Sequences, Series, and Recursion

Pavel Panchekha *

April 21, 2011

Formalisms

Sequences are a common topic of math competition questions and, in general, are something you should know about. Formally, a sequence a or $\{a\}_n$ is a ordered sequence of numbers $a_1, a_2, a_3, \ldots$, which can be infinite. Depending on who you ask, a sequence starts at a_0 or a_1 (we'll be using a_1 here, as that is the common mathematical convention). Often, elements of a sequence are integral, but that isn't necessary. Usually, a sequence is generated by a rule, but neither is that necessary. Most of you already know all of this.

We can also define the sum of the first n elements of a sequence as $a_1 + a_2 + a_3 + \cdots + a_n$, denoted by $\sum_{k=1}^{n} a_k$. Lastly, we can define the infinite sum $\sum_{k=1}^{\infty} a_k$, or just $\sum_k a_k$, as the limit of the partial sums.

Basic Sequence Types

The most basic sequence types are arithmetic and geometric sequences. An arithmetic sequence is given recursively by $a_n = a_{n-1} + d$, and a geometric sequence by $a_n = a_{n-1}r$, with a_1, r, and d given. It's pretty easy to see the closed-form (non-recursive) way to represent these: in an arithmetic sequence, we add d each time we go to the next element, so we see that $a_n = a_1 + (n-1)d$, and for a geometric sequence, we use the same reasoning to arrive at $a_n = a_1 r^{n-1}$.

Now, what about sums of these sequences? Well, for an arithmetic sequence, we can take the elements two at a time, with the first and last being a pair, then the second and second-to-last, and so on. Notice that all of these pairs have the same sum. Now, if we have an

*Thanks to Ben Alpert and Ben Kraft for reading and revising drafts.

even number of elements, all the numbers will be paired up, and each pair will have sum $a_1 + (a_1 + (n-1)d) = 2a_1 + (n-1)d$. There will be $n/2$ such pairs, so the final sum is

$$a_1 n + d\frac{n(n-1)}{2}.$$

On the other hand, if we have an odd number of terms, we have $(n-1)/2$ pairs of size $2a_1 + (n-1)d$ and a singleton of half that size, so the final result is

$$\frac{n-1}{2}(2a_1 + (n-1)d) + \frac{1}{2}(2a_1 + (n-1)d) = a_1 n + d\frac{n(n-1)}{2},$$

which is the same as what we got for even n. Cool!

Now, for geometric sequences. Consider the infinite sum $1 + r + r^2 + \ldots$, and let it equal S. Then $rS = r + r^2 + \ldots$, and $S - rS = 1$. Thus, the infinite sum S is equal to $1/1-r$. We can use this to find the partial sum: $1 + r + \cdots + r^n$ is equal to the infinite sum S minus $r^{n+1}S$, or $\frac{r^n-1}{r-1}$. From this, we see that the sum of a geometric sequence is $a_1\frac{r^n-1}{r-1}$. An important note is that in the case of an infinite sum, this is equal to $a_1/(1-r)$.

But anyway, geometric and arithmetic series are boring. Let's move on to something more fun!

Finite Differences

Another good tool to have to attack sequences, one that is much simpler than characteristic polynomials, is finite differences. Basically, the finite difference sequence of a sequence $\{a\}_n$ is the sequence $b_n = a_{n+1} - a_n$. It can give insight into how a sequence works.

For example, if you have the sequence $1, 4, 9, 16, 25, \ldots$, the finite differences are the sequence of odd numbers $3, 5, 7, 9, \ldots$ (prove it!). The finite differences of this are just $2, 2, 2, \ldots$. Now, there are a few obvious properties. Firstly, any linear sequence has constant finite differences (prove it!). Also, if you have the geometric sequence $a_n = a_0 r^n$, the finite difference is $b_n = a_0(r-1)r^n$ (prove it!).

