

SHORT REMARKS ON THE SOLUTION OF THE SIAM 100-DIGIT CHALLENGE

FOLKMAR BORNEMANN

ABSTRACT. In the January 2002 issue of SIAM News [14] Nick Trefethen challenged the mathematical community with ten numerical problems. Each problem, stated in a sentence or two and answered by a single real number, had to be solved for at least ten significant digits to get a maximal score. This winning score, worth \$100, was obtained by 20 teams—out of 94 teams from 25 countries [15]. In this e-print the author, one of the successful twenty, shortly remarks on his solutions, sharing the fun with prospective readers.

PROBLEM 1

Question: What is $\lim_{\epsilon \rightarrow 0} \int_{\epsilon}^1 x^{-1} \cos(x^{-1} \log x) dx$?

Answer (12 significant digits¹): 0.323367431678e+0

Method: The general idea is to transform the highly oscillatory integral into an alternating series,

$$\lim_{\epsilon \rightarrow 0} \int_{\epsilon}^1 \frac{\cos(x^{-1} \log x)}{x} dx = \sum_{k=0}^{\infty} \underbrace{\int_{x_{k+1}}^{x_k} \frac{\cos(x^{-1} \log x)}{x} dx}_{=(-1)^k a_k}$$

where $x_0 = 1$ and, for $k \geq 1$, the zeroes of the integrand,

$$-x_k^{-1} \log(x_k) = \left(k - \frac{1}{2}\right) \pi, \quad \text{i.e.} \quad x_k = \frac{W(u)}{u} \Big|_{u=(k-\frac{1}{2})\pi}.$$

Here, $W(u)$ denotes Lambert's W -function [4], a Matlab implementation of which can be found in the web [10]. Each of the a_k can be computed by a high order quadrature formula using, e.g., Matlab's `quadl` command. The sequence $(a_k)_{k \geq 1}$ turns out to be *totally monotone*. Thus, there are *stable* and efficient methods available to accelerate the convergence of the alternating series. For instance, Wynn's ϵ -algorithm could be used [2, Thm. 2.25][11]; or because of a simple and very useful error estimate, and in my experience often even a better convergence, the recent linear method of Cohen, Rodriguez Villegas and Zagier [3]. It guarantees D decimal digits using $\lceil 1.31D \rceil$ terms of the sequence only.

Date: April 19, 2002; last revision November 5, 2002.

¹With the exception of Problem 2 all the stated 12-digits solutions were obtained using *double precision* and Matlab. Problem 2 required extended precision as realized in most of the general purpose computer algebra systems.

PROBLEM 2

Question: A photon moving at speed 1 in the x - y plane starts at $t = 0$ at $(x, y) = (0.5, 0.1)$ 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 the origin is the photon at $t = 10$?

Answer (12 significant digits): 0.995262919443e+0

Method: Such problems are in general notoriously badly conditioned. For instance, with respect to the y -coordinate of the starting position, the relative condition number is $6.2 \cdot 10^9$. Thus, it is extremely unlikely that there is any algorithm to solve the problem to 10 digits using double precision only. Instead, we use *extended precision* in Maple and a straightforward code for calculating the points of intersections with the circles, searching for the first hit, and reflecting at the tangent. Increasing the precision shows that we loose approximately 12 digits in the final result.




FIGURE 1. The path of the photon from $t = 0$ to $t = 10$ [solid green]. Path for perturbed initial value $(0.5, 0.1 + 10^{-9})$ [dotted red].

PROBLEM 3

Question: 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 $\|A\|$?

