Newsjournal of the Society for Industrial and Applied Mathematics

siam news

Volume 48/ Number 5
June 2015
sinews.siam.org

A Minimalist Minimizes an Integral

MATHEMATICAL CURIOSITIES

By Mark Levi

In this issue we present a solution that is shorter than Johann Bernoulli's famous optics-based idea of minimizing

$$\int_{\gamma OA} F(y)\,ds \quad (1)$$

over smooth curves connecting two given points A and B; here $F(y) > 0$ is a given function and ds is an element of arc length. Bernoulli based his beautiful solution on the equivalence between Fermat's principle and Snell's law.

The following solution, in addition to being shorter, substitutes a mechanical analogy for Bernoulli's optical one—and thus could have been given by Archimedes.

The sum

$$P_N = \sum F(y_k)\Delta s_k$$
$$\approx \int_{\gamma AB} F(y)\,ds$$

can be interpreted mechanically as the potential energy of the system of rings and springs shown in Figure 1. Each of the N rings slides without friction on its own line; the neighboring rings are coupled by constant-tension springs whose tensions are given by the discretized values of $F_k = F(y_k)$. If P_N is minimal, each ring is in equilibrium, implying the balance of horizontal forces on the ring:

$$F(y_k)\sin\theta_k = F(y_{k+1})\sin\theta_{k+1},$$
$$k = 1,\ldots N;$$

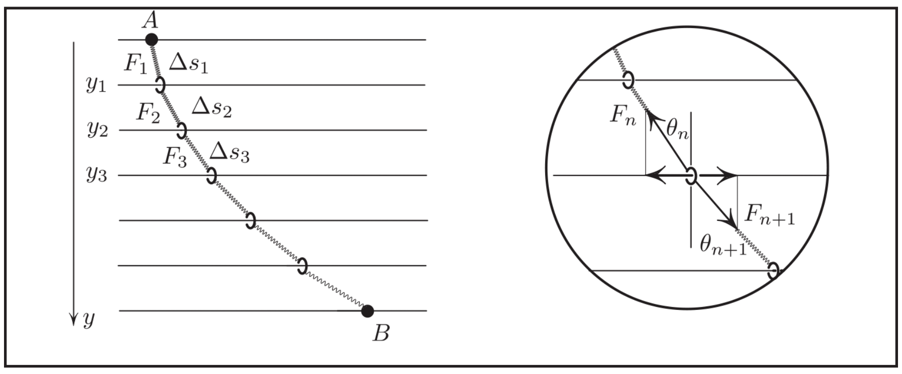

Figure 1. *Each spring has a prescribed tension F_k independent of its length Δs_k. The endpoints A and B are held fixed.*

in the continuous limit this gives

$$F(y)\sin\theta = \text{constant},$$

or, equivalently,

$$\frac{F(y)}{\sqrt{1+(y')^2}} = \text{constant}.$$

A discussion of this idea (along with some others in a similar spirit) can be found in [1].

References

[1] M. Levi, *Classical Mechanics with Calculus of Variations and Optimal Control*, AMS, Providence, Rhode Island, 2014.

Stay tuned: A related topic is explored in the next Mathematical Curiosities column (July/August issue).

Mark Levi (levi@math.psu.edu) is a professor of mathematics at The Pennsylvania State University. *The work from which these columns are drawn is funded by NSF grant DMS-1412542.*

Congratulations to New NAE/NAS Members

Members of the SIAM community honored this year by election to the U.S. National Academies of Engineering and Sciences include David Srolovitz of the University of Pennsylvania (NAE), pictured above, and Alan Hastings of the University of California, Davis (NAS). See page 5.

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

On Markov Decision Processes

By Nicole Bäuerle and Viola Riess

Sequential planning under uncertainty is a basic optimization problem that arises in many different settings, ranging from artificial intelligence to operations research. In a generic system, we have an agent who chooses among different actions and then receives a reward, after which the system moves on in a stochastic way. Usually the aim is to maximize the expected (discounted) reward of the system over a finite or, in certain cases, as described below, an infinite time horizon.

To obtain a tractable problem, it is often assumed that the transition law of the underlying state process is Markovian, i.e., that only the current state has an influence on future states. Such a situation leads to a Markov decision process (MDP); textbooks on MDPs include [1, 3, 5, 7]. MDPs differ from general stochastic control problems in that the actions are taken at discrete time points, rather than continuously. Stochastic shortest-path problems, consumption and investment of money, allocation of resources, production planning, and harvesting problems are a few examples of MDPs.

The formulation and development of MDPs started in the 1950s with Shapley, Bellman, and, later, Howard, Dubins, Savage, and Blackwell. The early achievements are closely related to the study of stochastic dynamic games. Subtle mathematical problems in the theory include measurability issues with arbitrary Borel state spaces, which naturally arise in, for example, partially observable Markov decision processes. However, with a problem that has enough structure, the solution algorithm is a very simple backward induction algorithm, which works as follows:

Suppose that the decision time points are numbered $n = 0, 1, \ldots, N$, the one-stage rewards are r_n, the terminal reward is r_N, the transition kernel of the state process is $Q_n(\cdot \mid x,a)$, and the set of admissible actions at stage n given state x is $D_n(x)$. The value functions $V_n(x)$ that represent the maximal expected reward starting in state x from stage n up to the final stage N can then be recursively computed by the following optimality equation:

$$V_N(x) = r_N(x). \quad (1)$$

$$V_n(x) = \sup_{a\in D_n(x)} \Big\{ r_n(x,a) + \int V_{n+1}(y) Q_n(dy \mid x,a) \Big\}. \quad (2)$$

When we finally reach V_0 we have computed the maximal expected reward of the system up to the time horizon N; the maximizers in this recursion yield the optimal strategy. This algorithm is very straightforward; the real challenge in applying the theory comes in the "curse of dimensionality." Going forward in time, the number of states that can be reached often increases dramatically with the number of admissible actions. For large n, the number of optimization problems to be solved is vast. If, for example, every state can have just two possible successors, then at stage n there are already 2^n different states for which we have to solve (2). This is not feasible for most interesting applications.

But we do have hope of being able to solve such problems. If we have a stationary model and a long time horizon, one trick is to approximate the problem by one with an infinite time horizon. With an infinite time horizon, we are left with a fixed-point problem to solve; various algorithms are available for such problems, including policy iteration and linear programming.

Many problems are not stationary, however. In such cases we need to use tools from approximate dynamic programming (see, for example, [6]). Loosely speaking, the solution methods combine backward induction with a forward simulation of states. The idea is to improve a given approximation of the value functions along a simulated path of the state process.

■■■

In the remainder of this article, we look at a specific application: valuation of a gas storage facility. A storage facility, which is often a depleted reservoir in an oil or gas field or a salt cavern, can be used not only to balance supply and demand but also to create profit through active storage management on a mark-to-market basis. Contracts for natural gas storage essentially represent real options on natural gas prices.

To find a fair price for storage, we first need a stochastic model for the gas pricing process. Most gas price models are continuous. One of the first models is the Schwartz model (see [8]), in which the log-price dynamics for gas is given by

$$d\log(P_t) = \alpha(\mu_t - \log(P_t))\,dt + \sigma_t dW_t,$$

where W_t is a Brownian motion, σ_t is the time-dependent volatility of the process, μ_t is its mean, and α is the mean-reversion factor.

