from *SIAM News*, Volume 35, Number 6, July/August 2002

The $100, 100-digit Challenge

Chastened Challenge Sponsor: "I Misjudged"

By Lloyd N. Trefethen

I've marked many problem sets in my time, but never a set like these. There's no other word for it. What *joy* the contes-tants took in this "Decimal Decathlon"!

The story began last January, when *SIAM News* printed a collection of ten mathematical problems. For each one, the solution was a single real number, and the challenge to readers was, Can you determine these numbers to many digits of accuracy? I announced that I would give a point for each correct digit, up to ten digits for each problem. Thus, the maximum possible score was 100, and the top scorer would win $100. The deadline was set at May 20, giving problem solvers four months of working time.

News of the challenge spread, as happens nowadays. SIAM's Web site and my own were amplified by announcements on half a dozen online discussion groups and in *Science* magazine. Copies appeared in coffee rooms. People got hooked. My mailbox started buzzing.

The contestants were graduate students and undergraduates, professors and teachers, mathematicians and physicists, industrial employees and retirees. They were singletons and pairs and triplets and teams of up to the maximum permitted size of six. All told, I received entries from 94 teams in 25 countries, from Chile to Canada and South Africa to Finland. The USA dominated, followed by Germany, but per capita, I was most struck by two outstanding entries from Slovenia, four from Switzerland, and four from the Netherlands.

In the January announcement I wrote, "If anyone gets 50 digits in total, I will be impressed." Well, I misjudged. In the end, along with plenty of scores in the twenties, forties, and sixties, at the top we had *twenty* perfect scores. After them came *five teams who scored 99 out of 100*. Six further teams scored between 90 and 98.

I am proud to announce the twenty first-prize teams:

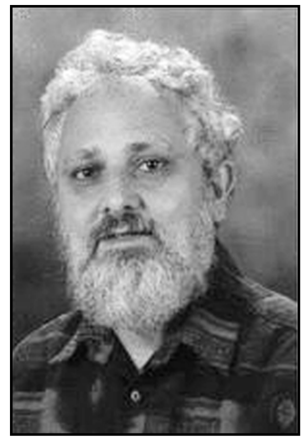

Among the first-prize winners was Dirk Laurie of the University of Stellenbosch, South Africa.

1. Peter Robinson of Quintessa, Ltd., in Henley-on-Thames, England.
2. J. Boersma, J.K.M. Jansen, F.H. Simons, and F.W. Steutel of the Eindhoven University of Technology, the Netherlands.
3. Ruud van Damme, Bernard Geurts, and Bert Jagers of the University of Twente, the Netherlands.
4. Gerhard Kirchner, Alexander Ostermann, Mechthild Thalhammer, and Peter Wagner of the University of Innsbruck, Austria.
5. Gaston Gonnet of the Eidgenössische Technische Hochschule (ETH) in Zurich and Robert Israel of the University of British Columbia, Canada. (Gonnet is one of the creators of Maple.)
6. Rolf Strebel and Oscar Chinellato, also of ETH.
7. Folkmar Bornemann of the Technical University of Munich, Germany.
8. Thomas Grund of the Technical University of Chemnitz, Germany.
9. Gerd Kunert and Ulf Kähler, also of TU Chemnitz.
10. Dirk Laurie of the University of Stellenbosch, South Africa.
11. "The Blue Hens": Carl Devore, Toby Driscoll, Eli Faulkner, Jon Leighton, Sven Reichard, and Lou Rossi of the University of Delaware.
12. "The Compuserve SCIMATH Forum Team": Marijke van Gans, Brian Medley, and Bernard Beard. (More on these below.)
13. "L'Equipe McGill": Martin Gander, Felix Kwok, Sebastien Loisel, Nilima Nigam, and Paul Tupper of McGill University in Montreal, Canada.
14. Danny Kaplan and Stan Wagon of Macalester College in St. Paul, Minnesota. Wagon is a veteran problem solver: He runs the Problem of the Week Web site at Macalester and is a co-author of a book (*Which Way Did the Bicycle Go?*) that compiles some of those problems.
15. Kim McInturff of Raytheon and Peter S. Simon of Space Systems/Loral, both in California.
16. Glenn Ierley, Stefan Llewellyn Smith, and Robert Parker of the University of California, San Diego.
17. "The Rice Team": Eric Dussaud, Chris Husband, Hoang Nguyen, Daniel Reynolds, and Christiaan Stolk of Rice University.
18. Jingfang Huang, Michael Minion, and Michael Taylor of the University of North Carolina.
19. Eddy van de Wetering of West Hartford, Connecticut. Van de Wetering is a physicist who works in financial derivatives.
20. "Team Mathematica": Paul Abbott of the University of Western Australia, together with Brett Champion, Yifan Hu, Daniel Lichtblau, and Michael Trott of Wolfram Research, Inc.