Problem Section 1

1. Above, we saw that n^2, a quadratic, has a linear finite difference sequence. Prove this for all quadratic sequences.

2. What about cubics?
3. What about an nth order polynomial?

Recursively Defined Sequences

In the very general case, recursively defined sequences are those of the form

$$a_n = f(a_1, a_2, a_3, \ldots, a_{n-1}) + g(n).$$

We call such a recursive definition homogeneous if $g(n) = 0$. If f doesn't depend on anything except the last k terms, we say the definition is kth-order. In general, however, we only care about two very specific classes of recursive definitions: kth-order linear recurrences and everything else. A kth-order linear recursive equation is one of the form:

$$a_n = b_1 a_{n-1} + b_2 a_{n-2} + \cdots + b_k a_{n-k} + g(n).$$

I'll add a quote here, by Richard Feynman: Classification of mathematical problems as linear and nonlinear is like classification of the Universe as bananas and non-bananas. Alas, it's true. But linear problems are usually the only solveable ones, so we have to deal[1].

Generally, we try as best we can to avoid non-homogeneous relations. Now, let's look at a problem that has to do with linear recurrence relations of this sort. This problem is rather old (first stated in Europe in 1202), but it is nonetheless a good example of the techniques that we'll be using.

You have the following model of rabbit growth: At first, you start with one immature rabbit, and each month, all immature rabbits mature and all mature rabbits give birth to an immature rabbit (somehow, they reproduce asexually, but that may have something to do with the fact that they're immortal).

Now, we can get the first few terms (always a good idea): $1, 1, 2, 3, 5$. In fact, we notice that each month, the number of mature rabbits is the number of rabbits the month before. Thus, we quickly derive the recurrence relation $R_{n+2} = R_{n+1} + R_n$. Now, this is cool, and we can quickly get lots of terms, but how about a closed form?

To find a closed form, we assume that the closed form expression is of the form $R_n = \lambda^n$. Actually, R_n clearly is *not* just an exponential, because then $\lambda^1 = \lambda^2$, giving us $\lambda = 1$, or

[1]Not convinced that non-linear problems are hard? Try the following recurrence: $a_{n+1} = 4a_n(1 - a_n)$, with $a_1 = 1/3$. No, try it for $a_1 = 1/3 + 10^{-10}$, and note how after a while, the solutions diverge completely. This is the so-called butterfly effect: errors will grow unboundedly in most non-linear systems. You can't get this in a linear system.

$R_n = 1$, but what we're doing here is called ansatz — making an educated guess and building on the results. Anyway, substituting in our guess gives us

$$
\begin{aligned}
\lambda^n &= \lambda^{n-1} + \lambda^{n-2} \\
\lambda^2 - \lambda - 1 &= 0 \\
\lambda &= \frac{1 \pm \sqrt{1 - 4(-1)}}{2} \\
&= \frac{1 \pm \sqrt{5}}{2}.
\end{aligned}
$$

Now we think. If we have some sequence $\{a\}_n$ which satisfies our recurrence, and some other sequence ($\{b\}_n$, say), then $\{a + b\}_n$ clearly also does (check it!), as does $\{ka\}_n$ for some constant k. So if we call our two solutions above ϕ (that's the positive one) and ϕ' (we'll note that it's really just $\frac{1}{-\phi}$), we can say that $A\phi^n + B(-\phi)^{-n}$ is the solution to our recurrence. Since we know the starting values ($R_1 = 1$, $R_2 = 1$), we can plug in and solve for A and B:

$$
\begin{aligned}
A\phi + B(1 - \phi) &= 1 \\
A(1 + \phi) + B(2 - \phi) &= 1 \\
B + (A - B)\phi &= 1 \\
A + 2B + (A - B)\phi &= 1 \\
A + B &= 0 \\
B - 2B\phi &= 1 \\
B = \frac{1}{1 - 2\phi} &= \frac{1}{\sqrt{5}}
\end{aligned}
$$