The valuation problem can now be solved as an MDP. The state of the problem is the current storage level x, together with the current gas price p; the action a is the change in the amount of gas. The set of admissible actions is quite complicated (Figure 1 shows a typical set)—the capacity

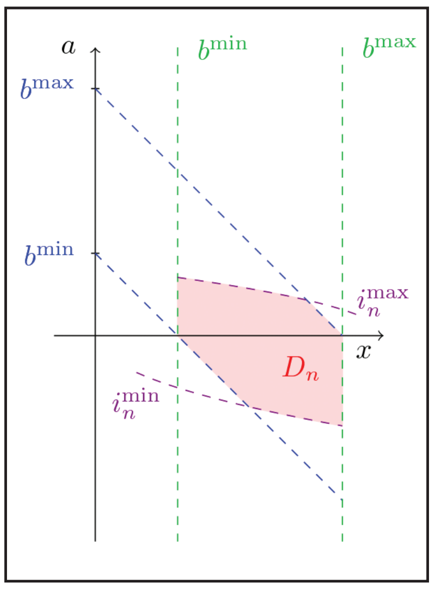

Figure 1. *Typical set of admissible actions.*

See **MDPs** *on page 3*

siam news
Volume 48/ Number 5/ June 2015

4 **Numerical Notation Systems as Cultural Artifacts**
Numerical notation systems may have been in use as early as 30,000 BC and have been developed in cultures worldwide. Nonetheless, writes reviewer Ernest Davis, because systems are well defined and not very numerous, it is possible to present an account that is "essentially complete and definitive"—a case in point being the book under review.

5 **U.S. National Academies Elect New Members**

5 **European Students Gather at TU Delft for Krylov Day**

6 **The Moody's Mega Math Challenge Marks 10th Year**
The Moody's/SIAM math modeling contest for high school students celebrated its 10th anniversary with a field just short of every state in the U.S. (the goal for next year). Some 5,000 students (1,128 teams) tackled this year's question: Is college worth it? Rachel Levy (who traces her own career path to an undergraduate experience in math modeling) brings the milestone events to life.

8 **New Mathematics for Extreme-Scale Computational Science?**
With the approach of the extreme-scale computing era, Ulrich Rüde dispels some misconceptions about the development of algorithms and software for modern computer systems. What's needed, he writes, "is not simply a few extra weeks for converting Matlab to Fortran and MPI. Designing efficient HPC software requires extensive creative research."

2 **Obituaries**

7 **Professional Opportunities**

7 **Announcements**

Obituaries

Michael James David Powell, who passed away on April 19, was one of the giants who established numerical analysis as a major discipline and created its current intellectual landscape. From his life work have emerged both mathematical foundations and practical algorithms of nonlinear optimization, as well as decisive contributions to approximation theory.

Mike Powell was born in London and went to school at Eastbourne. In 1957, having completed his National Service, he went to Cambridge to read mathematics. The standard duration of studies for the Mathematical Tripos at Cambridge is three years, but in those more flexible times, Mike accomplished this in just two years, followed by a one-year diploma in Numerical Analysis and Computing. Then, instead of staying in academia and working toward a doctorate, he joined the Atomic Energy Research Establishment Harwell, where he stayed for seventeen years.

In his first few years at AERE Harwell, Mike worked on questions in computational chemistry. Then, in 1962, came his first paper on optimization, a subject he would make his own. Historically, there were two "master methods" for optimization without constraints: firstly, the Newton algorithm, clearly impractical for large-scale computations with many variables because of the prohibitive cost of the evaluation of the Jacobian matrix in each iteration and the consequent linear algebra; secondly, the method of steepest descent, iterating locally in a direction determined by the gradient, and representing the ultimate demonstration that locally optimal decisions can be disastrous globally.

In a 1959 paper, Bill Davidon proposed an algorithm that used an approximate Jacobian, now called a "variable-metric method." For Mike the paper was a revelation. In 1962, he and a younger colleague, Roger Fletcher, published an extremely influential paper on what is now known as the DFP algorithm, acknowledging Davidon's pioneering contribution. This augured the start of a life journey for Mike and arguably the beginning of modern optimization, and was followed by extensive further research into many aspects of (mostly, but not always, unconstrained) optimization: from the convergence of DFP and BFGS (Broyden–Fletcher–Goldfarb–Shanno) variable-metric methods, to trust-region methods, local line searches, conjugate gradient methods for nonlinear problems, augmented Lagrangian functions, sequential quadratic programming, derivative-free methods, and so forth. Mike was engaged in this work until the last week of his life.

Numerical analysts tend to divide into two classes: those who subject numerical algorithms to hard analysis and full mathematical treatment, yet regard practical programs as an afterthought, best left to others, and those who focus on software issues and practicalities of implementation, while regarding analysis as an often unnecessary encumbrance—if it works, who needs a proof? Mike Powell was an exception. He firmly believed that hard analysis and beautifully written programs go hand in hand and that his responsibility, as a numerical analyst, was both to produce deep and challenging mathematics (his proofs of the convergence of the DFP and BFGS algorithms for convex functions are a striking example of a truly difficult, nonintuitive—and often counterintuitive—rigorous mathematical proof) and to create (and freely share with the community) professionally written software of the highest quality.

The Harwell terms of engagement, "peaceful use of atomic energy," allowed Mike a great deal of freedom to plough his own furrow, first and foremost in optimization, but also in approximation theory, and he was instrumental in setting up the Harwell library of numerical subroutines. Then, in 1976, he returned to Cambridge (receiving a Doctor of Science degree in 1979) as the John Humphrey Plummer Professor of Applied Numerical Analysis.

This was a momentous change in many ways. At Harwell Mike spent all his time on research, surrounded by kindred souls—Roger Fletcher, Alan Curtis, John Reid, Iain Duff, and others. At Cambridge he was expected to undertake the numerous duties of a "proper" professor—teaching, supervision of research students, administration, committee work—which he often regarded as a drain on time best spent doing research. Still worse, while Cambridge has had a glorious tradition in numerical analysis, from Isaac Newton onwards, by the 1960s this tradition had essentially died out. Thus, Mike was expected to establish numerical analysis from scratch in the Department of Applied Mathematics and Theoretical Physics, in an atmosphere in which anything but fluid dynamics was often seen as an improper occupation for a true applied mathematician. It is fair to say that Mike was an outlier in a large department, in what was then a wasteland betwixt the pure and applied mathematics departments at Cambridge. Until his retirement in 2001, Mike led a small group—ultimately, just two "teaching officers" (Cambridgese for "faculty") and a small cohort of research students, postdocs, and visitors.

Mike's interest in approximation theory started in Harwell, first in connection with least-squares calculations, ℓ_1 and ℓ_∞ approximations, and then in his very influential work on splines. But he did what may be his most memorable and influential work in this area, on radial basis functions, at Cambridge. The spur was a beautiful paper of Charlie Micchelli proving that, regardless of dimension, the problem of approximation by radial functions is nonsingular subject to fairly broad conditions. This created the promise of an exceedingly powerful interpolation method for multivariate scattered data but also opened a host of questions about the quality of such approximation. These questions have been addressed—and in large measure answered—by Mike and his research students, thereby creating the groundwork for the many subsequent applications of radial basis functions, not least in the computation of partial differential equations.

Michael J.D. Powell, 1936–2015. Photo by Lin Wang, courtesy of the Chinese Academy of Sciences.

Mike Powell was a unique character. He readily confessed to disliking administration, bureaucracy, paperwork, committees, and even teaching—anything that ate into valuable research time. Indeed, he retired early to focus more on his research (and on his golf handicap). Yet his sense of duty was such that, once unhappily compelled to spend time on any of these chores (and although a perfect English gentleman, Mike was never good at hiding his dislike or impatience), he discharged them with total commitment, in an exemplary fashion. In particular, his teaching (like his talks) was always crystal-clear and immaculately prepared: not a word, not a symbol out of place, everything logical and in the right sequence.

This sense of duty and Mike's total integrity made him a terrible academic politician: Everybody knew that, push come to shove, Mike Powell would support what he believed was *right*: There was little point to horse-trading or exchanging favours with him. He did not believe that his role as an academic was to build an empire or demonstrate formal "academic leadership": He led strictly by example, producing world-class research, educating his students well, and inspiring others.