All these groups are equal winners and deserve to receive $100. Funds are short, however, consisting of the original $100 plus an additional $200 kindly contributed by Oxford Scientific Consulting, Ltd. Thus, I have quite arbitrarily chosen teams 10, 11, and 12 to receive checks. Here are a few words about this arbitrary subset of our set of winners.

Dirk Laurie was one of five contestants who showed that a single remarkable individual can do it all. Laurie is one of South Africa's leading numerical analysts, a senior figure at the University of Stellenbosch. This is not the first time he has been in the public eye. Twice in past years he won second place on South Africa's television quiz show "Flinkdink," and to this day, he says, some strangers recognize him from those moments in the limelight.

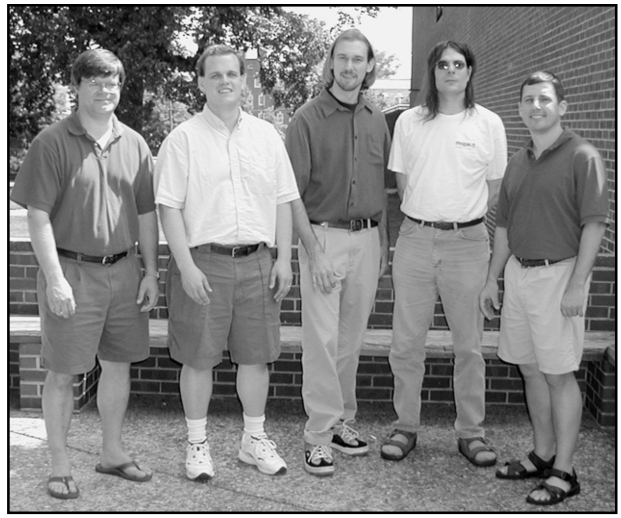

Delaware first-prize winners (left to right) Jon Leighton, Toby Driscoll, Eli Faulkner, Carl Devore, and Lou Rossi of the University of Delaware. Teammate Sven Reichard is not shown.

The Delaware Team exemplifies collaboration between students and faculty at a university. Toby Driscoll and Lou Rossi are assistant professors, Sven Reichard and Carl Devore are graduate students, and Eli Faulkner is an undergraduate about to begin his senior year; all are in the Department of Mathematical Sciences. Jonathan Leighton is a continuing education student in computer science.

The Compuserve SCIMATH Forum Team represents collaboration of a modern kind—these three have never met! They are Internet friends and regular participants in this online mathematics group. Marijke van Gans lives on the Isle of Bute in Scotland and intends in the future to go to graduate school. Brian Medley is a sixth-form (high school) mathematics teacher in Wigan, not far from Manchester in northern England. Bernard B. Beard is an associate professor at Christian Brothers University in Memphis, Tennessee. This remarkable trio exchanged ideas and computer files for months until they were sure they had nailed all ten answers. Beard collected the files into a 55-page pdf document full of zest and creativity and some beautiful illustrations. My friends, I hope to meet you one day, and I hope one day you meet each other!

The 99-digit second-prize winners were Niclas Carlson of the Åbo Akademi University, Finland; Michel Kern of INRIA in Rocquencourt, France; David Smith of Loyola Marymount University, in Los Angeles; Craig Wiegert of the University of Chicago; and Katherine Hegewisch and Dirk Robinson of Washington State University.

All right, then, what about these problems?

People solved them by all kinds of methods—implemented in Mathematica, Matlab, Maple, C/C++, Fortran, and, occasionally, in PARI, Visual Basic, UBASIC, Octave, Java, Pascal, or GSL. (One contestant crunched his numbers in Excel, but he was not among the high scorers.) Some contestants are world-famous mathematicians and favored advanced solution techniques. Others are more junior academics, students, or amateurs and tended to invent more techniques for themselves.

There's a pattern in how the work on such problems goes. You can solve almost any problem to two- or three-digit accuracy by some kind of brute force method. Most of us start out this way. To get more digits, you have to be cleverer, which in many cases can be achieved by some kind of extrapolation. This might be extrapolation in mesh size, limit of integration, or matrix dimension, so long as the number to be extrapolated depends regularly on the parameter. This approach can often be made to yield ten digits—and once you've got ten, if your software system permits, you can often get 50 or even 500. For most of the problems, I received solutions from several teams that agreed to 100 digits or more.

