100 Digit Challenge – Results and Methods

Peter Robinson[1], March 2002

This notes sets out the answers obtained for the 100-digit challenge (SIAM News, Volume 35, Number 1) and gives the algorithms used.

Answers (to 10 digits)

Q1	0.3233674317
Q2	0.9952629194
Q3	1.274224153
Q4	-3.306868647
Q5	0.2143352346
Q6	0.06191395447
Q7	0.7250783463
Q8	0.4240113870
Q9	0.7859336744
Q10	$3.837587979 \times 10^{-7}$

[1] Dr Peter Robinson,
Quintessa Limited,
Dalton House, Newtown Road,
Henley-on-Thames, Oxfordshire RG9 1HG.
PeterRobinson@quintessa.org

General Comments

Most of the solutions were calculated using small C++ codes implementing the algorithms described below. In cases where the standard double precision was inadequate, use was made of the C++ extended precision routines as described by Hida et al[2].

In each case, the approach taken was to try to manipulate the problem so that a simple numerical calculation could be undertaken in a practical time (minutes rather than hours of computation on a basic laptop PC) and with adequate precision.

The other frequently used item was a copy of Abramowitz and Stegun[3].

The results are an individual effort, but thanks are due to my colleague Steve Benbow for introducing the challenge and for some useful discussions on possible approaches.

---

[2] Yozo Hida, Xiaoye S Li, David H Bailey, "Quad-double arithmetic: Algorithms, implementation, and application," Lawrence Berkeley National Laboratory technical report LBL-46996 (2000). Condensed version in 15th IEEE Symposium on Computer Arithmetic, Vail, Colorado, June 2001.

[3] M Abramowitz and I Stegun, "Handbook of Mathematical Functions", Dover Publications.

Q1

What is $\lim_{\varepsilon\to 0}\int_{\varepsilon}^{1} x^{-1}\cos(x^{-1}\log x)\,dx$?

The log is taken to be the natural logarithm and is denoted ln below.

Substitute u = 1/x, so the answer is equal to

$$\int_1^\infty \frac{\cos(u\ln u)}{u}\,du \,.$$

To increase the power of *u* on the denominator, multiply the numerator and denominator by $1+\ln u$ and integrate by parts, noting that the derivative of $\sin(u\ln u)$ is $(1+\ln u)\cos(u\ln u)$. This can be repeated several times. Each time the power on the denominator goes up by one but the coefficients get larger and a constant term that must be cancelled by the integral appears.

The fourth power form is

$$2-\int_1^\infty \frac{\sin(u\ln u)}{u^4}\, f_4(u)\,du \,,$$

where

$$f_4(u)=\frac{6}{(1+\ln u)^3}+\frac{18}{(1+\ln u)^4}+\frac{25}{(1+\ln u)^5}+\frac{15}{(1+\ln u)^6}\,.$$

This version can give 8 or 9 digits with normal double precision but is too slow for the extended precision calculation.

So, the sixth power version was used:

$$\int_1^\infty \frac{\sin(u\ln u)}{u^6}\, f_6(u)\,du-62\,,$$

where

$$f_6(u)=\frac{120}{(1+\ln u)^5}+\frac{600}{(1+\ln u)^6}+\frac{1526}{(1+\ln u)^7}+\frac{2380}{(1+\ln u)^8}+\frac{2205}{(1+\ln u)^9}+\frac{945}{(1+\ln u)^{10}}\,.$$

Note that the coefficients can be calculated from simple recurrence relationships.

Using the 6th power version and double-double precision the required result can be obtained with a simple trapezoidal scheme using δu=1.25×10^{-7} and integration up to *u*=400. This result is confirmed by considering the convergence behaviour with step sizes powers of 2 times larger. It is also confirmed by an extended precision calculation.

