Newsjournal of the Society for Industrial and Applied Mathematics

siam news

Volume 46/ Number 6
July/August 2013

CSE 2013

How Will the Fast Multipole Method Fare in the Exascale Era?

By Lorena A. Barba and Rio Yokota

"Complexity trumps hardware," according to the conventional wisdom. It is natural that as computers get faster, people want to solve larger problems, eventually bringing fast algorithms to the forefront. This thought was likely in SIAM past-president Nick Trefethen's mind when, in his 1998 "Predictions for Scientific Computing, Fifty Years from Now," he declared that fast multipole methods would be ubiquitous. With the irreversible move toward parallel computing of the last decade, however, much of the conventional wisdom began to expire. How, we wonder, will FMM fare in the exascale era?

The non-stop trend in computer architectures today sees the breadth of parallelism increasing rapidly, while memory bandwidth grows much more slowly than processing speed. Overall byte-to-flop ratios (machine balance) are declining—relentlessly—at every level of memory. For now, algorithm designers must respond by striving to reduce data movement and synchronization points, even at the cost of more arithmetic operations. But in some cases, the only way forward may be either a change in the numerical engine or a complete mathematical reformulation of the problem. Such is the pressure from evolving computer architectures that algorithms with large communication and synchronization requirements could become obsolete.

Machine balance is converging: Almost all current architectures saturate at between 0.2 and 0.285 byte/flop (the exception, Fujitsu, at 0.36 byte/flop, is very expensive hardware). Unfortunately, a conflict exists between algorithmic complexity and arithmetic intensity (measured in flop/byte, the inverse unit of machine balance). Thus, algorithms with $\mathcal{O}(N^2)$ complexity, such as direct N-body or molecular dynamics, have high flop/byte ratios and high levels of parallelism. Conversely, the ideal $\mathcal{O}(N)$ algorithms require few operations per data point and tend to have more dependencies, e.g., the stencil operations of finite-difference or finite-element methods. FMM distinguishes itself by offering $\mathcal{O}(N)$ complexity as well as high flop/byte ratios. In Figure 1, which shows a roofline model of several current architectures along with the arithmetic intensity of representative algorithms, FMM computations tend to lie well under the flat part of the roof. This may explain why the computational science community is more and more interested in FMM and related methods.

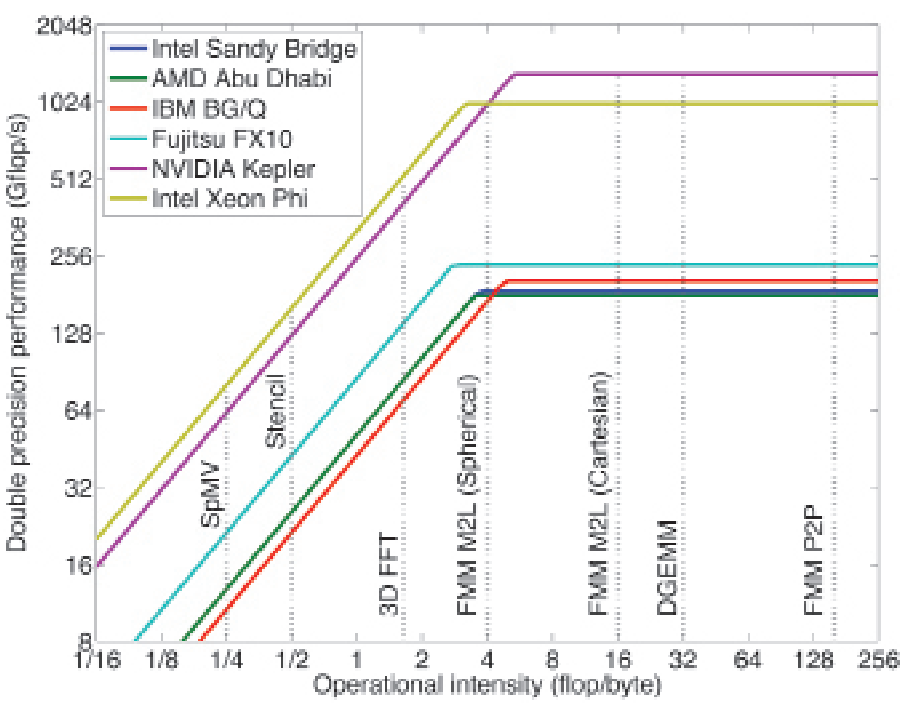

Figure 1. *Roofline model of current microarchitectures and operational intensity of representative algorithms. SpMV is a sparse matrix–vector multiplication, Stencil is a multigrid method with a seven-point finite-difference stencil, 3D FFT is a three-dimensional fast Fourier transform, and DGEMM is a dense matrix–matrix multiplication. FMM M2L (spherical) and FMM M2L (Cartesian) are FMM kernels corresponding to interactions between distant clusters that use, respectively, spherical and Cartesian expansions; FMM P2P is an FMM kernel corresponding to interactions between nearby particles.*

At the 2013 SIAM Conference on Computational Science and Engineering, topics related to FMM were the focus of at least four minisymposia, a set of 32 talks (MS 11, 31, 72, and 132 in the program). Participants included Leslie Greengard and Vladimir Rokhlin, the creators of FMM, as well as many of their past students and descendants. The sessions covered the latest mathematical developments—e.g., high-order quadrature and well-conditioned formulations for non-smooth geometries—as well as novel applications, including magnetic resonance imaging and algebraic preconditioning. Calculations of acoustic [2], electromagnetic [3], and electrostatic [10] fields continue to be the major application areas for multipole methods. Ongoing efforts are also under way to construct more general fast direct solvers, from both geometric [5] and algebraic [4] points of view. We organized a session that emphasized not only applications of FMM, but also the computer science aspects of autotuning, optimization for GPU hardware, and task-based approaches to parallelism.

In a recent paper we discussed the features of FMM that we believe make it a favorable algorithm for future architectures [8]. In brief, its use of a tree data structure extracts spatial data locality from nonuniform and multiscale resolutions and results in benign synchronization requirements. What's more, the various FMM kernels are either purely local or present a hierarchical communication pattern. These features have allowed several groups to achieve highly efficient massively parallel computations, using hundreds of thousands of CPU cores

See **Fast Multipole Method** *on page 3*

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

CSE 2013

Cardiac Modeling: The Road from Equations to the Clinic

By Natalia A. Trayanova

Advances in computer modeling, as they transform many traditional areas of physics and engineering, are also transforming our understanding of the function of the heart in health and disease. Modern cardiac researchers are increasingly aware that appropriate models and simulation can help interpret an array of experimental data and dissect important physiological mechanisms and interrelationships. Decades of developments in cardiac simulation have rendered the heart the most highly integrated example of a "virtual organ." These developments are firmly anchored in the long history of cardiac cell modeling, which dates back more than 50 years, to publication of the first model of the cardiac cell action potential, and are rooted in iterative interactions between modeling and experimentation.

Cardiac cell (myocyte) action potential models often take the form of coupled systems of nonlinear ordinary differential equations describing current flow through ion channels, pumps, and exchangers, as well as subcellular calcium cycling; model equations are solved to observe how states (concentrations of molecules) evolve in time as they interact with one another and respond to inputs. Over the years, models of myocyte action potential have improved rapidly, with the incorporation of validated biophysical relations and descriptions of many subcellular processes and pathways regulating the electrical function of the cell. The result has been dramatic enhancement of the physiological relevance of heart cell models.

Over the last two decades, cardiac modeling has also progressed to the level of the tissue and the whole heart, where the propagation of a wave of action potential

See **Cardiac Modeling** *on page 8*

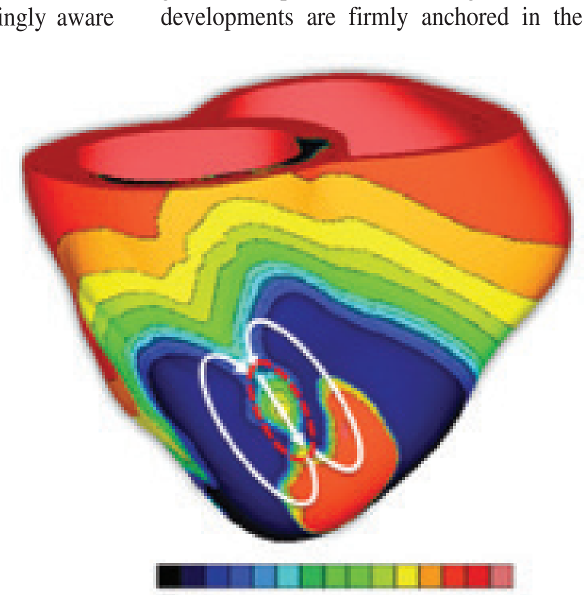

Simulated arrhythmia in a model of a patient heart with myocardial infarction. White arrows indicate direction of propagation of the re-entrant wave sustaining the arrhythmia. The dashed red circle encloses the predicted optimal target of catheter ablation. Reprinted from H. Ashikaga, H. Arevalo, F. Vadakkumpadan, R.C. Blake 3rd, J.D. Bayer, S. Nazarian, M. Muz Zviman, H. Tandri, R.D. Berger, H. Calkins, D.A. Herzka, N.A. Trayanova, and H.R. Halperin, "Feasibility of Image-Based Simulation to Estimate Ablation Target in Human Ventricular Arrhythmia," Heart Rhythm, *April 19, 2013, Epub ahead of print; doi:10.1016/j.hrthm.2013.04.015; pii:S1547-5271(13)00437-2; http://www.ncbi.nlm.nih.gov/pubmed/23608593.*

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

siam news
Volume 46/ Number 6/ July/August 2013

What Does It Take to Succeed as a Data Scientist in the Big Data Era?

