Newsjournal of the Society for Industrial and Applied Mathematics sinews.siam.org

siam news

Volume 51/ Issue 1
January/February 2018

Geosciences Special Issue

*In this **special issue**, read about the application of mathematics and computational science to various topics in the geosciences.*

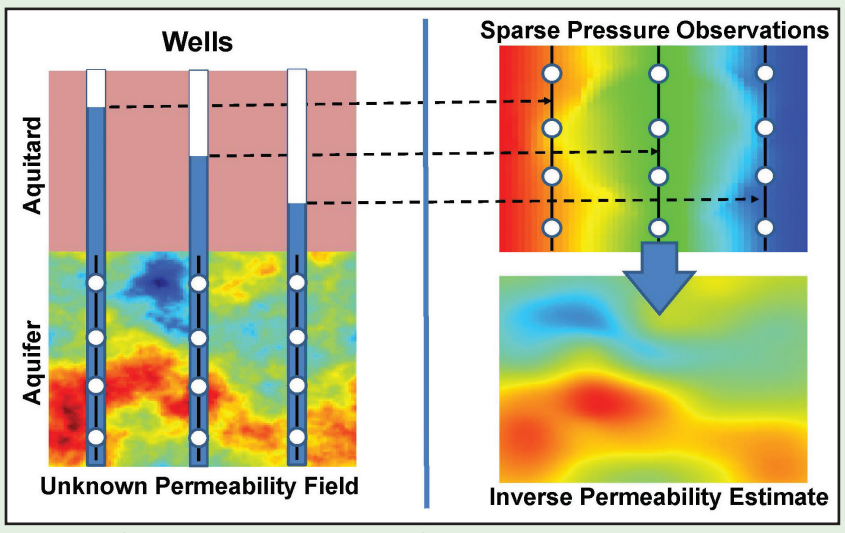

Figure 1. *Schematic representation of a typical hydrologic inverse problem where observations of hydraulic heads at wells are used to estimate aquifer permeability. Image credit: Youzuo Lin.*

In the article "Randomization in Characterizing the Subsurface" on page 5, Youzuo Lin, Daniel O'Malley, Velimir V. Vesselinov, George Guthrie, and David Coblentz describe how randomized matrix algorithms can provide fairly accurate, economical, and robust computational approaches to characterize Earth's subsurface.

Never Heard of Methane Hydrate? That Might be Good News

By Malgorzata Peszynska

Methane hydrate is an ice-like crystalline substance (gas clathrate) made of water molecules encasing a molecule of methane. Low temperatures and high pressure favor the hydrate's existence (e.g., entrapment of methane molecules in the ice cages), making hydrate deposits abundant in deep ocean sediments and in Alaska (see Figure 1, on page 4). Upon disturbance of these favorable conditions, methane might escape to the atmosphere from hydrate-bearing sediments, contributing to the overall balance of other greenhouse gases. Methane hydrate is considered a "smoking gun" in environmental and (paleo-)climate studies[1] [8]. Hydrates are also known for the nuisance they cause plugging up wells and pipes that carry other hydrocarbons to their destinations. They are a significant drilling hazard and contributed to the Deepwater Horizon explosion and oil spill. Researchers study the presence of hydrates due to their impact on slope stability of submarine formations [4]. If you haven't yet heard about methane hydrate, it has likely not caused any recent high-profile disasters.

However, exciting (non-disastrous) scientific news related to the study of gas hydrates is also possible, and mathematicians can play a role. Geophysicists have been studying hydrates for a long time, but many scientific mysteries surrounding their origin and evolution remain. Computational models exist, but require interdisciplinary collaboration and access to first-rate information and data to be meaningful (see Figure 2, on page 4). Computational and applied mathematicians can build virtual laboratories for hydrates, but their real impact requires investment in the language, methodology, priorities, and funding models of the allied fields.

For example, the hydrate "lives" in its host rock. How did it get there? Basin modeling, which predicts hydrate deposit formation from upward migration of gas over timescales of thousands of years and spatial scales of hundreds of meters, provides the answer. Comprehensive basin models [5] offer convincing scenarios of gas hydrate formation[2] (see Figure 1, on page 4). Notwithstanding the enormous uncertainty in basin modeling, these models are very complex and

See **Methane Hydrate** *on page 4*

[1] https://woodshole.er.usgs.gov/project-pages/hydrates/

[2] http://math.oregonstate.edu/~mpesz/gallery.html#hydrate

Physics-based Probing and Prediction of Extreme Events

By Mohammad Farazmand and Themistoklis P. Sapsis

Extreme events arise spontaneously in a variety of natural and engineering systems. They are usually unexpected, transient phenomena that take place over short time scales and have very large magnitudes when compared to typical system responses. In this sense, extreme events are a subclass of the so-called rare events. But while rare events are only associated with low probability, extreme events are also generally of very large magnitude [6]. Examples include extreme weather patterns, aeroacoustic instabilities in combustion engines, earthquakes, rogue waves on the ocean surface, and power grid overloads (see Figure 1). Since these events often have undesirable economic, environmental, and humanitarian consequences, their study is of great interest.

The four fundamental questions related to extreme events include the following:

1. Mechanism: What conditions lead to the occurrence of an extreme event?

2. Quantification: What is the likelihood or frequency of an extreme event taking place?

3. Prediction: Are there indicators or triggers whose measurement would signal a forthcoming extreme event?

4. Mitigation: What control strategies are best suited for suppressing extreme events? Can controlling the trigger lead to suppression? While mitigation of extreme events in nature appears to be out of reach, specialized control protocols may be devised to suppress those in engineered systems (such as power grids or fluid flows).

The first question is of primary focus, since its answer is the cornerstone for the subsequent three questions.

We concentrate on problems for which a model, in the form of a deterministic or stochastic dynamical system, is available. An extreme event is associated with unusual growth (or decay) of the time series of a particular observable of the system. These dramatic fluctuations occur when the system's trajectory visits a subset of the state space that acts as the basin of attraction to the extreme events. These basins harbor certain instabilities that momentarily repel the trajectory from its background dynamics (see Figure 2, on page 3). To understand the underpinning mechanism, one must detect the extreme event instability regions of the state space.

In high-dimensional chaotic systems, determining the instability regions is complicated. We propose a constrained optimization method to probe the state space of high-dimensional dynamical systems in search of the subsets that underpin extreme events [5].

Let the governing equations $\partial_t u = N(u)$ describe a system where $u(t) \in X$ denotes the state of the system at time t and X is an appropriate function space. Also, let $I : X \to \mathbb{R}$ denote an observable whose relatively large values constitute an extreme event. In order to probe the onset of extreme events in the state space, we seek initial states $u_0 \in X$ whose corresponding trajectory $u(t)$ maximizes the growth of the observable I over a given finite time τ. To obtain physically relevant maximizers, one must further constrain the initial states u_0 to ensure that they belong to the system's attractor; this rules out exotic maximizers that have zero probability of being observed under the system's long-term natural dynamics. This constrained optimization problem can be written more precisely as

$$\sup_{u_0 \in X} \Big[I(u(\tau)) - I(u(0)) \Big], \qquad (1a)$$

See **Extreme Events** *on page 3*

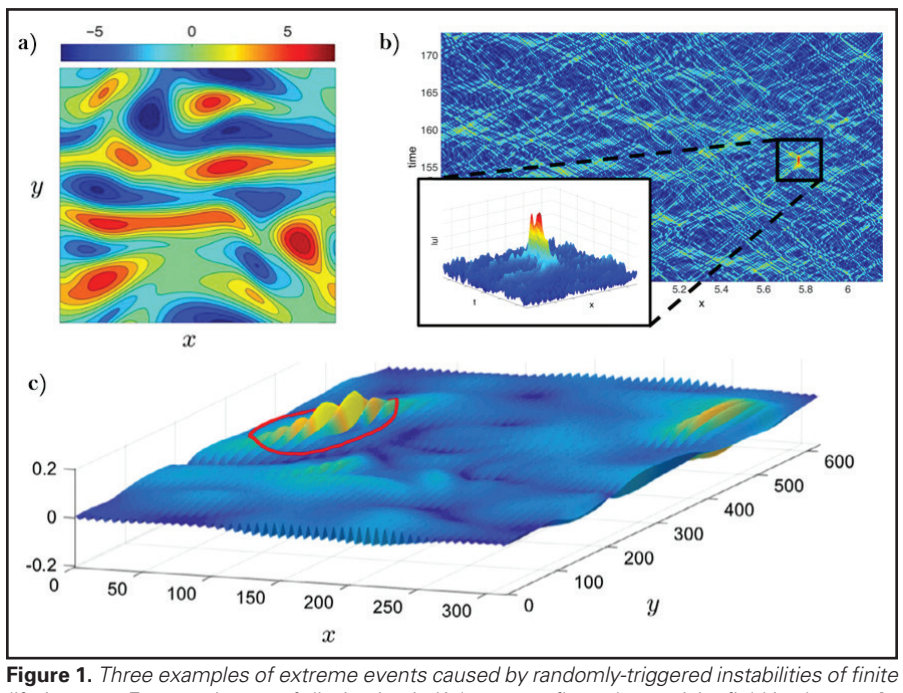

Figure 1. *Three examples of extreme events caused by randomly-triggered instabilities of finite lifetime.* **1a.** *Extreme bursts of dissipation in Kolmogorov flow; the vorticity field is shown.* **1b.** *Extreme event in dispersive wave turbulence; the magnitude of the wave field is shown.* **1c.** *A rogue wave in two-dimensional directional seas. The surface shows the wave elevation. Image credit: Mohammad Farazmand and Themistoklis Sapsis.*

Nonprofit Org
U.S. Postage
PAID
Permit No 360
Bellmawr, NJ

siam
SOCIETY for INDUSTRIAL and APPLIED MATHEMATICS
3600 Market Street, 6th Floor
Philadelphia, PA 19104-2688 USA

siam news
Volume 51/ Issue 1/ January/February 2018

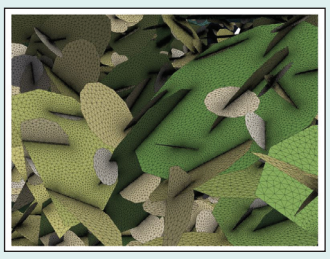

Snap to Structure

FROM THE SIAM PRESIDENT
By Nicholas Higham

The Adobe Photoshop image manipulation program has a feature called "snap to," which snaps edges of a selection into alignment with grid lines superimposed over the image as users move the selection around. This feature is an invaluable tool for producing perfectly aligned, professional-quality graphics. An analogous operation in mathematics is what Alan Edelman calls "snap to structure." Here, a mathematical object that is required to have a particular property, but fails to do so, is perturbed so that it has the property. An orthogonal projection onto the set of interest typically defines the perturbation.

A ubiquitous example of snap to structure occurs in floating-point arithmetic. When we compute the sum or product of two double-precision floating-point numbers, the exact result may have more significant digits than we started with. The IEEE standard requires the exact product to be snapped back to double precision according to precisely defined rounding rules, with round to nearest the default.

In recent work, Jörg Liesen and Robert Luce [1] consider the question of whether a given $m \times n$ matrix is a Cauchy matrix — whether its (i, j) element has the form $1/(s_i - t_j)$ for all i and j and some m-vector s and n-vector t. Various computations can be done asymptotically fast with such matrices (of which the famous Hilbert matrix is a special case). The authors also treat the problem of approximating a matrix by a Cauchy matrix, that is, snapping to Cauchy structure.

Snapping to a matrix structure is commonly described as solving a matrix nearness problem. Here, I take distance to be measured in the Frobenius norm. One of the most familiar examples is forming the nearest rank-k matrix. The Eckart-Young theorem states that to solve this problem one simply computes a singular value decomposition (SVD) of the given matrix and sets all singular values beyond the kth largest to zero. In the context of solving a linear least squares problem, snapping the coefficient matrix to rank k corresponds to forming a truncated SVD solution.

Loss of definiteness of a symmetric matrix is a common problem in many applications. One can find the nearest positive semidefinite matrix to a given matrix by computing a spectral decomposition and setting any negative eigenvalues to zero.

Another example, which I first encountered in aerospace computations, concerns orthogonality. A 3×3 direction cosine matrix drifts from orthogonality because of rounding errors and thus must be orthogonalized, with practitioners favoring the nearest orthogonal matrix over other orthogonalization techniques. For a complex scalar $z = re^{i\theta}$, the nearest point on the unit circle is $e^{i\theta}$. The matrix case generalizes this observation: if $A = UH$ is a polar decomposition of a real, nonsingular matrix A (U orthogonal, H symmetric positive definite) then U is the nearest orthogonal matrix to A.

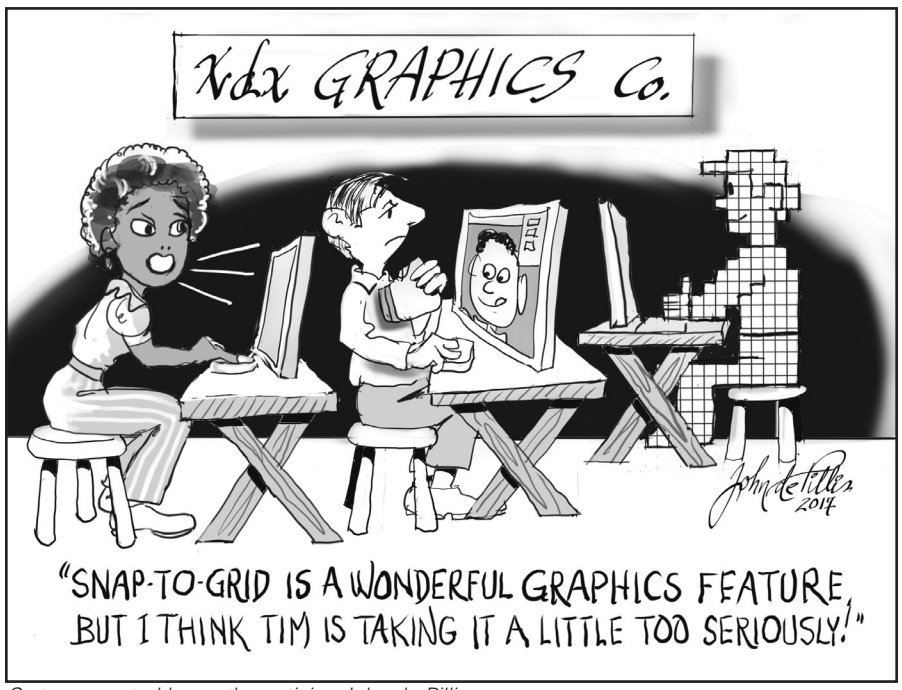

Cartoon created by mathematician John de Pillis.

