Newsjournal of the Society for Industrial and Applied Mathematics

siam news

Volume 47/ Number 4
May 2014

How Did Inmarsat Deduce Possible Flight Paths for MH370?

By John Zweck

Malaysia Airlines Flight 370 disappeared in the early hours of March 8 en route to Beijing from Kuala Lumpur with 239 people on board. As this issue of *SIAM News* goes to press, no confirmed remnants of the plane and no survivors have been found. This great tragedy has saddened people the world over. We sympathize with the family members of the passengers, crew, and pilots of MH370.

On March 28, the search area for MH370 was refocused. The new region is in the southern Indian Ocean, about 1850 km west of Perth, Australia; it is 1100 km northeast of the area of the Indian Ocean initially searched, starting on March 17. Both search areas were identified by a team of engineers from the British satellite company Inmarsat, in collaboration with the US National Transportation Safety Board (NTSB) and Britain's Air Accident Investigation Branch (AAIB). An enormous amount of engineering expertise—much more than can be addressed in this article—has been devoted to solving the mystery. As stated in an article in *The New York Times* on Sunday, March 23,[1]

> By Sunday afternoon [March 9], a team of Inmarsat engineers set to work using the principles of trigonometry to determine the distance between the satellite and the plane at the time of each ping, and then to calculate two rough flight paths.

In this article I consider some of the mathematics used by the Inmarsat engineering team to deduce the flight path of MH370. In addition, using a simulated fictitious flight (XX123), I illustrate some of the inherent difficulties of the team's search.

Publicly Available Data

The only publicly available reliable data for MH370 are (1) the location of the INMARSAT 3-F1 satellite[2] in a geostationary orbit at an elevation of 35,800 km over the point on the earth at 1.5º N, 64.5º E; (2) the data released by AAIB on March 25,[3] which includes the times at which the satellite pinged the aircraft, together with the measured burst frequency offsets (i.e., the Doppler-shift frequencies) at those times; and (3) a map released by the Malaysian government[4] showing that the angle between the 8:11 AM arc and the satellite position was about 40º. The initial search area is accurately shown on a map posted on March 17 on the website of *The Washington Post*.[5] That map also includes the approximate location of the satellite and four red arcs marked with the times 5:11 AM, 6:11 AM, 7:11 AM, and 8:11 AM.

See **Flight Path** *on page 8*

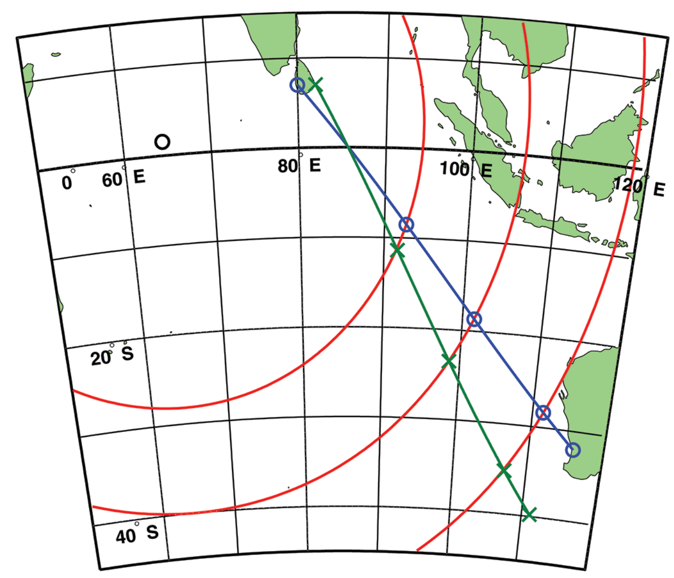

Figure 1. *Map of the path of* fictitious flight XX123. *For a departure from Colombo, Sri Lanka, at 12 noon and arrival in Perth, Australia, at 7:48 PM, the three red arcs show the location of the plane at 3 PM, 5 PM, and 7 PM. Based only on this information, the blue curve shows the path of the plane traveling at 400 knots deduced by the methods described in this article. The green curve shows how the deduced flight path changes when the starting location is assumed to be 220 km due east of Colombo. A plane traveling along the green path would have a speed of 416 knots and would be 970 km from Perth at 7:48 PM.*

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

Parallel Processing 2014

Heightened Interest in Fault-tolerant Algorithms As Exascale Era Draws Near

By Mark Hoemmen

For the 2014 SIAM Conference on Parallel Processing for Scientific Computing (Portland, Oregon, February 18–21), Keita Teranishi, Jaideep Ray, Mike Heroux, and I (all of Sandia National Laboratories) organized a four-part (17-speaker) minisymposium on fault-tolerant numerical algorithms. This was the third in an annual sequence of minisymposia on the topic offered at SIAM conferences, alternating between CS&E and Parallel Processing.

For our purposes, a "fault" occurs when a computer system behaves incorrectly in a way that *could* affect the running application's correctness or ability to make progress. The application "fails" if it actually produces incorrect results or fails to complete. We distinguish further between "hard" and "soft" faults. Hard faults cause the application to end prematurely or to wait forever. Soft faults let the application keep running (for a while, at least) but corrupt intermediate states; this *could* make the application return an incorrect result. Speakers in our minisymposium considered faults of both types.

Algorithms studied in our sessions include iterative sparse linear solvers, nonlinear multigrid, finite-difference stencils, time integration, dense matrix factorizations, fast Fourier transforms, and multilevel Monte Carlo. Talks also covered fault injection (both by software simulation and by an actual neutron beam directed at real hardware), programming models capable of identifying code and data that can go awry, and current trends in system reliability.

The growing interest of the parallel algorithms community in this topic is reflected in the breadth of our minisymposium, as well as in the Wednesday-evening panel discussion (Resilience at Exascale: Should it Worry Application Developers?) and related talks in other minisymposia. This interest has waxed and waned over the years, leading many to question whether we should expend so much effort on the problem. In particular, talk of application fault tolerance seems to return with each approach of parallel "scale" transitions—as in terascale, petascale, and the upcoming exascale. At previous scale transitions, systems stepped up to protect applications, with no need for contributions from the latter. However, experience has established that providing this protection incurs hardware costs (a high-bandwidth global file system) and requires significant "system middleware" effort to make current resilience mechanisms (mainly checkpoint-restart) efficient at scale.

