APPLIED MATH CENTER INVESTMENTS-SUCCESS AT SCALE

Office of Advanced Scientific Computing Research

Abani K Patra

ADVANCED SCIENTIFIC COMPUTING ADVISORY COMMITTEE

April 18-19, 2017

Supporting **Applied** Mathematics Research

- ***DOE/ASCR****: Support the* ***research and development*** *of applied mathematical* ***models, methods and algorithms*** *for understanding natural and engineered systems related to DOE's mission* ***with a focus on***
  - *discovery of new applied mathematics, for the ultra-low power, multicore-computing, and data-intensive future;*

- **Traditional Mathematics Research Support**
  - "Individual" Awards →Base program
  - Centers – facilitation of research in the domain …
  - *Large multi-year integrated efforts of multiple investigators, multiple institutions and larger focus areas, …*

Mathematical Multifaceted Integrated Capability Centers (MMICCS) -- New Paradigm in 2012

Long-term goals:

- Mathematics research that **5-10+ years out will impact DOE mission efforts**: DOE Applications, SciDAC Program, and Exascale Co-Design
- New Mathematical Multifaceted Integrated Capability Centers (MMICCs) directly enhances impact of applied math on DOE mission

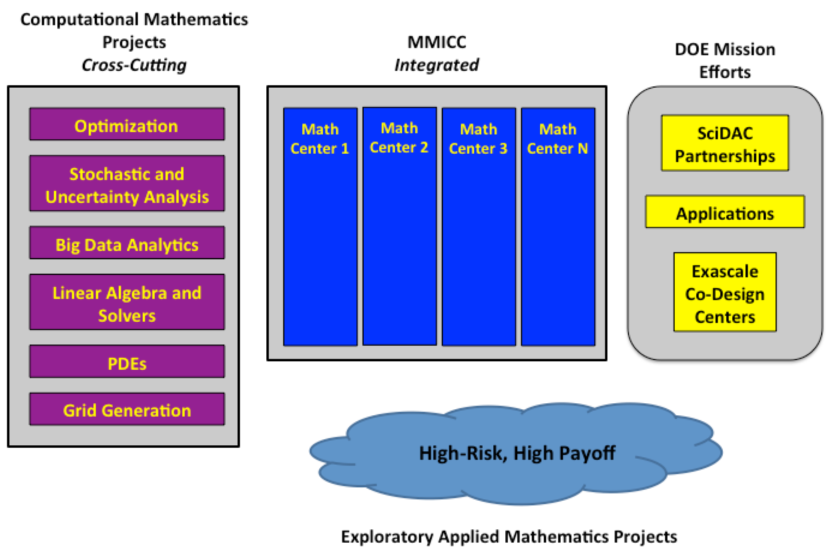

- **Cross-cutting mathematics projects**: addresses foundational, algorithmic and extreme-scale mathematical challenges
- **High-risk, high-payoff**: new mechanism to bring in highly innovative research

Mathematical Multifaceted Integrated Capability Centers (MMICCs)

Mathematical Multifaceted Integrated Capability Center must:

- Address the *long-term mathematical challenges* for one or more DOE grand challenges and that require new integrated, iterative processes across multiple mathematical disciplines.

- Identify *a set of interrelated mathematics research challenges that represent abstractions* of the grand challenges. These abstractions would then be optimally addressed through a multifaceted, integrated approach.

- Have *impact* to the DOE mission *in the 5-10+ year timeframe*

Mathematical Multifaceted Integrated Capability Centers (MMICCs)

MMICCs Organization:

- Integrated collection of sub-projects at multiple laboratory and university sites; Consistent with major theme.

- Center Director :
  - Provide overall direction for the center ensuring internal coordination and collaboration as well as appropriate external outreach; Lean management structure

- Must be sufficiently **flexible** to adapt to changing technical challenges and scientific needs.

- Identify key senior personnel; Sub-project goals and outcomes; integrative mechanisms; outreach to communities beyond the center

Portfolio Summary

- **Projects with Broad DOE mission relevance:**
  - **Proposals in** Complex Energy Systems; Subsurface Flows; Materials; Data driven methods …
  - **~50 preproposals → 14 full proposals → 3 awards**
  - **Anitescu:** ***The Multifaceted Mathematics Center for Complex Energy Systems (M2AC2S)*** : Complex energy systems such as power grid and renewables integration
  - **Karniadakis:** ***Modeling Mesoscale Processes of Scalable Synthesis*****:** Mesoscale modeling applicable to materials, chemistry, and biofuels
  - **Ghattas & Willcox:** ***DiaMonD: An Integrated Multifaceted Approach to Mathematics at the Interfaces of Data, Models, and Decisions*** : Multiscale, multiphysics challenges related to subsurface flows and materials for energy storage
  - Total expenditures of ~$5M/year in laboratories and $4M/year in universities

Mathematical Multifaceted Integrated Capability Centers (MMICCs)

<u>Goal:</u>
Address long-term mathematics research challenges with impact to the DOE mission *in the 5-10+ year timeframe.*

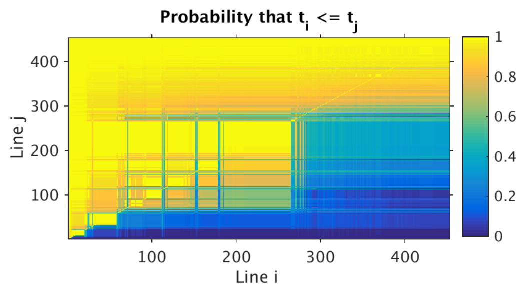

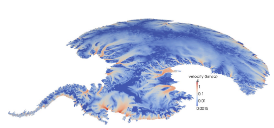

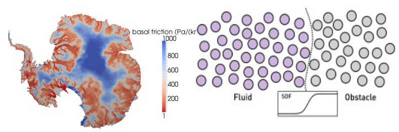

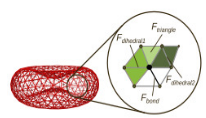

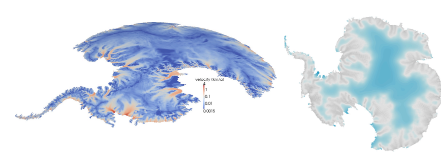

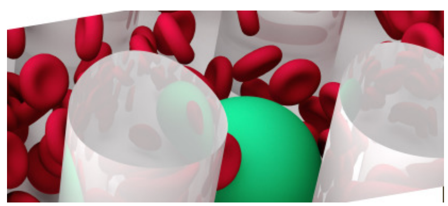

MMICCS have been very successful and can serve to anchor DOE Office of Science investments in Applied Math.

<u>Outcomes/Performance Measures:</u>
~(86+95+148)=~329 publications in peer reviewed literature in 4 years!

Greater than 30 faculty and lab researchers trained.

Gordon Bell, SIAM fellows,ECRP, keynotes, SIAM Best poster…

<u>Long-Term DOE Impact:</u>
New mathematics at the intersection of multiple mathematical sub-domains – data driven discovery, multi-scale modeling, grid optimization, large scale inversion, rare events …

- Several high impact “application transitions” − Grid Modernization Laboratory Consortium (GMLC), Exascale Application Project, partnerships with Center for Integrated Nanotechnologies, Material Synthesis and Simulation Across Scales …
- DOE’s Quadrennial technology review (QTR) feature.

Mathematical Multifaceted Integrated Capability Centers (MMICCs)

Program Review:

Program went through careful review process in Oct-Nov 16.

- 1.5 day reverse site visit format with detailed report prior to review and written Q&A
- 1 day study group with selected researchers, PIs and other agency program managers on best practices and lessons learnt in such focused investments for Applied Mathematics

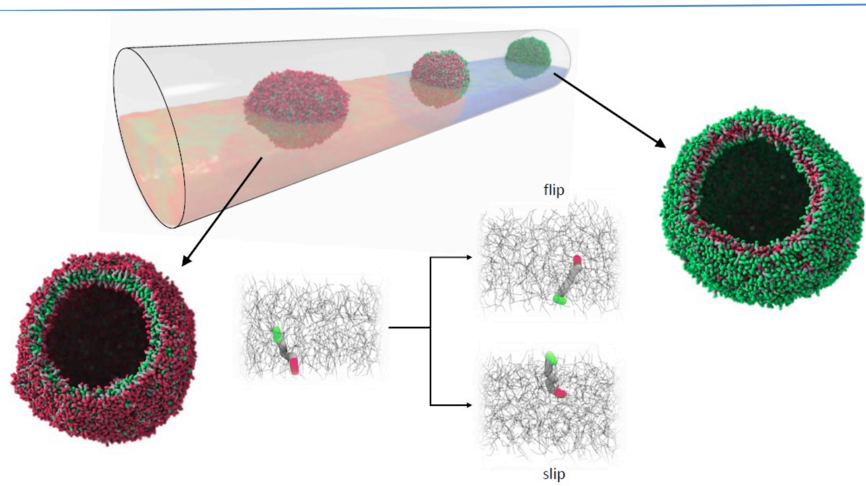

Outcomes:

Reviews and study group were very positive on the program and in addition to the many research highlights, notes on best practices and kudos, suggestions for improvement commented thus:

"The collective group of people involved in this round of MMICCs **would never had embarked on this successful line of research without the MMICCs program**."

Long-Term DOE Impact:

Core groups of organized researchers from multiple laboratories and leading university groups with great intellectual ability, diverse skills and experience levels have assembled to successfully, tackle grand challenge problems.

Challenge is to sustain and adapt groups to evolving needs in mission related research.

