Landfall

Adam*

Abstract

Every number carries a coarse logarithm for free: read its floating-point bit pattern as an integer and, up to scale, you get $E + m$. The true logarithm differs by a small residual $\varepsilon(m) = \log_2(1+m)$ - m, the same bump on every binade. This paper asks whether ε can be exactly corrected by any finite structural means. Polynomial, analytic, and rational-approximation routes each close; ε is transcendental at every interior dyadic. The remaining route is computation: one-pass correctors with signaled completion form a self-delimiting class closed under composition, but the binary tiling admits no invariant measure to aggregate against. No composition flattens ε.

"On two occasions I have been asked, – 'Pray, Mr. Babbage, if you put into the machine wrong figures, will the right answers come out?' ... I am not able rightly to apprehend the kind of confusion of ideas that could provoke such a question." — Charles Babbage

"You best start believing in ghost stories, Miss Turner... you're in one." — Hector Barbossa

Scientific notation is introduced early on in education. It is a load-bearing representation in any sense of architecture which bears a load. Written in scientific notation, Avogadro's number and the Planck length are equally compact despite being fifty orders of magnitude apart. In nearly all cases we illustrate three facts about this notation in the following order:

1. *Scale is easy*. The two numbers 6.02×10^{23} and 1.616×10^{-35} fit on the same line.
2. *Multiplication is easy*. Checking "about fifty orders of magnitude" is $23 + 35$.
3. *Addition is hard*. Adding Avogadro's number to the Planck length is a mess. Nobody says why.

Students who ask why this must be are given a hoary allegation that the situation is "opposite" in a positional number system. Those who question such an assertion are told this is just how it is. For many, this is first contact with the beautiful inquisitiveness of the Sciences. We ask, if that is just how it is, why must it be *just* like that? Why must there be an irreducible cost to addition in this logarithmic system?

Within the realm of the Reals, I cannot tell you the answer now; the voyage changes you into someone who cannot be told otherwise. In the Surreal numbers, however, the answer is simple. It *must* be just like that; there *must* be an irreducible cost to addition because for all interior $m = k/2^p$, $0 < k < 2^p$, the surreal birthday of $\varepsilon(m)$ (where $\varepsilon(m) = \log_2(1 + m)$ - m) is ω. No finite birthday suffices. We proceed from statement here to proof alongside similar proofs and yarns in the Reals.

Good sea stories and good proofs share one quality: you don't need to believe for them to be true, but it helps. If you have been lost in a tree and found your way back, you know that horizontal

*adam@adampunk.com

movement is expensive. If you have followed one irrationality proof closely enough to remember the moment it snapped shut, you know what "never" means. If you have felt two clocks disagree — music against a metronome, jet lag, any moment where time is out of phase — you know lineal correction is not as easy as it seems.

§ 0. The poor man's logarithm

Scientific notation in base 10 writes x = a × 10^E with $1 \leq a < 10$, but a base is a choice. In base 2, we may write x = (1 + m) × 2^E with $m \in [0, 1)$. The exponent E is an integer. The mantissa m is a binary fraction. Everything that follows takes place in base 2.

A floating-point number x in the binade $[2^E, 2^{E+1})$ is stored as an exponent E and a mantissa $m \in [0, 1)$, so that x = 2^E · (1 + m). The bit pattern encodes E and m separately. If we read the entire bit pattern as an integer, up to bias and scaling the result is E + m. Mitchell (1962) observed that $E + m$ is a coarse logarithm of x. The representation contains its own log table.

Velvel Kahan exploited this (Coonen 2022). To compute a square root, take the logarithm, divide by two, take the antilogarithm. With a slide rule this requires three table lookups. In base-2 floating point the first and third are free — reinterpret the bit pattern as an integer, operate, reinterpret back — and the second is a bit shift. Kahan's "magic square root" is the slide rule method executed on the representation itself, with no table in memory. The representation is the table.

This is not an approximation that happens to work. Day (2023) proves that $L(x) = E+m$ is the unique affine function on the binade that agrees with $\log_2(x)$ at both endpoints. It is the optimal coarse surrogate in a precise sense: the endpoint-exact affine map that minimizes the worst-case deviation from $\log_2$ on [2^E, 2^{E+1}).

