from *SIAM News*, Volume 36, Number 10, December 2003

SIAM Conference Highlights Work in Progress on Problems of Lasting Interest to Industry

By James Case

The first-ever SIAM Conference on Mathematics for Industry, postponed because of last summer's widely publicized outbreak of SARS, was held October 13–15, in the Metropolitan Hotel in Toronto. Current, former, and prospective employees of U.S. corporations, government laboratories, independent research institutes, and the military were in attendance, along with academics interested or involved in industrial problems and research. Ongoing collaborations between teams of industrial and academic investigators were responsible for an overwhelming proportion of the findings reported at the conference, and for much of the work in progress.

Following the usual SIAM format, the Toronto conference featured a series of invited plenary lectures on subjects of general interest, followed by more specialized sessions. The invited speakers discussed topics ranging from the frontiers of scientific computing to the modeling of the global HIV/AIDS epidemic to the battlefield needs and capabilities of soldiers yet unborn.

By far the most discussed of the six invited lectures was "A General Framework for the Validation of Computer Models and Simulations," by James C. Cavendish of the General Motors Research and Development Center. In it, he distinguished between the verification and the validation of a complex model, with intent to popularize a systematic six-step Bayesian validation procedure being developed as part of a research agreement between General Motors and the National Institute for Statistical Sciences.

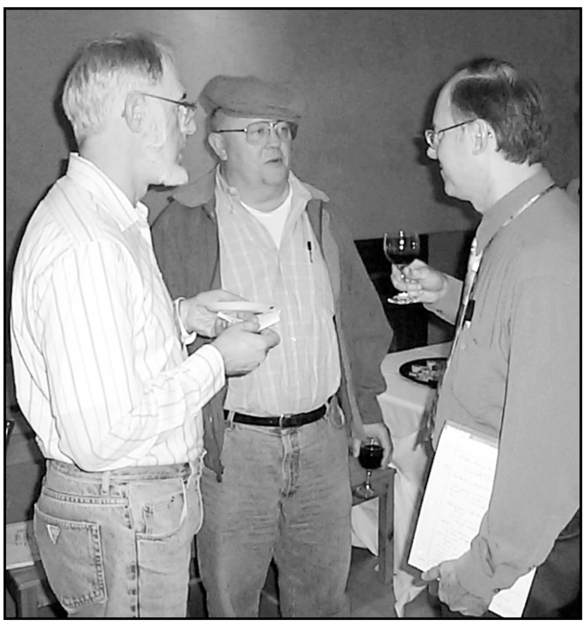

David Field of the General Motors Corporation (left) and David Ferguson of The Boeing Corporation (center) co-chaired the organizing committee for the Toronto conference. With them is Ken Davidson, director of the Fields Institute, which hosted an evening session on mathematics in the biotech industry.

To *verify* a vehicle crash-test model, one must demonstrate that the code is internally consistent and bug-free, so that the model equations are effectively solved. That is typically the cheap and easy part. The harder and more expensive part, more often than not, is to *validate* the model by demonstrating that its predictions agree—to within pre-set tolerances—with the actual measured performances of (costly) crash-test vehicles. Cavendish emphasized that the same model may be valid for one class of simulations, and invalid for another. Thus, for instance, commercially available crash-test models are well validated as they apply to low-speed impacts, in which metal parts bend but do not break, but not as they apply to high-speed impacts, in which fracture is all but certain to occur.