Then there's a third stage, attained by some contestants for some problems: the dazzlingly slick. An idea comes along, advanced or elementary, that simply busts the problem open. It's thrilling when this happens; but in all ten cases, it is entirely possible to get those ten digits by less specialized means.

I had thought of writing a learned article about these problems, but the barriers to entry have been raised too high! As I write these words, eight of the first-prize teams have already posted solutions online (follow the pointers from my Web page, www.comlab.ox.ac.uk), and several of them are preparing technical reports. Surely few numerical exercises have ever been examined in such minute detail as these ten!

So here goes, just a casual summary.

Problem 1. What is $I = \lim_{\varepsilon\to 0}\int_\varepsilon^1 x^{-1}\cos(x^{-1}\log x)dx$? This integrand is devilish because it diverges to ∞ and oscillates infinitely often. Unless you find a way to use them very cleverly, the automatic integrators in systems like Maple cannot get ten digits. Many people made use of the Lambert W function, and many changed variables with $t = x^{-1}$ to get $I = \lim_{L\to\infty}\int_1^L t^{-1}\cos(t\log t)dt$. But the problem of infinitely many oscillations must still be dealt with. A popular approach was to separate the oscillations to obtain an alternating series of contributions to the integral, which can then be summed and extrapolated by Aitken extrapolation or other methods. Many people also used integration by parts one or more times to make the integrand better behaved.

Here's the slickest solution I know: Since the cosine is the real part of the complex exponential, we can write $I = \mathrm{Re}\lim_{L\to\infty}\int_1^L t^{-1}\exp(it\log t)dt$, that is, $I = \mathrm{Re}\lim_{L\to\infty}\int_1^L t^{-1}t^{it}dt$. Now, since the integrand is an analytic function, we can rotate contours in the complex plane to get $I = \mathrm{Re}\lim_{L\to\infty}\int_1^{1+iL} t^{-1}t^{it}dt$. The oscillations are gone! For 15-digit accuracy, $L = 15$ is big enough, and standard numerical integrators give the answer in one second of computer time. *Solution:* 0.323367431677778 . . .

Problem 2. A photon moving at speed 1 *in the x–y plane starts at* $t = 0$ *at* $(x,y) = (1/2, 1/10)$*, heading due east. Around every integer lattice point (i,j) in the plane, a circular mirror of radius* 1/3 *has been erected. How far from* (0,0) *is the photon at* $t = 10$? This problem is elementary; a little care with geometry gives you a convincing trajectory. But there is a catch: The problem is ill-conditioned, for the dynamics are *chaotic*. Small perturbations in the initial conditions, or small perturbations introduced by rounding errors along the way, grow exponentially with time. After ten time units, they have been amplified by about eleven orders of magnitude, and, as a result, if you are working in the standard IEEE 16-digit precision, you cannot compute the answer to more than five or six significant digits. So extended precision is needed for this problem. Once you realize this, you can solve it in any number of software systems. The photon bounces 14 times and ends up near (–0.736, 0.670). (I should have been bolder and asked for the solution at $t = 100$.)

First-prize winners (front to back) Brian Medley, Bernard Beard, and Marijke van Gans of the Compuserve SCIMATH Forum. This cyber-photo symbolizes their cyber-collaboration: The three have never met.

This is the only one of the ten problems that requires extended precision—and hence the only one that cannot be solved in standard Matlab. To put it another way, it's the only one for which rounding errors are the key issue. In general, fast convergence of algorithms is a much bigger focus of effort in scientific computing than control of rounding errors. (I argued this point in a *SIAM News* essay, "The Definition of Numerical Analysis," in November 1992.) *Solution:* 0.995262919443354 . . .

Problem 3. The infinite matrix A with entries $a_{11} = 1$, $a_{12} = 1/2$, $a_{21} = 1/3$, $a_{13} = 1/4$, $a_{22} = 1/5$, $a_{31} = 1/6$, *etc., is a bounded operator on* ℓ^2. *What is* $\|A\|$*?* The norm is the largest singular value of A or, equivalently, the square root of the largest eigenvalue of A^TA. Almost all teams approached this problem by working with $N \times N$ finite sections of the infinite matrix. For $\|A\|$ to be accurate to ten digits, N must be in the thousands, too high for most computer systems. Therefore, most contestants found the needed digits by extrapolation of, for example, results for dimensions $N = 2, 4, 8, \ldots, 512$. *Solution:* 1.27422415282122 . . .