Academic honours duly arrived. In 1982 Mike was awarded (jointly with Terry Rockafellar) the inaugural SIAM George Dantzig Prize and a year later was elected a Fellow of the Royal Society. He received, among others, both the Naylor and Senior Whitehead Prizes of the London Mathematical Society (becoming the only person ever to receive two senior LMS prizes), the IMA Gold Medal and its Catherine Richards Prize, foreign membership in the U.S. National Academy of Sciences, corresponding fellowship in the Australian Academy of Sciences, a PhD *honoris causa* from the University of East Anglia.

On a personal note, Mike was a colleague, a neighbour, and a friend for 37 years. He was fiercely competitive, but also generous to a fault. With Caroline, he was a wonderful host. He was also a mentor for a young and inexperienced Junior Research Fellow, and a shining example thereafter. His standards were always high and demanding, for those around him but in particular for himself. His students were relatively few, but he trained them ever so well and pushed them to excel themselves. But he also genuinely cared about them and their lives; in return, they demonstrated fierce loyalty, as did his many friends worldwide.

We take so much for granted because this is our reality as numerical analysts: from variable-metric algorithms to methods for multivariate approximation, but also

See **Michael Powell** *on page 3*

siam news

ISSN 1557–9573. Copyright 2015, 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 2015: 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. $70.00 subscription rate worldwide for print copies. 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.*

Should Your Research Be on YouTube?

This year's SIAM Conference on CSE featured a new attraction: "SIAM Communication Doctors," a booth for people wishing to craft effective messages about their research—for communication to future employers, at outreach events, or for the press. Graduate students, postdocs, and faculty visited the booth, hoping that booth doctors could turn their research summaries into good stories that would appeal to the public. At the booth, reporter Flora Lichtman, whose work has appeared in *The New York Times* and on NPR's Science Friday, joined conference attendees Nick Higham (University of Manchester), Jeff Humpherys (Brigham Young University), Rachel Levy (Harvey Mudd College), and Matt Parno (Massachusetts Institute of Technology) to offer feedback as people pitched their ideas.

Jesus Ramos of Mountain View College hones a story about his research with the assistance of science journalist Flora Lichtman.

"Having given many talks and workshops on science/math communication, I can tell you SIAM's communication booth was exceptionally effective," Lichtman said. "The one-on-one setup meant visitors walked away with communication tips tailored to their particular research. It also gave "doctors" a better sense of some of the challenges researchers face when communicating about their work. A great model." In an earlier *SIAM News* article, Levy, Lichtman, and David Hu (Georgia Tech University) (http://bit.ly/1AKYOAD) offered tips on scientist–reporter collaborations.

"I think one of the biggest things I took away from my conversation at the Communication Doctors booth was the tip to have a few overview pictures, slides or sketches handy to discuss my research in case anyone shows interest," said David Gleich, an assistant professor of computer science at Purdue University. "What you take for granted as 'standard knowledge' is often quite surprising to others."

Based on the level of interest at the CSE conference, SIAM invites interested readers to submit video clips in which they briefly describe their work and explain why they should be given help to produce a YouTube video communicating the work. The pitch video should answer these or similar questions: What is your topic? Why is it important? Who will benefit from your work? Why do you need help communicating about your research?

Submissions will be judged on prospective content, enthusiasm, and audience. The winner will receive professional advice and coaching on research communication, along with the services of a professional videographer to film and edit the video. To apply, please prepare your video pitch (we hope that, although unprofessional, submissions will be creative and interesting!) and send it to videopitch@siam.org by July 10. A winning entry will be selected by August 14; all coaching and production work will be completed by the end of 2015.

The winning clip will be posted on YouTube on the SIAM channel, included in the "SIAM Presents" pages, and promoted via SIAM's media distribution lists.—*Rachel Levy, SIAM VP for Education, and Michelle Montgomery, SIAM.*

AAAS & the SIAM Community

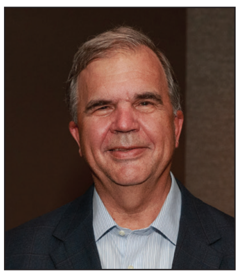

Photo by David Sytsma, Corporate Chicago Photography.

Jim Crowley

Kirk Jordan

At its 2015 meeting in San Jose, California, the American Association for the Advancement of Science recognized AAAS fellows elected in 2014. Among those from the Section on Mathematics are ***James Crowley****, executive director of SIAM ("for a distinguished record as a scientific administrator in the U.S. Air Force and for two decades of outstanding leadership as executive director of SIAM");* ***Charles Epstein****, Thomas A. Scott Professor of Mathematics and chair of the Graduate Group in Applied Mathematics and Computational Science at the University of Pennsylvania ("for distinguished contributions to applied analysis, especially microlocal analysis, index theory, and boundary value problems; and significant achievements in the mathematics of medical imaging"); and* ***Kirk Jordan****, an IBM Distinguished Engineer and associate program director in the Computational Science Center, Data Centric Systems, IBM T.J. Watson Research in Cambridge, Massachusetts ("for leadership and significant achievements in computational applied mathematics, especially in high-performance and parallel computing applied to fluid dynamics, systems biology, and high-end visualization").*

Selected and supported by SIAM to participate in the AAAS Mass Media Science & Engineering Fellows program in the summer of 2015 is Anna Lieb, a fourth-year PhD student in applied mathematics at the University of California, Berkeley. Shown here discussing her poster ("Optimizing Intermittent Water Supply") with AWM executive director Magnhild Lien at the 2013 SIAM Annual Meeting, Lieb will spend the 10-week fellowship period at NOVA.

Designed to strengthen connections between scientists and journalists, the AAAS program places students at the advanced undergraduate, graduate, and postgraduate levels at media organizations throughout the U.S. Fellows have worked as reporters, editors, researchers, and production assistants at radio and television stations, newspapers, and magazines. AAAS Mass Media Fellows, as described on the program website, "research, write and report today's headlines, sharpening their abilities to communicate complex scientific issues to non-specialists."

Students who consider this an appealing way to spend a summer are urged to apply for 2016 fellowships. The application process begins at the end of the year; details will appear at www.aaas.org or can be obtained from Jim Crowley (jcrowley@siam.org).

MDPs

continued from page 1

of the gas storage is of course restricted, and the maximal speed at which gas can be injected or withdrawn depends on the storage level.

To include transaction costs, a loss of gas at the pump, and/or a bid–ask spread in the market, we introduce for a quantity of gas the "ask price" k and the "bid price" e, which we assume to be affine functions of the price. The one-stage reward function $r_n(p,a)$ of the MDP is given by $-k(p) \cdot a$ if $a > 0$, by 0 if $a = 0$, and by $-e(p) \cdot a$ if $a < 0$.

The terminal reward function depends on the contract. Having specified the data, we can easily write down the optimality equation. The first step then is to get as much information as possible about the value function and the maximizers from the optimality equation in order to simplify the numerical algorithms. Here, indeed, it is possible to figure out (by induction) the structure of the optimal injection and withdrawal strategy (see, for example, [2, 9]), which can be characterized by three regions that depend on the gas price p: When the current gas storage level is below a certain bound $\underline{b}(p)$, it is optimal to inject gas, either as much as possible or up to $\underline{b}(p)$, whichever occurs first. If the current gas storage level is above a certain bound $\overline{b}(p)$, it is optimal to withdraw gas, either as much as possible or down to $\overline{b}(p)$, whichever occurs first. When the level is in between, the optimal strategy is to do nothing.