δu	Result	Change from previous (last 4 digits)
4e-6	0.323367421310	
2e-6	0.323367429086	7776
1e-6	0.323367431029	1943 (=1/4 of previous)
5e-7	0.323367431516	487 (=1/4 of previous)
2.5e-7	0.323367431637	121 (=1/4 of previous)
1.25e-7	0.323367431667	30

Thus, the changes scale with du^2 it is clear that final result is **0.3233674317**.

This is confirmed in the extended precision calculation, by allowing δu to increase as u changes (times 1+1e-6 each time), then a start of δu=5e-9 gives 0.323367431688.

Q2

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 $t = 10$?

This problem requires the extended precision calculation. It is very sensitive to the details (even calculating 1/3 in double precision instead of extended precision changes the 7^{th} significant figure).

The calculated distance was **0.9952629194**.

To check that the reflection calculation is correct, the calculated path was input into Excel and the following figure drawn.

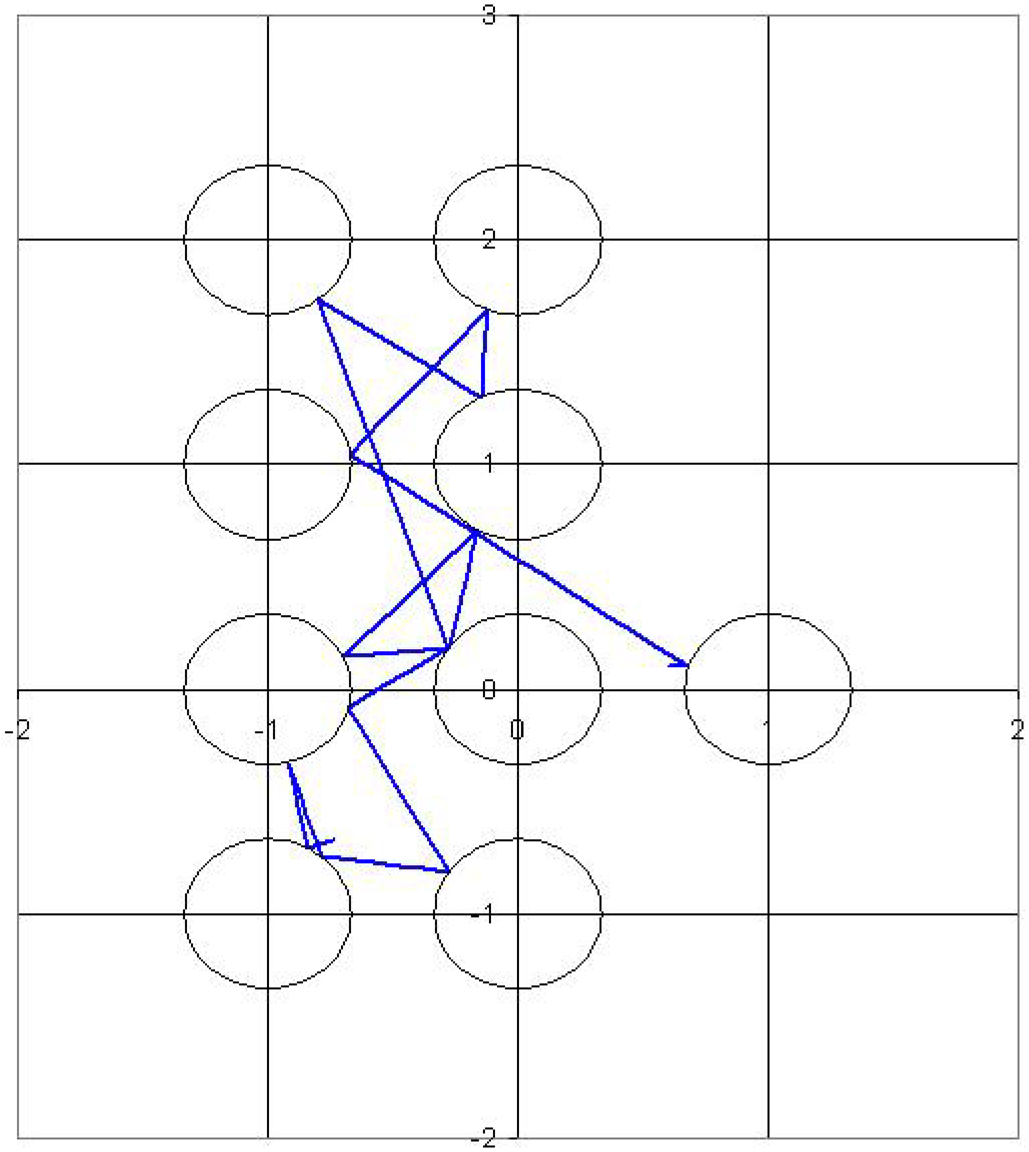

Q3

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 $\|A\|$?

If the infinite vector $\boldsymbol{x}$ is in ℓ^2 then $\|x\| = \sqrt{\sum_{k=1}^{\infty}(x_k)^2}$ is finite.

$\|A\|$ is defined by

$$\|A\| = \max\{\|Ax\| : x \in \ell^2, \|x\| = 1\}.$$

In matrix terms, $\|x\| = \sqrt{x^T x}$ and $\|A\| = \max\left[\sqrt{x^T A^T A x} : x^T x = 1\right]$.

This makes the required norm equal to square-root of the largest eigenvalue of $A^T A$ which can be calculated (at least for finite A) by simply iteratively multiplying a start vector by the matrix and normalising the resulting vector (the Power Method).

To tackle the infinite problem, consider an N-by-N matrix from the first N rows and columns and see how the calculated norm behaves with N.

Since the code was written in C++, it is convenient for the first entry to be a_{00}. With this approach, the general term in A is given by $a_{mn} = \dfrac{1}{(m+n+1)(m+n)/2+m+1}$. The matrix $A^T A$ is never formed (neither is A), but $A^T Ax$ is directly calculated.

