xCMT image. In the multiscale simulation, the composition is assumed uniform in the matrix, according to the X-ray powder composition. Two minerals are included—calcite and dolomite—that dissolve according to transition-state-theory-type rate expressions²

$$r_{CaCO_3} = \left(k_1 \gamma_{H^+} C_{H^+} + k_2 \gamma_{H_2CO_3^*} C_{H_2CO_3^*} + k_3 \right) \left(1 - \frac{Q_{CaCO_3}}{K_{CaCO_3}} \right) \quad (15)$$

$$r_{CaMg(CO_3)_2} = \left(k_4 \gamma_{H^+}^{0.5} C_{H^+}^{0.5} \right) \left(1 - \frac{Q_{CaMg(CO_3)_2}}{K_{CaMg(CO_3)_2}} \right) \quad (16)$$

where K_k is the solubility constant and Q_k is the ion activity product, for reaction k ($CaCO_3$, $CaMg(CO_3)_2$).

Pore-scale simulation

In the pore-scale simulation, the mineralogical composition of the fracture surface was obtained from the xCMT grayscale images by assuming that the core was composed of calcite and dolomite. A threshold value was chosen such that the fractions measured by X-ray diffraction were obtained (see Figure 2a). Calcite represents a small fraction of the surface but is the fastest reacting mineral.

The initial geometry of the domain is shown in Figure 2b along with the concentrations of calcite on the fracture surface. These concentrations can be used to identify the fast flow paths within the fracture plane. Concentrations of calcium in the influent solution are low and increase as the solution flows in the fracture due to dissolution. In slow-flowing regions, the solution can reach equilibrium with the calcite dissolving faster and, therefore, calcium concentrations are greater. Low concentrations, in turn, indicate fast flow paths, which, in this case, appear near the right side of the domain. This is caused by the fracture closing completely on the left side, between sections A and B (Figure 2a), as indicated by the white areas of the domain.

The concentrations in Figure 2b correspond to the steady-state conditions before the geometry evolves. When the reactive transport problem, 3- 7, reaches steady-state the reaction rates are used to move the fluid-solid interface. After simulating dissolution for over 33 hours, the geometry in Figure 2c is obtained. Here, two main aspects can highlighted. One is that the model captures the channelization of the flow. The preferential flow path initially dissolves faster. While the apertures near the inlet (section A) grow similarly across the entire section (Figure 3a), they only grow significantly in localized areas, where the initial aperture was larger (Figure 3b). The second aspect is the increase in surface roughness. This is driven by the mineralogical composition of the matrix (Figure 3c,d).

In spite of the model's ability to capture these key experimental observations, differences are also apparent. In particular, the porous structure of the altered layer is less tortuous (Figure 3) than observed




FIGURE 2. Duperow fractured core pore-scale simulation with dimensions as in Figure 1: (a) 3D depiction of the fracture surface (fracture-matrix interface) showing the initial mineralogical composition, with blue indicating dolomite and yellow indicating calcite, and displaying two cross-sections used in Figure 3; (b) Calcium concentrations on the fracture surface for the initial fracture geometry before any evolution; (c) Calcium concentrations on the fracture surface at 33 hours, showing localized widening of the fracture aperture, including the development of a preferential flow path, and the development of a rougher fracture surface.

