Asymptotic Approximation

An approximation like the one on the "Arctan sum" page gives the first few terms of an asymptotic series. The series usually diverges for a given n, but using a fixed number of terms the formula becomes more and more accurate as n increases.

Generate a formula that approximates S_n for this sequence:

$$S_n = \sum_{k=1}^{n} \sqrt{k}$$

Start with a few moderately large values of n and look for patterns.

Let $\mathrm{f1}(x) = \sqrt{x}$ and then look at S_n for $n = 10^4$, 10^6, and 10^8.

f1: 1, func, $\sqrt{\mathrm{x}}$

30, sci, 7, func, 1, enter, e4, enter, 1, enter, 1, sum, 11, sto	6.66716459197108355926683982802e+5
1, enter, e6, enter, 1, enter, 1, sum, 12, sto	6.66667166458822108355978766795e+8
1, enter, e8, enter, 1, enter, 1, sum, 13, sto	6.66666671666458784608355978767e+11

It seems like $S_n \approx (2/3)n^{3/2}$, which we might have guessed from the fact that the integral of $x^{1/2}$ is $(2/3)x^{3/2}$, and the integral gives a rough approximation for the sum. Next, subtract $(2/3)n^{3/2}$ in each of these three cases.

1, func, 11, rcl, e4, enter, 1.5, $\mathrm{y^x}$, 2, $*$, 3, /, $-$, 21, sto	4.97925304416892600173161351149e+1
12, rcl, e6, enter, 1.5, $\mathrm{y^x}$, 2, $*$, 3, /, $-$, 22, sto	4.99792155441689312100128526608e+2
13, rcl, e8, enter, 1.5, $\mathrm{y^x}$, 2, $*$, 3, /, $-$, 23, sto	4.99979211794168931210064935473e+3

These look close to $\sqrt{n}/2$, which would give 5e+1, 5e+2, and 5e+3.
Subtract $\sqrt{n}/2$ in each of these three cases.

30, fix, 21, rcl, e4, $\sqrt{\mathrm{x}}$, 2, /, $-$, 31, sto	$-0.207469558310739982683864885120$
22, rcl, e6, $\sqrt{\mathrm{x}}$, 2, /, $-$, 32, sto	$-0.207844558310687899871473392064$
23, rcl, e8, $\sqrt{\mathrm{x}}$, 2, /, $-$, 33, sto	$-0.207882058310687899350645267064$

The good news is that this sequence now seems to be converging. The bad news is that the limit isn't an easily recognizable fraction like 2/3 or 1/2.

To continue our asymptotic analysis, we would like to have this constant to fairly high precision. Try using $n = 10^{20}$ terms and computing (sum $- (2/3)n^{3/2} - \sqrt{n}/2$):

$-0.207886224973191232683973391800$

We might estimate that the three numbers starting in register 31 are correct to 2, 3, and 4 digits, so if that pattern continues this last value might be good to about 10 digits. But the sum key's estimated relative error for the $n = 10^{20}$ sum was about 10^{-28}.

The sum itself is about $(2/3)\,10^{30}$, with 30 digits left of the decimal. The sum key estimates that about 28 significant digits may be correct, so we can't be confident of many of the digits in the $-0.207...$ constant.

When n gets too big, S_n is so large that cancellation error keeps us from being able to get an accurate value for $S_n - (2/3)n^{3/2} - \sqrt{n}/2$. See the "Cancellation error" page for more on cancellation.

To get more digits for this constant, define a new sequence by subtracting off the part we think we know.

$$D_n = \left(\sum_{k=1}^{n} \sqrt{k}\right) - \frac{2n^{3/2}}{3} - \frac{\sqrt{n}}{2} = \left(\sum_{k=1}^{n} \sqrt{k}\right) - \left(\frac{2n}{3} + \frac{1}{2}\right)\sqrt{n} = \left(\sum_{k=1}^{n} \sqrt{k}\right) - \left(\frac{4n+3}{6}\right)\sqrt{n}$$

The form of one term subtracted from a divergent series is like example 3 from the "Extrapolation" page. Instead of extrapolating, this time we will find a sneaky way to express D_n as the n^{th} partial sum of a convergent series, without any extra terms on the outside.

Let $g(n) = \left(\frac{4n+3}{6}\right)\sqrt{n}$ be the term to be subtracted from the series.

We can turn it into a single convergent series by moving the $g(n)$ from the outside to the inside of the sum.

$$\left(\sum_{k=1}^{n} \sqrt{k}\right) - \left(\frac{4n+3}{6}\right)\sqrt{n} = \left(\sqrt{1} - g(1) + g(0)\right) + \left(\sqrt{2} - g(2) + g(1)\right) + \cdots + \left(\sqrt{n} - g(n) + g(n-1)\right)$$

$$= \sum_{k=1}^{n} \left(\sqrt{k} - g(k) + g(k-1)\right)$$

This is a telescoping series, where the $g(k)$ in each term of the sum cancels with the $g(k-1)$ in the next term when k is one higher. The only two bits that don't cancel are the $g(0)$ in the first term, which is zero already, and the $g(n)$ in the last term, which is the part outside the sum in the original version.

Before we sum this new convergent series, we will want to deal with the cancellation error in each term. As k gets bigger, $g(k)$ and $g(k-1)$ are both about $k^{3/2}$ in size, but the final k^{th} term is between -1 and 0. That means much precision would be lost in computing the terms from this formula when k is large.