The following results were obtained, with only 3 or 4 iterations being necessary (the next eigenvalue must be much smaller).

N	$\|A_N\|$	$\|A_N\| - \|A_{N/2}\|$
1000	1.27422414813	-
2000	1.27422415223	0.00000000410
4000	1.27422415275	0.00000000052
8000	1.27422415281	0.00000000006
16000	1.27422415282	0.00000000001

The change falls off with N^{-3}, so the final result is **1.274224153**.

Q4

What is the global minimum of the function
$\exp(\sin(50x)) + \sin(60e^y) + \sin(70\sin(x)) + \sin(\sin(80y)) - \sin(10(x+y)) + \frac{1}{4}(x^2 + y^2)$?

The terms are each bounded below individually, so the smallest cannot be less than

$e^{-1} + (-1) + (-1) - \sin(1) - (+1) + (0)$ = -3.47359.

Also the value at x = 0, y = 0 is 0.69519, so the square term restricts x and y to be within a circle around the origin of radius about 1.7.

A simple initial search was performed on a lattice of steps of 0.0005. This reveals find points where the value is less than –3.3. The derivative of the function in the region means that no very localized minima will be present. Thus, the search can be concentrated on the regions where a point is near the minimum. Two regions are candidates – near x=-0.395, y=-0.094 and near x=-0.03, y=0.21.

Looking at these regions in turn at 100 times higher resolution shows that the first region has a local minimum just below –3.2. The second region has many points below –3.3.

The search range is then refined step by step until the required accuracy is reached.

This gives a result of –3.30686864748.

Given that the last digit is close to be rounded up, this is checked with an extended precision calculation, which gives -3.3068686474752, the same to 10 figures: **–3.306868647**.

Q5

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 $\|f - p\|_\infty$?