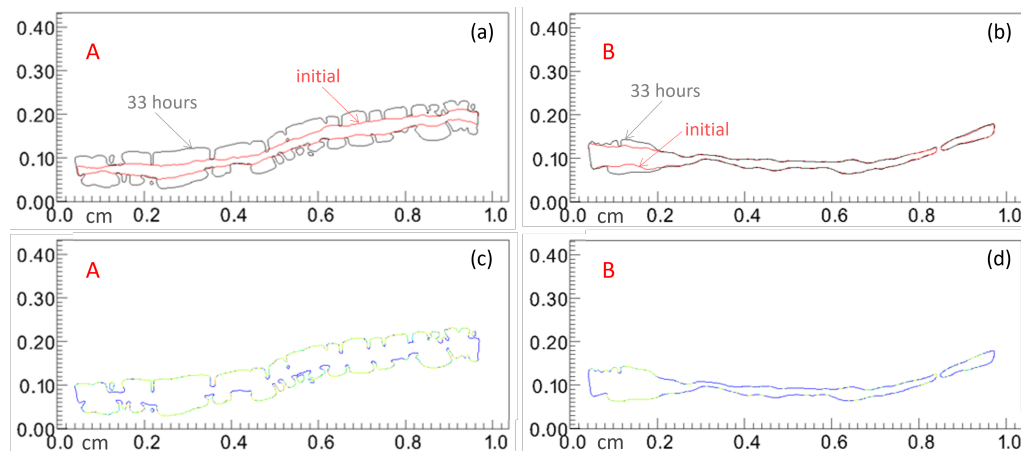


FIGURE 3. Duperow fractured core pore-scale simulation: (a)-(b) Simulated geometry of the fracture initially (red) and at 33 hours (black) at the A and B cross-sections, respectively (Figure 2, showing surface roughening near the inlet (A) and preferential dissolution away from the inlet (B)); (c)-(d) Composition of the fracture surface at 33 hours after the evolution of the fracture surface exposes part of the rock matrix, showing preferential dissolution in the calcite-rich areas.

experimentally.⁴ As a result, diffusion of the reactant and product species to and from the receding calcite surface is relatively unimpeded. Effluent calcium concentrations in the experiment decrease over time reflecting an increase in diffusion limitation to rates,² but this is not captured fully by the pore-scale model, where effluent calcium concentrations remain relatively constant. The pore-scale model, however, is able to capture an early increase in effluent calcium concentration² as it is able to reproduce the early increase in surface area via surface roughening, as noted above.

Multiscale simulation

The inability to capture the evolution of the matrix stems from the lack of resolution of the pore-scale model in capturing the mineral heterogeneity present in the core. This lack of resolution originates in the limited resolution of the xCMT data and is carried in the model. This justifies using a multiscale approach, whereby the fracture matrix is represented as a porous medium. The average mineralogical composition is used to initialize the porosity and bulk surface areas in the Darcy-scale sub-domain, assuming a uniform distribution over the altered layer near the fracture surface (see Figure 4a).

Comparing multiscale model results to experimental data in terms of effluent concentrations, it can be seen that the multiscale model can overcome the two issues in the pore-scale results.³ The multiscale model retains the ability of the pore-scale model to capture the existence of a preferential flow path, i.e., because it is solving the flow field in the complex geometry of the fracture as obtained from xCMT segmentation. This means that the model also captures the transport limitations to rates that develop in the fracture when concentration gradients develop (see Figure 4b). At the same time, the multiscale model also captures the transport limitation that controls effluent calcium concentrations, which take place in the fracture matrix. This is apparent from the concentration gradients shown in Figure 4c, which are mostly perpendicular to the local fracture surface.

Mesh refinement was used in the simulations to match the resolution between the Darcy-scale domain and the pore-scale domain. The pore-scale problem requires fine resolution to capture the complex flow dynamics with the Navier-Stokes equations 1-2, while the Darcy-scale model requires fine resolution where concentration gradients develop so as to accurately capture the diffusion limitation discussed above. Here, AMR has dual use in that, not only does it facilitate the discretization to match boundary fluxes and values along the interface, but it is also used in the traditional sense to refine in other areas of the domain where extra gridding is needed, namely, near the surface and where concentration gradients were steep (Figure 4c) while leaving the rest of the domain coarse as is appropriate. In this simulation, this implies that AMR naturally captures the thickness of the altered layer in the porous matrix, which varies over space (see Figure 4c).