"Big Data" was a pervasive theme at the 2013 SIAM International Conference on Data Mining, held in Austin, Texas, May 2–4.* Along with discussions of algorithms, machine learning techniques, and statistical methods and their applications, the program included a panel discussion, held the evening of May 3, on the implications of Big Data for the practice of data science.

The panel session, "On Being a Successful Data Scientist in the Big Data Era," was moderated by Srinivasan Parthasarathy of Ohio State University; the panelists provided perspectives from academia (Chris Jermaine, Rice; Raymond Mooney, University of Texas–Austin; and Peter Szolovits, MIT); government (Chandrika Kamath, Lawrence Livermore National Laboratory); and industry (L.V. Subramaniam, IBM Research–India). The panelists addressed three general questions: What has been the impact of the widespread publicity surrounding Big Data and its potential applications? What research challenges does Big Data present? How can we prepare students to be successful data scientists?

Just one example of the recent flood of publicity on Big Data was a 16-page special section devoted to the topic in *The New York Times* on June 20.[†] The lead article summarizes the promise of Big Data for *Times* readers as follows:

"The story is the same in one field after another, in science, politics, crime prevention, public health, sports and industries as varied as energy and advertising. All are being transformed by data-driven discovery and decision-making. . . .

"Big Data is the shorthand label for the phenomenon, which embraces technology, decision-making and public policy. Supplying the technology is a fast-growing market, increasing at more than 30 percent a year and likely to reach $24 billion by 2016, according to a forecast by IDC, a research firm. All the major technology companies, and a host of start-ups, are aggressively pursuing the business."

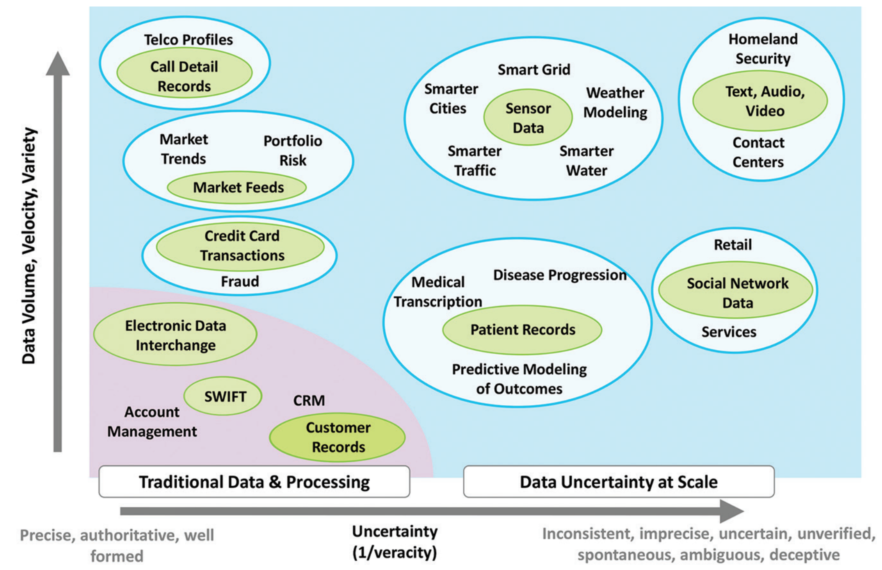

The need for managing "4V data" is widespread. Image courtesy of L.V. Subramaniam, IBM Research–India.

The panelists noted positive aspects of the extensive publicity: It increases the visibility of the field, expands the set of applications, and draws talent into the field. At the same time, they expressed reservations, citing increasing concerns about privacy and a possible backlash if exaggerated predictions are not met; they also mentioned ways in which the challenges posed by Big Data have been distorted—in particular, the emphasis placed on technology at the expense of data analysis.

Many of the research challenges discussed stem from what one panelist termed "4V data"—data with volume, velocity, variety, and veracity. Velocity is a consideration because some or all of the data may not be static, as in the case of streaming data. Variety is a reference to the different formats or modes in which data can appear. Veracity is 1/uncertainty, which means that analysis implies reasoning under uncertainty. Another challenge arises from the variety of applications. As one panelist put it, a data scientist has to have a 360-degree view, as customers do not know what to do with big data. Finally, various instruments (surveys, experiments, simulations) need to be designed so that the results of the analysis can be compared to real-world data.

The panel, having identified these challenges, turned to the question of designing training programs for data scientists. The panelists agreed that some elements of training are straightforward, including courses in algorithms, data mining, machine learning, programming, databases, and statistics. Other elements are somewhat more difficult to achieve. For example, students should have experience working with real data, which implies, in turn, that they need to understand the domain from which the data is derived. Finally, data scientists need to master tools like visualization, and they must have the ability to present their results clearly in presentations and in writing.—*William G. Kolata, SIAM Technical Director.*

*Links to the program, abstracts, and online proceedings can be found at http://www.siam.org/meetings/sdm13/.

[†]http://bits.blogs.nytimes.com/2013/06/19/sizing-up-big-data-broadening-beyond-the-internet/.

Predicting the Financial Crisis—View from the Real World

LETTERS TO THE EDITOR

To the Editor:

I liked the review of Nate Silver's book by James Case ("Foxes, Hedgehogs, and the Art of Prediction," *SIAM News*, April 2013), although I feel that Silver's comments on the financial crisis, as related by Case, were too academic. I'd like to provide some real-world perspective.

First, the paragraph of concern in the book review:

"Many of humankind's misfortunes, Silver writes, result from failure to predict clearly foreseeable events. The recent financial crisis is a case in point. He identifies four separate failures: that of homeowners and investors to foresee that housing prices could not continue to rise indefinitely, that of the ratings agencies to foresee that a fall in housing prices would cause a crisis in U.S. financial markets, that of economists to foresee that a financial crisis in the U.S. was tantamount to a global financial crisis, and that of leadership to foresee that a financial crisis would produce an unusually long and deep recession. None of the four qualifies as a "black swan"—all were readily apparent to anyone who cared to look "under the hood" of the models then in use."

Second, my perspective:

1. Whether the average homeowner thought that housing prices would rise indefinitely is surely debatable. But it is not debatable that many people who could not afford to buy a home were able to do so because the mortgage brokers did not perform the proper due diligence. The mortgage brokers knew that they were going to sell off the mortgage to a financial institution and collect a commission. The risk would be passed along to someone else.

2. Some of the financial institutions passed these questionable mortgages and their derivatives to investors and collected the commissions. The risk was passed on to the investors.

3. There were people at the ratings agencies who knew of the flaws in the models. But when they voiced their concerns to "senior management," they were told: "Shut up, we're making lots of money."

4. The U.S. Congress dismantled much of the previous regulatory structure (think Glass–Steagall Act).

5. The regulators failed to prevent the fraud in the system. However, to some extent their power to do so was inadequate.

6. Whether economists and leadership could have foreseen a global financial crisis is debatable, but vision is always more clear in hindsight.—*Steve Hellinger*

The writer spent nine years at JP Morgan Chase, most recently in Risk Technology.

siam news

ISSN 1557–9573. Copyright 2013, 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 ten times in 2013: 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: $66.00 U.S., Canada, and Mexico, $76.00 elsewhere. SIAM members and subscribers should allow 8 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 6 months or more.

Printed in the USA.

siam, is a registered trademark.

New Chinese Center Bridges Mathematics and Science/Technology

The National Center for Mathematics and Interdisciplinary Sciences was founded in November 2010 to create a bridge between mathematics and the natural sciences, engineering, technology, and the social sciences. Part of the "Innovation 2020" program launched by the Chinese Academy of Sciences, NCMIS has six interdisciplinary divisions, each linking mathematics with another field: information technology; economics and finance; advanced manufacturing, materials, and the environment; biology and medicine; physics; and engineering. The center's research program is affiliated with the Academy of Mathematics and Systems Science of the CAS.

NCMIS organizes interdisciplinary forums and has graduate and postgraduate programs in industrial and applied mathematics. It currently supports about 160 researchers from the Chinese Academy of Sciences and other institutions, for work on 21 research projects. The projects address such topics as computational fluid dynamics and applications, economic forecasting and policy simulation, mathematical methods for complex disease diagnosis and drug design, mathematical methods in digital design and manufacturing, and quantum information, computing and control.

