Oresme Blocks and Prime-Valuation Columns of the Harmonic Numbers

Certified prime clocks, carry compression, hyperrational scales, and Euler–Mascheroni arithmetic

Michael T. M. Emmerich

Technical report, revised 29 September 2026

Abstract

The harmonic numbers admit three complementary readings. First, the signed valuation matrix $(v_p(H_n))_{n,p}$ keeps each finite harmonic number inside $\mathbb{Q}$: rows encode rational factorization, while columns isolate one prime. Starting from Oresme's block proof of divergence, Part I develops several certified realizations of one prime column. Boyd's direct fixed-prime p-adic scan is retained as the natural baseline. In normalized coordinates $X_p(n) = p^{\ell_p(n)} H_n$, the visible output is the exact signed valuation and the internal state is a finite residue. A predecessor-only implementation can recover missing precision by replaying a prefix; a Prime Clockwork minutes address $m = \lfloor n/p \rfloor$ instead gives an exact parent recursion. The main state compression developed here goes one step further: at normalized width R, the carry

$$\kappa_{p,R}(n) = X_p(n) \bmod p^R$$

depends, once n has at least R base-p digits, only on its first R digits. A nonzero carry therefore freezes on an entire suffix cylinder, while a zero carry at an R-digit prefix is equivalent to that prefix lying in $J_p = \{n : v_p(H_n) > 0\}$. This exposes a direct time–space tradeoff between prefix replay, memoized minutes ancestry, and a bounded carry plus scale counter. It also makes repeated states, valuation plateaux, and the -1-per-added-digit descent transparent. The denominator split and the positive-prefix tree are classical; the report compares the carry formulation carefully with Eswarathasan–Levine, Boyd, Kamano, Sanna, and later work rather than attributing their underlying recursions anew.

Part II transfers the exact rational identities to hyperinteger indices. The values ${}^*H_K$ remain hyperrational, so ordered divergence can be studied together with prime-local denominator depth, without replacing the sequence by a continuous interpolant. At every fixed standard precision, the normalized state of an infinite row sees only a finite leading base-p prefix. Under finiteness of J_p, an infinite valuation column decomposes into an infinite scale term plus a bounded arithmetic carry correction.

Part III asks what such hyperrational arithmetic can reveal about the real standard part γ_{EM}. The rational sequence $G_n = 2H_n - H_{n^2}$ satisfies $G_n \to \gamma_{\mathrm{EM}}$ while retaining exact prime columns. Along $n = p^k$ one has $v_p(G_{p^k}) = -2k$, and the normalized units $p^{2k} G_{p^k}$ converge p-adically to a prime-dependent limit $-\Lambda_p$. We separate this arithmetic shadow from the real standard part, show by examples exactly what standard-part projection forgets, and analyze the reduced denominators of G_n. In particular, $B_n = \exp(n^2 + o(n^2))$ for the reduced denominator B_n of G_n, so this specific rational

approximation does not meet the elementary denominator-scaled criterion that would imply irrationality of γ_{EM}. The purpose is not to assign a p-adic valuation to γ_{EM}, but to keep the real and prime-local information of exact approximants simultaneously visible.

Part I
Rational harmonic arithmetic

1 One sequence, two directions of reading

Oresme's fourteenth-century proof of the divergence of the harmonic series is discrete from beginning to end: a suitable grouping of rational summands reveals persistent growth. Powers of 2 are the familiar choice, but powers of any integer greater than one work just as well [1, 16]. When that integer is a prime, the same scale also organizes divisibility in the partial sums.

Our object is the signed valuation matrix

$$H_n = \sum_{j=1}^{n} \frac{1}{j}, \qquad \mathcal{V} = (v_p(H_n))_{n\geq 1,\ p \text{ prime}}.$$

A complete row reconstructs one positive rational. A column is an arithmetic sequence that can be generated without consulting another prime column. This row–column organization is the viewpoint developed here, in connection with the Prime Clockwork representation [6].

The word *autonomous* needs an operational meaning. A valuation alone does not usually determine the next one: addition can cancel leading p-adic digits. We therefore separate the *displayed valuation* from the *local arithmetic state*. Boyd's direct fixed-prime method [3, §5.1] supplies the basic sequential model. We retain a normalized predecessor scan and its replay-based certification because it makes the information and precision issues explicit. Section 7 then introduces two progressively more clock-like organizations. First the minutes reading $m = \lfloor n/p \rfloor$ addresses an earlier state of the same column. Second, the minutes recursion is compressed into a small normalized carry: at fixed width R only the first R base-p digits can affect the carry, while all later suffix digits change only the scale counter. This turns the time cost of replay and the space cost of memoized history into an explicit algorithmic tradeoff.

The carry viewpoint is particularly useful for analysis. It identifies repeated states, constant valuation cylinders, and the exact monotone descent under appended base-p digits. It also shows that the only prefixes that require further refinement are precisely the positive harmonic prefixes in J_p. That exceptional tree is not new—it is central to the Eswarathasan–Levine/Boyd recursion—but the normalized carry formulation makes its complement, namely the full signed valuation column, collapse into simple frozen subtrees. Section 9 compares these statements with Boyd's direct scan, his sparse accelerated tree computation, and later work. The distinction between certified finite prefixes and one uniform precision for an infinite column will be particularly important.

Finally, the same identities can be evaluated at hyperinteger indices. This preserves the exact harmonic sums in the hyperrationals and places ordered growth beside prime-local denominator depth. These are two measurements, not two proofs of the same kind of

divergence.

2 Signed prime valuations of rational numbers

We begin with the coordinate system.

Definition 2.1 (Signed p-adic valuation)**.** Let p be a prime and let $x = a/b \in \mathbb{Q}^\times$ with $a, b \in \mathbb{Z}$ coprime and $b > 0$. Define

$$v_p(x) = v_p(a) - v_p(b),$$

where, for an integer $m \neq 0$, $v_p(m)$ is the largest $e \geq 0$ such that $p^e \mid m$.

Thus positive values record numerator factors, negative values record denominator factors, and zero means that p occurs in neither after reduction. For a positive rational x,

$$x = \prod_p p^{v_p(x)},$$

with only finitely many nonzero exponents. The complete valuation row therefore determines x uniquely.

Example 2.2. The seventh harmonic number is

$$H_7 = \frac{363}{140} = \frac{3 \cdot 11^2}{2^2 \cdot 5 \cdot 7}.$$

Hence

$$(v_2, v_3, v_5, v_7, v_{11})(H_7) = (-2, 1, -1, -1, 2).$$

The signs are essential: the same coordinate system simultaneously records numerator and denominator structure.

Multiplication and division are coordinatewise:

$$v_p(xy) = v_p(x) + v_p(y), \qquad v_p(x/y) = v_p(x) - v_p(y).$$

For completeness set $v_p(0) = +\infty$. Addition is subtler. The ultrametric inequality gives

$$v_p(x + y) \geq \min\{v_p(x), v_p(y)\},$$

with equality whenever the two valuations are unequal. Equal valuations are exactly where cancellation may occur. This simple fact will drive the column dynamics of the harmonic numbers.

3 Rows and columns of the harmonic table

Table 1 shows the first fifteen harmonic numbers and a window of their signed prime coordinates. It is only a window: numerator primes larger than 13 are suppressed, so a displayed row is not by itself a complete factorization vector. The full row, containing all primes, has finite support and reconstructs H_n.

Two ways of reading the table are mathematically natural.

Table 1: A prime-coordinate window for H_n, $1 \le n \le 15$.

n	H_n	v_2	v_3	v_5	v_7	v_{11}	v_{13}
1	1	0	0	0	0	0	0
2	3/2	-1	1	0	0	0	0
3	11/6	-1	-1	0	0	1	0
4	25/12	-2	-1	2	0	0	0
5	137/60	-2	-1	-1	0	0	0
6	49/20	-2	0	-1	2	0	0
7	363/140	-2	1	-1	-1	2	0
8	761/280	-3	0	-1	-1	0	0
9	7129/2520	-3	-2	-1	-1	0	0
10	7381/2520	-3	-2	-1	-1	2	0
11	83711/27720	-3	-2	-1	-1	-1	0
12	86021/27720	-3	-2	-1	-1	-1	2
13	1145993/360360	-3	-2	-1	-1	-1	-1
14	1171733/360360	-3	-2	-1	-1	-1	-1
15	1195757/360360	-3	-2	-1	-1	-1	-1

- **Row view.** Fix n. The full vector $(v_p(H_n))_p$ is a multiplicative coordinate description of the rational number H_n.
- **Column view.** Fix p. The sequence

$$a_n^{(p)} := v_p(H_n), \qquad n = 1, 2, 3, \ldots,$$

  can be generated by the certified prime-local state of Section 6.

The geometry of this distinction is summarized in Figure 1.
The 2-column already hints at a rigid rule:

$$0, -1, -1, -2, -2, -2, -2, -3, \ldots$$

Its proof uses the same powers of 2 as Oresme's grouping, but a different principle: uniqueness of the term of least valuation.

4 Oresme's proof and its prime-power variant

Oresme's argument groups the harmonic series into blocks whose endpoints are powers of 2:

$$1 + \frac{1}{2} + \left(\frac{1}{3} + \frac{1}{4}\right) + \left(\frac{1}{5} + \frac{1}{6} + \frac{1}{7} + \frac{1}{8}\right) + \cdots.$$

The block

$$\frac{1}{2^k + 1} + \cdots + \frac{1}{2^{k+1}}$$

contains 2^k terms, each at least $2^{-(k+1)}$, and hence contributes at least $1/2$. Therefore

$$H_{2^K} \ge 1 + \frac{K}{2}$$

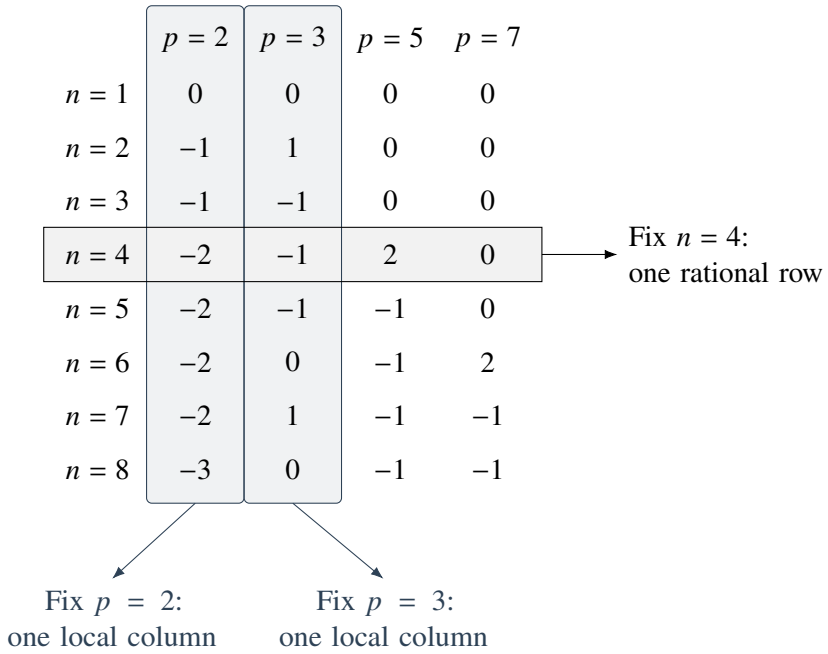

The columns do not need to consult one another. A full row, by contrast, combines all prime coordinates to describe one rational number.

Figure 1: The row–column viewpoint. "Autonomous" means prime-local computation, not that the valuation alone always determines the next valuation.

for every positive integer K, so the harmonic series diverges. The historical provenance and pedagogical significance of this proof are discussed, for example, by Struik and Babb [16, 1].

The number 2 is not essential to the comparison argument.

Definition 4.1 (Oresme p-block)**.** For an integer $p > 1$ and $k \geq 0$, define

$$B_{p,k} := \{p^k + 1, p^k + 2, \ldots, p^{k+1}\}.$$

Thus the left endpoint p^k is excluded and the right endpoint p^{k+1} is included. Consecutive blocks are disjoint and, together with $\{1\}$, partition the positive integers.

Proposition 4.2 (Prime-power Oresme estimate)**.** *For every integer $p > 1$ and $k \geq 0$,*

$$\sum_{j=p^k+1}^{p^{k+1}} \frac{1}{j} \geq \frac{p-1}{p}.$$

Consequently,

$$H_{p^K} \geq 1 + K\frac{p-1}{p}.$$

Proof. The block contains

$$p^{k+1} - p^k = (p-1)p^k$$

terms, and every term is at least $1/p^{k+1}$. Hence its sum is at least

$$(p-1)p^k \frac{1}{p^{k+1}} = \frac{p-1}{p}.$$

Summing the first K blocks gives the second assertion. □

The same Oresme block has an exact p-adic description when p is prime.

Proposition 4.3 (Arithmetic Oresme block)**.** *Let p be prime and put*

$$\Delta_{p,k} := H_{p^{k+1}} - H_{p^k} = \sum_{j=p^k+1}^{p^{k+1}} \frac{1}{j}.$$

Then

$$\boxed{\frac{p-1}{p} \le \Delta_{p,k} < p-1, \qquad v_p(\Delta_{p,k}) = -(k+1).}$$

Proof. The lower bound is Proposition 4.2. For the upper bound, every summand is strictly smaller than $1/p^k$, and there are $(p-1)p^k$ summands, hence $\Delta_{p,k} < p-1$.

For the valuation, the final summand $1/p^{k+1}$ has valuation $-(k+1)$. It is the unique term in the block with that valuation: every other denominator $j < p^{k+1}$ satisfies $v_p(j) \le k$, so $v_p(1/j) \ge -k$. A sum with a unique term of minimal p-adic valuation has the same valuation as that term. Therefore

$$v_p(\Delta_{p,k}) = -(k+1).$$

□

For $p = 2$, the valuation assertion is contained in the more general harmonic-difference results of Lengyel [9, Lemma 1 and Theorem 3]. This proposition is a useful bridge. In the real order each block contributes an amount bounded below by the same positive constant; in the p-adic column the very same block penetrates one level deeper at every step. Both statements concern an exact rational number, with no continuous approximation introduced.

For $p = 3$ this reads

$$\left(\frac{1}{4} + \cdots + \frac{1}{9}\right) \ge \frac{2}{3}, \qquad \left(\frac{1}{10} + \cdots + \frac{1}{27}\right) \ge \frac{2}{3},$$

and similarly at every higher power of 3.