Snap to structure is also natural when a real solution is expected but the relevant algorithm has made an excursion into the complex plane, potentially leaving an imaginary part of rounding errors. In this case, snapping means setting the imaginary part to zero. This is done in the MATLAB function `funm` (for computing a function of a matrix) when the matrix is real and the result's imaginary part is of the order of the unit roundoff.

Of course, much mathematical research is concerned with preserving and exploiting known structure in a problem, making snapping to structure unnecessary. For example, geometric numerical integration is about structure-preserving algorithms for the numerical solution of differential equations. For a simple illustration, suppose we decide to plot circles by numerically solving the differential equations $u'(t) = v(t)$ and $v'(t) = -u(t)$ with initial values $u(0) = 1$ and $v(0) = 0$, for which $u(t) = \cos t$ and $v(t) = -\sin t$. The forward Euler method produces solutions spiralling away from the circle, while the backward Euler method produces solutions spiralling into the origin. We could apply one of these methods and project the solution back onto the circle at each step. However, a better strategy would be one that produces approximations guaranteed to lie on the circle; the trapezium method has this property. For this problem, $u(t)^2 + v(t)^2$ is an invariant of the differential equations, which the trapezium method automatically preserves.

Despite the large body of theory on structure-preserving algorithms, we live in a world of imperfect data, use finite-precision arithmetic, and have a limited choice of available software, all of which can result in errors that destroy structure. So enforcing structure at some point during a computation may seem sensible. But is it always a good thing to do?

Consider the problem of computing the determinant of a matrix of integers in floating-point arithmetic. Using the definition of determinant by expansion by minors will give an integer result, excluding the possibility of overflow. But this approach is too costly except for very small dimensions. The determinant is normally computed via a pivoted LU factorization: $PA = LU$ implies $\det(A) = \pm\det(U)$. The computed determinant is likely to have a nonzero fractional part because of rounding errors. Rounding the result to the nearest integer might seem natural, but consider this MATLAB example involving a matrix whose elements are from Pascal's triangle:

```
>> n = 18;  A = pascal(n);  det(A)

ans =

  1.8502e+00.
```

If we round to the nearest integer, we declare the determinant to be 2. But the determinant is 1 (as for all Pascal matrices) and A is formed exactly here, so the only errors are within the determinant evaluation. Early versions of MATLAB used to snap to structure by returning an integer result after evaluating the determinant of a matrix of integers. This behavior was changed because, as we have just seen, it can give a misleading result. Interesting rounding error effects can also occur in the evaluation of the determinant of a matrix with entries ± 1.[1]

Photoshop sensibly allows users to disable snap to structure, as one does not aways want elements to line up with a grid. Likewise, snapping to a mathematical structure is not always the correct way to handle a loss of that structure. Understanding why is one of the things that makes life interesting for an applied mathematician.

Acknowledgments: I am grateful to Alan Edelman for stimulating discussions on the topic of this article.

References

[1] Liesen, J., & Luce, R. (2016). Fast Recovery and Approximation of Hidden Cauchy Structure. *Linear Algebra Appl., 493*, 261-280.

Nicholas Higham is the Richardson Professor of Applied Mathematics at The University of Manchester. He is the current president of SIAM.

[1] https://nickhigham.wordpress.com/2017/12/11/the-strange-case-of-the-determinant-of-a-matrix-of-1s-and-1s/

siam news

ISSN 1557–9573. Copyright 2018, all rights reserved, by the Society for Industrial and Applied Mathematics, SIAM, 3600 Market Street, 6th Floor, Philadelphia, PA 19104-2688; (215) 382-9800; siam@siam.org. To be published 10 times in 2017: January/February, March, April, May, June, July/August, September, October, November, and December. The material published herein is not endorsed by SIAM, nor is it intended to reflect SIAM's opinion. The editors reserve the right to select and edit all material submitted for publication.

Advertisers: For display advertising rates and information, contact Kristin O'Neill at marketing@siam.org.

One-year subscription (nonmembers): Electronic-only subscription is free. $73.00 subscription rate worldwide for print copies. SIAM members and subscribers should allow eight weeks for an address change to be effected. Change of address notice should include old and new addresses with zip codes. Please request address change only if it will last six months or more.

Printed in the USA.

Extreme Events

Continued from page 1

such that

$$\begin{cases} \partial_t u = N(u), \quad u(0) = u_0 \\ \underline{c}_i \le C_i(u_0) \le \overline{c}_i, \quad i = 1,2,\cdots,k \end{cases} \tag{1b}$$

where $\tau \in \mathbb{R}^+$ is the typical timescale of the extreme events. The constraints involving $C_i : X \to \mathbb{R}$ and $\underline{c}_i, \overline{c}_i \in \mathbb{R}$ are enforced to ensure that the maximizer belongs to the attractor (or at least to a small neighborhood of the attractor).

We apply this approach to the Kolmogorov flow, a two-dimensional Navier-Stokes equation driven by sinusoidal shear [5]. The energy dissipation rate is known to exhibit intermittent bursts along the system's trajectories [3]. Applying an instantaneous version of the optimization problem (1)—and exploiting the energy-conserving nature of the nonlinear term in the Navier-Stokes equation—reveals that the spontaneous transfer of energy to the mean flow from a large-scale Fourier mode causes the dissipation bursts.

Discovery of this mechanism led to an indicator (the energy of the large-scale Fourier mode) whose low values signal an upcoming burst of energy dissipation. Using long-term simulations and Bayesian statistics, we quantify the probability of future extreme events P_{ee} in terms of the indicator's current value. This results in short-term prediction of extreme energy dissipation in the Kolmogorov flow (see Figure 3).

The developed framework is also valuable in the identification of precursors for extreme events in nonlinear water waves, commonly referred to as rogue waves. Our approach employs the wave field's decomposition into a discrete set of localized wave groups with optimal length scales and amplitudes. These wave groups do not interact due to the prediction's short-term character; therefore, their dynamics can be characterized individually. Using direct numerical simulations of the governing envelope equations [2], we precomputed the expected maximum elevation for each wave group. The combination of the wave field decomposition algorithm and the precomputed map for expected wave group elevation allows one to (i) understand how the probability of rogue wave occurrence changes as the spectrum parameters vary, (ii) compute a critical length scale characterizing wave groups with high probability of evolving to rogue waves, and (iii) formulate a robust and parsimonious reduced-order prediction scheme for large waves.

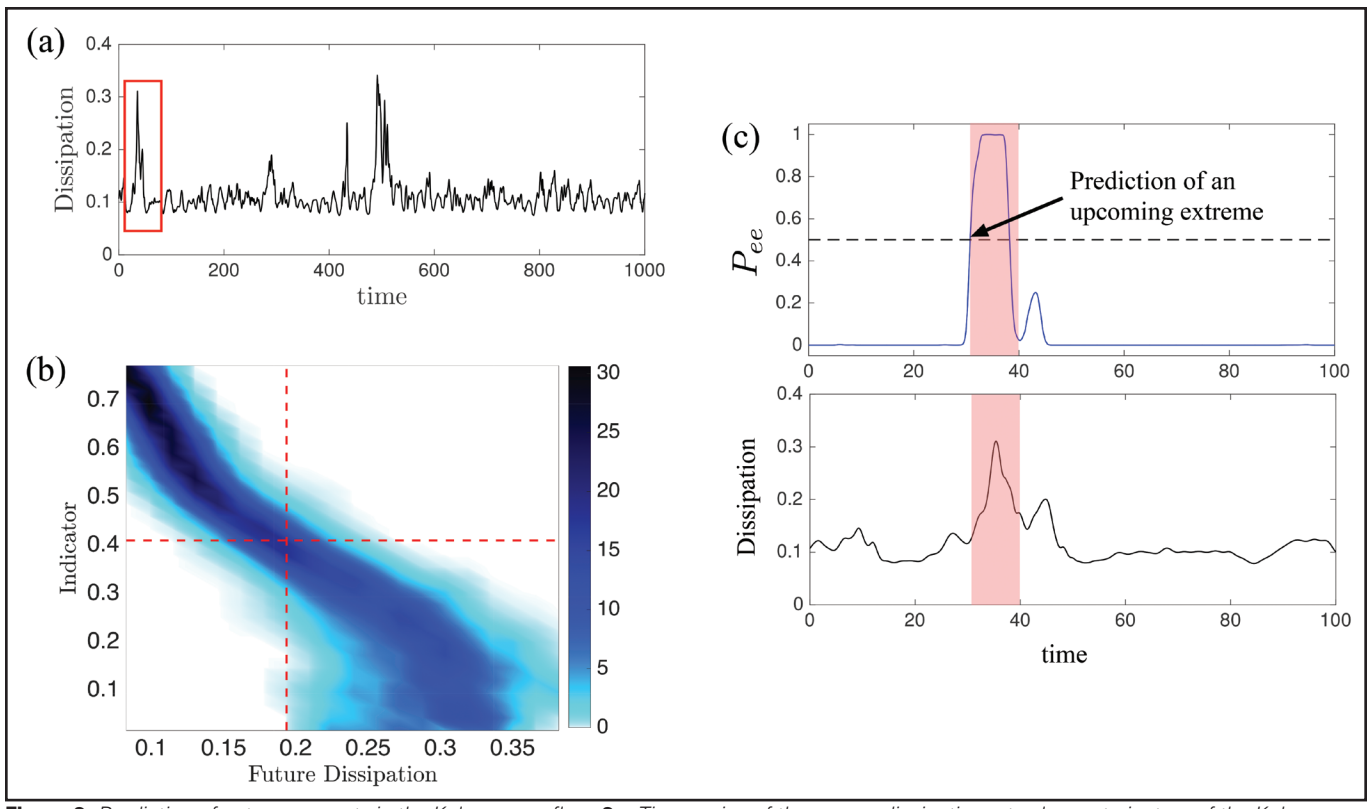

Figure 3. *Prediction of extreme events in the Kolmogorov flow.* **3a.** *Time series of the energy dissipation rate along a trajectory of the Kolmogorov flow.* **3b.** *Conditional probability density of the future energy dissipation and the indicator. Note that large future dissipation correlates strongly with small present values of the indicator, and vice versa.* **3c.** *Prediction of the extreme event marked with a red box in panel 3a.* P_{ee} *measure the probability of upcoming extreme events. See the online version of this article for a time evolution of the plot shown in panel 3c. Adapted from [5].*

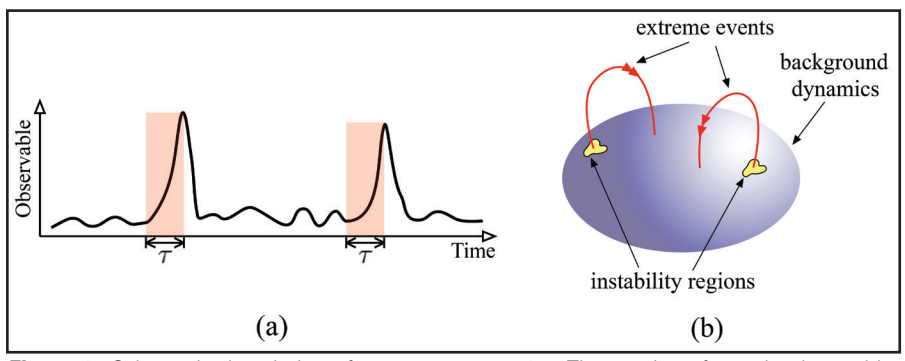

Figure 2. *Schematic description of extreme events.* **2a.** *Time series of certain observables show intermittent bursts.* **2b.** *Observed bursts correspond to the transient deviation of the system trajectory from the background attractor when it visits extreme event instability regions. Adapted from [5].*

Mathematics at the Upcoming AAAS Annual Meeting

The American Association for the Advancement of Science (AAAS) is the world's largest general scientific society. At its upcoming 2018 Annual Meeting, to be held in Austin, Texas this February, the mathematics section of AAAS will sponsor a symposium on mathematical approaches to major challenges in public health, environmental stewardship, and ecology. The symposium, titled "Mathematics of Planet Earth: Superbugs, Storm Surges, and Ecosystem Change," is organized by Hans Kaper and Hans Engler of Georgetown University.

Symposium speakers Glenn Webb (Vanderbilt University), Corina Tarnita (Princeton University), and Clint Dawson (University of Texas, Austin) will demonstrate how mathematical modeling and computation—coupled with new methods of gathering data—can predict, explain, or reconstruct phenomena such as the spread of diseases, storm surges, and vegetation patterns in landscapes under ecological stress. The same approaches can also assess the effectiveness of vaccination campaigns or evacuation plans, and identify early warning signals of ecological change.

During his talk on new developments in mathematical epidemiology, Webb will discuss models for diseases in individuals and communities that can track the dynamics of infectious pathogens through time and space. These models help evaluate interventions such as vaccination, quarantine, or medical treatment.

Tarnita will speak about the formation of spatial patterns in ecological systems like savannahs and shrublands. Mathematical models and data can help identify transitions of characteristic vegetation patterns that may act as early warning indicators for impending catastrophic changes, such as collapse to desert in arid regions.

Dawson will report on efforts to improve the resilience of coastal ecosystems to challenges like storm surges and land loss. Modeling resulting from collaborations spanning various fields—including mathematics, computational science, engineering, and environmental science—can help plan for the protection of coastal communities against increasing threats due to natural, ecological, and socioeconomic causes.

Collectively, the session will demonstrate successful applications of mathematical modeling and computational science, and joint efforts across disciplines to protect and improve human life and wellbeing — all part of an emerging effort to develop the "Mathematics of Planet Earth."

