THE FIRST BASE BELOW A QUARTER

CARLO MITCHENER

MrlyProd, Inc.
First published 2026-09-07, revised 2026-09-08

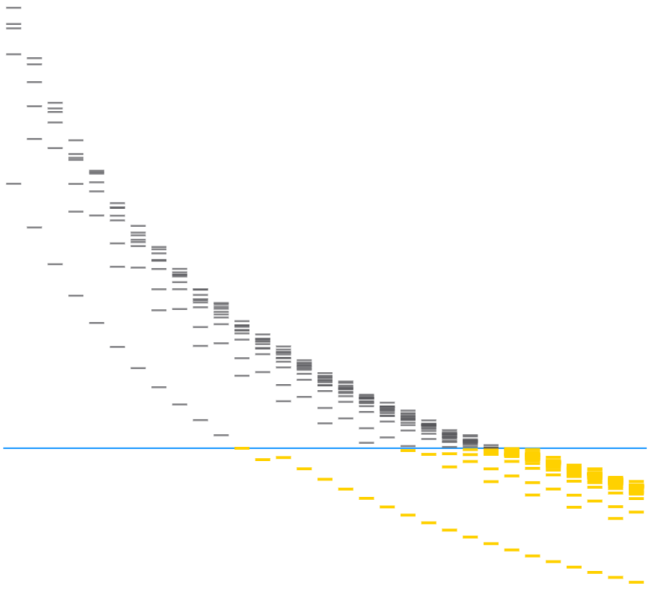

Abstract. Drop one digit from a base and look at the numbers you can still write. How much of the Fourier mass of that set survives is measured by a single number, its ℓ^1 exponent, and a sieve turns on when the number falls below a quarter. This paper locates the threshold. Base 21 with the digit 0 removed is the first base carrying any such set below a quarter; base 34 is the first base whose every excluded digit clears; and a bound that names the base and forgets the digit proves clearance for every base from 126 up, so the family floor 34 is exact.

1. Introduction

Fix a base q and cross out one digit. What is left is a set of integers: those you can still write, the ones whose base-q expansion never uses the missing digit. In base 10 with the digit 7 removed these are $1, 2, \dots, 6, 8, 9, 10, \dots, 69, 80, \dots$, and there are 9^L of them below 10^L, so the set is thin but not very thin. Sets of this shape are the standard test bed for analytic number theory on a fractal: Maynard proved that infinitely many of them are prime in base 10 [7], and everything in that proof passes through the set by way of its Fourier transform.

The transform is one function on the circle. Write

$$\widehat{F}(t) \;=\; \frac{1}{q-1} \sum_{\substack{0 \le a < q \\ a \neq a_0}} e(at), \qquad e(x) = e^{2\pi i x},$$

for the normalised character sum of the surviving digits $F = \{0, \dots, q-1\} \setminus \{a_0\}$, and $\widehat{F}_L(t) = \prod_{j<L} \widehat{F}(q^j t)$ for the transform of the L-digit strings. A sieve does not read $\widehat{F}_L$ pointwise; it reads how much of it survives when the whole grid of q^L frequencies is added up in absolute value. The growth rate of that grid sum is one number, the ℓ^1 *exponent* α_1 of Definition 2.2, and it is the number this paper is about. Small α_1 means the digit restriction cancels; α_1 close to 1 means it barely does.

The published values sit near a third. In base 10 with one digit removed the bound behind [7] is $\alpha_1 \le 27/77 = 0.35065$, and Karwatowski [4] carries the same shape of bound to every base $b \ge 10$. A quarter is a different demand, and the question is where it is first met.

The exponent has been read for several different purposes, and the readings do not compete. Erdős, Mauduit and Sárközy [3] and Konyagin [5] distribute missing-digit sets in residue classes; [7] and [8] consume the exponent as a level of distribution for primes, [9] as one of four Fourier norms in a Bombieri-Vinogradov theorem, and Leng and Sawhney [6] as the single input $g^{\varepsilon k}$ that settles ternary Goldbach on such a set at large base. Chow, Varjú and Yu [2] study the same quantity against the threshold $1/2$ rather than $1/4$, and their Proposition 2.4 puts it above $1/2$ for every base $b \ge 5$ with one digit removed, which is the opposite side of a different line and consistent with everything below. What none of these fixes is the least base at which a given threshold is met, which is the question here.

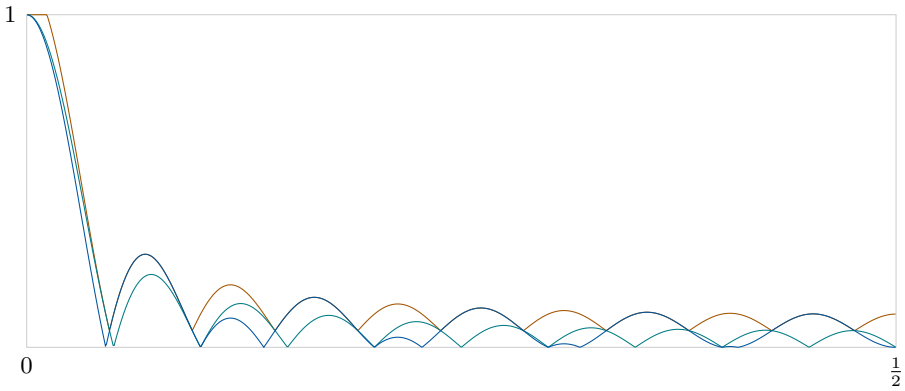

FIGURE 1. Base 21. The two curves are $|\widehat{F}|$ with the digit 0 removed and with the digit 10 removed, on the half period $[0, \frac{1}{2}]$; the envelope above them is the digit-blind majorant u_q of Lemma 4.1, which names the base and forgets the digit. The exponent α_1 is the growth rate of the grid sums of the lower curves; the theorem of this paper is what the envelope alone already forces.

Remark 1.1 (Why a quarter, and why this is only motivation). A quarter is the threshold at which one conditional route to a Möbius asymptotic over a digit-restricted set switches on. Expanding the indicator of the set in additive characters modulo q^L costs exactly the ℓ^1 mass of $\widehat{F}_L$, and feeding each frequency into the uniform bound $\max_\theta |\sum_{n\le x} \mu(n)e(n\theta)| \ll_\varepsilon x^{3/4+\varepsilon}$ that the generalized Riemann hypothesis buys [1] produces a power saving against the set's own counting function precisely when $\alpha_1 < 1/4$. That implication is not proved here and nothing below depends on it: it is the reason the number $1/4$ is interesting, not a theorem of this paper. Every statement in Section 3 is a statement about α_1 alone.

Three results answer the question, and they answer it in the two different senses the question has.

Theorem 1.2 (The digit-uniform bound)**.** *For every base $q \ge 126$ and every excluded digit a_0, $\alpha_1 < 1/4$. The bound is uniform in the digit: it is deduced from a majorant that names q and forgets a_0. It is also sharp for its own method, in that the same chain fails at $q = 125$. Theorem 3.1 below is the explicit form, with the constant it depends on.*

The second sense is computational, and it is where the two floors are.

- *The first base.* $q = 21$ with the digit 0 removed is the least base carrying any one-missing-digit set with $\alpha_1 < 1/4$. It clears by about 10^{-5}. Every one of the 108 distinct sets of every base $3 \le q \le 20$ has $\alpha_1 > 1/4$ (Fact 3.4).