Interesting and complex pattern of interactions.   M2ACS vs CM4

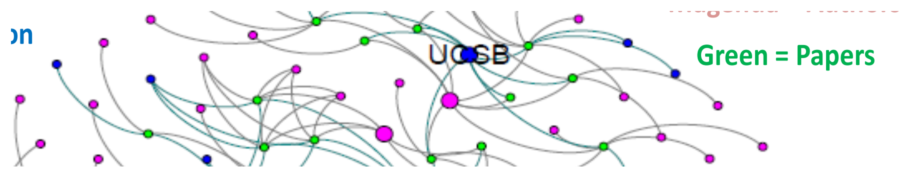

HIGHLIGHTS

M2ACS: Multifaceted Mathematics for Complex Energy Systems

Application subchallenges guiding integrative mathematics

- Integration is guided by application challenges that *span representatively both the domain set and the difficulty set addressed by mathematical themes*:
  - ~~Integrative Math via Application Challenge 1: Temporal/Network-Spatial Multiscale Approaches For High Impact Power Electronic Grid Controllers, Shuai Lu~~
  - Integrative Math via Application Challenge 2: Probabilistic Modeling for Complex Energy Infrastructure, Henry Huang
  - Integrative Math via Application Challenge 3: Model-driven boundary conditions for future US energy infrastructure, Chris De Marco,
  - Integrative Math via Application Challenge 4: Infrastructure - Decision Interdependencies in Electric Energy, Air Quality, Water, B. Lesieutre.

Mathematics Themes

- predictive modeling, a bottom-up theme
  - dynamics and stochastics,
  - Lyapunov function structure,
  - Stability,
  - Graph tools and concepts;

- scalable algorithms
  - dynamic simulation,
  - multi-level methods,
  - decomposition methods

- mathematics of decision, a top-down theme
  - hierarchical optimization,
  - equilibrium problems,
  - mathematics of data,
  - online optimization

- integrative mathematical frameworks – approximation,
  - relaxation,
  - rare event simulation
  - model reduction
  - Integrative frameworks

Fundamental (and unexpected?) Math Insights Stemming from Integrative Math Contemplation

- Identification of Stochastic Resonance Phenomenon in Power Systems (Tartakovsky).
- “First-principles” demonstration that cascade failures occur in groups of lines (observed in 2003 blackout; Weare; figure)
- Provably Exponentially Accurate Temporal Decomposition for long horizon problems (LQR w. bounds; Anitescu)

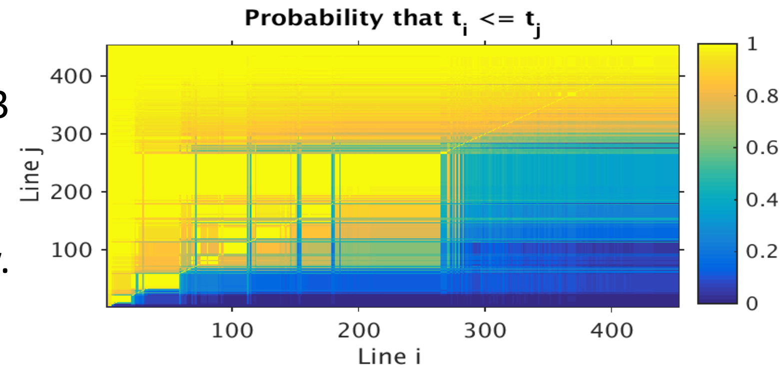

- COAP 2013 best paper prize. (Fall 2014; out of 91 eligible papers Miles Lubin, an alum of pre-M2ACS, Petra, Anitescu– for parallel simplex)
- 4 personnel with Early Career
- Grid Modernization Laboratory Consortium (GMLC)
- "Optimizing Stochastic Grid Dynamics at Exascale” was recommended in the Seed Exascale Application category of the Exascale Computing Project (ECP).

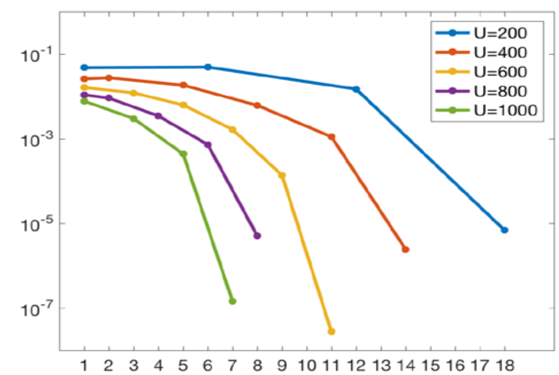

DOE's Quadrennial technology review (QTR) feature of PIPS ALCF runs.

- Three types of solvers (for stochastic programming; Petra, Lubin, Zavala, Chiang; Anitescu as initiator, cheerleader, and PR)
  - PIPS-IPM: Quadratic Program, Parallel Interior Point (PIP)
  - PIPS-S: Parallel Simplex
  - PIPS-NLP: Nonconvex, nonlinear programming, PIP
  - Featured in 2015 QTR ("mathematics" was mentioned 7 times, 2 for this)
  - *The connection of that work to M2ACS is explicitly stated.*

Improving the Energy Grid

The electrical grid has been described as "the largest and most complex machine ever made."[148] Accurately simulating this system requires combining the behavior of millions of consumers, the operation of thousands of power plants, weather events, and the decision-making processes of the utilities themselves. Simulating a system with this level of complexity requires high-performance computing. Accurate grid simulation has become even more complex due to changes in the grid, such as the increasing use of weather-dependent solar and wind resources, and sophisticated and highly localized, high-speed decision making at the consumer level. The complexity and range of conditions required for these simulations require stochastic optimization, where the response of the grid to a large sample of random inputs is computed.

High-performance computing can be used to address a key challenge in planning for the future of the electric grid: increasing penetration of wind and solar energy resources. All power plants, conventional or renewable, are subject to outages or changes in power, requiring reserves and other power sources that can be ramped up or down quickly.[149] These changes in output are both more frequent, and less predictable, for weather-dependent renewables such as solar and wind energy. Because reserves are expensive to maintain and operate, finding the minimum required reserves for the expected penetration of these technologies is crucial to affordable deployment.

In 2012, an INCITE-supported team led by ANL used ALCF supercomputing capabilities to demonstrate that up to 20% wind penetration could be accommodated on some configurations without the need for a significant increase in reserves (Figure 9.14).[150] This result showed that new reserves would not be needed to prepare for increased penetration of wind resources, removing another impediment to greater adoption.

Figure 9.14 The features and implied energy prices of the stochastic programming formulation are shown for the state of Illinois. The model contains approximately 2,000 transmission nodes, 2,500 transmission lines, 900 demand nodes, and 300 generation nodes. The needs to be considered over twenty-four successive hourly time periods can reach billions of variables and constraints once the uncertainty in the supply is taken into account.

Credit: Argonne National Laboratory

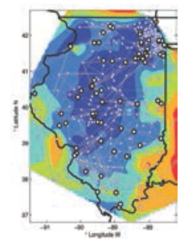

These results could only be obtained using the newer stochastic methods, and demonstrate the benefits of improved computational tools for grid simulation. SC-ASCR has continued work in this area through the Multifaceted Mathematics for Complex Energy Systems (M2ACS) project, which includes researchers from ANL, PNNL, SNL, the University of Wisconsin, and the University of Chicago.[151] New grid simulation capabilities can be used to plan for the future of the grid, develop new operational approaches, and predict the impact of grid disruptions due to physical and cyber attacks and natural disasters.[152]

Parallel distributed-memory simplex for large-scale stochastic linear programming

Argonne NATIONAL LABORATORY | Mathematics and Computer Science

The Challenge

- Simplex algorithms are a key computational tool in optimization.
- They are important both for solving linear programs and linear mixed-integer programs
- However, their reliance on direct linear algebra and their very sequential nature makes them very difficult to parallelize.
- There have been very few algorithms proposed for scalable, distributed memory simplex.

The opportunity/Novel Ideas

- In energy systems, stochastic dispatch and relaxation of stochastic unit commitment needed large-scale simplex methods.
- To solve them at scale, we developed a new parallel approach for the revised simplex method for dual block-angular linear programs.
- The key observations are that the (a) nonsquare scenario components of the basis matrix can be factored in parallel and (b) subsequently, a permutation makes the lower right block invertible.

Key idea: Permuted Pattern (b)

- Basis matrix after scenario factorizations, the permutation makes the lower right block invertible.

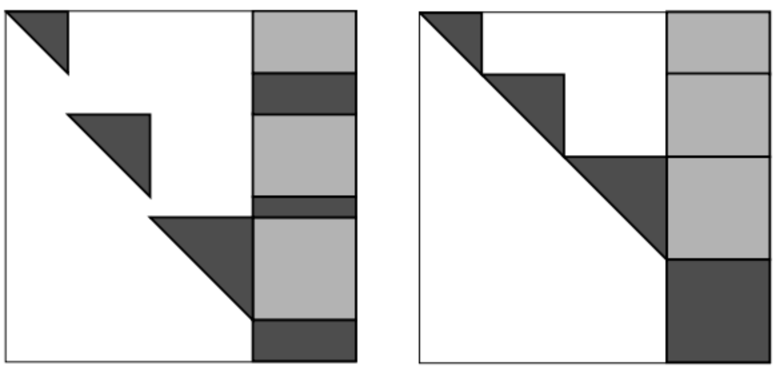

Principal Investigator(s): Mihai Anitescu, ANL

Impact

- We solve relaxation of 12-hour unit commitment problems (with 8,192 scenarios, 463,113,276 variables and 486,899,712 constraints) in less than 5 hours on a BG/P architecture.
- This is beyond what commercial simplex solvers could handle.
- The paper* describing this work has received the best 2013 paper award from the Computational Optimization and Application Journal (in 2014, out of 93 eligible entries)

U.S. DEPARTMENT OF ENERGY | Office of Science

* Lubin M, Hall JJ, Petra CG, Anitescu M. Parallel distributed-memory simplex for large-scale stochastic LP problems. Computational Optimization and Applications. 2013 Jul 1;55(3):571-96.