It is useful to distinguish the Oresme summation block from a second interval convention used below and also common in p-adic harmonic-number work:

$$I_{p,k} := \{p^k, p^k + 1, \ldots, p^{k+1} - 1\}.$$

Thus $B_{p,k}$ and $I_{p,k}$ contain the same number of integers but differ at their endpoints. The first is convenient for partitioning the summands of the harmonic series; the second is convenient because the integer scale exponent is k throughout $I_{p,k}$.

5 The leading p-adic layer inside a prime-power scale

Define the integer scale $\ell_p(n)$ by

$$p^{\ell_p(n)} \le n < p^{\ell_p(n)+1}, \qquad n \ge 1. \tag{1}$$

It can be found by integer comparisons; the notation $\lfloor \log_p n \rfloor$ is equivalent but not needed for computation. Write $a_p(n) = v_p(H_n)$. The localization $\mathbb{Z}_{(p)}$ consists of rationals with denominator prime to p. Their reduction modulo p^R uses the inverse of the denominator in $\mathbb{Z}/p^R\mathbb{Z}$.

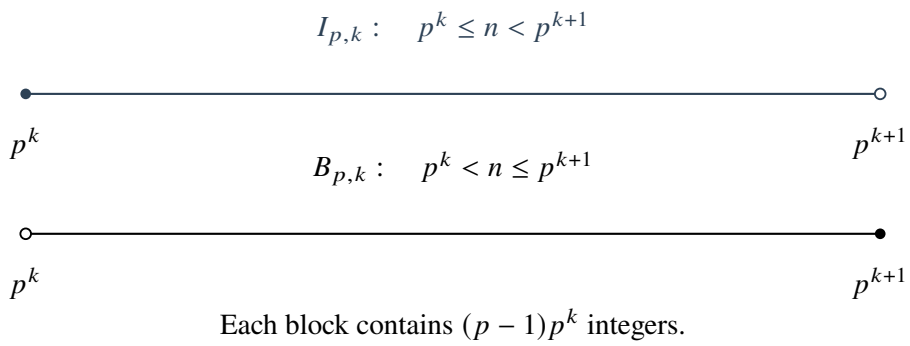

Figure 2: Two endpoint conventions on the same prime-power scale. A filled endpoint is included; an open endpoint is excluded. Both blocks consist of integer indices.

The interval $I_{p,k}$ has an immediate p-adic meaning. Let $n \in I_{p,k}$ and write

$$q = \left\lfloor \frac{n}{p^k} \right\rfloor, \qquad 1 \le q \le p-1.$$

Among the denominators $1, \dots, n$, those with maximal p-adic order k are exactly

$$p^k, 2p^k, \dots, qp^k.$$

They determine the leading p-adic digit of the scaled sum. The noncancelling case below is also Sanna's Lemma 3.2 [12]; a more general same-prime reduction appears in Kamano's Proposition 4 [7].

Lemma 5.1 (Leading block congruence)**.** *Let p be prime, $k \ge 0$, and $p^k \le n < p^{k+1}$. Put $q = \lfloor n/p^k \rfloor$. Then, in the localization $\mathbb{Z}_{(p)}$,*

$$\boxed{p^k H_n \equiv H_q \pmod{p}.}$$

Here $1 \le q < p$, so the denominator of H_q is prime to p and its residue modulo p is well defined.

Proof. After multiplying by p^k, a summand $1/j$ contributes a nonzero residue modulo p only if $v_p(j) = k$. Since $n < p^{k+1}$, such denominators are precisely $j = mp^k$ for $1 \le m \le q$. Their contribution is

$$\sum_{m=1}^{q} \frac{p^k}{mp^k} = \sum_{m=1}^{q} \frac{1}{m} = H_q.$$

Every other denominator has $v_p(j) \le k-1$, so $v_p(p^k/j) \ge 1$ and its contribution vanishes modulo p. □

Corollary 5.2. *With the notation of Lemma 5.1, if*

$$H_q \not\equiv 0 \pmod{p},$$

then

$$v_p(H_n) = -k.$$

The simplest case is $q = 1$, for which $H_q = 1$. We immediately obtain a fact that will later transfer to hyperinteger indices.

Corollary 5.3 (Exact valuation at prime powers)**.** *For every prime p and every integer $k \geq 0$,*

$$\boxed{v_p(H_{p^k}) = -k.}$$

More generally,

$$p^k \leq n < 2p^k \quad \Longrightarrow \quad v_p(H_n) = -k$$

whenever $p > 2$; for $p = 2$ this first subinterval is the whole scale interval.

The dyadic staircase

When $p = 2$, the quotient q can only be 1. Hence no leading cancellation is possible anywhere in $I_{2,k}$.

Proposition 5.4 (Dyadic valuation rule)**.** *For every $n \geq 1$,*

$$\boxed{v_2(H_n) = -\ell_2(n).}$$

Thus the 2-column is a deterministic staircase: it falls by one precisely when n reaches a power of 2. This arithmetic counterpart of Oresme's dyadic grouping is recorded, for example, in Boyd [3, p. 292] and Kamano [7, p. 1]. The same powers of 2 organize both an archimedean lower bound and the exact 2-adic denominator depth.

The first odd prime: $p = 3$

For $p = 3$ there are two possible values of q.

$$q = 1 : \qquad H_1 = 1 \not\equiv 0 \pmod 3,$$

so

$$3^k \leq n < 2 \cdot 3^k \quad \Longrightarrow \quad v_3(H_n) = -k.$$

But

$$q = 2 : \qquad H_2 = 1 + \frac{1}{2} \equiv 1 + 2 \equiv 0 \pmod 3.$$

Therefore the leading terms cancel on the second subblock

$$2 \cdot 3^k \leq n < 3^{k+1},$$

and the valuation rises above $-k$. At the first scale this gives

$$(v_3(H_3), \ldots, v_3(H_8)) = (-1, -1, -1, 0, 1, 0).$$

The leading congruence predicts exactly where the simple plateau must end; finer 3-adic unit information decides what happens after the cancellation.

6 A certified autonomous column recursion

6.1 Why a residue, rather than only a valuation?

The elementary transition

$$a_p(n+1) = \min\{a_p(n), -v_p(n+1)\}$$

is exact when the two entries of the minimum differ. At a tie it is only a lower bound. For example, $v_5(H_3) = v_5(1/4) = 0$, but

$$H_4 = \frac{11}{6} + \frac{1}{4} = \frac{25}{12}, \qquad v_5(H_4) = 2.$$

The missing information is local residue information, not another prime coordinate. Keeping an unrestricted exact rational unit would merely repackage the entire rational number. Instead we retain only a residue with a precise certification rule.

6.2 Normalization and exact transition

Define

$$X_p(n) = p^{\ell_p(n)} H_n, \qquad D_p(n) = v_p(X_p(n)) = \ell_p(n) + a_p(n). \tag{2}$$

Each summand $p^{\ell_p(n)}/j$ is p-integral, so $X_p(n) \in \mathbb{Z}_{(p)}$ and $D_p(n) \geq 0$. This is an *index-scale normalization*: the multiplier is known from n, before the harmonic valuation has been computed. It is not the unit normalization $H_n = p^{a_p(n)} u_n$, whose exponent is precisely the unknown output. Nor need $X_p(n)$ be an ordinary integer: $X_3(7) = 1089/140$ is a rational with denominator prime to 3. Such rationals have well-defined residues because their denominators are invertible modulo p^R.

Fix a precision $R \geq 1$ and write

$$x_n^{[R]} = X_p(n) \bmod p^R, \qquad 0 \leq x_n^{[R]} < p^R.$$

The arithmetic state consists of this residue and the known index scale; there is no harmonic numerator or denominator to factor. The residue is an incomplete representation of $X_p(n)$, but it can certify an exact valuation. The distinction between the internal representation and the certified output is essential.

Theorem 6.1 (Normalized transition and invariant)**.** *Let $n+1 = p^e m$, where $p \nmid m$, and put $\delta = \ell_p(n+1) - \ell_p(n) \in \{0, 1\}$. Then*

$$\boxed{X_p(n+1) = p^{\delta} X_p(n) + p^{\ell_p(n+1)-e} m^{-1}.} \tag{3}$$

Consequently, the initial state $x_1^{[R]} = 1$ and the update

$$x_{n+1}^{[R]} = \left(p^{\delta} x_n^{[R]} + p^{\ell_p(n+1)-e} \operatorname{inv}_{p^R}(m)\right) \bmod p^R \tag{4}$$

maintain $x_n^{[R]} \equiv X_p(n) \pmod{p^R}$ for every $n \geq 1$. Here $\operatorname{inv}_{p^R}(m)$ is the modular inverse of m.

Proof. Multiply $H_{n+1} = H_n + 1/(n+1)$ by $p^{\ell_p(n+1)}$. Since $p^e \leq n+1$, the exponent $\ell_p(n+1) - e$ is nonnegative. All terms therefore lie in $\mathbb{Z}_{(p)}$, and reduction modulo p^R gives (4). Induction establishes the invariant. □

The inverse requires no factorization of m: since $p \nmid m$, the extended Euclidean algorithm finds integers u, v with $um + vp^R = 1$, and $u \bmod p^R$ is the required residue.

This normalization makes the precision bookkeeping transparent. The transition never divides a stored residue by p; multiplication by p^δ cannot move an unknown high-order digit into a lower-order position. Thus the guaranteed precision modulo p^R in $X_p(n)$ is preserved. At a scale boundary it is even possible to retain one additional normalized digit; a fixed-R implementation simply does not retain that extra digit (see Section 9.4). The corresponding absolute precision in H_n is $p^{R-\ell_p(n)}$, which does change with the scale. No claim of precision independent of normalization is intended.

Proposition 6.2 (Certified readout)**.** *If $x_n^{[R]} \neq 0$, set $t = v_p(x_n^{[R]}) < R$. Then*

$$\boxed{a_p(n) = t - \ell_p(n).} \tag{5}$$

If $x_n^{[R]} = 0$, the certified information is only

$$a_p(n) \geq R - \ell_p(n). \tag{6}$$

In particular, a zero residue is not evidence that $H_n = 0$ or that its valuation is infinite.

Proof. In the nonzero case, the difference $X_p(n) - x_n^{[R]}$ has valuation at least $R > t$, so $v_p(X_p(n)) = t$. In the zero case that difference is $X_p(n)$ itself and its valuation is at least R. □

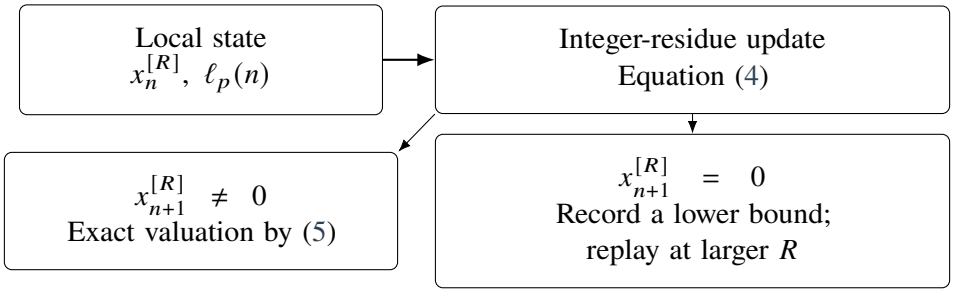

Figure 3: A finite residue has an explicit contract: an exact valuation or a certified lower bound. Precision recovery replays the column; it does not invent missing digits.

6.3 Algorithm, stopping rule, and termination

A *pass* scans $1 \leq n \leq N$ at one precision. The known counter $P = p^{\ell+1}$ signals the next scale change. Algorithm 1 uses only integer operations, modular inversion, and division of an index by the single prime p.

Precision recovery. Start with $R = 1$. Run ColumnPass(p, N, R). If every entry is resolved, return the column. Otherwise replace R by $2R$ and replay the entire pass from $n = 1$. Continue until no unresolved entry remains. A previous nonzero readout is already correct, but a previous residue cannot be extended merely by appending zero digits at the same row.

Theorem 6.3 (Unconditional finite-prefix termination)**.** *For every prime p and finite $N \geq 1$, the precision-recovery procedure terminates and returns exactly $(a_p(1), \ldots, a_p(N))$. More explicitly, any R satisfying*

$$p^R > p^{\ell_p(N)} N\, N! \tag{7}$$

Algorithm 1 ColumnPass(p, N, R)

Require: A prime p, integers $N \geq 1$ and $R \geq 1$.
Ensure: A list A of exact valuations or the marker unresolved.

$M \leftarrow p^R$, $\ell \leftarrow 0$, $P \leftarrow p$, $x \leftarrow 1$, $A \leftarrow [0]$
for $j = 2, \ldots, N$ **do**
  if $j = P$ **then**
    $\ell \leftarrow \ell + 1$; $P \leftarrow pP$; $x \leftarrow px \bmod M$
  end if
  $e \leftarrow 0$; $m \leftarrow j$
  while $p \mid m$ **do**
    $m \leftarrow m/p$; $e \leftarrow e + 1$
  end while
  $x \leftarrow (x + p^{\ell - e} \operatorname{inv}_M(m)) \bmod M$
  if $x = 0$ **then**
    Append unresolved to A ▷ $a_p(j) \geq R - \ell$
  else
    Append $v_p(x) - \ell$ to A
  end if
end for
return A

resolves the entire prefix in one pass.

Proof. For $1 \leq j \leq N$, write $j = p^{e_j} m_j$ with $p \nmid m_j$, and set $Q = \prod_{j=1}^{N} m_j$. This is an integer prime to p, with $Q \leq N!$. For each $n \leq N$,

$$Y_n = QX_p(n) = \sum_{j=1}^{n} p^{\ell_p(n) - e_j} \frac{Q}{m_j}$$

is a positive integer. Moreover,

$$0 < Y_n \leq Q p^{\ell_p(N)} N \leq p^{\ell_p(N)} N\, N!.$$

Since $p \nmid Q$, we have $v_p(Y_n) = D_p(n)$. If (7) holds, the positive integer $Y_n < p^R$ cannot be divisible by p^R. Thus $D_p(n) < R$ for every $n \leq N$. Proposition 6.2 resolves every entry. Doubling R eventually reaches such a value. □

The factorial bound is deliberately conservative: it proves termination without knowing a harmonic numerator, a different prime column, or the finiteness of J_p. It need not be computed in the adaptive algorithm. If

$$R_*(p, N) = 1 + \max_{1 \leq n \leq N} D_p(n),$$

then one pass succeeds exactly when $R \geq R_*(p, N)$. Starting at one and doubling uses final precision less than $2R_*(p, N)$ and at most $1 + \lceil \log_2 R_*(p, N) \rceil$ passes. The logarithm here only counts algorithmic stages; it is not a normalization of H_n.