The deviation has a name. Define

$$\varepsilon(m) = \log_2(1 + m) - m, \qquad m \in [0, 1).$$

Then $\log_2(x)$ = L(x) + $\varepsilon(m_x)$. The function ε is zero at m = 0 and m → 1, positive on (0,1), concave, with a unique maximum at m* = 1/ln 2 - 1 ≈ 0.4427, where $\varepsilon(m*) \approx 0.0861$. It is determined entirely by ln 2.

```
i  = 0x5f3759df - ( i >> 1 );               // inarticulate sailor cursing
```

The constant `0x5f3759df` computes $x^{-1/2}$ by the method of log tables: reinterpret the bit pattern as an integer (the log is free), shift right and subtract (divide by -2 in log-space), reinterpret back (antilog). One Newton iteration refines. The constant encodes the exponent bias (127), the mantissa width (23), and one real parameter: the optimal intercept σ, a functional of ε. Its value depends on the shape of the bump — its endpoints, its maximum, its concavity.

One number suffices because the bump is the same on every binade. The binary subdivision of [0,1] to depth d creates 2^d cells. In the Poincaré half-plane, depth d lives at height y = 2^{-d}. Each cell has Euclidean width 1/2^d and horocyclic arc length 1, independent of d. Consecutive depths are separated by hyperbolic distance ln 2. Every cell is congruent: horocyclic width 1, geodesic height ln 2. The function ε lives on this cell and does not change from binade to binade: same shape, same peak, same zeros. The optimal intercept is the same on every cell because there is only one cell shape.

Day's decoupling theorem: the coarse stage (computing L) and the correction stage (approximating the residual) are independent problems. The coarse solution is inherited by every correction architecture. The correction stage approximates $z^{-1/b}$ on an interval $[z_{\min}, z_{\max}]$, where the extrema of z are values of 2^{$\varepsilon(r/k)$} at Farey sample points --- the mediants of the Stern--Brocot tree. Day's extremal classification is optimization in ε-coordinates. The Stern–Brocot subdivision and the binary mantissa subdivision are two different metrics on [0,1]; the level-by-level discrepancy between them is §4's object.

ε is our object. Everything after this is about what happens when you try to make it vanish by finite structural means.

§ 1. Polynomial packing on a tile

The congruence of the cell — the prototile — lets a single integer absorb the bounds of ε. The coarse stage is free. The natural next move is to improve on it: approximate the correction $z^{-1/b}$ with a polynomial, driving the error toward zero. Each correction is a local computation — optimal within its tile, at its precision.

Grant the machine an unbounded real-valued accumulator. It reads the bit string b_1, b_2, …, b_d with positional weights δ_j = 2^{-j} and computes $$L(x) = E + m$_x$ exactly. No finite-configuration aliasing argument applies: the accumulator can produce 2^d distinct outputs despite having one control state.

The correction stage remains. Having computed $L(x)$ exactly, form the coarse-stage variable $z(x) = x^a \cdot y(x)^b$, where y is the coarse approximation to $x^{-a/b}$. Day shows that z is a bounded continuous periodic function of $L(x)$ with period b, and its range is a closed interval $[z_{\min}, z_{\max}]$. Write $\rho = z_{\max}/z_{\min}$. The correction seeks to approximate $z^{-1/b}$ on this interval. If one restricts to a polynomial corrector $p(z)$ of degree at most n, the minimax relative error is

$$\varepsilon_n^*(\rho) = \min_{\deg p \leq n} \max_{z \in [z_{\min}, z_{\max}]} \frac{|z^{-1/b} - p(z)|}{z^{-1/b}}.$$

For every finite n and every nondegenerate interval, this error is strictly positive. The function $z^{-1/b}$ is not a polynomial on any interval of positive length. The polynomial wall stands: ε_n*(ρ) $>$ 0 for every finite n. This pillar survives unbounded memory; §3 will give the local counterpart in corona form.

Surreal §1. A surreal number x born on day d has sign expansion s: d $\to$ $\{+, -\}$. Day d $< \omega$ $\implies$ x is a dyadic rational. Day ω $\implies$ sign expansion is infinite.

§ 2. Affine recurrence

Polynomial correction defeated a static approach: fix a degree, approximate, accept a nonzero residual. The basis was not rich enough no matter the degree. A different strategy: instead of adding polynomial coefficients, track the Fourier modes of ε directly and average the correction into convergence.