CM4 – Overview and interactions

- Stochastic Lagrangian equations (Mori-Zwanzig & DPD) for mesoscale reactive and charged transport
- Compatible meshless methods with spectral-like properties
- Renormalized Mori-Zwanzig formulation for reduced-order modeling
- Unified theory for the design and analysis of algebraic multigrid methods

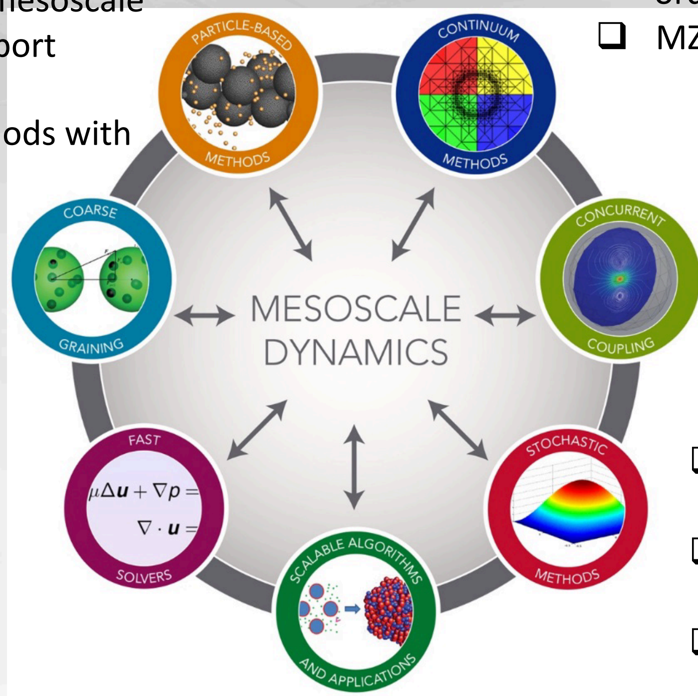

- Immersive boundary method of arbitrary high-order on unstructured grids
- MZ-based refinement

- Concurrent coupling of multi-fidelity models
- Multiscale universal interface algorithm and software

- Stochastic Eulerian-Lagrangianmethods for fluid-structure interactions
- Uncertainty quantification for molecular systems
- Concurrent sampling algorithm for molecular systems

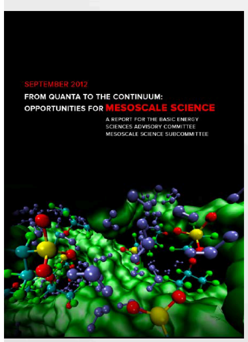

- Multi-fidelity framework & optimization of lithium-air battery
- Universal mesoscopic model of surface tension

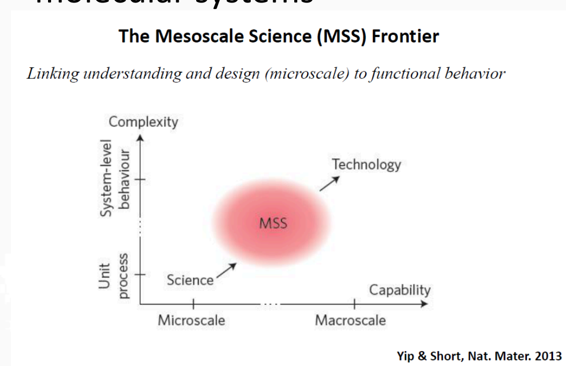

science push w/ technology pull

Meshless surface physics: solving PDE on manifolds using GMLS

N.Trask (SNL), P. Atzberger, B. Gross (UCSB), M. Maxey (Brown)

How to generate local high-order finite difference-like stencils for surface derivatives:

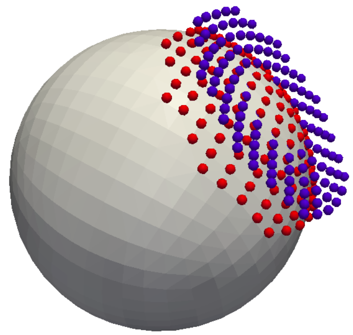

From points in neighborhood, locally parameterize manifold over tangent space using GMLS to approximate metric tensor

$$\Gamma(\chi_1,\chi_2) = \langle \chi_1, \chi_2, p^*(\chi_1,\chi_2)\rangle$$
$$p^* = \mathrm{argmin}_{p\in P_2} \sum_j ||\Gamma - \mathbf{x}_j||^2 W_{ij}$$
$$\mathbf{G}_{ij} = \partial_{\chi_i}\Gamma \cdot \partial\chi_j \Gamma$$

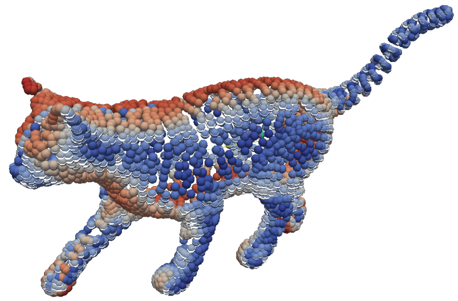

Use local reconstruction of metric tensor to get stencil approximating surface derivatives

$$\Delta_{\mathcal{M}} u = \frac{1}{\sqrt{det(G)}} \partial_{x_i} \sqrt{det(G)} \mathbf{G}_{ij} \partial_{x_j} u_j$$
$$= \sum_j \alpha_{ij} u_j$$

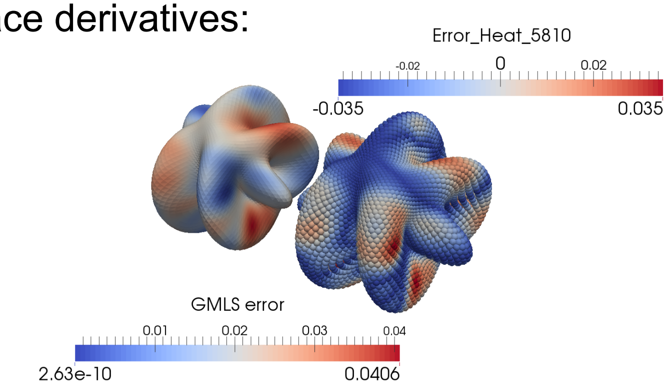

Accuracy competitive with spectral methods[1], particularly for surfaces with high curvature. O(N) solves using standard AMG

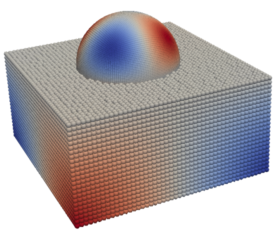

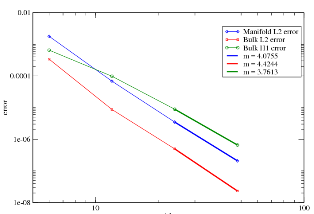

$$\begin{aligned} \partial_t u - \nabla\cdot\kappa\nabla u + u &= f \quad \text{in } \Omega \\ \alpha u - \beta v + \partial_n u &= 0 \quad \text{on } \mathcal{M} \\ \partial_t v - \Delta_{\mathcal{M}} v + \partial_n u + v &= g \quad \text{on } \mathcal{M}. \end{aligned}$$

High-order convergence when coupling bulk to manifold diffusion processes

Target applications:

Mesoscale problems involving surfaces physics + deformation of interface

Meshless → easy to handle complex moving interfaces

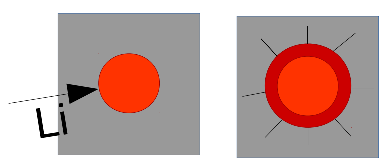

Li-ion battery failure: surface deposition at interface during lithiation drives fracture process

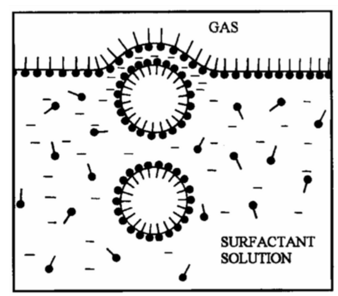

Surfactant transport driving bubble dynamics (Brown)

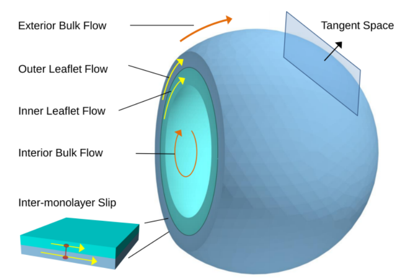

Hydrodynamics within lipid bilayer as vesicle deforms under flow (UCSB)

1. “Spectral Numerical Exterior Calculus Methods for Differential Equations on Radial Manifolds” B. Gross P.J. Atzberger (arxiv.org/pdf/1703.00996.pdf)