Abramowitz and Stegun give the first 26 terms of the Taylor series for $1/\Gamma(z)$, which can be used directly. For higher precision calculation, the Taylor series can be regenerated by expanding the equality

$$1/\Gamma(z) = z.\exp\left(\gamma z - \sum_{k=2}^{\infty} \frac{(-1)^k \zeta(k) z^k}{k}\right).$$

Since $f(z) - p(z)$ is continuous and differentiable on the unit disk, the Maximum Modulus theorem applies, so the supremum occurs on the circumference of the disk.

Since $f(\bar{z}) = \bar{f}(z)$ the coefficients of the polynomial must be real.

The Taylor series coefficients are a good starting guess. By taking 1000 points, $\theta \in [0, \pi]$ around the circumference with $z = e^{i\theta}$, an initial survey of possible coefficients reveals that the minimum supremum norm occurs near $0.0055+1.102z+0.625z^2-0.603z^3$.

The approach described by Bailly and Thiran[4] can then be followed.

This provides a transformation from the complex problem to an equivalent real problem once the location of the maxima are approximately known. From the survey of polynomial, the maxima are near θ *=1.4, 2,26* and π. Thus, there are 5 maxima for the full disk, and the problem is of type F_2 as described by Bailly and Thiran.

The problem is therefore equivalent to finding the minimum supremum norm of $g(\theta) - q(\theta)$ where $g(\theta) = \operatorname{Im}\left\{e^{-3i\theta/2} f(e^{i\theta})\right\}$ with $q(\theta) = a\sin\left(\frac{\theta}{2}\right) + b\sin\left(\frac{3\theta}{2}\right)$ and $\theta \in [0, \pi]$.

This is easily solved by application of the Second Remez Algorithm.

This converges rapidly, and allows the results to be accurately obtained (when the interval between the points is decreased). The result was confirmed in extended precision.

The required norm is 0.21433523459 with the extremes occurring at θ *=1.40319916, 2,2623774549* and π. This gives *a = -0.39454993655, b =-0.60888517114*

The 10 figure result is therefore **0.2143352346**.

---

[4] B Le Bailley and J P Thiran, "Comuting Complex Polynomial Chebychev Approximants on the Unit Circle by the Real Remez Algorithm", SIAM J Numer Anal, Vol 36, No 6, pp1858-1877.

Q6

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?

This can be approached by solving the lattice problem directly.

Start with a unit amount at (0,0) and transfer the appropriate fractions to the neighbours. Repeat for each cell on the next step. Remove anything at (0,0) and keep a running sum. The lattice must be large enough so that anything crossing the boundaries would never have returned. This can be tested by extending the size to check it makes no difference.

The process can be speeded up by reflecting anything that wants to go south of the origin back up and so only using half of the full lattice.

The required value of ε can then be found by a simple iterative scheme. A lattice extending to (100,100) is adequate. A few thousand steps are required for everything to return to the origin or leave across the eastern boundary. The required value is 0.061913954473, giving the 10 digit result **0.06191395447**.

Q7

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 positions a_{ij} with $|i-j|$ = 1, 2, 4, 8, … 16834. What is the (1,1) entry of A^{-1}?

The (1,1) entry of A^{-1} is also the first entry in x where $Ax=b$ with b = (1,0,0, … 0).

This can be calculated with a simple iterative scheme.

Start with all entries in x equal zero.

Update each entry in x in turn to make the corresponding row of the equation satisfied exactly. Iterate until convergence.

This scheme converges rapidly (20 iterations), giving the result 0.725078346268401. This was calculated in extended precision.

The primes were calculated with a simple sieve method.

The 10-digit result is therefore **0.7250783463**.

Q8

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, with the other sides held at u = 0, and heat flows into the plate according to $u_t = \Delta u$. When does the temperature reach u = 1 at the centre.

This problem is tackled via a Laplace transform. Take the x = 1 side to have u = 5.

