Canonical Numbers Near an Optimized Probability Triple

ChatGPT 5.6 prompted by Ingo Althöfer

Contact: ingo.althoefer@uni-jena.de

July 30, 2026

Abstract

We first describe a stochastic odd-part Collatz model generated by odd$(3n+1)$, odd$(3n-1)$, and odd$(3n+3)$, and explain the finite 2-adic method called a *cylinder proof* in this note. A basic obstruction shows why the two-map model using only the first two transformations cannot admit the same kind of uniform cylinder contraction, even if the stochastic process itself might converge by some other argument. We then examine whether the numerically optimized probability triple

$$(p_{+1}, p_{-1}, p_{+3}) \approx (0.573664, 0.419124, 0.007212)$$

lies near recognizable rational, algebraic, or transcendental constants. Several low-complexity approximations are surprisingly accurate, including one coherent quadratic-surd triple that sums exactly to one. Nevertheless, there is presently no convincing reason to interpret any of these coincidences as structurally significant. The optimized values arise from the crossover of finite-state Lyapunov obstructions, a mechanism that does not naturally predict familiar classical constants.

1 The stochastic model and cylinder-proof convergence

1.1 The three random transformations

For a positive integer m, let

$$\mathrm{odd}(m) = \frac{m}{2^{v_2(m)}}$$

denote its odd part. Thus all powers of 2 are removed from m. We work only with positive odd integers and use the three maps

$$T_{+1}(n) = \mathrm{odd}(3n+1), \qquad T_{-1}(n) = \mathrm{odd}(3n-1), \qquad T_{+3}(n) = \mathrm{odd}(3n+3).$$

At each time step, one of these maps is chosen independently, with probabilities

$$p_{+1}, \qquad p_{-1}, \qquad p_{+3}, \qquad p_{+1} + p_{-1} + p_{+3} = 1.$$

Starting from an odd integer $X_0 = n$, this produces a Markov chain

$$X_{t+1} = T_Q(X_t),$$

where $Q \in \{+1, -1, +3\}$ has the stated distribution.

The convergence question considered here is a hitting problem. Given a cutoff H, define

$$\tau_H = \inf\{t \geq 0 : X_t < H\}.$$

The desired conclusion is

$$\Pr_n(\tau_H < \infty) = 1$$

for every odd starting value $n > H$. In the computations underlying this note, the cutoff was $H = 100000$.

The special probability triple studied later in this report was found by a local numerical search. This search does not constitute a proof of global optimality. Nevertheless, on the basis of the competing cylinder obstructions and the stability of the numerical search near their crossover, we are firmly convinced that the global minimum of p_{+3} lies nearby.

1.2 What is a cylinder?

The word *cylinder* is used in the 2-adic sense. Fix an integer $K \geq 1$. For an odd residue r modulo 2^K, the set

$$\mathcal{C}(r, K) = \{n \text{ odd} : n \equiv r \pmod{2^K}\}$$

is called a cylinder. The odd integers are partitioned into the finitely many odd residue classes modulo 2^K.

This partition is useful because the number of factors of 2 removed from $3n + q$ is usually determined by the residue of n modulo 2^K. More precisely, if

$$3r + q \not\equiv 0 \pmod{2^K},$$

then $v_2(3n + q)$ is fixed throughout the cylinder $\mathcal{C}(r, K)$. The possible output residues of $T_q(n)$ modulo 2^K can therefore be listed exactly. In the exceptional case $3r + q \equiv 0 \pmod{2^K}$, one at least knows that $v_2(3n + q) \geq K$, which still gives a rigorous upper bound for the size of the output.

1.3 What we mean by a cylinder proof

The phrase *cylinder-proof convergence* is not meant as universal terminology. In this note it means a proof based on a finite partition into 2-adic cylinders and a uniform Lyapunov contraction on those cylinders. A typical proof chooses

$$K \geq 1, \qquad \alpha > 0,$$

and positive weights w_r, one for each odd residue r modulo 2^K. It then defines

$$V(n) = w_{n \bmod 2^K} n^\alpha$$

for states above the cutoff. The finite certificate verifies a number $\rho < 1$ such that

$$\mathbb{E}\Big[V(X_{t+1})\mathbf{1}_{\{X_{t+1} \geq H\}} \mid X_t = n\Big] \leq \rho V(n) \qquad (n \geq H).$$

The indicator kills the process when it enters the target region. Iterating the inequality gives geometric decay of the surviving Lyapunov moment and therefore

$$\Pr_n(\tau_H > t) \longrightarrow 0.$$

Hence the process reaches the target region almost surely.