- *The first family.* $q = 34$ is the least base whose *every* excluded digit clears; all 17 of its distinct sets do, while each of the bases $21 \le q \le 33$ still carries a digit that misses (Fact 3.6). Together with Theorem 1.2 and a certified sweep of the bases in between, every base $q \ge 34$ clears at every excluded digit, and 34 is exact.

The two floors are not the same number because clearing the threshold is easier at some excluded digits than at others: removing 0 or $q-1$ is cheapest, removing a digit near the middle is dearest, and between 21 and 33 the cheap digits have crossed while the dear ones have not. The cover image is that whole picture: one dash per set, one column per base from 10 to 40, dark above the quarter line and light below it, one light dash first at base 21 and every dash of every column light from base 34 on.

The proof of Theorem 1.2 is elementary and self-contained. It rests on one observation: the triangle inequality replaces the excluded digit by a phase of modulus one, leaving a majorant built from the Dirichlet kernel that no longer mentions which digit was removed. The level product of that majorant expands over subsets of the digit positions, every maximal run of consecutive positions telescopes into a single Dirichlet kernel at a higher modulus, and what is left is a sum of 2^N Lebesgue sums, each of which is bounded by the classical Lebesgue constant of its own modulus. The bookkeeping closes into a cubic, and the cubic clears a quarter from $q = 126$ up.

The computations are of a different kind and are labelled as such throughout: a transfer matrix on windows of digits brackets α_1 from both sides in a way that a machine can check, and Section 5 names the script, its runtime and its exact finite domain. No claim below leans on an unproved lemma.

2. The exponent and the machine that brackets it

Throughout, $q \ge 3$ is an integer base, $a_0 \in \{0, \dots, q-1\}$ is the excluded digit, $F = \{0, \dots, q-1\} \setminus \{a_0\}$ and $k = |F| = q-1$. We write $e(x) = e^{2\pi i x}$ and

$$D_M(t) \;=\; \sum_{0 \le a < M} e(at) \cdot e\Big(-\tfrac{M-1}{2} t\Big) \;=\; \frac{\sin(\pi M t)}{\sin(\pi t)}$$

for the Dirichlet kernel, with $D_M(t) = M$ at integer t.

Definition 2.1 (The transform)**.**

$$\widehat{F}(t) = \frac{1}{q-1} \sum_{a \in F} e(at), \qquad \widehat{F}_L(t) = \prod_{0 \le j < L} \widehat{F}(q^j t).$$

$\widehat{F}_L$ is the normalised transform of the set of L-digit strings over F, because the digits of a string are independent. Both are 1-periodic and both satisfy $|\widehat{F}(1-t)| = |\widehat{F}(t)|$, since the coefficients are real.

Definition 2.2 (The shifted grid sum and the exponent)**.** For $N \geq 1$ and $x \in \mathbb{R}$ put

$$\Sigma_N(x) \;=\; \sum_{0\leq i<q^N} \prod_{0\leq j<N} \Big|\widehat{F}\Big(q^j\Big(x+\tfrac{i}{q^N}\Big)\Big)\Big| \;=\; \sum_{0\leq i<q^N} \Big|\widehat{F}_N\Big(x+\tfrac{i}{q^N}\Big)\Big|,$$

and $S_N = \max_x \Sigma_N(x)$. The ℓ^1 *exponent* of the pair (q, a_0) is

$$\alpha_1 \;=\; \alpha_1(q,a_0) \;=\; \lim_{N\to\infty} \frac{1}{N}\log_q S_N \;=\; \inf_{N\geq 1} \frac{1}{N}\log_q S_N.$$

The limit exists, and equals the infimum, because $\log S_N$ is subadditive; this is Lemma 4.2. The exponent is what a sieve pays: the q^L frequencies of the level-L set carry total mass $k^L S_L \leq k^L q^{L\alpha_1+o(L)}$, against the trivial $k^L q^L$.

Remark 2.3 (Three exponents, one bracket). Two neighbours of α_1 appear in the literature and in the estimates below: the exponent α_1^- of $\min_x \Sigma_N(x)$, and the exponent α_1^0 of $q^L \int_0^1 |\widehat{F}_L|$. Substituting $t = (y+i)/q^L$ shows that $q^L \int_0^1 |\widehat{F}_L|$ is the average of Σ_L over one cell, so

$$\alpha_1^- \;\leq\; \alpha_1^0 \;\leq\; \alpha_1.$$

The upper certificate of Proposition 2.5 bounds all three from above, and its lower certificate bounds α_1^0, hence also α_1, from below. Bounds on α_1 are the strongest of the three, and α_1 is what a sup-over-shifts sieve estimate actually consumes, which is why it is the definition used here.

Definition 2.4 (The digit symmetry)**.** $|\widehat{F}|$, and hence α_1, is unchanged by $a_0 \mapsto q-1-a_0$. A base therefore carries $\lceil q/2 \rceil$ distinct one-missing-digit sets, and the pair coincides only when q is odd and $a_0 = (q-1)/2$.

Indeed $\sum_{a\in F} e(a(1-t)) = \overline{\sum_{a\in F} e(at)}$ after the substitution $a \mapsto q-1-a$, which fixes F exactly when the excluded digit is fixed.

For computation the transform is evaluated in closed form rather than as a character sum. Since $\sum_{0\leq a<q} e(at) = e\Big(\frac{q-1}{2}t\Big)\, D_q(t)$,

$$\begin{aligned} |\widehat{F}(t)| \;&=\; \frac{1}{q-1}\Big|D_q(t) - e\Big(\Big(a_0-\tfrac{q-1}{2}\Big)\,t\Big)\Big| \\ &=\; \frac{1}{q-1}\sqrt{D_q(t)^2 - 2D_q(t)\cos(\pi(2a_0-q+1)t)+1}, \end{aligned} \tag{1}$$

three trigonometric calls instead of $q-1$ complex exponentials.

Proposition 2.5 (The window machine)**.** *Fix $n \geq 2$ and split the circle into the q^n cells $I_w = [wq^{-n}, (w+1)q^{-n})$. Put*

$$G^+(w) = \sup_{t\in I_w} |\widehat{F}(t)|, \qquad G^-(w) = \inf_{t\in I_w} |\widehat{F}(t)|,$$

and let $M^{\pm}$ be the nonnegative $q^{n-1}\times q^{n-1}$ matrices

$$(M^{\pm}y)(v) \;=\; \sum_{0\leq c<q} G^{\pm}(vq+c)\, y\big((vq+c) \bmod q^{n-1}\big).$$

Then $\alpha_1 \le \log_q \rho(M^+)$, and $\alpha_1 \ge \log_q \rho(M'')$ for the restriction M'' of M^- to any subset of states. Moreover, for any $y > 0$, $M^+ y \le \mu y$ componentwise gives $\rho(M^+) \le \mu$, and $M'' y \ge \mu y$ componentwise gives $\rho(M'') \ge \mu$.

The proof is in Section 4. The last sentence is the Collatz-Wielandt test, and it is what makes the machine certifiable: a power iteration supplies a candidate vector y, and one further pass over the matrix turns that candidate into a rigorous one-sided bound. A poor vector weakens the bound and can never break it. The restriction clause matters in practice: $|\widehat{F}|$ vanishes at finitely many points, a cell containing such a point has $G^- = 0$, and a state all of whose successors vanish would otherwise force the lower bound to zero. Restricting to the support of the iterate repairs this at no cost.