Integrated Mathematical Approaches

ISPH2: Parallel implicit SPH implementation using LAMMPS/Trilinos

Improved SPH method: standard SPH fails to converge under refinement

Used for realistic problem on complex geometries

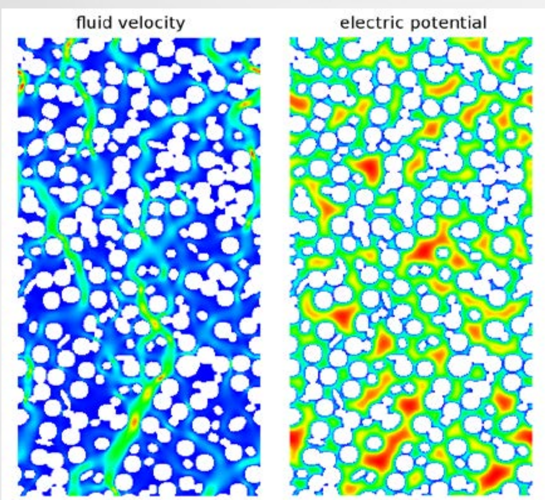

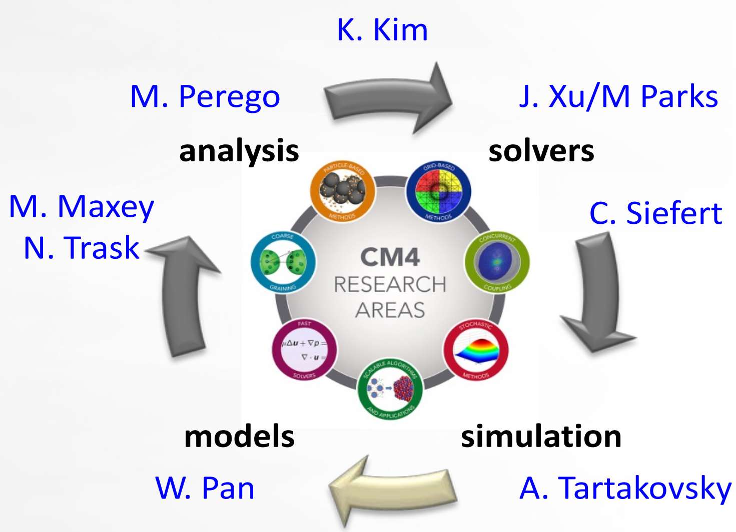

Very good scalability up to 134M particles on 32K cores (largest implicit SPH run ever)

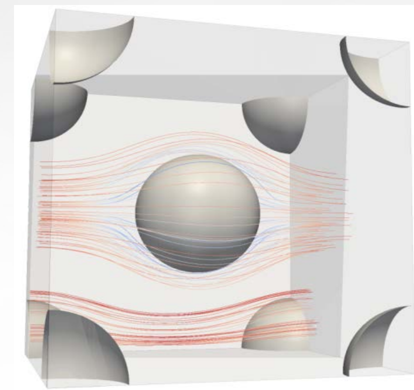

Trask, Maxey, Kim, Perego, Parks, Yang and Xu. Res, CMAME 2014.

Sandia National Laboratories

Pacific Northwest NATIONAL LABORATORY

BROWN

PENNSTATE 1855

CM4 – Integrative Applications

Multi-fidelity Modeling & Optimization of Li-O_2 Batteries

Objective: Maximize discharge capacity

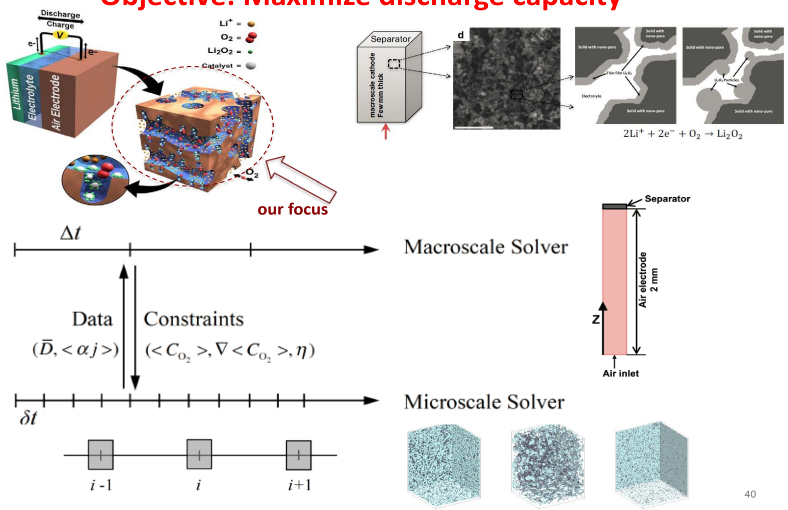

Multi-fidelity modeling

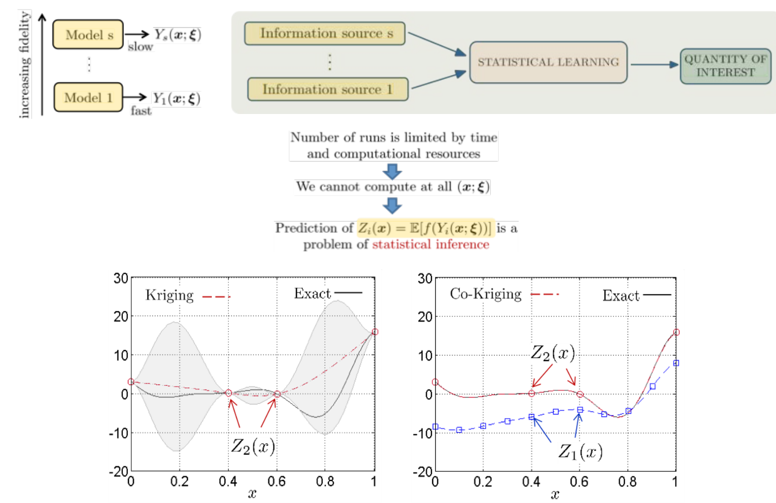

Nonlinear Multi-Fidelity Surrogates via Deep Learning

Pacific Northwest NATIONAL LABORATORY

We generalize the classical linear scheme of Kennedy and O'Hagan

AR1

$$f_2(\boldsymbol{x}) = \rho f_1(\boldsymbol{x}) + \delta_2(\boldsymbol{x})$$
$$f_1 \sim \mathcal{GP}(f_1|0, K_1(\boldsymbol{x}, \boldsymbol{x}'; \theta_1))$$
$$\delta_2 \sim \mathcal{GP}(\mu_{\delta_2}, K_2(\boldsymbol{x}, \boldsymbol{x}'; \theta_t))$$

to a *compositional and robust* model inspired by *deep learning* that can learn complex *nonlinear and space-dependent cross-correlations*

NARGP

$$f_2(\boldsymbol{x}) = g_2(\boldsymbol{x}, f_1(\boldsymbol{x}))$$
$$f_1 \sim \mathcal{GP}(f_1|0, K_1(\boldsymbol{x}, \boldsymbol{x}'; \theta_1))$$
$$g_2 \sim \mathcal{GP}(g_2|0, K_2((\boldsymbol{x}, f_1(\boldsymbol{x})), (\boldsymbol{x}', f_1(\boldsymbol{x}')); \theta_2))$$
$$K_2((\boldsymbol{x}, f_1(\boldsymbol{x})), (\boldsymbol{x}', f_1(\boldsymbol{x}')); \theta_2)) = k_{t_\rho}(\boldsymbol{x}, \boldsymbol{x}'; \theta_{t_\rho})) \times k_{t_f}(f_1(\boldsymbol{x}), f_1(\boldsymbol{x}'); \theta_{t_f}) + k_{t_\delta}(\boldsymbol{x}, \boldsymbol{x}'; \theta_{t_\delta})$$

Space-dependent nonlinear map from the low- to the high-fidelity model

*example for 2 levels of fidelity — extension to more levels and deeper networks is straightforward.

P. Perdikaris, M. Raissi, A. Damianou, N.D Lawrence, G. E Karniadakis, Nonlinear information fusion algorithms for robust multi-fidelity modeling, 2016.

42

Device-scale Performance and Design of Electrode Structures

J. Phys. Chem. C, 119, 14851–14860, 2015

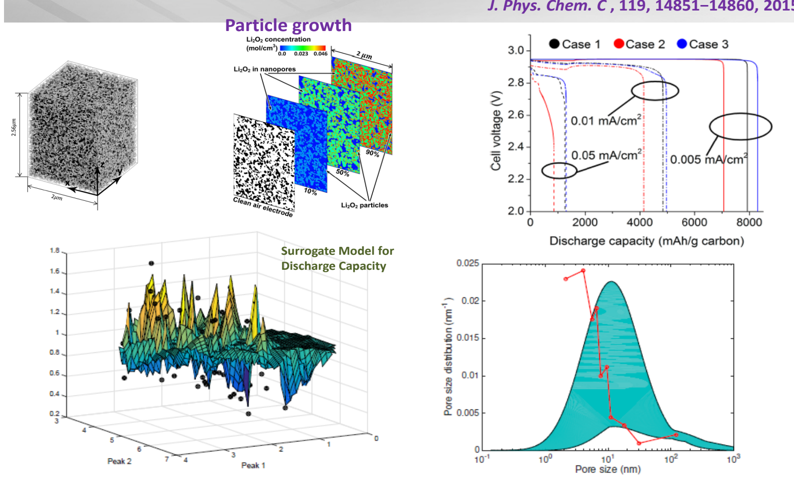

CM4 – Integrative Applications

FROM MICELLES TO CELLS

Modeling Mesoscopic Phenomena using Particle-Based Models: Theory, Implementation, and Applications

Yu-Hang Tang, George Em Karniadakis

Division of Applied Mathematics, Brown University

The successful application of computer simulation techniques for solving problems in physical sciences requires interdisciplinary effort spanning theory, software implementation, and application. In this poster I present various aspects of mesoscopic particle-based simulation methods, including model construction and parameterization, code and algorithm design, and conducting large-scale simulations using state-of-the-art supercomputers.

Dissipative Particle Dynamics

Dissipative Particle Dynamics (DPD) is a stochastic, particle-based simulation technique that was specifically designed for modeling mesoscopic systems. The DPD pairwise interaction consists of three terms, i.e. a conservative term, a dissipative term, and a random term.

