Small values in recurrence relations

Tom Womack

October 2005

This started off as an experiment with recurrence relations in Excel. Consider

$$a_n = a_{n-1} - a_{n-3}.$$

The a_i grow quite quickly, but (because of the subtraction) in a much less uniform way than the Fibonacci numbers grow; in particular, if you start $a_0 = 0$, $a_1 = 1$, $a_2 = 2$, you see that a_{64} through a_{66} and a_{68} through a_{70} are all five-figure numbers, but $a_{67} = 295$.

Recurrence relations and polynomials

It's a standard A-level exercise to see what value x has to take on for $a_n = x^n$ to be a solution to a recurrence relation; x has to be a root of the associated polynomial f,

$$a_{n+3} - a_{n+1} + a_n = 0 \equiv x^3 - x + 1 = 0,$$

and (at least if the associated polynomial has unique roots $\theta_1, \ldots, \theta_n$), the series θ_i^n are a basis for the solutions to the relation.

0.1 Some strange, and possibly instructive, graphs

Define

$$b_n = \frac{\sqrt{|a_{n-1}a_{n+1}|}}{|a_n|},$$

the ratio by which a_n differs from the value predicted as the geometric mean of its neighbours, and plot it for a large number of n. You observe that the peaks appear to be very regularly spaced; for example, b_n is noticeably higher than its neighbours for $n = 191, 374, 557, 740\ldots$ a series with entries differing by 183.

With different start conditions, the position of the peaks changes, but the spacing between consecutive peaks does not; changing the recurrence relation changes the spacing.

We could also, knowing that the trend is to have $|a_n| \approx \theta^n$ where θ is the root of the polynomial with largest absolute value, plot $c_n = \log|a_n\theta^{-n}|$ against n. This is somehow a cleaner definition than b_n (since we're not making an arbitrary choice as to the number of neighbours to take), and we get (at suitable scales) a cleaner pattern looking like the superposition of a number of pointed arches.

Contriving to make a point small

Let $r_0(n)$ be the 'impulse response sequence' generated by $a_0 = 1, a_1 = 0, a_2 = 0$, $r_1(n)$ be generated by $a_0 = 0, a_1 = 1, a_2 = 0$, and $r_2(n)$ likewise. Then we certainly have

$$a_n = a_0 r_0(n) + a_1 r_1(n) + a_2 r_2(n).$$

For any particular n, we can easily compute $r_i(n)$; by picking an appropriate λ and reducing the lattice

$$\begin{pmatrix} 1 & 0 & 0 & \lambda r_0(n) \\ 0 & 1 & 0 & \lambda r_1(n) \\ 0 & 0 & 1 & \lambda r_2(n) \end{pmatrix}$$

we can pick $a_0 \ldots a_2$ to make a_n either equal to zero (if we took a large enough λ) or at least unusually small for a of the relevant size (if we took a very small λ). For instance, the sequence beginning $43\ -41\ 69$ has $a_{100} = -9$ while its neighbours are in the tens of millions. The sequence beginning $19\ -112\ -45$ has a_{200} in the tens of millions with neighbours in the hundreds of trillions.

If we take one of these contrived cases, we see echoes further up the sequence at intervals, once again, of 183; so the obvious experiment is to hit a_{183}. This doesn't cause anything interesting to happen; on the other hand, if we hit a_{193}, and use large coefficients in $a_0 \ldots a_2$ to make $a_{193} = -1$, we do see that a_{10} is small in comparison to its neighbours. We also see that the arithmetic in Excel is not capable of distinguishing 1 and 0 for numbers around $a_{192} \approx 5 \times 10^{16}$.

What's really going on?

Seeing an obvious effect, which occurs at the points of an arithmetic progression and with an amplitude that clearly peaks somewhere and drops off in what looks like a $x^{-\epsilon}$ curve, makes me think of resonances; we've got something which is oscillating with a period which is very nearly an integer multiple of $\frac{1}{183}$.

So, let's look more closely at the θ_i. The recurrence relation is only of degree 3, so we could in principle solve it by radicals; on the other hand, we'll end up with quantities involving the cube root of expressions involving $\sqrt{-23}$, which sounds like something to be avoided. We have

$$\begin{aligned} \theta_1 &\approx 0.75488 \\ \theta_2 &\approx -0.87744 - 0.74486i \\ \theta_3 &\approx -0.87744 + 0.74486i \end{aligned}$$

We can ignore θ_1 because its absolute value is less than 1. The argument of θ_3 is about 0.77595π; `magma` takes half a second to compute θ_3/π to a thousand decimal places and to write down its continued fraction

$$0, 1, 3, 2, 6, 3, 25, 1, 1, 7, 1, 3550, \ldots$$

whose convergents turn out to be

$$\frac{3}{4}, \frac{7}{9}, \frac{45}{58}, \frac{142}{183}, \frac{3595}{4633}, \frac{3737}{4816}, \frac{7332}{9449}, \frac{55061}{70959}, \frac{62393}{80408}, \frac{221550211}{285519359} \dots$$

So there's the answer; 183 is the denominator of a convergent in the continued fraction which appears before a relatively-large term, so θ_2^{183} has an imaginary part small in relation to its real part; the expression for a_n must be symmetric in θ_2 and θ_3 because it's a real number, and thus it will have an approximate period of 183. If we compute tens of thousands of terms of the recurrence (which is not completely trivial because the terms get quite large: but c_n is what we're interested in, and it's defined in terms of $\log |a_n|$, so we can work in the logarithmic domain throughout), we do indeed see peaks in c_n at 4464, 13913, 23362, with common difference the convergent 9449.

Polynomial recurrences with small $\max |\theta_i|$

In the long run, the largest contribution to the root of a_n will be given by the root of f of the largest absolute value. So if one wants slowly-growing recurrence relations, it would be nice to find polynomials with integer coefficients (we need to insist they're monic otherwise the problem's poorly-defined, and it's probably helpful to insist on constant coefficient 1) such that $\max |\theta_i|$ is small. I coded this up in `magma` (requiring also that there are at most two non-zero middle coefficients, and those are taken from $\{\pm 1, \pm 2\}$) and left it running all evening.

If the root of largest absolute value has absolute value 1, then all the roots have absolute value 1 (since the constant coefficient of a monic polynomial is, up to sign, the product of the roots), and the solutions to the recurrence relation just consist of repetitions of an initial segment; this is uninteresting.

There's one other degenerate case: if $f(x) = g(x^n)$, we get a recurrence relation whose solution is the interleaving of n independent solutions to the recurrence relation g. This is uninteresting, so I restricted to f with the GCD of the exponents equal to 1.

After an evening of computation, the best value I found was just under 1.0422925, for the polynomial $x^{19} - x^{10} + x + 1$, along with a general trend that higher-degree polynomials offered values nearer to 1; it may of course be that this is because there's a larger space of higher-degree polynomials to sample. The **largest** I observed was just over 2.831177, for $x^3 + 2x^2 - 2x + 1$, but I'm sure larger values would appear if I allowed the middle terms to have coefficients larger than 2.

Obvious questions

1. Can you get $\max(|\theta_i|)$ arbitrarily close to 1 for a sparse polynomial with small coefficients?

2. Does this have anything to do with Pisot numbers and their ilk, the stuff James McKee was studying at Oxford at the end of the nineties?