6.4 A cancellation followed by a scale reset

For $p = 5$ the first five normalized values are

$$X_5(1) = 1, \quad X_5(2) = \frac{3}{2}, \quad X_5(3) = \frac{11}{6}, \quad X_5(4) = \frac{25}{12}, \quad X_5(5) = \frac{137}{12}.$$

Table 2 shows why three digits suffice to distinguish the cancellation at $n = 4$.

Table 2: Normalized residues for $p = 5$. A zero residue requires more precision; it is not assigned a valuation.

n	$\ell_5(n)$	$X_5(n) \bmod 5$	mod25	mod125	$a_5(n)$
1	0	1	1	1	0
2	0	4	14	64	0
3	0	1	6	106	0
4	0	0	0	75	2
5	1	1	1	1	−1

At $n = 4$, $75 = 3 \cdot 5^2$ certifies $a_5(4) = 2$ modulo 125. At $n = 5$, the scale changes and

$$X_5(5) = 5X_5(4) + 1 \equiv 1 \pmod{125},$$

so $a_5(5) = 0 - 1 = -1$. In general, at every scale boundary

$$X_p(p^{k+1}) = pX_p(p^{k+1} - 1) + 1 \equiv 1 \pmod{p}. \tag{8}$$

The prime-power valuation $-(k+1)$ is visible directly in the state transition, without consulting another column.

6.5 Computational scope and verification

A pass performs $N - 1$ modular inversions and $N - 1$ residue updates. The displayed powers can be evaluated by repeated squaring; modular inverses use the extended Euclidean algorithm. Index factor removal and valuation readout use repeated division by p. With $b = \lceil R \log_2 p \rceil + \lceil \log_2(N+1) \rceil$, the operands and scale counters have $O(b)$ bits. A streaming pass needs $O(b)$ working bits besides its emitted outputs; the supplied implementation additionally stores the returned list and residues. This is not an assertion of constant precision for all n, nor a claim of superiority to sparse tree searches at enormous indices.

The accompanying Python implementation was checked against exact `Fraction` arithmetic for $1 \le n \le 2000$ and all 26 primes $p \le 101$: 52,000 valuation comparisons. An additional 12,480 tests checked normalized residues and their readout at $R = 1, 2, 3, 4$; 49,964 eligible parent–child transitions checked the descent law below. For the persistent clock, 10,400 output snapshots (26 primes, 200 rows, and starting widths one and three) were checked against exact values and residues; those runs exercised 51 precision-recovery replays. A further 144 checks verified the one-digit precision gain at scale boundaries described in Section 9.4. The tests are reproducible checks, not substitutes for the invariant and termination proofs. No rational harmonic sums are used by the column algorithm itself; they occur only in the independent test oracle.

7 Constructing a column one tick at a time

The normalized recurrence can be operated as a persistent arithmetic clock, rather than restarted for each new requested upper limit. Its visible output is a valuation; its internal mechanism stores more than that one integer.

7.1 The complete local state and the meaning of autonomy

For fixed p and R, use the state

$$S_n = (n, \ell, P, x), \qquad \ell = \ell_p(n), \quad P = p^{\ell+1}, \quad x = X_p(n) \bmod p^R. \tag{9}$$

The scale counter ℓ and next boundary P are redundant functions of n, cached to avoid repeated scale searches. Initialize at

$$S_1 = (1, 0, p, 1), \qquad a_p(1) = 0.$$

For a tick, set $j = n + 1$. If $j = P$, replace (ℓ, P, x) by $(\ell + 1, pP, px \bmod p^R)$. Then remove the factors of p from j, writing $j = p^e m$, and perform

$$x \longleftarrow \left(x + p^{\ell-e} \operatorname{inv}_{p^R}(m)\right) \bmod p^R, \qquad n \longleftarrow j. \tag{10}$$

Theorem 6.1 proves that this state transition preserves (9). If $x \neq 0$, the displayed output is $v_p(x) - \ell$; otherwise the display is unresolved and the lower bound is $R - \ell$.

There are three distinct kinds of information here: the shared index n, the prime-local residue x, and the displayed valuation. An integer divisibility clock reporting only $e = v_p(j)$ does not provide the whole input for the harmonic tick: one also needs the unit residue $m \bmod p^R$. Both are computed from j using the single prime p. No factorization of m, and no other harmonic valuation column, is required. For example, knowing only that $3 \mid j$ but $9 \nmid j$ does not distinguish the inverse contributions of $j = 3$ and $j = 6$.

Thus *autonomy* means absence of cross-prime state dependencies. With n included in the state, the fixed-p, fixed-R transition is a deterministic autonomous map. With n viewed as an external counter, it is a locally driven recurrence. It is not an input-free clock or a claim of statistical independence. Nor is the entire signed output bounded-state: the scale $\ell_p(n)$ grows without bound. There is, however, a stronger finite-state statement about the *arithmetic correction*. Section 7.5 proves that at width R the normalized carry depends only on the first R base-p digits. Thus a fixed-width column is a finite carry mechanism plus an unbounded but elementary scale counter.

For any finite set of selected primes, one can either run each clock separately or broadcast the same successive index to all of them. The same certified entries result. In the predecessor-only implementation, a precision increase may replay that p-clock only, never another prime column. The memoized minute-hand refinement of Section 7.4 removes full-prefix replay by querying ancestors in the same column. The carry compression of Section 7.5 removes the need to retain that ancestry for valuation queries: at fixed width it keeps only a short leading prefix, a residue carry, and the scale. As Table 1 emphasizes, a preselected finite collection of columns need not contain every prime in a harmonic numerator.

7.2 A complete 3-adic trace

Take $p = 3$ and $R = 3$, so the internal arithmetic is modulo 27. Table 3 is generated by the tick rule itself, without constructing the reduced fractions H_n.

Table 3: The 3-clock through row 9. The state is normalized; subtracting the scale from its valuation gives the displayed output.

n	$\ell_3(n)$	$x_n^{[3]}$	$v_3(x_n^{[3]})$	$a_3(n)$
1	0	1	0	0
2	0	15	1	1
3	1	19	0	−1
4	1	13	0	−1
5	1	19	0	−1
6	1	6	1	0
7	1	18	2	1
8	1	15	1	0
9	2	19	0	−2

At the tick $2 \to 3$, the scale changes. The new denominator is $3 = 3^1 \cdot 1$, so

$$x_3^{[3]} \equiv 3x_2^{[3]} + 1 = 46 \equiv 19 \pmod{27}.$$

The output is $v_3(19) - 1 = -1$. At the tick $6 \to 7$ there is no scale change, but cancellation matters. Since $7^{-1} \equiv 4 \pmod{27}$,

$$x_7^{[3]} \equiv 6 + 3 \cdot 4 = 18 \pmod{27}, \qquad a_3(7) = 2 - 1 = 1.$$

The stored residue resolves the cancellation automatically; the old valuation alone would not have done so.

This precision does not suffice indefinitely. At row 22 one obtains

$$x_{22}^{[3]} = 0, \qquad a_3(22) \geq 3 - \ell_3(22) = 1.$$

The predecessor-only baseline can replay at $R = 4$, giving $x_{22}^{[4]} = 27 \pmod{81}$ and therefore $a_3(22) = 3 - 2 = 1$ exactly. A doubling policy starting from $R = 3$ instead uses $R = 6$; its residue is $513 \pmod{729}$, with the same valuation three. Section 7.4 shows how the same refinement is obtained without scanning rows $1, \ldots, 22$ again: the minutes addresses form the short chain $22 \mapsto 7 \mapsto 2$.

7.3 Baseline: persistent output with prefix replay

A certified clock need not know the final index N in advance. We first record the simplest recovery policy because it makes the certification argument transparent; it will then serve as a baseline for the no-replay minutes-hand method. Assume that all outputs through row n have been certified. Attempt the next tick at the current precision. If it is resolved, keep the new state and emit the next valuation. If it is unresolved, enlarge the precision and reconstruct the state at that row from the initial state. Algorithm 2 makes the policy explicit.

Proposition 7.1 (Persistent certified computation)**.** *Starting at S_1, successive calls to* CertifiedTickReplay *return $a_p(2), a_p(3), \ldots$ exactly. Each call terminates, and only finitely many precision increases occur before any prescribed finite row. Previously emitted valuations never need correction.*

Algorithm 2 CertifiedTickReplay: baseline exact output by prefix replay

Require: Prime p, precision R, a state S_n satisfying (9), and certified outputs through n.

$j \leftarrow n + 1$
$T \leftarrow$ the update of S_n by (10), including its scale check
while the residue in T is zero **do**
    $R \leftarrow 2R$
    $T \leftarrow (1, 0, p, 1)$ at modulus p^R
    for $i = 2, \ldots, j$ **do**
        Advance T once by the same scale check and (10)
    end for
end while
Keep T and the final R as the new state
return $v_p(x) - \ell$ from $T = (j, \ell, P, x)$

Proof. Every ordinary tick and every replay satisfies the invariant in Theorem 6.1. A nonzero readout is therefore exact. For any finite target row j, Theorem 6.3 supplies a finite precision resolving every row through j; successive doublings reach it. Previously certified readouts are properties of the exact harmonic numbers and remain valid at all higher precisions. □

A tick means one new *output index*; it does not promise one arithmetic operation or bounded delay. At unchanged precision the clock continues from its saved state without revisiting old rows. Recovery, however, may revisit an entire prefix. This makes the tradeoff very clean: the rolling replay implementation has small working state, but an unresolved row is paid for in *time* by another scan of $1, \ldots, n$. The minutes-hand and carry implementations below spend additional local computation or retained state in order to avoid that scan. Increasing R at the same row cannot in general be achieved from the old residue alone: many higher-precision residues have the same reduction. A checkpoint already stored at the required higher precision could replace some replay, but that is additional stored information, not a consequence of autonomy.

One may alternatively allow an unresolved display and keep ticking at fixed R: the modular state remains correct and later rows can again be resolved. Algorithm 2 chooses the stronger output contract that every emitted entry is exact. Neither policy requires another prime column. A finite prefix is always computable; the separate question of a fixed precision for *every* row is answered by Theorem 8.2 below.

7.4 The minutes hand: refinement without prefix replay

The replay policy above deliberately uses only the predecessor state H_{n-1}. Prime Clockwork suggests a second address that is already present in each p-clock. At time n, write

$$n = pm + r, \qquad m = \left\lfloor \frac{n}{p} \right\rfloor, \qquad 0 \le r < p. \tag{11}$$

The integer clock calls r the seconds reading and m the minutes reading [6]. For integer factorization the minutes address is enough to generate the valuation recursively when the clock rings. For harmonic numbers the visible valuation alone is not enough, but the same address gives an exact recursion for the *normalized residue state*.

Define the p-integral correction

$$U_p(m,r) := \sum_{a=0}^{m-1}\sum_{b=1}^{p-1}\frac{1}{pa+b} + \sum_{b=1}^{r}\frac{1}{pm+b} = \sum_{\substack{1\le j\le pm+r\\ p\nmid j}}\frac{1}{j} \in \mathbb{Z}_{(p)}. \tag{12}$$

Splitting the harmonic sum into denominators divisible and not divisible by p gives the exact identity

$$\boxed{H_{pm+r} = \frac{1}{p}H_m + U_p(m,r).} \tag{13}$$

For $m \ge 1$, put $k = \ell_p(pm+r) = \ell_p(m) + 1$. Multiplication by p^k turns (13) into

$$\boxed{X_p(pm+r) = X_p(m) + p^k U_p(m,r).} \tag{14}$$

Thus the minutes hand points to an earlier state in the *same* column, and the seconds hand selects only the short tail in the correction.

Proposition 7.2 (Low digits are inherited from the minutes address)**.** *Let $n = pm + r \ge p$ and $k = \ell_p(n)$. For every $R \le k$,*

$$X_p(n) \equiv X_p(m) \pmod{p^R}.$$

More generally, to determine $X_p(n) \bmod p^R$ with $R > k$, it is enough to know $X_p(m) \bmod p^R$ and $U_p(m,r) \bmod p^{R-k}$.

Proof. This is immediate from (14). The correction is multiplied by p^k, so its first possible contribution occurs at normalized digit k. □

This is the precise sense in which the minutes hand can replace prefix replay. If a residue at row n vanishes modulo p^R, request a larger width R' and refine the state at $m = \lfloor n/p \rfloor$ to that same normalized width. If the parent is not available at that width, refine its minutes address, and so on. The ancestry

$$n,\ \left\lfloor \frac{n}{p} \right\rfloor,\ \left\lfloor \frac{n}{p^2} \right\rfloor, \dots$$

has only $\ell_p(n) + 1$ nonzero entries. Memoization retains every refinement, so no already traversed prefix of the column needs to be scanned again.

A finite correction evaluator. The identity above is useful only if the correction can itself be evaluated at the requested precision without replaying all previous harmonic rows. For $s \ge 1$ and $p \nmid b$, the finite geometric expansion gives

$$\frac{1}{pa+b} \equiv \sum_{q=0}^{s-1}(-p)^q a^q b^{-q-1} \pmod{p^s}. \tag{15}$$

Let

$$P_q(m) = \sum_{a=0}^{m-1} a^q, \qquad S_{p,q}^{(s)} = \sum_{b=1}^{p-1} b^{-q-1} \pmod{p^s}.$$

Summing (15) yields the finite formula

$$\boxed{U_p(m,r) \equiv \sum_{q=0}^{s-1} (-p)^q P_q(m) S_{p,q}^{(s)} + \sum_{b=1}^{r} (pm+b)^{-1} \pmod{p^s}.} \tag{16}$$

All inverses here are of p-units. The power sums need not be formed by an m-term loop; for example they satisfy

$$P_0(m) = m, \qquad P_q(m) = \frac{m^{q+1} - \sum_{j=0}^{q-1} \binom{q+1}{j} P_j(m)}{q+1}, \tag{17}$$

where the division is exact in the integers. Thus, for fixed p and a requested finite precision, the correction can be evaluated from the minutes reading m and seconds reading r without traversing the intervening harmonic rows. Formula (16) is an elementary finite p-adic expansion; it is not asserted as a new analytic formula. In Boyd's accelerated method the closely related quantity $H_{pm} - H_m/p = U_p(m,0)$ is evaluated by his convergent series of Theorem 5.2 [3, pp. 294–295].