Plans are well under way for ICIAM 2015, which will be held in Beijing, August 10–14. NCMIS is the main local organizer of the congress, and NCMIS director Lei Guo is the congress director. The China Society for Industrial and Applied Mathematics, together with the Chinese Mathematical Society, the Chinese Computational Mathematical Society, the Operations Research Society of China, the Chinese Association for Applied Statistics, and the Systems Engineering Society of China, will host ICIAM 2015. At NCMIS, we look forward to acquainting the world applied mathematics community with some of the exciting interdisciplinary work under way in which mathematics is a central player.—*Li Wang, National Center for Mathematics and Interdisciplinary Sciences, Chinese Academy of Sciences.*

Also looking ahead to the summer of 2015 is SIAM, which will be organizing minisymposia for the Beijing congress. As usual, ICIAM will serve as the 2015 Annual Meeting.

On May 13, SIAM president Irene Fonseca, a professor of the mathematical sciences at Carnegie Mellon University, gave a public lecture, "Variational Methods in Materials and Image Processing," in the NCMIS Distinguished Lecture Series. She is shown here with Xiao-Shan Gao, vice president of the Academy of Mathematics and Systems Science, who awarded her a certificate of recognition.

Visiting the Great Wall during an ICIAM 2015 planning meeting were Irene Fonseca, and ICIAM treasurer José Alberto Cuminato and president Barbara Keyfitz.

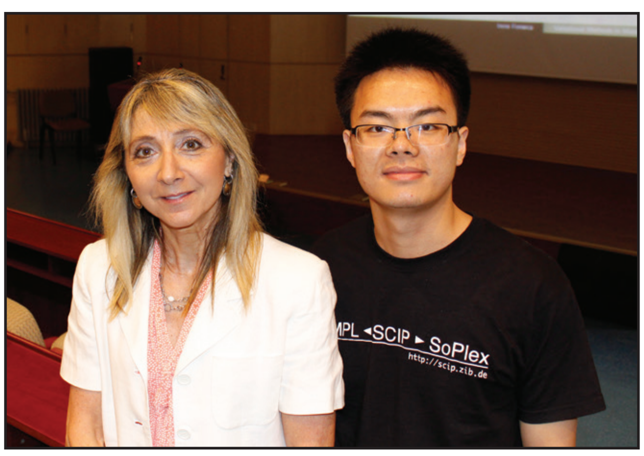

Zhenli Sheng, president of the Chinese Academy of Sciences Chapter of SIAM, helped make Fonseca's talk a successful activity of the student chapter.

In Vancouver, at the closing of ICIAM 2011, NCMIS director Lei Guo accepted a "talking stick" from congress director Arvind Gupta—setting the stage, according to Canadian First Nations tradition, for plenty of good talks at ICIAM 2015.

Fast Multipole Method

continued from page 1

[6,7] or thousands of GPUs [9]. Our own results suggest that applications that use an FFT-based numerical engine are at a clear disadvantage compared with FMM in terms of parallel scalability beyond several thousand processes (see Figure 2).

Some of our colleagues have conveyed a more gloomy outlook, even for FMM [1]. Projecting current architectural trends and offering a new analysis of communication that is calibrated to the state-of-the-art implementation of the kernel-independent version of FMM (known as `kifmm`), they predicted that FMM will become memory-bound by 2020. Re-inspection of Figure 1 shows that with our FMM code (called `exaFMM`), we obtain an arithmetic intensity of 4 flop/byte on the box–box interactions (far-field) using spherical expansions. This lies on the edge of the compute-bound "flat roof." If machine balance worsens as peak performance increases and the inclined part of the roofline extends to the right, this kernel can indeed become bandwidth-bound. Yet this evolution will shift the crossover point at which the more flop-intensive kernel using Cartesian expansions becomes competitive; we will then have the option of switching to this kernel for far-field calculations. Nevertheless, we are more optimistic than our colleagues about the hardware projections and think that architecture researchers will make a breakthrough before DGEMM is at risk of becoming memory-bound!

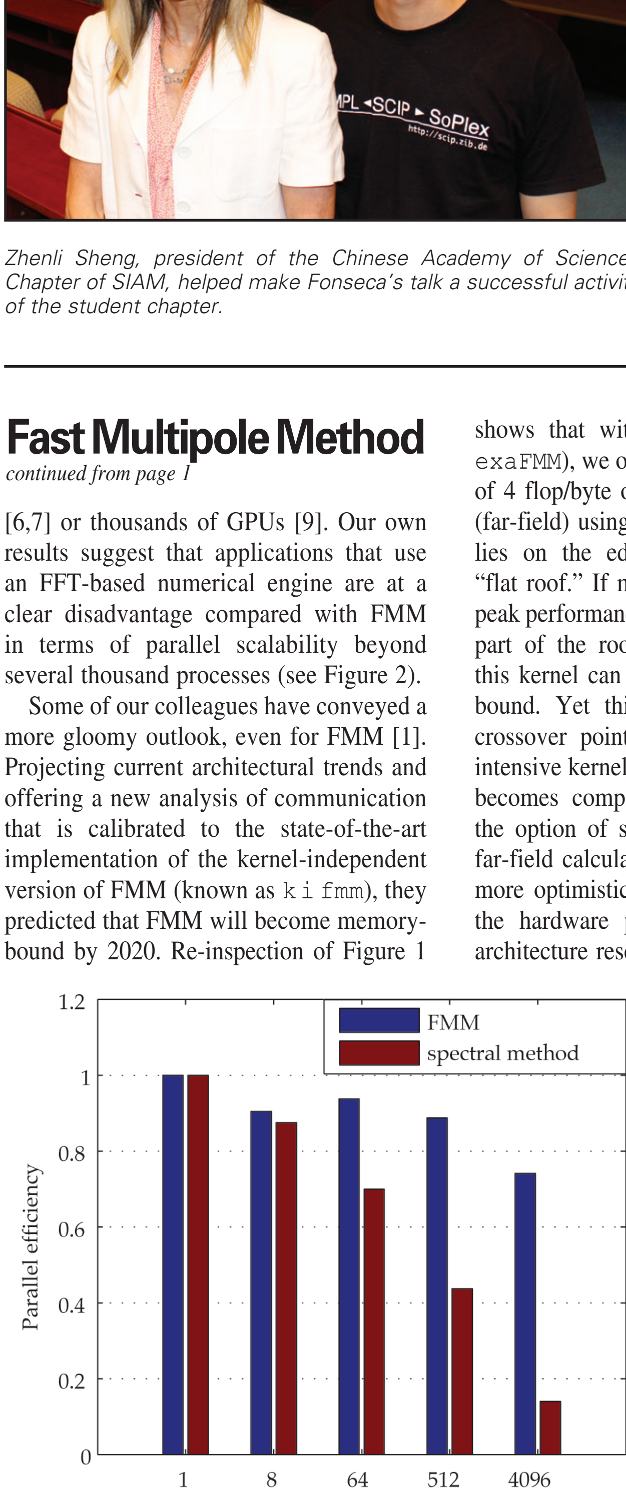

Figure 2. *Weak scaling from 1 to 4096 processes of two parallel application codes for fluid turbulence, one using an FMM-based solver on GPUs (one GPU per MPI process), the other an FFT-based solver on CPUs. Figure used under CC-BY license; doi:10.6084/m9.figshare.92425.*

Another cause for optimism is the new spirit of collaboration emerging in the field. Leading up to the SIAM CSE conference, several groups were exchanging ideas about building a set of standard benchmark tests for FMM codes and, in due course, developing a community software library. A handful of open-source codes are already available, but adoption of multipole algorithms would thrive if we had a BLAS-like collection of highly optimized inner kernels for FMM and an FFTw-like auto-tuning framework that selects the best inner kernels for the application given the available hardware. An open collection of FMM kernels in a flexible framework would not only make it easy to compare various mathematical and implementation ideas, it would also promote collaboration and reduce duplication of efforts. Add to this a set of micro-benchmarks, a website that hosts the data sets and benchmark performance reports, and a mechanism for anyone to add a new data set or benchmark problem, and we would accelerate our collective pace.

The 2013 SIAM CSE conference was a milestone for the number of FMM researchers who came together in one place. On the basis of the handshakes and face-to-face conversations that took place in Boston, we may be ready to start building a virtual community for exascale FMM.

References

[1] A. Chandramowlishwaran, J.W. Choi, K. Madduri, and R. Vuduc, *Brief announcement: Towards a communication optimal fast multipole method and its implications at exascale*, Proceedings of the ACM Symposium on Parallel Algorithms and Architectures (SPAA), Pittsburgh, Pennsylvania, June 2012.

[2] M. Darbas, E. Darrigrand, and Y. Lafranche, *Combining analytic preconditioner and fast multipole method for the 3-D Helmholtz equation*, J. Comput. Phys., 236 (2012), 289–316.

[3] Z. Gimbutas and L. Greengard, *Fast multiparticle scattering: A hybrid solver for the Maxwell equations in microstructured materials*, J. Comput. Phys., 232:1 (2012), 22–32; preprint on http://arxiv.org/abs/1104.5293.

[4] K.L. Ho and L. Greengard, *A fast direct solver for structured linear systems by recursive skeletonization*, SIAM J. Sci. Comput., 34:5 (2012), A2507–A2532.

[5] H. Langston, L.F. Greengard, and D. Zorin, *A freespace adaptive FMM-based PDE solver in three dimensions*, Commun. Appl. Math. Comput. Sci., 6:1 (2011), 79–122.