There's a helpful trick here, however: for a sequence of the second order, with starting values $a_0 = 0$ and $a_1 = 1$, both coefficients are the same, and both are equal to $\frac{1}{\sqrt{D}}$, where D is the determinant of the characteristic polynomial (the polynomial in λ that we get after dividing out λ^{n-2}). Now, our sequence does not have an R_0 term, but we can extend it backwards. Since $R_2 = 1$ and $R_1 = 1$, we know $1 = 1 + R_0$ and $R_0 = 0$. Now, we easily see that the equation for the sequence is

$$
R_n = \frac{\phi^n - (-\phi)^{-n}}{\sqrt{5}}.
$$

By the way, you should all recognize this as the Fibonacci sequence. We can make the formula we derived even nicer by noticing that $(-\phi)^{-1}$ term is less than 1/2 for $n \geq 2$, so we can just ignore it and get the simpler formula

$$R_n \approx \frac{1}{\sqrt{5}}\phi^n.$$

This formula is never off by more than 1/2, so we can just round the result to the nearest integer to get our answer.

Problem Section 2

1. The "Tribonacci" sequence is defined by $T_{n+3} = T_{n+2} + T_{n+1} + T_n$ and the starting values $T_1 = T_2 = T_3 = 1$. Find the smallest n for which T_n is over 9000. A computer would be helpful, but don't just brute force it.

2. Given a set of 3 initial values, what does the sequence $a_{n+3} = 3a_{n+2} - 3a_{n+1} + a_n$ do?

Multiple Roots and Polynomial Approximations

In general, given a linear, homogeneous sequence, we can form the characteristic polynomial by rewriting the sequence but replacing a_{n+k} with λ^k. Finding the roots of the polynomial $r_0, r_1, \ldots, r_k$ lets us write the closed form expression in the form $a_n = C_0 r_0^n + C_1 r_1^n + \cdots + C_k r_k^n$. One important thing to note is that if one of the roots is a double root or kth-order root, the corresponding coefficient will instead be a $(k-1)$-order polynomial. By the way, some of your roots may be complex. For example, consider the recurrence $a_{n+2} = -a_n$, with initial values $a_0 = 0$, $a_1 = 1$. This is the sequence $\sin \frac{\pi}{2} n$, and has characteristic polynomial $\lambda^2 = -1$. Clearly, this has no real roots. But you can solve it to get $a_n = \frac{1}{2i}(i^n - i^{-n})$.

Here's an example where the roots can have multiplicities. Given a sequence a where $a_{n+2} = 2a_{n+1} - a_n$, describe what this sequence does for starting values a_1 and a_2.

Well, we can try an example or two (always a good idea!). If the starting values are $a_1 = 2$ and $a_2 = 7$, we have the sequence $2, 7, 12, 17, \ldots$. If the starting values are $a_1 = 3$ and $a_2 = 5$, we have $3, 5, 7, 9, 11, \ldots$. In general, the sequence looks like it's making an arithmetic sequence from the starting values. Let's see why.

We'll use characteristic polynomials. Forming the characteristic polynomial, we get $\lambda^2 = 2\lambda - 1$. We instantly see that this factors, giving $(\lambda - 1)^2 = 0$, so 1 is a double root. Now, since it's a double root, the closed form we get will have a coefficient in front of r_0^n that is linear (the double root means the coefficient will have two terms, meaning it's linear). So, our closed form looks like $a_n = (An + B)1^n$, which obviously simplifies to $a_n = An + B$. Thus, we've proven that our sequence from before produces a linear extrapolation of the starting terms.

In fact, you can generalize this to get a quick way to do nth-order extrapolations of a few data points. If we want a cubic approximation to $1, 2, 5, 3$, we simply use the recursive sequence $a_{n+4} = 4a_{n+3} - 6a_{n+2} + 4a_{n+1} - a_n$. So we can quickly calculate that the next term would be $4 \cdot 3 - 6 \cdot 5 + 4 \cdot 2 - 1 = 12 - 30 + 8 - 1 = -11$.

This also proves the nontrivial property that if the value of an nth-order polynomial is integral at $n + 1$ consecutive points, it is integral for any integer argument.

Problem Section 3

1. Find a recursive definition for the sequence whose closed form is $a_n = (n^2 + 1)2^n + 1$.
2. A 3rd order polynomial P has the property that $P(1) = 1$, $P(2) = 18$, $P(4) = 17$, and $P(5) = 23$. Find $P(3)$.
3. Check whether there exists a quintic P such that $P(0) = 0$, $P(1) = 1$, $P(2) = -2$, $P(3) = 3$, $P(4) = -4$, $P(5) = 5$, and $P(6) = -3$.