As the exascale era approaches, the confluence of several factors suggests that algorithm developers need to get involved soon. First, the current processor-building technology, CMOS (complementary metal-oxide semiconductor), is close to its physical size limits. As CMOS transistor gates shrink to the thickness of a few atoms, it gets harder and harder to make them consistently correct. Second, for about ten years, commodity processors have been achieving performance improvements not through frequency increases but through parallelism. More parallelism means more components, which makes faults more likely. Third, very large parallel computers are now limited mainly by power. Megawatts of electricity have to get into the facility, and megawatts of heat have to leave it. Earlier trends, however, suggest that hardware will only become less reliable. Counteracting this by

See **Fault Tolerance** *on page 7*

William Gropp (right) of the University of Illinois at Urbana–Champaign received the SIAM Activity Group on Supercomputing Career Prize at the 2014 SIAM Conference on Parallel Processing for Scientific Computing. Shown here with organizing committee co-chair and SIAG/SC chair Ali Pinar, Gropp was honored "for his excellence in research, contributions to numerical software for linear algebra, high-performance parallel and distributed computation, and exceptional service to the scientific community." In a prize talk titled "A Tale of Two Timelines," he used two projects from his career—the parallel numerical library PETsc and MPICH, an implementation of the Message-Passing Interface—to illustrate revolution and evolution in computer science.
Established in 2009, the SIAG/SC Career Prize is awarded every two years to an outstanding senior researcher who has made broad and distinguished contributions to the field of algorithms research and development for parallel scientific and engineering computing.

siam news
Volume 47/ Number 4/ May 2014

NIBIB Supports Modeling, Analysis, and Simulation in Clinical Medicine and Biomedical Research

By Margot Kern and Grace C.Y. Peng

When determining whether to open a blocked artery with a stent, cardiologists calculate the fractional flow reserve (FFR), a number that predicts the likelihood that a blockage is significantly impeding blood flow to the heart. Current methods for calculating FFR involve threading a catheter from the arteries in a patient's groin or wrist all the way into the coronary arteries. It's an invasive test, and one that costs several thousand dollars.

Recently, a team of surgeons and biomedical engineers has created an alternative method for determining FFR that doesn't require catheterization. The method—called HeartFlow—uses anatomical data from a CT scan of a patient's heart to create a realistic 3D multiscale computational model of the patient's coronary arteries. Virtual blood, obeying the physical laws of fluid dynamics, runs through the arteries; the FFR can be calculated noninvasively at all points along the arteries. Three recent clinical trials indicate that HeartFlow shows promise for impacting current state-of-the-art diagnostic methods [1–3]. In the future, doctors could use the model to try out virtual stents, determining optimal size and location for stents to be placed in a patient.

HeartFlow is just one example of how computational modeling can help solve complex problems in clinical medicine and biomedical research. Biomedical modeling provides a powerful platform for improving the efficiency of scientific experimentation, data acquisition, knowledge transfer, and technology development. Furthermore, models can be used to test the safety and efficacy of medical devices and drugs in silico, reducing the dependency on human and animal studies and accelerating translation from basic science to clinical medicine.

Support for Biomedical Modeling, Simulation, and Analysis Research

NIBIB—the National Institute of Biomedical Imaging and Bioengineering—is one of 27 institutes and centers that make up the US National Institutes of Health, the federal government's premier biomedical research agency. The mission of NIBIB is to support the development of novel biomedical technologies to help diagnose, treat, and prevent illness.

Among many other focus areas, NIBIB provides support for quantitative research projects that contribute to its mission. These projects generally fall into three categories:

1. Modeling: biomedical models and computational algorithms with potential clinical or biomedical applications. Research in this area includes mathematical, statistical, transport, network, population, mechanical, electrical, and electronic models applied to a broad range of biomedical fields. Particular emphasis is placed on multiscale modeling and the methods used to bridge scales.

2. Analysis: mathematical, statistical, and signal processing methods for the analysis of complex biomedical systems, clinical diagnosis, and patient monitoring.

3. Simulation: technologies for training and education in clinical practice and biomedical research, and simulation methods for emulating work flow, and understanding and predicting health and disease.

Interagency Modeling and Analysis Group and the Multiscale Modeling Consortium

NIBIB coordinates the Interagency Modeling and Analysis Group, known as IMAG, which is made up of more than 80 program staff from ten government agencies who are involved in managing research programs in biomedical, biological, and behavioral systems; the programs have in common a need for novel modeling and analysis methods. From its inception in 2003, the group has provided an open forum for sharing updates on individual programs from the various IMAG agencies and for planning transagency activities that will have a broad impact on the communities served by IMAG.

Early in its existence, IMAG recognized that the modeling community was at the forefront of efforts to develop computational methods across the biological continuum. In addition, IMAG identified a strong desire among modelers to form multidisciplinary partnerships across varied research communities. This led to a 2004 solicitation: Interagency Opportunities in Multiscale Modeling in Biomedical, Biological, and Behavioral Systems. In 2006, the 24 awardees from the solicitation formed the Multiscale Modeling Consortium, which has grown to include more than 100 projects relevant to multiscale modeling.

Funding announcements, meetings, workshops, tutorials, and other information relevant to the multiscale modeling community can be found on the IMAG Wiki page: www.imagwiki.org/mediawiki.

How Can the SIAM Community Help Promote Biomedical Modeling, Analysis, and Simulation?

Although computational models are increasingly recognized as an important platform for understanding complex biological systems and facilitating technology development, some in the biomedical community continue to question their usefulness. A 2012 article in *Biomedical Computation Review* titled "Meet the Skeptics" explored the hesitancy of biologists and physicians to incorporate biomedical models into their work [4]. The article suggested that the issue stems in part from a combination of biologists who are intimidated by complex math and modelers who are dismissive of those not quantitatively trained. In addition, critics argue that modelers often don't fully understand the biological conditions they are trying to model and, as a result, base their models on inaccurate assumptions. Some biomedical researchers, while acknowledging the promise of computational modeling, either believe that their biological problems are too complex for a model to be helpful or assume that they don't have enough data to build an accurate model. Finally, a lack of well-known examples of biomedical models that have had a major impact on human health has led some to question their ultimate usefulness.

Members of the SIAM community can help assuage these concerns and become champions of computational modeling, analysis, and simulation in the following ways:

■ Be prepared to describe and explain your models and analytical methods as simply as possible, working to make them accessible to potential users.

See **NIBIB** *on page 7*

Funding Opportunities in Biomedicine

NIBIB encourages the SIAM community to apply for funding of work in which quantitative methods are used to further biomedical research. Initiatives (from NIH and other US agencies) that seek novel biomedical modeling, simulation, and analysis techniques include:

■ **Multiscale Modeling Initiative.** This initiative sponsored by the Interagency Modeling and Analysis Group (IMAG) seeks new predictive computational models that encompass multiple biological and behavioral scales to accelerate biological, biomedical, behavioral, environmental, and clinical research. More info at http://www.nibib.nih.gov/research/featured-programs/interagency-modeling-and-analysis-group-imag.

■ **BRAIN Initiative.** The goals of the NIH Brain Research through Advancing Innovative Neurotechnologies Initiative are to: (1) map the circuits of the brain; (2) measure the fluctuating patterns of electrical and chemical activity flowing within those circuits; and (3) understand how their interplay creates our unique cognitive and behavioral capabilities. More info at http://www.nih.gov/science/brain.

■ **BD2K.** The NIH Big Data to Knowledge initiative aims to develop the new approaches, standards, methods, tools, software, and competencies that will enhance the use of biomedical Big Data by supporting research, implementation, and training in data science and other relevant fields. More info at http://www.bd2k.nih.gov. Proposals for research in the development of software and methods for biomedical big data in targeted, high-need areas (http://grants.nih.gov/grants/guide/rfa-files/RFA-HG-14-020.html) are due June 19, 2014.

■ **Biomedical Information Science and Technology Initiative.** BISTI, a consortium of representatives from each of the NIH institutes and centers, serves as the focus of biomedical computing issues at NIH. The mission of BISTI is to make optimal use of computer science and technology to address problems in biology and medicine by fostering new basic understandings, collaborations, and transdisciplinary initiatives between the computational and biomedical sciences. More info at http://www.bisti.nih.gov.