Exascale capability

With the emergence of exascale computing resources as a result of the Exascale Computing Project (ECP), Chombo-Crunch is able to take advantage of performance and memory bandwidth improvements afforded by accelerator-based computing architectures. The flow and transport solvers that support the Chombo-Crunch suite of application codes have undergone performance engineering and portability for GPU-based exascale computing resources. Similarly, the Chombo-PETSc interface, with access to hypre, has adapted to this paradigm, enabling flow and transport simulations in heterogeneous media at unprecedented scale and resolution on the Oak Ridge Leadership Class Facility (OLCF) exascale machine, Frontier.¹¹ The pore scale

Chombo-Crunch application code has been scaled to full machine on Frontier (8192 nodes) to run the baseline exascale challenge problem of fracture evolution due to flow, transport, reactions and mechanics on domains up to 10 cm long ideal fractures. The pore-scale simulations presented here represent portions of the full exascale challenge problem based on available image data for the purpose of visualization of scientific results. Resolution limitations with the original experimental data did not warrant increasing model resolution. However, this would be possible with the GPU-based computational resources. Hence, these simulations were run on up to 1000s of nodes, not full machine. While fine pore-scale resolutions are now possible, resolution would impose additional time stepping constraints, as noted earlier. In general, we experience 2.5x performance improvement of the GPU over the CPU, theoretically allowing for doubling the resolution without increasing run time.

Furthermore, the memory footprint and bandwidth improvements on GPU-based machines are evident in our computations on Frontier. Previous state-of-the-art for pore-scale simulations was load balancing of one domain-decomposed box, or patch (32,768 grid cells per box), per core per MPI rank on CPU-based machines. The Frontier supercomputer is a GPU-based architecture with 9,402 compute nodes, each node containing 4 AMD MI250X GPUs, with 2 Graphics Compute Dies (GCDs) per GPU for a total of 8 GCDs per node, as well as 64 AMD EPYC CPU cores per node. This architecture allows for load balancing of 64 boxes per GCD, with 1 MPI rank per GCD, due to significant gains in memory bandwidth on GPU-based nodes over CPU-based nodes. The overall result is additional capability to simulate over 400 cm³ of porous medium volume on GPUs compared to on the order of 10 cm³ on CPUs.¹¹ The significance of this gain in capability on GPUs is a new regime of brute force pore-scale simulation of flow and transport with resolved reactions on much larger domains, challenging conventional understanding of what it means to be a representative elementary volume (REV) at the Darcy scale. We note that similar memory bandwidth improvement is available on other GPU-based machines (e.g., NERSC Perlmutter). In other words, exascale resources enable us to tackle much larger pore-scale domains than previously.

CONCLUSION

Resolution requirements to capture small-scale processes in fractured porous media, such as those described by pore-scale models, are generally very strict.

Scientific Impact of the Exascale Computing Project (ECP)




FIGURE 4. Duperow fractured core multiscale simulation with dimensions as in Figure 1: 3D depiction of the fractured Duperow core represented as the combination of (b) a pore scale sub-domain and (c) a Darcy-scale porous matrix, showing the mineralogical composition of the matrix in (a), where light gray indicates the altered layer, which is depleted in calcite and has a porosity of about 10%. Simulated calcium concentrations at 80 hours are shown in (b) for the fracture sub-domain and in (c) for the matrix sub-domain.

In the Duperow pore-scale simulation, where the scale of mineral heterogeneity falls well below the resolution of the model or below the resolution of the characterization technique, capturing the evolution accurately is challenging. To some extent, this limitation is stronger in the characterization methods, where there is a trade-off between field-of-view and resolution. The limitation for numerical models is one of computational cost. With the emergence of the new exascale computational capabilities, however, the ability to resolve small scale features in larger domains has improved substantially.