Here's another cool use of characteristic polynomials. Let's say we have a sequence which is periodic with period p. Then your recurrence is $a_{n+p} = a_n$, and your characteristic polynomial $\lambda^p - 1 = 0$. The roots of this are the p-th roots of unity (which is, honestly, pretty cool by itself). It also means that a complete characterization of periodic functions (periodic on the integers, that is) is just that: a linear combination of terms involving $e^{2\pi i l/p}$. This has cool connections to number theory and group theory, but let's get back to the main topic of this talk.

So, let's do another problem. We have the sequence $1, 2, 4, 8, 16, \ldots, 2^n$. However, we don't recognize that the sequence is just the powers of two (silly us!), so when we are asked for the next term, we just do a polynomial approximation. How far off are we?

Let's try a few simple examples. If we call the $P(n)$ the value we're trying to find, we can use that polynomial extrapolation formula we got in the last section to see that:

$$\begin{aligned}
P(1) &= 1 = 1 \\
P(2) &= 2 \cdot 2 - 1 = 3 \\
P(3) &= 3 \cdot 4 - 3 \cdot 2 + 1 = 7 \\
P(4) &= 4 \cdot 8 - 6 \cdot 4 + 4 \cdot 2 - 1 = 15
\end{aligned}$$

Note that in each case, $P(n)$ is exactly one less than 2^n. Let's try to prove it.

Now, since we have a shiny new hammer (ooh, characteristic polynomials… shiny!), let's try to hit this problem on the head with it. We know (by again using our polynomial extrapolation trick) that we're going to need the value of

$$\sum_{k=0}^{n}(-1)^n\binom{n+1}{n-k}2^{n-k}.$$

How do we do it? Let's simplify a bit, by subtracting this from 2^n:

$$2^{n+1}-\sum_{k=0}^{n+1}(-1)^k\binom{n+1}{k}2^{n+1-k}.$$

Notice anything? It looks like the binomial theorem: the above sum is equivalent to $(-1+2)^k = 1$, so in every case, we are off by exactly 1, in that we always end up with something one less than the correct value. This proof, by the way, though slick, is rather opaque – the reason the value is one away is rather non-obvious.

Rewind

So, let's go back to the problem we had before, on the estimation of 2^n. We'll be using the result that you end up proving in the problem section just before, so you should probably do it. It's not too hard. (Really. It's not. Go do it.)

Now then, since we know that an nth order polynomial will become an $(n-1)$-th order polynomial when you take its finite differences, we can "work backwards" to get a polynomial approximation to a sequence: if you have n terms, $n-1$ finite differences will reduce you to a single number, which obviously has a constant polynomial approximation. You can then extend that sequence, and work backwards.

This actually leads to a trivial proof of what we had above. Consider the finite differences of $1, 2, 4, \ldots, 2^n, 2^{n+1}-1$. What are they? $1, 2, 4, \ldots, 2^{n-1}, 2^n-1$ (check it!)! So, all we need is some trivial induction to prove our claim above.

Inhomogenous Recurrence Relations

You know how you're only supposed to buy milk that's been homogenized? Well, the same applies to linear recurrence relations. Only solve homogenous ones.

Well, sort of. Because there's actually a lot of cool things you can do with inhomogenous ones as well. To start with, here's an example:

$$a_{n+1} = a_n + d$$

Now, we've seen this before. In fact, we solved it in the very beginning of the talk: it's an arithmetic series. But our shiny cool new toys get pwned by this really trivial stupid problem. What to do! Aha, don't worry! There's a cool trick involved here. If $a_{n+1} = a_n + d$, then $a_{n+2} = a_{n+1} + d$, right? And that means that $a_{n+2} - a_{n+1} = a_{n+1} - a_n$. We rearrange this and find the characteristic polynomial, which is $(\lambda - 1)^2$. Then, we can just solve our recurrence the usual way. We know that $a_n = An + B$, and from $a_1 = a$, $a_2 = a + d$, we see that $A = a$, $B = d$. Now, note that $a_2 = a + d$ was not one of the initial conditions; we needed to use our recurrence to get that. In general, when solving inhomogenous relations, you'll need to apply your recurrence several times to get extra "initial conditions".