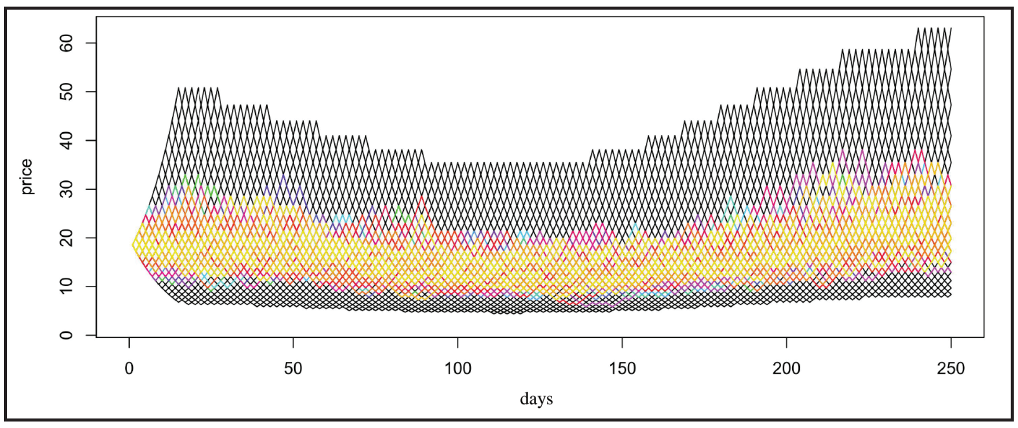

Figure 2. *Price grid for the recombining tree (black) with simulated price paths.*

As described in [2], we solved the gas storage problem using different numerical algorithms, all based on a combination of the backward induction algorithm and knowledge of the structure of the optimal strategy. One algorithm used a recombining (linearly growing) tree to approximate the gas pricing process (see Figure 2) and hence avoid the curse of dimensionality. Another algorithm combined backward induction with a forward simulation of the gas pricing process, with linear regression on a finite number of basis functions to approximate the value function (see also [4]). The resulting optimal strategy can be seen in Figure 3. These algorithms work rather well and can also be used for more complicated gas pricing process models.

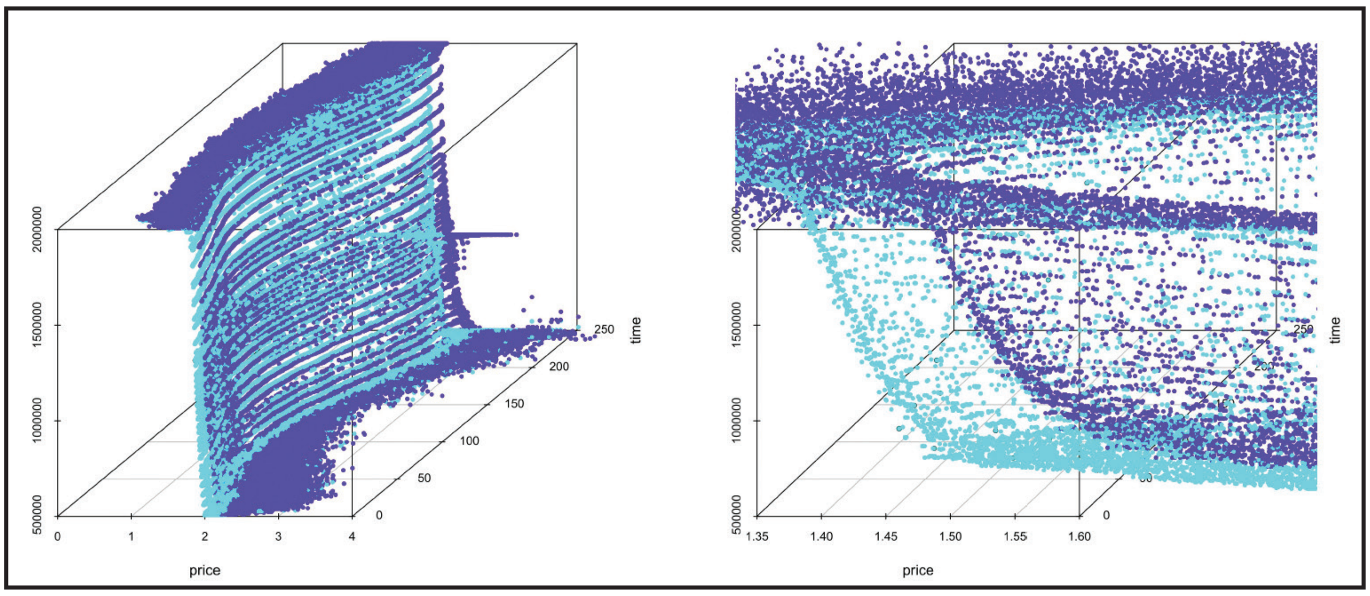

Figure 3. *Strategy bounds for a regression method:* $\underline{b}_n(p)$ *(light blue) and* $\overline{b}_n(p)$ *(dark blue).*

References

[1] N. Bäuerle and U. Rieder, *Markov Decision Processes with Application to Finance*, Springer, Berlin, 2011.

[2] N. Bäuerle and V. Riess, *Gas storage valuation with regime switching*, (2014), arXiv: 1412.1298.

[3] D.P. Bertsekas, *Dynamic Programming and Optimal Control, Vol.* 2, Athena Scientific, Belmont, MA, 2012.

[4] A. Boogert and C. de Jong, *Gas storage valuation using a multifactor price process*, J. Energy Markets, 4 (2011), 29–52.

[5] O. Hernández-Lerma and J.B. Lasserre, *Discrete-time Markov Control Processes*, Springer, Berlin, 1996.

[6] W.B. Powell, *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, John Wiley & Sons, Hoboken, NJ, 2011.

[7] M.L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, John Wiley & Sons, Hoboken, NJ, 2009.

[8] E. Schwartz, *The stochastic behavior of commodity prices: Implications for valuation and hedging*, J. Finance, 52 (1997), 923–973.

[9] N. Secomandi, *Optimal commodity trading with a capacitated storage asset*, Management Science, 56 (2010), 449–467.

Nicole Bäuerle is a professor of mathematics and Viola Riess is a junior research scientist at Karlsruhe Institute of Technology, Germany.

This article was contributed by the SIAM Activity Group on Control and Systems Theory, via Francois Dufour of the University of Bordeaux, the group's SIAM News *liaison. The biennial conference of SIAG/CST will be held in Paris this summer, July 8–10.*

Michael Powell

continued from page 2

the very idea that a numerical algorithm is a creature with a double personality—a mathematical entity on call for rigorous mathematical analysis and a computational scheme that must be programmed and implemented with a similarly high level of cleverness. Mike Powell, in his life's work and attitudes, demonstrated these twin motives of numerical analysis and their underlying unity at their very best. He will be missed.—*Arieh Iserles, University of Cambridge.*

Numerical Notation Systems as Cultural Artifacts

Numerical Notation: A Comparative History. *By Stephen Chrisomalis, Cambridge University Press, New York, 2010, 496 pages, $114.99.*

BOOK REVIEW

By Ernest Davis

The construction of written notation systems for numbers, closely following the development of words for numbers, is one of the most fundamental and most widespread initial steps in the development of mathematics. Numerical notation can be found on artifacts from Egypt and Mesopotamia dating back to about 3200 BC, in Egypt concurrent with the earliest written records of language, and in Mesopotamia predating written records by centuries. But it is possible that numerical notation is even older: Artifacts with markings that are believed to be numerical tallies have been found from as early as 30,000 BC.

Among cultural artifacts, numerical notation systems are unusual in two respects. First, they are extremely well defined and limited semantically; they represent some or all of the natural numbers, and sometimes certain fractions. Second, they are not very numerous; records document only about 100 distinct systems that have ever been in any kind of general use. (The exact count depends on how you individuate.) It is therefore possible to present an account of numerical notation systems throughout history that is essentially complete and definitive, up to the limits in the historical record.

Stephen Chrisomalis's *Numerical Notation: A Comparative History* is such an account. He gives complete descriptions of all known numerical notation systems: how they work and how they have been used. Each system is illustrated with a clear hand-drawn table of symbols; some are also accompanied by photographs showing their use on historical or archaeological artifacts.* Chrisomalis also traces the evolutionary history of these systems. When, as often, this history is obscure, he surveys the scholarly literature and makes his own judgment of the probable historical relations between notations, taking into account such considerations as similarity of structure, similarity of symbols, known contacts between cultures, and proximity in time and space.† Considering the complete survey, he analyzes a substantial number of regularities and near regularities that govern individual systems, and a much smaller number of regularities that govern how one numerical system can evolve from another. Finally, he discusses how these regularities relate to characteristics of human cognition and human society.