Answer (12 significant digits): 0.127422415282e+1

Method: The general term of the infinite matrix A is given by

$$a_{ij} = \frac{2}{i(i+1) + (j-1)(j+2i-2)}, \quad i, j = 1, 2, \dots,$$

which makes the matrix a Hilbert-Schmidt operator on ℓ^2 with

$$\|A\| \leq \|A\|_{\text{HS}} = \sqrt{\sum_{k=1}^{\infty} \frac{1}{k^2}} = \sqrt{\pi^2/6} = 1.28254983016\dots$$

Let $A_n \in \mathbb{R}^{n \times n}$ be the upper-left quadratic $n \times n$ -corner, then there is

$$\lim_{n \rightarrow \infty} \|A_n\| = \|A\|,$$

a very slow convergence of monotonic logarithmic type with the asymptotic expansion

$$\|A_n\| \sim \|A\| + n^{-3}(c_0 + c_1 n^{-1} + \dots), \quad c_0 \approx -4.8.$$

Thus, for 10 digits accuracy one has to choose $n = 3348$, better $n = 4000$. The norm $\|A_n\|$ is calculated by the power method applied to $A_n^T A_n$. However, this approach would not be useful for higher accuracies, say 16 digits, because this would mean hundred times a bigger n , thus a ten thousand times more work.

Due to the asymptotic expansion Wynn's ρ -algorithm is accelerating the convergence [8]. Because of cancellation effects one should use extended precision. *Mathematica's* precision model is especially well suited to cut of the noise in the ρ -algorithm. The number of correct digits is cross checked by applying the same procedure to the even slower monotonely convergent sequence

$$\|A_n\|_{\text{HS}}^2 - \|A_n\|^2 = \frac{\pi^2}{6} - \|A\|^2 + n^{-2}(d_0 + d_1 n^{-1} + \dots).$$

PROBLEM 4

Question: What is the global minimum of the function

$$\begin{aligned} f(x, y) = & \exp(\sin(50x)) + \sin(60e^y) + \sin(70 \sin(x)) \\ & + \sin(\sin(80y)) - \sin(10(x + y)) + \frac{1}{4}(x^2 + y^2)? \end{aligned}$$

Answer (12 significant digits): -0.330686864748e+1

Method: Here, I used a combination of a grid-based direct search and a gradient descent. The main point is that any guess $f_0 = f(x_0, y_0)$ for the global minimum allows to restrict the search to the domain

$$\frac{1}{4}(x^2 + y^2) \leq f_0 + 3 + \sin(1) - \exp(-1).$$

Bounds on the derivatives in this domain allow to separate the local minima on a fine enough grid.

PROBLEM 5

Question: 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 $\|f - p\|_{\infty}$?

Answer (12 significant digits): 0.214335234590e+0

Method: I choose the algorithm of Tang [13] for three reasons: it is quite fast (in this case linearly convergent), easy to code in Matlab, and because of relying on

the dual problem, it gives *lower* and *upper bounds* for the error $\|f - p\|_\infty$. This way it is simple to control the number of correct digits. Matlab implementations of the *complex* gamma function $\Gamma(z)$ and its logarithmic derivative, the digamma function $\psi(z) = \Gamma'(z)/\Gamma(z)$, can be found in the web [5].

PROBLEM 6

Question: 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 $1/4 + \epsilon$, and west with probability $1/4 - \epsilon$. The probability that the flea returns to $(0,0)$ sometime during its wanderings is $1/2$. What is ϵ ?

Answer (12 significant digits): 0.619139544740e-1

Method: Let the probability of return for a given $\epsilon > 0$ be p_ϵ . It relates to m_ϵ , the expected number of visits to the origin (including the start), by

$$m_\epsilon = \frac{1}{1 - p_\epsilon},$$

cf. [12, P1.4]. Introducing the *characteristic function* $\phi_\epsilon(\theta)$ of the random walk, cf. [12, D6.2],

$$\begin{aligned} \phi_\epsilon(\theta_1, \theta_2) &= \frac{1}{4}e^{i\theta_1} + \frac{1}{4}e^{-i\theta_1} + \left(\frac{1}{4} + \epsilon\right)e^{i\theta_2} + \left(\frac{1}{4} - \epsilon\right)e^{-i\theta_2} \\ &= \frac{1}{2}(\cos(\theta_1) + \cos(\theta_2) + 4i\epsilon \sin(\theta_2)), \end{aligned}$$

allows one to calculate m_ϵ by a recurrence criterion of Chung and Fuchs [12, P8.1]

