Height gaps for Mahler series

Jason Bell (Joint with Boris Adamczewski and Daniel Smertnig)

March 30, 2021

Difference and differential equations

Most people are familiar with the study of solutions of homogeneous linear differential equations

$$q(z) = \sum_{i=0}^{m} p_i(t) D^i F(t) = 0,$$

where D is differentiation with respect to t and the $p_i(t)$ and $q(t)$ are rational functions (not all zero). A power series that satisfies such an equation is called D-finite and this class includes algebraic power series and famous series such as e^t.

Slightly less well-known are the class of solutions to difference equations, where one uses an endomorphism σ instead of a derivation. The class we'll look at is Mahler series, where one takes the continuous endomorphism of $K[[t]]$ induced by $t \mapsto t^k$ and studies power series solutions to equations

$$q(z) = \sum_{i=0}^{m} p_i(t) F(t^{k^i}).$$

Such a solution is called a k-Mahler power series and this includes many nice classes of series. Most notably, this class contains the "k-automatic power series", which we'll now define.

Let k be a positive integer ≥ 2. Then every nonnegative integer has a unique base-k expansion on the digits $0, \ldots, k-1$ with no "leading zeros".

For example, if $k = 3$, then the base-k expansion of 17 is 122, since $17 = 1 \cdot 3^2 + 2 \cdot 3^1 + 2 \cdot 3^0$.

In general, if w is a word on the alphabet $\{0, \ldots, k-1\}$ with no leading zeros that is the base-k expansion of n, we write $[w]_k = n$ and $(n)_k = w$.

Convention: $[\epsilon]_k = 0$ and $(0)_k = \epsilon$, where ϵ is the empty word (the identity in the free monoid on the set $\{0, \ldots, k-1\}$.

Definition

A *finite-state k-automaton with output* is a finite directed graph where each vertex has out-degree k with labels $0, 1, 2, \ldots, k-1$, along with a special distinguished vertex (the initial state) and each vertex has an “output value” from some finite set Δ.

Note: if we have a finite-state k-automaton with output, we can associate a function from $f : \mathbb{N} \to \Delta$ as follows. Given $n \in \mathbb{N}$, we find $(n)_k \in \{0, \ldots, k-1\}^*$; we start at the initial state, reading the digits of $(n)_k$ from right-to-left, using the arrows to tell us where to go at each step. Once this process finishes, we look at the output value of the vertex we finish at and that is $\tau(n)$.

Example: The Thue-Morse automaton

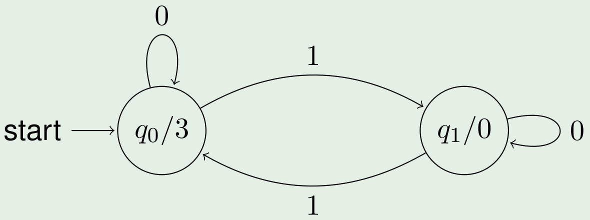

Here: $k = 2$; If we look $(13)_2 = 1101$ and the output is 0, so the associated map τ has $\tau(13) = 0$. In general, we see that $\tau(n) = 3$ if and only if n has an even number of 1's in its binary expansion.

Then we say that a power series $F(z) = \sum \tau(n) z^n \in \mathbb{C}[[z]]$ is a k-*automatic power series*.

Theorem

A k-automatic power series is k-Mahler.

In our example, we had $k = 2$ and the function $\tau(n)$ takes the values 0 and 3 and is 3 if and only if n has an even number of 1's in its binary expansion. So why is $F(z) = \sum \tau(n) z^n$ a 2-Mahler series in this case?

Notice that n has an even number of 1's in its binary expansion if and only if $2n$ does too and if and only if $2n+1$ does not, so $\tau(2n) = \tau(n)$ and $\tau(2n+1) = 3 - \tau(n)$.

$$F(z) = \sum_{j=0}^{\infty} \tau(2j) z^{2j} + \sum_{j=0}^{\infty} \tau(2j+1) z^{2j+1},$$

and we can rewrite the right side Notice that n has an even number of 1's in its binary expansion if and only if $2n$ does too and if and only if $2n+1$ does not. So

$$F(z^2) + \left(3z/(1-z^2) - zF(z^2)\right).$$

Well, that doesn't prove the theorem, but in fact one can show that one can do similar constructions more generally by using the fact that the k-*kernel* of an automatic sequence is finite.