Figure 4a displays the contours of the probability density function describing the occurrence of wave groups with length scales in the x- and y-direction—L_x and L_y respectively—and amplitude A_0. These wave groups are the direct consequence of the dispersive mixing between different wavenumbers. The gray surface marks the parametric boundary above which individual wave groups become unstable. Figure 4b shows the maximal growth occurring over finite time. The parameter space's low-dimensionality allows us to easily identify a critical pair of length scales that can be tracked to forecast waves with realistic probability of occurrence but also significant growth over finite times. The resulting prediction scheme permits the data-driven prediction of rogue waves occurring due to nonlinear effects without solving any wave equations. This strategy predicts the rogue waves on average approximately 100 wave periods ahead of time [1,4].

We conclude with the following remarks:

(i) For complex high-dimensional systems, knowledge of the physical model does not imply knowledge of the mechanism underlying extreme events. Constrained optimization (1) offers one systematic method for the discovery of precursors to extreme events by carefully probing the dynamical system's state space.

(ii) Direct numerical simulations, although insightful, are not adequate for understanding extreme events. Only a low-dimensional subset of the many interacting degrees of freedom in high-dimensional systems contributes to extreme event formation. However, because of the complex coupling among all degrees of freedom, it is unclear how one would implement a data-driven approach to isolate the ones that underpin extreme events. Nevertheless, such a data-driven analysis merits further investigation.

(iii) Knowledge of the mechanism underpinning extreme events may enable the construction of indicators whose measurement permits the data-driven prediction of upcoming extreme events. In other words, upon discovering an indicator of extreme events from the model-based optimization (1), prediction can be accomplished in a completely data-driven fashion without resorting to the model.

Acknowledgments: We acknowledge support from Office of Naval Research grant N00014-15-1-2381, Air Force Office of Scientific Research grant FA9550-16-1-0231, and Army Research Office grant 66710-EG-YIP.

References

[1] Cousins, W., & Sapsis, T.P. (2016). Reduced-order precursors of rare events in unidirectional nonlinear water waves. *J. Fluid Mech, 790*, 368-388.

[2] Dysthe, K., Krogstad, H.E., & Müller, P. (2008). Oceanic rogue waves. *Annu. Rev. Fluid Mech., 40*, 287-310.

[3] Farazmand, M. (2016). An adjoint-based approach for finding invariant solutions of Navier-Stokes equations. *J. Fluid Mech., 795*, 278-312.

[4] Farazmand, M., & Sapsis, T.P. (2017). Reduced-order prediction of rogue waves in two-dimensional deep-water waves. *J. Comput. Phys., 340*, 418-434.

[5] Farazmand, M., & Sapsis, T.P. (2017). A variational approach to probing extreme events in turbulent dynamical systems. *Sci. Adv., 3*(9), e1701533.

[6] Lucarini, V., Faranda, D., de Freitas, A.C.G.M.M., de Freitas, J.M.M., Holland, M., Kuna, T.,...,Vaienti, S. (2016). *Extr. Recur. Dynam. Syst*. New York, NY: Wiley.

Mohammad Farazmand is a postdoctoral associate at the Massachusetts Institute of Technology. Themistoklis Sapsis is an associate professor of mechanical engineering at the Massachusetts Institute of Technology.

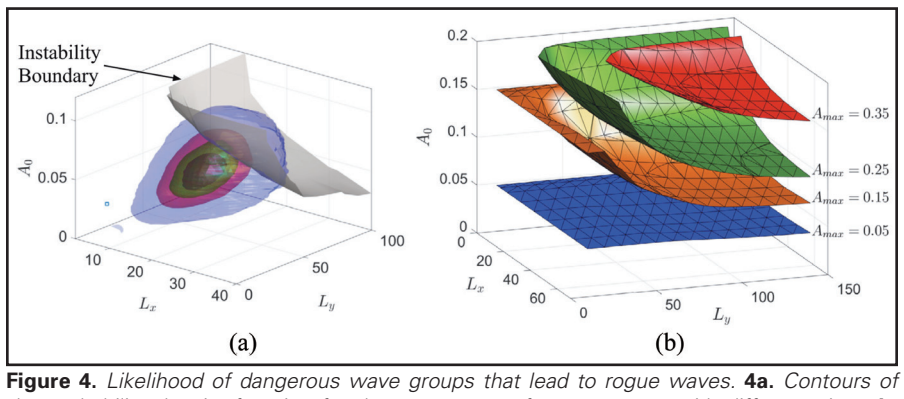

Figure 4. *Likelihood of dangerous wave groups that lead to rogue waves.* **4a.** *Contours of the probability density function for the occurrence of wave groups with different size.* **4b.** *Contours of maximum finite-time growth of wave groups due to nonlinear focusing phenomena. The combination of statistics and dynamics offers a critical set of length scales relevant to the prediction of extreme wave groups. View the online version of this article to see how a random wave field is generated and propagated under the modified nonlinear Schrödinger equation. Adapted from [4].*

Methane Hydrate

Continued from page 1

sensitive. While many mathematicians are versed in Stefan nonlinear free-boundary problems, it takes an investment to understand, analyze, and build multicomponent models in which the temperature is at the backstage rather than at the center (where mass fractions rule). Further, computational models for hydrates are quite sensitive to their data due to both their singularity and the relatively narrow envelope of physically-meaningful solutions. To work around the sensitivity, one must analyze the models and their computational counterparts.

The study of gas hydrates is relatively uncharted territory for computational mathematics. While investigating solutions to the comprehensive, coupled, multiphase multicomponent models is out of reach for well-posedness analysis, one can study subproblems that focus on the primary difficulties while freezing the model's other elements. This strategy works particularly well for basin modeling at large timescales in which the hydrate evolution (almost) fits in the traditional framework of free boundary problems. However, there is a snag. Unlike freezing of water (melting of ice), which always occurs around 0°C, the phase behavior for hydrate formation (dissociation) is associated with variable pressure and temperature conditions. This challenge requires

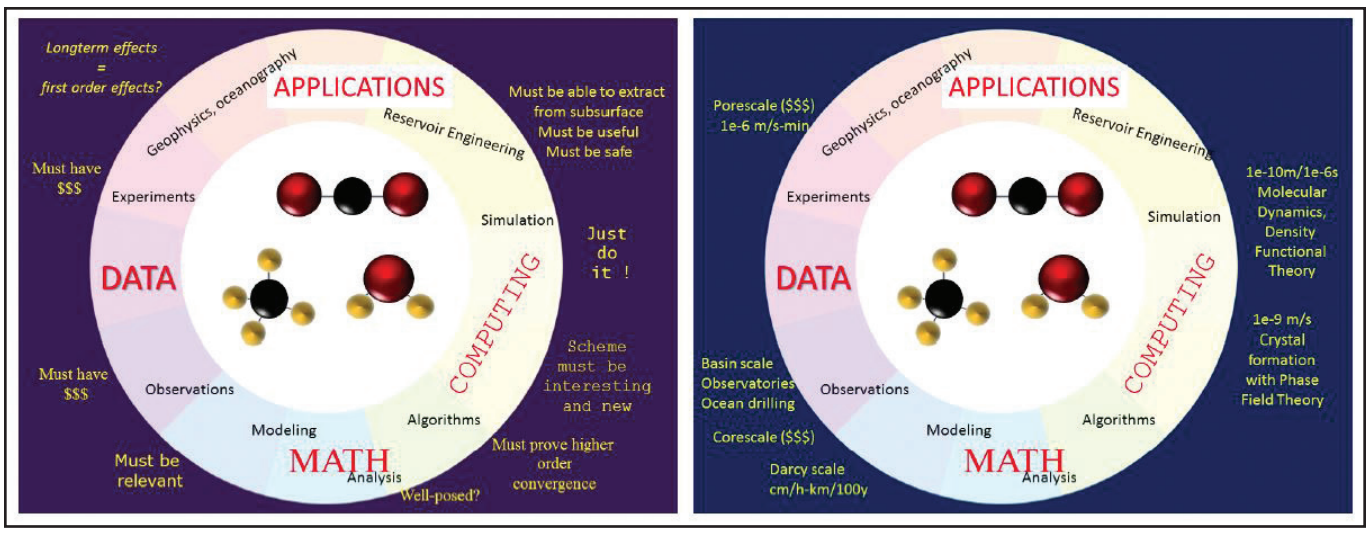

Figure 2. *Interdisciplinary objectives and multiple scales in gas hydrate studies and modeling. The molecules pictured in the middle include water (H_20), methane (CH_4), and carbon dioxide (CO_2); the latter is present because of its relevance to methane production and carbon sequestration by the process of molecule exchange [10]. Image credit: Malgorzata Peszynska.*

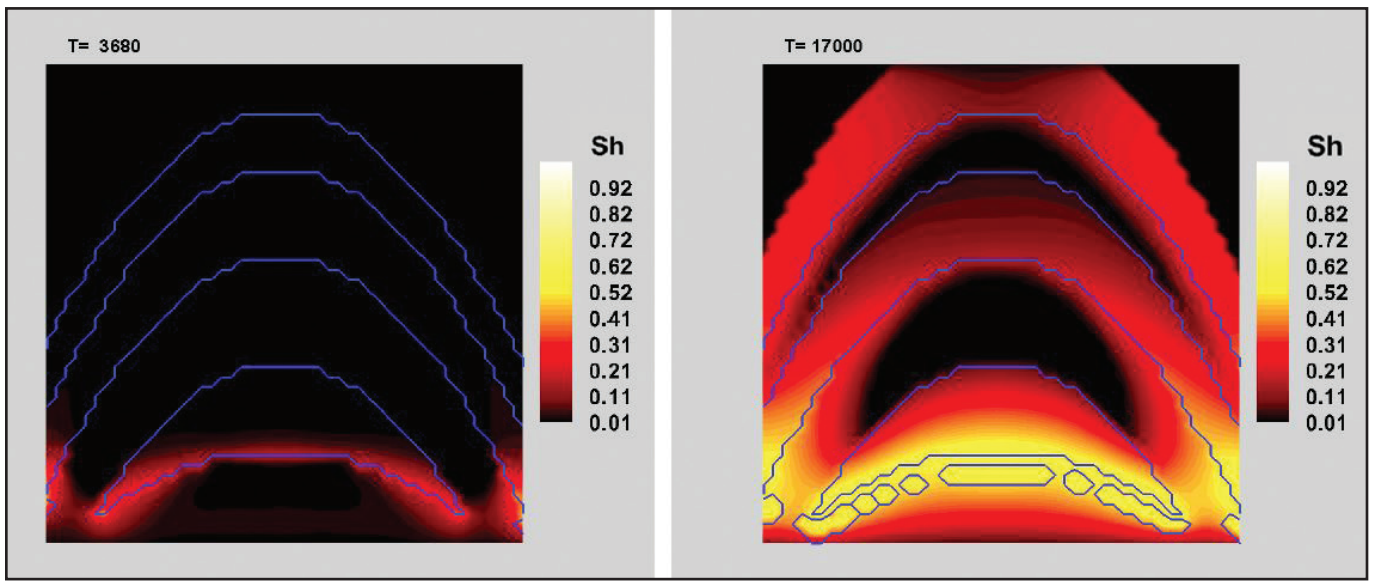

Figure 1. *This two-dimensional simulation is inspired by a one-dimensional case study [2], while geometry is modeled after that of Hydrate Ridge. The red-to-yellow colors show the increase of methane hydrate saturation (Sh) caused by the migration of methane upwards from the bottom of the reservoir into the hydrate zone. The blue lines delineate rock types; initially, they differentiate between fine-grained (clay) and coarse-grained (silt) sediment. Fractures form at the bottom over time. The uneven migration and hydrate formation are due to the heterogeneity of the rock and the complicated dependence of thermodynamics on the host rock. Image credit: Malgorzata Peszynska.*

a special convex analysis construction [3, 7]. In addition, while transport by diffusion is unlikely to have singular solutions (similarly to the temperature in the Stefan problem), strong advection fluxes lead to large discontinuous deposits observed in nature and require a very weak notion of solutions to the partial differential equations.

Another challenge accompanies the burden (and beauty) of multiple scales. In many places around the world, the presence of hydrates has been inferred rather than confirmed. Unlike the free gas deposits visible in seismic images of the subsurface, the hydrate cannot be easily seen because its density is close to that of water, necessitating modeling to assist the interpretation of well core data [6]. Furthermore, the response of sediment to seismic waves depends on the microscopic (pore-scale) distribution of the hydrate and whether it cements the distance between host rock grains, thus altering the elastic response. Without detailed models at the pore-scale, the unusual pockmark or chimney-like morphology of some deposits is hard to explain. However, at and below the pore-scale, the mathematical modeling of methane hydrate is still an art rather than a science. While some phase-field and molecular dynamics models have emerged [9], they are still far from the comfortable framework in which we prove theorems.

Going down in pore-scales from centimeter to meter or kilometer to micron or nanometer has become fairly established in petroleum engineering and hydrology. Geoscientists and imaging experts design flow experiments and carry out X-ray computed tomography imaging that displays the evolution of fluids in the void space between rock grains in either the fully-opaque three-dimensional porous media or two-dimensional micromodels, which look like pinball structures. The virtual laboratory for hydrates at the pore-scale is nearly there. Unfortunately, hydrates are unstable in standard conditions, rendering a bank of images and experimental data both costly and difficult to acquire. First-principles modeling can therefore play a significant role.

Methane hydrate is also of interest as a potential energy source. In the last decade, the U.S. (at the Ignik-Sikumi well in Alaska), Japan, Korea, and India have implemented and explored hydrate projects [1]. However, technological difficulties and many open questions remain — the models used for production require a very different timescale compared to basin modeling, and demand another degree of complexity.