Algorithm 3 MinuteRefine: normalized residue by minutes addressing

Require: Prime p, row $n \geq 1$, target width $R \geq 1$, memoized residues.

if a cached residue for $X_p(n)$ is known to width at least R **then**
  return its reduction modulo p^R
end if
if $n < p$ **then**
  $x \leftarrow \sum_{j=1}^{n} j^{-1} \pmod{p^R}$
else
  $(m,r) \leftarrow \mathrm{divmod}(n,p)$ and $k \leftarrow \ell_p(n)$
  $y \leftarrow \textsc{MinuteRefine}(p,m,R)$
  if $R \leq k$ **then**
    $x \leftarrow y \pmod{p^R}$
  else
    $u \leftarrow U_p(m,r) \pmod{p^{R-k}}$ using (16)
    $x \leftarrow y + p^k u \pmod{p^R}$
  end if
end if
cache (n,R,x), replacing any lower-precision entry
return x

To obtain an exact valuation, begin at any positive width R. If the returned residue is nonzero, output $v_p(x) - \ell_p(n)$. If it is zero, increase R (for example by doubling) and call MinuteRefine again. Higher digits are obtained by recursive refinement of cached ancestors rather than by a fresh scan from row one.

Proposition 7.3 (Certified no-replay recovery)**.** *For every prime p and finite row n, the minute-hand refinement terminates with the exact value $v_p(H_n)$. A precision increase need revisit only the finite minutes ancestry of n together with the required local correction evaluations; it need not replay the prefix $1, \ldots, n$.*

Proof. Equation (14) proves the recursive invariant, and (16) computes each correction to any prescribed finite precision. The ancestry decreases strictly until it reaches $n < p$. Since $X_p(n)$ is a nonzero rational p-adic integer, its valuation $D_p(n)$ is finite. Any width $R > D_p(n)$ exposes a nonzero residue, so repeated precision increases terminate. Memoization preserves previously obtained higher-precision ancestor states. □

Example 7.4 (Resolving H_{22} by the minutes hand)**.** For $p = 3$, the width $R = 3$ gives $X_3(22) \equiv 0 \pmod{27}$ and is unresolved. Request $R = 4$. Since

$$22 = 3 \cdot 7 + 1,$$

the minutes hand points to row 7. To refine row 7, its minutes hand points again to row 2:

$$7 = 3 \cdot 2 + 1.$$

The base state is

$$X_3(2) = H_2 = \frac{3}{2} \equiv 42 \pmod{81}.$$

The finite correction formula gives

$$U_3(2, 1) \equiv 10 \pmod{27}, \qquad U_3(7, 1) \equiv 4 \pmod{9}.$$

Hence

$$X_3(7) \equiv 42 + 3 \cdot 10 = 72 \pmod{81},$$

and then

$$X_3(22) \equiv 72 + 9 \cdot 4 = 108 \equiv 27 \pmod{81}.$$

Therefore $v_3(X_3(22)) = 3$ and

$$\boxed{v_3(H_{22}) = 3 - \ell_3(22) = 1.}$$

Only the ancestry $22 \mapsto 7 \mapsto 2$ is refined; rows $1, \ldots, 21$ are not replayed.

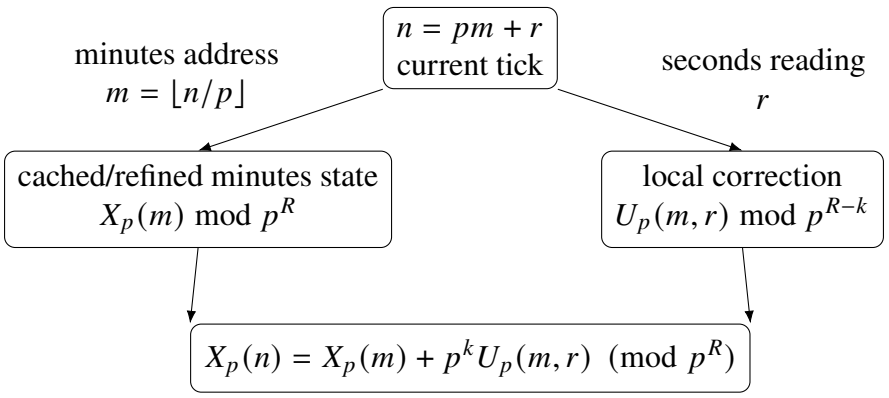

Figure 4: Minutes-hand precision refinement. The common time still advances by +1, but an unresolved row is refined by addressing the earlier minutes state and a p-integral local correction, not by replaying the full prefix.

7.5 From memoized ancestry to a small carry

The memoized minutes recursion is already a no-prefix-replay algorithm, but it still sounds as though a growing history of ancestors is the essential state. For valuation columns this is unnecessary. The normalization itself implies a stronger compression.

Let

$$L_p(n) = \ell_p(n) + 1$$

be the number of base-p digits of n. For $1 \le R \le L_p(n)$ define the R-digit leading prefix

$$\pi_{p,R}(n) = \left\lfloor \frac{n}{p^{L_p(n)-R}} \right\rfloor, \tag{18}$$

and define the normalized carry of width R by

$$\boxed{\kappa_{p,R}(n) := X_p(n) \bmod p^R.} \tag{19}$$

The terminology *carry* is intended literally: it is the bounded p-adic information that must be carried in addition to the unbounded scale counter. It is not the whole rational H_n and it is not the displayed valuation.

Theorem 7.5 (Leading-prefix carry theorem)**.** *If $L_p(n) \ge R$, then*

$$\boxed{\kappa_{p,R}(n) = \kappa_{p,R}(\pi_{p,R}(n)).} \tag{20}$$

Equivalently, after the first R base-p digits have been read, every remaining suffix digit leaves the normalized residue modulo p^R unchanged.

Proof. Write the current integer as $N = pm + r$. As long as N has more than R base-p digits, its scale $k = \ell_p(N)$ satisfies $k \ge R$. The exact minutes identity (14) therefore gives

$$X_p(N) = X_p(m) + p^k U_p(m, r) \equiv X_p(m) \pmod{p^R}.$$

Replacing N by its minutes reading m removes the last base-p digit. Repeat until exactly R leading digits remain. The result is (20). □

The same fact can be written as a direct digit algorithm. Let

$$n = (d_1 d_2 \cdots d_L)_p, \qquad q_j = (d_1 \cdots d_j)_p$$

be its successive leading prefixes. At fixed width R, put $c_j = X_p(q_j) \bmod p^R$. Then

$$c_1 = H_{d_1} \bmod p^R, \tag{21}$$

$$c_j = c_{j-1} + p^{j-1} U_p(q_{j-1}, d_j) \pmod{p^R} \qquad (2 \le j \le \min\{L, R\}), \tag{22}$$

and, when $L > R$,

$$c_{R+1} = c_{R+2} = \cdots = c_L = c_R. \tag{23}$$

Thus one does not need a memoized chain of arbitrary old rows in order to answer a valuation query. At width R, the working prefix satisfies $q_j < p^R$, the carry satisfies $0 \le c_j < p^R$, and at most the first R digits can alter it. Increasing R may require recomputing this short leading calculation because a residue modulo p^R does not contain its next p-adic digit, but it never requires replaying the rows $1, \ldots, n$.

Algorithm 4 CarryRead(p, n, R): normalized carry from leading digits

Require: Prime p, row $n \geq 1$, width $R \geq 1$.

Write $n = (d_1 \cdots d_L)_p$ and set $M \leftarrow p^R$
$q \leftarrow d_1$; $c \leftarrow H_{d_1} \bmod M$
for $j = 2, \ldots, \min\{L, R\}$ **do**
    $s \leftarrow R - j + 1$
    $u \leftarrow U_p(q, d_j) \bmod p^s$ using (16)
    $c \leftarrow (c + p^{j-1}u) \bmod M$
    $q \leftarrow pq + d_j$
end for
return (c, q)

For an exact valuation query, start with $R = 1$. If the returned carry is nonzero, output

$$v_p(H_n) = v_p(c) - \ell_p(n). \tag{24}$$

If $c = 0$, increase R and repeat. When $R \leq L_p(n)$ this reveals one more leading digit of the index; if $R > L_p(n)$ the whole index has already been read and only additional p-adic precision is being requested. Since $X_p(n) \neq 0$, the process terminates for every finite n.

7.6 Repeated states, valuation cylinders, and monotone descent

The carry theorem converts repeated leading prefixes into exact repeated arithmetic states. Let q be an R-digit base-p integer and let

$$c = \kappa_{p,R}(q), \qquad c \neq 0, \qquad d = v_p(c) < R.$$

For every digit length $L \geq R$ define the suffix cylinder

$$I_{q,L} = \{n : L_p(n) = L,\ \pi_{p,R}(n) = q\} = \left[qp^{L-R}, (q+1)p^{L-R} - 1\right]. \tag{25}$$

Corollary 7.6 (Frozen carry cylinders)**.** *For every $n \in I_{q,L}$,*

$$\boxed{D_p(n) = d, \qquad v_p(H_n) = d - (L-1).} \tag{26}$$

Hence the signed valuation is constant on each fixed-length cylinder, whose length is p^{L-R}. Appending one further base-p digit while keeping the same leading carry decreases the valuation by exactly one.

Proof. Theorem 7.5 gives $X_p(n) \equiv c \not\equiv 0 \pmod{p^R}$. The certified readout therefore gives $D_p(n) = v_p(c) = d$. Since $\ell_p(n) = L - 1$, the second formula follows. Increasing the digit length by one increases ℓ_p by one while the carry remains unchanged. □

This is the monotonicity supplied by the clockwork: it is monotonicity down a base-p suffix subtree, not monotonicity of $v_p(H_n)$ as an ordinary function of the consecutive integer n. In a $+1$ scan the same theorem appears as long constant plateaux: the leading R digits and the scale are unchanged throughout each interval $I_{q,L}$.

Repeated carry states can be compressed further. If two R-digit prefixes q and q' have the same nonzero carry, their suffix cylinders have the same normalized state at every later length. For the valuation output alone, it is enough that the two carries have the same valuation d;

after certification the full residue need not be retained. A zero carry is different: two zero states cannot be merged solely because both residues are zero, since their future corrections depend on their actual prefix addresses. Those are precisely the states that remain active.

The raw state reduction is already visible in small examples. The following counts were obtained by the carry implementation; "addresses" means all R-digit leading prefixes and "distinct carries" means distinct residues modulo p^R. The table is descriptive, not a claim of minimal automata.

Table 4: Illustrative carry compression at fixed width.

p	R	prefix addresses	zero carries	distinct carries
2	1	1	0	1
3	1	2	1	2
3	2	6	1	4
3	3	18	1	9
3	4	54	0	22
5	1	4	1	3
5	2	20	2	10
5	3	100	0	37

7.7 Zero carries and the exceptional positive tree

The active zero states have an exact arithmetic meaning.

Proposition 7.7 (Zero carry equals a positive prefix)**.** *Let q have exactly R base-p digits. Then*

$$\boxed{\kappa_{p,R}(q) = 0 \quad \Longleftrightarrow \quad v_p(H_q) > 0 \quad \Longleftrightarrow \quad q \in J_p.} \tag{27}$$

Proof. Here $\ell_p(q) = R - 1$. The residue is zero exactly when $D_p(q) = v_p(X_p(q)) \geq R$. Since $D_p(q) = R - 1 + v_p(H_q)$, this is equivalent to $v_p(H_q) \geq 1$. □

Thus adaptive carry refinement has a particularly simple shape. Read the base-p digits from the left. As soon as the width-R carry is nonzero, the entire suffix subtree is certified by Corollary 7.6; no later digit can alter its normalized depth. If the carry is zero, the R-digit prefix lies in J_p and one more digit or one more precision digit must be inspected. The active refinement trie is therefore exactly the positive-prefix tree already studied by Eswarathasan–Levine and Boyd. The carry formulation does not rediscover that tree; it identifies the enormous complement of that tree with frozen signed-valuation cylinders.

For $p = 3$ the exceptional path is especially transparent:

$$2_3 \quad \longrightarrow \quad 21_3 = 7 \quad \longrightarrow \quad 211_3 = 22 \quad \longrightarrow \quad \text{no four-digit zero carry.}$$

At widths $R = 1, 2, 3$ the unique zero carries occur at 2, 21_3, and 211_3 respectively; at $R = 4$ there are none. Thus four normalized digits resolve every row, but most rows exit much earlier. For $p = 2$ the single one-digit prefix already has nonzero carry, so the dyadic staircase is the one-carry-state extreme.

7.8 Time, space, and which implementation is actually simpler

The different algorithms should be compared by task rather than by a single headline complexity. We count row or address operations here and deliberately do not claim a complete bit-complexity comparison of the correction evaluators.

Table 5: Four useful organizations of the same fixed-prime arithmetic.

Method	retained state	when precision is insufficient	natural use
Boyd direct scan, §5.1	one rolling truncated H_n	choose a larger cutoff; a previously discarded higher digit is not recoverable from the old truncation alone	simplest dense sequential prefix when adequate precision is known
Normalized replay baseline	one rolling $X_p(n) \bmod p^R$	enlarge R and rescan $1, \ldots, n$	simplest explicit certification proof; low space, potentially repeated time
Memoized minutes ancestry	residues for addressed parents at the highest precision reached	refine $n, \lfloor n/p \rfloor, \ldots$	adaptive random access without full-prefix replay; exchanges time for cache space
Compressed carry	leading prefix $q < p^R$, carry $c < p^R$, and scale	increase R and inspect/recompute only the first $\min\{R, L_p(n)\}$ digits	structural analysis, repeated-state detection, plateaux, and random-access valuation queries

This makes the time–space issue explicit. Prefix replay uses almost no history but may pay an $O(n)$ row-scan cost when a cancellation demands more precision. Memoized minutes refinement avoids that replay by storing earlier states, so its natural cost is space proportional to the distinct addressed states and precisions retained. The carry theorem shows that this history is not intrinsic to a valuation query. At fixed width R the relevant address and arithmetic state are bounded by p^R, independently of how far to the right the row n lies. A full lookup table would contain $(p-1)p^{R-1}$ leading-prefix addresses, but one need not build it: the adaptive sparse version refines only zero-carry prefixes, and nonzero branches collapse immediately to terminal depth labels.