[6] I. Lashuk, A. Chandramowlishwaran, H. Langston, T. Nguyen, R. Sampath, A. Shringarpure, R. Vuduc, L. Ying, D. Zorin, and G. Biros, *A massively parallel adaptive fast-multipole method on heterogeneous architectures*, Proceedings of the Conference on High Performance Computing, Networking, Storage, and Analysis, SC 2009, Portland, Oregon, 1–12; Commun. ACM, 55:5 (2012), 101–109; doi:10.1145/2160718.2160740.

[7] M. Winkel, R. Speck, H. Hübner, L. Arnold, R. Krause, and P. Gibbon, *A massively parallel, multidisciplinary Barnes-Hut tree code for extreme-scale* N*-body simulations*, Comput. Phys. Commun., 183:4 (2012), 880–889.

[8] R. Yokota and L.A. Barba, *A tuned and scalable fast multipole method as a preeminent algorithm for exascale systems*, Int. J. High-Perform. Comput. Applic., 2012; published online Jan. 24 (2012); doi:10.1177/1094342011429952; preprint on http://arxiv.org/abs/1106.2176.

[9] R. Yokota, L.A. Barba, T. Narumi, and K. Yasuoka, *Petascale turbulence simulation using a highly parallel fast multipole method*, Comput. Phys. Commun., 184:3 (2013), 445–455; preprint on arXiv:1106.5273.

[10] B. Zhang, B. Lu, X. Cheng, J. Huang, N.P. Pitsianis, X. Sun, and J.A. McCammon, *Mathematical and numerical aspects of the adaptive fast multipole Poisson-Boltzmann solver*, Commun. Comput. Phys., 13:1 (2013), 107–128.

Lorena A. Barba is an assistant professor in the Department of Mechanical Engineering at Boston University. Rio Yokota is a research scientist in the Strategic Initiative for Extreme Computing at KAUST in Saudi Arabia.

Quantitative Evidence Often a Tough Sell in Court

Math on Trial. *By Leila Schneps and Coralie Colmez, Basic Books, New York, 2013, 255 pages, $26.99.*

BOOK REVIEW
By Michael O. Finkelstein

This book discusses in detail ten legal cases in which mathematical testimony was introduced in evidence. Each case is assigned a "Math Error Number" and given a modestly catchy title. Most of the cases are well known and have been written about repeatedly in law reviews and books, as the authors acknowledge in their Sources section. So the book is a revisiting of some legal chestnuts (plus some other less well known cases) in which math appears to have played a role.

The authors have evidently done their homework, and they give us a raft of new details about the cases and their context. There is room here to discuss only three of the cases.

■ ■ ■

What appears to be the first case in which a mathematician testified to mathematical probability involved one Hetty Green, the niece of Sylvia Ann Howland, a wealthy unmarried woman living in New Bedford, Massachusetts. This case the authors call Math Error Number 9: Choosing a Wrong Model. When Sylvia died, in 1865, she left a will in which she bequeathed to Hetty a life interest in a trust of her large estate. But Hetty, unsatisfied with this generosity, produced a separate one-page document, purportedly signed by Sylvia, that gave her Sylvia's estate outright, and purported to invalidate any future wills to the contrary. Hetty claimed that Sylvia had signed the document immediately before she had signed an earlier will. The executor rejected the claim on the ground, among others, that Sylvia's signature was a forgery, and the case became a cause célèbre.

In the lawsuit that followed, mathematics entered the fray when Benjamin Peirce, a professor of mathematics at Harvard, working with his son, Charles Sanders Peirce, later a famous logician and philosopher, testified to the probability that so many of the downstrokes in the disputed and authentic signatures of Sylvia would coincide. To do this, they took 42 admittedly genuine documents signed by Sylvia and computed the number of matching downstrokes in each of 861 pairs of signatures. From this they calculated that the probability of coincidence of a downstroke in a pair of signatures was about 1/5. Since all 30 downstrokes in the disputed and authentic signatures coincided, Peirce père, using a binomial model, concluded that the probability of 30 coinciding downstrokes if both signatures were authentic was $(1/5)^{30}$ = 1 in 2,666 millions of millions of millions. The calculation is not quite right, as the authors and others have pointed out, but the major fault lies in the inappropriateness of the binomial model, with its assumption of independence of pairs of downstrokes and the constant p = 1/5 across pairs of documents.

This much is well-trodden ground. What the book adds of interest is its discussion of the case's rich context. However, I was troubled by the authors' statement that followed their description of Benjamin Peirce's testimony: "What he meant, *and what the jury took him to mean* . . . [italics added]" was that the chance of two signatures being identical was negligible. How could the authors know what the jury thought in this mid-19th-century case? And then I recollected that because the case was in equity—Hetty sought specific performance of the agreement recited in the one-page document—it would have been heard by a judge without a jury. In that case, the authors would have made up that "fact," pure and simple. They also concluded that the judge "simply opted to reject Hetty's testimony altogether, and the case ended with a settlement that was essentially identical to Sylvia's latest will." This coda leaves the impression that the court unfairly brushed aside Hetty's claim. But a federal statute of the time provided that, in actions by or against an executor, neither party was allowed to testify against the other as to any transactions with the deceased. Because Hetty was the sole witness to the purported signing of the agreement, the circuit judges correctly dismissed her case.

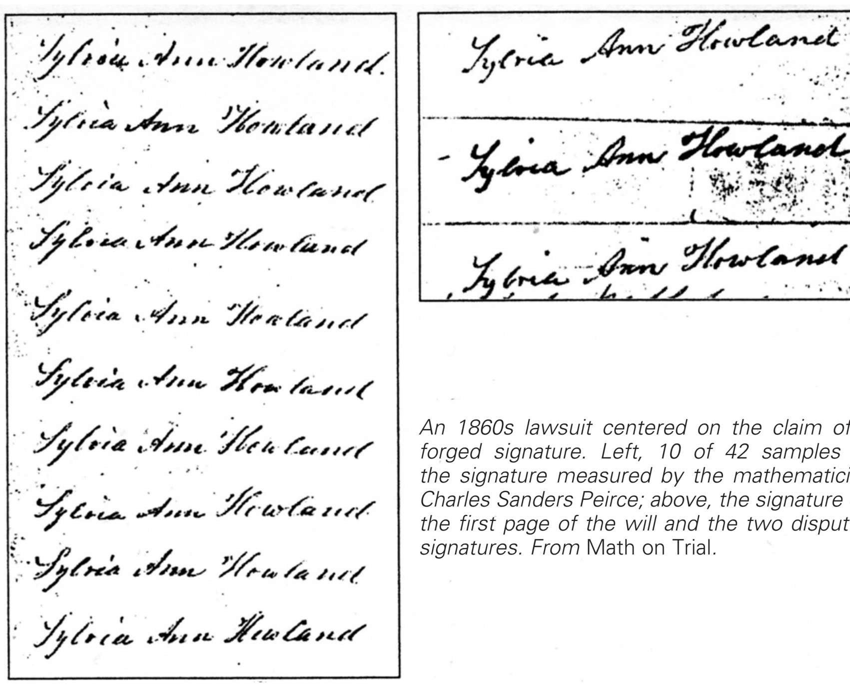

An 1860s lawsuit centered on the claim of a forged signature. Left, 10 of 42 samples of the signature measured by the mathematician Charles Sanders Peirce; above, the signature on the first page of the will and the two disputed signatures. From Math on Trial.

These errors and infelicities are evidence of some bias (and sloppiness) on the part of the authors: They wanted to portray the mathematics as potent and the parties in these cases as victims of the misuse of mathematical models. Some of them undoubtedly were. The case of Sally Clark (Math Error Number 1: Multiplying Non-independent Probabilities), tried in England for the murder of two of her babies, is a horrible example of a miscarriage of justice that makes painful reading. Even there, however, the role of mathematics may not have been as important as the authors suggest. Both deaths were cases of sudden infant death syndrome, or SIDS. In addition to adducing much medical testimony, the Crown called the eminent pediatrician Roy Meadow. He testified that the risk of one

See **Math in Court** *on page 7*

Geometry Meets Art in an Arcimboldesque Assemblage

Manifold Mirrors: The Crossing Paths of the Arts and Mathematics. *By Felipe Cucker, Cambridge University Press, New York and Cambridge, UK, 2013, 432 pages (including three indices: names, concepts, symbols), $29.99.*

BOOK REVIEW
By Philip J. Davis

Flipping through a publisher's catalog that just arrived, I find a selection of mathematics books described variously as:

- "For the mathematical novice who wants to explore the universe of abstract mathematics."
- "For teachers and students so that they can access the subject matter."
- "To activate student learning and comprehension."
- "To supplement a standard undergraduate course."
- "To add to the mathematics curriculum of life science students."

In publicizing their books, authors and publishers try to make clear the purpose of the books and their intended audience. As I read, or more likely skim, a new book, I usually ask myself the same questions: What is the book about and who is it for?

On the back cover of *Manifold Mirrors*, the author and the publisher assert clearly who and what the book is for:

"The book began life as a liberal arts course and is certainly suitable as a textbook. However, anyone interested in the power and ubiquity of mathematics will enjoy this revealing insight into the relationship between mathematics and the arts."