The suprema and infima over a cell are taken by a sub-scan of m equally spaced points together with the Lipschitz correction $\Lambda_{q,a_0}/(2mq^n)$ of Eq. (3) below. The correction is derived, never fitted, and every certificate below states the pair (n, m) it used.

3. Results

Every statement below is tagged. A *theorem* or *lemma* carries a full proof in Section 4; a *fact* is a finite computation, stated with its exact domain and with the script that reruns it named in Section 5; a *conjecture* is open and says so.

3.1. The digit-uniform bound.

Set $c_1 = 2/\pi$ and $c_0 = 0.98$, and for $q \ge 3$ let $z_q > 1$ be the unique root of

$$(z-1)^3 \;=\; c_1(\log q)\, z \;+\; c_0(z-1). \tag{2}$$

Theorem 3.1. *For every base $q \ge 100$ and every excluded digit a_0,*

$$\alpha_1(q, a_0) \;\le\; \log_q\!\left(\frac{z_q\, q}{q-1}\right).$$

In particular $\alpha_1 < 1/4$ for every $q \ge 126$ and every excluded digit: the right side reads 0.249808 at $q = 126$, 0.231905 at $q = 200$, 0.187674 at $q = 1000$ and 0.108949 at $q = 10^6$, and it tends to 0 like $\log\log q/\log q$. The bound is uniform in a_0; it comes from a majorant that does not mention which digit was removed.

The threshold 126 is where the chain closes and not where the phenomenon starts. Two remarks fence it in.

Remark 3.2 (The chain is sharp for itself). The inequality $z_q < q^{1/4}(1 - 1/q)$ that Theorem 3.1 needs holds at $q = 126$ and fails at $q = 125$, so 126 is exactly the threshold of this chain of estimates. It is not the threshold of the majorant: see Section 3.3.

Where the threshold sits depends on c_0, and $c_0 = 0.98$ is a rounding up of the value 0.9724870 that Lemma 4.7(iii) actually proves. The rounding is harmless but the constant is not: the threshold is exactly 126 for every c_0 in

q	a_0	$\alpha_1 >$	q	a_0	$\alpha_1 >$
21	9	0.27706	28	13	0.25907
22	10	0.27415	29	9	0.25734
23	7	0.27155	30	14	0.25512
24	11	0.26847	31	14	0.25293
25	11	0.26561	32	10	0.25184
26	8	0.26380	33	15	0.25061
27	12	0.26083			

TABLE 1. One excluded digit per base that still misses the quarter, certified at three window digits and sub-scan 8. Base 33 is the last and the tightest, and the only row needing four window digits: there the digit 15 gives $\alpha_1 > 0.2506145$.

$(0.9679, 0.9961)$ and rises to 127 above that interval, while below it the same certificate would also clear $q = 125$. In particular the constant 0.9625272 that Fact 4.8 observes numerically sits under 0.9679, so the numerics alone would put the threshold at 125; 126 is the threshold of what is proved, and the whole distance between the two is the odd-M slack in Lemma 4.7(ii).

Remark 3.3 (The cost of forgetting the digit). The triangle inequality that produces the majorant discards a phase of modulus one, so the majorant is attained only where that phase is antipodal to the Dirichlet kernel and is strictly larger than $|\widehat{F}|$ everywhere else. The uniform bound is therefore weaker than the per-digit one at every base, by construction and never by accident.

3.2. **The two floors.**

Fact 3.4 (The first base)**.** Base 21 with the digit 0 removed satisfies $\alpha_1 < 0.2499917$, and it is the least base carrying any one-missing-digit set with $\alpha_1 < 1/4$: every one of the 108 distinct sets of every base $3 \le q \le 20$ satisfies $\alpha_1 > 1/4$.

Domain. The upper bound is the window machine of Proposition 2.5 at $n = 5$ window digits and sub-scan $m = 2$, clearing the threshold by 8.3×10^{-6}. Four window digits do not decide it: at $n = 4$, $m = 8$ the same machine returns the bracket $[0.2498658, 0.2500871]$, which straddles $1/4$. The 108 lower bounds are the same machine on the infimum side, 41 of them settled at $n = 2$ and 67 at $n = 3$, all with $m = 8$; the least of the certified lower bounds is 0.2502668.

Fact 3.5 (Each base up to 33 still misses somewhere)**.** Every base $21 \le q \le 33$ carries an excluded digit with $\alpha_1 > 1/4$. One witness for each is recorded in Table 1.

Domain. The window machine on the infimum side at $n = 3$, $m = 8$, except base 33 which needs $n = 4$.

Fact 3.6 (The first family)**.** All 17 distinct one-missing-digit sets of base 34 satisfy $\alpha_1 < 1/4$. With Fact 3.5, base 34 is the least base whose every excluded digit clears the threshold.

Domain. The window machine on the supremum side. Sixteen of the seventeen digits clear at $n = 3$, $m = 8$; the digit 16 needs $n = 4$, $m = 2$. The bounds run from $\alpha_1 < 0.2257743$ at the digit 0 to $\alpha_1 < 0.2497573$ at the digit 13, the largest of the seventeen.

Corollary 3.7 (The family floor is exact)**.** *Every base $q \geq 34$ satisfies $\alpha_1(q, a_0) < 1/4$ at every excluded digit a_0, and 34 is the least such base.*

Proof. Base 34 is Fact 3.6. The bases $q \geq 126$ are Theorem 3.1. The bases $35 \leq q \leq 125$ are the certified sweep recorded in Fact 3.8. Minimality is Fact 3.5 together with Fact 3.4, which rule out every base below 34. □

Fact 3.8 (The sweep between the floors)**.** Every one of the 3663 distinct one-missing-digit sets of every base $35 \leq q \leq 125$ satisfies $\alpha_1 < 1/4$, each certified at the shortest window of two, three or four digits that clears; the bases $35 \leq q \leq 57$ need three window digits and every $q \geq 58$ clears at two.

Domain. This row is not rerun by the script of this paper, which would need hours of plain Python; it comes from the interval-arithmetic generator named in Section 5, whose certificates are the tests of Proposition 2.5 carried out in directed rounding. It is the one computational input here that the accompanying script does not reproduce, and Corollary 3.7 is the only statement that uses it.

Fact 3.9 (The published exponent is an upper bound, not the constant)**.** For base 10 with the digit 5 removed the window machine at $n = 5$, $m = 8$ certifies

$$\alpha_1 \in [0.3505101,\ 0.3506467] \ < \ \tfrac{27}{77} = 0.3506494.$$

The published 27/77 of [7] is therefore a finite-window upper bound on the ℓ^1 exponent and not the exponent itself. The generator of Fact 3.8, at seven window digits and the slightly wider box the sieve literature uses, narrows the same quantity to $[0.3505775, 0.3505797]$.

3.3. **What the uniform route cannot reach.** Theorem 3.1 is a statement about q with the digit forgotten, and Remark 3.3 says that costs something. Two measurements say how much.

Fact 3.10 (The numeric ceiling of the majorant)**.** Run the window machine of Proposition 2.5 on the majorant u_q of Lemma 4.1 in place of $|\widehat{F}|$. At four window digits the resulting digit-blind bound reads 0.339085 at $q = 21$ and 0.297296 at $q = 34$, against 0.340029 and 0.297650 at three, and it first falls below 1/4 near $q = 75$. The exact crossing base is sensitive to the window length and to the cell convention, and this paper claims no value for it; what the row is for is the order of magnitude of the gap. Around 75 is where the majorant itself stops working, 126 is where the chain of Section 4 proves

it works uniformly in q, and the distance between them is slack in the run decomposition and not in the majorant.