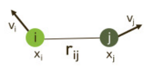

Pairwise force	$\mathbf{F}_i = \sum_{j \neq i} (\mathbf{F}_{ij}^C + \mathbf{F}_{ij}^D + \mathbf{F}_{ij}^R)$
Conservative	$\mathbf{F}_{ij}^C = a_{ij}(T_{ij}) w_C(r_{ij}) \mathbf{e}_{ij}$
Dissipative	$\mathbf{F}_{ij}^D = -\gamma_{ij} w_D(r_{ij}) (\mathbf{e}_{ij} \cdot \mathbf{v}_{ij}) \mathbf{e}_{ij}$
Random	$\mathbf{F}_{ij}^R = \sigma_{ij} w_R(r_{ij}) \xi_{ij} \delta t^{-\frac{1}{2}} \mathbf{e}_{ij}$

Our $_{USER}$**MESO** package for LAMMPS is a fully GPU-accelerated package for running Dissipative Particle Dynamics simulations. Instead of being merely a translation of the conventional molecular dynamics, the package integrates several innovations that specifically targets CUDA devices: an atomics-free neighbor list construction algorithm and a locally transposed storage layout; a new seeding scheme for in-situ random number generators; fully overlapped computation/transfer; and specialized transcendental functions.

An atomic-free algorithm was invented to construct the neighbor list for each DPD particle in parallel.

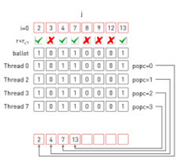

A locally column-major, globally row-major neighbor list can be fast both for writing and reading.

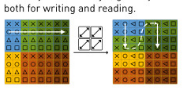

Transposition throughput (K20)
78.03/79.06 GB/s R/W
157.09 GB/s Aggregate

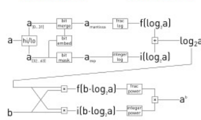

Custom transcendental functions were design to exploit the knowledge on parameter range to reduce instruction branches.

Speedup and Scaling

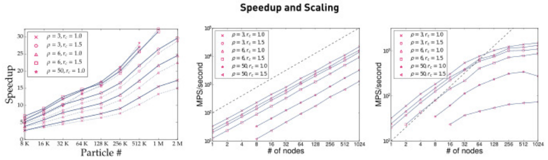

Mesoscopic Simulations of Cell Sorting Microfluidics

Circulating Tumor Cells and Microfluidic Devices

8 million cancer death each year | 90% attributed to metastasis

Metastasis is the process of cancer (cells) spreading to other body parts or organs which are not directly connected to the primary site.

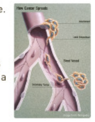

The process start when some cancer cells penetrate the walls of blood vessels and travel through the bloodstream as circulating tumor cells (CTCs) to other sites and tissues in the body. The CTCs can re-penetrate the vessel or walls at a different location and continue to multiply, eventually forming another clinically detectable tumor.

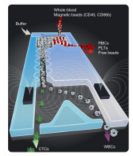

The CTC-iChip is a very effective microfluidic device to date to separate CTCs from RBCs. It consists of an array of obstacles that lead to cell separation by exploiting the concept of the Deterministic Lateral Displacement (DLD). An optimal "egg-looking" shape has been used as a building block in the first compartment of the CTC-iChip microfluidics device.

Karabacak, Nezihi Murat, et al. Nature protocols 9.3 (2014): 694-710.

Numerical Model

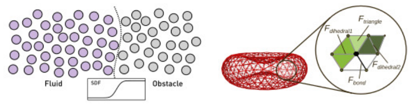

The microfluidic channels consists of stationary particles by freezing those inside the wall geometry after a equilibration phase. Boundary conditions are enforced through DPD interactions between wall and solvent particles. Non-penetrability is enforced by bouncing backing particles according to a Signed Distance Function (SDF).

Each RBC membrane is discretized as a biconcave mesh of 500 vertices. Three types of bonded forces are used to reproduce the dynamics of the membrane: bonds, angles, and dihedrals. The cancer cells are treated similarly on a spherical mesh.

HPC Software Design

Our simulator, uDeviceX, was developed targeting the Cray XK7 accelerated supercomputers. The per-node performance of XK7 is primarily contributed by the K20X GPU based on the Kepler microarchitecture and CUDA compute capability 3.5. The software implements asynchronous point-to-point communication and parallel I/O to achieve the best scalability possible.

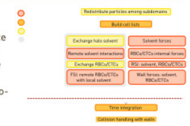

uDeviceX reaches 65.5% of the peak for the most computationally-intensive kernel and 34% overall peak performance of TITAN. It outperforms the LAMMPS by a factor of up to 45X. We simulate microfluidic devices with a volume of 132 mm^3 at submicron resolution and up to 1.43 billion deformable RBCs.

The most time consuming kernels are the ones computing the DPD, FSI and boundary interaction forces. The DPD force kernel is deeply optimized by enforcing the Newton's 3rd law, floating point *in situ* random number generator (RNG), and shared-memory work queue to minimize warp divergence and maximize temporal and spatial locality.

```
shared start[16], count[16], queue[64]
nq   = 0
lane = threadIdx.x mod 32
for cell i = 0 to Ncells stride Nwarps
  if lane < 14
    ncid = id of laneth neighbor cell
    count[lane] = cell_size[ncid]
    start[lane+1] = prefix( count[lane] )
  for j = 0 to count[14] stride 32
    k = trisection( count, j )
    s = start[k] + j - count[k]
  for each particle d in cell i
    p = bool( |xd-xs| < rc )
    hit = ballot( p )
    if p:
      queue[nq + popc( hit & lanemask.lt )] = (s,d)
    nq += popc( hit )
    if nq >= 32:
      is, id = queue[ lane ]
      f = force( is, id )
      atomic( force[id]+=f )
if       is < id:
       atomic( force[is]-=f )
      queue[ lane ] = queue[ lane + 32 ]
      nq -= 32
```

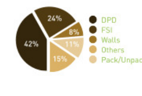

Outperforming factor over LAMMPS

Nodes	Ranks/node	Particles/rank	Outperforming factor $N_R = 4$	Outperforming factor $N_R = 1$
1	1	3,538,944	26.7	44.1
48	1	3,538,944	26.7	**44.6**
48	8	442,368	7.40	17.2
18600	1	3,538,944	22.2	**37.9**
18600	8	442,368	4.9	11.4

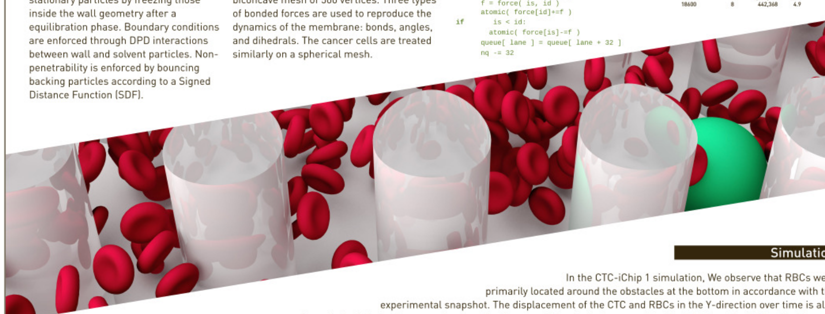

Simulation

In the CTC-iChip 1 simulation, We observe that RBCs were primarily located around the obstacles at the bottom in accordance with the experimental snapshot. The displacement of the CTC and RBCs in the Y-direction over time is also shown in the left plot for assessing the effectiveness of the device in separating the CTC from the RBC cells. We note that the CTC was drastically displaced in the Y-direction allowing for its separation from the RBCs in agreement with the experiment. In the CTC-iChip 2 simulation, we clearly observe cell separation even after the first cycle as illustrated in the right plot, but irreversibility is achieved only after the fifth cycle. Interestingly, due to the flow properties and random initial conditions, type 2 cells end up getting trapped between the funnels.

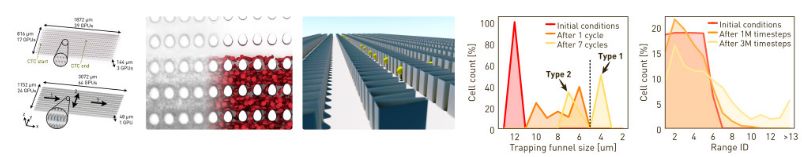

Massive Ensemble-Statistics of Non-Equilibrium Vesicle Dynamics

Data-Driven Mechanism Discovery

Diblock copolymer self-assemblies whose blocks have opposite thermoresponsivity can flip its surface/inside composition upon temperature change. To find out how each individual molecule behaves during thermally induced inversion, we employed the **proper orthogonal decomposition** (POD) method to extract patterns from the noisy molecular trajectories.

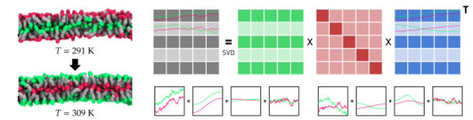

We discovered two dominant modes: **flips** and **slip**. Surprisingly, 21.5% of the molecules used the slip mode and did not invert orientation although the membrane inverted overall composition.

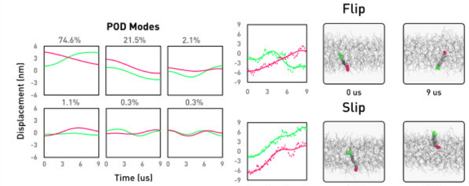

Quantifying Rare Events Using Petascale Simulations

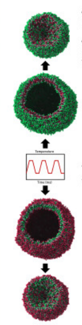

A unilamellar vesicle (ULV) formed by a $\mathbf{LCST_2N_5UCST_2}$ triblock copolymer may invert by switching its surface and internal thermorespon-sive blocks reversibly.

An interchange of the shapes of the block radial density distribution ρ(r) for the LCST and UCST blocks indicates that the molecules did not invert in a lock-step fashion but rather by diffusing through the wall of non-responsive blocks.

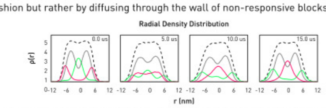

The collapse and subsequent release dynamics is quantified by a thermal loading test **using 576 ensembles each containing a single ULV running on 18,432 GPUs**.

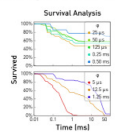

Table 1: Survival rates and estimated collapse probability of vesicles under different thermal loading frequencies. An asterisk indicates a field that cannot be estimated reliably due to the long tail of the failure probability.