■ **National Robotics Initiative.** The interagency NRI seeks the development of intelligent robots that can respond to users' needs and to changing environments, with the goal of achieving functional independence in humans. Future goals include robots that can help reduce human errors in healthcare settings, monitor symptoms, and dispense drugs. More info at http://www.nsf.gov/funding/pgm_summ.jsp?pims_id=503641 and http://robotics-vo.us/.

siam news

ISSN 1557–9573.

Assyr Abdulle Receives 2013 SIAM Dahlquist Prize

Assyr Abdulle of the Ecole Polytechnique Fédéral de Lausanne received SIAM's Dahlquist Prize at the 2013 SciCADE Conference. The selection committee cited his "contributions to the numerical analysis of stiff ordinary differential equations, to multiscale methods for partial and stochastic differential equations and to the implementation of numerical algorithms to problems in chemistry and biology." The work of Abdulle, a professor (Professeur Ordinaire) of mathematics at EPFL, was also recognized by SIAM in 2009, when he received the James H. Wilkinson Prize in Numerical Analysis and Scientific Computing.

Abdulle spent his formative years in Geneva, where he obtained a professional degree in violin, an instrument he still plays regularly. In 2001 he received a PhD in mathematics from the Université de Genève under the supervision of Gerhard Wanner and Ernst Hairer. He has held appointments in Princeton, Zurich (ETH), Basel, and Edinburgh.

The Dahlquist Prize is awarded every two years for contributions to fields associated with Germund Dahlquist. Presentation of the prize can take place at a SciCADE Conference, at ICIAM, or at a SIAM Annual Meeting. Chus Sanz-Serna, chair of the SciCADE scientific committee and recipient of the first Dahlquist Prize (in 1995), for his work in symplectic integration, presented the certificate and cash prize to Abdulle in Valladolid, Spain, where SciCADE 2013 was held, September 16–20.

In a plenary prize talk titled "From Darcy to Wave Equations," Abdulle gave a broad overview of his current work on multiscale numerical methods. These methods work simultaneously at different temporal and spatial scales; the aim is to account in a realistic way for the macroscopic effects of, say, molecular phenomena without having to resort to a full-fledged numerical simulation of the complex microscales involved. As the initial motivating example in his lecture, Abdulle chose flow in a porous medium, as studied by Henry Darcy in his spectacularly successful work on the water supply of the city of Dijon, France, in the mid-19th century.

SciCADE 2013, with a program devoted to scientific computing and the numerical solution of differential equations, attracted more than 250 researchers from about 30 countries, the largest contingent coming from Germany. The meeting was also the forum for the awarding of the John Butcher Prize of the New Zealand branch of ANZIAM (Australia and New Zealand Industrial and Applied Mathematics). Valladolid (41.6 N, 4.7 W) is almost exactly antipodal to Wellington (41.3 S, 147.8 E), and is thus the most distant venue the New Zealanders could have chosen. The Butcher Prize, which recognizes the best presentation by a PhD student, went to Yuto Miyatake of Tokyo. Also recognized at the conference was Ludwig Gauckler (Technical University of Berlin), who received the SciCADE New Talent Award.

The conference was organized by the Department of Applied Mathematics at the University of Valladolid to coincide with the celebration of Sanz-Serna's 60th birthday. Previous SciCADE meetings were held in Beijing (2009) and Toronto (2011); the next is scheduled for 2015 in Potsdam, Germany.—*Mari Paz Calvo, Universidad de Valladolid.*

Assyr Abdulle (left) received SIAM's 2013 Dahlquist Prize at SciCADE, in Valladolid, Spain. Awarding the prize was its first-ever recipient (1995), Chus Sanz-Serna, whose 60th birthday was also celebrated at the conference.

USGS Seeks Research Proposals for Earthquake Hazards

The US Geological Survey will award up to $5 million in grants for earthquake hazards research in 2015.

Interested researchers can apply online at grants.gov, under funding opportunity number G14AS00036; detailed information is available at http://www.grants.gov/web/grants/view-opportunity.html?oppId=252944. Proposals are due May 22, 2014.

The USGS awards research grants annually to universities, state geological surveys, and private institutions. Topics of previous awards include investigation of the Central Virginia Seismic Zone to develop a better understanding of this active seismic area; examination of the paleoseismic record in the area of Alaska's Prince William Sound to characterize earthquakes prior to the Great Alaska Earthquake of 1964 and to better understand future earthquakes in this hazard-prone region; and use of GPS to measure ground deformation in the greater Las Vegas area and to provide information on how faults rupture in large, damaging earthquakes.

A complete list of funded projects and reports can be found on the USGS Earthquake Hazards Program external research support website (http://earthquake.usgs.gov/research/external/research.php).

Highlights of a Career Grounded in Major Disasters

Six Sources of Collapse: A Mathematician's Perspective on How Things Can Fall Apart in the Blink of an Eye. *By Charles R. Hadlock, Mathematical Association of America, Washington, DC, 2012, 221 pages, $50.00.*

BOOK REVIEW
By James Case

Charles Hadlock is a mathematician who spent much of his career at the consulting firm Arthur D. Little. His assignments included post-event analyses of several front-page disasters, including those at Bhopal, Three Mile Island, and Prince William Sound. The expertise developed in such projects brought the firm, originally headquartered in Boston and now a multinational, an increasing flow of commissions from corporations—many of them among the Fortune 500—seeking to identify potential disasters lurking alongside their projected paths of expansion. The teams assigned to conduct such studies operated at the highest corporate levels, often reporting directly to the firms' boards of directors, rather than to management. In his book and in a related public lecture titled "Sustainability or Collapse? An Exploration of Key Dynamics That May Determine Our Future" that he gave at the MAA Carriage House in May 2013, Hadlock described memorable moments in his career, and a few of the methods he found useful.

The book is organized around two tables, one from the first chapter, the other from the eighth and final chapter. The former lists and briefly describes nine categories of entities vulnerable to collapse; the latter names and (with equal brevity) describes the titular "six sources of collapse." The more technical chapters between present tales of collapse brought on by one or more of the six sources; each chapter is devoted to a particular source of collapse and—because Hadlock intends the book to be read by non-mathematicians—to the mathematics needed for the analysis thereof. Each chapter $I \in \{2,\ldots,7\}$ fills (at least partially) row $I - 1$ of his imaginary matrix with stories explaining, for various $J \in \{1,\ldots,9\}$, how a source of type $I - 1$ can contribute to the collapse of an entity of type J. No attempt is made to fill in the entire matrix, or even to convince the reader that it could be done.

■■■

One chapter examines collapses resulting from apparently unlikely events, such as the June 2011 power outage in Montana caused by a deer carcass draped across a power line some thirty feet above ground level. How, Hadlock wonders, is one to compute the odds against that? Because he is convinced that most people—horseracing handicappers and insurance writers naturally excepted—are inept estimators and interpreters of probabilities, the chapter begins with a primer on statistics stressing extreme-value statistics. In it he identifies an extreme-value index (EVI) γ that, when positive, indicates a complementary cumulative distribution function (CCDF = 1 – CDF) asymptotic to x^{-r} as $x \to \infty$. Thus, when $\gamma > 0$, it is customary to say that the associated probability density function has heavy tails, meaning that it imputes far greater likelihood to large deviations from the median—there might not be a mean—than would any garden variety distribution, such as the normal, logistic, or exponential. Examples include the Cauchy distribution ($\gamma = 1$), the Fréchet distribution ($\gamma = 1$), the Pareto distribution ($\gamma = 1/r$), and the Student's t distribution ($\gamma = 1/n$).