Definition

Given a sequence $f : \mathbb{N} \to \Delta$, we define the k-*kernel* of f to be the set of distinct subsequences of the form

$$f(k^a n + b)$$

with $a \geq 0$ and $0 \leq b < k^a$.

In our Thue-Morse example, we see the 2-kernel is just the two maps $\tau(n)$ and $3 - \tau(n)$, which follows from induction and the fact that $\tau(2n) = \tau(n) = 3 - \tau(2n+1)$. In fact, finiteness of the k-kernel characterizes k-automaticity.

Allouche and Shallit extended the notion of automaticity to regular sequences. Here K is a field and $f : \mathbb{N} \to K$ is a K-valued sequence. Then we can look at the K-vector space spanned by the elements of the K-kernel.

Definition

We say that f is k-*regular* if the K-vector space spanned by the sequences in the k-kernel is finite-dimensional as a K-vector space.

We see that k-automatic $\implies$ k-regular, since if the k-kernel is finite, the K-vector space generated by the elements of the k-kernel will be finite-dimensional.

Example

Let $f(n)$ be the number of 1's in the binary expansion of n.

Then $f(2n+1) = f(n) + 1$ and $f(2n) = f(n)$, so now we see that the vector space spanned by the 2-kernel of f is spanned by the constant map $g(n) = 1$ and $f(n)$.

So f is 2-regular, but it is not 2-automatic, since it takes infinitely many values.

In general: $g(n)$ is k-automatic $\iff$ $g(n)$ is k-regular and the range of g is finite.

Just as with automatic power series, regular power series (generating series of k-regular sequences) are also k-Mahler. So we have the inclusions

$$k-\text{automatic series} \subseteq k-\text{regular series} \subseteq k-\text{Mahler series}.$$

These containments are all proper.

One of the most natural ways of studying and understanding integer power series is via their asymptotics. We use this to get both a sense of how the coefficients grow as well as how complicated the series is.

This approach has been especially fruitful in the study of solutions to differential equations, although there has apparently been less work in the realm of solutions to difference equations.

Still, it is a natural question as to what the asymptotics are of solutions to Mahler series. We'll look at a few examples of Mahler series.

- For automatic power series the coefficients are $\mathrm{O}(1)$.
- For regular series with integer values the coefficients are known to have a gap: either the coefficients are $\mathrm{O}(\log(n)^d)$ or infinitely often they are $\mathrm{O}(n^\kappa)$ and are at least Cn^κ for infinitely many n for some $C > 0$.

For example the sum of the 1's in the binary expansions of n jumps around, but it is at most $\log_2(n)$. On the other hand, the sequence $f(n) = n^2 + 1$ is k-regular for every k and it grows like a polynomial in n.

Notice that if $f(n) = n^2 + 1$, then $f(k^a n + b)$ is a quadratic polynomial in n, so the k-kernel is spanned by the maps 1, n, n^2, so f is indeed k-regular.

Let $F(z)$ denote the infinite product of cyclotomic polynomials

$$\prod_{n=0}^{\infty} \frac{1}{1 - z^{k^n}} = \sum_{n=0}^{\infty} a_n z^n \, .$$

Then $F(z^k) = (1 - z)F(z)$, so $F(z)$ is k-automatic.

The integer a_n is equal to the number of partitions of n into k-powers. The asymptotics of a_n were first studied by Mahler who proved that

$$\log a_n \sim \frac{\log^2 n}{2 \log k} \cdot$$

These results of Mahler have been refined and generalized by de Bruijn and most recently by Dumas–Flajolet.

The asymptotics immediately show that $F(z)$ cannot be regular.

There's one obvious type of Mahler series we haven't consider. Notice that $1/(1-2z)$ is k-Mahler for every k. Since we have

$$F(z) = 1/(1-2z).$$

Notice the coefficients are growing exponentially in this case.

The five behaviours

So we've shown there are some different possibilities of how coefficients in a Mahler series. So far we've seen examples with growth of the following types:

- $\mathrm{O}(1)$ (automatic series);
- Not $\mathrm{O}(1)$ but $\mathrm{O}((\log n)^d)$ (certain regular series);
- Not $\mathrm{O}((\log n)^d)$ and $\mathrm{O}(n^\kappa)$ (other regular series);
- Growth "like" $\frac{\log^2 n}{2 \log k}$;
- Exponential growth.