In summary, the overarching challenge in studying hydrates is arguably not just the complexity of the problem itself, but rather the ability to simplify and extract subproblems so as to move forward and make progress without compromising the results. For applied and computational mathematicians, the process of translating the hydrate model so it can fit into a mathematical framework amenable to analysis and simulation can be very rewarding. However, it requires vigilance from all participants of the interdisciplinary team. As we make simplifying assumptions that enable intricate mathematics, we must be careful not to render the model inapplicable. When implementing numerical schemes for the comprehensive model, it is hard to know if the visible discontinuities are real, attributable to deficiencies of the scheme, or due to poor resolution of the phase behavior data or solver. Hence, it is essential for all team members to understand each other's objectives, respectfully acknowledge the existence of knowledge gaps, and carefully fill them in. Progressing from basin to production timescales—or from sediment depths to the scale of microcracks and gas chimneys—requires more resources across several fields. Observations and data are needed, so modeling can at least be guided even if validation is impossible.

References

[1] Boswell, R. (2013). Japan completes first offshore methane hydrate production test — methane successfully produced from deepwater hydrate layers. *Cent. Nat. Gas Oil, 412*, 386-7614.

[2] Daigle, H., & Dugan, D. (2011). Capillary controls on methane hydrate distribution and fracturing in advective systems. *Geochem. Geophys. Geosyst., 12*(1).

[3] Gibson, N.L., Medina, F.P., Peszynska, M., & Showalter, R.E. (2014). Evolution of phase transitions in methane hydrate. *Journ. Math. Anal. App., 409*(2), 816-833.

[4] Hong, W., Torres, M.E., Carroll, J., Crémière, A., Panieri, G., Yao, H., & Serov, P. (2017). Seepage from an arctic shallow marine gas hydrate reservoir is insensitive to momentary ocean warming. *Nat. Comm., 8*, 15745.

[5] Liu, X., & Flemings, P.B. (2007). Dynamic multiphase flow model of hydrate formation in marine sediments. *Journ. Geophys. Res.: Solid Earth, 112*(B3).

[6] Peszynska, M., Hong, W., Torres, M.E., & Ji-Hoon, K. (2016). Methane hydrate formation in ulleung basin under conditions of variable salinity: Reduced model and experiments. *Transp. Porous Med., 114*(1), 1-27.

[7] Peszynska, M., Showalter, R.E., & Webster, J.T. (2015). Advection of methane in the hydrate zone: model, analysis and examples. *Math. Meth. Appl. Sci., 38*(18), 4613-4629.

[8] Ruppel, C.D., & Kessler, J.D. (2017). The interaction of climate change and methane hydrates. *Rev. Geophys., 55*(1), 126-168.

[9] Tegze, G., Pusztai, T., Tóth, G., Gránásy, L., Svandal, A., Buanes, T.,…, Kvamme, B. (2006). Multiscale approach to CO2 hydrate formation in aqueous solution: Phase field theory and molecular dynamics. Nucleation and growth. *Journ. Chem. Phys., 124*(23), 234710.

[10] White, M., & Suk Lee, W. (2014). Guest Molecule Exchange Kinetics for the 2012 Ignik Sikumi Gas Hydrate Field Trial. In *Offshore Technology Conference*. Houston, TX.

Malgorzata Peszynska is a professor of mathematics at Oregon State University. She received her Ph.D. from the University of Augsburg in 1992, and has held various positions at the Polish Academy of Sciences, Warsaw University of Technology, Purdue University, and The Institute for Computational Engineering and Sciences at the University of Texas at Austin.

Sign Up for SIAM's Science Policy Alerts

As most SIAM members are aware, science policy decisions greatly affect the state of scientific research. Hence, SIAM plays an active role in advocacy for federal support of mathematical and computational science research. SIAM conducts a wide range of activities to ensure sound scientific policy in the society's priority areas and the research community's interests.

Through its Committee on Science Policy, SIAM helps influence congressional legislation and federal programs related to applied mathematics and computational science, meets with relevant federal agency leaders and policymakers to better understand the research climate, and provides input on issues of concern to the SIAM community.

SIAM members are indispensable in this effort to safeguard the interests of the field. Community involvement and advocacy holds decision-makers accountable and ensures continued support for programs and policies vital to the discipline.

The scientific community's successful push to preserve graduate student tax benefits is a recent example. When the House-passed bill threatened to tax tuition waivers used by graduate students to offset educational costs, students across the country staged walkouts and protests. This motivated several Republican House members to urge that the provision be excluded from the bill's final version.

Over 60 scientific and engineering societies sent a letter to Congress opposing the tuition waver tax, emphasizing that the proposition would increase the financial burden on graduate students and hinder the pursuit of STEM degrees. Several organizations also issued action alerts encouraging members to implore their representatives to uphold the tax-exempt status of graduate tuition waivers. Ultimately, the provision was eliminated from the Senate version of the bill, which was signed into law by President Trump at the end of 2017.

In order to keep members informed and engaged regarding important issues affecting the discipline—and offer guidelines on how they can take action—SIAM provides timely information to the community via the science policy electronic mailing list.[1]

Sign up for the mailing list!

[1] http://www.siam.org/about/science/sci_policy_form.php

Graph Representations of Fractured Media in the Subsurface

By Gowri Srinivasan

Fractures are the primary pathways for fluid flow in otherwise low-permeability subsurface media, such as shale or granite. Applications where flow and transport through fractured media are central to informed decision-making via predictive capability have recently inspired a great deal of interest. Due primarily to the increased availability of natural gas, which produces 50 percent less carbon dioxide than coal, U.S. emissions have dropped to their lowest levels in 20 years. As a result, the production of hydrocarbons from shale formations—which involves hydrofracturing rock to establish fracture connectivity and extracting the natural gas flowing out of the fractures—is a topic of ongoing debate and research. Moreover, as countries like North Korea continue to improve their ability to conduct low-yield nuclear tests, chemical signature detection (e.g., xenon migration through fractured rock) provides the definitive smoking gun when used with conventional seismic methods.

Ignoring the topological and geometrical properties of these fracture networks and modeling subsurface fractures as a continuum with certain effective properties falls short of capturing key behaviors [6], including flow channeling and arrival times of particles advecting with the flow. It is thus desirable to retain the underlying structure of these fractured systems, resulting in the emergence of an alternative discrete fracture network (DFN) modeling methodology. Recognition that fracture geometry and network topology play a critical role in determining quantities of interest relating to flow and transport in fractured subsurface media is a key distinguishing factor of DFN modeling from standard continuum models.

Until recently, DFN models were limited to one-dimensional pipe-network approximations [2], two-dimensional systems, or relatively small three-dimensional (3D) systems [1]. However, recent advances in high-performance computing have enabled flow and transport simulations in large, explicit 3D DFN representations [5]. The 3D DFN software dfnWorks [5], which recently won an R&D 100 Award, assigns each network fracture a shape, location, aperture, and orientation by sampling distributions whose parameters are based on site characterization. Figure 1a shows a typical network consisting of 7,200 fractures of varying size, orientation, and permeability. Once meshed, this system has nearly 16 million nodes and 32 million triangular elements. Figure 1b depicts a close-up of the mesh representation; the flow equations are solved explicitly on this detailed representation. The increase in model fidelity comes at a huge computational cost because of the large number of mesh elements required to represent thousands of fractures with sizes that range over several orders of magnitude (from mm to km) and accurately resolve large pressure gradients at fracture intersections. Given that most sites can only be characterized statistically, computing uncertainty bounds for the purpose of decision-making requires thousands of runs.

One can think of a graphical representation of a fracture network, where nodes and edges inherit geophysical and geometric properties of that network, as a coarse-scalerepresentation that preserves the key topological properties [3]. Optimal assign-

See **Graph Representations** *on page 7*

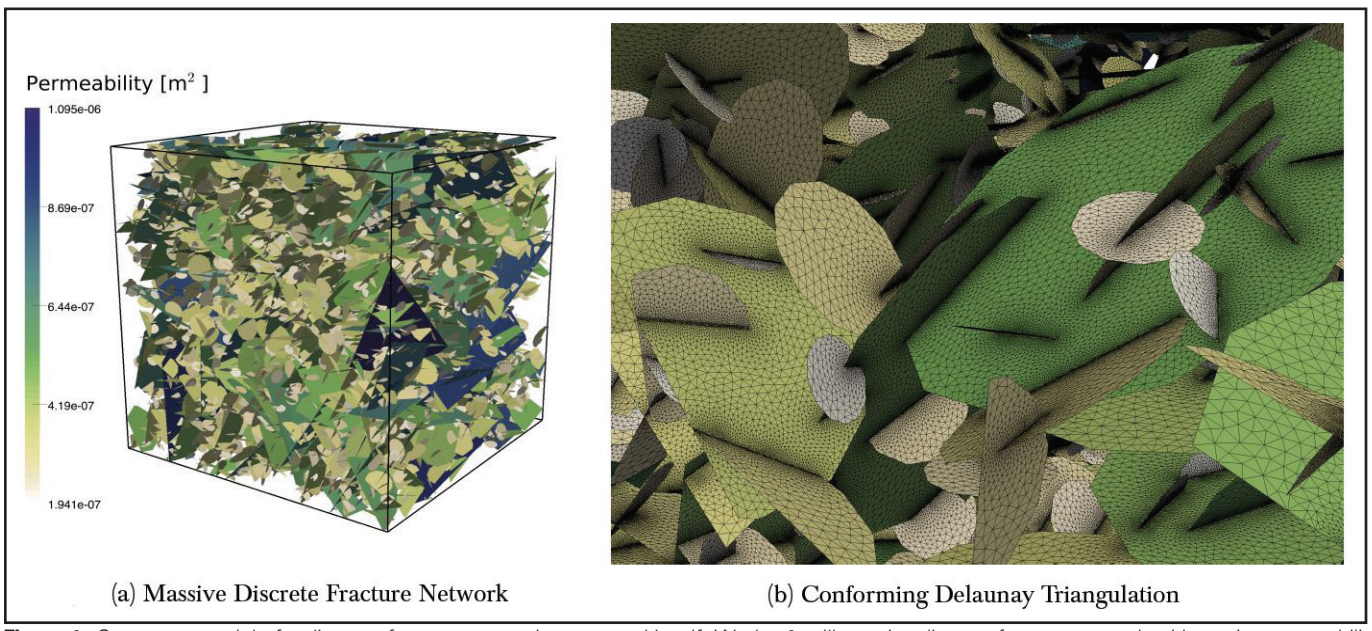

Figure 1. *Computer model of a discrete fracture network generated by dfnWorks.* **1a.** *Illustrative discrete fracture network with varying permeabilities.* **1b.** *Close-up of the meshing on the network. Image credit: Jeffrey Hyman.*

Randomization in Characterizing the Subsurface

By Youzuo Lin, Daniel O'Malley, Velimir V. Vesselinov, George Guthrie, and David Coblentz

Current methods for characterizing Earth's subsurface, such as standard inverse techniques, are not sufficiently accurate to meet the needs of modern applications in the fields of energy exploration, environmental management, and global security. While increasing the quantity of field measurements and robustness of the applied data-/model-analysis methods can improve accuracy, such approaches can be computationally impractical for large data sets and complex site conditions. Therefore, there is a need to develop economically-feasible and robust computational methods while maintaining accuracy. For example, in-field drilling for geothermal operations may yield high failure rates, resulting in unacceptably high costs; errors and/or large uncertainties in the estimated subsurface characteristics are the main impediment to the successful siting of an in-field well. This problem is not uniquely geothermal. Accurate characterization of uncertain subsurface properties is also critical for monitoring storage of carbon dioxide, estimating pathways of subsurface contaminant transport, and supervising ground-based nuclear-explosion tests.

We have developed various methods to characterize the subsurface, including efficient computational strategies to identify subsurface permeability given a set of hydraulic heads, as shown in Figure 1 (on page 1), and a data-driven subsurface geological feature detection approach using seismic measurements, as shown in Figure 2 (right). A major challenge for many subsurface applications is the large number of observations and high feature dimensionality.

Randomized matrix algorithms—which aim to construct a low-rank approximation of an input matrix—have received a great deal of attention in recent years. The low-rank approximation, often called a matrix "sketch," is usually the product of two smaller matrices, which yields a good approximation that represents the original output's essential information. Therefore, one can employ a sketching system as a surrogate for the original data to compute quantities of interest. We have employed randomization techniques to solve various large-scale computational problems. Here we provide examples to demonstrate two major applications in solving real-world subsurface problems.

Randomized Subsurface Permeability Estimation

A porous medium's permeability is a physical quantity needed to predict flow and transport of fluids and contaminants in the subsurface. The permeability's estimation is often posed as a regularized inverse problem

$$\tilde{x} = argmin_{\mathbf{x}}\{\|\mathbf{d} - f(\mathbf{x})\|_R^2 + \lambda \|\mathbf{x} - X\boldsymbol{\beta}\|_Q^2\}, \quad (1)$$

where f is the forward operator mapping from the permeability to the pressure (called "hydraulic head" in hydrology parlance), $\mathbf{d}$ is a recorded hydraulic head dataset, $\mathbf{x}$ is a vector of permeabilities, $\|\mathbf{d} - f(\mathbf{x})\|_R^2$ measures the data misfit, and $\|\mathbf{x} - X\boldsymbol{\beta}\|_Q^2$ is the regularization term.

The solution to (1) can be obtained as

$$\tilde{\mathbf{x}} = X\boldsymbol{\beta} + QH^T\boldsymbol{\varepsilon}, \quad (2)$$

where H is the Jacobian matrix of the forward modeling operator f, defined as

$$H = \frac{\partial f}{\partial \mathbf{x}}\Big|_{\mathbf{x}=\overline{\mathbf{x}}}. \quad (3)$$

One may obtain $\boldsymbol{\beta}$ and $\boldsymbol{\varepsilon}$ by solving the linear system

$$\begin{pmatrix} HQH^T + R & HX \\ (HX)^T & 0 \end{pmatrix}\begin{pmatrix} \boldsymbol{\varepsilon} \\ \boldsymbol{\beta} \end{pmatrix} = \begin{pmatrix} \mathbf{y} - f(\overline{\mathbf{x}} + H\mathbf{x}) \\ 0 \end{pmatrix}. \quad (4)$$

However, solving (4) can be both prohibitively expensive and memory demanding. To combat this problem, we developed a novel randomized technique that enables an efficient computational method [1].