A team from Little that included Hadlock was called on to estimate the probability of an admittedly unlikely event as part of an evaluation of a proposal for burying radioactive wastes in the numerous salt domes that underlie much of the Gulf Coast region, both on shore and off. The likelihood that groundwater might leach the wastes away from the domes into more sensitive areas seemed minimal—the salt was still there millions of years after deposition, whereas it would surely be gone if the domes had not remained dry throughout the intervening millennia. Yet the Little team worried that drilling for petroleum or other minerals might accidentally puncture a storage site, letting water in and (eventually) radioactive waste out. Advocates of the plan were adamant that the probability of such an accident was negligible, as the sites would be clearly marked on every official map and drillers, in any case, are well regulated. Their claims were corroborated, moreover, by operators of several salt mines located within such domes. In short, the probability of such an accident was widely agreed to be effectively zero! Nada!

Informed opinion did not change until November 20, 1980, when a Texaco rig drilled a hole, 14 inches in diameter, directly into an operating salt mine at Jefferson Island, Louisiana. The entire content of the shallow lake in which the rig was standing—including the rig itself and several nearby barges—disappeared forever through the rapidly expanding hole into the mine below. Apparently, the location of the salt mine had been miscalculated—a fact that made it into the Little team's final report. Amazingly, nobody was killed in the accident, as it played out slowly enough for everyone on the rig, and in several nearby fishing boats, to scramble to safety. In this case, the miscalculation of a small probability was costly in money and machines, but not in lives. It might easily have been otherwise.

Miscalculated probabilities also played a role in the demise of the hedge fund Long Term Capital Management, which almost precipitated a worldwide financial crisis in 1998, and in the crash of 2008 following the collapse of the US housing industry. Hadlock explains that LTCM, like many in the finance industry, was performing risk estimation in what is known as a Black–Scholes environment, wherein all random processes are presumed to be Wiener processes $w(t)$, $0 \le t < \infty$, driven by Gaussian white noise. An important assumption concerning such processes is that distinct increments $\Delta w(t) = w(t+h) - w(t)$ and $\Delta w(s) = w(s+h') - w(s)$ are stochastically independent as long as the intervals $[t,t+h]$ and $[s,s+h']$ do not overlap. When that assumption is violated, the processes of interest cease to be Wiener processes and risk estimates obtained by assuming otherwise lose whatever validity they may once have had.

For many failed financial institutions in 2007–08, the outcome had been made to seem, by inappropriate risk assessment procedures, no more than a vanishingly remote possibility.

What happened to LTCM in 1998, and to many financial institutions in 2007–08, was that successive increments in the values of certain assets suddenly became preponderantly negative, so that accounts expected to fluctuate in value around roughly stationary means began to drift steadily southward. And because those accounts tended to be highly leveraged—meaning that they had been purchased mainly with borrowed money—a host of nominal owners found themselves in bankruptcy, owing far more than their total net worth. In each case, the

See **Disasters** *on page 5*

New Directions: A Big Deal for SIAM

TALK OF THE SOCIETY

By James Crowley

The organizers of SIAM's Modeling Across the Curriculum project, under the leadership of VP for education Peter Turner, have conducted two workshops. The first resulted in a report (http://www.siam.org/reports/modeling_12.pdf) that focused on the development of math modeling-based high school courses in STEM subjects and on a larger role for math modeling in undergraduate curricula.

The second workshop, designed to develop some of the recommendations that arose in the first, was held in January of this year. In a keynote address, Joan Ferrini-Mundy, the National Science Foundation's assistant director for Education and Human Resources, questioned the 50 workshop participants about the growing field of big data, and specifically about the implications for education.

Big Data—What Is It?

Big Data has become a Big Deal. In industry, people are hired as data analysts or data scientists. In response, universities are creating programs to train students for jobs in these fields. And, of course, agencies, among them NSF, have created major research programs in big data (including the ongoing Big Data Initiative launched by the Obama administration in 2012; see NSF update on page 8). Some would claim that massive data sets and the information that can be mined from them will supplant traditional modeling and simulation as a means of discovery. Although we doubt that this will happen, big data is clearly emerging as a field of study. The tools created in that research will have an impact on the work of scientists and engineers; already, we are seeing an impact on business applications.

As is often the case with emerging fields, big data is not clearly defined. What do we mean by "data-enabled science" or "data science"? What courses are required for a degree in this field?

More particularly, what does mathematics have to say about data science? Obviously, statistics is a longstanding player here, as is computer science, with its supporting areas of machine learning, databases, and so on. What should a data science program look like? What topics should be covered? Should someone trained in data science know, say, linear algebra?

A measure of the growing importance of this area is the recent creation of several data science centers. Three prominent new ones come to mind—one supported by the Sloan Foundation, another by the Simons Foundation, and a third at UC Berkeley. With time, answers to the "What is it?" question will emerge as research is focused around this theme. Centers will define research topics; education programs will be established to serve the demand for people trained in this area. We can look at other emerging research areas for clues.

UQ

Another area that seems to invite the "What is it?" question is uncertainty quantification. I am reminded of the comment made at SIAM's first conference on UQ (2012) by a prominent senior SIAM member who teaches probability and statistics: "UQ—I thought that's what I had been doing for the past 30 years!"

The field as defined by the SIAM Activity Group on Uncertainty Quantification focuses on quantifying uncertainty in output from large computational simulations, typically governed by partial differential equations, given uncertainties in parameters, data, and even the model itself. Techniques like MCMC (Markov Chain Monte Carlo) play a major role in this area.

The growth of this area of research has been amazing. Fifteen years ago, John Guckenheimer led a study that produced a SIAM report on uncertainty quantification, which led in turn to a small SIAM workshop on the topic. These efforts attracted sufficient attention that the *Chronicle of Higher Education* ran an article, "To Improve Their Models, Mathematicians Seek a 'Science of Uncertainty' " (April 16, 1999), featuring Guckenheimer and Jim Glimm.

Nevertheless, UQ occupied a small corner of SIAM conferences—typically as a theme in the CS&E conferences—until 2011, when SIAG/UQ was created. The group's second conference, held at the beginning of April, shows how the research area has grown. The 560 participants came from a diverse set of disciplines. The program featured well-known speakers from applied mathematics, such as Andrew Stuart, and statistics, James Berger being one example, as well as various application disciplines. Geosciences was a featured application area, represented on the organizing committee by Marcia McNutt, who was director of the United States Geological Survey when the conference was organized and is currently editor-in-chief of *Science*.

We expect that more detailed answers to the "What is it?" question will be forthcoming. At the conference, SIAG/UQ chair Max Gunzburger raised the issue with regard to *SIAM/ASA Journal on Uncertainty Quantification*. The editorial board can be counted on to have a good understanding of what constitutes UQ, but questions remain about the criteria for articles appropriate for the journal, especially with regard to supporting areas. When, for example, is a paper on numerical methods for stochastic differential equations appropriate for the journal? We expect to hear more on this issue, and more about sessions at this lively conference, in a future issue of *SIAM News*.