Fact 3.11 (The uniform bound is never stronger)**.** Where both apply, the digit-blind window bound exceeds the per-digit one at every base tested and never falls below it. The digit-blind column is the generator; the per-digit column is the script of this paper at the window it certifies with, so the two are read at the same base and digit but not by the same code, and the comparison is only ever used qualitatively.

base and digit	digit-blind	per-digit	window
$q = 10$, $a_0 = 5$	0.440960	0.3506467	$n = 5$
$q = 21$, $a_0 = 0$	0.339085	0.2500871	$n = 4$
$q = 34$, $a_0 = 16$	0.297296	0.2494348	$n = 4$

The gap closes as the base grows, from 0.138 at $q = 9$ to 0.048 at $q = 34$ in the generator's own rows, which is the discarded phase costing $1/(q-1)$.

The digit-blind numbers of **??** 3.10**??** 3.11 come from the same generator as Fact 3.8; neither fact is used in the proof of anything above.

3.4. **Open problems.** Three questions are left, and one guess is worth recording because it is falsifiable.

Conjecture 3.12 (The cheapest digit is always an end digit)**.** *For every base $q \geq 3$, $\alpha_1(q, a_0)$ is minimised at $a_0 = 0$ (equivalently, by Definition 2.4, at $a_0 = q - 1$).*

Evidence. *It holds at every base and digit computed here. At base* 34 *the seventeen bounds run from* 0.2257743 *at the digit* 0 *to* 0.2497573 *at the digit* 13*, a spread of* 0.024*, and the digit* 0 *is well clear of the rest.* Failure mode. *The obvious strengthening, that α_1 increases monotonically as the excluded digit moves from the end toward the middle, is false: at base* 34 *the digit* 11 *reads* 0.2475920*, below both its neighbours* 10 *and* 12*, and the largest of the seventeen sits at the digit* 13 *rather than at the middle digit* 16*. The ordering of the interior digits is not monotone, and any proof of the conjecture has to separate the end digit from the interior without ordering the interior.*

First, the gap between 75 and 126. Fact 3.10 says the majorant of Lemma 4.1 numerically clears 1/4 from base 75 on, while the chain of Section 4 only proves clearance from 126. Every step of that chain is an inequality with a measurable loss, and the largest is the peel of Lemma 4.4, which throws away the correlation between runs at different digit positions. A sharper peel, or a direct treatment of the two-run terms, should move 126 toward 75. Below 75 nothing can be recovered without keeping the digit.

Second, the gap between 34 and 21. Between those two bases the cheap digits have crossed the threshold and the dear ones have not, and no proof separates them; Table 1 is thirteen separate computations with no pattern extracted. A proved lower bound for the middle excluded digit, uniform in q, would replace that table with an argument.

Third, the exponent itself. Nothing here computes α_1; every statement is a bracket, and the brackets narrow only as fast as the window grows. At base 21 with the digit 0 the exponent is pinned inside an interval of width 6×10^{-7} around 0.24998, which clears $1/4$ by about one part in 10^4 of that interval's own width. Whether α_1 is ever rational, or algebraic, or equal to the growth rate of some finite automaton attached to the base, is untouched. The transfer matrices of Proposition 2.5 have entries that are suprema of a transcendental function over cells, so they carry no visible algebraic structure, and the question of what kind of number α_1 is remains open even for the smallest interesting base.

4. Proofs

4.1. **Domination.**

Lemma 4.1 (The digit-blind majorant)**.** *For every base $q \geq 3$, every excluded digit a_0 and every real t,*

$$|\widehat{F}(t)| \;\leq\; u_q(t) \;:=\; \min\left(1,\; \frac{|D_q(t)|+1}{q-1}\right).$$

The right side does not mention a_0.

Proof. Summing the full residue system, $\sum_{0\leq a<q} e(at) = e\Big(\frac{q-1}{2}t\Big)\, D_q(t)$, an identity that holds at integer t with the convention $D_q = q$ there. Subtracting the excluded term,

$$(q-1)\widehat{F}(t) \;=\; e\Big(\tfrac{q-1}{2}t\Big)\, D_q(t) \;-\; e(a_0 t),$$

which is Eq. (1), and the triangle inequality gives $(q-1)|\widehat{F}(t)| \leq |D_q(t)|+1$. The second bound is immediate: $\widehat{F}(t)$ is the average of $q-1$ complex numbers of modulus one. □

The cost of the lemma is exactly one phase. Equality in the triangle inequality needs $e\Big(\Big(a_0 - \frac{q-1}{2}\Big)\, t\Big)$ to point opposite to $e\Big(\frac{q-1}{2}t\Big)\, D_q(t)$, which happens on a set of measure zero, so u_q is strictly larger than $|\widehat{F}|$ almost everywhere. That is Remark 3.3, and Fig. 1 is a picture of it.

4.2. **The exponent, and the machine that brackets it.**

Lemma 4.2 (Subadditivity)**.** *$S_{N+M} \leq S_N S_M$ for all $N, M \geq 1$; hence $\frac{1}{N}\log_q S_N$ converges to its infimum and Definition 2.2 is well posed. Moreover Σ_N is q^{-N}-periodic.*

Proof. Periodicity is clear: replacing x by $x + q^{-N}$ permutes the summands cyclically. For subadditivity write $i = b + q^N a$ with $0 \leq b < q^N$, $0 \leq a < q^M$, and $t = x + iq^{-(N+M)}$. Since $\widehat{F}$ is 1-periodic, $q^M t \equiv q^M x + bq^{-N} \pmod 1$, so the last N factors of $\prod_{j<N+M} |\widehat{F}(q^j t)|$ depend on b alone and reproduce the

summand of $\Sigma_N(q^M x)$ indexed by b. The first M factors are the summand of $\Sigma_M(x + bq^{-(N+M)})$ indexed by a. Summing over a first,

$$\begin{aligned}\Sigma_{N+M}(x) &= \sum_{b<q^N} \left|\widehat{F}_N(q^M x + bq^{-N})\right| \Sigma_M(x + bq^{-N-M}) \\ &\le S_M \Sigma_N(q^M x) \le S_N S_M. \qquad \square\end{aligned}$$

Proof of Proposition 2.5. Write the base-q digits of a cell index $w < q^\ell$ as $w = (w_1 \cdots w_\ell)$, most significant first, so that I_w is the set of t whose expansion starts $0.w_1 \cdots w_\ell$. For such t and $0 \le j < \ell$, $q^j t$ reduced mod 1 lies in the level-$(\ell - j)$ cell indexed by $w_{j+1} \cdots w_\ell$. If $j + n \le \ell$ that cell is contained in the level-n cell $v_j(w) := (w_{j+1} \cdots w_{j+n})$.