Our approach aims to construct a sketching matrix, the elements of which are drawn randomly from a Gaussian distribution. We then replace the data $\mathbf{d}$ with $S\mathbf{d}$ and the forward $f(\mathbf{x})$ with $Sf(\mathbf{x})$. Therefore, the linear system in (2) and (4) can be substituted correspondingly with

$$\hat{\mathbf{x}} = X\boldsymbol{\beta} + QH^TS^T\boldsymbol{\varepsilon} \quad (5)$$

and

$$\begin{pmatrix} SHQH^TS^T + R & SHX \\ (SHX)^T & 0 \end{pmatrix}\begin{pmatrix} \boldsymbol{\varepsilon} \\ \boldsymbol{\beta} \end{pmatrix} = \begin{pmatrix} S(\mathbf{y} - f(\overline{\mathbf{x}}) + H\,\overline{\mathbf{x}}) \\ 0 \end{pmatrix}. \quad (6)$$

(2) and (5) and (4) and (6) seem almost identical, except for the introduction of matrix S. However, a simple computational cost analysis can reveal the significant impacts of the randomized matrix. Assume that the number of model parameters is $\tilde{m}$; the number of observations is $\tilde{n}$, which yields the size of the Jacobian matrix $H \in \mathbb{R}^{\tilde{n}\times\tilde{m}}$; and the covariance matrix is $Q \in \mathbb{R}^{\tilde{n}\times\tilde{m}}$. We also denote the rank of the sketching matrix by k, and $k \ll \tilde{n}$. The drift matrix $X \in \mathbb{R}^{\tilde{m}\times\tilde{p}}$, where $\tilde{p}$ is small. The dimension of the original system matrix in (4) is $((\tilde{n}+p)\times(\tilde{n}+p))$, while the dimension of the randomized system in (6) is $((k+p)\times(k+p))$ — much smaller than the original system. Therefore, the computational cost of solving (6) is significantly lower than that of solving (4); this is the power of randomization in solving traditional inverse problems, as illustrated in [1].

See **Randomization** *on page 6*

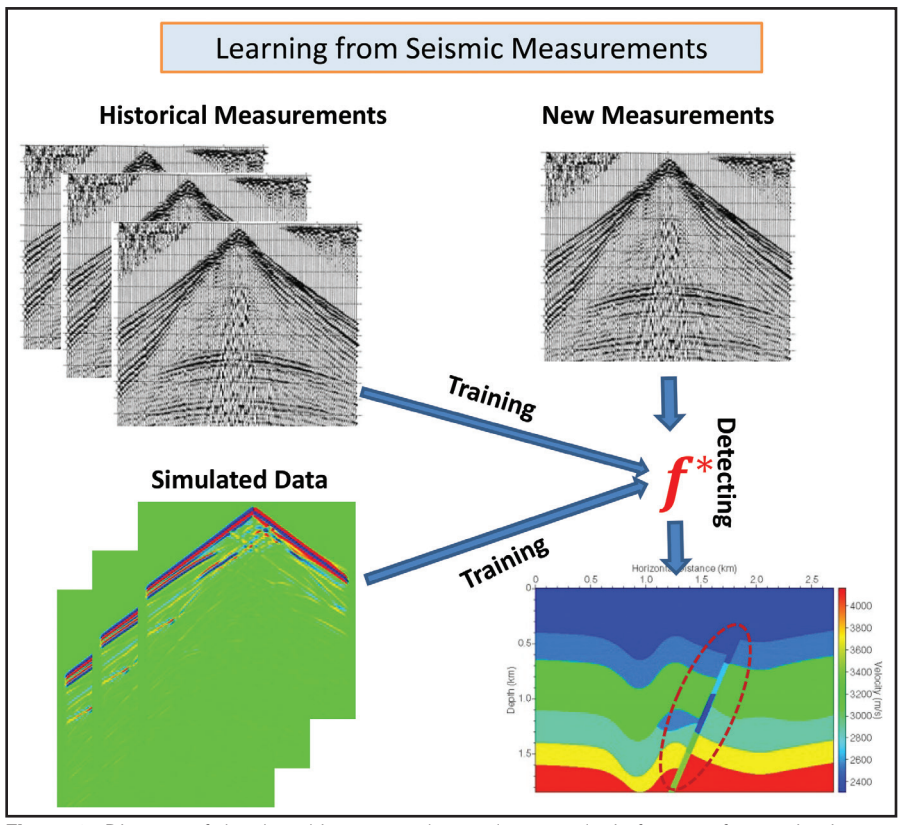

Figure 2. *Diagram of the data-driven procedure to learn geologic features from seismic measurements. Image courtesy of [2] and [3].*

Slings, Bullets, Blow-up, and Linearity

MATHEMATICAL CURIOSITIES
By Mark Levi

The sling is a simple weapon — essentially a pendulum with two strings instead of one, each of which holds a side of a cradle enclosing a projectile. When one spins the pendulum and releases a string, the discharged projectile can travel over 400 meters. The sling was the world distance champion for millennia until the English long bow, and then the firearm, surpassed it. Speaking of the latter, a bullet shot into water has an unexpectedly short killing range despite its enormous speed.

Both of these phenomena—the surprisingly long killing distance of a sling and the surprisingly short killing distance of a water-bound bullet—involve near-infinite (on the human scale) velocities. Both can be "explained" by the fact that solutions of the simple ordinary differential equation (ODE) $\dot{x}=x^2$ blow up in finite time, something usually presented as a mere scholastic curiosity.

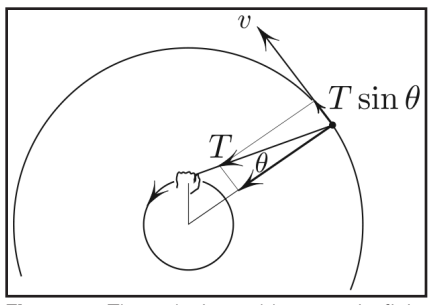

Figure 1. *The velocity v blows up in finite time $T=$ the tension of the string, giving rise to the accelerations $a_{\text{tangential}}$ and $a_{centripetal}$.*

Launching the Sling

The caricature of the sling in Figure 1 shows the hand moving in a circle *so as to make the cradle describe a concentric circle* (to achieve this, the hand must accelerate a certain way). I claim that the speed v then satisfies the ODE mentioned above: $\dot{v}=kv^2$, with a constant k. Indeed, Figure 1 gives us $a_{\text{tangential}}=a_{centripetal}\tan\theta$; and since $a_{\text{tangential}}=\dot{v}$ and $a_{centripetal}=v^2/R$, this yields

$$\dot{v}=kv^2, \quad k=\tan\theta/R.$$

And the solution

$$v=\frac{v_0}{1-v_0kt}$$

does in fact approach infinity as $t\to t_\infty=1/v_0k$ in the idealized non-relativistic world of infinitely fast hands and infinitely strong strings. The "shadow" of this infinite speed is seen in the fact that a good slinger can launch at over $1/5$th the speed of sound (computed from the aforementioned >400 m distance. In fact, the actual speed is higher, given that the calculation ignores the air resistance).

Shooting into the Water

Continuing with the weapons-related theme, let us ask: what depth renders harmless a bullet shot down into the water? Assume that (i) the water drag on the bullet is proportional to the square of the speed; (ii) the terminal sinking velocity of the bullet is 1 m/sec, and (iii) the safe velocity of the bullet is $v_s=10$ m/sec (the speed gained in dropping 15 feet, painful but probably not dangerous). With these assumptions, the safe depth turns out to be

$$x_s\approx 0.1\ln\Big(1+\frac{v_0}{v_s}\Big), \qquad (1)$$

in meters. Before deriving this expression, let us find safe depths for various bullet speeds. For $v_0=10^3$ m/sec (about three times the speed of sound), $x_s\approx 46$ cm. For the escape velocity $v_0\approx 11$ km/sec, $x_s\approx 69$ cm. Every increase in the order of magnitude simply adds about 23 cm to the safety depth. And for a bullet travelling at the speed of light (here I am abandoning the last touch with reality), the safe depth is just 2 m! This drastic loss of speed of the bullet is identical, up to the time-reversal, to the sling projectile's drastic gain of speed.

Derivation of (1)

The bullet shot straight down obeys Newton's law: $m\dot{v}=-cv^2+mg$, or

$$\dot{v}=-kv^2+g, \quad k=\frac{c}{m}, \qquad (2)$$

where $v=\dot{x}$ and the x-axis points down. From the terminal velocity condition $-kv_{\text{term}}^2+g=0$, we find $k=g/v_{\text{term}}^2\approx 10m^{-1}$. Neglecting g in (2), we get

$$\dot{v}=-kv^2, \quad k\approx 10m^{-1}.$$

Substituting the initial condition $v(0)=v_0$ and the time t_s of reaching velocity v_s into the solution of this ODE gives

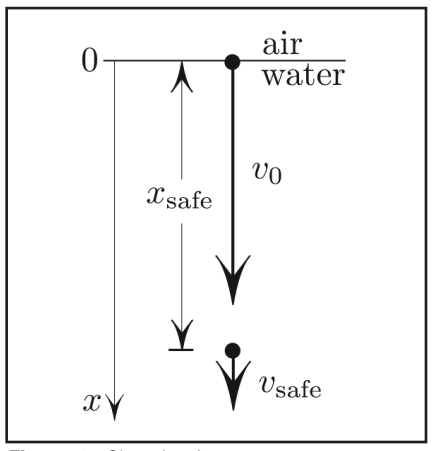

Figure 2. *Shooting into water.*

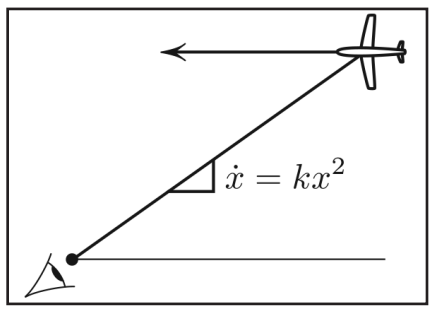

Figure 3. *Geometry of the finite-time blow-up. The slope x of the line of sight satisfies the "sling" ODE.*

$$v_s=\frac{v_0}{1+kv_0t_s}\Rightarrow t_s=k^{-1}(v_s^{-1}-v_0^{-1})<\frac{1}{kv_s}.$$

Integrating v gives $x=k^{-1}\ln(1+kv_0t)$, and substituting $t_s<1/kv_s$ into this expression gives (1).

The Hidden Linearity

The ODE $\dot{x}=x^2$ (a special case of the Riccati equation) hides linear growth, which can be expressed in two equivalent ways. Algebraically, the equation simply amounts to the linear growth of $1/x$, namely to $\frac{d}{dt}\frac{1}{x}=1$. Geometrically, this ODE governs the evolution of the slope $x=v/u$ of solution vectors in the (u,v)-plane of the trivial system $\dot{u}=-1$, $\dot{v}=0$. Figure 3 offers an essentially equivalent realization. The blow-up of the solution occurs when the plane is overhead.

As a concluding remark, the assumption $dv/dt=-kv^2$ amounts to a more natural-sounding statement: the kinetic energy E decays exponentially with the distance $dE/dx=-2kE$.

The figures in this article were provided by the author.

Mark Levi (levi@math.psu.edu) is a professor of mathematics at the Pennsylvania State University.

Randomization

Continued from page 5

The developed methods are available in the open source code Mads.[1]

Subsurface Geological Feature Detection Using Randomized Data-Driven Methods

Seismic waves are more sensitive to the acoustic/elastic impedance of the subsurface than other geophysical measurements (see Figure 2, on page 5). Hence, seismic exploration has been widely used to infer heterogeneities in media impedance, which indicate geologic structures.

Analyzing and interpreting seismic measurements for identifying prospective geological features is challenging. The difficulties arise from the processing of large amounts of seismic data and the incorporation of subjective human factors. Different geologic features play different roles in characterizing subsurface structure. In particular, identifying geological fault zones is essential to many subsurface energy applications. In carbon sequestration, potential leaks of stored carbon dioxide can create geologic faults, so knowing fault locations is necessary to monitor carbon dioxide storage. We have developed a novel data-driven geological feature detection method and successfully applied it to seismic measurements [2, 3], as illustrated in Figure 2 (on page 5). Both historical and simulated seismic data are fed into learning algorithms. A detection function $f^*(\mathbf{x})$ is the output of the training process, where $\mathbf{x}$ represents the pre-stack seismic measurements. The function creates a link from the seismic measurements to the corresponding geological features.

Suppose one has n historical feature vectors $\mathbf{X}=[\mathbf{x_1^T}\ \cdots\ \mathbf{x_n^T}]^\mathbf{T}\in\mathbb{R}^{n\times d}$, which are from seismic measurements and $\mathbf{x}_i\in\mathbb{R}^d$, and the associated labels $\mathbf{y}=[y_1\ \cdots\ y_n]^T\in\mathbb{R}^{n\times 1}$, which in this example denote the location of the dipping angle of geologic faults. The kernel ridge regression (KRR) is utilized to learn the mapping function [2, 3]. We directly state the dual problem of KRR without derivation

$$\alpha=argmin_\alpha\Big\{\frac{1}{2}\Sigma_{i=1}^n\parallel y_i - (K\alpha)_i\parallel_2^2+\frac{\lambda}{2}\alpha^TK\alpha\Big\}, \qquad (7)$$

where K is a kernel function and $\lambda>0$ is a regularization parameter. The problem in (7) has a closed-form solution

$$\alpha^\star=(K+\lambda I_n)^{-1}\mathbf{y}\in\mathbb{R}^\mathbf{n}, \qquad (8)$$

where I_n is a $n\times n$ identity matrix. Finally, for any unknown data $\mathbf{x}'\in\mathbb{R}^d$, the prediction made by KRR can be obtained by

$$f(\mathbf{x}')=\Sigma_{i=1}^n\alpha_i^\star\kappa(\mathbf{x}',\mathbf{x}_i). \qquad (9)$$

However, the direct utilization KRR prediction in (7) is computationally expensive, because of the inversion of the large-scale matrix in (8). We employ the Nyström method—a randomized kernel matrix approximation tool—to the geologic detection task, aiming to solve large-scale problems using modest computational resources.