There is therefore a clear answer to the question "is the minutes-hand algorithm simpler?" For a dense list $H_1, \ldots, H_N$ at a known sufficient precision, Boyd's direct sequential update is still the simplest implementation and is difficult to improve upon conceptually: one rolling state, one new summand, one output row. The replay/certification version in this report is intentionally more cumbersome; it is retained because it makes exactness and missing information explicit. The memoized minutes-hand version is simpler for adaptive random access but introduces cache management. The compressed carry version removes that cache

and, once the correction evaluator is available, is in my view the simplest formulation for *analysis*: only a short leading prefix can affect the arithmetic correction, repeated states are literal equal carries, and monotone descent is a one-line consequence of carry freezing.

Implementation simplicity therefore depends on the task. The carry reader has slightly more arithmetic machinery than Boyd's direct scan because it needs $U_p(m, r)$; formula (16) supplies an elementary evaluator, while Boyd's Theorem 5.2 can replace it with a more sophisticated accelerated evaluator. But the carry reader eliminates the predict–fail–replay control structure, global checkpoints, and an unbounded memoized history. For one distant valuation it processes at most the leading precision width until certification; for a full dense stream, no claim is made that this beats a sequential scan, since merely emitting N answers already costs N output steps.

The companion implementation `harmonic_minute_hand.py` retains the memoized ancestry version. A separate reference implementation `harmonic_carry_clock.py` implements CarryRead without a history cache. The latter was checked against exact rational arithmetic for 2,400 valuation queries (eight primes through 19 and $1 \le n \le 300$), plus 8,292 explicit carry-freeze checks. These computations test the programs; Theorem 7.5 and Proposition 7.7 supply the mathematical guarantees.

8 Digit recursion and the finite part of a column

The minutes-hand identity (13) is also the natural base-p digit recursion. The carry theorem above is its finite-precision form: a nonzero carry freezes the normalized depth of every suffix descendant, while a zero carry keeps the prefix on the exceptional positive tree. The present section now removes the fixed modulus and states the corresponding exact valuation consequences. If

$$C_p(M) = \sum_{\substack{1 \le j \le M \\ p \nmid j}} \frac{1}{j} \in \mathbb{Z}_{(p)},$$

then $U_p(n, r) = C_p(pn + r)$ and

$$\boxed{H_{pn+r} = \frac{H_n}{p} + C_p(pn + r),} \tag{28}$$

with normalized form

$$X_p(pn + r) = X_p(n) + p^{\ell_p(n)+1} C_p(pn + r). \tag{29}$$

This decomposition underlies the Eswarathasan–Levine construction [5]; it is displayed explicitly in [4, (1.1)]. The preceding section uses it as a memoized state-addressing rule; here we extract its consequences for the shape of the valuation column.

Proposition 8.1 (Exact descent after a nonpositive value)**.** *If $a_p(n) \le 0$, then, for every $0 \le r < p$,*

$$a_p(pn + r) = a_p(n) - 1, \qquad D_p(pn + r) = D_p(n). \tag{30}$$

Consequently, for every $t \ge 0$ and $0 \le s < p^t$,

$$a_p(p^t n + s) = a_p(n) - t. \tag{31}$$

Proof. In (28), the first term has valuation $a_p(n) - 1 < 0$, whereas the second has valuation at least zero. Hence the valuation of the sum is exactly $a_p(n) - 1$. Also $\ell_p(pn + r) = \ell_p(n) + 1$, proving the identity for D_p. Iteration proves (31). □

For example, $H_6 = 49/20$ has $a_3(6) = 0$. It follows immediately that $a_3(18) = a_3(19) = a_3(20) = -1$ and that $a_3(n) = -2$ for $54 \le n \le 62$. These conclusions need no unit calculation on the resulting subtrees. Kamano's more general Proposition 4(i) [7], specialized to ordinary harmonic numbers, gives, for integers $t \ge 0$,

$$a_p(n) = -t \quad \Longleftrightarrow \quad a_p(\lfloor n/p^t \rfloor) = 0 \qquad (n \ge p^t).$$

Thus same-column reduction has a direct predecessor in the literature.

8.1 Where a prefix first leaves the positive tree

Let $J_p = \{n \ge 1 : a_p(n) > 0\}$. Write n in base p with $L = \ell_p(n) + 1$ digits, and let $n_1, \dots, n_L = n$ be its successive nonempty leading prefixes. If $n \notin J_p$, let j be the first position with $n_j \notin J_p$. Then

$$\boxed{a_p(n_j) = 0, \qquad a_p(n) = -(L - j), \qquad D_p(n) = j - 1.} \tag{32}$$

Indeed a one-digit harmonic sum is p-integral; any child of a positive-valuation parent is also p-integral by (28). Thus the first nonpositive value on the path is zero. Apply Proposition 8.1 to the remaining digits.

Equation (32) explains exactly how the J_p tree encodes negative entries, rather than merely marking positive ones. In carry language, the prefixes $n_1, \dots, n_{j-1}$ have zero carries at their own digit widths, while n_j is the first prefix with a nonzero carry. Its valuation is zero, so its normalized depth is $j - 1$; every later suffix keeps that same depth. Thus the first nonzero carry and the first exit from J_p are the same event seen in two coordinate systems.

Given a *complete, certified* finite list J_p with the positive valuations attached, a column entry can be obtained by reading the base-p digits until the first exit. The rest of the calculation is an integer digit count. An incomplete list of observed positive entries is not a certificate and must not be used this way.

Theorem 8.2 (A uniform-precision characterization)**.** *For a fixed prime p, the following are equivalent:*

(i) *J_p is finite;*

(ii) *$D_p(n)$ is bounded for all $n \ge 1$;*

(iii) *some finite R makes $X_p(n) \bmod p^R$ nonzero for every $n \ge 1$.*

If J_p is finite, let h be the largest base-p digit length of one of its elements, or $h = 0$ when J_p is empty. Then for every $n \ge p^h$,

$$a_p(n) = -\ell_p(n) + d_p(n), \qquad d_p(n) \in \{0, \dots, h\}. \tag{33}$$

Proof. Under (i), a prefix must leave J_p within its first $h + 1$ digits. Equation (32) bounds $D_p(n)$ by h outside J_p. Its values on the finite set J_p are bounded as well, proving (ii) and (33). If $D_p(n) \le B$ for all n, choose $R > B$ to obtain (iii). Conversely (iii) implies $D_p(n) < R$. Whenever $n \in J_p$, we have $\ell_p(n) + 1 \le D_p(n) < R$, so such indices lie in a finite initial interval. This proves (i). □

The theorem expresses the finiteness problem as a question about a single column's required residue precision. It is deduced here from the same prefix mechanism as Boyd's finiteness/divergence criterion [3, Proposition 3.3]; it does not establish finiteness for an arbitrary prime. A usable global bound, given a certified J_p, is

$$R = 1 + \max\big(\{h\} \cup \{\ell_p(n) + a_p(n) : n \in J_p\}\big).$$

The contrast is important: finite-prefix termination needs no conjecture, whereas a fixed precision valid on the whole infinite column is equivalent to the finiteness of J_p.

Combined with Theorem 7.5, this gives a finite carry representation of the entire signed column whenever J_p is finite. Choose $R > \sup_n D_p(n)$. Apart from the finitely many rows with fewer than R digits, every n is evaluated from

$$\big(\pi_{p,R}(n),\ \kappa_{p,R}(\pi_{p,R}(n)),\ \ell_p(n)\big).$$

The first two entries range over finite sets; the only unbounded coordinate is the scale. If only the signed valuation is required, each nonzero carry may be compressed further to its depth $d \in \{0, \dots, R-1\}$. A full table of all R-digit addresses is finite but may be wasteful; the adaptive tree keeps only the positive prefixes and terminal exit labels. This is the precise sense in which the clock state can be a small carry rather than an ever-growing history.

8.2 A complete example: the ternary column

For $p = 3$, the positive tree is $J_3 = \{2, 7, 22\}$, each with valuation one [5, 7]. Its base-3 prefixes are 2, 21, and 211. Kamano's complete rule [7, Theorem 2(i)] can be displayed as a table of normalized depths. The short indices not covered by a displayed prefix are handled separately below.

Table 6: Ternary leading prefixes and normalized depth, with $k = \ell_3(n)$. A suffix may be empty. The last two lines apply for $k \geq 2$.

Leading ternary prefix	$D_3(n)$	$a_3(n)$
1	0	$-k$
20 or 22	1	$1-k$
210 or 212	2	$2-k$
211	3	$3-k$

The remaining short positive prefixes are $n = 2$ and $n = 7$, with $D_3(2) = 1$ and $D_3(7) = 2$. Hence $D_3(n) \leq 3$ for *every* n, and a fixed modulus $3^4 = 81$ resolves the whole ternary column. This bound is sharp for a uniform normalized residue readout: $D_3(22) = 3$, so modulus 27 is insufficient there. For the dyadic column $D_2(n) = 0$ for all n, and modulus 2 suffices. In carry language there is one one-digit address, 1_2, with carry 1; all suffixes inherit it. The ternary exceptional carry chain is

$$2_3\ (R = 1), \qquad 21_3\ (R = 2), \qquad 211_3\ (R = 3),$$

after which $R = 4$ has no zero address. The whole signed valuation is still not a finite-state output because its scale counter $-\ell_p(n)$ is unbounded, but its arithmetic correction is a finite-state carry when a uniform R exists.

9 Comparison with related work

The comparison with Boyd must distinguish four issues: the underlying addition identity, the information retained at a particular row, the precision policy along a sequence of rows, and the output task. A difference in any one of these does not establish a difference in the others.

9.1 What Boyd's first method already computes

In §5.1, pp. 293–294, Boyd [3] represents H_n by

$$H_n = \sum_{k=a_p(n)}^{s-1} b_k p^k + O(p^s), \qquad 0 \le b_k < p. \tag{34}$$

His notation $O(p^s)$ means a remainder of valuation at least s [3, §2, p. 290]; it is not a bound on real absolute error. The cutoff s is an *exponent*, not necessarily the number of stored digits. Negative-power digits are present when $a_p(n) < 0$.

Boyd explicitly updates successive rows by adding truncated expansions of H_{n-1} and $1/n$. He explains that carries move to higher powers of p, so they do not introduce uncertainty into already known lower powers. He then gives the condition

$$s > a_p(n) \tag{35}$$

for exact valuation readout [3, p. 294]. This does not require knowing the valuation in advance: the first visible nonzero digit certifies it. If no nonzero digit is visible, the representation certifies only $a_p(n) \ge s$.

Both methods therefore use incomplete numerical representations to produce exact valuation outputs. Boyd's computation is already prime-local and sequential. He does not need another prime column, and the addition of a new term is not an innovation of the present normalization.

Example 9.1 (The cutoff is not circular)**.** For $p = 3$, $H_7 = 363/140$ satisfies

$$H_7 = 0 + O(3), \qquad H_7 = 6 + O(9) = 2 \cdot 3 + O(3^2).$$

The first representation implies only $a_3(7) \ge 1$; the second reveals $a_3(7) = 1$. Since $\ell_3(7) = 1$, the same data in normalized coordinates are

$$X_3(7) = 0 + O(9), \qquad X_3(7) = 18 + O(27).$$

The resolved cases have respectively $s = 2$ and $R = 3$. Each shows the same single nonzero significant digit, at exponent one in H_7 and at exponent two in $X_3(7)$.

9.2 Boyd's prefix tree and the antecedent of the carry tree

The carry compression should also be compared with Boyd's Section 3, not only with his numerical Section 5. Following Eswarathasan–Levine, Boyd groups positive indices into generations

$$G_m = J_p \cap [p^{m-1}, p^m - 1]$$

and arranges $J_p \cup \{0\}$ as a rooted base-p tree. A node n has children $pn + r$, and Boyd attaches a mod-p *type* through an expansion $H_n = ap + O(p^2)$; the type determines which

next digits can remain in the positive tree [3, pp. 291–292]. This is a direct antecedent of any sparse prefix representation and must be credited as such.

Proposition 7.7 identifies exactly the same exceptional skeleton from the normalized side: an R-digit prefix has zero carry modulo p^R precisely when it belongs to J_p. What the carry formulation adds is not the parent map $n \mapsto pn + r$ or the discovery of a small mod-p node type. It adds a complementary description of the *full signed column*: once a prefix leaves Boyd's positive tree, its normalized carry freezes, every suffix cylinder is resolved at once, and the valuation thereafter is the frozen carry depth minus the growing scale. In other words, Boyd's tree organizes the exceptional states; the carry theorem simultaneously collapses all nonexceptional descendants.

There is a further algorithmic distinction. Boyd's tree construction is naturally a generation-by-generation search for J_p. The carry reader is a random-access transducer for an arbitrary row: read leading digits until the first nonzero carry, then stop. When a whole dense prefix is wanted, this random-access organization need not be faster than the direct scan. Its advantage is that state equivalence and suffix invariance are explicit.

9.3 A precision dictionary: equal information at a fixed row

Fix a row n and put $\ell = \ell_p(n)$. Multiplication by p^ℓ gives the exact equivalence

$$H_n = \widetilde{H}_n + O(p^s) \quad \Longleftrightarrow \quad X_p(n) = p^\ell \widetilde{H}_n + O(p^{s+\ell}). \tag{36}$$

Thus, whenever $R = s + \ell \geq 1$, the two conventions are related by

$$\boxed{R = s + \ell, \qquad s = R - \ell.} \tag{37}$$

In particular, the normalized readout criterion is exactly Boyd's criterion in shifted coordinates:

$$R > D_p(n) = \ell + a_p(n) \quad \Longleftrightarrow \quad s > a_p(n).$$

If the first nonzero digit is visible, the number of known significant digit positions is also identical:

$$R - D_p(n) = s - a_p(n).$$

Here R counts slots starting at exponent zero in $X_p(n)$; some of those slots can be zero. It is not automatically the relative precision starting at the first nonzero digit. Normalizing does not reveal an unknown digit, resolve a previously unresolved value at equivalent precision, or reconstruct H_n exactly from finitely many digits.

9.4 Different precision policies along the $+1$ dynamics

Pointwise equivalence does not mean that two fixed numerical parameters give the same data at every row. A fixed cutoff s for H_n corresponds to an increasing normalized width

$$R_n = s + \ell_p(n).$$

A fixed normalized width R, as used in a column pass here, corresponds to a moving cutoff

$$s_n = R - \ell_p(n).$$

The latter may eventually be negative; that is meaningful absolute p-adic precision. It can still resolve a valuation below that cutoff.