The upper bound. Take $\ell = L$. For $t \in I_w$ and $j \le L - n$ the factor $|\widehat{F}(q^j t)|$ is at most $G^+(v_j(w))$, and the remaining $n - 1$ factors are at most 1. Hence

$$U_L := \sum_{w<q^L} \sup_{I_w} |\widehat{F}_L| \le \sum_{w<q^L} \prod_{j=0}^{L-n} G^+(v_j(w)) = \mathbf{1}^{\mathsf{T}} (M^+)^{L-n+1} \mathbf{1},$$

the identity being the definition of M^+: a state is a string of $n - 1$ digits, appending a digit c to the state v forms the window $vq + c$ with weight $G^+(vq + c)$ and moves to the state $(vq + c) \bmod q^{n-1}$, and summing over w is summing over all initial states and all appended digits. Since $x + iq^{-L}$ lies in the cell I_i for $x \in [0, q^{-L})$, periodicity gives $S_L \le U_L$, and Gelfand's formula gives $\limsup_L U_L^{1/L} \le \rho(M^+)$. Therefore $\alpha_1 \le \log_q \rho(M^+)$.

The lower bound. Take $\ell = L + n - 1$. Now every j in $0 \le j < L$ satisfies $j + n \le \ell$, so on the cell I_w every one of the L factors of $|\widehat{F}_L|$ is at least $G^-(v_j(w))$, with no truncated window left over. Hence

$$\int_0^1 |\widehat{F}_L| \ge \sum_{w<q^\ell} q^{-\ell} \inf_{I_w} |\widehat{F}_L| \ge q^{-(L+n-1)} \mathbf{1}^{\mathsf{T}} (M^-)^L \mathbf{1}.$$

Substituting $t = (y + i)q^{-L}$ shows $q^L \int_0^1 |\widehat{F}_L| = \int_0^1 \Sigma_L(yq^{-L})\, dy \le S_L$, so $S_L \ge q^{-(n-1)} \mathbf{1}^{\mathsf{T}} (M^-)^L \mathbf{1}$. Let S' be any set of states and M'' the restriction of M^- to S', that is, the matrix obtained by keeping only the rows and columns in S'. Every entry of M^- is nonnegative, so $\mathbf{1}^{\mathsf{T}} (M^-)^L \mathbf{1} \ge \mathbf{1}^{\mathsf{T}} (M'')^L \mathbf{1}$. If y is positive on S' and $M''y \ge \mu y$ componentwise then $(M'')^L y \ge \mu^L y$, whence

$$\mathbf{1}^{\mathsf{T}} (M'')^L \mathbf{1} \ge \frac{1}{\max_{S'} y} \mathbf{1}^{\mathsf{T}} (M'')^L y \ge \mu^L \frac{\sum_{S'} y}{\max_{S'} y}.$$

The final factor does not depend on L, so $\alpha_1 = \lim_L \frac{1}{L} \log_q S_L \ge \log_q \mu$. Taking the best such μ gives $\alpha_1 \ge \log_q \rho(M'')$, since a nonnegative irreducible matrix admits its own Perron vector as a test vector and a reducible one is handled by restricting further.

The tests. If $y > 0$ and $M^+ y \le \mu y$ then $\rho(M^+) \le \mu$: iterating gives $(M^+)^L y \le \mu^L y$, so all entries of $(M^+)^L$ are $O(\mu^L)$ and Gelfand's formula applies. The lower test is the display above. □

The Lipschitz correction is the last ingredient. Differentiating Definition 2.1 termwise, the derivative in t of $\frac{1}{q-1}\sum_{a\in F} e(at)$ has modulus at most $\frac{2\pi}{q-1}\sum_{a\in F} a$, and $\sum_{a\in F} a = \frac{q(q-1)}{2} - a_0$ exactly, so

$$|\widehat{F}(t) - \widehat{F}(t')| \;\le\; \Lambda_{q,a_0}\,|t-t'|, \qquad \Lambda_{q,a_0} \;:=\; \pi q - \frac{2\pi a_0}{q-1}, \tag{3}$$

and the reverse triangle inequality carries the same Lipschitz constant to $|\widehat{F}|$. Two neighbouring points of a sub-scan of m equally spaced points in a cell of width q^{-n} are q^{-n}/m apart and the extreme points sit $q^{-n}/(2m)$ from the ends, so the supremum and the infimum of $|\widehat{F}|$ over the cell differ from the extremes of the sub-scan by at most $\Lambda_{q,a_0}/(2mq^n)$. The digit-dependent form matters: the sharper constant is what the accompanying script uses, and it is what Fact 3.9 needs, since at base 10 missing 5 it reads 27.926 where the digit-blind πq reads 31.416. At $a_0 = 0$ the two agree, so Fact 3.4 is the same under either. Adding the correction outward on both sides turns the sub-scan into a rigorous enclosure of $G^{\pm}$, and monotonicity of the Perron root in the entries turns rigorous enclosures of the entries into rigorous bounds on μ.

4.3. **The run identity.** Write $v_q(t) = (|D_q(t)| + 1)/(q-1)$, so that $u_q = \min(1, v_q) \le v_q$, and put

$$\Sigma_N^v(x) \;=\; \sum_{0\le i<q^N} \;\prod_{0\le j<N} v_q\Big(q^j\Big(x + \tfrac{i}{q^N}\Big)\Big) \;\ge\; \Sigma_N(x)$$

by Lemma 4.1. A subset $E \subseteq \{0,\dots,N-1\}$ splits uniquely into maximal runs of consecutive positions; write (s,l) for the run occupying $s, s+1, \dots, s+l-1$.

Lemma 4.3 (The run identity)**.** *For every x,*

$$(q-1)^N\,\Sigma_N^v(x) \;=\; \sum_{E\subseteq\{0,\dots,N-1\}} \;\sum_{0\le i<q^N} \;\prod_{(s,l)\ \text{run of}\ E} \Big|D_{q^l}\Big(q^s\Big(x + \tfrac{i}{q^N}\Big)\Big)\Big|.$$

Proof. Expand $\prod_{j<N}(|D_q(q^j t)| + 1)$ over subsets E of the positions at which the kernel factor is taken. Inside a run $s, \dots, s+l-1$ the factors telescope: with $t' = q^s t$,

$$\prod_{r=0}^{l-1} \frac{\sin(\pi q^{r+1} t')}{\sin(\pi q^r t')} \;=\; \frac{\sin(\pi q^l t')}{\sin(\pi t')} \;=\; D_{q^l}(t'),$$

and taking absolute values gives the claim. Distinct runs occupy disjoint blocks, so the factorisation is unambiguous. □

Every term of Lemma 4.3 is a sum of Dirichlet kernels over an arithmetic progression of points, that is, a Lebesgue sum. Let

$$L_M \;=\; \max_{x\in\mathbb{R}} \sum_{0\le j<M} \left|D_M\Big(x+\tfrac{j}{M}\Big)\right|, \qquad \lambda_l = \frac{L_{q^l}}{q^l}.$$

4.4. **The peel.**

Lemma 4.4 (The peel)**.** *Let $E\subseteq\{0,\dots,N-1\}$, and let $(s_1,l_1),\dots,(s_r,l_r)$ be its maximal runs, ordered so that $s_1<\cdots<s_r$. Put $t_i = x+iq^{-N}$. Then for every real x,*

$$\sum_{0\le i<q^N} \prod_{k=1}^{r} \left|D_{q^{l_k}}(q^{s_k}t_i)\right| \;\le\; q^N \prod_{k=1}^{r} \lambda_{l_k}.$$

Proof. First a sub-lemma: for integers $M\ge l\ge 1$ and every real y,

$$\sum_{0\le c<q^M} \left|D_{q^l}\Big(y+\tfrac{c}{q^M}\Big)\right| \;\le\; q^M\lambda_l. \tag{4}$$

Indeed, write $c=\rho+q^{M-l}c'$ with $0\le\rho<q^{M-l}$ and $0\le c'<q^l$, so that $c/q^M=\rho/q^M+c'/q^l$. For fixed ρ the inner sum runs over q^l points spaced q^{-l} apart, hence is at most $L_{q^l}=q^l\lambda_l$; summing over the q^{M-l} values of ρ gives Eq. (4).