What about a harder problem? Like this one:

$$a_{n+2} = a_{n+1} + 3a_n + F_n$$

where F_n is the nth Fibonacci number. Well, how are we going to solve this? We can use a trick similar to the one we just used. Take the three relations

$$\begin{aligned} a_{n+4} &= a_{n+3} + 3a_{n+2} + F_{n+2} \\ a_{n+3} &= a_{n+2} + 3a_{n+1} + F_{n+1} \\ a_{n+2} &= a_{n+1} + 3a_n + F_n \end{aligned}$$

Now we can subtract the second and third equations from the first, cancelling the Fibonacci terms (due to their recursive relation) and leaving us with

$$a_{n+4} - 2a_{n+3} - 3a_{n+2} + 4a_{n+1} + 3a_n = 0,$$

which we can solve in the usual way (the characteristic polynomial: $(\lambda^2 - \lambda - 1)(\lambda^2 - \lambda - 3)$) I won't go through the actual solution here, it's ugly; suffice it to say that the roots are approximately 1.6, $-.6$, -1.3, and 2.3.

OK, generalizing time! Let's say that the inhomogenizing term is some arbitrary thing, but it has a recurrence relation. So, basically, we have

$$c_k a_{n+k} + c_{k-1} a_{n+k-1} + \cdots + c_0 a_n = f(n),$$

where

$$d_m f(n+m) + d_{m-1} f(n+m) + \cdots + d_0 f(n) = 0.$$

What do we do now? Well, we want to cancel all of the f terms, so we can just add the first equation times d_m and shifted over m, times d_{m-1} shifted over $m-1$, and so on. We get:

$$\begin{aligned}
& d_m && (c_k a_{n+k+m} + c_{k-1} a_{n+k+m-1} + \cdots + c_1 a_{n+m+1} + c_0 a_{n+m}) && (1)\\
& +d_{m-1} && (c_k a_{n+k+m-1} + c_{k-1} a_{n+k+m-2} + \cdots + c_1 a_{n+m} + c_0 a_{n+m-1}) && (2)\\
& +\dots && && (3)\\
& +d_1 && (c_k a_{n+k+1} + c_{k-1} a_{n+k} + \cdots + c_1 a_{n+2} + c_0 a_{n+1}) && (4)\\
& +d_0 && (c_k a_{n+k} + c_{k-1} a_{n+k-1} + \cdots + c_1 a_{n+1} + c_0 a_n) && (5)
\end{aligned}$$

Now, we regroup this by the a_n terms:

$$\begin{aligned}
& (d_m c_k) && a_{n+k+m}\\
& +(d_{m-1} c_k + d_m c_{k-1}) && a_{n+k+m-1}\\
& +(d_{m-2} c_k + d_{m-1} c_{k-1} + d_m c_{k-2}) && a_{n+k+m-2}\\
& +\dots && \\
& +(d_1 c_0 + d_0 c_1) && a_{n+1}\\
& +(d_0 c_0) && a_n\\
& =0 &&
\end{aligned}$$

Form the characteristic polynomial:

$$(d_m c_k)\lambda^{k+m} + (d_{m-1} c_k + d_m c_{k-1})\lambda^{k+m-1} + \cdots + (d_1 c_0 + d_0 c_1)\lambda + (d_0 c_0) = 0$$

Now, compare this to the two polynomials $c_k \lambda^k + c_{k-1}\lambda^{k-1} + \cdots + c_1 \lambda + c_0$ and $d_m \lambda^m + d_{m-1}\lambda^{m-1} + \cdots + d_1 \lambda + d_0$. Who noticed it? The big polynomial is just the product of the two smaller ones! Cool!