Because coefficients can bounce around a lot, we find it convenient to use Landau notation: We write $a_n \in \Omega(b_n)$ when $a_n \notin o(b_n)$, that is, when there exists a positive number c such that $a_n > cb_n$ for infinitely many positive integers n.

Theorem

(B-Coons-Hare/B) We can refine the gaps for regular series and say that if a_n is an integer-valued k-regular series then either: there is some nonnegative integer d such that

$$a_n \in \mathrm{O} \cap \Omega((\log\, n)^d))$$

or $a_n \in \mathrm{O} \cap \Omega(n^\kappa)$, for some $\kappa = \log_k(\alpha)$, with $\alpha > 1$ a real algebraic number.

So in some sense we understand how coefficients of regular power series can grow.

Are there other growth types? No! We'll make this precise, but in general one is often interested in Mahler series with coefficients in $\bar{Q}$ and asymptotic information is not always so relevant as far as measuring the complexity of the coefficients. A much better way of capturing complexity in this context is via heights.

The height is a map $h : \bar{\mathbb{Q}} \to [0, \infty)$, which has the property that $h(n) = |n|$ for n an integer, so we can think of the height as giving us a coarse analogue of asymptotics for general algebraic numbers.

A crash course on heights

We give a quick overview of the ideas involved.

- Let K be a number field of degree $d := [K : \mathbb{Q}]$ that is Galois.
- Let $M(K)$ denote the set of places of K. Recall that each place $v \in M(K)$ is either finite or infinite and, in either case, determines a normalized absolute value $|\cdot|_v : K \to [0, \infty)$.

- (finite places) If $v \in M_{\text{fin}}(K) \subset M(K)$ is finite, it corresponds to a prime ideal $\mathfrak{p}$ of the ring of integers $\mathcal{O}_K$ of K, then the order $\text{ord}_{\mathfrak{p}} x$ of $x \in \mathcal{O}_K$ is the largest power $m \geq 0$ such that $x \in \mathfrak{p}^m$. If more generally $x \in K$, then one writes $x = a/b$ for some $a, b \in \mathcal{O}_K$ and $\text{ord}_{\mathfrak{p}} x := \text{ord}_{\mathfrak{p}} a - \text{ord}_{\mathfrak{p}} b$.

  $|x|_v := 0$ if $x = 0$, and $|x|_v := \text{N}(\mathfrak{p})^{-\text{ord}_{\mathfrak{p}}(x)}$ if $x \neq 0$, where $\text{N}(\mathfrak{p})$ is the cardinality of the finite field $\mathcal{O}_K/\mathfrak{p}$.

- If $v \in M_{\text{inf}}(K) \subset M(K)$ is an infinite place, v is either *real* or *complex*. In the first case, v corresponds to a real embedding $\tau : K \to \mathbb{R}$, and we take $|x|_v = |\tau(x)|$, where $|\cdot|$ is the ordinary absolute value on $\mathbb{R}$. In the second case, v corresponds to a distinct pair $\tau, \bar{\tau} : K \to \mathbb{C}$ of complex embeddings, and we take $|x|_v = |\tau(x)|^2 = |\bar{\tau}(x)|^2$.

With these definitions we then have a product formula

$$\prod_{v \in M(K)} |c|_v = 1$$

for $c \in K^*$. (And I should add that $|c|_v = 1$ for all but finitely many places when $c \in K^*$.)

Now we can define the height of an algebraic number $a \in K$ number as follows:

$$h(a) = \frac{1}{d} \sum_{v \in M(K)} \log(\max(|c|_v, 1))$$

Let's look at a few examples.

What is $h(2/9)$? Here $K = \mathbb{Q}$ and $|2/9|_v = 1$ except for the 2-adic place, the 3-adic place, and the infinite place (Euclidean absolute value). Notice $|2/9| < 1$ and $|2/9|_2 < 1$ so they don't contribute anything to

$$\sum_{v \in M(Q)} \log(\max(|2/9|_v, 1)),$$

so we just get $\log\,|2/9|_3 = \log(9)$. In general

$$h(a/b) = \max(\log|a|, \log|b|)$$

for $\gcd(a,b) = 1$, a, b nonzero integers.

What is $h(\sqrt{3})$?