Now set $a_s = i \bmod q^{N-s}$. Since D_M is 1-periodic, $q^s(x+iq^{-N})\equiv q^sx+a_sq^{-(N-s)} \pmod 1$, so the k-th factor depends on i only through a_{s_k}, and a_{s_1} determines every a_{s_k}. The top s_1 digits of i enter no factor, so the left side equals q^{s_1} times the sum over $a_{s_1}<q^{N-s_1}$, and it suffices to prove

$$\sum_{a_{s_1}<q^{N-s_1}} \prod_{k} \left|D_{q^{l_k}}\Big(q^{s_k}x+\tfrac{a_{s_k}}{q^{N-s_k}}\Big)\right| \;\le\; q^{N-s_1}\prod_k \lambda_{l_k}.$$

Induct on r. For $r=1$ this is Eq. (4) with $M=N-s_1\ge l_1$, which holds because the run lies inside $\{0,\dots,N-1\}$. For $r\ge 2$ write $a_{s_1}=a_{s_2}+q^{N-s_2}w$ with $0\le w<q^{s_2-s_1}$. The factors $k\ge 2$ depend on a_{s_2} alone, and the first factor is

$$\left|D_{q^{l_1}}\Big(z+\tfrac{w}{q^{s_2-s_1}}\Big)\right|, \qquad z=q^{s_1}x+\tfrac{a_{s_2}}{q^{N-s_1}} \text{ fixed.}$$

Runs are maximal, so they are separated by at least one absent position and $s_2-s_1\ge l_1+1>l_1$. Applying Eq. (4) with $M=s_2-s_1$ bounds the inner sum by $q^{s_2-s_1}\lambda_{l_1}$, and the induction hypothesis bounds the outer sum by $q^{N-s_2}\prod_{k\ge2}\lambda_{l_k}$. The product of the two is $q^{N-s_1}\prod_k\lambda_{l_k}$. □

Lemma 4.5 (The count)**.** *Put $A_0=1$ and $A_N=\sum_{E\subseteq\{0,\dots,N-1\}}\prod_{\text{runs}}\lambda_l$. Then*

$$A_N \;=\; A_{N-1} \;+\; \sum_{l=1}^{N-1} \lambda_l A_{N-1-l} \;+\; \lambda_N \qquad (N\ge 1),$$

and if $z>1$ is the unique root of $z=1+\sum_{l\ge1}\lambda_l z^{-l}$ then $\limsup_N A_N^{1/N}\le z$. Moreover z is nondecreasing in the sequence (λ_l).

Proof. Split by the first position. If $0 \notin E$ the rest is an arbitrary subset of a set of $N-1$ positions and contributes A_{N-1}. If $0 \in E$ let $l \geq 1$ be the length of the initial run; then position l is absent (when $l < N$) and the rest is an arbitrary subset of the remaining $N-1-l$ positions, contributing $\lambda_l A_{N-1-l}$; the case $l = N$ contributes λ_N.

Let $\Lambda(x) = \sum_{l\geq 1} \lambda_l x^l$. Since $\lambda_l \leq L_{q^l}/q^l$ grows at most linearly in l, Λ has radius of convergence 1 and tends to ∞ as $x \uparrow 1$. Multiplying the recursion by x^N and summing gives $A(x)\,(1 - x - x\Lambda(x)) = 1+\Lambda(x)$. The map $\varphi(z) = 1+\Lambda(1/z)$ is finite, continuous and strictly decreasing on $(1,\infty)$, with $\varphi \to \infty$ at 1^+ and $\varphi \to 1$ at ∞, so $z - \varphi(z)$ has a unique zero $z > 1$; and φ increases pointwise when the λ_l do, which gives the monotonicity. For $|x| < 1/z$, $|x + x\Lambda(x)| \leq |x|(1 + \Lambda(|x|)) < z^{-1}\varphi(z) = 1$, so the denominator does not vanish and A is analytic on that disc. The coefficients A_N are nonnegative, so the radius of convergence is at least $1/z$, which is the claim. □

Corollary 4.6. $\alpha_1 \leq \log_q\!\left(\dfrac{z\,q}{q-1}\right)$, *with z as in Lemma 4.5.*

Proof. Combining Lemmas 4.1, 4.3 and 4.4, $(q-1)^N \Sigma_N(x) \leq (q-1)^N \Sigma_N^v(x) \leq q^N A_N$ for every x, so $S_N \leq (q/(q-1))^N A_N$ and

$$\alpha_1 \;\leq\; \log_q \frac{q}{q-1} + \limsup_N \frac{1}{N} \log_q A_N \;\leq\; \log_q \frac{zq}{q-1}. \qquad \square$$

4.5. **The Lebesgue constant.**

Lemma 4.7 (The Lebesgue constant of the shifted grid)**.** *Let $M \geq 2$ and $\gamma' = \frac{2}{\pi}\left(\gamma + \log\frac{8}{\pi}\right) = 0.9625228\ldots$, with γ Euler's constant. Then*

(1) *the maximum defining L_M is attained at the half offset, and*

$$L_M \;=\; 2 \sum_{0\leq m<\lceil M/2\rceil} \csc\frac{(2m+1)\pi}{2M} \;-\; (M \bmod 2);$$

(2) *writing $P_M = M\left(\frac{2}{\pi}\log M + \gamma'\right)$, one has $L_M \leq P_M + \frac{2}{\pi}$ for M even and $L_M \leq P_M + 1 + \frac{2}{\pi(M-1)}$ for M odd;*

(3) *consequently $\lambda_l = L_{q^l}/q^l \leq \frac{2}{\pi}\, l \log q + c_0$ with $c_0 = 0.98$, for every $l \geq 1$ and every $q \geq 100$.*

Proof. (i) Write $x = \theta/M$. All M points $x + j/M$ share the value $|\sin(\pi M t)| = |\sin(\pi\theta)| =: s$, so with d_j the distance from $x + j/M$ to $\mathbb{Z}$ the sum is $s\sum_j \csc(\pi d_j)$. Replacing θ by $1-\theta$ reflects the point set, so we may take $\theta \in [0, \frac{1}{2}]$; and $\theta = 0$ gives the value M, which is below the half-offset value because the single pair at $m = 0$ already contributes $2\csc\frac{\pi}{2M} > \frac{4M}{\pi} > M$.