There is an instructive distinction at a scale boundary. From $x \equiv X_p(p^{k+1} - 1) \pmod{p^R}$ and (8) one gets

$$X_p(p^{k+1}) \equiv px + 1 \pmod{p^{R+1}}. \tag{38}$$

This extra normalized digit is justified: an old error divisible by p^R is multiplied by p, and the new summand is the exact integer one. Retaining modulus p^{R+1} preserves the old absolute cutoff $s = R - k$ in H_n. Retaining only modulus p^R, as in our fixed-width pass, discards the additional highest digit and changes the cutoff to $s - 1$. Neither choice corrupts the lower digits.

For example, $X_3(8) \equiv 15 \pmod{27}$ implies

$$X_3(9) \equiv 46 \pmod{81}.$$

The fixed-$R = 3$ trace in Table 3 retains only $46 \bmod 27 = 19$. Keeping $46 \bmod 81$ instead would preserve the cutoff $s = 2$ for H_n. This structural gain at a scale change must not be confused with arbitrarily increasing the precision of an already stored residue at a fixed row.

For a finite prefix, fixed absolute precision resolves all rows when $s > \max_{n \le N} a_p(n)$, whereas fixed normalized precision resolves them when $R > \max_{n \le N} D_p(n)$. For the entire column these become two different boundedness questions. Theorem 8.2 establishes the equivalence with finiteness of J_p for bounded $D_p(n)$; no such equivalence is asserted here merely from bounded $a_p(n)$. For example, the complete ternary rule gives $a_3(n) \le 1$ and $D_3(n) \le 3$ for all n. Consequently $s = 2$ suffices in Boyd's direct absolute-cutoff representation, and $R = 4$ suffices in the fixed-width normalized representation. These are different allocations of digits, not different exactness guarantees.

9.5 Recovery, carry compression, and the scope of the distinction

The report deliberately retains several implementations because their contrast clarifies what is arithmetic and what is architecture. Boyd's first method in §5.1 already performs successive fixed-prime addition and exact valuation readout from truncated expansions. If adequate absolute precision is chosen in advance, this rolling scan is arguably the simplest possible dense computation. If higher digits of an already truncated state are later required, those digits are not recoverable from the old truncation alone; the replay baseline in Section 7 makes that information loss explicit rather than attributing it to Boyd.

The minutes identity (13) is likewise classical. Boyd's Theorem 5.2 evaluates the parent correction

$$H_{pm} - \frac{1}{p} H_m = U_p(m, 0)$$

by a convergent p-adic series, and his accelerated algorithm computes children H_{pm+r} successively from selected surviving parents [3, pp. 294–295]. When the supported generations outlive the chosen series precision, he restarts with a larger truncation parameter [3, p. 295]. Thus neither scale parents, local corrections, sparse positive-prefix trees, nor higher-precision restarts are claimed here as new ideas.

The memoized minutes implementation changes the recovery policy: an unresolved row recursively requests missing precision only from its minutes ancestry instead of rescanning every predecessor. That avoids a time cost at the price of cached state. The carry theorem then removes most of that space cost for the valuation problem itself. At normalized width R, only the first R leading digits can alter the state; a nonzero carry certifies every suffix

descendant and can be replaced, for valuation output, by the small depth label $v_p(c)$. Only zero-carry prefixes remain active, and Proposition 7.7 shows that these are exactly the Boyd/ Eswarathasan–Levine positive prefixes. The resulting architecture is a finite carry transducer plus an unbounded scale counter whenever the uniform precision criterion holds.

I have not found the exact normalized freeze statement (20), or its use as a carry quotient for the complete signed column, stated in the sources cited here. That observation should be read narrowly. The theorem is an immediate consequence of the classical parent decomposition, and its active exceptional tree is already present in Eswarathasan–Levine and Boyd. The potentially distinctive point is therefore the *organization*: using the Prime Clockwork minutes address to turn the full signed harmonic column into frozen carry cylinders plus a small exceptional tree, and exploiting that organization for random access, precision localization, repeated-state analysis, and the hyperrational extension below. A priority claim stronger than this would require a wider literature search than undertaken for this technical report.

Formula (16) supplies a self-contained finite correction evaluator by geometric expansion and integer power sums. It is elementary p-adic algebra, not a competing analytic novelty claim. Boyd's coefficient series can be substituted for this evaluator and may be preferable at high precision. Rogers [11, §2.1] explains the precision loss in Boyd's scale transitions, and Carofiglio–Cherubini–Gambini [4, §§1.1, 2] develop the sparse computation further.

For a complete stream through N, the direct scan still has a decisive simplicity advantage: it already visits each requested row once. Carry compression is most compelling for random access and structural questions, where its bounded-prefix theorem replaces an n-step history by a small arithmetic state. Thus the report makes no universal speed claim. It claims a cleaner state decomposition and proves the invariances that follow from it.

Signed valuations and leading digits. Kamano's Proposition 4 and Theorem 2(i) [7] cover same-prime reduction and the full ternary column. Sanna's Lemma 3.2 and Theorem 1.3 [12] study leading-digit plateaus and how often minimal valuation occurs. Consequently the signed negative part of a column is not absent from prior work. Equation (32) makes its reconstruction from the positive prefix tree explicit in the present coordinates.

Coordinate organization and scope. Bell [2] studies p-regularity for valuation sequences of polynomial values; this does not establish it for harmonic numbers. Boyd also puts real and finite-prime sizes together through the product formula [3, p. 292]. Prime Clockwork [6] motivates both the row–column organization and the seconds/minutes address used here. The harmonic construction differs from the integer clock because a valuation alone is insufficient state: a normalized p-adic residue and a p-integral correction are required. Its combination with certified precision refinement, the uniform-precision criterion, and hyperrational Oresme blocks is the present focus. These comparisons identify the mathematical antecedents; they do not settle priority for this particular combined organization.

Part II
Hyperrational rows and divergence

10 Hyperinteger rows and exact arithmetic scales

10.1 What is extended

Let ${}^*\mathbb{R}$ be a nonstandard enlargement with transfer, and ${}^*\mathbb{N}$, ${}^*\mathbb{Q}$ its hyperintegers and hyperrationals [10, 8]. Transfer extends the finite harmonic sum to

$$ {}^*H_K = \sum_{j=1}^{K} \frac{1}{j} \in {}^*\mathbb{Q} \qquad (K \in {}^*\mathbb{N},\ K \geq 1). $$

A hyperfinite sum is an internal extension of a finite sum, not an additional summability rule for the infinite harmonic series. Fix a *standard* prime p. The maps v_p, ℓ_p, X_p, and D_p extend internally. We omit stars on the last three when the index is hyperinteger, but retain ${}^*v_p$ to distinguish its hyperinteger-valued output from an ordinary valuation on $\mathbb{Q}_p$. In particular,

$$ X_p(K) = p^{\ell_p(K)} {}^*H_K \in {}^*\mathbb{Z}_{(p)}, \qquad {}^*v_p({}^*H_K) = -\ell_p(K) + D_p(K). \tag{39} $$

No completion or continuous interpolation of H_n has been substituted for these exact values.

10.2 Ordered divergence without logarithmic subtraction

The Oresme estimates give for every finite $n \geq 1$

$$ 1 + \frac{p-1}{p}\ell_p(n) \leq H_n < 1 + (p-1)(\ell_p(n)+1). $$

For the lower bound use $p^{\ell_p(n)} \leq n$; for the upper bound compare with $H_{p^{\ell_p(n)+1}}$ and use the upper block bound. Both inequalities transfer. If K is infinite, so is $\ell_p(K)$, and hence ${}^*H_K$ is infinite in the ordered field.

At $K = p^L$, with L an infinite hyperinteger, the corresponding exact arithmetic measurements are

$$ {}^*H_{p^L} \geq 1 + \frac{p-1}{p}L, \qquad {}^*v_p({}^*H_{p^L}) = -L, \tag{40} $$

$$ \frac{p-1}{p} \leq \Delta_{p,L} < p-1, \qquad {}^*v_p(\Delta_{p,L}) = -(L+1), \tag{41} $$

where $\Delta_{p,L} = {}^*H_{p^{L+1}} - {}^*H_{p^L}$. Thus each next block is appreciable in the order—bounded above and below by positive standard constants—while its valuation has an infinite negative depth.

10.3 A finite correction at an arbitrary infinite row

Corollary 10.1 (Bounded arithmetic correction at hyperinteger indices)**.** *Suppose J_p is finite, with maximum digit length h as in Theorem 8.2. For every infinite hyperinteger K,*

$$ \boxed{{}^*v_p({}^*H_K) = -\ell_p(K) + d, \qquad d \in \{0, 1, \ldots, h\}.} \tag{42} $$

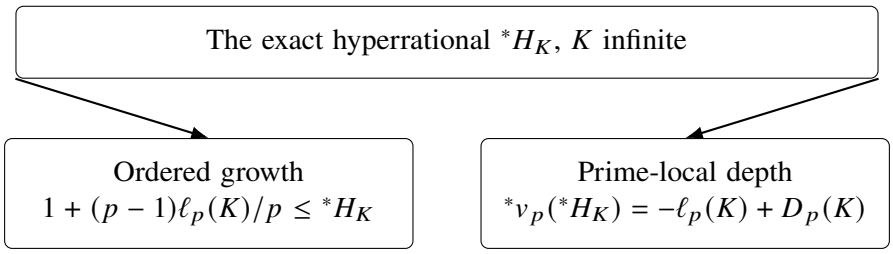

Figure 5: The shared scale does not identify two notions of size. Oresme's inequality proves ordered growth; the local state retains additional divisibility information.

The correction d is determined by the first exit of the leading base-p prefixes from J_p.

Proof. Transfer (33) and (32). Every infinite K exceeds the standard bound p^h. The transferred finite set $\{0, \dots, h\}$ has no additional nonstandard elements, so d is an ordinary integer. □

For $p = 2$, $d = 0$. For $p = 3$, Table 6 gives $d = 0, 1, 2, 3$ according to the first at most three ternary digits. Equivalently, a single nonzero residue

$$x = X_3(K) \bmod 81 \in \{1, \dots, 80\}$$

has $v_3(x) = d$, so

$${}^*v_3({}^*H_K) = -\ell_3(K) + v_3(x).$$

This is a finite local residue attached to an infinite scale counter. It specifies what arithmetic information accompanies divergence, not merely how large a real comparison function is.

For general p, the finiteness hypothesis must remain visible. Boyd's Proposition 3.3 [3] identifies it with $a_p(n) \to -\infty$; the unconditional ray statement (40) does not imply the same claim at every infinite index. Conversely, none of these valuation statements alone proves ordered divergence: p^{-k} has valuation $-k$ but tends to zero in the real order. Oresme supplies the separate order argument.

10.4 The minutes hand at an infinite row: finite carry without hyperfinite replay

The minutes-hand identity of Part I transfers particularly cleanly. Let $K \ge p$ be a hyperinteger and write

$$K = pM + r, \qquad M = \left\lfloor \frac{K}{p} \right\rfloor, \qquad 0 \le r < p. \tag{43}$$

Because r lies in the finite standard set $\{0, \dots, p-1\}$, it is an ordinary standard digit. If $k = \ell_p(K) = \ell_p(M) + 1$, transfer of (14) gives the exact internal identity

$$\boxed{X_p(K) = X_p(M) + p^k U_p(M, r).} \tag{44}$$

Thus the same two-hand address remains valid at a hyperinteger time: the seconds hand gives the last base-p digit, while the minutes hand points to the parent row in the same prime column.

Fix now a *standard* precision $R \ge 1$. If K is infinite, then $k \ge R$, and the correction in (44) is invisible modulo p^R:

$$X_p(K) \equiv X_p(M) \pmod{p^R}. \tag{45}$$

Repeatedly taking minutes readings removes suffix digits without changing this standard-width carry. Let

$$L = \ell_p(K) + 1$$

be the (infinite) base-p digit length and define the leading R-digit prefix internally by

$$\Pi_{p,R}(K) = \left\lfloor \frac{K}{p^{L-R}} \right\rfloor. \tag{46}$$

It satisfies

$$p^{R-1} \le \Pi_{p,R}(K) < p^R,$$

so it is bounded by a standard integer and is therefore itself standard. Iteration of (45), equivalently transfer of Theorem 7.5, yields

$$\boxed{X_p(K) \equiv X_p\big(\Pi_{p,R}(K)\big) \pmod{p^R}.} \tag{47}$$

Proposition 10.2 (Finite carry description of a standard-width hyperinteger shadow)**.** *For a fixed standard prime p, an infinite hyperinteger K, and a fixed standard width R, the normalized residue $X_p(K) \bmod p^R$ is completely determined by the ordinary finite prefix $\Pi_{p,R}(K)$. No hyperfinite prefix scan $1, 2, \dots, K$ and no stored hyperfinite ancestry are needed to describe this local shadow.*

This is stronger than saying that the hyperfinite forward algorithm transfers. The forward algorithm is an internal construction of length K; the minutes-hand theorem collapses every *standard-resolution observation* to a standard finite carry. At width R, all of the infinitely many suffix digits beyond the first R are provably irrelevant to the normalized residue.

If the carry

$$c_{p,R}(K) := X_p\big(\Pi_{p,R}(K)\big) \bmod p^R$$

is nonzero with $v_p(c_{p,R}(K)) = d < R$, then the exact hyperinteger-valued valuation is already certified:

$$\boxed{{}^*v_p({}^*H_K) = -\ell_p(K) + d.} \tag{48}$$

The infinite part is entirely the deterministic scale counter; the arithmetic correction is the ordinary valuation of a finite carry. If the carry is zero, standard precision R has not yet resolved the row and one inspects a longer leading prefix.

Under the finiteness hypothesis for J_p, Theorem 8.2 supplies one standard width R_0 that resolves every row. Corollary 10.1 can then be sharpened into the explicit carry formula

$${}^*v_p({}^*H_K) = -\ell_p(K) + v_p\big(\kappa_{p,R_0}(\Pi_{p,R_0}(K))\big), \tag{49}$$

where the second term belongs to the finite set $\{0, \dots, R_0 - 1\}$. The earlier first-exit description through the positive prefix tree and this carry description are equivalent: a zero carry is precisely a prefix in J_p by Proposition 7.7.