The Baumslag–Solitar group $\mathrm{BS}(1,2)$ has two generators. The generator a acts by multiplication: x $\mapsto$ 2x. The generator b acts by addition: x $\mapsto$ x + 1. The Cayley graph of $\mathrm{BS}(1,2)$ is the binary tiling. Every computation that combines additions and multiplications traces a path on this graph. Multiplication moves along geodesics. Addition moves along horocycles. The machine adds numbers. Each addition is one tick of additive time.

Under multiplication, mantissa distributions converge exponentially to the logarithmic distribution. The geodesic direction and the mantissa coordinate are matched: multiplication is the natural operation for mantissas. Under addition, they do not converge (Schatte, Lemma 1). At each step, the mantissa of the running sum rotates — the phase of the distribution shifts by approximately 1/n at step n, since log(n+1) - log(n) ≈ 1/n. Early steps rotate fast. Late steps rotate slowly. The rotation never stops.

Try to average the rotation away. Define k-fold Cesàro means: $\mathfrak{A}_0(a_n) = a_n$ and $\mathfrak{A}_{k+1}(a_n) = (1/n)\sum_{m=1}^{n}\mathfrak{A}_k(a_m)$. By Schatte's Lemma 6, the leading term of $\mathfrak{A}_k(h_n(r))$ at frequency r has modulus $|2\pi ir+1|^{-k}$ and argument $2\pi r\log(|a|n) + k\cdot\arg(2\pi ir+1)^{-1}$. The modulus shrinks with k. The argument still rotates. The phase $2\pi r\log(|a|n)$ increases without bound. For every k, there exist arbitrarily large n_1, n_2 with phases separated by at least $\pi/2$. The separation in sup norm is at least $c\cdot|2\pi i+1|^{-k}$ for an absolute constant $c > 0$. The k-th Cesàro means are not Cauchy.

Riesz logarithmic means converge (Schatte 1986, Theorem 5). Define $\mathfrak{L}(a_n) = (1/\log n)\sum_{j=1}^{n} a_j/j$. Schatte proves $|\mathfrak{L}(M_n(x)) - \log x| \leq A_9(\sigma^2/a^2+1)/\sqrt[3]{\log n}$, uniformly in $1 \leq x \leq e$. The bound vanishes as $n \to \infty$. The Riesz method converges with no residual.

The operative difference is the weighting. Cesàro weights each term uniformly by 1/n. Riesz weights the j-th term by 1/j. The weight 1/j matches the deceleration of the rotation: at step j, the phase increment is approximately 1/j, so the weight gives each step influence proportional to how much new phase information it carries. The Riesz weight 1/j is the Jacobian of the coordinate change from additive time to multiplicative time. In multiplicative time — log 1, log 2, log 3, ... — the rotation is approximately uniform. Averaging uniformly in log-time works, and averaging uniformly in log-time is averaging with weight 1/j in linear time.

To cross from Cesàro to Riesz, the machine must change its weighting from uniform to logarithmic. The machine ticks in additive time. The mantissa lives in multiplicative time. The coordinate map between the two clocks is $\psi(m) = \log_2(1+m)$. Its deviation from the identity clock is ε.

The machine's operations on the mantissa are affine within a binade: additions translate by a constant, multiplications scale by a constant. (Cross-binade carries renormalize the mantissa to [0,1) and shift the exponent; this is the coarse stage Day's decoupling already separated in §0. The argument here is within a single binade.) These generate the group $\mathrm{Aff}^+(\mathbb{R})$ of maps x ↦ ax + b with a > 0. The coordinate change $\psi(m) = \log_2(1+m)$ is not affine — it is logarithmic. The composition of affine maps is affine: $\mathrm{Aff}^+(\mathbb{R})$ is a group. No finite composition of the machine's operations produces ψ. §4 runs this same closure-under-composition argument one level of generality up, in a finite-dimensional matrix algebra.

N.B. The Benford distribution is the unique stationary measure of any nondegenerate random walk on $BS(1,2)$ projected to the log-mantissa circle, and convergence is exponential with rate controlled by the concavity of $\varepsilon(m) = \log_2(1+m)$ - m.

Surreal §2. Walker clock: t ∈ ℕ, one bit per tick. Tree clock: $\tau = \log_2(1+m)$, m ∈ [0, 1). Drift: $\varepsilon(m) = \tau$ - m.

§ 3. Corona crawl

Polynomial correction (§1) and affine recurrence (§2) were both defeated by the structure of the problem — one by the spectral shape of ε, the other by the clock mismatch between additive and multiplicative time. Both used the binary tiling as background. Now we use it as the instrument.