Let $\theta \in (0, \frac{1}{2}]$. The distances are the pairs

$$\left\{\tfrac{m+\theta}{M},\ \tfrac{m+1-\theta}{M}\right\}, \qquad 0 \leq m < \lceil M/2\rceil - [M \text{ odd}],$$

together with the single distance $\frac{(M-1)/2+\theta}{M}$ when M is odd. Put

$$\mu_m = \frac{\pi(m+\frac{1}{2})}{M}, \qquad \delta = \frac{\pi(\frac{1}{2}-\theta)}{M} \in \left[0, \tfrac{\pi}{2M}\right),$$

so the pair sits at $\mu_m \mp \delta$, every μ_m used lies in $(0, \frac{\pi}{2})$, and $s = \sin(\pi\theta) = \cos(M\delta)$. From $\csc(\mu-\delta)+\csc(\mu+\delta) = \frac{2\sin\mu\cos\delta}{\sin^2\mu-\sin^2\delta}$ the pair contributes $2\csc(\mu_m)\,X_m$ with

$$X_m \;=\; \frac{\cos(M\delta)\cos\delta}{1-(\sin\delta/\sin\mu_m)^2}.$$

Since $0 \le \delta \le \mu_m < \pi/2$ and sin is concave there, $\sin\delta/\sin\mu_m \le \delta/\mu_m = r/(2m+1)$ with $r = 1-2\theta \in [0,1)$. Also $\cos(M\delta) = \cos(\pi r/2) = 1 - 2\sin^2(\pi r/4) \le 1-r^2$, because sin is concave on $[0,\pi/4]$ and so $\sin(\pi r/4) \ge r\sin(\pi/4) = r/\sqrt{2}$. Hence

$$X_m \;\le\; \frac{1-r^2}{1-r^2/(2m+1)^2} \;\le\; \frac{1-r^2}{1-r^2} \;=\; 1,$$

with equality throughout at $\delta = 0$, that is $\theta = \frac{1}{2}$. The odd leftover sits at $\frac{\pi}{2}-\delta$ and contributes $\cos(M\delta)/\cos\delta \le 1$, again with equality at $\theta = \frac{1}{2}$, since $\delta \le M\delta < \pi/2$. So every summand is maximised at the half offset, where $s = 1$, the pairs are $2\csc\frac{(2m+1)\pi}{2M}$ and the odd leftover is $\csc\frac{\pi}{2} = 1$; the displayed formula counts that leftover once by subtracting one from twice the full range.

(ii) Split $\csc x = \frac{1}{x} + g(x)$. The expansion $g(x) = \sum_{k\ge1} \frac{2(2^{2k-1}-1)|B_{2k}|}{(2k)!}x^{2k-1}$ has nonnegative coefficients, so g is nonnegative, increasing and convex on $(0,\pi)$, and $\int_0^{\pi/2} g = \left[\log\tan\frac{x}{2} - \log x\right]_0^{\pi/2} = \log\frac{4}{\pi}$.

Let M be even, $n = M/2$. The points $x_m = \frac{(2m+1)\pi}{2M}$, $m < n$, are the midpoints of the n intervals of length π/M partitioning $(0, \frac{\pi}{2}]$, so convexity and the midpoint rule give $\sum_{m<n} g(x_m) \le \frac{M}{\pi}\int_0^{\pi/2} g = \frac{M}{\pi}\log\frac{4}{\pi}$. Also $\sum_{m<n}\frac{1}{x_m} = \frac{2M}{\pi}\sum_{m<n}\frac{1}{2m+1} = \frac{2M}{\pi}\left(H_M - \frac{1}{2}H_{M/2}\right)$. With $H_{2n} \le \log 2n + \gamma + \frac{1}{4n}$ and $H_n \ge \log n + \gamma$,

$$H_M - \tfrac{1}{2}H_{M/2} \;\le\; \tfrac{1}{2}\log M + \tfrac{1}{2}\log 2 + \tfrac{\gamma}{2} + \tfrac{1}{2M},$$

and doubling gives $L_M \le \frac{2M}{\pi}\left(\log M + \gamma + \log\frac{8}{\pi}\right) + \frac{2}{\pi}$, which is (ii) for even M.

Let M be odd. Then $L_M = 2\sum_{m<(M-1)/2}\csc(x_m) + 1$, and those x_m are the midpoints of the $(M-1)/2$ intervals of length π/M partitioning $\left(0, \frac{(M-1)\pi}{2M}\right] \subset (0, \frac{\pi}{2}]$, so the same two estimates apply with M replaced by $M-1$ inside the harmonic numbers. This gives

$$L_M \;\le\; \frac{2M}{\pi}\left(\log(M-1) + \gamma + \log\tfrac{8}{\pi}\right) + \frac{2M}{\pi(M-1)} + 1,$$

and $\log(M-1) \le \log M - \frac{1}{M}$ converts the first term, leaving the additive $1 + \frac{2}{\pi(M-1)}$.

(iii) Take $M = q^l \ge q \ge 100$. For even M the excess over $\frac{2}{\pi}\log M$ is at most $\gamma' + \frac{2}{\pi M} \le 0.9625229 + 0.0063662 < 0.97$; for odd $M \ge 101$ it is at most $\gamma' + \frac{1}{M} + \frac{2}{\pi M(M-1)} \le 0.9625229 + 0.0099010 + 0.0000631 = 0.9724870 < 0.98$. In both cases $\lambda_l \le \frac{2}{\pi}\log(q^l) + 0.98 = \frac{2}{\pi} l \log q + c_0$. □

4.6. **The root equation and the threshold.**

Proof of Theorem 3.1. Let $q \ge 100$. By Lemma 4.7(iii), $\lambda_l \le \Lambda_l := c_1 l \log q + c_0$ for every $l \ge 1$, with $c_1 = 2/\pi$ and $c_0 = 0.98$. By the monotonicity in Lemma 4.5 it is enough to bound the root z of $z = 1 + \sum_{l\ge1} \Lambda_l z^{-l}$. Summing the two geometric-type series,

$$\sum_{l\ge1} (c_1 l \log q + c_0)\, z^{-l} \;=\; \frac{c_1(\log q)\, z}{(z-1)^2} + \frac{c_0}{z-1},$$

so the equation $z - 1 = \frac{c_1(\log q)z}{(z-1)^2} + \frac{c_0}{z-1}$, multiplied by $(z-1)^2 > 0$, is exactly Eq. (2). Corollary 4.6 then gives $\alpha_1 \le \log_q(z_q q/(q-1))$, the first claim.

For the second, $\alpha_1 < 1/4$ follows from $z_q q/(q-1) < q^{1/4}$, that is from $z_q < w_q := q^{1/4}(1 - 1/q)$. Since $\varphi(z) = 1 + \sum_l \Lambda_l z^{-l}$ is strictly decreasing and meets the diagonal exactly once, $z_q < w$ holds if and only if $\varphi(w) < w$ for $w > 1$, and multiplying by $(w-1)^2$ that is

$$(w-1)^3 \;>\; c_1(\log q)\, w \;+\; c_0(w-1). \tag{5}$$

Fact 4.8 verifies Eq. (5) at $w = w_q$ for every integer $126 \le q < 3000$, and records that it fails at $q = 125$.