Without the finiteness hypothesis, every fixed standard width still has the finite representation (47). What can fail is the existence of one standard width that certifies the *exact* valuation on every hyperinteger path. An internal path may remain in the exceptional zero-carry tree through every standard depth, in which case exact resolution requires nonstandard width. This is the precise place where the Eswarathasan–Levine/Boyd finiteness problem enters the hyperrational picture.

10.5 What precision and completeness mean here

The finite-horizon algorithm and its correctness theorem also transfer to internal hyperfinite computations. This is a logical statement, not an assertion that an ordinary computer can execute infinitely many steps. For an arbitrary prime a sufficiently large transferred precision may itself be hyperinteger. A fixed *standard* precision valid at every row is available precisely under the uniform-precision condition of Theorem 8.2.

There is a second qualification. All standard-prime columns need not reconstruct a hyperrational: an infinite hyperprime has valuation zero at every standard prime, just as 1 does. A transferred full factorization requires the internal prime coordinates, including nonstandard primes. The fixed-standard-prime columns used here are therefore local observables, not a claim of complete hyperrational reconstruction.

Why the Euler–Mascheroni viewpoint is postponed. The classical comparison $H_n - \log n \to \gamma_{\mathrm{EM}}$ is analytically natural: it isolates the limiting discrepancy between a discrete sum and its integral model. It is nevertheless a different measurement from the signed factorization row of the exact rational H_n. Subtracting a logarithm is not by itself an irreversible loss of information—one may keep both terms—but replacing an exact rational sequence by its limiting real description forgets arithmetic data of the representatives. Part III studies that loss directly, using rational and hyperrational approximants whose real standard part is γ_{EM}.

Part III
Euler–Mascheroni through prime-column arithmetic

11 A hyperrational representative of the real constant

The Euler–Mascheroni constant is classically defined by

$$\gamma_{\mathrm{EM}} = \lim_{n\to\infty} (H_n - \log n). \tag{50}$$

This comparison is one of the most useful analytic descriptions of the harmonic series; see Lagarias for historical and modern perspectives [14]. The present question is narrower and deliberately arithmetic: can one approach the same real constant through exact rational or hyperrational objects while retaining their prime columns?

A particularly simple answer is obtained without inserting a logarithm into the approximating object.

Definition 11.1 (A rational harmonic approximant to γ_{EM})**.** For $n \geq 1$ define

$$G_n := 2H_n - H_{n^2} = H_n - \sum_{j=n+1}^{n^2} \frac{1}{j} \in \mathbb{Q}. \tag{51}$$

The construction of G_n is purely rational. The logarithm enters only when one identifies its real limit. Write

$$\varepsilon_n := H_n - \log n - \gamma_{\mathrm{EM}}.$$

The classical bounds

$$0 < \varepsilon_n < \frac{1}{2n} \tag{52}$$

and the monotonic decrease of ε_n are standard consequences of the sum–integral comparison or Euler–Maclaurin expansion [14, 13]. Since

$$G_n - \gamma_{\mathrm{EM}} = 2\varepsilon_n - \varepsilon_{n^2},$$

we obtain the following rational bracketing statement.

Proposition 11.2 (A logarithm-free approximating sequence)**.** *For every $n \geq 1$,*

$$0 < G_n - \gamma_{\mathrm{EM}} < \frac{1}{n}. \tag{53}$$

Consequently $G_n \to \gamma_{\mathrm{EM}}$.

Proof. Because $\varepsilon_n > \varepsilon_{n^2} > 0$, $2\varepsilon_n - \varepsilon_{n^2} > 0$. By (52), $2\varepsilon_n - \varepsilon_{n^2} < 2\varepsilon_n < 1/n$. □

The usual Euler–Maclaurin expansion gives more detail:

$$G_n = \gamma_{\mathrm{EM}} + \frac{1}{n} - \frac{2}{3n^2} + \frac{1}{10n^4} + O(n^{-6}). \tag{54}$$

This expansion is used below only to compare real rates of approximation; the prime-column statements will be proved directly from exact rational identities.

Let N now be an infinite positive hyperinteger. Transfer of (53) yields

$$G_N := 2^*H_N - {}^*H_{N^2} \in {}^*\mathbb{Q}, \qquad 0 < G_N - \gamma_{\mathrm{EM}} < \frac{1}{N}. \tag{55}$$

Hence G_N is finite as a hyperreal and

$$\boxed{\mathrm{st}_\infty(G_N) = \gamma_{\mathrm{EM}}.} \tag{56}$$

Here st_∞ denotes the real standard-part map on finite hyperreals. The phrase "real part" will mean this standard part; it is not the integer part of a number and not the real component of a complex number.

Equation (56) is the first separation of roles: G_N is an exact hyperrational carrying internal factorization data, whereas γ_{EM} is its real shadow after infinitesimal distinctions are identified.

Remark 11.3 (The error need not be hyperrational). It is important not to write a p-adic valuation of $G_N - \gamma_{\mathrm{EM}}$ without an additional hypothesis. In the usual superstructure enlargement, ${}^*\mathbb{Q} \cap \mathbb{R} = \mathbb{Q}$. Since $G_N \in {}^*\mathbb{Q}$,

$$G_N - \gamma_{\mathrm{EM}} \in {}^*\mathbb{Q} \quad \Longleftrightarrow \quad \gamma_{\mathrm{EM}} \in \mathbb{Q}.$$

The rationality of γ_{EM} is unknown. Thus the unconditional prime-column object is G_N itself, not its real infinitesimal error.

12 Exact prime depth along a prime-power ray

The power indices used throughout Parts I–II make the arithmetic of G_n unusually transparent.

Proposition 12.1 (Prime depth of the harmonic approximant)**.** *Let p be any prime and $k \geq 1$. Then*

$$\boxed{v_p(G_{p^k}) = -2k.} \tag{57}$$

Moreover,

$$\boxed{p^{2k}G_{p^k} \equiv -1 \pmod{p}.} \tag{58}$$

Proof. The unique denominator of maximal p-power in H_{p^k} is p^k, so $v_p(H_{p^k}) = -k$ and $p^k H_{p^k} \equiv 1 \pmod{p}$. Similarly, $v_p(H_{p^{2k}}) = -2k$ and $p^{2k}H_{p^{2k}} \equiv 1 \pmod{p}$. Now

$$v_p(2H_{p^k}) = v_p(2) - k > -2k = v_p(H_{p^{2k}}),$$

including the case $p = 2$. The ultrametric equality for unequal valuations gives $v_p(G_{p^k}) = -2k$. After multiplication by p^{2k}, the term $2p^{2k}H_{p^k}$ is divisible by p, whereas $p^{2k}H_{p^{2k}} \equiv 1 \pmod{p}$, proving (58). □

Remark 12.2 (The same proof in clock coordinates). The proposition can be read directly from the minutes/carry state. Since $C_{p,k} = X_p(p^k)$ and the prime-power address is a chain of zero seconds digits, (61) implies $C_{p,k} \equiv 1 \pmod{p}$ for every k. Then

$$p^{2k}G_{p^k} = 2p^k C_{p,k} - C_{p,2k} \equiv -1 \pmod{p}.$$

Thus the exact valuation $-2k$ is already a one-digit carry statement. The unique-denominator proof above is the most elementary proof; the clock proof is useful because the same mechanism extends immediately to arbitrary finite precision in Section 13.

By transfer, if L is an infinite positive hyperinteger and p remains a standard prime, then

$$\boxed{\mathrm{st}_\infty(G_{p^L}) = \gamma_{\mathrm{EM}}, \qquad {}^*v_p(G_{p^L}) = -2L.} \tag{59}$$

The same exact hyperrational is therefore finite and infinitely close to a fixed positive real number in the ordered field, while its p-column lies at an infinitely negative valuation depth.

This is not contradictory. If

$$U_{p,L} := p^{2L} G_{p^L},$$

then ${}^*v_p(U_{p,L}) = 0$, but $U_{p,L}$ is real-infinite, of the order p^{2L}. A p-adic unit need not have bounded ordinary size. The two notions of size are independent measurements.

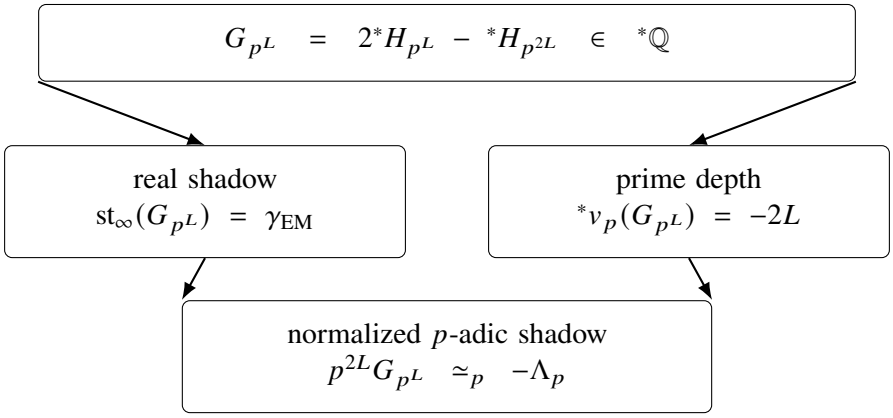

Figure 6: One exact hyperrational supports several projections. The real standard part, the signed valuation, and the normalized p-adic shadow answer different questions. None should be identified with another.

13 A normalized p-adic limit from the minutes hand

The valuation in Proposition 12.1 records only the first nonzero p-adic digit. The minutes recursion shows that the entire normalized state along a prime-power ray stabilizes one digit at a time.

Define

$$C_{p,k} := p^k H_{p^k} = X_p(p^k). \tag{60}$$

At the next power of p, the clock address is

$$p^{k+1} = p \cdot p^k + 0.$$

Thus (14) gives directly

$$\boxed{C_{p,k+1} = C_{p,k} + p^{k+1} U_p(p^k, 0).} \tag{61}$$

This is the minutes-hand form of the denominator splitting

$$H_{p^{k+1}} = \frac{1}{p} H_{p^k} + \sum_{\substack{1 \le j \le p^{k+1} \\ p \nmid j}} \frac{1}{j}.$$

Equivalently,

$$C_{p,k+1} - C_{p,k} = p^{k+1} \sum_{\substack{1 \le j \le p^{k+1} \\ p \nmid j}} \frac{1}{j}, \tag{62}$$

so

$$v_p(C_{p,k+1} - C_{p,k}) \ge k + 1. \tag{63}$$

The carry interpretation is immediate. At any fixed width R,

$$\boxed{C_{p,k} \equiv C_{p,R-1} \pmod{p^R} \qquad (k \ge R-1).} \tag{64}$$

Once the prime-power ray has reached digit length R, all later minutes-hand steps add corrections beyond the visible carry. In particular, the first R normalized digits are not merely Cauchy; they have literally frozen.

Definition 13.1 (Prime-local harmonic limit)**.** For each prime p put

$$\Lambda_p := \overset{p\text{-adic}}{\lim_{k\to\infty}} p^k H_{p^k}. \tag{65}$$

Then $\Lambda_p \in 1 + p\mathbb{Z}_p$.

Equation (64) gives the finite certificate

$$\boxed{\Lambda_p \equiv C_{p,R-1} \pmod{p^R}.} \tag{66}$$

Thus the width-R p-adic shadow of Λ_p is exactly the carry at the single finite prefix

$$10\cdots0 \quad (R \text{ base-}p \text{ digits}),$$

that is, at the row p^{R-1}. No long ray needs to be generated in order to obtain those R digits.

Now define

$$U_{p,k} := p^{2k} G_{p^k}.$$

The exact identity

$$U_{p,k} = 2p^k C_{p,k} - C_{p,2k} \tag{67}$$

can also be read entirely in carry language. Fix R. Once $k \ge R$, the first term vanishes modulo p^R, while the second has already frozen by (64). Hence

$$\boxed{U_{p,k} \equiv -C_{p,R-1} \pmod{p^R} \qquad (k \ge R).} \tag{68}$$

This proves the normalized limit without any sequential passage through the intermediate rows.

Proposition 13.2 (Normalized prime shadow of the γ_{EM} approximant)**.** *For each fixed prime p,*

$$\boxed{p^{2k} G_{p^k} \longrightarrow -\Lambda_p \quad \text{in } \mathbb{Q}_p.} \tag{69}$$

At precision p^R, the limiting residue is already the finite carry $-C_{p,R-1} \bmod p^R$.

In a sufficiently saturated nonstandard enlargement of $\mathbb{Q}_p$, the transferred sequence gives, for infinite L,

$$\boxed{\mathrm{st}_\infty(G_{p^L}) = \gamma_{\mathrm{EM}}, \qquad \mathrm{st}_p(p^{2L} G_{p^L}) = -\Lambda_p.} \tag{70}$$

Here st_p means the p-adic standard part of this near-standard normalized hyperrational. The minutes-hand statement behind the second equality is especially concrete: for every standard R,

$$p^{2L} G_{p^L} \equiv -C_{p,R-1} \pmod{p^R}.$$

Thus its standard p-adic shadow is assembled from a compatible family of ordinary finite carries.

Equation (70) does *not* identify γ_{EM} with $-\Lambda_p$: the two standard parts are taken in different topologies and after different normalizations. There are independently defined p-adic Euler constants in the literature, notably those attached to the p-adic digamma functions of Morita and Diamond; see Murty–Saradha [15]. No identification of the elementary limit Λ_p above with one of those constants is asserted here. Determining whether a useful relation exists is a separate question.

14 What the real standard part forgets

The standard-part map deliberately identifies all finite hyperreals that differ by a real infinitesimal. Prime-column data need not be constant on such equivalence classes.

14.1 Same real part, radically different valuation depth

Let $N = p^L$ with L infinite and set

$$\widetilde{G}_N := G_N + p^{-L^2}.$$

For every fixed standard $m \geq 1$,

$$N^m(\widetilde{G}_N - G_N) = p^{mL - L^2} \approx 0.$$

Thus G_N and $\widetilde{G}_N$ have the same real standard part and agree beyond every fixed algebraic order in N^{-1}. Nevertheless,

$$^*v_p(G_N) = -2L, \qquad ^*v_p(\widetilde{G}_N) = -L^2,$$

because the perturbation has strictly smaller valuation. Hence neither the real standard part nor any fixed-order real asymptotic expansion determines the prime depth of its hyperrational representative.