Back to Math Modeling and Education

Education—beginning at K–12 and continuing through postsecondary levels—has become an increasing emphasis for the math sciences community in recent years. This is due, in part, to the interest of the Obama administration in STEM education, with its stated goal of training more people in the STEM disciplines.

For SIAM, the focus has been on aspects of education clearly related to mathematics and its applications, especially modeling. Along with the Modeling Across the Curriculum workshops mentioned at the beginning of this article, SIAM is the publisher of a new guidebook on modeling (downloadable at https://m3challenge.siam.org/about/mm/ and soon to appear in print). Undertaken to help prospective mentors of high school teams competing in the SIAM-run Moody's Mega Math Challenge, the guidebook is a resource for anyone seeking a background in mathematical modeling. In the words of authors Karen Bliss, Katie Fowler, and Ben Galluzzo, the guidebook is intended for students and teachers who want "to learn how to model." It is designed to "demystify the process of how a mathematical model can be built."

Undergraduate education (at least in the early years) and K–12 education are relatively new topics for SIAM. Driven by calls for the inclusion of applications and modeling in math curricula, and spurred by the activities of the energetic SIAM Education Committee, such topics are likely to see continuing emphasis in upcoming years.

Disasters

continued from page 4

outcome had been made to seem, by inappropriate risk assessment procedures, no more than a vanishingly remote possibility.

■ ■ ■

The remaining five "source" chapters follow a similar pattern: a motivational paragraph or two, followed by a thumbnail sketch of some useful branch of mathematics—statistics, the Monte Carlo method, game theory, differential equations, catastrophe theory, network analysis—followed by a series of war stories explaining how a particular source could cause or at least contribute to the collapse of some important entity, be it a bridge, a communications network, or a global financial system.

It would be interesting to know how Hadlock's analysis applies to the most worrisome collapse of all, the one of which even chronic doom sayers hesitate to speak, namely the collapse of modern civilization predicted by the 1972 Club of Rome report, "The Limits to Growth." That oft vilified document identified three possible paths of global development, corresponding to three general classes of policy alternative: effective political action, misguided political action, and continued political inaction. Forty years later, the predictions corresponding to the third alternative are right on track as to population growth, food production, and several other fundamental quantities. They imply, among other things, that global food production will peak around the year 2030, and that (amid untold pain and suffering) global population will do likewise a few years later. Could clear thinking of the sort Hadlock describes prevent such an outcome?

Hadlock's writing is invariably clear and concise, yet informal and appealing, as befits a mathematical memoirist. The book is by no means a textbook, although it could be regarded as a compendium of modules suitable for presentation to undergraduates at various stages of development by instructors seeking to enhance their lectures with a smattering of genuine applications. For that reason alone, every teaching mathematician should own a copy of this book!

James Case writes from Baltimore, Maryland.

Textbooks, Videos, MOOCs . . . Where's a Student to Turn?

*With his new textbook—*Differential Equations and Linear Algebra*—to be published in May, Gilbert Strang of MIT recently visited SIAM headquarters to plan for distribution of the book. (All of Strang's books with Wellesley-Cambridge Press are available from SIAM.) An informal interview about the new book with* SIAM News *soon evolved into a broader conversation about teaching and learning basic mathematics today.*

To begin, a question about the new book. Why have you combined those two big subjects?

This is first of all a full text on differential equations. That course has a well-established structure, but many of the textbooks are a generation old. The linear algebra part is often limited to matrix operations. That misses the transformation that has moved systems to a central place in our work. One spring or one loop or one equation—that is a start, but it's not enough. You need coupling and interaction and a matrix.

Do you see a desire for more linear algebra in this basic course? Where do the subjects connect?

A big connection is through eigenvalues and eigenvectors. Networks are everywhere, and students have to understand matrices—this is the natural format and the right language for understanding systems and data.

Many departments (including engineering) want students to know more linear algebra. But the curriculum is often quite full. A textbook can make the connections between matrix operations and linear differential equations. To start, their solutions have the same form: complete solution = particular solution plus all null solutions. Without help, most students won't see it.

Do students still rely on textbooks? So much is available online. How do they actually learn a new subject?

Students should answer that question themselves! I will ask the class.

[On returning to MIT, Strang created a new gmail account and posed that question to his linear algebra students.]

The replies were fascinating—I should have done this before. Almost all the students report that they read the text for a first picture of a new topic (like orthogonal matrices). The problem sets bring out "details" that they missed in the reading. I wouldn't always call them details! But their other courses make demands on their time, and engineering courses really crowd them. They try their best to be efficient.

Back when you were SIAM president, you mentioned in SIAM News *that your class was being videotaped for MIT's OpenCourseWare. Do students still watch those tapes?*

Students do go to ocw.mit.edu (or YouTube), and they often skip actual class lectures. New tapes are being made this semester in linear algebra (18.06). For differential equations (18.03), Arthur Mattuck's lectures are a tremendous resource on OCW. He has taught more students than anyone in the history of MIT! The upgrades at 18.06SC and 18.03SC added problem solving and more—closer to a MOOC, but not constrained by a fixed schedule.

Videotapes can be extremely useful. The 18.06 tapes have had four million viewers; it was the right idea to make it freely available, to give it away. Videotapes are not difficult to make! I will add selected topics in differential equations, like delta functions and impulse responses—this is the ODE analog of an inverse matrix. These are such fundamental ideas that it amazes me how often they are missed.

A full-scale MOOC would be much more demanding, with careful assessment every minute. I love teaching but I hate grading.

Overall, does the web help you teach?

Absolutely! For the new book, solutions will be posted on math.mit.edu/dela along with sample sections and codes for Euler and Runge–Kutta. The time for printing solutions in the back of the book is over! The beauty of a website is that it grows and improves. Everyone can contribute. Ideas and requests are already coming to diffeqla@gmail.com.

Books are still exciting. They develop in new ways. So does our research, and our teaching, and so does SIAM. Probably the SIAM book department will turn this into my first e-book.

Gilbert Strang's new textbook may well become his first (SIAM) e-book.

siam e-books

Announcing the 2014 Class of SIAM Fellows

siam Fellow

SIAM is pleased to announce the newly selected Class of SIAM Fellows, a group of distinguished members of SIAM who were nominated by their peers for exceptional contributions to the fields of applied mathematics and computational science. Please join us in congratulating these 32 members of our community.