Here we are working in $K = \mathbb{Q}(\sqrt{3})$ and most places of $\mathbb{Q}$ lift to exactly two places of K. For example, we have two Archimedean absolute values on K:

$$|a + b\sqrt{3}|_1 := |a + b\sqrt{3}|, \quad |a + b\sqrt{3}|_2 = |a - b\sqrt{3}|.$$

In this case, $|\sqrt{3}|_v \leq 1$ for all finite places and so the height of $\sqrt{3}$ is $\log\,|\sqrt{3}|$.

This is somehow a more "democratic" notion of asymptotics. For example $17 - 12\sqrt{2}$ is small, but this is because we've chosen the positive square root of $\sqrt{2}$. Had we chosen the negative square root, the number would be pretty big, and the height is in some sense attempting to average over all conjugates.

If we recast our examples of different types of growths of integer Mahler series

$$\sum a_n z^n$$

in terms of heights, we see we get the following possibilities:

- $a_n = \mathrm{O}(1) \iff h(a_n) = \mathrm{O}(1)$.
- $a_n = \mathrm{O} \cap \Omega((\log\, n)^d),\ d > 0 \iff h(a_n) \in \mathrm{O} \cap \Omega(\log\log\, n)$.
- $a_n = \mathrm{O} \cap \Omega(n^\kappa) \iff h(a_n) \in \mathrm{O} \cap \Omega(\log\, n)$;
- $h(a_n) \in \mathrm{O} \cap \Omega((\log\, n)^2)$.
- exponential growth $\implies h(a_n) \in \mathrm{O} \cap \Omega(n)$

In fact, we're able to show that one has to fall into one of these cases.

Theorem

(Adamczewski-B-Smertnig) Let $F(z) = \sum_{n=0}^{\infty} a_n z^n$ *be a* k*-Mahler series with coefficients in* $\bar{\mathbb{Q}}$*. Then exactly one of the following properties holds.*

1. $h(a_n) \in O \cap \Omega(n)$.
2. $h(a_n) \in O \cap \Omega(\log^2 n)$.
3. $h(a_n) \in O \cap \Omega(\log n)$.
4. $h(a_n) \in O \cap \Omega(\log \log n)$.
5. $h(a_n) \in O(1)$.

We make a remark that it is not in general possible to replace lower bounds of the type Ω by stronger ones. For instance, the 2-Mahler function $\sum_{n=0}^{\infty} 2^n z^{2^n}$ belongs to class (3), but most of its coefficients vanish.

Interlude

How did this project get started?

What gave really gave impetus to this project is that at the time, I had been interested in the concept of “height gaps” due to noting they held in many contexts: differential equations, arithmetic dynamics, etc. Daniel Smertnig was a postdoc at the University of Waterloo and each time I would suggest some idea, he was always able to do incredible things. In this case the idea I suggested we look at was ...

The Mahler Denominator

Every k-Mahler series with coefficients in K satisfies a homogeneous Mahler equation with polynomial coefficients:

$$p_0(z)F(z) + p_1(z)F(z^k) + \cdots + p_d(z)F(z^{k^d}) = 0.$$

Notice that

$$p_0(z)F(z) \in \sum_{j\geq 1} K[z]F(z^{k^j}),$$

but in general there might be smaller degree polynomials that do this as well.

Definition

Let $F(z)$ be a k-Mahler power series, and let

$$\mathfrak{I} = \left\{ p(z) \in K[z] : p(z)F(z) \in \sum_{i=1}^{\infty} K[z]F(z^{k^i}) \right\}.$$

The k-*Mahler denominator* of $F(z)$ is the unique generator $\mathfrak{d}(z) \in K[z]$ of the ideal $\mathfrak{I}$, with the lowest nonzero coefficient of $\mathfrak{d}$ being 1.

This somehow seems to what really controls the asymptotics/height growth of Mahler series.

Example

The equation

$$(z - 1/2)F(z) - (z - 1/8)(z^3 - 1/2)F(z^3) = 0$$

has a unique nonzero power series solution (up to a scalar) and is minimal with respect to this solution. However, we can see the Mahler denominator is 1.

$$F(z) = (z - 1/8)(z^2 + 1/2z + 1/4)(z^9 - 1/2)F(z^9)\,.$$

The expected pole at $1/2$ disappears after one iteration of the equation.

The Mahler Denominator controls asymptotics

If we look at some examples ...