In his book, Felipe Cucker, who is Chair Professor of Mathematics at City University of Hong Kong, and whose specialty is foundational aspects of numerical algorithms, has interpreted "arts" broadly. He considers, at the very least, painting, printmaking, photography, design, decoration, the weaving of rugs and carpets, dance, poetry, architecture, and music.

The mathematics discussed is perhaps less sweeping in scope, but the reader will find in Cucker's book mini or micro lectures on linear transformations, symmetries, groups, tilings, projections, perspective, the seven frieze and the 17 wallpaper design groups, axioms, formal languages, elements of music, and hyperbolic, projective, and spherical geometries. Here and there the reader will find proofs of some of the theorems mentioned; a prefatory user's manual offers vague suggestions as to how a reader might deal with the proofs. Cucker gives a critique of the axiomatic method as set out in Euclid's *Elements* and expounded by Hilbert and later writers, e.g., Marvin J. Greenberg.

Manifold Mirrors is liberally illustrated, with images from the worlds of both geometry and art. Many of the works of art are in color; a number of them were quite unfamiliar to me. As to how works of famous artists are connected to mathematics, Cucker uses as an example the emergence ("eclosion") of pictorial arrangements in elliptical form in Renaissance art.

Cucker's text is liberally strewn (and decorated) with quotations from famous authors. He includes—confesses, in fact, that he cannot resist including—an often quoted mistranslation of Euripides on mathematics and art: "Mighty is geometry; joined with art, resistless." He quotes Leonard Bernstein on music as mathematics, Charles Dodgson (Lewis Carroll) on existence, and Gödel—in his incompleteness theorem—on consistency. Cucker points out that the "greatest English poets," among them Shakespeare, Milton, and Wordsworth, preferred iambic pentameter for their poetry.

You can also read here a metrical analysis of the canons of Johann Sebastian Bach, and you will find a detailed explanation of how Poincaré visualized hyperbolic geometry. Cucker quotes Richard Feynman on a particularly beautiful gate in Neiko, Japan:

". . . when one looks closely he sees that in the elaborate and complex design along one of the pillars, one of the small design elements is carved upside down; otherwise the thing is completely symmetrical. If one asks why this is, the story is that it was carved upside down so that the gods will not be jealous of the perfection of man."

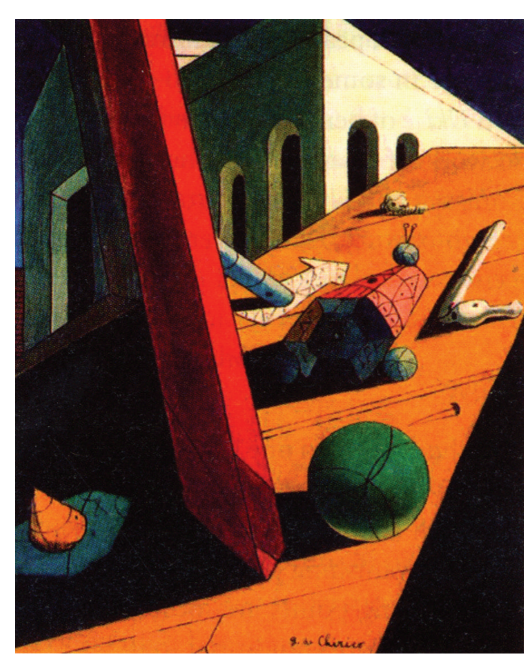

The Evil Genius of a King, Giorgio de Chirico, 1914–15. From Manifold Mirrors.

Mini and micro biographies pepper the text. We meet, among many others, Gauss, Felix Klein, and George Birkhoff, along with J.S. Bach, John Milton and Jorge Borges, Vincent Van Gogh and Maurits C. Escher, art critic and historian Ernst Gombrich, Bohemian music critic Eduard Hanslick.

In short, Cucker has produced a pot au feu, an eclectic catch-all. There is much that can be learned from Cucker's presentation of the marriage of mathematics and art.

Giuseppe Arcimboldo (1526–1593), an Italian painter to the Hapsburg Court in Prague and Vienna, was famous for creating portraits from assemblages of fruit, vegetables, flowers, grains, animals. I consider *Manifold Mirrors* Arcimboldesque in that it is an assemblage of many basic mathematical ideas and constructs, adding up to . . . well, to a unique work. Answering the questions I mentioned at the outset of this review, the book is a fine source book for teachers of mathematics who, in laying out a curriculum, want to go beyond the dry sequences of definitions, theorems, proofs and yet who have little interest in standard applications—of analytical mechanics, say. The book is a Newton-free environment.

Philip J. Davis, professor emeritus of applied mathematics at Brown University, is an independent writer, scholar, and lecturer. He lives in Providence, Rhode Island, and can be reached at philip_davis@brown.edu.

The Interdisciplinary Quest for New Ideas in Economics

The Physics of Wall Street: The History of Predicting the Unpredictable. *By James Owen Weatherall, Houghton Mifflin Harcourt, New York, 2013, 304 pages, $27.00.*

BOOK REVIEW
By James Case

Anyone wishing to understand the world of finance must be aware of the role that mathematical models now play in it. Twenty years ago, Peter Bernstein's instant classic *Capital Ideas: The Improbable Origins of Modern Wall Street* was the place to learn. But because financial thinking and practice continue to evolve, the story requires periodic updating. Weatherall's ambitious new book begins with an attempt to fill this need.

Several of the early chapters (particularly the third and sixth) summarize material described elsewhere at book length. The third describes the life and times of the late Benoît Mandelbrot, who covered much of the same ground several years ago in a book [3] written with Richard L. Hudson, and whose autobiography [2] has since appeared. The sixth is a brief history of the Prediction Company of Santa Fe, New Mexico, which has been the subject of at least one full-length book [1]. Although the firm's original workforce—including founders Doyne Farmer and Norman Packard—has long since moved on, "Predco" remains an active (and perhaps wildly successful) subsidiary of UBS, the Union Bank of Switzerland.

It is difficult to gauge the degree of Predco's success because—perhaps in an effort to discourage imitation—management has chosen to remain secretive about the firm's actual earnings. The best assessment Weatherall could elicit is a quote from a "knowledgeable source" to the effect that, over the firm's first 15 years, "its risk-adjusted return was almost a hundred times larger than the S&P 500 return over the same period." If so, the firm's performance compares favorably with Warren Buffett's at Berkshire–Hathaway, or Ed Thorp's at Princeton–Newport Partners, though none of the three (according to Weatherall) can match the achievement of James Simons's Renaissance Technologies.

James Simons graduated from MIT in 1958, completed his Berkeley PhD (under S.S. Chern) in 1962, and worked as a code breaker at IDA during the run-up to the Vietnam war. He became chair of the Stony Brook mathematics department in 1968, and received the 1976 AMS Oswald Veblen Prize in Geometry, largely on the strength of the book *Characteristic Forms and Geometrical Invariants* he co-authored with Chern. The theories presented lie at the forefront of theoretical physics, especially string theory. In 1982 Simons took leave of academic life to found Renaissance Technologies, a hedge fund management firm that uses computer modeling to find inefficiencies in highly liquid securities. Its flagship product, the Medallion Fund, is a high-risk high-return vehicle in which only the firm's executives may invest. It regularly outperforms the firm's other funds, which are open to outside investors. According to Weatherall, Renaissance employs many mathematicians and physical scientists capable of understanding and improving the algorithms on which the firm relies, but no one trained in finance or economics. In 2006, Simons was named Financial Engineer of the Year by the International Association of Financial Engineers.

Benoît Mandelbrot is a central character in the book, due in part to his admiration for Louis Bachelier—the founder of mathematical finance and, arguably, of the entire theory of stochastic processes—which caused him to ferret out much of what is known of Bachelier's family background and early years. But Mandelbrot is significant mainly, at least in Weatherall's eyes, as a skeptic who warned tirelessly of the perils of over-reliance on derivative securities evaluated in a "Black–Scholes environment." In that environment, all stochastic processes are Gaussian normal processes, which exhibit significantly fewer extreme variations (either extremely large or extremely small) than do the historical price series they are intended to model. So it is hardly surprising that derivatives evaluated in such an environment appear less risky than they really are. Mandelbrot championed the substitution of "Levy-stable distributions," a one-parameter class that includes the slender-tailed normal distribution (for parameter value $\alpha = 2$), the fat-tailed Cauchy distribution (for $\alpha = 1$), and an entire continuum of intermediate distributions for intervening values of α.

Mandelbrot was not alone in observing the inadequacy of valuations performed in a Black–Scholes environment. The highly secretive firm O'Connor and Associates was founded in 1977 to exploit opportunities then becoming available on the newly opened Chicago Board of Trade. Realizing that the assumptions underlying the Black–Scholes formula were a first approximation at best, the founders set out to develop—in collaboration with first-generation quants Michael Greenbaum and Clay Struve—a modified Black–Scholes model that would be able to account for the sudden and dramatic price changes that can lead to fat-tailed distributions. The firm was famously successful, first in options and later in other types of derivatives, in part because their modification of the Black–Scholes model tended to outperform its parent.