Problem 4. What is the global minimum of the function $\exp(\sin(50x)) + \sin(60e^y) + \sin(70\sin(x)) + \sin(\sin(80y)) - \sin(10(x+y)) + \frac{1}{4}(x^2+y^2)$*?* This function is cooked up to have infinitely many local minima—667 of them in the unit square alone—but only one global minimum. But it isn't hard to figure out which local minimum is the right one; if you look closely at a contour plot you'll find that it is near (–0.024,0.211). Starting with this initial guess, you can find the minimum to high precision with all kinds of numerical methods. Alternatively, the `NMinimize` command in Mathematica finds the global minimum without too much trouble if you use its `SimulatedAnnealing` or `DifferentialEvolution` options. *Solution:* –3.30686864747523 . . .

Problem 5. Let $f(z) = 1/\Gamma(z)$, *where* $\Gamma(z)$ *is the gamma function, and let* $p(z)$ *be the cubic polynomial that best approximates* $f(z)$ *on the unit disk in the supremum norm* $\|\cdot\|_\infty$. *What is* $\|f - p\|_\infty$*?* This problem of complex Chebyshev approximation was the one that gave people the most trouble, perhaps because the complex functions have an advanced flavor that not all contestants are comfortable with. Certainly one needs to make use of at least one fact of complex analysis: By the maximum modulus principle, $f - p$ will achieve its largest magnitude on the boundary of the domain of approximation, and thus we need only consider z on the unit circle. It also helps to note that it's enough to consider polynomials with real coefficients—for if p has complex coefficients, then $p(z) + (\bar{p}(\bar{z}))/2$ will be an equally good approximant (in fact, better) that is real.

I admire the hard work and ingenuity that so many teams expended to find their solutions to this problem. After trying one method, for example, the SCIMATH people wrote, "Every trick in the book gets stuck in the mud of this gully!" Undaunted, they persevered until they found another method that worked. Half a dozen experts made use of specialized algorithms and software created for such problems by Fischer and Modersitzki, Le Bailly and Thiran, or Tang. *Solution:* 0.214335234590459 . . .

Problem 6. A flea starts at (0,0) *on the infinite 2D integer lattice and executes a biased random walk: At each step it hops north or south with probability* 1/4, *east with probability* $1/4 + \varepsilon$, *and west with probability* $1/4 - \varepsilon$. *The probability that the flea returns to* (0,0) *sometime during its wanderings is* 1/2. *What is* ε*?* If no wind is blowing, i.e., if $\varepsilon = 0$, then the flea returns to the origin infinitely often with probability 1. For any nonzero ε, however, the probability of return falls below 1. Many people simulated this problem by numerical experiments to get a few digits. In such a simulation, as in more advanced methods, one must be careful not to confuse the *expected number of returns* ν with the *probability of return* $p = \nu/(\nu + 1)$.

The high-precision methods devised for this problem were wide-ranging. Some contestants used 1D or 2D recurrence relations and infinite series; some used generating functions and Fourier analysis; some formulated a linear algebra problem to determine probabilities at each lattice point on the grid, a "lattice Green function" that can in turn be related to a double integral. The probability of return, given ε, can in fact be written down explicitly in terms of elliptic functions; but in any case one ends up solving a nonlinear equation to find the value of ε that gives $p = 1/2$. *Solution:* 0.0619139544739909 . . . (or its negative)

Problem 7. Let A be the 20,000-by-20,000 matrix whose entries are zero everywhere except for the primes 2, 3, 5, 7, . . . , 224737 along the main diagonal and the number 1 in all the positions a_{ij} with $|i - j| = 1, 2, 4, 8, \ldots, 16384$. What is the (1,1) entry of A^{-1}*?* When I posed this problem, I was thinking of matrix iterations. The matrix is too big for a direct solution on most computers, but by the conjugate gradient method with a diagonal preconditioner, one can find the answer in fewer than twenty iterations. The required number is the first entry in the solution vector x of the system $Ax = e_1$, where $e_1 = (1,0,0,\ldots)^T$. Virtually every contestant found the answer by an iterative method like this, often using simpler iterations, such as Gauss–Seidel or a method devised ad hoc.