The Nyström method computes a low-rank approximation $K\approx\psi\psi^T$ in $\mathcal{O}(nds+ns^2)$ time. Here, $s\ll n$ is user-specified; larger values of s lead to better approximation but incur higher computational costs. We can compute the tall-and-skinny matrix $\psi\in\mathbb{R}^{n\times s}$ as follows. First, we sample s items from $\{1,\cdots,n\}$ uniformly at random without replacement; let the resulting set be $\mathcal{S}$. Subsequently, we construct a matrix $C\in\mathbb{R}^{n\times s}$ as $c_{il}=\kappa(\mathbf{x}_i,\mathbf{x}_l)$ for $i\in\{1,\cdots,n\}$ and $l\in\mathcal{S}$; let $W\in\mathbb{R}^{s\times s}$ contain the rows of C indexed by $\mathcal{S}$. Figure 3 illustrates the approximation. Finally, we compute the low-rank approximation $\psi=C(W^\dagger)^{1/2}$.

With the low-rank approximation obtained via the Nyström method, we can efficiently calculate an approximated solution

$$\tilde{\alpha}=(\psi\psi^T+\lambda I_n)^{-1}\mathbf{y}, \qquad (10)$$
$$=\lambda^{-1}\mathbf{y}-\lambda^{-1}\psi(\lambda I_s+\psi^T\psi)^{-1}\psi^T\mathbf{y},$$

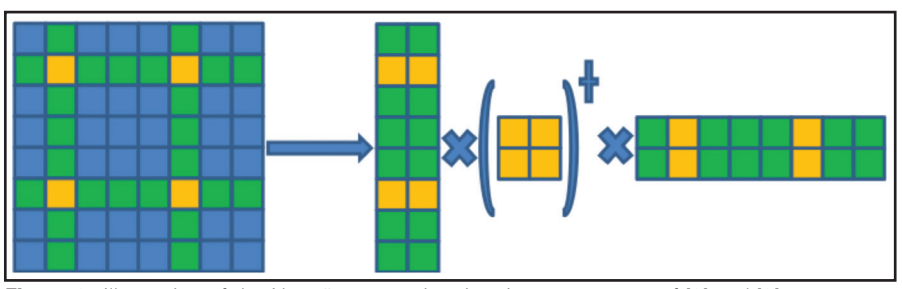

Figure 3. *Illustration of the Nyström approximation. Image courtesy of [2] and [3].*

where the latter equality follows from the Sherman-Morrison-Woodbury matrix identity. It is worthwhile mentioning that the $n\times n$ matrix of $\psi\psi^T$ in (8) has been replaced by the matrix of $\psi^T\psi\in\mathbb{R}^{s\times s}$, which is much smaller. This significantly reduces the computational costs. More details and results can be found in [2, 3].

Acknowledgements: This work was co-funded by the U.S. Department of Energy (DOE) Office of Fossil Energy's Carbon Storage program, Los Alamos National Laboratory Environmental Programs, and the DOE Office of Science (DiaMonD project: An Integrated Multifaceted Approach to Mathematics at the Interfaces of Data, Models, and Decisions, Grant #11145687).

References

[1] Lin, Y., Le, E., O'Malley, D., Vesselinov, V., & Bui-Thanh, T. (2017). Large-Scale Inverse Model Analyses Employing Fast Randomized Data Reduction. *Wat. Resc. Res, 53*(8), 6784-6801.

[2] Lin, Y., Wang, S., Thiagarajan, J., Guthrie, G., & Coblentz, D. (2017). Efficient Data-Driven Geologic Feature Detection from Pre-stack Seismic Measurements using Randomized Machine-Learning Algorithm. Preprint, *arXiv:1710.04329*.

[3] Lin, Y., Wang, S., Thiagarajan, J., Guthrie, G., & Coblentz, D. (2017). Towards Real-Time Geologic Feature Detection from Seismic Measurements using a Randomized Machine-Learning Algorithm. In *SEG Technical Program Expanded Abstracts 2017* (pp. 2143-2148). Houston, TX: Society of Exploration Geophysics.

Youzuo Lin and Daniel O'Malley are staff scientists in the Earth and Environmental Sciences Division at Los Alamos National Laboratory (LANL). Velimir Vesselinov is a staff scientist in the Earth and Environmental Sciences Division at LANL and a principle investigator of several Department of Energy-funded projects related to environmental management. George D. Guthrie is a geochemist in the Earth and Environmental Sciences Division at LANL. David Coblentz is a R&D Manager and staff scientist in the Earth and Environmental Sciences Division at LANL.

[1] http://mads.lanl.gov

Graph Representations

Continued from page 5

ment of edge weights, a topic of research itself, can enhance the effectiveness of the reduced graph representation to emulate the DFN. Furthermore, the weighted graph framework seamlessly lends itself to numerically solving flow and transport equations on the nodes, resulting in efficient and accurate DFN emulators to use instead of the computationally expensive DFN model. The reduced computational burden is particularly valuable in the context of uncertainty quantification.

Graphs provide a simple and elegant way to characterize, query, and utilize connectivity — one of the fundamental aspects of a fracture network. The suitable choice of mapping—which nodes represent attributes of the DFN and which attributes are edges connecting those nodes—depends on the question being asked. Figure 2 illustrates three possible graph representations that result from a small eight-fracture network. The most general representation in Figure 2a is a bipartite graph with two node sets, one representing the fractures (top) and the other representing the intersections (bottom), with edges connecting a fracture node to an intersection node if the fracture includes that intersection.

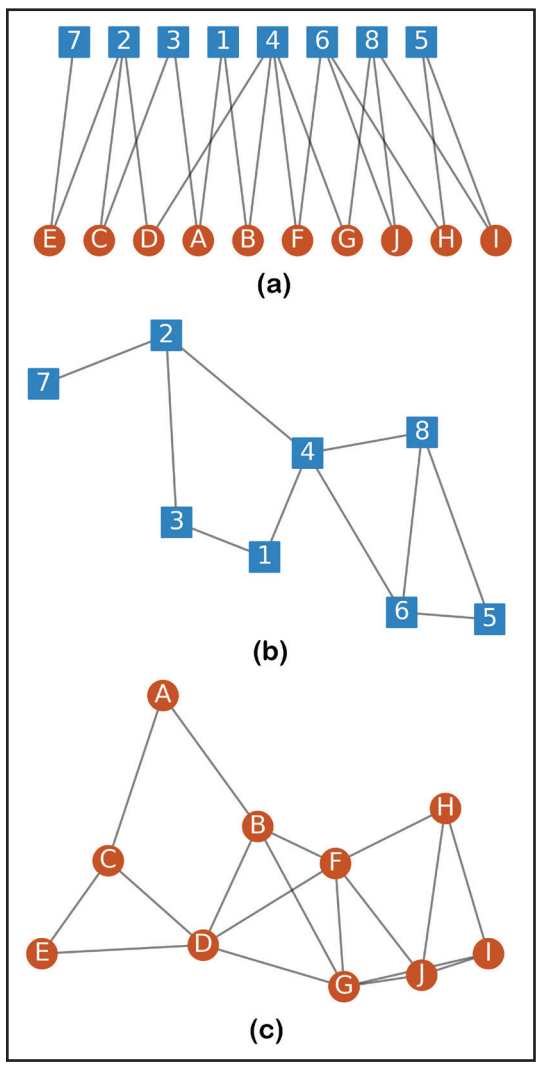

Figure 2. *Graph representations of a small eight-fracture network.* **2a.** *A bipartitite graph representation.* **2b.** *Each fracture is represented by a node, and each edge is an intersection of fractures.* **2c.** *Each fracture intersection is represented by a node, and fractures are a clique of edges. Image credit: Aric Hagberg.*

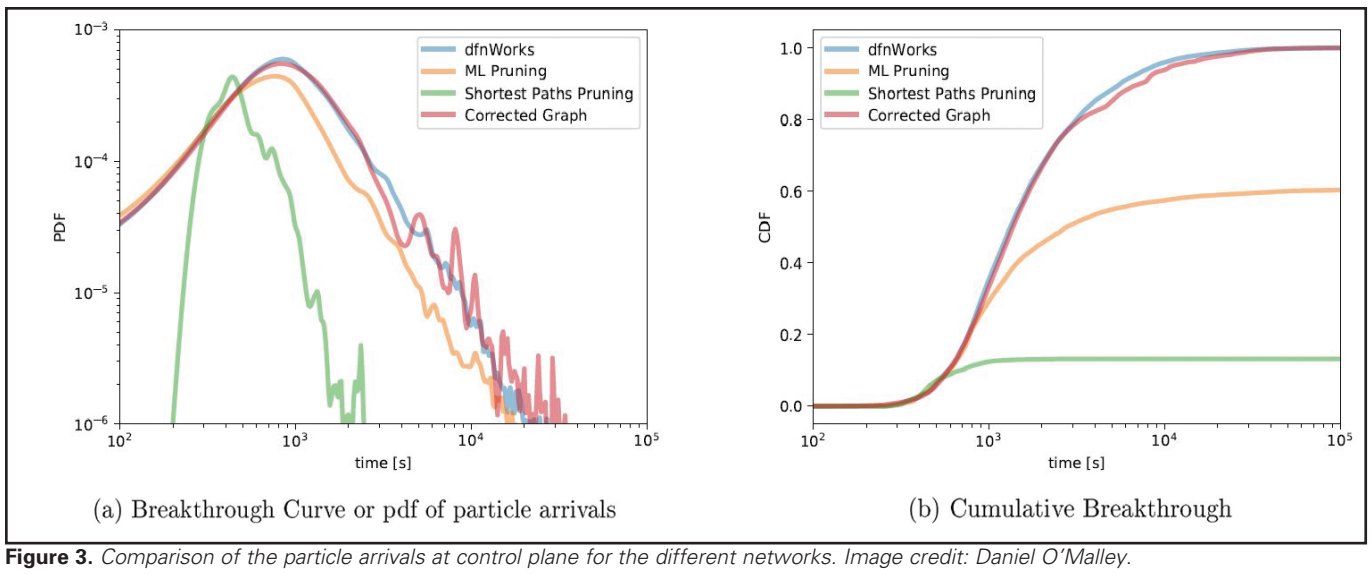

(a) Breakthrough Curve or pdf of particle arrivals

(b) Cumulative Breakthrough

Figure 3. *Comparison of the particle arrivals at control plane for the different networks. Image credit: Daniel O'Malley.*

The mapping in Figure 2b represents each DFN fracture as a graph node and each fracture intersection as an edge connecting the nodes. This can be viewed as a projection of the bipartite representation on the fracture nodes. One can easily extract measures of network connectivity (degree) and importance of individual fractures (centrality) under this mapping. Predictions of first arrivals are critical for applications such as nuclear nonproliferation, since trace gases like xenon migrate in high concentrations early on. In this case, one can use a mapping based solely on the topology shown in Figure 2b to determine the shortest paths between source and sink through a fracture network [4]. Figure 2c depicts a mapping that represents fractures as collections of edges and each intersection as a graph node; this is a projection of the bipartite representation onto the intersection nodes. Such mapping allows graph edges to inherit hydrological properties, such as permeability and in-plane geometry, and can simulate flow and transport in a computationally efficient manner. For DFN transport modeling, the breakthrough curve is a typical quantity of interest (QOI) that describes the mass breakthrough of the transported species as a function of time as it crosses a control plane at the outlet. The appropriate mapping depends on the QOI.

Figure 3 illustrates both the probability distribution function of arrival times at the outlet and the cumulative distribution function for a few of the relevant graph-based reduced-order models. A comparison to the high-fidelity DFN model is also shown to gauge the effectiveness of these models in emulating QOIs. The aggressive shortest path, which represents 10 percent of the entire network, can serve as a surrogate at very early times but deviates significantly beyond one percent of the total simulation time. This further underscores large-scale structure's dominance among the factors controlling transport within the network, especially at early times. Furthermore, one can use the structure to identify flow channels. Machine learning (ML) approaches have shown initial success in reducing the network size to 25 percent by classifying fractures as participating in the flow [7]. The pruned network predicted by ML techniques is a less aggressive reduction in fractures and tracks the DFN model well in regard to peak arrival.

In field applications such as hydraulic fracturing or environmental remediation, the entire breakthrough curve is typically of interest to measure production or remediation efficiency, respectively. The graph-based flow and transport models can run with up to five orders of magnitude of computational savings, even though the tradeoff is a systematic—but remediable—discrepancy from the mesh-based DFN solution. The corrected graph transport solution provides the closest match to the DFN solution at both early and peak arrival times, although discrepancies exist in the tail behavior. As mentioned previously, each of these methods is an acceptable surrogate for the full DFN model, and choosing the appropriate model depends heavily on the questions being asked.

It is important to note that one can always compare the aforementioned graph-based reduced models with the original higher fidelity DFN representation. The recent explosion of ML in geosciences and other fields raises the question of whether the developed emulators can predict QOIs outside of the training data, and what might be considered an extrapolation beyond the training regime. This criticism is particularly well-founded when learning from limited experimental data or already approximated models. However, 3D DFN models have been extensively used and validated against field observations, e.g., at the Äspö Hard Rock Laboratory in southern Sweden. This affords reasonable confidence that the "ground truth" furnished by the high-fidelity DFN models provides a solid foundation for additional insight into fractured systems.

Acknowledgments: The author thanks the following collaborators at Los Alamos National Laboratory for contributing to the research presented in this article: Aric Hagberg, Jeffrey Hyman, Satish Karra, Daniel O'Malley, David Osthus, Jamal Mohd-Yusof, Shriram Srinivasan, and Hari Viswanathan. The author also thanks Los Alamos National Laboratory's Laboratory Directed Research and Development Program for support throughout this project.

References

[1] Bogdanov, I., Mourzenko, V., Thovert, J.-F., & Adler, P. (2007). Effective permeability of fractured porous media with power-law distribution of fracture sizes. *Phys. Rev. E, 76*(3), 036309.