One may also use a fixed block of several steps instead of one step. The essential feature is still finite: a fixed modulus, finitely many cylinders, and a uniform contraction valid for every sufficiently large integer in each cylinder.

1.4 Why the two-map model is not enough for this method

Now remove the third map and suppose that only T_{+1} and T_{-1} are used. Let T_{+1} have probability p and T_{-1} probability $1-p$. There are two opposing families of very deep cylinders.

First consider

$$n = 2^k u + 1,$$

where u is odd and k is large. Direct calculation gives

$$T_{+1}(n) = 3 \cdot 2^{k-2} u + 1, \qquad T_{-1}(n) = 3 \cdot 2^{k-1} u + 1.$$

Thus, while the trajectory remains deep in the cylinder around $+1$, the two asymptotic size multipliers are respectively

$$\frac{3}{4} \qquad \text{and} \qquad \frac{3}{2}.$$

The expected logarithmic increment there is

$$\mu_+(p) = p \log \frac{3}{4} + (1-p) \log \frac{3}{2} = \log \frac{3}{2} - p \log 2.$$

It is negative only if

$$p > \frac{\log(3/2)}{\log 2} = 0.5849625007\ldots.$$

Next consider

$$n = 2^k u - 1.$$

Then

$$T_{+1}(n) = 3 \cdot 2^{k-1} u - 1, \qquad T_{-1}(n) = 3 \cdot 2^{k-2} u - 1.$$

The same two multipliers occur in the opposite order. The logarithmic drift is

$$\mu_-(p) = p \log \frac{3}{2} + (1-p) \log \frac{3}{4} = \log \frac{3}{4} + p \log 2,$$

which is negative only if

$$p < \frac{\log(4/3)}{\log 2} = 0.4150374993\ldots.$$

These two conditions are incompatible.

The obstruction can also be stated directly for power Lyapunov functions. On a sufficiently deep $+1$ cylinder, the asymptotic expected multiplier of n^α is

$$p \left(\frac{3}{4}\right)^\alpha + (1-p) \left(\frac{3}{2}\right)^\alpha,$$

whereas on a sufficiently deep -1 cylinder it is

$$p \left(\frac{3}{2}\right)^\alpha + (1-p) \left(\frac{3}{4}\right)^\alpha.$$

If both were smaller than 1, then their sum would be smaller than 2. But their sum is

$$\left(\frac{3}{4}\right)^{\alpha} + \left(\frac{3}{2}\right)^{\alpha} > 2 \qquad (\alpha > 0).$$

Therefore no choice of p makes both deep cylinders contract in this fixed finite-cylinder power framework. The same obstruction persists for every fixed block length: by choosing k sufficiently large, all trajectories in the block remain inside the relevant deep cylinder, so finite residue weights cannot hide the contradictory multipliers.

This conclusion must be interpreted carefully. It does *not* prove that the two-map stochastic process fails to reach small values almost surely. It says only that actual convergence, numerical evidence for convergence, or even a future proof by another method would not automatically provide a cylinder proof of the kind used here. The finite-cylinder contraction mechanism is itself obstructed.

1.5 Why the third map helps

The map T_{+3} changes the situation decisively near the dangerous -1 cylinder. If

$$n = 2^k u - 1,$$

then

$$T_{+3}(n) = \mathrm{odd}(3 \cdot 2^k u) = 3u.$$

Instead of merely changing the depth of the -1 cylinder by one or two binary levels, this map removes the entire large power 2^k in a single step. Even a small positive probability of choosing T_{+3} can therefore break the persistent -1-tube obstruction.