We can algebraically simplify this formula to eliminate the cancellation.

$$\sqrt{k} - g(k) + g(k-1) = \sqrt{k} - \left(\frac{4k+3}{6}\right)\sqrt{k} + \left(\frac{4k-1}{6}\right)\sqrt{k-1}$$

$$= \frac{1}{6}\Big(-(4k-3)\sqrt{k} + (4k-1)\sqrt{k-1} \Big)$$

$$= \frac{1}{6}\Big(-(4k-3)\sqrt{k} + (4k-1)\sqrt{k-1} \Big) \frac{\Big((4k-3)\sqrt{k} + (4k-1)\sqrt{k-1}\Big)}{\Big((4k-3)\sqrt{k} + (4k-1)\sqrt{k-1}\Big)}$$

$$= \frac{-1}{6\left((4k-3)\sqrt{k} + (4k-1)\sqrt{k-1}\,\right)}$$

This version of the formula for the k^{th} term has no cancellation, so it will be stable for the summation. Define it to be f2 and do the infinite sum. Save the result in register 0.

f2: 1, func, 2, sto, 4, ∗, 3, −, 2, rcl, $\sqrt{\text{x}}$, ∗, 2, rcl, 4, ∗, 1, −, 2, rcl, 1, −, $\sqrt{\text{x}}$, ∗, +, 1/x, 6, /, chs

40, fix, 7, func, 1, enter, e9999, enter, 1, enter, 2, sum, 0, sto

−0.20788622497735456601730672539704930222632

Call this number c_0. It is the constant term in the asymptotic expansion. Pressing x↔y shows the sum function's estimated error is about 2e−55, so we should have the constant term to over 50 digits.

It turns out that $c_0 = \zeta(-1/2)$, from the Riemann zeta function, so those of us who are not experts in analytic number theory can be forgiven for not immediately recognizing it.

What comes next? We would like to know what further powers of n occur in the asymptotic expansion, and what coefficients go with them. There is some theory that can help when the function we are summing is simple enough that we can easily find its integral and several of its derivatives.

But when the function is more complicated, like the "Arctan sum" example, we can continue crunching numbers and trying to guess the coefficients as we did in that case.

Here, $f(n) = \sqrt{n}$ is easy to integrate and differentiate, so the theory can help.

This theory comes from the Euler-Maclaurin formula (check Wikipedia to see more details). The B_i values are Bernoulli numbers from screen 3. Here are the first few terms (general form, then specific for this f):

$$\text{term 1} = \int f(n)\, dn = \frac{2\,n^{3/2}}{3}$$

$$\text{term 2} = \frac{f(n)}{2} = \frac{\sqrt{n}}{2}$$

$$\text{term 3} = \frac{B_2\, f'(n)}{2!} = \frac{1}{24\sqrt{n}}$$

$$\text{term 4} = \frac{B_4\, f'''(n)}{4!} = \frac{-1}{1920\, n^{5/2}}$$

$$\text{term 5} = \frac{B_6\, f^{(5)}(n)}{6!} = \frac{1}{9216\, n^{9/2}}$$

$$\text{term 6} = \frac{B_8\, f^{(7)}(n)}{8!} = \frac{-11}{163840\, n^{13/2}}$$

The first two terms confirm what we had already guessed. The Euler-Maclaurin formula doesn't give the constant term directly — there are constants included in all the other terms. So even with this shortcut the constant term often has to be computed separately, as we did above.

When the integer numerators and denominators in the coefficient fractions get big quickly, as they do here, it is harder to guess the proper coefficients and powers of n if we have to do the number crunching method instead of using Euler-Maclaurin.

In either case, it is a good idea to check the formula using a few different values for n. Here is the formula:

$$S_n = \sum_{k=1}^{n} \sqrt{k} \;\approx\; \frac{2n^{3/2}}{3} + \frac{\sqrt{n}}{2} + c_0 + \frac{1}{24\sqrt{n}} - \frac{1}{1920n^{5/2}} + \frac{1}{9216n^{9/2}} - \frac{11}{163840n^{13/2}}$$

$$= c_0 + \sqrt{n}\left(\frac{2n}{3} + \frac{1}{2} + \frac{1}{24n} - \frac{1}{1920n^3} + \frac{1}{9216n^5} - \frac{11}{163840n^7}\right)$$

Let f3(n) evaluate this formula.

f3: 1, func, 3, sto, 2, $*$, 3, /, 0.5, +, 3, rcl, 24, $*$, 1/x, +, 3, rcl, 3, $\mathrm{y^x}$, 1920, $*$, 1/x, $-$, 3, rcl, 5, $\mathrm{y^x}$, 9216, $*$, 1/x, +, 3, rcl, 7, $\mathrm{y^x}$, 163840, $*$, 1/x, 11, $*$, $-$, 3, rcl, $\sqrt{\mathrm{x}}$, $*$, 0, rcl, +

f4(n) will check the formula by using the sum key directly and then computing the relative error between that and the formula.

f4: 4, sto, 10, sci, 7, func, 1, x↔y, 1, enter, 1, sum, 41, sto,
4, rcl, 3, $\mathrm{f_n}$, 42, sto, 41, rcl, $-$, 41, rcl, /, 3, func, |x|

Our prediction is that the absolute error in the formula should be proportional to $1/n^{17/2}$, since that would be the power of n on the next term after the last one we used.

The size of the number produced by the formula is about $n^{3/2}$, so dividing the absolute error by that gives $1/n^{10}$ as the predicted relative error. For example, in the table below, $n = 100$ should have a relative error of about $1/(10^2)^{10} = 10^{-20}$ or less.

n	relative error
10^1	1.140196559e−14
10^2	1.230669339e−24
10^3	1.238857073e−34
10^4	1.239683995e−44
10^5	1.238022431e−54

If we had derived this formula for use in a program with 16 significant digit accuracy (64-bit double precision format), then this formula gives full 16 digit accuracy for any n bigger than 15.