Let's do a quick check, for $a_{n+1} = a_n + d$. The polynomial for the left is $\lambda - 1$, and for the inhomogenizing term it's got to be $\lambda - 1$ (check it!). K, does that work? Well, the product is $(\lambda - 1)^2$, which gives us the closed form $An + B$, which is indeed correct. Win!

Now, what if the inhomogenizing term has no nice recurrence? Say we want $a_{n+2} = 4a_{n+1} - 7a_n + \sin n$? Well, honestly, in that case you're completely screwed no matter what, so don't sweat it. Get out your handy dandy calculator and start programming.

Problem Section 4

1. What happens if we add the solution to our recurrence back into the recurrence? As in, what if we have $a_{n+2} = a_{n+1} + a_n + F_n$, where the inhomogenizing term has the same recurrence relation as the rest of the recurrence?

2. "Verify" the formulae for sums of arithmetic and geometric series using a cool application of inhomogenous recurrence relations.

3. Find a way to get the partial sums of a recurrence relation in explicit form. This is really cool, so I highly suggest you do it.

Hmm... This Needs More Calculus

Doesn't everything?

Now, if we have a function $f(n)$, we can define the *finite derivative* $\frac{\Delta f}{\Delta n}$ to be $f(n+1) - f(n)$, that is, the finite differences of the sequence $a_n = f(n)$. You'll note that this definition is somewhat similar to the calculus definition:

$$\frac{df}{dx} = \lim_{h \to 0} \frac{f(x+h) - f(x)}{h},$$

but when we say $\lim_{a \to b}$, we mean "get a as close as you can to b without getting there". And in the integers, the closest you can get is one away, so the normal derivative becomes the finite derivative.

Why is this useful? Well, it's really cool, and there's a lot more to say on this subject than I will[2], but there are a few applications we want to use on our problem above, to actually explain why the answer is $2^{n+1} - 1$.

What is $\frac{\Delta}{\Delta n} 2^n$? It's 2^n – we proved that above (or at least, you should have). That's actually pretty remarkable, especially if you know regular calculus, where the equivalent function is e^x (which is part of why e is so important).

What about n^2? We have $(n+1)^2 - n^2 = 2n + 1$. Hmm. We'd like it to be $2n$, so that it's similar to regular calculus. Note that instead, $n^2 - n$'s finite derivative is $2n$. What about n^3? What's its replacement? Do we have to get these experimentally? Is there a general rule? There is. Consider $n^{\underline{k}} = n(n-1)(n-2)(n-3)\cdots(n-k+1)$. Prove for yourself that $\frac{\Delta}{\Delta n} n^{\underline{k}} = k n^{\underline{k-1}}$.

Now, I really want to get back to resolving our problem, but I'll return to other cool applications of finite calculus. For now, though, I'll make an unqualified statement: Taylor's theorem works in finite calculus (the proof is straightforward if you know the proof in regular calculus; it is, however, tedious). So, what does that mean? It means there's an operator

[2] See: `http://www.stanford.edu/~dgleich/publications/finite-calculus.pdf`

T_k, where $T_k f(n)$ is defined as

$$T_k f(n) = \sum_{l=0}^{\infty} \frac{(\frac{\Delta}{\Delta n} f)(k)}{l!} (n-k)^{\underline{l}},$$

and that (provided certain conditions are met) $T_k f(n)$ is equal to $f(n)$ (at least, in a region near k). Now, simplifying the above a bit and setting k to be 0, we get

$$\sum_{l=0}^{\infty} \binom{n}{l} \left(\frac{\Delta}{\Delta n} f \right) (0).$$

Isn't it cool that the binomial coefficients just popped out of nowhere? It's also crucial to our problem. As in regular Taylor's theorem, there's a sense in which $Tf(n)$, if you cut it down to the first $k+1$ terms, is the best kth order polynomial approximation to $f(n)$. So, to solve the above problem, we just have to consider the Taylor expansion of 2^{n+1}, with the $(n+2)$-th term chopped off.

What's the Taylor expansion of 2^{n+1}? Since $\frac{\Delta}{\Delta n} 2^n = 2^n$, and $2^0 = 1$, it's $\binom{n}{0} + \binom{n}{1} + \binom{n}{2} + \ldots$, which is indeed equal to 2^n. Now, this series cuts itself off at some point (at n terms), which is a very nice feature – you don't have to chop off infinitely many terms. Just one will do.