[2] Cacas, M.C. (1990). Modeling fracture flow with a stochastic discrete fracture network: Calibration and validation: 1. the flow model. *Wat. Resc. Res., 26*(3), 479-489.

[3] Djidjev, H., O'Malley, D., Viswanathan, H., Hyman, J.D., Karra, S., & Srinivasan, G. (2017). Learning on graphs for predictions of fracture propagation, flow and transport. In *2017 IEEE International Parallel and Distributed Processing Symposium Workshops* (pp. 1532-1539). Lake Buena Vista, FL.

[4] Hyman, J. D., Hagberg, A., Srinivasan, G., Mohd-Yusof, J., & Viswanathan, H. (2017). Predictions of first passage times in sparse discrete fracture networks using graph-based reductions. *Phys. Rev. E., 96* (1), 013304.

[5] Hyman, J.D., Karra, S., Makedonska, N., Gable, C.W., Painter, S.L., & Viswanathan, H.S. (2015). dfnWorks: A discrete fracture network framework for modeling subsurface flow and transport. *Comp. Geosci., 84*, 10-19.

[6] Painter, S., & Cvetkovic, V. (2005). Upscaling discrete fracture network simulations: An alternative to continuum transport models. *Wat. Resc. Res., 41*(2).

[7] Valera, M., Guo, Z., Kelly, P., Matz, S., Cantu, A., Percus, A.,…,Viswanathan, H. (2017). Machine learning for graph-based representations of three-dimensional discrete fracture networks. *Comput. Geosci.* In press.

Gowri Srinivasan is an applied mathematician and scientist in the Theoretical Division of Los Alamos National Laboratory. Her research interests include reduced-order models, uncertainty quantification, and machine learning.

Pursuing Science and Logic in an Age of Excitement and Turmoil

Exact Thinking in Demented Times: The Vienna Circle and the Epic Quest for the Foundations of Science. *By Karl Sigmund. Basic Books, New York, NY, December 2017. 480 pages. $32.00.*

BOOK REVIEW
By Ernest Davis

In 1920s Vienna, a group of philosophers, mathematicians, and physicists called "the Vienna Circle" embarked on a formidably ambitious project to investigate the foundations of science. Just as Alfred North Whitehead and Bertrand Russell had shown that mathematics can be built up from set theory in *Principia Mathematica*, Vienna Circle members aimed to demonstrate that one could logically build up scientific theory from basic observations. Their philosophy, known as "logical positivism" or "logical empiricism," idolized science, logic, and empiricism. It dismissed as meaningless any theorizing or speculation that was not empirically based, with a particular contempt for "metaphysics."

Exact Thinking in Demented Times, by mathematician Karl Sigmund, is a group biography of Vienna Circle participants and their predecessors, associates, and adversaries. It is both deeply researched and enormously entertaining, with vivid personal portraits, remarkable incidents and anecdotes, and a dramatic interpretation of an exciting and tragic historical period. Douglas Hofstadter helped polish Sigmund's English translation and contributed a fascinating preface describing his own intellectual encounters with the Vienna Circle and their idol/nemesis, Ludwig Wittgenstein.

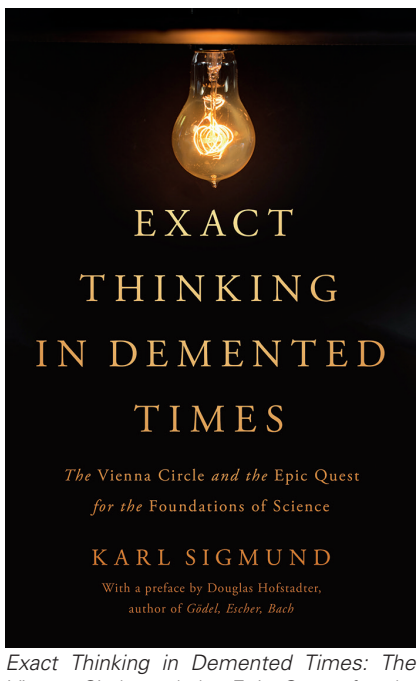

Exact Thinking in Demented Times: The Vienna Circle and the Epic Quest for the Foundations of Science. By Karl Sigmund. Courtesy of Basic Books.

At any given time, 10 to 20 people were involved with the Vienna Circle, which held regular meetings from 1924 to 1936. Central figures at the start included physicist Moritz Schlick, who served as chair; sociologist Otto Neurath; and mathematicians Otto Hahn and Philipp Frank. Rudolf Carnap joined the Circle in 1926 and became the leading exponent of logical positivism — his book, *The Logical Structure of the World*, became a bible for the movement. In 1929, the Circle announced themselves to the world by publishing a manifesto entitled "The Scientific Conception of the World: The Vienna Circle," which laid out their views and goals.

Looming over the Vienna Circle were four of the great intellectual giants of the first half of the twentieth century: Albert Einstein, Kurt Gödel, Karl Popper, and Wittgenstein. Sigmund wisely gives himself much latitude to digress outside the Circle proper, and includes substantial additional accounts of the lives and works of these four figures.

Einstein was not directly involved with the Circle, but he interacted with many participants over the years; his revolutionary discoveries about the nature of time, space, atoms, and light inevitably and profoundly impacted any study of physics. While Gödel was a member of the Circle for some years and attended meetings, he apparently stuck firmly to a Platonist conception of mathematics. His great discoveries in logic were met with excitement by the group's mathematicians, but hardly affected the overall philosophy. Popper was never invited to join the group or attend their meetings, and first came to prominence with a characteristically hard-hitting attack on logical positivism. Nonetheless, he was on collegial terms with the Circle and considered them as philosophical comrades-in-arms, fighting the same good fight.

Wittgenstein's impact on the Circle was enormous, and a substantial portion of *Exact Thinking in Demented Times* tracks the complex evolution of the Circle's relations with him. To many in the Circle, including Schlick and Hahn, Wittgenstein's *Tractatus Logico-Philosophicus* seemed a revelation; Neurath, however, found many of its oracular pronouncements to be mere "metaphysics." The Circle spent several semesters of meetings working through it line by line, and naturally wanted to bring Wittgenstein into their discussions. This led to complications since Wittgenstein never believed that anyone but himself actually understood him properly, and had radically changed his views since *Tractatus*.

While many of Wittgenstein's acolytes viewed him as a secular saint, Sigmund depicts him as somewhere between utterly horrible and completely impossible. The following characteristic anecdote serves as an example. In 1944, Wittgenstein held a philosophy chair at the University of Cambridge. While there, he was approached by Rose Rank (formerly a dedicated member of the Circle and always extremely poor), who was then working in Britain on a factory assembly line. She asked him if he could recommend her for a fellowship. Wittgenstein replied that there was nothing he could do for her, and added that he saw no disgrace in earning a living from manual labor.

Early chapters of the book describe the Circle's intellectual predecessors, starting with Ernst Mach and Ludwig Boltzmann and their debate about the reality of atoms. *Exact Thinking in Demented Times* is full of fascinating figures and incidents, some close to the narrative's main thread and others more distant. Sigmund painstakingly describes the struggles of Circle participants to obtain regular academic appointments, and offers charming vignettes of the characters' personal romances. He recounts the remarkable story of Gustav Klimt's scandalous Faculty Paintings (1900-07, destroyed in 1945) at the University of Vienna, which depicted the arts and sciences as "naked men and women drifting in forlorn trances through an uncanny void." He gives an account of three Viennese novelists—Robert Musil, Hermann Broch, and Leo Perutz—who were each influenced by their study of mathematics and wrote works with mathematician protagonists.

In the final chapters, the Vienna Circle and logical positivism come to an end in two distinct ways. The rising tide of Nazism, culminating in the Anschluss of 1938—the "demented times" of the title—

See **Science and Logic** *on page 12*

Seeing Through Rock: Mathematics of Inverse Wave Propagation

The following is a short introduction to an invited lecture to be presented at the upcoming 2018 SIAM Annual Meeting (AN18) in Portland, Ore., from July 9-13.

Most of what we know about Earth's interior comes from measuring physical fields on or near the surface and inferring subsurface structures from these measurements. Apart from inspecting outcrops, the optical band of the electromagnetic spectrum is useless—you cannot actually see through rock, despite the title of this talk—but there are many other choices for probing fields. Seismic fields (roughly speaking, elastic waves in Earth) have a few key advantages as subsurface probes. They cover relatively long distances—measurable waves of the lowest frequency excite all of Earth—and combine this range with resolution, propagating a few hundred wavelengths before disappearing into Earth's background seismic noise. To some extent, they exhibit space-time locality and directionality (that is, locality in phase space) and definite velocity, which makes them waves.

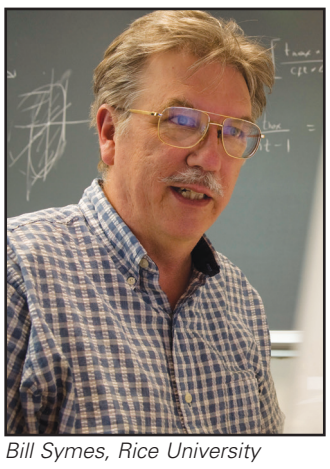

Bill Symes, Rice University

Whether elastodynamics or a more complex model of wave physics is adequate to predict measured seismic data, and what information about Earth's structure might result from doing so, is an *inverse problem*. One can only answer these questions affirmatively by exhibiting a distribution of elastic (or other) parameters for which the solution of the associated partial differential equations (PDEs) actually predicts the seismic measurements to satisfactory accuracy. Joe Keller proposed this "find the question given the answer" definition of "inverse problem" in his 1976 *American Mathematical Monthly* article. However, Keller's examples do not clarify how instances of such problems involving nonlinear relations between billions of data points and billions of parameters, mediated by complex systems of PDEs, might be wrestled into computational practicality.

As I will explain in my 2018 SIAM Annual Meeting talk, answers to this question began taking shape in the 1980s, as Albert Tarantola and many others elaborated the "outer loop" point of view mentioned in Bruce Hendrickson's November 2017 *SIAM News* piece.[1] The outer loop perspective posits that posing data fit as an optimization problem over choices of elastic parameters invites selection by iteration over solutions of the elastodynamic (or similar) wave equations. Tarantola's classic text, *Inverse Problem Theory and Methods for Model Parameter Estimation*, republished by SIAM in 2005, shows how tools from control theory make the computation of artifacts needed for local (descent-based) optimization more feasible.

The first practicable algorithms required another 20 years of computational and mathematical advances and inspired a very active research area known as full waveform inversion, with branches in industrial and academic seismology. The computational needs were obvious consequences of the world's three-dimensionality. The mathematical issues were more surprising — already evident in the 1980s, they stem from the simple observation that perturbation of wave speed effectively differentiates the waveform. Combined with the nonlinearity of the elastic wavefield as a function of the coefficients in the (linear) wave equation, this hypersensitivity with respect to wave velocities—which are functions of the coefficients—can cause local optimization to stagnate far from a useful elastic model of Earth, *no matter how many Pflops are expended*. I will review some of the numerous ideas advanced to overcome this obstacle, and show why they might work via simple examples.

— Bill Symes, Rice University

[1] https://sinews.siam.org/Details-Page/the-future-of-scientific-computation

How Paradoxes Shape Mathematics and Give Us Self-verifying Computer Programs

The following is a short introduction to the I. E. Block Community Lecture, to be presented at the upcoming 2018 SIAM Annual Meeting (AN18) in Portland, Ore., from July 9-13.

A paradox is a seeming contradiction, and the liar's paradox is among the best-known. "This statement is false" — the statement is true exactly when it is false. A paradox is often self-referential, making a statement about itself.

Thomas Hales, University of Pittsburgh

Paradoxes can be so amusing that we might be tempted to believe they are nothing more than a game. However, they became a serious business more than a century ago, through a paradox similar to the barber paradox: a barber named Bertie shaves exactly those who do not shave themselves. Does Bertie shave himself? If he does, then he doesn't; if he doesn't, then he does. This paradox triggered a deep crisis in the foundations of mathematics. Bertrand Russell and Alfred North Whitehead spent years rebuilding mathematics from the ground up; after a decade of effort, they finally produced a paradox-free proof that $1+1=2$.

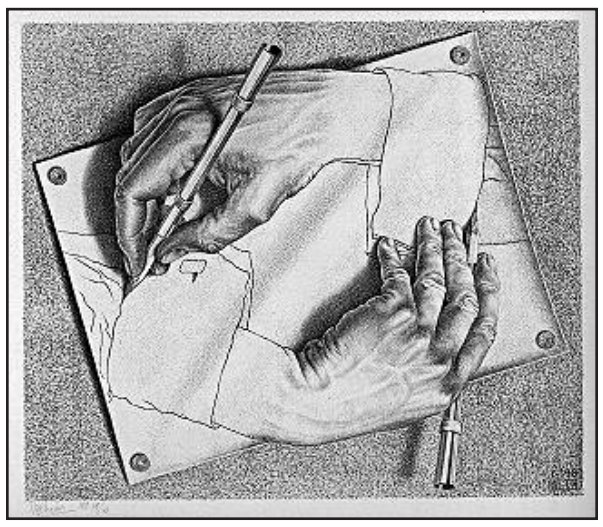

Drawing Hands, a lithograph by M.C. Escher, is an example of a paradox. Thomas Hales will explore the relation between paradoxes and mathematics at the 2018 SIAM Annual Meeting. Public domain image.

Other clever paradoxes expose the disturbing limits of computation and mathematics. Researchers such as Alan Turing and Kurt Gödel discovered these mathematical bombshells.

To avoid paradox, most current mathematicians work within restrictive systems that banish self-reference. Rebelling against these restrictions, some have avoided paradox in mathematics by placing the entire mathematical universe inside a "nutshell" that sits within a still larger universe, and then continuing with an ever increasing expanse of universes, each a nutshell in the next.

Today, we design computer programs that verify the lack of bugs in other computer programs. Can computer programs be fed into themselves to verify their own correctness? Or does paradox stop us in our tracks? One approach to self-verification is the design of computer languages that avoid paradox through nutshell universes. Other ideas for program self-verification come from the century-old proof that $1+1=2$.