Period	Survived	$p(t;a,b) = \frac{t^{a-1}}{\Gamma(a)b^a} e^{-\frac{t}{b}}$ a	b	Mean lifetime (ms)	Half-life (ms)
5 μs	0.0%	1.170	0.239	0.28	0.21
12.5 μs	0.0%	0.394	23.10	9.11	3.25
25 μs	47.2%	*	*	*	*
50 μs	58.3%	*	*	*	*
125 μs	58.3%	*	*	*	*
0.25 ms	54.2%	*	*	*	*
0.50 ms	70.8%	*	*	*	*
1.25 ms	4.2%	0.461	29.07	20.36	8.59

References

1. **Tang Y.-H.,** Li, Z., Li, X., Deng, M., & Karniadakis, G. E. *Macromolecules*, **2016**, 49 (7), pp 2895–2903.
2. Li, Z.; **Tang Y.-H.,** Li, X. & Karniadakis, G. E. *Chem. Comm.*, **2015**, 51, 11038 - 11040.
3. **Tang Y.-H.** & Karniadakis, G. E. *Comput. Phys. Comm.*, **2014**, 185.11, 2809-2822.
4. Li, X.; **Tang Y.-H.;** Liang, H. & Karniadakis, G. E. *Chem. Comm.*, **2014**.
5. Li, Z.; **Tang Y.-H.;** Lei, H.; Caswell, B. & Karniadakis, G. E. *J. Comput. Phys.*, **2014**, 265, 113-127.
6. Rossinelli, D., **Tang, Y.-H.**, Lykov, K., Alexeev, D., Bernaschi, M., Hadjidoukas, P., Bisson, M., Joubert, W., Conti, C., Karniadakis G., Fatica, M., Pivkin, I., Koumoutsakos, P. The In-Silico Lab-on-a-Chip: Petascale and High-Throughput Simulations of Microfluidics at Cell Resolution. ACM 2015 Gordon Bell Award Finalist, Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis. 2015.

CM4

Work supported by the Department of Energy Collaboratory on Mathematics for Mesoscopic Modeling of Materials (CM4). Simulations were carried out at the Oak Ridge Leadership Computing Facility through the INCITE program under project BIP102 and BIP118. YHT acknowledges partial financial support from an IBM Ph.D. Scholarship Award.

DiaMonD highlight: Bayesian inversion for large-scale complex models with application to flow of the Antarctic ice sheet

O. Ghattas (UT-Austin), T. Isaac (U. Chicago -> GaTech), N. Petra (UC Merced), G. Stadler (NYU)

Driving forward problem: modeling flow of the Antarctic ice sheet

- Nonlinear Stokes with temperature- and strain-rate-dependent (non-Newtonian) viscosity
- Strong nonlinearities, complex rheology, Highly ill-conditioned due to orders-of-magnitude variation in viscosity & basal friction, Wide range of spatial scales: $O(10^2$ m) to $O(10^6$ m)
- To meet these challenges, we have developed a new state-of-the-art numerical model based on **locally-mass conserving and high-order discretization, aggressive parallel AMR, physics-based algorithmically optimal multilevel linear and nonlinear solvers** with demonstrated scalability to $O(10^6)$ cores (on LLNL Sequioia BG/Q)

Driving inverse problem: inferring Antarctic basal friction from satellite observations of surface velocity

- The parameter field to be inferred is basal friction, which is an infinite-dimensional field
- Observational data from InSAR
- Bayesian formulation to invert not only for mean of basal friction, but uncertainty in inversion as well
- Developed novel optimal low-rank-based Bayesian inversion method that exploits Hessian structure and scales to $O(10^6)$ parameters (largest & most complex Bayesian inverse problem solved)

Awards, recognition, training of next generation researchers:

- Underlying multilevel implicit solver received 2015 Gordon Bell Prize, 2016 Copper Mountain Best Student Paper Award, SC14 Best Poster (from among 193 submissions)
- Junior researchers on project received 2016 SIAM Supercomputing Early Career Prize (T. Isaac), 2016 ACM/IEEE-CS George Michael Memorial HPC Fellowship (J. Rudi), and assumed faculty positions at NYU, UC Merced, & GaTech

Bayesian inversion for basal friction field in Antarctica

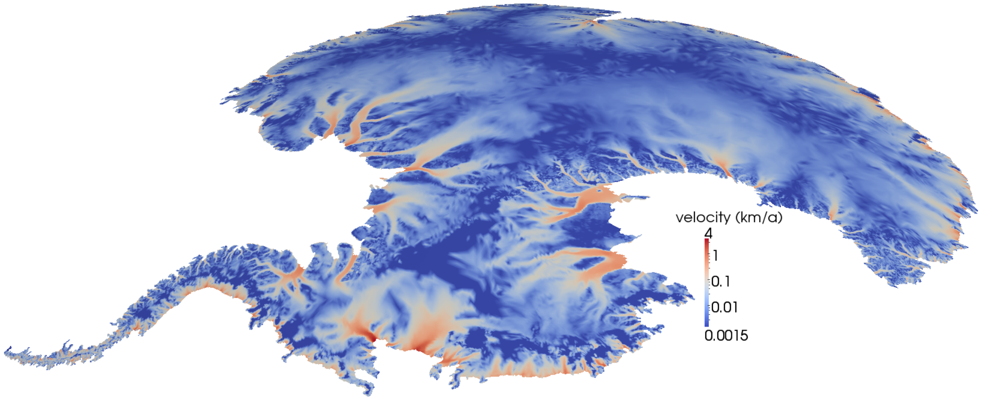

InSAR-based ice surface velocity observations

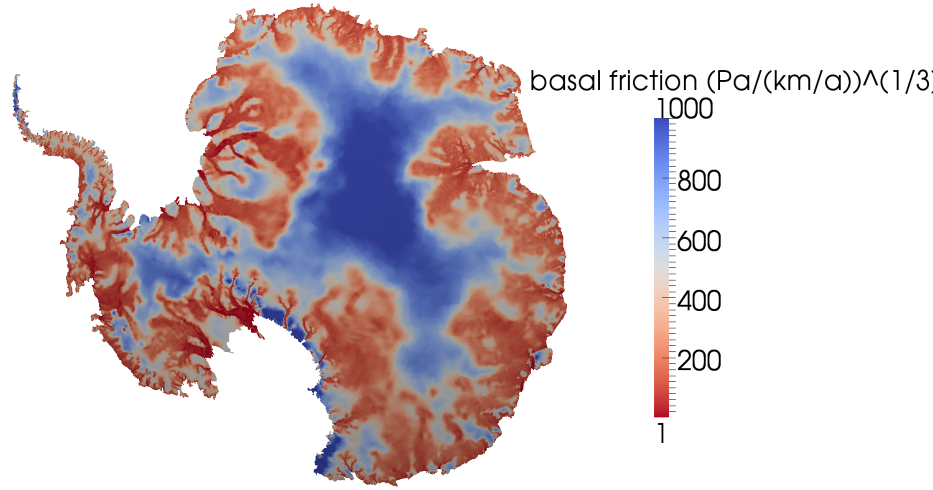

Inferred mean of basal friction field

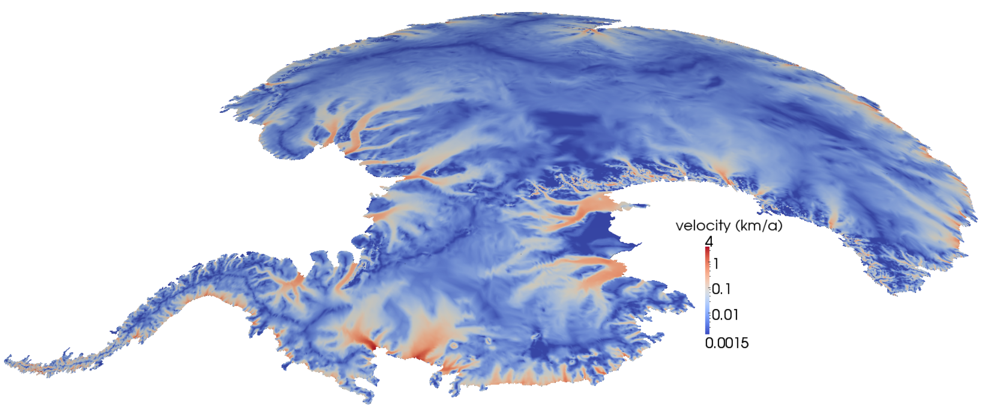

Reconstructed ice surface velocity field (based on inferred mean of basal friction field)

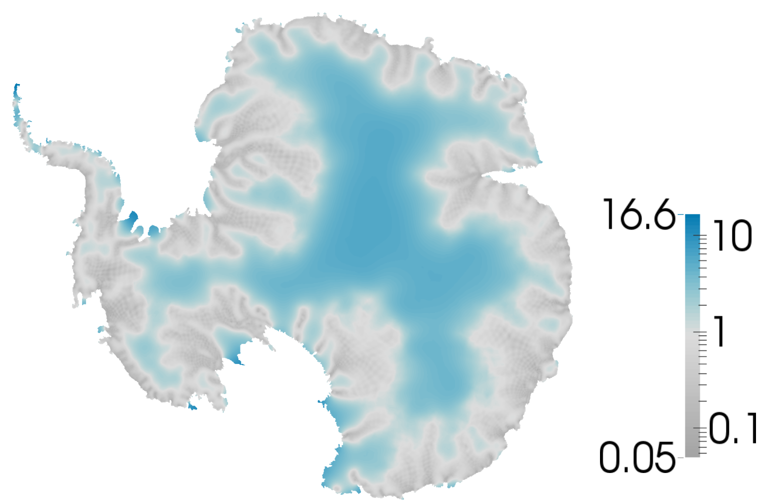

Inferred uncertainty in basal friction field (standard deviation of Gaussianized posterior of log basal friction)