Mark Ainsworth
Brown University

John S. Baras
University of Maryland, College Park

Lorenz T. Biegler
Carnegie Mellon University

Åke Björk
Linköping University, Emeritus

Alfred M. Bruckstein
Technion – Israel Institute of Technology

Suncica Canic
University of Houston

Inderjit S. Dhillon
The University of Texas at Austin

Vladimir L. Druskin
Schlumberger-Doll Research

Leah Edelstein-Keshet
University of British Columbia

Donald Estep
Colorado State University

Bengt Fornberg
University of Colorado Boulder

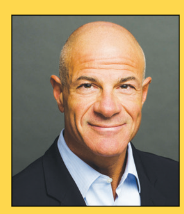

Omar Ghattas
The University of Texas at Austin

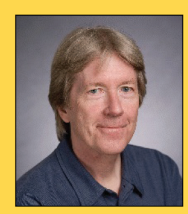

Philip E. Gill
University of California, San Diego

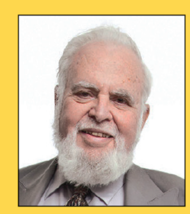

Solomon W. Golomb
University of Southern California

Jan S. Hesthaven
École Polytechnique Fédérale de Lausanne

Dorit S. Hochbaum
University of California, Berkeley

Masakazu Kojima
Tokyo Institute of Technology and JST Crest

Jeffrey C. Lagarias
University of Michigan

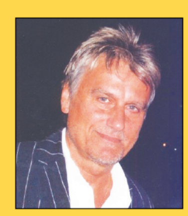

Jean B. Lasserre
Centre National de la Recherche Scientifique and Institute of Mathematics, University of Toulouse

Tai-Ping Liu
Academia Sinica

Mitchell B. Luskin
University of Minnesota

Nancy K. Nichols
University of Reading

Peter J. Olver
University of Minnesota

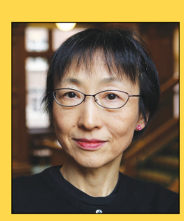

Yuriko Yamamuro Renardy
Virginia Polytechnic Institute and State University

L. Ridgway Scott
University of Chicago

Mikhail Shashkov
Los Alamos National Laboratory, X-Computational Physics Division

Christine A. Shoemaker
Cornell University

Valeria Simoncini
Università di Bologna

Zdeněk Strakoš
Charles University in Prague

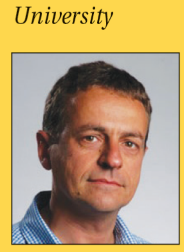

Bernd Sturmfels
University of California, Berkeley

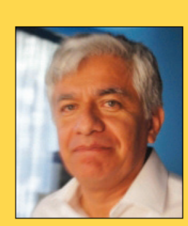

Jorge X. Velasco-Hernandez
Instituto de Matemáticas UNAM

Michael S. Vogelius
Rutgers, The State University of New Jersey

siam® SOCIETY for INDUSTRIAL and APPLIED MATHEMATICS

www.siam.org/prizes/fellows

Nominate a colleague for the 2015 Class of Fellows. Nominations close November 3, 2014.

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 $2.85 per word (minimum $350.00). Display advertising rates are available on request.*

Advertising copy must be received at least four weeks before publication (e.g., the deadline for the July/August 2014 issue is June 13, 2014).

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

Read
Professional Opportunities online:
jobs.siam.org/

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 $1.85 per word (minimum $275.00). Announcements must be received at least one month before publication (e.g., the deadline for the July/August 2014 issue is June 13, 2014).*

Institute for Applied Mathematics and its Applications

Nominations for Board of Governors

The Institute for Applied Mathematics and its Applications invites nominations for its Board of Governors. Applicants can self-nominate, or they can be nominated by others.

The IMA's board consists of 15 distinguished members from academia, industry, and government. The board is the principal governing body of the IMA. Incoming members of the board will serve a five-year term, beginning on January 1, 2015.

The role of the board is twofold: first, to provide oversight and advice on matters of institute management, development, and institutional relationships; and second, for its members to play an active role in the scientific planning for and development of annual program themes, as well as the identification of lead program organizers. The board meets for two days annually; subcommittees meet several times annually via conference call.

Submission of Nominations: Nominations for prospective members of the board should be submitted via the online form at http://www.ima.umn.edu/bog. All nominations will be reviewed by the Nominations Committee. Applicants will be notified of the committee's decision no later than December 1, 2014.

Closing Date: Nominations are due no later than July 31, 2014.

Readers with questions should contact: Fadil Santosa, santosa@ima.umn.edu, the IMA director, or Dana Randall, randall@cc.gatech.edu, chair of the IMA Board of Governors.

NIBIB

continued from page 2

■ Explain that models are platforms for iterative research efforts, and not necessarily the end solution.

■ Be prepared to explain how modeling can efficiently and economically drive scientific experimentation, data acquisition, technology development, and even policy. Use examples to show how modeling has the potential to accelerate translation from basic science to clinical medicine.

■ Share success stories. Do you have or know of a successful biomedical model that has predicted outcomes that are difficult or impossible to obtain through experiments? If so, NIBIB wants to know about your story (which can be sent to SuccessStories@mail.nih.gov).

■ Finally, members of the SIAM community are encouraged to apply for NIH funding and to participate in the NIH Peer Review process. Information on serving as a reviewer can be found at: http://grants.nih.gov/grants/peer_review_process.htm#Overview.

References

[1] B. Koo et al., *Diagnosis of ischemia-causing coronary stenoses by noninvasive fractional flow reserve computed from coronary computed tomographic angiograms: Results from the prospective multicenter DISCOVER-FLOW (Diagnosis of Ischemia-Causing Stenoses Obtained Via Noninvasive Fractional Flow Reserve) study*, J. Am. Coll. Cardiol., 58:19 (2011), 1989–1997.

[2] J.K. Min et al., *Diagnostic accuracy of fractional flow reserve from anatomic CT angiography*, JAMA, 308:12 (2012), 1237–1245.

[3] B.L. Nørgaard et al., *Diagnostic performance of non-invasive fractional flow reserve derived from coronary CT angiography in suspected coronary artery disease: The NXT trial*, J. Am. Coll. Cardiol., January 2014.

[4] K. Sainani, *Meet the skeptics: Why some doubt biomedical models—and what it takes to win them over*, Biomedical Computation Review, June 5, 2012.

Grace C.Y. Peng is a program director at the National Institute of Biomedical Imaging and Bioengineering, National Institutes of Health, and chair of the Interagency Modeling and Analysis Group. Margot Kern is a science writer for NIBIB.

Fault Tolerance

continued from page 1

introducing redundancy will lead to costs in power or performance. Thus, vendors will find it more and more tempting to relax system correctness, but only as long as applications can make up for it. Finally, it is not clear that existing system mechanisms for recovering from a parallel process loss—in particular, global coordinated checkpoint-restart—will work at extreme scales. This is due to their high disk bandwidth requirements, which introduce hardware and power costs poorly related to actual needs of applications. Together, these trends explain why we algorithm developers need to look at this problem now, so that we can be involved in system design discussions.

During the Resilience at Exascale panel discussion, the mood of both panel and audience shifted audibly as soon as Shekhar Borkar (Intel Corp.) hinted that computer hardware might start to need the cooperation of system software for correctness. Getting good performance on *correct* parallel computers is already hard. Consider the difficulty of porting an MPI-only library (the Message Passing Interface, a model for distributed-memory parallelism) to add shared-memory parallelism. Even worse, consider that on future computers, arithmetic or storage could be silently incorrect, or a process could fail at any time and algorithms (rather than the checkpointing library) would be asked to deal with it. The growing uncertainty about system reliability causes special concern, given that many of the algorithms our community develops are used in simulations to support high-consequence decisions. Silent wrong answers could cause death or injury, or could have unacceptable financial or environmental costs. Our minisymposium offered some hope, however, that algorithms will be able to address this concern.