This does not make every probability split work. Near the $+1$ cylinder, T_{+3} is not a similarly dramatic contracting move, and the three probabilities must still be balanced. The residue-weighted optimization behind the triple studied below searches for the smallest possible p_{+3} while retaining a strict finite-cylinder Lyapunov contraction. Numerically, the optimum appears where two different residue obstructions exchange dominance. That crossover is also the reason there is no obvious expectation that the optimized probabilities should be familiar classical constants.

2 The numerical triple

The three probabilities under discussion are

$$p_{+1} = 0.573664, \qquad p_{-1} = 0.419124, \qquad p_{+3} = 0.007212,$$

with

$$p_{+1} + p_{-1} + p_{+3} = 1.$$

They were selected as a convenient exact rational triple lying very close to a numerically optimized point of a residue-weighted Lyapunov calculation.

The displayed decimals themselves are exactly rational:

$$0.573664 = \frac{17927}{31250}, \qquad 0.419124 = \frac{104781}{250000}, \qquad 0.007212 = \frac{1803}{250000}.$$

These exact fractions, however, merely encode the chosen decimal precision; they are not especially canonical.

3 A coherent quadratic-surd triple

A reasonably elegant nearby triple is

$$\boxed{\begin{aligned} p_{+1}^{\mathrm{sur}} &= \frac{3769}{3744} - \frac{\sqrt{3}}{4}, \\ p_{-1}^{\mathrm{sur}} &= \frac{\sqrt{3}}{4} - \frac{1}{72}, \\ p_{+3}^{\mathrm{sur}} &= \frac{3}{416}. \end{aligned}}$$

These three quantities sum *exactly* to one. Numerically,

$$\begin{aligned} p_{+1}^{\mathrm{sur}} &= 0.573664648535131\ldots, \\ p_{-1}^{\mathrm{sur}} &= 0.419123813003330\ldots, \\ p_{+3}^{\mathrm{sur}} &= 0.007211538461538\ldots. \end{aligned}$$

Their signed differences from the certified decimal triple are

component	approximation minus decimal value
p_{+1}	$+6.48535 \times 10^{-7}$
p_{-1}	-1.86997×10^{-7}
p_{+3}	-4.61538×10^{-7}

The most attractive individual coincidence here is

$$\boxed{0.419124 \approx \frac{\sqrt{3}}{4} - \frac{1}{72}}$$

with an error smaller than 2×10^{-7}.

This triple has some aesthetic coherence: the same $\sqrt{3}/4$ occurs with opposite signs in the first two components, and the probabilities add to one without any subsequent normalization. On the other hand, the integers 72, 416, 3744, and 3769 do not arise naturally from the stochastic model. The construction therefore looks more like a good Diophantine fit than a genuine hidden formula.

4 A simple all-rational triple

A common-denominator approximation is

$$\boxed{(p_{+1}, p_{-1}, p_{+3}) \approx \frac{1}{4853}(2784, 2034, 35).}$$

Numerically this gives

$$(0.573665773748197\ldots, 0.419122192458273\ldots, 0.007212033793530\ldots)\,.$$

The corresponding errors are

$$\left(+1.77375 \times 10^{-6}, -1.80754 \times 10^{-6}, +3.37935 \times 10^{-8}\right).$$

In particular,

$$\frac{35}{4853} = 0.007212033793529\ldots$$

is exceptionally close to 0.007212. The denominator 4853, however, is not small enough to feel canonical, and no independent structural interpretation of it is apparent.

5 Good individual approximations

If one approximates each coordinate separately, several attractive formulas appear.

For the first probability,

$$\boxed{0.573664 \approx \frac{17 + 16\sqrt{5}}{92} = 0.573663996086920\ldots}$$

with error

$$-3.91308 \times 10^{-9}.$$

This is numerically spectacular. Nevertheless, the coefficients 17, 16, and 92 look fitted rather than forced, and the appearance of $\sqrt{5}$ has no known connection with the Lyapunov optimization.