Figure 4.1 from Numerical Notation*—The Etruscan "abacus-gem" (CII 2578 ter) showing a figure seated at a board working with Etruscan numerals. Source: Fabretti 1867: 224. Courtesy of Cambridge University Press.*

Chrisomalis identifies five main structures for notational systems. All numerical notation systems are built around powers of a fixed base, always a multiple of 10, and generally 10. A *cumulative-additive* system, such as Roman numerals, has a symbol for each power of the base (I, X, C, M); these are repeated and the values are then added. (The Roman numerals also have a *subbase* of 5 (V, L, D), common in additive systems, and a *subtractive* feature (IX for 9), which is extremely rare.) In a *ciphered-additive* system, such as the Greek or Hebrew numerals, each multiple of a power of 10 has its own symbol, and the values of these are added together. For instance, in the Greek alphabetic system, ν represents 400, λ 30, and δ 4; $\nu\lambda\delta$ thus represents 434.

In a *multiplicative-additive system*, signs for the digits 1 through 9 alternate with signs for the power of 10. Traditional Chinese numerical notation works this way; so (to some extent) does the English language, e.g., "two *thousand* three *hundred* for*ty*-seven." In *ciphered-positional* notation, such as the Western numerals, there are symbols for the numbers from either 0 or 1 up to the base minus 1; the power of the base is then indicated by the position of the symbol in the numeral. To be unambiguous (not all such systems are), a system must have either a symbol for zero or some other way of indicating powers with a zero coefficient. Finally, *cumulative-positional* systems represent powers of ten positionally, as in the Western numbers, but the coefficients cumulatively; the famous base-60 ancient Babylonian system followed this principle. (A ciphered base-60 system would of course need 60 distinct symbols for the digits.)

About 30% of the systems that Chrisomalis discusses are hybrids that combine different principles for different ranges of numbers—often, a cumulative or ciphered-additive system for lower powers of the base and a multiplicative-additive system for higher powers. However, no naturally arising systems of pure numbers use any other principles. One can imagine a system that represents numbers by their prime factorization, or that uses division (as in "a half-dozen"), or that uses the factorials as a base (e.g., representing 301 as [2,2,2,0,1], since $301 = 2 \cdot 5! + 2 \cdot 4! + 2 \cdot 3! +1 \cdot 1$), etc.; but such systems do not actually arise.

■■■

Chrisomalis's historical accounts are always impeccably clear, but unavoidably somewhat dry; after 100 numerical notation systems, one's eyes begin to glaze. However, he provides all kinds of fascinating historical and cultural tidbits along the way. Large numbers go back to the very earliest days of numerical notation; an Egyptian macehead from 3100 BC records the supposed capture of 120,000 prisoners, 1,422,000 goats, and 400,000 cattle. Some numerical systems were used only for counts of quite specific categories; in fact, in ancient Uruk in

See **Notation** *on page 5*

---

*In the online copy that I was reading, these photographs were not always very clear, but the reproduction quality appears to be better in the print version.

†Some accounts of the scholarly literature in the area suggest wildly conjectured relations between notational systems that are similar in some respects, but separated by millennia.

U.S. National Academies Elect New Members

Among the most important honors accorded to scientists and engineers in the U.S. is election to the National Academies of Engineering and Sciences. The National Academies were created (NAS in 1863 by President Abraham Lincoln; NAE in 1964) to advise the federal government in matters of science and technology. At their annual meetings, each announces the names of newly elected members. Distinguished members of the SIAM community appear regularly on both lists, and 2015 is no exception.

New members of NAE, announced in February, include **Ingrid Daubechies**, James B. Duke Professor of Mathematics at Duke University. Cited "for contributions to the mathematics and applications of wavelets," research for which she is well known in SIAM circles, Daubechies has made works of art a recent focus of her research. These efforts include the development and use of a new method for the virtual restoration of digital paintings, such as the Ghent Altarpiece of 1432.

NAE recognized **Michael Todd**, Leon C. Welch Professor in the School of Operations Research and Information Engineering at Cornell University "for contributions to the theory and application of algorithms for continuous optimization." His research interests are in algorithms for linear and convex programming, particularly semidefinite programming and ellipsoid optimization. He also works in the development and analysis of interior-point methods. A long-time editor of *SIAM Journal on Optimization* (1997–2007), he is also a former chair of the SIAG on Optimization (2011–14).

David Srolovitz, pictured on the first page, is the inaugural Joseph Bordogna Professor of Engineering and Applied Science at Penn. He is a member of several departments, including Materials Science, and is co-chair of the organizing committee for SIAM's 2016 Conference on Mathematical Aspects of Materials Science.

David D. Yao, Piyasombatkul Family Professor and professor of industrial engineering and operations research at Columbia University, was cited "for understanding of stochastic systems and their applications in engineering and service operations."

Elected a foreign associate of NAE was **Martin Vetterli**, a professor of communication systems at Ecole Polytechnique Federale de Lausanne; he was cited "for development of time-frequency representations and algorithms in multimedia signal processing and communications."

New members of the National Academy of Sciences, announced this spring, include **Donald Geman**, **Alan Hastings**, and **Moshe Vardi**. Geman, a professor of applied mathematics at Johns Hopkins University, offers an appealing introduction to his main research interests (computational vision and computational medicine) on his website. Hastings, pictured on the first page, is a professor in the Department of Environmental Science and Policy at UC Davis; his research interests include mathematical biology, with a focus on theoretical ecology and population biology. Vardi is a professor of computer science at Rice University, where he is also director of the Ken Kennedy Institute for Information Technology. In 2015 he was also named a SIAM fellow.

Notation

continued from page 4

Sumeria, there were 15 different numerical systems for different kinds of quantities, including "the regular Š system [for] barley, the Š' system for germinated barley for brewing beer, and the Š* system for barley groats." In modern China, six numerical systems are to some degree active, depending on the region and the particular use. One of these, the "accountants' system," uses deliberately complex symbols in order to avoid falsification.

Chrisomalis emphasizes strongly that the use of numerical notation varies significantly from one culture and time to another—we should not make the mistake of supposing that our own uses of numbers apply universally. In particular, in most times and places, written numbers were not used for calculation; calculations were done by some method of finger calculation or with tools, such as an abacus or counting sticks. It is therefore a mistake to suppose that the inefficiency of a notational system for calculation was any kind of drawback.

By way of analogy (mine, not Chrisomalis's), consider the numerical notation for dates, e.g., 2/24/2015 (American style) for February 24, 2015. What is it good for? Well, it makes it easy to approximate the time between dates, particularly if they don't cross a boundary: 7/15/2015 is about 5 months after 2/24/2015; 9/26/1898 was about 117 years earlier. It is also easy to judge the relation of dates to yearly events: 12/25/1898 was Christmas, was about 4 days after the solstice, and was the beginning of winter in New York and of summer in Sydney. That's about it. Calculating the day of the week for 12/25/1898 or the exact number of days between 12/25/1898 and 2/24/2015 is laborious by hand, and requires several lines in a computer program (and you have to be very careful to avoid off-by-one errors).

The irregularity of the calendar is a constant source of inefficiency and trouble for the construction of calendars, either printed or automated. Why do we put up with this? First, we rarely have to compute the number of days that have elapsed since some date in the past (though we do often have to determine the day of the week of a future date). Second, the costs of changing it would be prohibitive. Third, because an earth year happens to be 365.2425 earth days, no calendar that incorporates both years and days can possibly be very elegant.[‡] If our descendants living on seasonless space stations make fun of us for measuring time in such an obviously awkward way, they will simply be missing the point.

Even within the basic Western numerical notation, there are suboptimalities that we tend to overlook because we are so used to them. Who knows how many man-hours and dollar-equivalents have been lost over the last five centuries because the handwritten digits 4 and 9, and 1 and 7, are easily confused. The Roman numerals are much clearer in that respect.