In closing, we point out that some of the techniques that can make algorithms get the right answer despite hardware faults can also improve their resistance to other problems—including software bugs, input that violates an algorithm's assumptions, and other issues that affect correctness (such as loss of energy conservation). It was inspiring to read the software development history of the Mars rover Curiosity in the *Communications of the ACM* (February 2014). If large-scale computers are indeed becoming more and more sensitive to the physics of the real world, we scientific and engineering algorithm developers can take lessons from the embedded and real-time control communities. Moreover, principled use of the least expensive and most generally applicable resilience techniques can make codes safer and easier to use, even if our guesses about fault rates turn out to be exaggerated.

Slides with synchronized audio for many sessions from the 2014 parallel processing conference, including the minisymposium discussed here, can be accessed at www.siam.org/meetings/presents.php.

With exascale computing on the horizon, Mark Hoemmen of Sandia National Laboratories makes the case in the accompanying article that algorithm developers need to get involved in system design now. He is shown here at the conference with Malin Kallen of Uppsala University (left) and Sarah Knepper of Intel.

Ali Pinar awarded the SIAG/SC Junior Scientist Prize to Jed Brown of Argonne National Laboratory. Brown, cited "for his work on algorithms research and impact on software development for parallel scientific computing," titled his prize talk "Exploits in Implicitness: Hardware, Problem Structure, and Library Design."
The prize was created to recognize outstanding junior researchers for significant research contributions to algorithms for parallel computing in science and engineering in the three calendar years prior to the year of the award. Conference photos by Susan Whitehouse.

Flight Path

continued from page 1

Although the map is qualitatively correct, the data it seeks to represent does not agree quantitatively with the three reliable data sets described previously.

Because the methodology used by the Inmarsat team to deduce possible flight paths depends on accurate knowledge of the locations of the arcs, it is not currently possible for their calculations to be accurately reproduced by outsiders. Instead, in this article I illustrate the mathematical methods used by the Inmarsat team by considering fictitious flight XX123, starting at 12 noon from Colombo, Sri Lanka, and moving toward Perth, Australia, at a constant speed of 400 knots (741 km/h) along an arc of a great circle. The only data assumed to be reliable are that the plane was in the southern hemisphere somewhere on three known arcs at 3 PM, 5 PM, and 7 PM, and that the starting location was at 7º N, 80º E. The three arcs, which are shown in red in Figure 1, go through (92.6º E, 8.5º S), (101.3º E, 18.5º S), and (111.6º E, 28.4º S). The goal, using only this data, is to deduce the speed and flight path of the fictitious aircraft, assuming that it was traveling on a great circle.

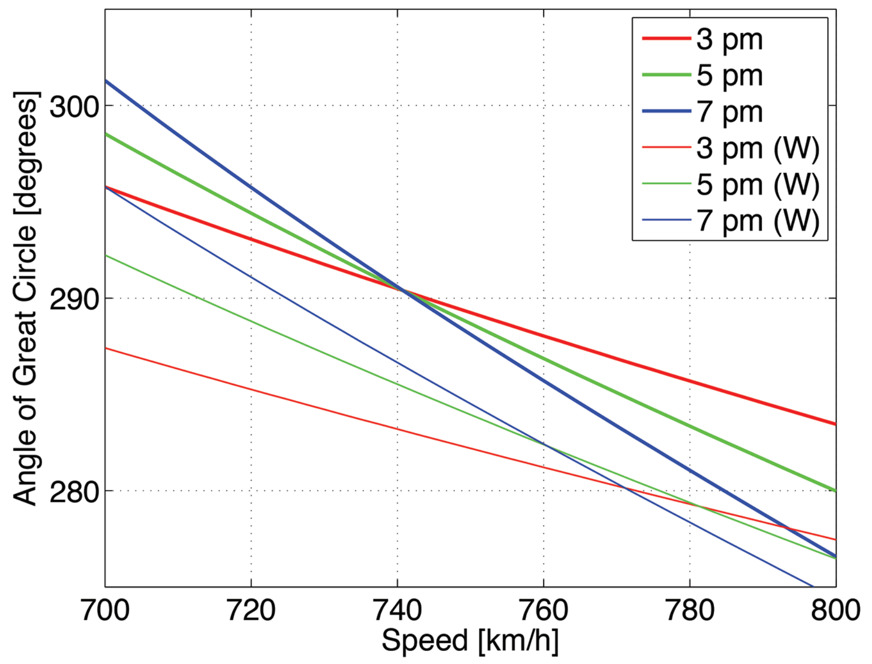

Figure 2. *Direction of XX123's flight path as a function of speed. For each speed on the horizontal axis, the three heavier curves show the angle on the vertical axis required for fictitious flight XX123 to cross the three arcs in Figure 1 at the required times, assuming that the plane starts from Colombo. A match can be detected at a speed of 740 km/h (400 knots). If the starting location is moved 220 km east, however, the thin solid curves indicated with a* W *in the legend exhibit an approximate match at a speed of 770 km/h (416 knots).*

Question 1: Did the Plane Head North or South?

The Inmarsat team used a sophisticated analysis of Doppler-shift data to conclude that MH370 headed south into the Indian Ocean, rather than north into central Asia. To understand why Doppler data could be used to reach this conclusion, let's perform the following thought experiment:

Your friend holds a basketball at eye level, with her hands outstretched. Her eye is the satellite; the basketball is the earth. Your friend is looking straight at a point on the equator of the basketball. This is the position of the satellite, indicated by the open black circle in the top left corner of Figure 1. (In reality, this point is just above the equator, but that doesn't matter for the purposes of our thought experiment.)

We represent the flight of the plane by the motion of an ant walking on the basketball. The ant starts just above the equator. You want to see whether the ant is getting closer to or farther from your friend's eye. Suppose that the ant is walking due south on the basketball. While it is still in the northern hemisphere, the distance between the ant and your friend's eye decreases. The ant is closest to your friend's eye at the instant it crosses the equator. Once it passes into the southern hemisphere, the ant starts to move away from your friend's eye.

The Doppler effect is a property of an electromagnetic signal that is sent by one object (the aircraft) and received by another (the satellite). If the aircraft is moving toward the satellite, the light received by the satellite is more blue than the light that was sent; if the aircraft is moving away from the satellite, the received light is more red. The change in the frequency of the transmitted light, whether toward (high-frequency) blue or (low-frequency) red, is called the Doppler shift. The Inmarsat engineers were able to calculate the Doppler shift at each ping. As I explain in the Epilogue, their results showed that the plane took the southerly route. Before releasing this finding to the public, they validated their results by comparison with satellite data gathered by other planes that were flying in a known direction.

Question 2: How Did Inmarsat Deduce Possible Flight Paths?

Where in the Southern Hemisphere did MH370 go? The shape and location of the four red arcs in the *Washington Post* map of March 15 provide an important clue. As in the map shown in Figure 1 for fictitious flight XX123, the *Washington Post* map for MH370 suggests that the arcs are on concentric circles whose common center is given by the position of the satellite.

This fact can be established from the following quote from an AAIB press report of March 25[3]:

The position of the satellite is known, and the time that it takes the signal to be sent and received, via the satellite, to the ground station can be used to establish the range of the aircraft from the satellite. This information was used to generate arcs of possible positions from which the Northern and Southern corridors were established.