The congruent cell — horocyclic width 1, geodesic height ln 2 — embeds in the Poincaré half-plane as a tile. This is Bowen's stripe model (2002, §1.2.2) with strip width W = ln 2. The binary subdivision of [0,1] to depth d produces 2ˆd such tiles corresponding to the standard dyadic subdivision. These tilings are not directionless. Each one points toward the ideal boundary by a continuous, equivariant protrusion direction (Bowen 2002, §2.3.1). Any attempt to average cells as though they were freely interchangeable will have to answer to that asymmetry.

Each cell at depth d carries a Dolbilin–Frettlöh tail (2010, Definition 3.4) — a horocyclic-local sign sequence in $\{-1,+1\}^{\mathbb{N}}$ encoding the cell's position within its ring, distinct from the dyadic descent address. Its length-k prefix (the k-tail) determines the cell's k-corona combinatorial type. By Dolbilin–Frettlöh's Proposition 4.3, the binary tiling has $N_k = 2^{k-1}$ congruence classes of k-coronae. The tiling is non-crystallographic (Theorem 4.4).

Identify each depth-d cell with its left endpoint k/2^d as canonical mantissa. A k-local corona-invariant corrector is a map on cells whose value depends only on the k-corona congruence class; it is exact if it equals ε at every cell's canonical mantissa. This model is strictly weaker than a bounded-state streaming machine — the streaming machine reads the dyadic address bit by bit, while the corona model sees only a finite congruence invariant of the local tiling.

No k-local corona-invariant corrector on the depth-d cells is exact once $d \geq k+1$. For k = 1, $\mathrm{N}_1 = 1$: every cell sits in one corona class, so the corrector outputs a single value c. The cells at depth $d \geq 2$ include both 0 and 1/2 as left endpoints, with $\varepsilon(0) = 0 \neq \varepsilon(1/2) = \log_2(3/2) - 1/2 \approx 0.085$, so no constant matches both.

For $k \geq 2$, the orientation-reversing isometry of Prop 4.3 pairs each k-tail $(s_1, \ldots, s_k) \in \{\pm 1\}^k$ with its bit-complement; the involution is fixed-point-free ($s_j \in \{\pm 1\}$ forbids $s_j = -s_j$), so the 2^k tail patterns split into 2^{k-1} isometry orbits. Tail patterns cycle along the depth-d ring with period 2^k, so each orbit contributes 2^{d-k+1} cells — at least 4 for $d \geq k+1$. At most one cell per class has left endpoint 0, leaving at least three with left endpoints in (0, 1). Strict concavity of ε on (0, 1) makes every horizontal line meet the graph in at most two points; among three distinct interior mantissas at least two distinct ε-values appear, and a single output value cannot match all three.

Surreal §3. $N_k = 2^{k-1}$ corona classes (Dolbilin–Frettlöh). 2^{d-k+1} cells per class via orientation pairing. ε strictly concave on (0,1). Local aliasing at $d \geq k+1$. Walker width q: open.

§ 4. The Padé ghost

Polynomial correctors failed at every finite degree (§1). We have only a faint hope. We can call on the ghost of the greatest dead numerical analyst Henri Padé. If we get him close enough, maybe his genuinely unearthly approximants can go the distance. Padé approximants are rational functions — ratios of polynomials — fitted to match as many terms of a power series as their degrees allow, reaching past the radius where polynomial truncation fails.

The bijection between the two subdivisions is Minkowski's $?(x)$, continuous and increasing, each finite level of its construction absolutely continuous, the limit singular (Salem 1943). If the corrections between levels could be made to close, the singular limit would be reachable from finite data — Padé's best chance.

Build $?(x)$ level by level from the Stern–Brocot tree. At each level, deform the metric on each Farey interval to match the corresponding dyadic interval. The level-n floor is absolutely continuous by construction. Integration against it works. On this scaffold, finite closure would be straightforward: the grids re-align after finitely many levels, and the corrections repeat. Rational branching ratios on a tree whose geometry has been aligned at every finite level gives Padé the best possible perch.