Cavendish illustrated the elements of the new six-step Bayesian validation procedure by applying them to the problem of predicting the number of small rigid bodies—such as BBs or M&Ms—that can be squeezed into a quart jar or other irregularly shaped glass container. The model for making such predictions simply multiplies the volume of the container by a "space utilization factor" m and is thus easily verified. Validation, as usual, is more difficult. For (spherical) BBs, m cannot exceed 74.04%, long suspected and proved in 1998 by Thomas Hales and Samuel Ferguson to be the largest fraction of 3-space that can be occupied by identical spheres. Moreover, it is not hard to exceed 67% utilization of capacity, even in comparatively small and irregularly shaped containers. Hence, validation is not difficult for BBs and other small spheres of the same size. For less symmetric solids of rotation like M&Ms, however, the task is more difficult, because there is no comparably applicable theory. Caven-dish walked the audience through all six stages of the new Bayesian validation procedure, beginning with a (loose) prior distribution on the range of foreseeable outcomes and ending with a significantly tightened posterior distribution. While his main example concerned BBs and M&Ms in little glass jars, Cavendish stressed that the same six-step Bayesian procedure applies to crash-test models, atmospheric models, groundwater models, earthquake models, and virtually any other class of models in which uncertain model inputs cloud the validity of model output.

On the evening of October 14, a number of SIAM conferees walked over to the Fields Institute for a pair of talks on mathematics in the biotechnology industry, followed by a bountiful reception. The speakers, Daniel Kolber of TM Bioscience and Elisabeth Tillier of the Ontario Cancer Institute and the University of Toronto, considered the massive data sets generated by the biotech industry. Tillier focused on the discrete sequence data obtained from large-scale genome-sequencing projects; Kolber discussed the continuous data generated by "microarrays," an essential tool in the industry. Both had more to say about the data itself than about the data-mining techniques used to gain useful insight about the data.

In the final plenary talk, John Par-mentola, director for research and laboratory management, U.S. Army, showed a video purporting to depict a small-scale military action circa 2015. While the opposition was armed with rocks, small arms, and a few cold-war-era tanks, coalition forces availed themselves of the smaller, lighter-weight—and thus more rapidly deployable to distant continents—weapons the military is now developing to meet expected future needs. Parmentola described several of what seems to be a staggering number of research and design problems that need to be solved before such weapons can be developed. One example is the positioning and protection of drone aircraft deployed to maintain communication between units advancing on, say, opposite sides of a mountain ridge.

Perhaps the most striking feature of Parmentola's presentation was the implied contrast between the military's 15-year planning horizon and that of the corporations represented at the conference, few of which seemed to be looking even five years into the future. Time after time, industrial speakers emphasized that they were seldom concerned with the optimality or long-term utility of their solutions, and stressed the need for robust "quick and dirty" prescriptions effective under existing conditions. Most professed to operate on the assumption that tomorrow's conditions will differ, in unforeseeable ways, from today's, and that anything beyond the immediate future will pretty much have to take care of itself. It is not that they have no interest in the long-term future—they simply feel that conditions in U.S. industry are changing too fast and too unpredictably to reward long-range planning.

By no means were all the gems on display in Toronto confined to the plenary sessions. In a contributed paper session on image and data analysis, Jacek Turski of the University of Houston described his efforts to facilitate the construction of an "active vision system" consisting of a moving camera head, a hardware image processor, and an image-analyzing computer. Such an apparatus would mimic the initial stage in the process by which we primates process visual information—the stage in which signals received by the eye, in the form of photons incident on the light-sensitive (rod and cone) cells of the retina—are transmitted to the visual cortex at the back of the brain. While effective electronic eyes remain far in the future, Turski seems to have taken an important step toward that goal.

Experimental evidence points to the existence of a one-to-one correspondence (retinotopy) between the light-sensitive cells of the retina and associated cells in the visual cortex. While the latter act as "pixels of the brain," and are presumably arranged in rows and columns, the former function as "pixels of the eye," observably arranged in concentric rings about a point in the middle of the retina. This suggests that, if only for the sake of simplicity, the retinotopy should associate the point (pixel) $z = re^{i\theta}$ of the retina with the point (pixel) $w = \ln z = \ln r + i\theta = u + i\theta$ of the visual cortex. Turski suggests that the eye processes "patterns" $f \in L^2(\hbar)$ by decomposing them into linear combinations of basic patterns in that space, akin to the complex exponentials of ordinary Fourier analysis, and proposes to automate the procedure by which it does so. To that end, he has developed an entire "projective Fourier analysis," based on a readily invertible "projective Fourier transform." These developments are explained in detail in a paper scheduled to appear in *SIAM Review* next year.