■■■

To me, the most interesting part of Chrisomalis's book is his analysis of the regularities that govern numerical systems. He adduces 14 principles that hold in all the systems he has studied, 8 that hold in nearly all. Among the universals: "Every base is a multiple of 10"; "Any system that can represent $N + 1$ can also represent N." For the near-universals: "No numerical notation explicitly represents arithmetic operations such as addition and multiplication" (in contrast are linguistic forms, such as "a thousand *and* fourteen," "vingt-*et*-un"). The single exception occurs in the Shang Chinese numerals: "All numerical notation systems are ordered and read from the highest to the lowest power of the base." This is, of course, necessarily true for positional systems, but would not have to be true for additive systems. One could imagine, in Roman numerals, writing IICVCX to mean 217, but in fact this is not allowed. There are a few exceptions in some alphabetic systems, where the notation follows the word order for the lexical number. Chrisomalis's explanations of these in terms of human cognitive capacity, such as the limited size of working memory, and the relation of numerical notations to language, are thought-provoking and no doubt an important part of the true explanation.

All in all, Chrisomalis's book is an impressive accomplishment and a valuable contribution to our understanding of the fundamentals of mathematics as a cultural activity.

[‡] It seems to me, by the way, that this is a clear counter-example to the common theory that we find mathematical regularities in nature because we impose them as a conceptual framework. If we could impose preferred mathematical regularities on nature, we wouldn't be dealing with this.

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

European Students Gather at TU Delft for Krylov Day

"Before starting to read everything on a new subject, I always try to think about it unbiased, and so I started with (probably) re-inventing the wheel." —Peter Sonneveld, speaking of the early development of IDR(s).

On February 2, the SIAM Student Chapter at TU Delft held a one-day workshop on Krylov subspace methods. The speakers, 12 PhD students in numerical linear algebra, gave overviews of their current work and its relation to Krylov subspaces. Participants came from different universities in The Netherlands and other European countries; among them were representatives of SIAM Student Chapters at Magdeburg, Manchester, and Prague.

Although Krylov methods are often associated with the iterative solution of large-scale linear systems, workshop participants described the application of Krylov subspaces in a wide variety of fields. Topics discussed included polynomial eigenvalue problems, estimation of matrix condition numbers, approximation of matrix functions, and applications in seismic wave propagation and flow control.

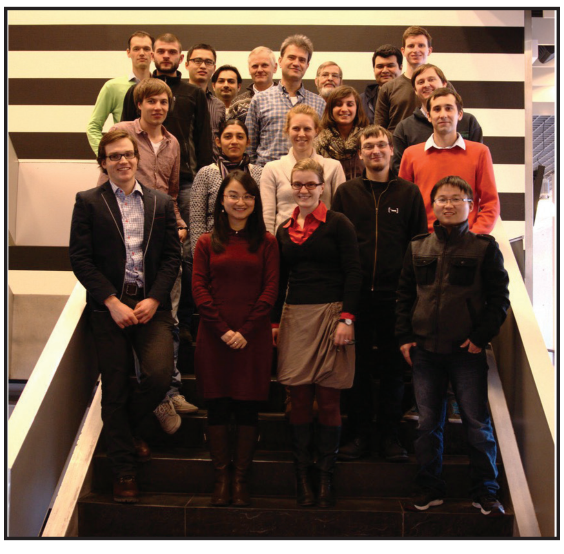

Participants in Krylov Day 2015.

As the main speaker, Peter Sonneveld of TU Delft gave a historical talk about the development of the induced dimension reduction (IDR) method, a short-recurrence Krylov method for the efficient iterative solution of linear systems with general system matrices. In collaboration with Martin van Gijzen, Sonneveld has translated theoretical work he did in the 1980s into the IDR(s) algorithm.

More information can be found at http://sscdelft.github.io/kd15 (on the Krylov Day) and http://ta.twi.tudelft.nl/nw/users/gijzen/IDR.html (on IDR(s)).

The Moody's Mega Math Challenge Marks 10th Year

By Rachel Levy

The Top Six Teams in 2015

1. *North Carolina School of Science and Mathematics* (Team 4902); Durham, North Carolina: $20,000
2. *North Carolina School of Science and Mathematics* (Team 4904); Durham, North Carolina: $15,000
3. *Elk River High School* (Team 5560); Elk River, Minnesota: $10,000
4. *Staples High School* (Team 5057); Westport, Connecticut: $7,500
5. *Maggie Walker Governor's School* (Team 4892); Richmond, Virginia: $5,000
6. *South County High School* (Team 4187); Lorton, Virginia: $2,500

Can you recall the first time you worked with a team on a significant mathematical modeling problem? For me, it was as a senior at Oberlin College, in a project for NASA in an operations research course taught by Professor Bruce Pollack-Johnson (now at Villanova). I am certain that the experience played a large role in my decision to become an applied mathematician and to join the faculty at Harvey Mudd College, which provides industrial mathematical modeling experiences through its senior capstone clinic projects. SIAM provides students with mathematical modeling experiences through the Moody's Mega Math (M^3) Challenge, which, like COMAP's Hi-MCM, makes the experience of team-based modeling available to U.S. high school students. M^3 is entirely Internet-based, and carries no entry or participation fees.

This year M^3 celebrated its 10th anniversary. The competition is organized by Michelle Montgomery's marketing and outreach team at SIAM, in collaboration with Frances Laserson, president of The Moody's Foundation. A charitable organization established by Moody's Corporation, the Foundation sponsors M^3 as part of its commitment to supporting education, in particular the study of mathematics, finance, and economics. The competition began in the New York City metropolitan area and has expanded each year; the 2016 competition will be open to students anywhere in the U.S.

"STEM education is imperative to continue the robust pipeline of talent at Moody's and elsewhere in our industry," Laserson comments. "This year, we reached more than 5,000 future applied mathematicians, economists, and computational scientists across the country via this contest, and are proud to have a part in motivating young people to study and pursue careers in these important fields."

This year, 1,128 three- to five-member teams of juniors and seniors from 45 states participated in M^3. They had only 14 hours and 20 pages to develop and communicate their solutions to this year's question: "Is college worth it?" In their math models, competitors were asked to determine the cost of earning a degree, account for the impact of President Obama's recent free two-year community college proposal, and contrast potential financial outcomes for those pursuing STEM and non-STEM degrees. Students also quantified factors that would influence a graduate's overall quality of life after school. The problem was written by SIAM member Eric Eager of the University of Wisconsin, LaCrosse.

The M^3 challenge gives students practice in cooperation and project management. They are allowed to employ any mathematical techniques they choose and to use data and other information from the web to develop their models. Communication plays a key role, both between team members and in the writing of the report. The students work in a situation familiar to many professionals in BIG (business, industry, and government) mathematics jobs: Given a new problem and a tight deadline, they must develop an insightful and useful solution.

Several former M^3 champions attended the 10th-anniversary event; among them were three of the four members of the inaugural winning team from Staples High School in Westport, Connecticut. Speaking to this year's finalists, the returning 2006 winners discussed the impact of the competition on their careers, including the benefit of internships at Moody's, the value of mathematical modeling experience as a talking point in job interviews, and the influence of the competition on their decisions to pursue careers that involve mathematical modeling. Miles Lubin is now in a PhD program in operations research at MIT, Elizabeth Marshman is in a master's program in biomedical engineering at Stanford, and Andrew Tschircart works in the U.S. Office of the Comptroller of the Currency. The fourth member, Vikas Murali, who was not able to be at the ceremony, is the founder and CEO of ActvContent.

Mark Zandi, chief economist, Moody's Analytics, and co-founder of Economy.com, gave the keynote talk at the event. Zandi spoke candidly to the young competitors, providing inspiration and advice based on his own experiences:

> "I am a forecaster and I forecast really good things for you. All those good things can be even better if you soak up as much education as you possibly can. Take some risks. Do something that makes you feel really uncomfortable. It's when you take chances that cool and interesting things happen. Surround yourself with people who complement you and stick with them."