In simpler terms, every point on a given arc is at the same distance from the satellite. If you think of the satellite as being located at the tip of an ice cream cone, then the arcs are on the intersection of the cone with a spherical ball of ice cream in the cone.

Let's imagine that the satellite antenna looks like the familiar dish antennas we often see here on earth. When the satellite antenna points straight down to earth, it points at the position of the satellite designated by the little black circle in the top left corner of Figure 1. For the purposes of this discussion, let's call this point on the earth the *satellite pole*. The role it will play is similar to that of the North Pole on familiar maps of the earth. We can also talk about longitude and latitude measured with respect to the satellite pole. We refer to these notions of longitude and latitude as being in the *satellite coordinate system*. The red arcs are circles of latitude measured in the satellite coordinate system.

To estimate the flight path of MH370, the Inmarsat team initially assumed that after a certain time the plane switched to auto-pilot and was therefore traveling along the arc of a great circle at close to constant speed and at a constant elevation. From radar data, we know that the speed of the plane an hour after takeoff was 540 mph. In the approach presented here, by contrast, rather than specifying the speed of the plane at the beginning of the calculation, we infer it in the course of the calculation.

Before going into technical details, let's do a second thought experiment, also using a basketball and, because we don't know the exact locations of the arcs for MH370, using instead the data for flight XX123. For this experiment:

Begin by drawing the three XX123 arcs on the basketball with a red marker. Now put the basketball in a hoop that fits snugly around the ball. Pin the rim onto the basketball at the starting location of the plane at 12 noon, and put another pin at the exact opposite point on the ball. Now we consider all the possible great circles through this point. Just as we can tell which line of longitude we are on by specifying how far east or west of the Greenwich meridian we are, we can tell which great circle we are on using an angle that I call the *angle of the great circle*. The rim of the hoop can now define any of the great circles that pass through the starting point. Imagine that the rim has a ruler marked on it.

To continue, you need to pick a speed for the plane. Because you can't assume any knowledge of the actual speed of the plane, you need to choose a reasonable range of speeds and do the experiment for each speed in that range. For each choice of speed, draw marks on the rim that are equally spaced and that correspond to the distance the plane would travel in two hours at the speed you picked. Now rotate the rim around the basketball. Do the three marks on the ruler line up exactly with the three red arcs? If so, you have a *match*. If you find a match, you have found a possible flight path.

You could actually do this experiment with a real basketball and rim. Unless someone had told you ahead of time that there was an answer to be found, and had shown you roughly where to look for it, you might find it extremely challenging to find a match. This is what makes the original work of the Inmarsat team so impressive. For XX123, I worked out some mathematical equations that enabled me to use a computer simulation to search for a match. The derivation of the equations uses mathematics—trigonometry, vector algebra, and appropriate choices of coordinates on the sphere—that could easily be explained to high school students. I presume that the Inmarsat team made similar calculations. I then coded my formula up in Matlab and ran the code to find the blue flight path for XX123 shown in Figure 1.

Back in our thought experiment, if you didn't pin the rim on at exactly the right starting place, you might have come up with a very different solution. This, in essence, is the problem faced by the Inmarsat engineers, who didn't know the exact location of the plane at the time it went on auto-pilot. In fact, the March 28 move of the search area was motivated by a more careful analysis of Malaysian radar data; the analysis suggested that before going on auto-pilot, the plane had traveled faster and hence farther west than thought initially. This sensitivity to initial conditions is illustrated by the green flight path shown in Figure 1, which starts 220 km due east of the "true" starting location of XX123.

Recall that to perform this experiment you had to pick a speed for the plane. If you chose the speed incorrectly, you may not have found a match. Figure 2 illustrates the matches I found for flight XX123. The horizontal axis of the plot shows the range of speeds I used in my search for a match. For each of the three arcs and for each speed, I computed the angle of the great circle along which the plane would have to fly to get from its starting location to that arc in the required time. The vertical axis shows those angles. You have a match when there is an angle and a speed at which all three curves meet.

Epilogue

From a given starting location there are two possible flight paths that cross the arcs at the correct times: one headed south and the other north.

On March 25, AAIB released the results of a sophisticated analysis of Doppler-shift data[6] for the relative motion of MH370 and the satellite, which provides strong confirmation that the aircraft took the southerly route. As I will explain in a forthcoming article, the Inmarsat engineering team was able to reach this conclusion only by taking into account a slight elliptical rotation of the satellite about its nominal location above the equator.

This new technique developed by the Inmarsat team—a tour de force—is central to their ongoing investigation of the fate of MH370. Recently, I analyzed the technique mathematically and derived a formula for the velocity of a plane at each ping time in terms of the Doppler shift and the angles of the arcs. (Formulae like this are surely well known to experts in the important field of Doppler-based tracking.) In a future pedagogical article, I plan to demonstrate how, with arc angles and Doppler shifts measured at suitably chosen ping times, this formula can be used to determine the flight path of a plane, even if it is not flying along a great circle or at constant speed. It has been challenging for the Inmarsat team to use such a method to completely eliminate the uncertainty in the flight path of MH370 because, at times during the first three hours, the path of the plane was extremely erratic and the plane was not pinged at sufficiently short intervals during that time.

Acknowledgments

Thanks to Joel Achenbach, Alan Boyle, Matthew Goeckner, Brian Marks, and Yannan Shen for helpful conversations. Yanping Chen has produced movies illustrating the thought experiments in this article; see http://www.utdallas.edu/~zweck/.

Notes

[1] http://www.nytimes.com/2014/03/23/world/asia/a-routine-flight-till-both-routine-and-flight-vanish.html.

[2] http://www.n2yo.com/?s=23839.

[3] http://apps.washingtonpost.com/g/documents/world/heres-an-analysis-of-satellite-data-from-britains-aaib/892/.

[4] http://www.slate.com/blogs/future_tense/2014/03/15/flight_370_disappearance_missing_airliner_apparently_flew_to_central_asia.html.

[5] http://www.washingtonpost.com/national/health-science/satellite-locates-malaysian-flight-370-still-flying-seven-hours-after-takeoff/2014/03/15/96627a24-ac86-11e3-a06a-e3230a43d6cb_graphic.html?hpid=z2.

[6] http://apps.washingtonpost.com/g/documents/world/images-from-britains-aaib-on-mh370/893/.

John Zweck is a professor of mathematics at the University of Texas at Dallas.

NSF Updates Big Data Program

The National Science Foundation has issued an updated solicitation for its Critical Techniques and Technologies for Advancing Big Data Science & Engineering (BIGDATA) program. Details can be found at http://www.nsf.gov/funding/pgm_summ.jsp?pims_id=504767.

The deadline for full proposals is June 9, 2014.

The solicitation invites two types of proposals: "Foundations" (F), for those developing or studying fundamental techniques, theories, methodologies, and technologies of broad applicability to big data problems; and "Innovative Applications" (IA), for those developing techniques, methodologies, and technologies of key importance to a big data problem that directly impacts at least one specific application.

All proposals must address critical challenges for big data management, big data analytics, or scientific discovery processes impacted by big data. These techniques, methodologies, and technologies can be computational, statistical, or mathematical in nature; proposals can focus on novel theoretical analysis or experimental evaluation of these techniques and methodologies.