These limitations are not exclusive of pore-scale models. Most applications must answer questions at much larger scales, e.g. reservoir scales, where relatively small features can determine the success or failure of the application. For example, enhanced geothermal systems rely on creating permeability (thus, enabling flow) by injecting fluids to develop distributed fracture networks (e.g. Figure 5), and carbon sequestration reservoirs must be overlain by formations that do not have or develop leakage pathways. Multiscale models not only provide a compromise for computational cost and performance, but also the flexibility to simulate different sub-problems separately, but coupled explicitly. This allows, in practice, for the use of complementary models with different time steps and spatial resolutions in each sub-domain. This is not limited to different process descriptions (e.g. pore-scale vs Darcy-scale), and is generally not possible in models that combine multiscale processes in a single equation such as the Darcy-Brinkman-Stokes.¹² This feature of the multiscale model presented here could be especially useful when considering larger domains




FIGURE 5. Conceptual extension of the hybrid multiscale model to capture fracture networks, bringing the embedded boundary method to explicitly capture fracture intersections, variable apertures, rough fracture surfaces and the residual proppant spherules of $250 \mu\text{m}$ in diameter that keep the fractures open and sustain permeability in enhanced geothermal applications.

in which small features, that require high resolution and small time steps, control overall behavior but occupy only a small portion by volume of the larger porous domain. In this view, the multiscale approach described here enables bridging spatial scales and process representations where needed or warranted.

Ultimately, it is the combination of significant improvements in computational capability with multiphysics capability that represents a consequential step forward in the modeling of coupled processes in com-

plex media for subsurface applications and beyond. We can envision that in the near future it will be possible to simulate domains with multiple regions, each described with different scale or process models, and using different spatial resolutions with or without adaptive mesh refinement. An example would be a porous media intersected by multiple fractures and subject to coupled THCM processes, where flow takes place preferentially in the fractures, while the generation, growth or sealing of these fractures is controlled by chemical-mechanical processes at contact points, and heat is transferred by the intact mass of rock matrix (see Figure 5). While this work focused on exascale capability, conceptually the multiscale-multiphysics framework is broadly applicable to the development of computational models for a range of computational platforms.

ACKNOWLEDGMENTS

This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231 and resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725.