$$\begin{aligned} m_\epsilon &= \lim_{t \uparrow 1} \frac{1}{4\pi^2} \int_{-\pi}^{\pi} \int_{-\pi}^{\pi} \Re \frac{1}{1 - t\phi_\epsilon(\theta_1, \theta_2)} d\theta_2 d\theta_1 \\ &= \frac{1}{\pi} \int_0^\pi \frac{2}{\sqrt{3 - 4\cos\theta_1 + \cos^2\theta_1 + 16\epsilon^2}} d\theta_1 \\ &= \frac{1}{\pi} \int_{-1}^1 \frac{2}{\sqrt{1 - t^2} \sqrt{3 - 4t + t^2 + 16\epsilon^2}} dt \\ &= \frac{2\sqrt{2} K\left(1 - \frac{1+8\epsilon^2 - \sqrt{1-16\epsilon^2}}{1+8\epsilon^2 + \sqrt{1-16\epsilon^2}}\right)}{\pi \sqrt{1 + 8\epsilon^2 + \sqrt{1 - 16\epsilon^2}}} \\ &= \frac{\sqrt{2}}{\text{AGM}(\sqrt{1 + 8\epsilon^2 - \sqrt{1 - 16\epsilon^2}}, \sqrt{1 + 8\epsilon^2 + \sqrt{1 - 16\epsilon^2}})}. \end{aligned}$$

Here, $K(m)$ denotes the complete elliptic integral of the first kind of parameter m (*not* modulus k), and AGM denotes the arithmetic geometric mean. This way, the equation $p_\epsilon = 1/2$ is equivalent to

$$\text{AGM}(\sqrt{1 + 8\epsilon^2 - \sqrt{1 - 16\epsilon^2}}, \sqrt{1 + 8\epsilon^2 + \sqrt{1 - 16\epsilon^2}}) = \sqrt{2}/2,$$

which can be solved by bisection since $\text{AGM}(\dots)|_{\epsilon=0} = 0$ and $\text{AGM}(\dots)|_{\epsilon=1/4} = \sqrt{6}/2$.

PROBLEM 7

Question: Let A be the $20,000 \times 20,000$ matrix whose entries are zero everywhere except for the primes $2, 3, 5, 7, \dots, 224737$ along the main diagonal and the number 1 in all the positions a_{ij} with $|i - j| = 1, 2, 4, 8, \dots, 16384$. What is the $(1, 1)$ entry of A^{-1} ?

Answer (12 significant digits): 0.725078346268e+0

Method: The $(1, 1)$ entry under consideration is given as x_1 of the solution x of the linear system $Ax = e_1$, where $e_1 = (1, 0, 0, \dots, 0)^T$. The diagonally preconditioned conjugate gradient method yields a reasonably small a posteriori estimate of the componentwise backward error. I used LAPACK's forward error estimate [7, (7.27)] to decide about the accuracy of the solution. Here, I relied on the Higham-Hager 1-norm estimator [7, Algorithm 14.4] as implemented in Matlab's `normest1` command, using again the diagonally preconditioned CG iteration to supply this command with the appropriate `AFUN` routine.

PROBLEM 8

Question: 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 flows into the plate according to $u_t = \Delta u$. When does the temperature reach $u = 1$ at the center of the plate?

Answer (12 significant digits): 0.424011387034e+0

Method: Because of $\lim_{t \rightarrow \infty} u(t, 0, 0) = 5 \cdot 0.25 = 1.25$, the boundary value $u = 5$ along one of the sides is the "simplest" one making the question meaningful. Solving the initial boundary value problem by Fourier series in the x - y -coordinates yields

$$\begin{aligned} u(t, 0, 0) &= \frac{5}{4} - 10 \sum_{j=0}^{\infty} \sum_{k=0}^{\infty} (-1)^{j+k} \frac{\lambda_j e^{-(\lambda_j^2 + \lambda_k^2)t}}{\lambda_k (\lambda_j^2 + \lambda_k^2)} \\ &= \frac{5}{4} - 5 \sum_{j,k=0}^{\infty} (-1)^{j+k} \left(\frac{\lambda_k}{\lambda_j} + \frac{\lambda_j}{\lambda_k} \right) \cdot \frac{e^{-(\lambda_j^2 + \lambda_k^2)t}}{\lambda_j^2 + \lambda_k^2} = \frac{5}{4} - 5 \left(\sum_{k=0}^{\infty} \frac{(-1)^k e^{-\lambda_k^2 t}}{\lambda_k} \right)^2, \end{aligned}$$