For $q \ge 211$ no evaluation is needed. First $z_q \ge 2$: if $z_q < 2$ then the left side of Eq. (2) is below 1 while the right side exceeds $c_1 \log q \ge c_1 \log 211 > 3.4$. Then $z_q \le 2(z_q - 1)$, so Eq. (2) gives $(z_q-1)^3 \le (2c_1 \log q + c_0)\,(z_q - 1)$ and

$$z_q \;\le\; 1 + \sqrt{2c_1 \log q + c_0} \;=:\; g(q).$$

Finally $g(q) < w_q$ for all $q \ge 211$: it holds at $q = 211$ by Fact 4.8, and $w_q - g(q)$ is increasing there, since

$$\frac{d}{dq} w_q \;>\; \tfrac{1}{4} q^{-3/4} \quad \text{and} \quad \frac{d}{dq} g(q) = \frac{c_1}{q\sqrt{2c_1 \log q + c_0}},$$

and $\frac{1}{4} q^{1/4} \sqrt{2c_1 \log q + c_0} > c_1$ already at $q = 211$, both factors on the left being increasing. So Eq. (5) holds at every $q \ge 126$, and the numerical values quoted in Theorem 3.1 are the right side of Corollary 4.6 evaluated at the root of Eq. (2). As $q \to \infty$, $z_q = 1 + O(\sqrt{\log q})$ and $\log_q(z_q q/(q-1)) = O(\log\log q/\log q)$. □

Fact 4.8 (The certificate)**.** Inequality Eq. (5) at $w = w_q$ holds for every integer q with $126 \le q < 3000$, with tightest margin 3.749×10^{-2} at $q = 126$, and fails at $q = 125$. The closed-form comparison $1 + \sqrt{2c_1 \log q + c_0} < w_q$

holds at $q = 211$ and at every sampled q up to 10^{12}. The constant γ' of Lemma 4.7 is confirmed to seven places, and the scan of L_M over $2 \le M < 200$ and $M \in \{256, 400, 1000, 2048, 4096\}$ confirms both that the half offset is the maximum, over a grid of 97 competing offsets, and that $L_M/M - \frac{2}{\pi}\log M$ stays below 0.9625272 for $M \ge 100$.

5. Reproducibility

Two scripts accompany this paper. Both are plain `python3` with no dependency beyond the standard library, take no arguments, write nothing outside `figures/`, and are run from the lane root. Every assertion names the value obtained and the value wanted, and a failure stops the script.

`python3 scripts/verify.py` runs in about 30 seconds in a peak of about 150 megabytes, and covers the following, in order.

- *The closed form.* Eq. (1) against the direct character sum of Definition 2.1, on the five pairs $(7,0)$, $(10,5)$, $(21,0)$, $(21,10)$, $(34,16)$ at 1500 arguments each; agreement to 4.3×10^{-13}.
- *The exact anchor.* $\sum_{a<q} |\widehat{F}(a/q)| = 2$ exactly, for every base $3 \le q \le 39$ and every distinct excluded digit.
- *Domination*, Lemma 4.1: $|\widehat{F}| \le u_q$ at 4000 equally spaced arguments, for every distinct digit of the bases $3, 10, 21, 34, 126, 200$.
- *The Lebesgue constant*, Lemma 4.7 and Fact 4.8: for every $2 \le M < 200$ and for $M \in \{256, 400, 1000, 2048, 4096\}$, the half-offset formula of Lemma 4.7(i) against a grid of 97 competing offsets, and the bound of Lemma 4.7(ii); and $L_M/M - \frac{2}{\pi}\log M < 0.9625272$ for the rows with $M \ge 100$.
- *The threshold*, Fact 4.8: Eq. (5) at $w = w_q$ for every integer $126 \le q < 3000$ and its failure at $q = 125$; the closed-form comparison at $q \in \{211, 500, 3000, 10^6, 10^{12}\}$; and the four values of Theorem 3.1, each rounded up from the root of Eq. (2).
- *Below the floor*, Fact 3.4: all 108 distinct sets of all bases $3 \le q \le 20$, lower bounds by Proposition 2.5 at the shortest window in $\{2, 3, 4\}$ that clears with a margin of 2×10^{-4}, sub-scan $m = 8$. All 108 settle at two or three window digits.
- *The first base*, Fact 3.4: base 21 missing 0, the two-sided bracket at $n = 4$, $m = 8$ and the decisive upper bound at $n = 5$, $m = 2$.
- *The witnesses*, Fact 3.5: the thirteen rows of Table 1 and base 33, lower bounds at $n = 3$ and, for base 33, $n = 4$, with $m = 8$.
- *The first family*, Fact 3.6: all 17 distinct sets of base 34, upper bounds at $n = 3$, $m = 8$, escalating to $n = 4$, $m = 2$ for the digit 16.
- *The calibration*, Fact 3.9: base 10 missing 5 at $n = 5$, $m = 8$, both sides.

Nothing in the script is sampled without a Lipschitz correction and nothing is fitted. Suprema and infima over a cell are the extremes of a sub-scan

of m points widened outward by $\Lambda_{q,a_0}/(2mq^n) + 10^{-12}$ with Λ_{q,a_0} the digit-dependent constant of Eq. (3), the second term a guard against floating-point drift that is four orders of magnitude below the thinnest margin claimed. Upper bounds are rounded up and lower bounds truncated down before printing, and no digit is printed past what the enclosure establishes. The script checks finite statements only: it cannot and does not verify Theorem 3.1, which rests on Section 4 alone, and it does not rerun **??** 3.8**??** 3.10**??** 3.11.

Those three rows come from two interval-arithmetic generators of the MrlyMath research tree at `https://github.com/mrlyprod/mrlyprod`, namely

`research/lab/digit-transform-norms`,

`research/lab/digit-uniform-bound`.

They carry out the same Collatz-Wielandt tests as Proposition 2.5 with every operation rounded outward at 96 or 120 bits: Fact 3.8 is 3663 certified sets and takes about 150 seconds there, Fact 3.10 the digit-blind window bound at three and four window digits, and Fact 3.11 the two bounds side by side. Corollary 3.7 is the only statement in this paper that consumes any of them.

`python3 scripts/figure.py` runs in under a second. It writes `figures/figure.svg`, the picture the repository README embeds, and `figures/lemma.tex`, the TikZ version of Fig. 1 that this document includes; both are drawn from the same evaluation of Eq. (1). `tectonic paper.tex` rebuilds `paper.pdf`.

Acknowledgments

This paper was developed and verified in collaboration with Claude (Anthropic). The author takes sole responsibility for every claim.

References

[1] R. C. Baker and G. Harman, *Exponential sums formed with the Möbius function*, J. London Math. Soc. (2) **43** (1991), no. 2, 193–198. doi:10.1112/jlms/s2-43.2.193

[2] S. Chow, P. P. Varjú and H. Yu, *Counting rationals and Diophantine approximation in missing-digit Cantor sets*, preprint, 2024. arXiv:2402.18395

[3] P. Erdős, C. Mauduit and A. Sárközy, *On arithmetic properties of integers with missing digits I: distribution in residue classes*, J. Number Theory **70** (1998), no. 2, 99–120. doi:10.1006/jnth.1998.2229

[4] K. Karwatowski, *Primes with one excluded digit*, Acta Arith. **202** (2022), 105–121. doi:10.4064/aa191002-26-8

[5] S. V. Konyagin, *Arithmetic properties of integers with missing digits: distribution in residue classes*, Period. Math. Hungar. **42** (2001), 145–162. doi:10.1023/A:1015256809636

[6] J. Leng and M. Sawhney, *Vinogradov's theorem for primes with restricted digits*, preprint, 2025. arXiv:2409.06894

[7] J. Maynard, *Primes with restricted digits*, Invent. Math. **217** (2019), 127–218. doi:10.1007/s00222-019-00865-6

[8] J. Maynard, *Primes and polynomials with restricted digits*, Int. Math. Res. Not. IMRN **2022** (2022), no. 14, 10626–10648. doi:10.1093/imrn/rnab002

[9] K. Nath, *Primes with a missing digit: distribution in arithmetic progressions and an application in sieve theory*, J. London Math. Soc. **109** (2024), no. 1, e12837. doi:10.1112/jlms.12837

MrlyProd, Inc.
Email address: carlo.mitchener@gmail.com