REFERENCES

1. K. T. B. MacQuarrie and K. U. Mayer, "Reactive transport modeling in fractured rock: A state-of-the-science review," *Earth-Science Reviews*, vol. 72, no. 3, pp. 189–227, Oct. 2005. [Online]. Available: [dx.doi.org/10.1016/j.earscirev.2005.07.003](https://doi.org/10.1016/j.earscirev.2005.07.003)
2. H. Deng, S. Molins, C. Steefel, D. DePaolo, M. Voltolini, L. Yang, and J. Ajo-Franklin, "A 2.5d Reactive Transport Model for Fracture Alteration Simulation," *Environmental Science & Technology*, Jun. 2016. [Online]. Available: [dx.doi.org/10.1021/acs.est.6b02184](https://doi.org/10.1021/acs.est.6b02184)
3. S. Molins and P. Knabner, "Multiscale Approaches in Reactive Transport Modeling," *Reviews in Mineralogy and Geochemistry*, vol. 85, no. 1, pp. 27–48, Sep. 2019. [Online]. Available: [dx.doi.org/10.2138/rmg.2019.85.2](https://doi.org/10.2138/rmg.2019.85.2)
4. J. Ajo-Franklin, M. Voltolini, S. Molins, and L. Yang, "Coupled Processes in a Fractured Reactive System," in *Geological Carbon Storage*. American Geophysical Union (AGU), 2018, pp. 187–205. [Online]. Available: [dx.doi.org/10.1002/9781119118657.ch9](https://doi.org/10.1002/9781119118657.ch9)
5. D. Trebotich and D. Graves, "An adaptive finite volume method for the incompressible Navier–Stokes equations in complex geometries," *Communications in Applied Mathematics and Computational Science*, vol. 10, no. 1, pp. 43–82, 2015. [Online]. Available: [dx.doi.org/10.2140/camcos.2015.10.43](https://doi.org/10.2140/camcos.2015.10.43)
6. C. I. Steefel, C. a. J. Appelo, B. Arora, D. Jacques, T. Kalbacher, O. Kolditz, V. Lagneau, P. C. Lichtner, K. U. Mayer, J. C. L. Meeussen, S. Molins, D. Moulton, H. Shao, J. Šimůnek, N. Spycher, S. B. Yabusaki, and G. T. Yeh, "Reactive transport codes for subsurface environmental simulation," *Computational Geosciences*, vol. 19, no. 3, pp. 445–478, Sep. 2014. [Online]. Available: [dx.doi.org/10.1007/s10596-014-9443-x](https://doi.org/10.1007/s10596-014-9443-x)
7. S. Molins, D. Trebotich, C. I. Steefel, and C. Shen, "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, Mar. 2012. [Online]. Available: [dx.doi.org/10.1029/2011WR011404](https://doi.org/10.1029/2011WR011404)
8. D. Trebotich, M. F. Adams, S. Molins, C. I. Steefel, and C. Shen, "High-Resolution Simulation of Pore-Scale Reactive Transport Processes Associated with Carbon Sequestration," *Computing in Science & Engineering*, vol. 16, no. 6, pp. 22–31, 2014. [Online]. Available: [dx.doi.org/10.1109/MCSE.2014.77](https://doi.org/10.1109/MCSE.2014.77)
9. S. Molins, D. Trebotich, G. H. Miller, and C. I. Steefel, "Mineralogical and transport controls on the evolution of porous media texture using direct numerical simulation," *Water Resources Research*, vol. 53, no. 5, pp. 3645–3661, 2017. [Online]. Available: [dx.doi.org/10.1002/2016WR020323](https://doi.org/10.1002/2016WR020323)
10. S. Molins, D. Trebotich, B. Arora, C. I. Steefel, and H. Deng, "Multi-scale Model of Reactive Transport in Fractured Media: Diffusion Limitations on Rates," *Transp Porous Med*, vol. 128, pp. 701–721, 2019. [Online]. Available: <https://doi.org/10.1007/s11242-019-01266-2>
11. D. Trebotich, "Exascale Computational Fluid Dynamics in Heterogeneous Systems," *Journal of Fluids Engineering*, vol. 146, no. 4, p. 041104, 02 2024. [Online]. Available: [dx.doi.org/10.1115/1.4064534](https://doi.org/10.1115/1.4064534)
12. S. Molins, "Reactive interfaces in direct numerical simulation of pore scale processes," *Rev Mineral Geochem*, vol. 80, pp. 461–481, 2015. [Online]. Available: [dx.doi.org/10.2138/rmg.2015.80.14](https://doi.org/10.2138/rmg.2015.80.14)

Sergi Molins is a Staff Scientist in the Energy Geo-

Scientific Impact of the Exascale Computing Project (ECP)

sciences Division at Berkeley Lab with extensive experience in the field of reactive transport. He uses reactive transport models to address problems related to water, energy and the environment in applications in natural and engineered systems. Model development in Chombo-Crunch is an integral part of his work. Contact him at smolins@lbl.gov

David Trebotich is a Staff Scientist in the Applied Math and Computational Research Division at Berkeley Lab. His research involves end-to-end development of high resolution algorithms for complex flows in multi-scale, multiphysics systems using adaptive, finite volume methods. He is the lead developer of Chombo-Crunch and the Subsurface Exascale Project. Contact him at dptrebotich@lbl.gov.

Carl I. Steefel is a Senior Scientist in the Energy Geosciences Division at Berkeley Lab, and the Principal Investigator of the Subsurface Exascale Project. He has over 30 years of experience in developing models for multicomponent reactive transport in porous media and applying them to topics in reactive contaminant transport and water-rock interaction. He is an American Geophysical Union Fellow and the original author of CrunchFlow. Contact him at cisteefel@lbl.gov.