Judging the competition is fun and rewarding. For the past three years I have served as one of the M^3 triage judges—applied math professionals, mostly SIAM members, who use an online platform to read, score, and make brief constructive comments on papers.

I enjoy seeing what high school students can do with the big, messy, real-world challenge problems, and it is a nice bonus that M^3 compensates judges for their time. This year a record 225 PhDs from academia, business, industry, and government participated in the judging; M^3 will need even more judges as western states join the competition in 2016. If you would like to join us in the fun, please contact Michelle Montgomery (montgomery@siam.org).

As SIAM VP for Education I also had the honor of serving as a finalist judge and giving a short talk at the award ceremony held in the Moody's building in New York. I was impressed by the high quality of the student presentations, the poise of the team members, and the insightful answers to our tough questions. The subtle communication between teammates as they chose who would answer a particular question gave us a glimpse of the camaraderie within the teams. In my talk I discussed the fallacy of the genius stereotype (that great mathematicians work alone, and without benefit of the ideas of others); I also shared the "Mathematician's Happy Dance" that my colleagues and I do when one of us makes a breakthrough after a lengthy struggle. Video interviewer Adam Bauser of Bauser Media Group assures me that I was the first person to dance on SIAM livestream, and you can see the winners doing the dance at the end of the 2015 overview video (http://bit.ly/1DZDlE5).

Excellent well-crafted problems are critical to the success of M^3. A good problem is one that has not been solved, that can be approached in many ways using a variety of high school-level mathematics, and that matters—both to the students and to society. A problem-development team reviews

See **M^3 Challenge** *on page 7*

M^3 Challenge

continued from page 6

Members of the first-place NCSSM team with Moody's Foundation president Frances Laserson. From left, Guy Blanc, Laserson, Sandeep Silwal, Michael An, Jenny Wang, and Evan Liang.

problems, which can be submitted by anyone. M^3 pays $150 for a problem that is accepted for the test bank and $1,000 for a problem used for the challenge. To learn more about problem submission, see the "suggest problems" page at http://m3challenge.siam.org/resources/suggest-problems, or attend an M^3 information session the next time you go to the Joint Mathematics Meetings, MathFest, or a SIAM Annual Meeting.

In addition to an archive of past problems and example solutions, the M^3 website provides guidance for teachers who would like to work with students on mathematical modeling. Resources, including a free mathematical modeling handbook with connections to the Common Core State Standards in mathematics, science, and language arts, can be downloaded, along with a set of reference cards. Although developed with coaching for M^3 in mind, the materials are used by K–12 teachers to get ideas about how to engage their students in authentic mathematical modeling activities.

This year's six finalists—from Connecticut, Minnesota, North Carolina, and Virginia—took home a combined $60,000 in scholarships (see sidebar). An additional $65,000 in scholarships was distributed among the six semi-finalists and the 53 honorable mention teams. For the 10th anniversary, the M^3 organizers produced a series of retrospective videos highlighting past competitors' successes during their college careers and on into the workplace. Readers can watch videos, including the finalist team presentations and award ceremony highlights on SIAM Connect: http://bit.ly/1GVWskM.

This year's first- and second-place teams were both from the North Carolina School of Science and Mathematics (NCSSM), in Durham, North Carolina, coached by mathematics instructor and phenomenal mathematical modeling coach Dan Teague. NCSSM, the first residential public school in the U.S., attracts juniors and seniors from across North Carolina. On the school's website, Teague discusses the value of mathematical modeling as an inherently interdisciplinary experience:

> "Mathematical modeling requires much more than mathematics. It requires knowledge of how things work, which comes from the students' experiences in the sciences, both natural and social, their programming ability, and their ability to write clearly and persuasively and explain complicated ideas in written form. We all share in these students' accomplishments, because we all contributed."

Rachel Levy, SIAM vice president for education, is an associate professor of mathematics at Harvey Mudd College.

Extreme-scale HPC

continued from page 8

The deficiencies just listed are enough to drive a multi-decade mathematical research program. But underlying the deficiencies are some great opportunities in the form of novel mathematical research directions. Here are a few of them:

■ With 10^9 parallel threads (in future extreme-scale systems) we will have to avoid all unnecessary communication and synchronization. Research is already under way in some areas, including dense linear algebra, although the problem is wide open elsewhere, e.g., for iterative solvers. New asynchronous, communication-avoiding algorithms must be designed. Lower bounds must be found on the amount of communication/synchronization necessary for a particular problem. Chaotic relaxation strategies or stochastic and nondeterministic algorithms—possibly among the key innovations needed for extreme-scale computing—could also provide greater robustness and built-in fault tolerance overall.

■ Extreme-scale systems will provide the computational power to move from qualitative simulation to predictive simulation, and from predictive simulation to optimization, parameter identification, and inverse problems; they will make stochastic simulations possible and allow us to better quantify uncertainties.

■ Extreme-scale computing will enable us to bridge the gap between the mesoscale and the human scale. "Mesoscale" refers to certain physical scales, such as a cell in a biological system, a grain of sand, or a pore in an aquifer. A living human has around 10^{11} neurons and 10^{13} red blood cells; a pile of sand may consist of 10^{10} grains. The mesoscale is halfway between the atomic and the human scales. Mesoscale computing entails dealing with large numbers of objects, but such ensembles may become tractable on extreme-scale systems—with 10^{18} flop/s we can perform $O(10^5)$ flop/s for each human blood cell.

Thus, the extreme scale may offer new possibilities for simulation science. To exploit this capability, however, we need new methods for modeling and simulating large mesoscopic ensembles for long enough times. New algorithms must be invented, new modeling paradigms devised. We also need new techniques for validation and verification: We are not interested in accurate predictions of each individual blood cell in a human being, but the ensemble behavior must be physically meaningful and must provide insight, e.g., physiological, beyond that offered by classical techniques. Such multiphysics scenarios and multiscale modeling paradigms will gain increased momentum with the advent of extreme-scale computing and will become even more interesting research topics.

In summary, I believe that the advent of extreme-scale computing is forcing mathematical scientists to address the growing performance abyss between existing mathematical theory and the practical use of HPC systems. Tweaking codes is not enough—we must look back and perform new analyses in areas in which we have not thought deeply enough, in order to develop a new methodology for interdisciplinary algorithm and performance engineering. Beyond this, extreme-scale computing opens fascinating new opportunities in fundamental research that far surpass increased mesh resolution. Opportunities for the development of asynchronous algorithms and large-scale mesoscopic modeling are just two examples.

Ulrich Rüde is a professor in the department of computer science at the University of Erlangen-Nuremberg.

You can read this issue at

sinews.siam.org

You can also access the *SIAM News* archives
(currently from November 2012 to the present).

Professional Opportunities

Send copy for classified advertisements to: Advertising Coordinator, SIAM News, *3600 Market Street, 6th Floor, Philadelphia, PA 19104–2688; (215) 382–9800; fax: (215) 386–7999; marketing@siam.org. The rate is $3.00 per word (minimum $375.00). Display advertising rates are available on request.*

Advertising copy must be received at least four weeks before publication (e.g., the deadline for the September 2015 issue is July 31, 2015).

Advertisements with application deadlines falling within the month of publication will not be accepted (e.g., an advertisement published in the September issue must show an application deadline of October 1 or later).

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

Announcements

Send copy for announcements to: Advertising Coordinator, SIAM News, *3600 Market Street, 6th Floor, Philadelphia, PA 19104–2688; (215) 382–9800; marketing@siam.org. The rate is $1.95 per word (minimum $275.00). Announcements must be received at least one month before publication (e.g., the deadline for the September 2015 issue is July 31, 2015).*

Call for Nominations for the 2016 Vasil A. Popov Prize

University of South Carolina

The Vasil A. Popov Prize is awarded every three years for outstanding research in fields related to the work of Popov, best known for his contributions to approximation theory.