The firm O'Connor and Associates was famously successful, first in options and later in other types of derivatives, in part because their modification of the Black–Scholes model tended to outperform its parent.

The true worth of the O'Connor model went undemonstrated until the stock-market crash of 1987, during which the S&P 500 index lost more than 20% of its value in a single afternoon. By anticipating the conditions under which the Black–Scholes model would fail, the firm not only survived but prospered. By 1992, it was among the largest players on the Chicago commodities market, with more than 600 employees and billions of dollars under management. Around that time, two O'Connor associates became aware of Predco, and persuaded their partners to invest in it. Whereas Farmer and Packard had rejected other suitors, they welcomed acquisition by O'Connor and Associates, in part because they needed the money to increase the scope of their operations, and in part because the acquiring firm seemed technically sophisticated enough to understand what they were up to. The deal was unexpectedly sweetened when, later in 1992, O'Connor was purchased outright by the giant Swiss Bank Corporation, and sweetened again when (in 1998) SBC merged with the even larger UBS.

The algorithms brought by Farmer and Packard to the analysis of financial data began as data-mining techniques designed by physicists like themselves to find small islands of predictability amid the vast

See **Wall Street** *on page 6*

Wall Street

continued from page 5

oceans of apparently patternless variation characteristic of chaotic time series. By allowing different algorithms to operate simultaneously on the same data stream, and to "vote" on proposed trades, they were frequently able to identify opportunities worth pursuing. They were thus among the first to employ the "black box" and "algorithmic" trading methods on which the hedge fund industry has come to rely.

■ ■ ■

The last three chapters in the book are the ambitious ones. The first of the three describes, in decidedly nontechnical language, Didier Sornette's efforts to discover warning signs that might predict impending catastrophes. One of Sornette's earliest projects involved ruptures in the high-pressure Kevlar tanks in which rocket fuel is sometimes stored during space flight. Such tanks performed admirably most of the time, but had a distressing tendency to explode, on rare occasions, without warning.

Sornette observed that a sequence of small earthquakes is more likely to presage a large one if the intervals between the small ones become log-periodically shorter. And in the summer of 1997, he observed log-periodic fluctuations in various stock market indices.

Lab tests revealed that, under pressure, the tanks routinely developed microscopic fractures that sometimes grew, by connecting with other fractures, to potentially dangerous size. Even the smallest fracture could potentially do this, much as liquids can sometimes percolate through marginally porous media. Sornette and his colleagues noticed that the tanks would begin to "rumble" as fractures began to appear, and that if the intervals between rumblings became "log-periodically" shorter—meaning that the peaks and valleys were roughly coincident with those of a function of the form $\cos(\omega \log(t_c - t))$—a critical event was likely to occur at or about the critical time t_c.

Later, in collaboration with his wife, Anne Sauron-Sornette, a geophysicist, Sornette observed that a sequence of small earthquakes is more likely to presage a large one if the intervals between the small ones become log-periodically shorter. And in the summer of 1997, in collaboration with friend and management scientist Olivier Ledoit, he observed log-periodic fluctuations in various stock market indices. The two responded by buying up a host of cheap "far-out-of-the-money" put options, entitling them to sell the underlying securities in October at August and/or September prices. When the Dow fell 554 points on Monday, October 27, and both the NASDAQ and S&P 500 indices dropped by comparable amounts, the two earned a 400% profit and documented the fact by releasing their Merrill Lynch trading statement. Sornette makes no claim that all catastrophes can be predicted in this way. But some of them obviously can, and the rewards for so doing are often substantial.

Weatherall's penultimate chapter concerns a second husband–wife team, consisting of mathematical physicist Eric Weinstein, now a hedge fund manager and financial consultant in Manhattan, and economist Pia Malaney. At Weinstein's suggestion, Malaney wrote her Harvard thesis on the applicability of gauge theory to the construction of economic indices, such as the NASDAQ, the S&P 500, the Dow, and (especially) the Consumer Price Index. The CPI is particularly important to ordinary citizens, as a variety of pension, health care, and Social Security benefits are adjusted annually to keep pace with the CPI. Gauge theory has not caught on in the economics community, despite the prominence and support of Malaney's thesis adviser, Eric Maskin (a 2007 Nobel laureate). Yet Weinstein remained convinced that the idea has merit, and eventually persuaded Lee Smolin, a physicist at Canada's Perimeter Institute (and author of a book titled *The Trouble with Physics*), that the proposal deserved a hearing.

Smolin responded by organizing a conference, held at Perimeter in 2009, and attended by physicists, mathematicians, biologists, mainstream economists, and financiers. His hope was to persuade all in attendance that economic theory is inadequate and in need of an extensive overhaul. Furthermore, he believed, nothing less than a giant research effort—on the scale of the World War II-era Manhattan Project—could reasonably be expected to produce the necessary revision. The conference itself, Weatherall writes, was quite successful. Those in attendance readily agreed that economic theory is inadequate. They failed, however, to reach consensus as to where the key deficiencies lie or how to fix them. Then too, lurking in the background, was the funding problem. How, if it were even possible to obtain support for something as vast as a new Manhattan Project, would the funds be split among cooperating disciplines? In the absence of any assurance that theirs would get its due, attendees lost interest and returned to their accustomed activities. Even Smolin has gone back to doing physics, having apparently decided that economics is a waste of scientific resources. Though economic problems seem reasonably tractable, he found the economics profession hostile to new and unfamiliar ways of thinking!

Today, Weatherall writes, even as Weinstein, Malaney, Sornette, Farmer, and a handful of others continue to develop nontraditional economic models, the world economy remains a shambles. The so-called recovery from the collapse of 2007–2008 lags far behind schedule. What, he asks, can be done to expedite real recovery and delay the next recession? His answer appears in the book's final chapter "Epilogue: Send Physics, Math, and Money!" The world, he submits, needs new and more fruitful economic ideas. To discover them, he sees no alternative to a monumentally large, copiously funded, interdisciplinary research initiative. Good luck selling that inside the beltway!

References

[1] T.A. Bass, *The Predictors*, Allen Lane, New York, 2000.

[2] B. Mandelbrot, *The Fractalist: Memoir of a Maverick Scientist*, Pantheon, New York, 2012.

[3] B. Mandelbrot and R.L. Hudson, *The Misbehavior of Markets*, Basic Books, New York, 2004.

James Case writes from Baltimore, Maryland.

NSF Announces Funding Opportunities

The Directorate for Computer and Information Science and Engineering at the National Science Foundation and the Office of Financial Research at the U.S. Department of Treasury invite proposals to the CIFRAM (Computational and Information Processing Approaches to and Infrastructure in support of Financial Research and Analysis and Management) program. The program is designed to identify and fund a small number of exploratory but potentially transformative research proposals in the area of finance informatics.

Principal investigators interested in seeking research support through the program should submit two-page white papers to nsfofr@nsf.gov. Proposals that involve collaborations between computer scientists, mathematicians, statisticians, and experts in financial risk analysis and management are especially welcome.

A full list of topics of research interest and complete details for the submission of proposals can be found at http://www.nsf.gov/pubs/2013/nsf13093/nsf13093.jsp. Additional information is also available from NSF program directors Vasant Honavar (vhonavar@nsf.gov) and Frank Olken (folken@nsf.gov), as well as from the OFR liaison, Mark Flood (Mark.Flood@treasury.gov).

NSF also invites proposals for its Catalyzing New International Collaborations program. The CNIC program is designed to promote the professional development of U.S. STEM researchers and to advance their research through international engagement. Grants will be made from the Office of International Science and Engineering to support research and education activities that present unique opportunities and offer potentially high benefits through collaboration with scientists and engineers abroad.

NSF will consider proposals from researchers at U.S. institutions for collaborative work with researchers in any country that is not explicitly proscribed by the Department of State. Activities can be in any field of science and engineering supported by NSF. This solicitation offers support for the initial phases of an international collaboration, with the strong expectation that in the next phase, U.S. investigators will submit proposals to an NSF Directorate for continued funding of the research initiated with a CNIC grant.

Details can be found at http://www.nsf.gov/funding/pgm_summ.jsp?pims_id=12815&org=OISE&from=home.

Principal investigators must communicate with NSF program directors prior to submission to this solicitation. For more information, prospective PIs should contact program directors Nancy Sung or R. Clive Woods at oise-cnic@nsf.gov.

MOR

continued from page 8

structural and thermal analysis. Currently, the approach makes extensive use of the recent static condensation reduced basis element (SCRBE) method for real-time simulations of complex structures/geometries, using libraries of components and providing these online for download. This platform has been designed for both companies and universities.

Despite the advances presented at SIAM CSE 2013 and briefly mentioned here, several challenges remain to be addressed. Among them are the handling of large numbers of parameters, and the definition of efficient parameter sampling and better exploration techniques for the construction of reduced order models. Fluid mechanics continues to offer many interesting challenges for the MOR community, such as the model order reduction of turbulent viscous flows at high Reynolds numbers and the efficient handling of features propagating in flows.

Overall, and in particular from the perspective of MOR, SIAM CSE 2013 can be considered a very successful event, gathering the computational science and engineering community to exchange information about recent developments, open questions, new insights, and interests that could result in further leading applications.