with

$$\lambda_k = \frac{2k+1}{2} \pi.$$

The solution of $u(t, 0, 0) = 1$ is easily located at $t \geq 0.4$. For these values of t the sum converges extremely rapidly, e.g., only *three* terms are sufficient to get 16 correct digits. The solution can now be determined by a simple bisection to high accuracy.

PROBLEM 9

Question: 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 of $\alpha \in [0, 5]$ at which $I(\alpha)$ achieves its maximum?

Answer (12 significant digits): 0.785933674350e+0

Method: As in Problem 1 the integral is transformed to an alternating series, the convergence of which is accelerated by the linear method of Cohen, Rodriguez Villegas and Zagier [3]. The x^α -singularity at $x = 0$ is easily dealt with by the




FIGURE 2. A plot of $I(\alpha)$ [blue] and $I'(\alpha)$ [green]

adaptive quadrature as provided by, e.g., Matlab's `quadl` command. A plot of the function $\alpha \mapsto I(\alpha)$ shows that the maximum is achieved somewhere in $[0.5, 1]$. A bisection is used to find the zero of the derivative $I'(\alpha)$ within this interval. The derivative is of the form

$$I'(\alpha) = \int_0^2 (\dots) \sin(\alpha/(2-x)) dx + \int_0^2 (\dots) \cos(\alpha/(2-x)) dx$$

and can be computed to high accuracy by applying the method introduced in Problem 1 to each of the two terms.

Alternative Method: The integral

$$\int_0^2 x^\alpha \sin(\alpha/(2-x)) dx = \int_0^2 (2-x)^\alpha \sin(\alpha/x) dx$$

is the Euler integral transform of a function which can be expressed in terms of Meijer's G -function [9, 07.34.03.0055.01],

$$\sin(\alpha/x) = \sqrt{\pi} G_{2,0}^{0,1} \left(\frac{4x^2}{\alpha^2} \middle| \frac{1}{2}, 1 \right).$$

The integral has therefore itself a closed form expression in terms of Meijer's G -function [9, 07.34.21.0084.01]:

$$\int_0^2 (2-x)^\alpha \sin(\alpha/x) dx = \sqrt{\pi} \Gamma(\alpha+1) G_{2,4}^{3,0} \left(\frac{\alpha^2}{16} \middle| \frac{\alpha+2}{2}, \frac{\alpha+3}{2} \right)$$

Mathematica and *Maple* both provide a `MeijerG` command allowing to calculate Meijer's G -function efficiently to high accuracy. A routine for minimization without derivatives can be used to solve the problem. Because of numerical differentiation,

however, obtaining 10 significant digits requires extended precision of about at least twice as many digits.

PROBLEM 10

Question: 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 one of the ends rather than one of the sides?

Answer (12 significant digits): 0.383758797925e-6

Method: Let $p(x)$ be the probability that a particle starting at the point x of the $a \times b$ rectangle R , $a \leq b$, hits Γ_a (the two “ends”, i.e., the edges of length a) rather than Γ_b (the two “sides”, i.e., the edges of length b). By [6, Theorem 13.7(5)] the function p is the solution of the elliptic boundary value problem on R

$$\Delta p = 0, \quad p|_{\Gamma_a} = 1, \quad p|_{\Gamma_b} = 0.$$

Calculating a Fourier series solution of this boundary value problem finally gives the alternating series

$$p(x_{\text{center}}) = \sum_{k=0}^{\infty} \frac{4(-1)^k}{\pi(2k+1) \cosh\left(\frac{(2k+1)\pi b}{a}\right)}.$$

For $b/a = 10$ this series converges extremely fast, the first term gives an approximation which is accurate to at least 14 digits:

$$\left| p(x_{\text{center}}) - \underbrace{\frac{4}{\pi \cosh(5\pi)}}_{=3.83758797925 \cdot 10^{-7}} \right| \leq \frac{4}{3\pi \cosh(15\pi)} \approx 2.9 \cdot 10^{-21}.$$