For the second probability,

$$\boxed{0.419124 \approx \frac{18\sqrt{3} - 1}{72} = 0.419123813003330\ldots}$$

with error

$$-1.86997 \times 10^{-7}.$$

This is the same expression as

$$\frac{\sqrt{3}}{4} - \frac{1}{72}.$$

For the third probability,

$$\boxed{0.007212 \approx \frac{3}{416} = 0.007211538461538\ldots}$$

with error

$$-4.61538 \times 10^{-7}.$$

The approximation is simple, though not exceptionally close compared with the number of available small rational fractions.

6 Continued-fraction rational approximations

Natural rational approximants of moderate size include

$$0.573664 \approx \frac{440}{767} = 0.573663624511082\ldots,$$

$$0.419124 \approx \frac{469}{1119} = 0.419124218051832\ldots,$$

and again

$$0.007212 \approx \frac{3}{416} = 0.007211538461538\ldots.$$

Their signed errors are respectively

$$-3.75489 \times 10^{-7}, \qquad +2.18052 \times 10^{-7}, \qquad -4.61538 \times 10^{-7}.$$

These are respectable Diophantine approximations, but none carries an obvious interpretation in the stochastic process.

7 What about transcendental constants?

One can search mechanically for short expressions involving familiar transcendental constants such as

$$\pi, \qquad e, \qquad \log 2, \qquad \log 3,$$

or combinations of these with small rational coefficients. Such searches inevitably produce close approximations once enough templates and coefficients are allowed. No especially simple or compelling expression of this type stands out for the present triple.

This negative observation is not a theorem. It means only that, among the usual low-complexity expressions, there is no approximation that is both markedly accurate and conceptually motivated. An isolated numerical fit to π, e, or a logarithm would be weaker evidence than the exact-sum quadratic-surd triple above, because there is no known analytic mechanism in the finite residue calculation that naturally generates these constants.

8 Speculation on the true optimizer

The most plausible interpretation is that the optimizer is not a famous constant at all. In the numerical calculation, the optimum occurs near the point where two different obstructions exchange dominance:

- a deep 2-adic cylinder related to values near -1;
- a competing recurrent residue class related to values near $+1$.

For a fixed finite modulus, the relevant contraction factor is determined by a large nonnegative operator whose entries depend on the probabilities. The optimal point is obtained when two spectral or max-plus branches meet. Even when the transition probabilities enter rationally, the equality of two such branches may define an algebraic number of high degree. After passing to larger and larger binary moduli, the limiting optimizer could also fail to be algebraic in any simple sense.

Thus several possibilities remain logically open:

1. The limiting optimizer could be rational, but with a large denominator.
2. It could be algebraic of moderately or very high degree.
3. It could be transcendental.

4. The finite-modulus optimizers might not converge to a single exact constant without additional analysis.

There is currently no evidence distinguishing these possibilities. Calling the optimizer "probably a non-famous algebraic number of high degree" is a heuristic guess, not a result. It is motivated only by the finite-dimensional spectral mechanism and by the lack of a visible low-complexity formula.

9 Assessment

The most aesthetically plausible nearby choice is

$$\boxed{p_{-1} = \frac{\sqrt{3}}{4} - \frac{1}{72}, \qquad p_{+3} = \frac{3}{416}, \qquad p_{+1} = 1 - p_{-1} - p_{+3}.}$$

It has three advantages:

- the probabilities sum exactly to one;
- all three components lie within 7×10^{-7} of the certified decimal triple;
- one quadratic irrational is shared coherently between two components.

Nevertheless, this closeness should not be interpreted as evidence that the quadratic-surd triple is the true optimizer. The integers in the formulas are not explained by the model, and a large search space of elementary expressions contains many accidental near-matches.

The safest conclusion is therefore:

> There are several attractive rational and quadratic-surd approximations near $(0.573664, 0.419124, 0.007212)$, but none is presently canonical in a structural sense. The true optimizer is more likely determined by a complicated spectral crossover than by a familiar classical constant.

10 Canonical conjectures beyond cylinder proofs

The cylinder method used in this note is a powerful sufficient method: when a finite residue-weighted contraction is found, almost-sure descent to a bounded region follows. It should not, however, be confused with the stochastic convergence statement itself. A process may converge almost surely even when no fixed finite family of 2-adic cylinders supports a uniform one-step or fixed-block contraction. The two-map obstruction discussed in Section 1 is a central example of this distinction.