Nominees must have received their PhD within the previous six years.

Nominations, which must include a brief description of the relevant work and the nominee's curriculum vitae, should be sent to Pencho Petrushev, Chair, Popov Prize Selection Committee, Interdisciplinary Mathematics Institute, University of South Carolina, Columbia, SC 29208; popov.prize@gmail.com. The deadline for nominations is November 15, 2015.

The prize will be awarded in May 2016 at the Fifteenth International Conference in Approximation Theory in San Antonio, Texas. For further information, visit http://imi.cas.sc.edu/popov-prize-call-nominations/.

Fifteenth International Conference in Approximation Theory

San Antonio, Texas, May 22–25, 2016

Organizers: Greg Fasshauer and Larry Schumaker

Invited speakers: Josef Dick (New South Wales), Simon Foucart (Georgia), Elisabeth Larsson (Uppsala), Doron Lubinsky (Georgia Tech), Carla Manni (Rome), Mike Neamtu (Vanderbilt), and Ulrich Reif (Darmstadt).

The eighth Vasil A. Popov Prize will be awarded at the meeting (for nominations visit http://imi.cas.sc.edu/popov-prize-call-nominations/).

Papers in all areas of approximation theory will be organized into contributed sessions, and the organizers invite suggestions for minisymposia.

Travel support: The organizers especially encourage students and postdocs to attend and to present their work. They hope to be able to provide some support for these groups and for members of other under-represented groups. An application form is available on the website.

Information: For details on the conference, see http://www.math.vanderbilt.edu/~AT15.

New Mathematics for Extreme-scale Computational Science?

By Ulrich Rüde

Let's start with the good news: Mathematics continues to be the most important contributor to any work in large-scale computational science. This is so because computational complexity becomes ever more important with faster computers. Once systems are large enough, the best algorithms will always be the ones with the best asymptotic complexity. High-quality journals, such as the *SIAM Journals on Numerical Analysis* and *Scientific Computing*, regularly publish papers in this area that advance the research frontier.

That said, many of these novel algorithms underperform by many orders of magnitude. Contrary to the belief of some in the mathematics community, what is needed is not simply a few extra weeks for converting Matlab to Fortran and MPI. Designing efficient HPC software requires extensive creative research. The deficiencies of currently available software, as outlined in the following list, are much more fundamental.

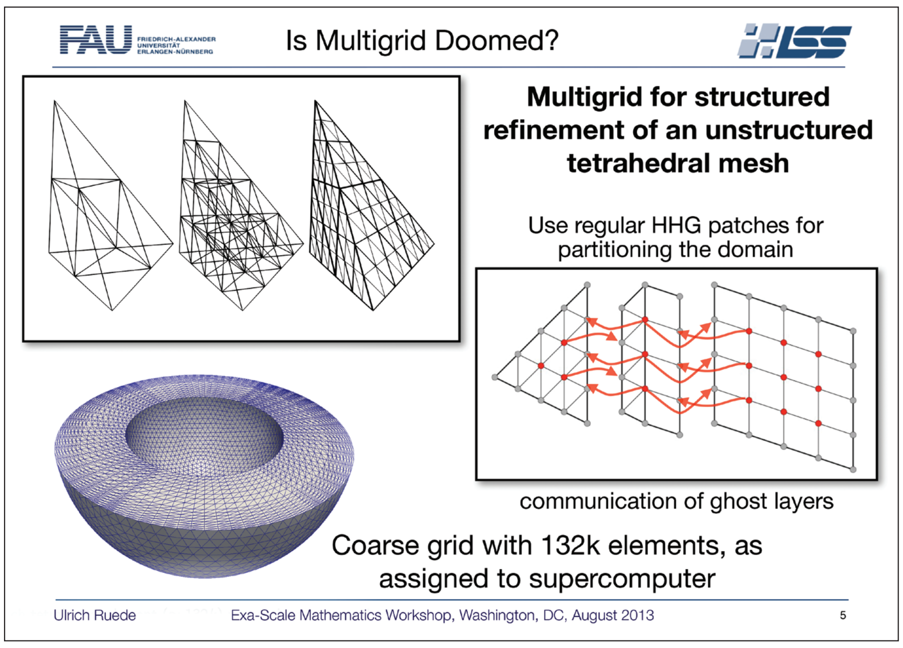

Is Multigrid Doomed?

Multigrid for structured refinement of an unstructured tetrahedral mesh

Use regular HHG patches for partitioning the domain

communication of ghost layers

Coarse grid with 132k elements, as assigned to supercomputer

Ulrich Ruede Exa-Scale Mathematics Workshop, Washington, DC, August 2013 5

From a talk given by the author at the Exa-Scale Mathematics Workshop in Washington, DC, August 2013.

■ There is nothing so practical as a good theory,[*] but a misconception about the role of rigorous theory seems to have taken root in the math community. For example, a rigorous asymptotic bound of the form $|e| \le C\, h^p$ has only heuristic implications if we are assessing the quality of a discretization for all finite values of h (that is, in any practical computation). Such theorems are a poor basis for comparing one discretization to another of the same or even different order, as long as the constants remain unspecified. We need more extensive quantitative theory. In its absence, systematic numerical experiments become as important as or even more important than rigorous theory.

■ Some areas of contemporary applied mathematics have an underdeveloped tradition in systematic algorithmic benchmarking. This starts with the lack of generally accepted standard test examples, which means that the numerical cost of an algorithm (i.e., the number of flop/s required for a specific discretization or by a specific solver) is frequently left unquantified. Consequently, relatively inefficient algorithms sometimes remain in use even when better alternatives exist.

■ On modern computer systems, the traditional cost metric of numerical mathematics (i.e., the number of flop/s needed to solve a problem) increasingly fails to correlate with truly relevant cost factors, such as time to solution and energy consumption. It will be necessary to quantify much more complex algorithmic characteristics, such as memory footprint and memory-access structure (e.g., cache re-use, uniformity of access, utilization of block-transfers), processor utilization, and communication and synchronization requirements. These effects must be built into better complexity models—models that are simple enough to be used, but that capture the true nature of computational cost far better than a simple count of flop/s.

■ For extreme-scale computational science, we need a more systematic integrated methodology that we can use to engineer algorithms. Starting from a mathematical model, we want to be able to predict a priori the performance that can be achieved; afterward, we will evaluate the actual performance with respect to the prediction, accounting for the discrepancies. This must be done on all levels of the "simulation pipeline"—the mathematical model, the discretization, the solver, its sequential implementation, and eventually its parallelization.

It is important that this be seen not simply as tweaking a given algorithm to run fast on a particular architecture, but as true co-design. In particular, the process includes the design and development of the algorithms and data structures. If it is known that, say, a 2D multigrid Poisson solver reaches h^2-discretization accuracy in fewer than 30 operations per unknown, we must be able to justify the use of a more complicated discretization and more expensive solver for the same problem class. We may have good reasons, but such algorithmic choices must be based on clear arguments that account for the accuracy achieved relative to the cost.

■ On the implementation side, even the sequential version of an algorithm often reaches only a fraction of the peak performance of a core. We should be able to explain why this is so. It may turn out that memory or communication bandwidth is the relevant bottleneck. Generally, theory must concisely quantify the performance bounds for a given computer system, and the design process must be based on a systematic accounting for the limiting resources. To achieve this, we need realistic a priori cost predictions throughout the development process. And in general, we should be more honest in assessing parallel performance. David Bailey's "Twelve Ways to Fool the Masses …"[†] are still too much in use.

See **Extreme-scale HPC** *on page 7*

---

[*] According to Kurt Lewin (1890–1947).

[†] A modernized version can be found at http://blogs.fau.de/hager/category/fooling-the-masses/. For the original, see http://www.davidhbailey.com/dhbpapers/twelve-ways.pdf.