Researchers in artificial intelligence are carrying paradoxes into new realms. Can beneficial artificial intelligence turn malevolent when it starts modifying its own computer code?

In my talk at the 2018 SIAM Annual Meeting, I will present some of my favorite paradoxes and explain how they play into self-verifying computer programs.

— Thomas Hales, University of Pittsburgh

Look for feature articles by other AN18 invited speakers introducing the topics of their talks in future issues.

Professional Opportunities and Announcements

Send copy for classified advertisements and announcements to ***marketing@siam.org****. For rates, deadlines, and ad specifications visit* ***www.siam.org/advertising****.*

*Students (and others) in search of information about careers in the mathematical sciences can click on "Careers and Jobs" at the SIAM website (***www.siam.org****) or proceed directly to* ***www.siam.org/careers****.*

Georgetown University

Department of Mathematics and Statistics

The Department of Mathematics and Statistics at Georgetown University announces a search for the position of Director of Graduate Studies (DGS) in Mathematics and Statistics. The DGS is a full-time administrative position with responsibility for managing the M.S. degree program in mathematics and statistics at Georgetown. The DGS is expected to teach at least one graduate course per year and oversee all operations involving the graduate program. Responsibilities include—but are not limited to—recruiting, admissions, curriculum development, and scheduling of courses. There is a full-time assistant who handles many of the administrative and clerical duties and reports to the DGS.

To apply, go to the following link, where a full description of the position is posted: **https://georgetown.wd1.myworkdayjobs.com/en-US/Georgetown_Faculty/job/Main-Campus/Director-of-Graduate-Studies--Mathematics---Statistics---Georgetown--College_JR03031**. Establish an account and follow the prompts.

California Institute of Technology

Department of Computing and Mathematical Sciences

The Department of Computing and Mathematical Sciences (CMS) at the California Institute of Technology invites applications for the position of lecturer in computing and mathematical sciences. This is a (non-tenure-track) career teaching position with full-time teaching responsibilities. The start date for the position is ideally September 1, 2018, and the initial term of appointment can be up to three years.

The lecturer will teach introductory computer science courses—including data structures, algorithms, and software engineering—and work closely with the CMS faculty on instructional matters. The ability to teach intermediate-level undergraduate courses in areas such as software engineering, computing systems, and/or compilers is desired. The lecturer may also assist in other aspects of the undergraduate program, including curriculum development, academic advising, and research project monitoring. The lecturer must have a track record of excellence in teaching computer science to undergraduates. In addition, the lecturer will have opportunities to participate in research projects in the department. An advanced degree in computer science or a related field is desired but not required.

Applications will be accepted on an ongoing basis until the position is filled.

Please view the application instructions and apply online at **https://applications.caltech.edu/job/cmslect.**

The California Institute of Technology is an equal opportunity/affirmative action employer. Women, minorities, veterans, and disabled persons are encouraged to apply.

The High-Performance Conjugate Gradients Benchmark

By Jack Dongarra, Michael A. Heroux, and Piotr Luszczek

The High-Performance LINPACK (HPL) Benchmark has been a measure of supercomputing performance for more than four decades, and the basis for the biannual TOP500 list of the world's fastest supercomputers for over 25 years. The benchmark is one of the most widely recognized and discussed metrics for ranking high-performance computing (HPC) systems. When HPL gained prominence as a performance metric in the early 1990s, there was a strong correlation between its predictions of system rankings and the ranking realized by full-scale applications. In these early years, computer system vendors pursued designs that would increase HPL performance, thus improving overall application function.

Similarity of computations was partially responsible for this correlation. For example, frontal matrix solvers were commonly used in engineering applications and often consumed a large fraction of compute time. The computational and data access patterns of these solvers are similar to HPL. Moreover, memory system and floating-point computation performance was much more balanced. For example, the Cray YMP and C90 (1990s systems) could perform two reads and a write per clock cycle, enabling near-peak performance for writing the weighted sum of two vectors as another vector, the so-called AXPY operation. On today's modern processors, that operation executes at about one to two percent of peak speed. HPL has a computational complexity of $O(n^3)$ and a data access complexity of $O(n^2)$, so simply running larger problems meant that its performance was minimally impacted by this trend. Many real applications have moved to more efficient algorithms with computational complexity closer to $O(n)$ or $O(n \log_2 n)$ and similar data access complexity, and have realized a much smaller performance gain from computer system improvements. Even so, the net performance improvement in time to solution of new algorithms on new platforms has far exceeded HPL improvements. In contrast, time to solution for HPL is now measured in days, and is a serious concern for benchmarkers on leadership platforms.

We expressed the following in [1]:

> HPL remains tremendously valuable as a measure of historical trends and as a stress test, especially for the leadership class systems that are pushing the boundaries of current technology. Furthermore, HPL provides the HPC community with a valuable outreach tool, understandable to the outside world. Anyone with an appreciation for computing is impressed by the tremendous increases in performance that HPC systems have attained over the past few decades in terms of HPL. At the same time, HPL rankings of computer systems are no longer so strongly correlated to real application performance, especially for the broad set of HPC applications governed by differential equations.

These tend to strictly demand high bandwidth and low latency as they possess the aforementioned lower computational complexity. In fact, we have reached a point where designing a supercomputer for good HPL performance can lead to design choices that are either ill-suited for the real application mix or add unnecessary components or complexity to the system. Left unchecked, we expect the gap between HPL predictions and real application performance to increase in the future.

Many aspects of the physical world are modeled with partial differential equations, which help predictive capability, thus aiding scientific discovery and engineering optimization. The High-Performance Conjugate Gradients (HPCG) Benchmark is a complement to the HPL Benchmark and now part of the TOP500 effort. It is designed to exercise computational and data access patterns that more closely match a different yet broad set of important applications, and to encourage computer system designers to invest in capabilities that will impact the collective performance of these applications.

Rank	Site	Computer	Cores	HPL Rmax (Pflop/s)	TOP500 Rank	HPCG (Pflop/s)	Fraction of Peak
1	RIKEN Advanced Institute for Computational Science *(Japan)*	K computer – SPARC64 VIIIfx 2.0GHz, Tofu interconnect, Fujitsu	705,024	10.51	10	0.603	5.30%
2	National Supercomputer Center in Guangzhou *(China)*	Tianhe-2 (MilkyWay-2) – TH-IVB-FEP Cluster, Intel Xeon 12C 2.2GHz, TH Express 2, Intel Xeon Phi 31S1P 57-core, NUDT	3,120,000	33.86	2	0.58	1.10%
3	Department of Energy/National Nuclear Security Administration/Los Alamos National Laboratory/Sandia National Laboratories *(USA)*	Trinity – Cray XC40, Intel Xeon E5-2698 v3 300160C 2.3GHz, Aries, Cray	979,072	14.13	7	0.546	1.80%
4	Swiss National Supercomputing Centre (CSCS) *(Switzerland)*	Piz Daint – Cray XC50, Intel Xeon E5-2690v3 12C 2.6GHz, Cray Aries, NVIDIA Tesla P100 16GB, Cray	361,760	19.59	3	0.486	1.90%
5	National Supercomputing Center in Wuxi *(China)*	Sunway TaihuLight – Sunway MPP, SW26010 260C 1.45GHz, Sunway, NRCPC	10,649,600	93.01	1	0.481	0.40%
6	Joint Center for Advanced High Performance Computing *(Japan)*	Oakforest-PACS – PRIMERGY CX600 M1, Intel Xeon Phi Processor 7250 68C 1.4GHz, Intel Omni-Path Architecture, Fujitsu	557,056	13.55	9	0.385	1.50%
7	Department of Energy/Office of Science/Lawrence Berkeley National Laboratory/National Energy Research Scientific Computing Center *(USA)*	Cori – XC40, Intel Xeon Phi 7250 68C 1.4GHz, Cray Aries, Cray	632,400	13.83	8	0.355	1.30%
8	Department of Energy/National Nuclear Security Administration/Lawrence Livermore National Laboratory *(USA)*	Sequoia – IBM BlueGene/Q, PowerPC A2 1.6 GHz 16-core, 5D Torus, IBM	1,572,864	17.17	6	0.33	1.60%
9	Department of Energy/Office of Science/Oak Ridge National Laboratory *(USA)*	Titan – Cray XK7, Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x, Cray	560,640	17.59	5	0.322	1.20%
10	Global Scientific Information and Computing Center, Tokyo Institute of Technology *(Japan)*	TSUBAME3.0 – SGI ICE XA (HPE SGI 8600), IP139-SXM2, Intel Xeon E5-2680 v4 15120C 2.9GHz, Intel Omni-Path Architecture, NVIDIA TESLA P100 SXM2 with NVLink, HPE	136,080	12.12	13	0.189	1.60%

High-Performance Conjugate Gradients (HPCG) Benchmark: Top 10 systems as of November 2017. *The chart lists rank according to HPCG, computer location, computer name and specifications, core (processor) count, HPL performance, TOP500 rank, HPCG performance, and the fraction of the theoretical peak performance obtained for the HPCG Benchmark.*

We articulated the subsequent ideas in [1]:

> The setup phase [of HPCG] constructs a logically global, physically distributed sparse linear system using a 27-point stencil at each grid point in the 3D domain, such that the equation at point (i, j, k) depends on the values of its location and 26 surrounding neighbors. The matrix is constructed to be weakly diagonally dominant for interior points of the global domain, and strongly diagonally dominant for boundary points, reflecting a synthetic conservation principle for the interior points and the impact of zero Dirichlet boundary values on the boundary equations. The resulting sparse linear system has the following properties:
>
> • A sparse matrix with 27 nonzero entries per row for interior equations and seven to 18 nonzero terms for boundary equations
>
> • A symmetric, positive definite, nonsingular linear operator
>
> • The boundary condition is reflected by subtracting one from the diagonal
>
> • A generated known exact solution vector with all values equal to one
>
> • A matching right-hand-side vector
>
> • An initial guess of all zeros.
>
> The central purpose of defining this sparse linear system is to provide a rich vehicle for executing a collection of important computational kernels. However, the benchmark is not about computing a high fidelity solution to this problem. In fact, iteration counts are fixed in the benchmark code and we do not expect convergence to the solution, regardless of problem size. We do use the spectral properties of both the problem and the preconditioned conjugate-gradient algorithm as part of software verification.

The HPCG reference code is complete, standalone, and derived from mini-applications developed in the Mantevo project.[1] It measures the performance of basic operations in a unified code:

- Sparse matrix-vector multiplication
- Vector updates
- Global dot products
- Local symmetric Gauss-Seidel smoother
- Sparse triangular solve (as part of the Gauss-Seidel smoother).

The code is also driven by a multigrid preconditioned conjugate gradient algorithm that exercises the key kernels on a nested set of coarse grids. The reference implementation is written in C++ with MPI and OpenMP support.

Computer system vendors have invested significant resources to optimize HPCG, including the adaptation of their math kernel libraries to provide optimized functionality that can benefit the broader communities using these libraries.

HPL follows the peak performance of the machine relatively closely — a fact that is well known to benchmarking practitioners and most HPC experts. The performance levels of HPCG are far below those seen by HPL. This should not be surprising to those in the high-end and supercomputing fields and is attributable to many factors, including the commonly-cited "memory wall."

HPCG has already been run on many large-scale supercomputing installations in China, Europe, Japan, and the U.S. (and off-planet in an orbiting satellite). The above chart shows the top 10 systems on the current HPCG Benchmark list, as of November 2017. A full list and more details are available online.[2] HPCG results have been integrated into the TOP500 list.[3]

Acknowledgments: The authors thank the Department of Energy National Nuclear Security Administration for funding this work.

[1] https://mantevo.github.io

[2] http://www.hpcg-benchmark.org/

[3] https://www.top500.org/

References

[1] Dongarra, J., Heroux, M.A, & Luszczek, P. (2016). High-performance conjugate-gradient benchmark: A new metric for ranking high-performance computing system. *Int. J. High Perf. Comp. App., 30*(1), 3-10. http://journals.sagepub.com/doi/abs/10.1177/1094342015593158

Jack Dongarra is a university distinguished professor in the Department of Electrical Engineering and Computer Science at the University of Tennessee, and a distinguished research staff member in the Computer Science and Mathematics Division at Oak Ridge National Laboratory. Michael A. Heroux is a senior scientist at Sandia National Laboratories, director of SW Technologies for the U.S. Department of Energy's Exascale Computing Project, and scientist-in-residence at St. John's University, Minn. His research interests include all aspects of scalable scientific and engineering software for new and emerging parallel computing architectures. Piotr Luszczek is a research assistant professor at the University of Tennessee and a research director at the university's Innovative Computing Laboratory. His research interests are in performance engineering and numerical linear algebra. He teaches courses in these subjects at the graduate level and internationally at invited lectures.

Science and Logic

Continued from page 9

drove the Circle far from Vienna. Some Circle members were Jewish, some were active in left-wing politics, and all were considered undesirable in the Third Reich. Schlick was murdered by a psychopathic stalker in 1936; though he was neither Jewish nor politically active, Nazi writers defended the murder as a justifiable reaction to his perverse philosophy. The rest of the Circle fled, sooner or later, into exile, ending up in either the U.S. or the U.K.

Ultimately, the philosophers were more deadly than the politics. A series of trenchant, unanswerable critiques of logical positivism by Popper, Willard Van Orman Quine, Thomas Kuhn, and others left the Circle entirely demolished — to a degree unusual in philosophy, where victories and defeats are typically partial and provisional. The word "positivist" became a pejorative, akin to narrow-mindedly ignoring everything except superficial, easily-quantifiable measurements. Positivism also became associated with the behaviorist theory of psychology, which flourished and perished over much the same time period. Strikingly, the word "positivism" is almost never used in Sigmund's book.

Nonetheless, the problem of characterizing the relation of observational evidence to scientific theory is real, important, and unsolved. No one questions that the Vienna Circle very much overestimated the scope and power of the type of analysis it was pursuing. However, that does not negate the possibility that this kind of analysis is worth pursuing when analyzing the foundations of physical science.

Ernest Davis is a professor of computer science at the Courant Institute of Mathematical Sciences, New York University.