Speaking in a minisymposium on organizational complexity, Daniel Reaume of the GM R&D Center described a game-theoretic approach to the scheduling and selection of projects that rely on a common pool of resources. On the face of it, given that the Nash equilibrium solutions of many-player games ordinarily fall far short of (Pareto) optimality, the approach seems less than promising. But a growing body of evidence suggests that when the players' objective functions are more or less identical, the shortfall is frequently insignificant. Moreover, the Brown–Robinson technique—a.k.a. the "method of fictitious play"—appears to converge rapidly and inexpensively in such situations.

The fact that the strategies so obtained are typically of the "mixed" variety—participants are advised to take action #1 with probability p_1, action #2 with probability p_2, and so on—has traditionally limited their appeal to management. Responsible decision makers rarely welcome the suggestion that they should toss coins and roll dice. But a newly discovered technique identifies vertices of the simplex of potential mixed strategies that, when substituted for the calculated mixed-equilibrium strategies, yield identical performance without stochastic mixing.

Another broadly significant presentation was that of John Betts of The Boeing Company. Speaking in a minisymposium on optimization technologies in engineering design, he described a toy problem of the form

$$\begin{aligned}
&\text{minimize } J = \int_0^T \int_0^1 u^2(x,t)\,dx\,dt \\
&+\int_0^T \left[g^2(t) + h^2(t)\right]dt \\
&\text{subject to } u_t = u_{xx} \\
&\text{and } u(x,0) = f(x), \\
&u(0,t) = g(t), \\
&u(1,t) = h(t) \\
&\text{and } x(1-x)\,t/(1+t) \le u(x,t) \\
&\forall\ 0 \le x \le 1 \text{ and } 0 \le t \le T
\end{aligned}$$

in which *f(x)* is given, and *g(t)* and *h(t)* are to be chosen by the minimizer. Constraints of the form $\phi(x,t) \le u(x,t)$, in which ϕ is a given function while u is unknown, are known in the optimization community as state-space constraints. Although they turn up in all manner of optimal design problems, and frequently contribute to numerical instability, little exists in the way of applicable theory. Such problems are often described as "infinite-dimensional energy-minimization problems." Whatever novel challenge the problem presents is a result of the lower limit on $u(x,t)$. It can be further assumed, with but little loss of generality, that $f(x) \equiv 1$ and $g(t) \equiv h(t)$.

Betts explained that the usual attack on problems of this sort involves discretizing the space variable x, which reduces the whole to a finite-dimensional energy-minimization problem that, but for the state-space constraint on the unknown function $u(x,t)$, would be easy to solve as a linear-quadratic optimal control problem. The presence of the state constraint, however, seems to necessitate discretization of the time variable t as well, in order to exploit finite-dimensional nonlinear programming methods. And that's where the real difficulties arise. As Betts explained, the results seem to be extraordinarily sensitive to the manner in which the discretization is done, and to the nature and degree of contact between the unknown function $u(x,t)$ and the known lower limit $x(1 - x)t/(1 + t)$.

Each of the presentations mentioned here—along with many others given at the conference—described work in progress on problems of lasting concern to one or more of the industries represented. This was particularly true of the sessions devoted to supply-chain management, a contentious issue in almost every industry because of the rapid pace of globalization. This was exactly what the conference organizers hoped to achieve—a forum in which people and organizations with problems to solve would meet others with relevant skills.

SIAM intends to build on the success of the Toronto conference by making it a regular event. Perhaps, in addition to the industries represented in Toronto—auto, aircraft, computer, computer chip, and biotech—future conferences will draw participants from an expanded pool, including the energy, airline, communications, food, and entertainment industries.

James Case writes from Baltimore, Maryland.