U.S. DEPARTMENT OF ENERGY | Office of Science

Inference via low-dimensional couplings

DiaMonD highlight: Y. Marzouk, A. Spantini, D. Bigoni, MIT

- ***Integration* against an intractable probability measure is the core task of statistical inference**
- **New & principled approach: *couple* the target measure to a tractable measure via a transport map**
  - Yields **independent & unweighted Monte Carlo samples**
  - Enables sparse quadrature, quasi-Monte Carlo, etc. for arbitrary distributions
  - Provides easy-to-evaluate error measures (unlike MCMC)
- Key impacts:
  - **Tractable and fully Bayesian computation:** bridges the gap between variational inference (common in *machine learning*) and sampling methods
  - Generalizes important statistical models (e.g., Markov random fields) and algorithms (e.g., Gaussian smoothers, EnKF) to the fully nonlinear and non-Gaussian case
  - A new foundation for **sequential data assimilation** algorithms…

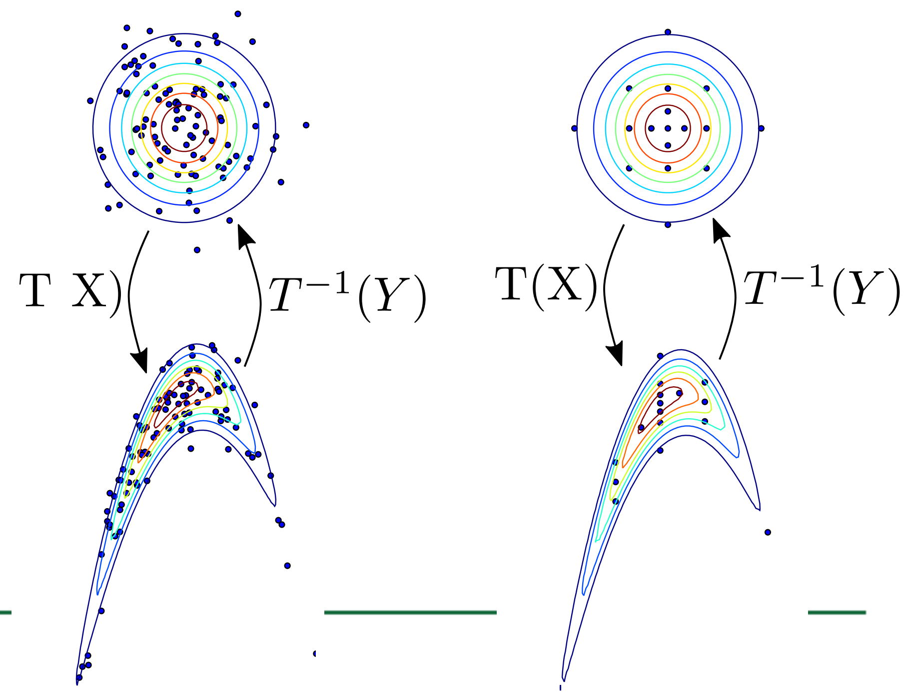

- **New theory and algorithms identify and exploit low-dimensional structure in transport maps:**
  – Outcomes: sparse, decomposable, and/or low-rank maps
  – Yields new **sequential inference algorithms** for streaming data
  – Generalizes previous **dimension reduction** strategies for inverse problems
- **Efficient Bayesian inference in large-scale applications: *no importance weights, resampling, or MCMC***
  – High-dimensional spatial problems *(below, 4096 dimensions)*

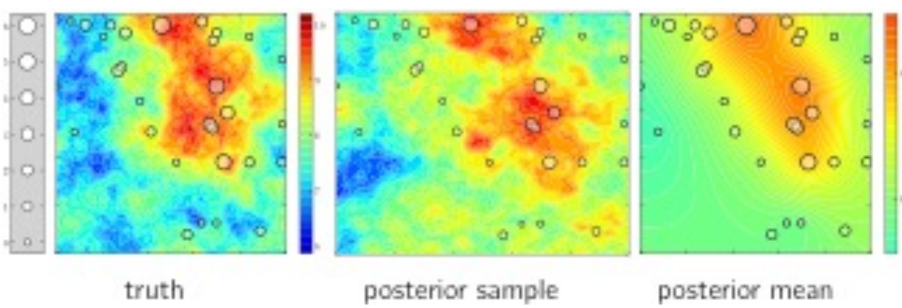

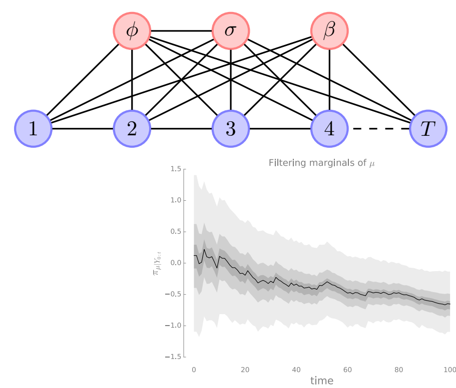

  – Data assimilation: new *variational* algorithms for Bayesian filtering, smoothing, and **online** parameter estimation, with no particle degeneracy

- Multiscale inference [Parno, Moselhy, & Marzouk *JUQ* 2016]

DiaMonD Future Development: Learning from Data with Quantum Computing (LANL)

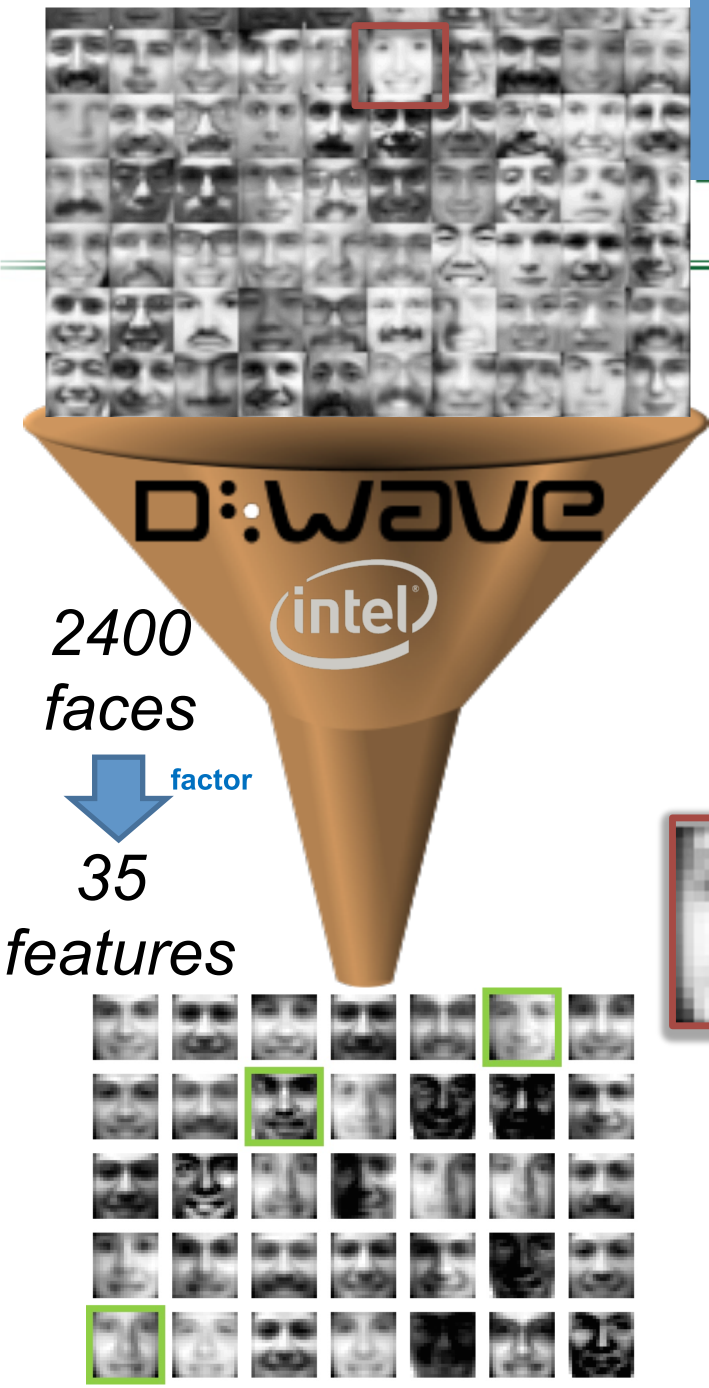

Feature extraction using non-negative matrix factorization:

- Accurate, low-rank representations of big data
- Up to 60x speed-up over Gurobi (a state-of-the-art classical tool)
- Next-gen D-Wave hardware will be much more performant for this problem
- Next step: **inferring subsurface flow models from observational data with the D-Wave**

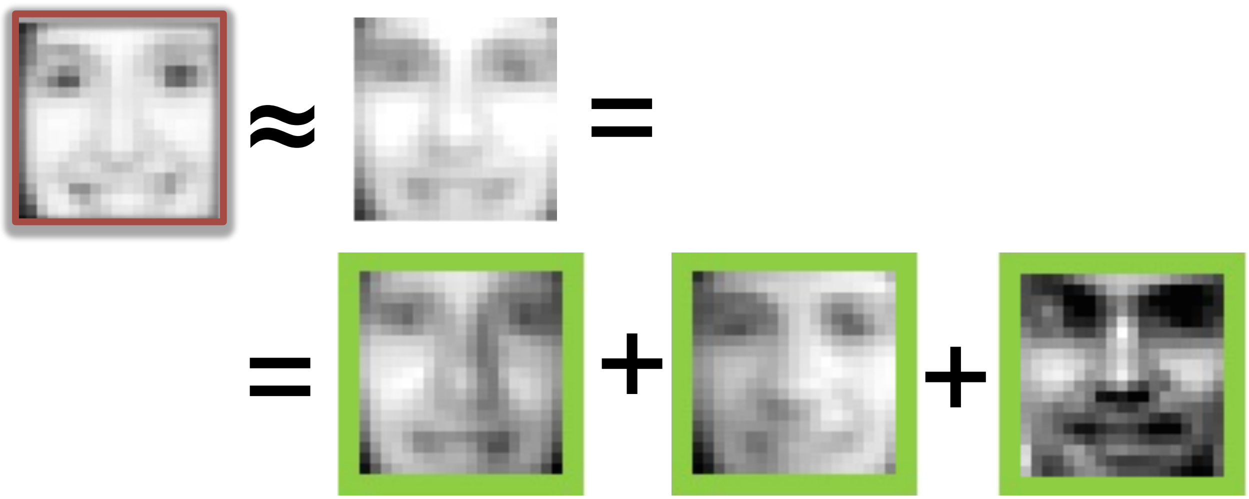

O'Malley, Vesselinov, Alexandrov, Alexandrov (LANL)

Study Group on Applied Mathematics Center Like Investments

- Post-review study – PIs, senior lab personnel, Other agency program managers, …
- Ten Well Defined Questions – five to capture best practices and five to explore futures

"The collective group of people involved in this round of MMICCs would never had embarked on this successful line of research without the MMICCs program."