David Amsallem is a researcher in the Department of Aeronautics & Astronautics at Stanford University. Bernard Haasdonk is a junior professor in the Department of Mathematics at Universität Stuttgart. Gianluigi Rozza is a researcher at SISSA, International School for Advanced Studies, MathLab, Trieste.

Professional Opportunities

Math in Court

continued from page 4

such death in a family like Sally Clark's was 1 in 8543, that of two such deaths 1 in 8543×8543 = 1 in 73 million. The authors omit mention of the rebuttal testimony and the judge's comments on the evidence, stating simply that Meadow's figure "was accepted without question by judge and jury."

This is somewhat misleading. There was rebuttal testimony, by a Professor Berry, who challenged the 1:73 million figure on the ground that "familial factors" could lead to two SIDS deaths in a family. More significantly, the judge—in commenting on the evidence, as English judges are allowed to do—cautioned the jury not to put much reliance on the statistics: "We do not convict people in these courts on statistics." He also noted, "If there is one SIDS death in a family, it does not mean there cannot be another one in the same family." The jury evidently questioned the evidence, because two members voted against conviction (the jury convicted because 10 members of the jury found her guilty). On the first appeal, the court dismissed objections to the 1:73 million figure, pointing out that the arguments against squaring were known to the jury and that the trial judge repeated them in his summing up. On the second appeal, the prosecution announced that it would no longer defend the verdict, citing newly discovered evidence that one of the babies had died of an infection, as well as misleading statistics. The appellate court then reversed the conviction on the ground of the newly discovered evidence, and, in dictum, excoriated the statistics as "manifestly wrong" and "grossly misleading."

A third case is the redoubtable *People v. Collins*, decided by the California Supreme Court in 1968 (Math Error Number 2: Unjustified Estimates). This is a granddaddy of legal chestnuts; almost every law school evidence course includes it. The story is quickly told: An elderly woman, while walking in an alley, was assaulted from behind and robbed. A witness said that a Caucasian woman with dark-blond hair in a ponytail ran out of the alley and entered a yellow automobile driven by a black man with a mustache and a beard. A couple answering roughly to that description (there were some variances) was subsequently arrested and tried. At the trial the prosecutor called an instructor of mathematics from a nearby college to testify to the product rule of elementary probability. He then had the witness assume the individual probabilities of six relevant characteristics, which were evidently plucked by the prosecutor from thin air (e.g., interracial couple in car—1/1000; black man with beard—1/10; woman with blond hair—1/3), and then apply the product rule to multiply them together. The witness came up with 1/12,000,000 as the probability of a couple answering to that description. The jury convicted, but on appeal the Supreme Court of California reversed the conviction on the ground that there was no evidence to support the individual probabilities and multiplying them was clearly improper as they were not independent. The authors describe the mathematical exercise in greater detail than is usual and approve the Supreme Court's rejection of the multiplication exercise because of lack of foundation and the obvious lack of independence of the factors. Nothing new here.

A deeper point is raised by an appendix to the Supreme Court's opinion (evidently written by Harvard law professor Laurence Tribe, who was then a clerk to the chief judge). Using a mathematical model (which I have criticized elsewhere), the appendix purports to show that if the rate of such couples was indeed 1/12,000,000 in some population, there was a 40% chance there would be two such couples. From this, the appendix concludes that "the prosecutor's computations . . . imply a very substantial likelihood that a couple other than the Collinses was the one at the scene of the robbery."

The authors call the arguments in the appendix "flawless and convincing." I do not agree. First, the rate of such couples in the population, by itself, does not give us probabilities of guilt or innocence, but only a likelihood ratio for the evidence, which is not the same thing. Second, the Bayesian case for guilt is overwhelming, and it is surprising that the authors, who belong to the Bayes in Law Research Foundation, did not at least allude to it. If there were two such couples, the likelihood ratio for this evidence would be on the order of 12,000,000 (assuming a probability of 1 that the witnesses would have so described the Collinses if they were guilty). With that enormous figure (or any figure in the millions), it would take only a wisp of other evidence, when combined with the likelihood ratio, to make the posterior odds of guilt overwhelming. And there was more than a wisp; indeed, the other evidence was rather significant (although the authors labor to disparage it). There is an important Bayesian teaching here that the appendix, the authors, and others have failed to recognize: A tell-tale trace doesn't have to identify a suspect uniquely to make a powerful case when combined with other evidence.

■ ■ ■

In these and other cases, the authors paint mathematical evidence as extremely powerful. But in all these cases, other evidence made the role mathematics actually played in the decisions unclear. On that subject I must report my experience with a large group of moot court cases in which mathematics was the sole evidence. For many years, at the end of my course Statistics for Lawyers, at Columbia Law School (and other law schools), we have put on a moot court trial in which the students are divided into two teams and given the facts of a case and a data set. Each side also gets a professor of statistics as an expert witness. The students' job is to prepare their witness for direct testimony, prepare to cross-examine the opposing witness, and give opening and closing arguments to the jury. The data set is designed to support a finding for the plaintiff, although with some weaknesses. Over many years of these trials, with different data sets, different expert witnesses, as well as generations of students and juries, a common theme has emerged: *The proponent of statistics almost always loses*. On questioning, jurors say in various ways that they didn't believe the statistics. Evidently, quantitative evidence (particularly if presented in complex statistical models) is not an easy sell to legal decision makers and may have been less influential in actual cases than the authors suggest in this book.

Applied mathematicians who read the book should not take away the impression that the legal chestnuts discussed fairly represent the role of mathematics in legal proceedings today. Much has changed in the law, in addition to DNA evidence, which the book does discuss. Well-qualified statisticians and economists from academia and consulting organizations now routinely testify to multiple regression models in a variety of disputes, including those involving employment discrimination, voting, the death penalty, and antitrust. Binomial models are used in jury discrimination challenges. The U.S. Supreme Court has approved a number of such models and in a landmark opinion has required the district courts to determine whether proposed scientific testimony is sufficiently valid and reliable to be heard in evidence. To help the judges, the Federal Judicial Center produced the fat *Reference Manual on Scientific Evidence*, which has proved to be a best seller and is now in its third edition. Law reviews devoted to quantitative studies of the legal system have made their appearance, and there are textbooks on law and statistics (I am the co-author of one). This is a growing field, with variety and interest extending far beyond the world of the ten cases discussed in *Math on Trial*.

Michael O. Finkelstein, a retired lawyer, has taught statistics for lawyers at the law schools of Columbia, Harvard, Yale, and the University of Pennsylvania.

CSE 2013

A Conference Within a Conference for MOR Researchers

By David Amsallem, Bernard Haasdonk, and Gianluigi Rozza

Model order reduction (MOR) is increasingly important in the mathematical and engineering sciences. With the availability of highly accurate discretization and simulation techniques for systems of ordinary and partial differential equations, the demand for accelerated solution of these high-dimensional models has become more and more urgent. This is particularly so in multi-query or real-time simulation scenarios, such as simulation-based control and optimization, statistical investigation, and interactive design, to name just a few.

Acceleration of simulation models is typically achieved through surrogate modeling, based on the projection of high-dimensional quantities, such as the state vector, onto low-dimensional subspaces. MOR procedures aim to achieve good approximation of states or the input–output behaviour of the full model. A typical feature of MOR is a decoupling of the computational procedure: In a possibly expensive "offline" step, the reduced order model is generated using significant computational resources. This can include, for instance, full model simulations for carefully chosen configurations, selected either by numerical experts or by automatic procedures. Simulation of the reduced model is then carried out in an inexpensive "online" step; this enables rapid simulations for many different configurations on less powerful computational platforms.

As a topic inherently linked to simulation science, MOR was discussed at previous SIAM CSE conferences, but its presence in the program increased dramatically at the 2013 conference. MOR was the subject of an invited plenary presentation, by Jan Hesthaven of Brown University, who gave a broad introduction to certified MOR by reduced basis methods (a recording of the talk can be accessed at http://www.siam.org/meetings/presents.php). It was also the theme of 22 minisymposia, a total of 88 presentations, held in the course of the week; eight of the minisymposia were organized by the authors of this article. With a track consisting of several minisymposia held in the same room throughout the week, the sessions constituted something of a conference within a conference. The remainder of this article highlights noteworthy recent developments and research directions featured in the many conference sessions devoted to MOR.