$F_1(z) = 1/(1-2z)$	$\mathfrak{d}(z) = 1-2z$
$F_2(z) = \prod_j (1-z^{2^j})^{-1}$ $(k=2)$	$\mathfrak{d}(z) = 1-z$
$F_3(z) = \prod_j (1+z^{2^j})^{-1}$ $(k=2)$	$\mathfrak{d}(z) = 1+z$
$F_4(z) = (1+z)^2(1+z^k)^2(1+z^{k^2})^2\cdots$	$\mathfrak{d}(z) = 1$

Maybe you've seen this trick before:

$$F_3(z) = \prod (1+z^{2^j})^{-1} = \prod \frac{(1-z^{2^j})}{(1-z^{2^{j+1}})} = (1-z).$$

Theorem

Let $f(z) = \sum_{n=0}^{\infty} a_n z^n$ be a k-Mahler function. The following statements are equivalent.

1. *We have $h(a_n) \in o(\log^2 n)$.*
2. *Every non-zero root of the k-Mahler denominator of f is a root of unity that is not periodic under repeated iteration of the map $\zeta \mapsto \zeta^k$*
3. *The power series f is k-regular.*
4. *We have $h(a_n) \in O(\log n)$.*

Theorem

Let $F(z) = \sum_{n=0}^{\infty} a_n z^n$ be a k-Mahler function. The following statements are equivalent.

1. *We have $h(a_n) \in o(n)$.*
2. *Every non-zero root of the k-Mahler denominator of f is a root of unity.*
3. *The power series f has radius of convergence one with respect to all places.*
4. *We have $h(a_n) \in O(\log^2 n)$.*

It's not hard to show that every Mahler series $\sum a_n z^n$ satisfies $h(a_n) = \mathrm{O}(a_n)$, so this really explains the first three gaps:

- $h(a_n) \in O \cap \Omega(n)$.
- $h(a_n) \in O \cap \Omega(\log^2 n)$.
- $h(a_n) \in O(\log n)$.

So we're left to deal with classifying height gaps for k-regular series. As mentioned before, there were already some precedents, which indicated that there should be a sort of trichotomy: growth like $\log n$; growth like $\log \log n$; and $\mathrm{O}(1)$ growth.

Growth of k-regular series

Although Allouche and Shallit introduced k-regular series, they existed in an equivalent form, studied in depth by Berstel and Reutenhauer, called noncommutative series. As far as the relevance to k-regular series, one takes $k \geq 2$, K a field, and a monoid homomorphism $\Phi : \{0, \ldots, k-1\}^* \to M_d(K)$, and column vectors $v, w \in K^d$. Then one can associate a function

$$f : \mathbb{N} \to K$$

via the rule

$$n \mapsto w^T \Phi((n)_k) v.$$

Example

Let

$$A_0 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix},$$

and

$$A_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

Then one can make a map $f : \mathbb{N} \to \mathbb{Z}$ via

$$f([i_d \cdots i_0]_2) = e_1^T A_{i_d} \cdots A_{i_0} e_1.$$

What is $f(n)$?

Theorem

A series is k-regular if and only if f can be realized as a map

$$n \mapsto w^T \Phi((n)_k) v$$

as above.

The spectral radius controls asymptotics

Now what we can see occurs is that if we have some nonzero eigenvalue that is not a root of unity then $h(a_n) \in O \cap \Omega(\log\ n)$ and if all eigenvalues are roots of unity or zero then $h(a_n) = O(\log\log\ n)$.

Now luckily Michael Coons, Kevin Hare, and I had shown that there was a fundamental gap for the growth of coefficients of integer regular series between $\log(n)$ and $\mathrm{O}(1)$ (automatic case) and a variation of this argument can be used to show that if $h(a_n) = O(\log\log\ n)$ then either $h(a_n) \in O \cap \Omega(\log\log\ n)$ or $h(a_n) = \mathrm{O}(1)$ and is automatic.

One last corollary

Theorem

Let K be a field of characteristic 0 and let $F(z) = \sum a_n z^n$ be a k-Mahler power series. Then a_n is k-automatic if and only if $\{a_n : n \geq 0\}$ is finite.

This is not true in characteristic p.

Why is this true? You can first use a specialization argument to assume that the a_n are in $\bar{Q}$. Then $h(a_n) = \mathrm{O}(1) \implies F(z)$ is automatic!

Thanks!