So, we have our polynomial approximation to 2^{n+1}:

$$2^{n+1} \approx \sum_{k=0}^{n} \binom{n+1}{k}.$$

Note that we're not adding in the last term here. But what is that last term? $\binom{n+1}{n+1} = 1$, which is exactly why we're 1 off in our approximation. Now we, in a sense, know the *reason* for the theorem we've been proving. This actually also makes it very easy to extend our solution. What if we extend the approximation *two* terms ahead? Well, we're chopping off $\binom{n+1}{n} + \binom{n+1}{n+1}$, or $n+2$, so we'll be that far off. Extend $1, 2, 4, 8, 16, 32$ two terms out, and you're going to get 63, then 121. Isn't that cool?

Deleted Footage, Blooper Reals, and Bonus Material

So, some of you may protest that though the math works, the *reason* that inhomogenous relations work out the way they do is mystical. So let's explain what really happens. Since your recurrence is linear, we can consider it the application of a matrix, like so:

$$\begin{bmatrix} -\frac{c_{k-1}}{c_k} & -\frac{c_{k-2}}{c_k} & \cdots & -\frac{c_1}{c_k} & -\frac{c_0}{c_k} \\ 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & \cdots & 1 & 0 \end{bmatrix} \begin{bmatrix} a_{n+k-1} \\ a_{n+k-2} \\ a_{n+k-3} \\ \vdots \\ a_{n+1} \\ a_n \end{bmatrix} = \begin{bmatrix} a_{n+k} \\ a_{n+k-1} \\ a_{n+k-2} \\ \vdots \\ a_{n+2} \\ a_{n+1} \end{bmatrix}$$

Now, the basic assertion that we made here was that over the complex numbers (since our roots can be complex) this matrix is diagonalizeable. Its eigenvalues are just the roots of our characteristic polynomial (convince yourself!). Now, what if we have our matrix multiplication, but then a matrix addition? That's what you get for inhomogenous polynomials. Well, if the matrix you're adding is a linear recurrence itself, you can represent that by the recurrence's matrix operating in higher dimensions and then just offloading its results into the part that we're adding. And what would the eigenvalues of this big matrix be? Obviously, just the union of the eigenvalues (with associated multiplicities). Which gives you the overall polynomial being the product of the two smaller ones.

Now, no lecture would be complete without an open problem. In this case, the open problem is: say you have a linear recurrence $c_k a_{n+k} + \cdots + c_0 a_n = 0$. Is there an n such that $a_n = 0$? What we want is a decision procedure. Well, using the material we learned here, we can see that we reduce the problem to: given a linear combination of exponentials over the polynomials, can we determine whether the equation has a root? Go forth, solve it, and win the Fields medal. I dare you.

Now, if that's an open problem, I hope you'll agree that non-linear recurrences are just about impossible. Here's an example: the famous $3n+1$ problem. Take a number. I'll use 17 here, for demonstration. Now, if this number is odd (it is) multiply it by 3 and add 1. OK, now I have 52. If the number is even (now it is) divide it by 2. OK, now 26. Then 13. Then 40. Then 20, 10, 5. Then 16, 8, 4, 2, 1, 4, 2, 1, 4, 2, 1, ... The question is, do all numbers eventually reach that 4, 2, 1 cycle? Now, this recurrence is even almost linear — each of the two options at any point is linear, it's just that the thing overall is highly non-linear.

More

For more on finite calculus, including a great tutorial, I'm again going to suggest

`http://www.stanford.edu/~dgleich/publications/finite-calculus.pdf`.

On the subject of characteristic polynomials, there's a wonderful compilation of good problems (and much of the same material as here) at

`http://mathcircle.berkeley.edu/BMC3/Bjorn1/Bjorn1.html`.

Wikipedia is, as always, your friend. Its article on recurrences is pretty good; find it at

`http://en.wikipedia.org/wiki/Recurrence_relation`.