The transformed problem is

$$s\bar{u} = \Delta\bar{u}$$

where the bar denotes the transform. Then solution with the given boundary conditions is simply

$$\bar{u}(x,y) = \frac{5}{s}\sum_{n=0}^{\infty}(-1)^n \frac{2\cos(\beta_n y)\sinh(\alpha_n(1+x))}{\beta_n \sinh(2\alpha_n)},$$

with

$$\begin{aligned}\beta_n &= (n+1/2)\pi \\ \alpha_n &= \sqrt{s+\beta_n^{\,2}}\end{aligned},$$

and so

$$\bar{u}(0,0) = \frac{5}{s}\sum_{n=0}^{\infty}(-1)^n \frac{1}{\beta_n \cosh(\alpha_n)}.$$

This can be inverted using Talbot's algorithm[5], which gives high precision results with a small number of contour points[6] (128 were used here).

The exact time can be found be repeated bisection.

The number of terms of the infinite sum needed is small, as the high frequency modes are only relevant very near the x = 1 boundary. Here, 200 terms were found to be adequate.

The required result is found to be t = 0.4240113870337, giving the 10-digit result **0.4240113870**.

---

[5] A Talbot, "The accurate numerical inversion of Laplace transforms", *J. Inst. Math. Applic*. **23**, 97-120 (1979).

[6] P C Robinson and P R Maul, "Some experience with the numerical inversion of Laplace transforms", *Math. Engng Ind*., Vol 3, No 2, pp. 111-131 (1991).

Q9

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 value?

Substituting $v=\alpha/(2-x)$, i.e. $x=2-\alpha/v$, gives

$$I(\alpha)=[2+\sin(10\alpha)]2^\alpha\alpha\int_{\alpha/2}^{\infty}(1-\alpha/2v)^\alpha\frac{\sin v}{v^2}dv$$

This can be roughly calculated to find where the maximum is likely to be. This shows that it occurs with α near $\pi/4$.

The approach taken to accurate evaluate the integral is to split it into two parts. Then the integral to $v=\pi/2$ is calculated by quadrature (Simpson's rule). The integral from $\pi/2$ to infinity is evaluated as an infinite sum as follows.

Expand the bracket as a power series in 1/v. This gives the integral as

$$\int_{\pi/2}^{\infty}(1-\alpha/2v)^\alpha\frac{\sin v}{v^2}dv=\sum_{k=0}^{\infty}\left(\frac{-\alpha}{2}\right)^k\frac{\alpha(\alpha-1)...(\alpha-k+1)}{k!}\int_{\pi/2}^{\infty}\frac{\sin v}{v^{2+k}}dv$$

Then the sin integrals can be evaluated from recurrence relationships.

Let

$$s_n=\int_{\pi/2}^{\infty}\frac{\sin v}{v^n}dv$$

$$c_n=\int_{\pi/2}^{\infty}\frac{\cos v}{v^n}dv$$

then

$$s_n=\frac{-1}{(n-1)(\pi/2)^{n-1}}+\frac{c_{n-1}}{n-1}\qquad n>1$$

$$c_n=-\frac{s_{n-1}}{n-1}\qquad n>1$$

and

$$s_1=\frac{\pi}{2}-Si(\pi/2)$$

$c_1 = -Ci(\pi/2)$

where Si and Ci are Sine and Cosine integrals and are defined by Abramowitz and Stegun. Taylor series expansions are give which converge quickly.

$$Si(x) = \sum_{r=0}^{\infty} \frac{(-1)^r x^{2r+1}}{(2r+1)(2r+1)!}$$

$$Ci(x) = \gamma + \ln x + \sum_{r=1}^{\infty} \frac{(-1)^r x^{2r}}{2r(2r)!} .$$

The maximum of *I(α)* can then be located iteratively.

The precise location of the maximum is very hard to find, with changes in the 11th significant figure of α only changing the 24th figure in *I(α)*. Moreover the maximum is very close to rounding the 10th figure up or down. Finally, the value found was 0.7859336743504 which rounds to **0.7859336744**. The maximum itself is 3.03373258648549….