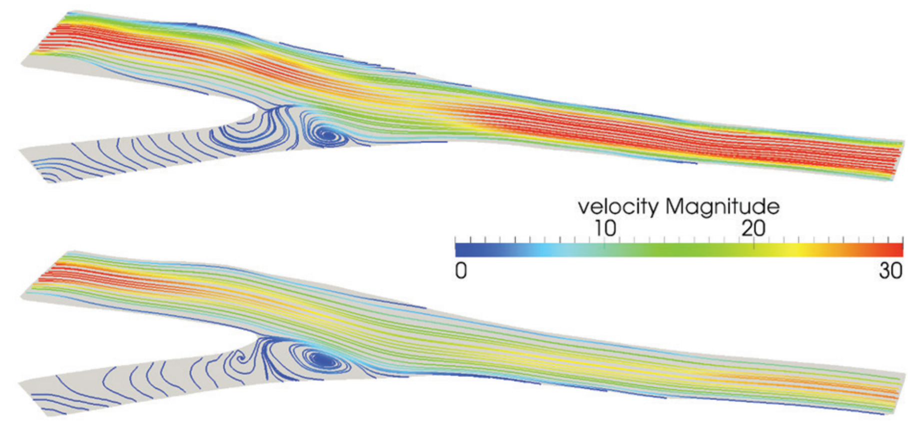

Figure 2. *Shape optimization of a bypass anastomosis (visualization of the velocity field) using the reduced basis method applied to viscous flow equations with the iterative optimization algorithm of vorticity reduction. From A. Manzoni, Reduced Models for Optimal Control, Shape Optimization and Inverse Problems in Haemodynamics, PhD Thesis 5402, EPFL, 2012; T. Lassila, A. Manzoni, A. Quarteroni, and G. Rozza,* A reduced computational and geometrical framework for inverse problems in haemodynamics, *International Journal for Numerical Methods in Biomedical Engineering, in press, 2013.*

■■■

Several talks were devoted to structure-preserving model reduction. Presenters demonstrated that an appropriate MOR technique should always preserve certain characteristics, such as passivity or Lagrangian structure of the underlying large-scale dynamical system, in order to obtain physical and accurate predictions.

While many MOR techniques are based on simulations that use high-fidelity models, such as finite element or finite volume schemes, speakers also presented novel data-based and data-driven MOR approaches. With these techniques, researchers use experimental data and observations to build reduced order models. Several speakers presented novel approaches based on solution snapshots generated from physical experiments, as well as the use of Loewner matrices, kernel interpolation schemes, and Kalman filtering.

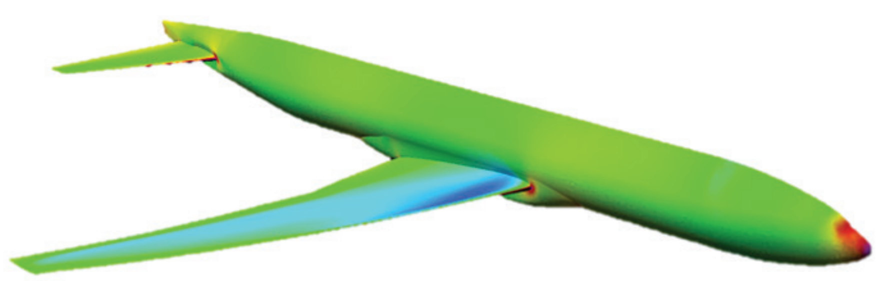

Figure 1. *Pressure at the surface of a commercial aircraft computed using model order reduction. From K. Washabaugh, D. Amsallem, M. Zahr, and C. Farhat, Nonlinear Model Reduction for CFD Problems Using Local Reduced-Order Bases, 42nd AIAA Fluid Dynamics Conference and Exhibit, June 25–28, 2012, New Orleans, Louisiana.*

Rapid generation of predictions is one objective of MOR, as is the certification of accuracy and reliability of those predictions by a priori and a posteriori error estimators. Several speakers discussed such error quantification. In particular, a new space–time norm analysis framework guarantees the effectivity and reliability of the resulting error estimators, and the stability of the reduced models for long-time integration of flow problems.

Many presentations highlighted difficulties associated with the efficient reduction of nonlinear problems. In such cases, an additional level of approximation is required to achieve large online speedups. Among the proposed approaches are empirical interpolation, gappy proper orthogonal decomposition, missing-point estimation, and hyper-reduction procedures.

Other speakers proposed novel approaches for reducing the complexity associated with expensive bottlenecks, such as spatial integral evaluations and the computation of euclidean inner products of large-scale vectors.

The parametric dependency of systems is an important question in MOR. Most parametric MOR procedures are limited to the case of small numbers of parameters, as high parameter dimensions are associated with the curse of dimensionality and issues related to effective sampling. Realistic applications, however, are typically characterized by high-dimensional parameter spaces. Three new approaches that can be effective in such scenarios were presented at SIAM CSE: (1) parameter space reduction for inference (e.g., in subsurface contaminant analysis), (2) state space partition by local reduced bases (e.g., in CFD), and (3) coupling of submodels by reduced basis elements, such as "LEGOs" (e.g., in stress analysis of component-based structures).

Many talks at CSE 2013 were devoted to optimal control and optimization with reduced order models.

Several sessions detailed the application of MOR to complex systems in general, and to computational fluid dynamics problems in particular. Examples include flow control problems applied to airplane design (Figure 1) and cardiovascular devices (Figure 2). In addition to CFD, application fields ranged from environmental sciences, building design, and computational neuroscience to the design of semiconductors and car engines. A notable effort to transfer MOR technology to industry was described by a speaker from Akselos, a startup company that provides an easy-to-use interactive platform for real-time simulations in

See **MOR** on page 6

Cardiac Modeling

continued from page 1

is simulated by a reaction–diffusion partial differential equation. The reaction–diffusion PDE describes current flow through tissue composed of myocytes that are electrically connected via low-resistance gap junctions. Cardiac tissue has orthotropic electrical conductivities that arise from the cellular organization of the myocardium (cardiac muscle) into fibers and laminar sheets. Global conductivity values are obtained by combining the fiber-and-sheet organization with myocyte-specific local conductivity values. Current flow in the tissue is driven by ionic exchanges across cell membranes during the myocyte action potential. Simulation of electrical wave propagation in the myocardium entails simultaneous solution of the PDE and the set of action potential ODEs over the tissue volume. In certain cases, such as simulation of the delivery of an external current to the myocardium, use of a system of coupled PDEs rather than a single PDE allows for an explicit representation of current flow in the extracellular space.

The progress made in simulating cardiac electrical behavior at the organ level is the most exciting, and it is there that the author's laboratory has made its most significant contributions. In general, many of the emergent, integrative behaviors in the heart result not only from complex interactions within a specific level but also from feed-forward and feedback interactions connecting a broad range of levels in the biological hierarchy. The ability to construct multiscale models of the electrical functioning of the heart, representing integrative behavior from the molecule to the entire organ, is of particular significance, as it paves the way for clinical applications of cardiac organ modeling. In the clinic, as in labs in which electrophysiological measurements are made in intact animals, assessment of function is conducted at the level of the entire organ. Validating the whole-heart simulations at that level ensures their predictive capabilities and the possibility of their entry into the realm of healthcare.

The key to attaining predictive capabilities for multiscale cardiac models at the organ level has been the use of individual models, either MRI- or CT-based, of the geometry of the heart, and the application of diffusion tensor MR imaging (DTMRI) to measure the anatomy, fiber, and sheet structure of the heart in cases of ex vivo studies. This has led to a new generation of whole-heart image-based models with unprecedented structural and biophysical detail. Clearly, modeling the function of the heart has benefitted significantly from the revolution in medical imaging.

Cardiac models have been used to gain insight into mechanisms of arrhythmia in many disease settings. In addition, a major thrust in computational cardiac electrophysiology is to use models as a testbed for the evaluation of new antiarrhythmic drug therapies. It is now possible to test hypotheses regarding mechanisms of drug action on the scale of the whole heart. Multiscale heart models of antiarrhythmic drug interactions with ion channels have provided insight into why certain pharmacological interventions result in drug-induced arrhythmia, whereas others do not. This work has the potential to help guide the drug development pipeline—a process well known for both high failure rate and high cost.

The use of heart models in personalized diagnosis, treatment planning, and prevention of sudden cardiac death is also slowly becoming a reality. The feasibility of subject-specific modeling has been demonstrated through the use of heart models reconstructed from clinical MRI scans to evaluate infarct-related ventricular tachycardia (fast and often lethal arrhythmia). Such applications of whole-heart modeling are expected to help predict optimal locations for the procedure known as catheter ablation in the hearts of individual patients, as well as to stratify patients for arrhythmic risk.

Nowadays, whole-heart image-based multiscale models of cardiac electrical behavior are not the whole story. Combined with models of cardiac mechanics, these models are being integrated into increasingly "multi-physics" whole-heart models of cardiac electromechanical behavior. Such multiscale whole-heart electromechanics models are being applied in a patient-specific manner to investigate improved methods for cardiac resynchronization therapy in dyssynchronous heart failure. Recently, in a step toward comprehensive modeling of cardiac function, efforts have been made to link electromechanical models of the heart with fluid dynamics models.

Computer simulations of the function of the diseased heart represent an important research avenue in the new discipline of computational medicine. Biophysically detailed models of the heart assembled with data from clinical imaging modalities that incorporate electromechanical and structural remodeling in cardiac disease could become a first line of screening for new therapies and approaches, new diagnostic developments, and new methods of disease prevention. Bedside implementation of patient-specific cardiac simulations could become one of the most thrilling examples of CSE-based approaches in translational medicine.

Natalia A. Trayanova is the Murray B. Sachs Professor in the Department of Biomedical Engineering and the Institute of Computational Medicine at Johns Hopkins University.

This article is based on the author's invited lecture at the SIAM CSE meeting in Boston, in February 2013. A recording of the talk has been posted at http://www.siam.org/meetings/presents.php.