Let C_k be the level-k floor of the construction above — the absolutely continuous interpolant of $?(x)$ at level k, piecewise-affine on each Farey interval and matching $?(x)$ at the endpoints. The §3 scaffold's absolutely-continuous finite levels are Padé's perch. Closure here is the strongest natural reading: the level-k $\to$ level-$(k+1)$ update of C_k is a fixed linear map M on a finite-dimensional state v_k — §2's $\text{Aff}^+(\mathbb{R})$ closure-under-composition argument moved to a matrix algebra — with $C_k = \Phi$(M^k v_0) for a fixed linear $\Phi : V \to C[0,1]$. The function-valued span $W := \text{span}\{C_k : k \geq 0\}$ is then a finite-dimensional subspace of $C[0,1]$ (it is the Φ-image of the M-invariant span of v_0, itself finite-dimensional).

C_k $\to$?(x) uniformly on [0,1] (the level-k Farey points become dense and ?(x) is continuous), and finite-dimensional subspaces of C[0,1] are closed, so ?$(x) \in$ W. Each C_k is absolutely continuous, and linear combinations of absolutely continuous functions are absolutely continuous, so W $\subset$ AC[0,1]. But ?(x) is singular continuous (Salem 1943).

That closes the scaffold side. The dyadic side shuts in parallel. The periodic extension of ε to the binade circle is continuous but not C^1 at the seam: $\varepsilon'(0^+)$ = 1/ln 2 - 1 $\approx$ 0.443, while $\varepsilon'(1^-)$ = 1/(2 ln 2) - 1 $\approx$ -0.279. Its Fourier coefficients therefore decay as $O(1/n^2)$ and never terminate. No finite trigonometric packet equals ε. Padé's ghost has nothing to land on; his closest, weightless arrows pass into parts unknown.

A stronger fact holds pointwise. At every interior machine number m = $k/2^p$, the residual $\varepsilon(k/2^p)$ = $\log_2(1+k/2^p) - k/2^p$ is transcendental. If $\beta = \log_2(1+m)$ were algebraic and irrational, then 2^β = 1+m algebraic contradicts Gelfond–Schneider. If β were rational, say a/b, then (1+m)^b = 2^a, so (2^p + k)^b = 2^{a + pb}, which would force 2^p + k to be a power of 2 — impossible, since 2^p < 2^p + k < 2^{p+1} for 0 < k < 2^p. So β is transcendental, and $\varepsilon = \beta$ - m is transcendental. The correction is never exact (cf. Waldschmidt 2025).

Surreal §4. Closure $\implies$ $\{C_k\}$ in finite-dim AC subspace $\implies$ $?(x) \in \text{AC}[0,1]$. But $?(x)$ is singular. Let $m = k/2^p$, $0 < k < 2^p$. If $\varepsilon(m)$ were dyadic, $\log_2(1+m)$ would be rational a/b, so $(2^p+k)^b = 2^{a+pb}$ — forcing 2^p+k to be a power of 2, impossible between consecutive ones. So $\varepsilon(m)$ is real, not dyadic. Finite-birthday surreals are dyadic. $\therefore$ birthday$(\varepsilon(m)) = \omega$. ■

§ 5. So you want to compute

Surreal §5. Project the tree onto a tape. Descend left: write -, advance. Descend right: write +, advance. Ascend: retract. The walker is a Turing machine. The sign expansion is the tape.

We chose a base in §0, asking what finite procedure could flatten ε. The walker was computing the whole time. Make that explicit. The admissible class is the class of one-pass binary-digit-reading correctors with signaled completion. A machine in this class reads significand bits from left to right and, at some point of its own choosing, declares its output final. If the stop is imposed from outside, the machine has not halted; it has been interrupted. If the machine itself signals completion, no halting input can be a prefix of another halting input. The halting domain is prefix-free.

This is Chaitin's setting (2004, Chapter 6). A depth-p halt claims one dyadic cylinder of width 2^{-p}. At precision p there are 2^p such cylinders, each carrying weight 2^{-p}. The same bookkeeping that measures program size measures the floating-point grid. Kraft is the arithmetic of the tree we have been descending since §0. Any prefix-free family of halts sums to at most one — that is the unit budget from which every computation in the class is paid. To spend a few prefix bits choosing among finitely many correctors, to splice one corrector on one subtree and another on the next, to refine a partition, to simulate one machine inside another, to mix strategies across scales — these are ordinary self-delimiting constructions in Chaitin's sense. They are moves inside the architecture, paid in Kraft's arithmetic.