But would you believe that the problem can be solved *exactly*? Not just directly, but truly exactly! The number we are looking for is the ratio of two determinants (Cramer's rule), each of which is an integer. "All" one has to do is to find these integers. I never dreamed of such a thing, but one of our contestants, the "LinBox team," actually achieved it. Jean-Guillaume Dumas of LMC_IMAG in Grenoble ran 186 processors for four days using LinBox software; the mathematics involved blackbox modular techniques and the Chinese remainder theorem. The numerator and denominator are relatively prime and each has exactly 97,390 digits: The quotient is 310164 . . . 7075/427766 . . . 5182, where each ". . ." represents 97,380 omitted digits. Zhendong Wan of the University of Delaware subsequently verified the result on one processor with a large memory (4GB), in 12 days of further computation. *Solution:* 0.725078346268401 . . .

Problem 8. A square plate $[-1,1] \times [-1,1]$ *is at temperature* $u = 0$. *At time* $t = 0$ *the temperature is increased to* $u = 5$ *along one of the four sides while being held at* $u = 0$ *along the other three sides, and heat then flows into the plate according to* $u_t = \Delta u$. *When does the temperature* *reach* $u = 1$ *at the center of the plate?* The heat equation is one of the most basic of all partial differential equations, and there are many ways to get a few digits of precision for this problem. A good start, taking advantage of linearity, is to reformulate it so that the temperature is 1 everywhere at $t = 0$ but 0 on all four boundaries; you now look for the time at which $u = 0.2$ at (0,0). From here, you can use a finite difference approximation after exploiting symmetry to cut the domain in four. If the formulation is clean enough, you can compute results for a sequence of mesh sizes, such as $2^{-1}, 2^{-2}, \ldots, 2^{-6}$, and then use Richardson extrapolation to enhance these results to ten or more digits.

There is no need for discretization, however. This problem has so much structure that we can solve it with Fourier series, which converge exponentially fast. Indeed, Joseph Fourier himself, who was concerned with keeping wine cellars cool during summer months, solved just this kind of problem by just this method. For the times of interest, the temperature at (0,0) is captured to 15-digit accuracy by four Fourier terms: $T \approx a^{-2}[e^{-2at} - (2/3)e^{-10at} + (1/9)e^{-18at} + (2/5)e^{-26at}]$, where $a = \pi^2/4$. It's not hard to find the value of t for which this function is equal to 0.2. *Solution:* 0.424011387033688 . . .

Problem 9. The integral $I(\alpha) = \int_0^2 [2 + \sin(10\alpha)]x^{\alpha} \sin(\alpha/(2-x))\,dx$ *depends on the parameter* α. *What is the value* $\alpha \in [0, 5]$ *at which* $I(\alpha)$ *achieves its maximum?* We have here another devilish integrand, oscillating infinitely often. Merely evaluating $I(\alpha)$ is a challenge, and we must do so in a way that is efficient and smooth enough to be the basis of a minimization. This can certainly be done numerically, and by a variety of clever methods, this is what most people did. One needs care and determination, but no specialized knowledge.

Mathematica has some relevant specialized knowledge, however: It knows how to evaluate this integral in closed form! The solution involves a special function dating from the 1930s, known as Meijer's G function, which Mathematica can evaluate numerically, and with the aid of arbitrary-precision arithmetic, the problem now becomes straightforward. *Solution:* 0.785933674350371 . . .

Problem 10. A particle at the center of a 10×1 *rectangle undergoes Brownian motion (i.e., 2D random walk with infinitesimal step lengths) till it hits the boundary. What is the probability that it hits at one of the ends rather than at one of the sides?* With a problem like this, one naturally tries to simulate some sample paths—the Monte Carlo method. One soon finds that fewer than one path in a million escapes out the ends. To find the probability to ten digits of relative accuracy would require something like 10^{26} samples!

So one must be cleverer, and this was a problem on which people showed great ingenuity. Contestants exploited all kinds of elegant reformulations, series, and reflections to discover the required number. Ultimately, what is going on here is a problem of partial differential equations: If the Laplace equation holds on the rectangle with boundary conditions 0 on the sides and 1 on the ends, what is the value at the origin? Experts call this number the *harmonic measure* of the ends with respect to the point (0,0). It can be calculated by conformal mapping or by Fourier series, and most teams ended up with a rapidly convergent series representation. In fact, the solution can be written explicitly in terms of the Jacobi elliptic sine function. An approximation good enough for the required ten digits, based on approximating the rectangle by a semi-infinite strip, is $p = (8/\pi)e^{-5\pi}$. *Solution:* 0.000000383758797925122 . . .

If you add together our heroic numbers, the result is $\tau = 1.497258836\ldots$. I wonder if anyone will ever compute the ten thousandth digit of this fundamental constant.

Lloyd N. Trefethen is a professor in numerical analysis and a fellow of Balliol College at Oxford University.