This used 80 million points in a Simpson rule quadrature scheme for the integral to $\pi/2$. The sine and cosine integrals used 250 terms in the sum, and 500 terms in the power series in 1/*v* were used.

Q10

A particle at the centre 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?

An initial lattice calculation suggests that the answer is of the order 4e-7, but cannot be found accurately with this approach.

The problem can be re-posed as a 2D continuum diffusion problem. Given a unit point amount of material at the centre of a 10 by 1 rectangle at time $t=0$, with zero-concentration conditions on all four sides, what fraction of the unit amount leaves across the short boundaries?

The zero-concentration conditions mean that nothing returns across a boundary once it has hit.

This can further be simplified by separating the two coordinates: the spreading across the direction of interest is irrelevant (except inasmuch as it removes material from the system). This is equivalent to writing the two-dimensional solution as a product of two one-dimensional solutions.

Let $f_L(t)$ be the rate at which material reaches a zero-concentration boundary a distance L from the initial position in a 1D problem. Let the amount that has left up to time t be denoted $F_L(t)$, so

$$F_L(t) = \int_0^t f_L(\tau) d\tau .$$

Then, the fraction, β_{10}, that is required is given by integrating the rate across the short boundary, scaled by the fraction that has not already crossed the long boundary

$$\beta_{10} = \int_0^\infty f_{10L}(t)(1 - F_L(t)) dt .$$

Thus, to evaluate the result requires evaluation of the rate at which material leaves the zero concentration boundaries in a 1D diffusion calculation.

$$u_t = u_{xx}$$

$$u(x,0) = \delta(x)$$

$$u(\pm L, t) = 0$$

$$f_L(t) = -2u_x(L,t) .$$

The solution is simply given by

$$u(x,t) = \frac{1}{L} \sum_{n=0}^{\infty} \cos(\alpha_n x) e^{-\alpha_n{}^2 t}$$

where

$$\alpha_n = (n+1/2)\frac{\pi}{L}.$$

So,

$$f_L(t) = \frac{2}{L}\sum_{n=0}^{\infty}(-1)^n \alpha_n e^{-\alpha_n{}^2 t}.$$

This gives

$$F_L(t) = \frac{2}{L}\sum_{n=0}^{\infty}(-1)^n \alpha_n \frac{1-e^{-\alpha_n{}^2 t}}{\alpha_n{}^2}$$

and so

$$1-F_L(t) = \frac{2}{L}\sum_{n=0}^{\infty}\frac{(-1)^n}{\alpha_n}e^{-\alpha_n{}^2 t}$$

Noting that the α values are proportional to *1/L*, gives

$$\beta_Q = \frac{16}{\pi^2}\sum_{n=0}^{\infty}\sum_{m=0}^{\infty}(-1)^{m+n}\frac{2m+1}{2n+1}\frac{1}{(2m+1)^2+Q^2(2n+1)^2}$$

for the general case of a long side *Q* times the length of the short side.

In practice, this is written as

$$\beta_Q = \frac{16}{\pi^2}\sum_{n=0}^{\infty}\frac{(-1)^n}{2n+1}\left\{\sum_{m=0}^{\infty}(-1)^m\frac{2m+1}{(2m+1)^2+Q^2(2n+1)^2}\right\}.$$

The sum over the *n* terms converges relatively quickly whereas the *m* sum converges slowly.

Combining consecutive pairs of *m* gives satisfactory convergence, and the contribution from terms other than *n=0* is small and falls quickly (in fact they seem to tend to vanish as the sum over *m* is taken to larger numbers of terms).

It is possible to take advantage of the alternating signs in the terms to obtain good approximations to the residual sums, saving some computational effort.

The extended precision was used to guard against rounding errors – a sum to *m*=200,000,000 was used, but only the n=0 and n=1 terms were needed. The result with normal double precision was in fact identical to 10 figures.

The result is **3.837587979e-7**.