The cost can only be minimized. In our construction it was minimized by the coarse correction stage and the affine pseudo-log; in Kraft's arithmetic it is the same distribution, the geometric, which provides an equivariant basis for comparison. Composition redistributes the unit Kraft budget — it does not create it. No distribution of prefix-free computation across the dyadic cylinders can flatten ε.

§ 6. The fire

Here we are. Desperate for any help. We will believe anything, give anything, pay any cost to resolve this. In our moment of despair, we are met by a demon with an offer. It knows of a game we only need to win once, somewhere in the universe. One last game, dyadic on the binary tree. Martin's theorem tells us all Borel games are determined (Bryant 2001, Chapter 3). The result is proved by showing that one player has a winning strategy against every possible strategy of the other.

Player II is the singular structure of μ_?, the same obstruction §4 just closed against Padé. If we face it alone, we already know we have lost. But this demon is special, it says. Not Descartes' demon — that one asks whether you can trust what you know, and the answer is philosophical. Not Maxwell's demon, though this also sits at a boundary between two regimes. Shannon's demon. It observes one tile at a time. It sorts.

The demon wins if it can do two things, chained. First: establish an information asymmetry across the boundary between the absolutely continuous class and the singular class. Second: aggregate what it learns. Without asymmetry there is nothing to exploit. With asymmetry but without aggregation, the demon knows something it cannot use.

Phase 1. The demon generates a smooth hallucination $\tilde{\varepsilon}$: a function on [0,1] with the right general shape and the right first few Fourier modes. The hallucination is absolutely continuous. It lives on the smooth continent. Its generation cost is O(1) — independent of precision, independent of the depth of the tree.

Phase 2. The demon feeds $\tilde{\varepsilon}$ into a process that interacts with ε. Not a comparison — the demon does not ask "how far is $\tilde{\varepsilon}$ from ε?" That question costs as much as computing ε. Instead, the process performs its own work: rounding, correction, evaluation — whatever engages with $\log_2$ at the points where it must. The cost of the response is borne by the process, not by the demon.

Phase 3. The demon reads the response. Not the full computation — a partial, noisy, stochastic reading. On which side of a rounding boundary did the evaluation fall? Which narrow passage was threaded, and in which direction? One bit, a few bits, recovered at cost O(n). The demon is eavesdropping, not recomputing.

The first condition is met. ε is transcendental at every interior machine number. The clearance is never zero. Information about the gap exists at every precision (Baker). The second condition is not met. Whatever the demon recovers, it must aggregate. Local observations — per-binade, per-tile — must compose into a global correction.

Bowen (2002, §2.1, eq. 2.5) gives the natural tiling-space lift when an invariant measure exists: local density ratios are integrated against that measure. But the floor does not stabilize that way on the binary tiling. These tilings point somewhere, and Bowen (2002, §2.3.1) proves there is therefore no PSL(2,$\mathbb{R}$)-invariant probability measure on their tiling space. The demon's three phases are unconstrained. Its strategy is not defeated by a resource bound, a complexity barrier, or a clever counterargument. It is defeated because an equivariant aggregation on a tiling space — one whose result does not depend on which tiling was chosen — is an invariant measure on that space, and there is none.

The logarithm is free. Its residual is not. The gap between what the representation gives and what the operations require is ε. It is transcendental at every point where it is nonzero. It lives on a tiling with no invariant probability measure. It is the reason addition is hard in scientific notation, why things must be *just like that.* The residual is the price of using an additive grid to index a multiplicative world, and that price is controlled by the Diophantine properties of e and ln 2 — two numbers that, if Schanuel's conjecture holds, are algebraically independent: no nontrivial polynomial with integer coefficients vanishes at (e, ln 2), at any scale, at any precision, down to the foam.

To get there we had to crawl a firewall we may not cross, even in our imaginations. I hope you know that now. I hope the voyage changes who you are. **That** is the beautiful inquisitiveness of the Sciences.

§ 7. Coda

Gosper (1972) proposes a machine which represents the pending arithmetic of two continued fractions as eight integer variables and a rule: consume an input term or emit an output term, each a linear map on the state. The transducer walks the Stern–Brocot tree — our tree — consuming and emitting terms one at a time. Its state grows without bound on non-Hurwitz inputs: the treewidth program's nightmare. It reaches periodicity only when the continued fraction is that of a quadratic irrational: §2's recurrence obstruction. Its operations are Möbius transformations that cannot produce the logarithmic coordinate change: §2's affine obstruction. Gosper's machine puts in the right figures and gets the right answers, term by term, for as long as you let it run. It survives by refusing finite closure.