Study Group Discussions

1. ***Was the scope of the project sufficiently well defined? Was it easy to define what was "in-scope" and what was not? If new challenges or ideas came about as the project progressed, were you able to exploit the opportunities and/or cope with the failures without losing focus?***
   - Proposals were well-defined at a high level, focusing on conceptual ideas
   - Specifics of the research developed organically. It was easy to try new specific thrusts because the scope was not overly prescribed. Many examples of new directions and projects shown.
   - Potential issue -- **no mechanism to bring in additional people from outside the project or remove people to adapt to changing priorities and research directions or poor performance**.

Study Group Discussions

2. The MMICCs projects are unique in DOE investments as "focused research" at large scale. Was the scale – too big, too small or just right? How did you decide on and integrate the components of the project and their evolution as the project progressed?

- *"Right sized" – not too small and not too big*
- Interactions within the projects were generally not pre-designed and emerged as the project progressed to address specific problems.
- Sub-teams organized around specific sub-problems and harnessed diverse skills present and encouraged team members to explore new areas.
- *Potential Weakness:* Flexibility needed to form cross-institutional teams

Study Group Discussions

3. ***Integration of laboratory and university skills, resources and personnel is always tricky because of the different priorities and processes. How did that work here?***

   – Academic members with strong existing DOE connections

   – Early-career investigators, who are talented but not yet tenured

   – Connecting the junior researchers—by providing travel budgets, lab visits, and cross-university visits …

   – *Potential Issue*: Faculty are not research intensive during semester. Engage graduate students and post-docs directly; use laboratory internships

4. Project Management practices – can you comment on best practices or drawbacks?

- Frequent meetings early on were identified as important to gets team members talking to each other.
- Defining mathematically detailed model problems early on in version controlled documents provided focus.
- One project used weekly webinars to maintain contact with and between the project participants.
- One project held frequent PI meetings where all senior PIs could provide feedback.
- For larger meetings with project leadership including track leads, it was important to have an agenda beforehand with topics solicited from team.

Study Group Discussions

5. ***What were the best practices in recruiting and developing talent that we should template and replicate?***
   - **Retention was more of a challenge than recruiting. Senior PIs departing soon after award creates problems**.
   - The scale of the centers and the possibility for collaborative opportunities were useful selling points.
   - The stability of having 5 years of guaranteed funding allowed the projects to recruit the best students.
   - Providing the opportunity to spend time at DOE labs was an attraction for students.
   - Having high-profile PIs also helped attract talent.

Study Group Discussions

1. How do we identify and sustain support for mission relevant fundamental research in Applied Mathematics and Statistics?

- Keep MMICCs centers organized around mission-driven problems.
- Maintain several diverse centers;
- Find ways to better disseminate results to highlight the successes. Make argument for support of Mathematics.

Study Group Discussions

2. How can we encourage investigators to explore new and risky emerging research paradigms inside these larger organized research units?

- Risky research → valuable research that would have not been funded otherwise.
- Flexibility is key. Provide opportunities for teams to self assemble. The lack of prescription allows centers to evolve.
- Make the centers easily reconfigurable
  - Give PIs the freedom to issue new subcontracts to add new people.
  - Fellows who can propose 1-year high-risk projects and give the MMICCs freedom to spend 10% of budget on high-risk projects.
  - Allow centers to add new people, perhaps within 2-3 years of taking a tenure-track job, who bring new research directions.

Study Group Discussions

3. ***How do we design these units to respond to the challenge of simultaneously supporting research on long standing hard problems like coupled high fidelity climate modeling (for example) and fast changing demands of things like new data intensive science and advanced architectures?***
   – Long-term commitment with the latitude to rebalance and reconfigure their MMICC will allow them to adapt to fast-changing demands while still addressing long-standing problems.
   – Combine this long-term horizon with supplemental add-on proposals: write 5-7 pages to bring in more people to work on tasks different from the original tasks. This would fund additional parallel work.
   – Any center does need to include people from both labs and academia; the cross-fertilization is essential.

Study Group Discussions

4. ***How do we incentivize these entities to integrate researchers at DOE laboratories and universities? How do we help them create and sustain the talent pipeline for DOE's current and future workforce needs?***
   – Centers need to be funded on a cycle that works with funding students and postdocs.

   – Guarantee student/post doc funding for fixed term

   – Try to synchronize the funding of lab and university partners.

Study Group Discussions

5. ***Interesting ideas for new structures to organize the research program. Governance structures, award mechanisms and monitoring?***
   - Add a fellows program as described above to foster new research directions and the inclusion of new team members.
   - Institute a program in which existing ASCR Applied Mathematics grant holders can write small supplemental proposals to collaborate with MMICCs
   - Major milestone at three years where serious feedback can be given.
   - Need to put a mechanism for renewal in the next call, where the renewal process is something that can be appropriately customized and not involve a full blown proposal

QUESTIONS/ COMMENTS?

DiaMonD research on DNS methods for complex fluids & porous media

(George Biros, UT Austin)

Goal: HPC, scalable, multi-fidelity, black-box (tuning- and parameter-free) solvers

Formulation: integral equations, implicit time-stepping, high-order discretization methods, preconditioners, O(N) fast N-body direct solvers

Reduced order models: Physics-based correction to small DOF; high-dimensional regression

Physics apps: transport & mixing, polymers, colloids, active matter, non-equilibrium statistical mechanics

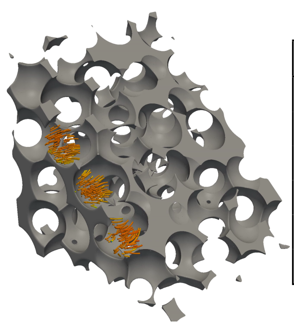

p	N_{dof}/p	N_{iter}	T_{solve}	TFLOPS	η
1	8.0E+6	155	477	0.36	1.00
6	7.8E+6	115	388	2.27	1.04
27	8.6E+6	101	401	10.3	1.05
125	8.5E+6	98	419	45.3	0.99
508	8.9E+6	92	444	173	0.94
2048	9.1E+6	90	474	656	0.88

2048 (x86+KNC) nodes, 18B DOFs, 88% efficiency

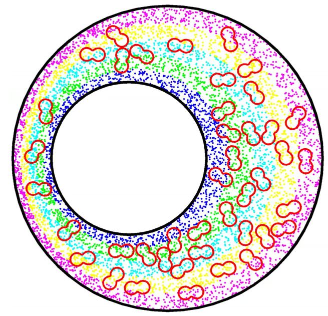

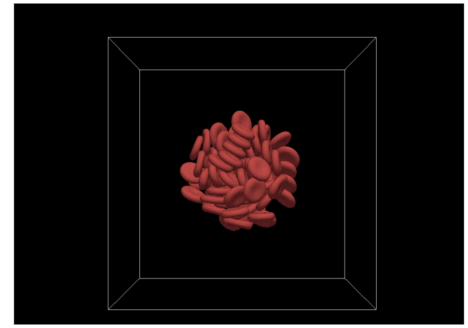

Particulate flows

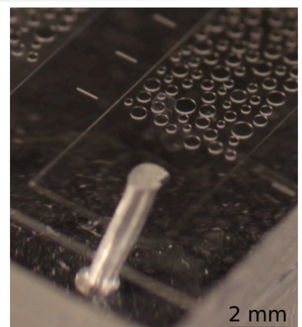

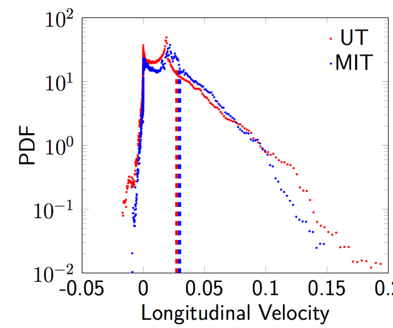

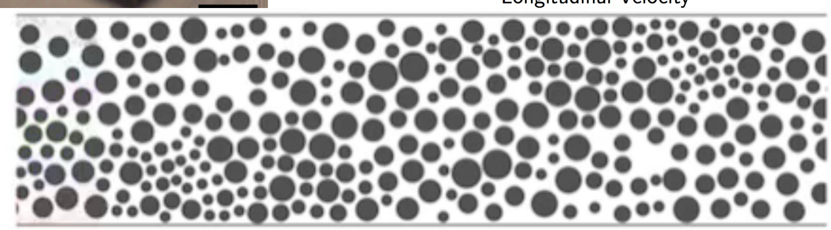

Porous media (collab. w/ R. Juanes, MIT