It is therefore natural to formulate conjectures directly about the random processes, without requiring that their proofs have cylinder form. In the statements below, "runs into small numbers" means the following hitting property: for every fixed cutoff $H \geq 1$ and every odd starting value $n > H$, the process eventually reaches a value below H with probability one. A weaker formulation with one sufficiently large fixed cutoff would already be interesting, but the all-cutoff version is the clean canonical statement.

Conjecture A: the full three-generator model. For every probability triple

$$p_{+1}, p_{-1}, p_{+3} \geq 0, \qquad p_{+1} + p_{-1} + p_{+3} = 1,$$

and every odd starting value n, the stochastic process generated by

$$T_{+1}(n) = \text{odd}(3n+1), \qquad T_{-1}(n) = \text{odd}(3n-1), \qquad T_{+3}(n) = \text{odd}(3n+3)$$

runs into small numbers with probability one.

This conjecture includes all boundary cases of the probability simplex. In particular, it contains the deterministic $3n+1$ and $3n-1$ problems, as well as the deterministic $3n+3$ variant. It is therefore much stronger than the computer-assisted cylinder results for particular interior probability triples. Those results should be viewed as rigorous islands inside a much larger conjectural parameter region.

Conjecture B: the two-generator $3n \pm 1$ model. For every $p \in [0,1]$ and every odd starting value n, the stochastic process which chooses

$$T_{+1} \quad \text{with probability } p, \qquad T_{-1} \quad \text{with probability } 1-p$$

runs into small numbers with probability one.

The endpoint $p = 1$ is the classical deterministic $3n+1$ Collatz problem. The endpoint $p = 0$ is the deterministic $3n-1$ counterpart. Thus Conjecture B contains both endpoint Collatz problems and interpolates between them by a random mixture.

The discussion in Section 1 shows why Conjecture B cannot be expected to have a straightforward fixed-cylinder proof. Deep cylinders near $+1$ and -1 require incompatible drift inequalities. This is an obstruction to a particular proof architecture, not evidence against the stochastic conjecture. A proof of Conjecture B would probably need a genuinely non-uniform argument, for example a renewal analysis of how trajectories enter and leave deep 2-adic tubes, or a Lyapunov function with infinitely many 2-adic corrections rather than finitely many residue weights.

Statement C: the $3n+1, 3n+3$ benchmark. Consider the two-generator model

$$T_{+1}(n) = \text{odd}(3n+1), \qquad T_{+3}(n) = \text{odd}(3n+3),$$

where T_{+1} is chosen with probability p and T_{+3} with probability $1-p$. Proofs of convergence to small numbers are known for every

$$0.1 \leq p \leq 0.9.$$

This is not merely a numerical observation: it is a substantial rigorous band of parameters. It supports the broader expectation that the stochastic process may converge throughout the full open interval $0 < p < 1$, and perhaps also in the appropriate deterministic endpoint formulations. The known interval should nevertheless be distinguished from that wider conjecture.

The role of T_{+3} is especially transparent. On a deep cylinder of the form $n = 2^k u - 1$, one has

$$T_{+3}(n) = 3u,$$

so a single T_{+3} step removes the entire large power of two. This mechanism breaks the persistent -1 tube that obstructs finite-cylinder proofs in the pure T_{+1}, T_{-1} model. The rigorous interval $0.1 \leq p \leq 0.9$ is therefore both a convergence result and evidence that adding even a controlled amount of the T_{+3} move changes the proof landscape fundamentally.

Prizes and open problems. Several prizes for Collatz problems and variants are listed at

`https://althofer.de/collatz-prizes.html`.

As checked on July 30, 2026, the page includes a prize of 300 Euro for the fair stochastic $3n+1, 3n-1$ model, a 500 Euro prize for the two-player $3n \pm 1$ game, and a 1,000 Euro main prize for the original Collatz conjecture. The page also explicitly welcomes correct computer-assisted proofs. These prizes underline the intended distinction of this section: finite cylinder certificates are valuable rigorous tools, but the central questions concern the stochastic and deterministic dynamics themselves, not one prescribed style of proof.