There is a nifty closed expression of the series in terms of the modular λ -function,

$$p(x_{\text{center}}) = \frac{2}{\pi} \arcsin \sqrt{\lambda(10i)},$$

which follows easily from [1, (3.2.29)][16, (T1.10)]. Using the theory of *singular moduli* for the λ -function [1, p. 139], this turns out to be equal to the “elementary” expression²

$$p(x_{\text{center}}) = \frac{2}{\pi} \arcsin k_{100} = \frac{2}{\pi} \arcsin \frac{\sqrt{2} - (\sqrt{5} - 2)(3 + 2 \cdot 5^{1/4})}{\sqrt{2} + (\sqrt{5} - 2)(3 + 2 \cdot 5^{1/4})}.$$

REFERENCES

- [1] Jonathan M. Borwein and Peter B. Borwein, *Pi and the AGM*, John Wiley & Sons Inc., New York, 1987. MR 89a:11134
- [2] Claude Brezinski and Michela Redivo Zaglia, *Extrapolation methods*, North-Holland Publishing Co., Amsterdam, 1991, Theory and practice, With 1 IBM-PC floppy disk (5.25 inch). MR 93d:65001
- [3] Henri Cohen, Fernando Rodriguez Villegas, and Don Zagier, *Convergence acceleration of alternating series*, Experiment. Math. **9** (2000), no. 1, 3–12. MR 2001m:11222
- [4] R. M. Corless, G. H. Gonnet, D. E. G. Hare, D. J. Jeffrey, and D. E. Knuth, *On the Lambert W function*, Adv. Comput. Math. **5** (1996), no. 4, 329–359. MR 98j:33015

²Thanks to J. Boersma for pointing out this relation.

- [5] Paul Godfrey, *Matlab implementation of the complex Gamma and DiGamma function*, <http://my.fit.edu/~gabdo/paulbio.html>, 2001.
- [6] G. R. Grimmett and D. R. Stirzaker, *Probability and random processes*, 2nd ed., Oxford University Press, New York, 1992. MR 93m:60002
- [7] Nicholas J. Higham, *Accuracy and stability of numerical algorithms*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1996. MR 97a:65047
- [8] Naoki Osada, *An acceleration theorem for the q -algorithm*, Numer. Math. **73** (1996), no. 4, 521–531. MR 97c:65013
- [9] Wolfram Research, *Wolfram Research's Mathematical Functions: Meijer's G-function*, <http://functions.wolfram.com/PDF/MeijerG.pdf>.
- [10] Nici Schraudolph, *A matlab implementation of the Lambert W-function*, <http://www.inf.ethz.ch/~schraudo/pubs/W.m>.
- [11] M. N. Senhadji, *On condition numbers of the Shanks transformation*, J. Comput. Appl. Math. **135** (2001), no. 1, 41–61. MR 2002f:40003
- [12] Frank Spitzer, *Principles of random walk*, second ed., Springer-Verlag, New York, 1976, Graduate Texts in Mathematics, Vol. 34. MR 52 #9383
- [13] Ping Tak Peter Tang, *A fast algorithm for linear complex Chebyshev approximations*, Math. Comp. **51** (1988), no. 184, 721–739. MR 89j:30054
- [14] Nick Trefethen, *A Hundred-dollar, Hundred-digit Challenge*, SIAM News **35** (2002), no. 1, 1.
- [15] ———, *The \$100, 100-digit Challenge. Chastened Challenge Sponsor: "I Misjudged"*, SIAM News **35** (2002), no. 6, 1–3.
- [16] I. J. Zucker, *Some infinite series of exponential and hyperbolic functions*, SIAM J. Math. Anal. **15** (1984), no. 2, 406–413. MR 85f:33002

TECHNICAL UNIVERSITY OF MUNICH, 80290 MUNICH, GERMANY

E-mail address: bornemann@ma.tum.de

URL: <http://www.ma.tum.de/m3/bornemann/>