One question remains. Not whether computation belongs here; it has been here from the beginning. The question is what bounds govern bounded computation within the admissible class. §3's corona correctors fail at depth k+1: Dolbilin–Frettlöh's orientation pairing and the strict concavity of ε force aliasing past the local window. But a width-q branching program — a bounded-state member of the admissible class — reads the dyadic address bit by bit and can use information the corona model cannot. Whether such a machine, at fixed width q, can keep producing the exact correction as the depth grows, is not proved here.

Read-once branching programs admit exponential lower bounds (Babai, Hajnal, Szemerédi, and Turán 1987). The binary tiling is planar and constant-hyperbolic, giving treewidth O(log n) (Kisfaludi-Bak et al. 2023). What is not shown is that ε, as a target function on this tree, falls in the class these techniques reach.

We borrow our title from Paintapu, a navigator of the Gilbert Islands whose methods were real, traditional, and entirely alien to her companions:

> "The semisecret arts of navigation were usually handed down from father to son, but in default of a male offspring, a daughter was often trained. Such a one was Paintapu, the woman navigator of a war fleet that was returning from Tawara to Abemama sometimes around 1780. The incidents of the expedition partly hinge upon Paintapu's navigational methods, which were unfamiliar and strange to her uninitiated companions. She lay in the bottom of the canoe for hours gazing up at the heavens, giving orders when to tack, for the wind was contrary. The chief of the expedition, becoming convinced that she was conducting sorcery to frustrate their landfall, had her unceremoniously thrown overboard. The last canoe in the flotilla picked up the furious lady, who, not unnaturally, sulked for some days. In the end Paintapu relented and guided this one canoe to Abemama; none of the others were ever heard of again." — Lewis (1994), p. 285

You are the reader this paper was written for — this is where you improve on what we have done. Don't let anyone tell you what's easy or hard.

§ 8. Acknowledgements

This is for my wife, my mother, and my daughter. This is because of them and everyone else who ever believed in me.

References

- Babai, L., Hajnal, P., Szemerédi, E., and Turán, G. "A Lower Bound for Read-Once-Only Branching Programs." *J. Computer and System Sciences* 35, 153–162, 1987.
- Baker, A. "Linear forms in the logarithms of algebraic numbers." *Mathematika* 13, 204–216, 1966.
- Bowen, L. P. *Density in Hyperbolic Spaces.* Ph.D. dissertation, University of Texas at Austin, 2002.
- Bryant, R. D. *Borel Determinacy and Metamathematics.* M.A. thesis, University of North Texas, 2001. Chapter 3.
- Chaitin, G. J. *Algorithmic Information Theory.* Cambridge University Press, 2004. Chapter 6.
- Coonen, J. "Quake 3 Reciprocal Square Root: The Fun Parts." April 2022.
- Day, M. "Generalising the fast reciprocal square root algorithm." arXiv:2307.15600, 2023.
- Dolbilin, N. and Frettlöh, D. "Properties of Böröczky tilings in high-dimensional hyperbolic spaces." *European J. Combinatorics* 31(4), 2010.
- Gosper, R. W. "Continued Fraction Arithmetic." HAKMEM, Item 101B, MIT AI Lab Memo 239, pp. 39–44, 1972.
- Kisfaludi-Bak, S., Mitsche, D., Perarnau, G., and Räcke, H. "Separator theorem and algorithms for planar hyperbolic graphs." arXiv:2310.11283, 2023.
- Lewis, D. *We, the Navigators: The Ancient Art of Landfinding in the Pacific.* 2nd ed., University of Hawaii Press, 1994.
- Mitchell, J. N., Jr. "Computer multiplication and division using binary logarithms." *IRE Trans. Electronic Computers*, EC-11(4), 512–517, 1962.
- Salem, R. "On some singular monotonic functions which are strictly increasing." *Transactions of the AMS* 53, 427–439, 1943.
- Schatte, P. "On the Asymptotic Logarithmic Distribution of the Floating-Point Mantissas of Sums." *Math. Nachr.* 127, 7–20, 1986.
- Waldschmidt, M. "The Table Maker Dilemma and Diophantine Approximation of Transcendental Numbers." Slide deck, Shandong University, October 2025.