The converse separation also occurs: improving the real approximation need not alter the valuation. For $p \geq 5$ and $n = p^k$, define

$$\begin{aligned} G_n^{(0)} &= G_n, \\ G_n^{(1)} &= G_n - \frac{1}{n}, \\ G_n^{(2)} &= G_n - \frac{1}{n} + \frac{2}{3n^2}. \end{aligned}$$

Then (54) gives successively errors of orders n^{-1}, n^{-2} and n^{-4}, but

$$v_p(G_n^{(0)}) = v_p(G_n^{(1)}) = v_p(G_n^{(2)}) = -2k. \tag{71}$$

Indeed, after multiplication by n^2, the three leading residues modulo p are respectively -1, -1, and $-1 + 2/3 = -1/3$, all nonzero.

Table 7: Real accuracy and p-adic depth are different observables ($p \geq 5$, $n = p^k$).

approximant	leading real error	v_p
G_n	n^{-1}	$-2k$
$G_n - 1/n$	$-2/(3n^2)$	$-2k$
$G_n - 1/n + 2/(3n^2)$	$1/(10n^4)$	$-2k$

14.2 All standard prime columns still need not determine a hyperrational

For ordinary positive rationals, the complete standard-prime valuation row determines the number. At a hyperinteger index one must distinguish that statement from its transferred internal version: a hyperrational may also contain infinite hyperprime factors.

Let T be an infinite positive hyperinteger, put $M = T!$, and define

$$Y := \frac{2M+1}{M+1} = 2 - \frac{1}{M+1}. \tag{72}$$

Then $\mathrm{st}_\infty(Y) = 2$. For every standard prime p and every standard r, the hyperinteger M is divisible by p^r, whence $2M + 1 \equiv M + 1 \equiv 1 \pmod{p^r}$. Therefore

$$^*v_p(Y) = 0, \qquad \mathrm{st}_p(Y) = 1 \quad \text{for every standard } p.$$

Thus the collection of all standard-prime shadows can fail to determine the real standard part of a general hyperrational. The missing internal factorization data live at nonstandard primes. This does not weaken the use of a chosen standard-prime column as a local observable; it clarifies what that observable does and does not encode.

15 Denominators: where the prime columns meet real approximation

A valuation column becomes directly relevant to the arithmetic quality of a rational approximation through the reduced denominator. Write

$$G_n = \frac{A_n}{B_n}, \qquad (A_n, B_n) = 1, \qquad B_n > 0. \tag{73}$$

Then

$$\boxed{B_n = \prod_q q^{\max(0,-v_q(G_n))}.} \tag{74}$$

Thus the negative columns, taken together, reconstruct the denominator.

This connects the local arithmetic with a basic irrationality test.

Lemma 15.1 (Elementary denominator-scaled criterion)**.** *Let $r_n = A_n/B_n$ be reduced rationals, $B_n > 0$, with $r_n \to \alpha$ and $r_n \neq \alpha$. If*

$$B_n|r_n - \alpha| \longrightarrow 0,$$

then α is irrational.

Proof. If $\alpha = a/b$ in lowest terms, then

$$\left|\frac{A_n}{B_n} - \frac{a}{b}\right| = \frac{|bA_n - aB_n|}{bB_n} \geq \frac{1}{bB_n},$$

so $B_n|r_n - \alpha| \geq 1/b$. □

The sequence G_n does not satisfy this criterion. Let b_m denote the reduced denominator of H_m. Boyd records the classical growth

$$b_m = \exp(m + o(m)), \tag{75}$$

based on the prime number theorem [3, p. 293]. Let d_n be the reduced denominator of $2H_n$. Then $d_n \leq b_n$. Since $H_{n^2} = 2H_n - G_n$, denominator divisibility gives

$$b_{n^2} \leq d_n B_n,$$

whereas a denominator of a difference divides the product, so

$$B_n \leq b_{n^2} d_n.$$

Consequently

$$\boxed{B_n = \exp(n^2 + o(n^2)).} \tag{76}$$

Together with $G_n - \gamma_{\mathrm{EM}} \sim 1/n$, this yields

$$\boxed{B_n|G_n - \gamma_{\mathrm{EM}}| \longrightarrow \infty.} \tag{77}$$

Thus G_n is an excellent conceptual bridge between real and prime-local arithmetic, but it is not an efficient irrationality approximant in the sense of Lemma 15.1.

The column picture shows where much of this denominator comes from. For $n \geq 3$ and any prime q satisfying

$$\frac{n^2}{2} < q \leq n^2,$$

one has $q > n$, so $2H_n$ is q-integral. In H_{n^2} the only summand whose denominator is divisible by q is $1/q$. Hence

$$\boxed{v_q(G_n) = -1.} \tag{78}$$

At a hyperinteger index, the analogous portion of the denominator lies at infinite hyperprimes near N^2, again explaining why fixed standard columns are only part of the complete internal factorization.

The minutes-hand compression clarifies the contrast. For any *fixed* standard prime p and fixed local width, the state of G_N can collapse to a small standard carry. The complete reduced denominator, however, ranges over all prime columns, including the large primes near N^2 and, at an infinite index, nonstandard primes. Local state compression therefore makes a chosen column inexpensive without pretending that it reconstructs the global factorization.

16 Local precision, exact reconstruction, and direct carry access

The arithmetic of G_n can be computed from one prime column without forming the large reduced fraction. Fix p and write

$$a = \ell_p(n), \qquad b = \ell_p(n^2), \qquad X_p(m) = p^{\ell_p(m)} H_m.$$

Then

$$\boxed{p^b G_n = 2p^{b-a} X_p(n) - X_p(n^2).} \tag{79}$$

There are two useful ways to exploit this identity, and their contrast mirrors Part I.

Forward implementation: save and continue

A Boyd-like +1 clock may save its state at row n, continue successively to n^2, and combine the two normalized residues according to (79). This is conceptually direct and is efficient when the whole dense interval is being generated anyway. If the resulting residue is zero at the chosen width, the forward/certification implementation increases precision according to the policies discussed in Part I.

Minutes-hand implementation: query the two rows directly

For random access there is no reason to traverse the interval $n+1, \dots, n^2$. The minutes-hand algorithm evaluates $X_p(n) \bmod p^R$ and $X_p(n^2) \bmod p^R$ independently from their leading-digit ancestry. The carry theorem compresses this still further.

Since $n \geq p^a$, one has $n^2 \geq p^{2a}$ and therefore

$$b - a \geq a. \tag{80}$$

Consequently, whenever $a \geq R$, the first term in (79) is invisible modulo p^R:

$$p^b G_n \equiv -X_p(n^2) \pmod{p^R}. \tag{81}$$

If n^2 has at least R base-p digits, Theorem 7.5 then gives the stronger random-access formula

$$\boxed{p^b G_n \equiv -X_p\big(\pi_{p,R}(n^2)\big) \pmod{p^R}, \qquad \ell_p(n) \geq R.} \tag{82}$$

Thus, at a fixed normalized width, sufficiently large rows no longer require *either* harmonic state H_n or H_{n^2} to be generated sequentially. The local state of the γ_{EM} approximant is the negative carry addressed by the first R base-p digits of n^2.

Proposition 16.1 (Carry compression for the γ_{EM} approximant)**.** *Fix a prime p and a width R. For every n with $\ell_p(n) \geq R$, the normalized residue $p^{\ell_p(n^2)} G_n \bmod p^R$ depends only on the first R base-p digits of n^2. If this carry is nonzero, it certifies $v_p(G_n)$ exactly after subtracting $\ell_p(n^2)$.*

This is the Part III analogue of the frozen cylinders of Part I. Different indices whose squares have the same R-digit leading prefix have the same normalized G-carry at width R. Repeated carries can be merged just as in the harmonic column itself. A zero carry is the only case that forces an increase of precision.

For the prime-power ray $n = p^k$, the formula becomes especially simple. If $k \geq R$, then the first R digits of $n^2 = p^{2k}$ are $10\cdots0$, so

$$p^{2k} G_{p^k} \equiv -X_p(p^{R-1}) = -C_{p,R-1} \pmod{p^R}. \tag{83}$$

This is exactly the frozen-carry formula (68) and makes the normalized limit $-\Lambda_p$ visible without following the ray to infinity.

At an infinite hyperinteger N, the simplification is stronger still. For every standard R one automatically has $\ell_p(N) \geq R$, and therefore

$$\boxed{p^{\ell_p(N^2)} G_N \equiv -X_p\big(\Pi_{p,R}(N^2)\big) \pmod{p^R}.} \tag{84}$$

The standard-width p-adic shadow of the hyperrational approximant is thus an ordinary finite carry read from the leading digits of the hyperinteger square. No hyperfinite replay from 1 to N^2 is part of this description.

Complexity interpretation. This is a structural row-complexity statement, not a universal bit-complexity claim. A dense forward scan to n^2 performs $O(n^2)$ row updates but stores only a rolling state. A minutes/carry query follows at most the requested number of leading digits once the carry has frozen; evaluating the local corrections still incurs modular-arithmetic and large-integer costs. The gain is that the number of harmonic *rows* touched is controlled by precision rather than by the size of the index. This is precisely the setting in which repeated states, invariances, and hyperinteger standard-width shadows become transparent.

Valuation certification versus rational reconstruction

Whichever local implementation is used, exact valuation certification must be distinguished from reconstructing the entire rational. Suppose a p-integral rational $x = a/b$ is known modulo p^R, with

$$|a| \leq U, \qquad 0 < b \leq V, \qquad p \nmid b.$$

If

$$p^R > 2UV, \tag{85}$$

then there is at most one such rational with the given residue. For if a/b and c/d are two candidates, then $p^R \mid (ad - bc)$ while $|ad - bc| < p^R$, forcing equality.

Thus three computational objectives are genuinely different:

$$\boxed{\text{certify a valuation} \;\neq\; \text{reconstruct the exact rational} \;\neq\; \text{certify its real error.}}$$

For G_{p^k}, one normalized digit already proves $v_p(G_{p^k}) = -2k$ by (58); reconstructing the full fraction can require precision commensurate with its exponentially large height. Conversely, the real interval

$$G_n - \frac{1}{n} < \gamma_{\mathrm{EM}} < G_n$$

requires no factorization at all. The minutes-hand construction sharpens the separation: a short local carry may settle the prime-column question even when both exact rational reconstruction and a dense forward scan would be very large computations.

17 Interpretation and open arithmetic questions

The construction above gives a precise answer to the question "what do the hyperrational prime columns reveal about the real approximation to γ_{EM}?"

1. They reveal exact denominator depth and normalized residue data of the chosen hyperrational representative. Along p^L this includes the exact depth $-2L$ and the p-adic standard shadow $-\Lambda_p$.
2. This information is not an invariant of the real standard part. Infinitesimally close hyperrationals with the same real part can have wildly different valuations, while substantial improvements in real accuracy can leave a valuation unchanged.
3. The collection of negative columns controls the reduced denominator, which is directly relevant to Diophantine approximation. For the simple sequence G_n, the denominator grows too fast relative to the real error to yield the elementary irrationality criterion.
4. The full internal row at a hyperinteger index involves nonstandard primes as well as standard primes. Fixed standard columns are therefore local arithmetic observables, not a complete description of an arbitrary hyperrational.

This viewpoint does not assign a p-adic valuation to the real number γ_{EM}, whose rationality is itself unknown [14]. It does something more concrete: it keeps an exact arithmetic representative, its real standard part, and its prime-local shadows available simultaneously. The resulting research question is not whether one of these projections replaces the others, but whether a more carefully designed rational harmonic approximant can combine unusually small real error with unusually favorable prime-column cancellation. That is the point at which the clockwork may interact with the classical Diophantine-approximation problem for Euler's constant.

18 Conclusion

The article has followed the harmonic numbers through three layers. Part I remains entirely in $\mathbb{Q}$. A complete row records the signed prime factorization of one H_n, while a fixed prime column can be advanced by a certified local +1 clock. The visible valuation is not sufficient state; a normalized residue supplies the cancellation information. Boyd's direct method already has fixed-prime successive addition and exact valuation readout, and his positive-prefix tree already supplies the sparse exceptional genealogy. The normalization used here reorganizes that information into a scale counter and residue state. A predecessor-only implementation may recover missing precision by replay; a memoized Prime Clockwork minutes address refines only the ancestry $n, \lfloor n/p \rfloor, \ldots$; and the carry theorem compresses this further. At width R, only the first R base-p digits can affect the normalized carry. Nonzero carries freeze whole suffix cylinders, zero carries are exactly positive prefixes, and the valuation falls by one whenever an additional suffix digit increases the scale. Thus the full signed column is a small arithmetic carry mechanism coupled to an unbounded scale counter. Every finite row is computable exactly without full-prefix replay; one fixed carry width for the entire column is equivalent to finiteness of J_p.

Part II transfers these exact identities to hyperinteger rows. Oresme's prime-power blocks prove ordered divergence without leaving the hyperrationals, and the prime columns retain a

second, arithmetic scale. The minutes-hand identity itself transfers: at every fixed standard width, an infinite row collapses to an ordinary finite leading-prefix carry together with its infinite scale counter. Under the finite-J_p hypothesis, one standard carry width resolves every row and the infinite negative scale differs from $-\ell_p(K)$ by an exact bounded correction. Without that hypothesis, the possible obstruction is precisely a path that remains in the zero-carry tree through every standard depth.

Part III then deliberately reintroduces the Euler–Mascheroni constant, but only after retaining an exact hyperrational representative. The sequence $G_n = 2H_n - H_{n^2}$ has real standard part γ_{EM} and, along prime powers, exact valuation $-2k$. Its normalized units converge to a separate p-adic limit $-\Lambda_p$. In clock coordinates the first R digits of Λ_p are already frozen at the finite carry $X_p(p^{R-1}) \bmod p^R$, and for sufficiently large n the width-R normalized state of G_n is read directly from the leading-prefix carry of n^2. Thus no sequential scan to n^2 is intrinsic to the local arithmetic question. The real and p-adic shadows remain simultaneously available but logically distinct. Their comparison shows both the strength and the limitation of valuation columns: they expose arithmetic structure erased by real standard part, yet they do not by themselves determine the real limit. Denominators provide the bridge to Diophantine approximation, and for this particular approximant they also show why convergence alone is far from an irrationality proof.

A A minimal implementation of the certified scan

The following code implements the full precision-recovery procedure with no external packages. Inputs satisfy the preconditions stated in Algorithm 1: p is prime and $N \geq 1$. The separately supplied file `harmonic_column_recursion_dynamic.py` also includes a persistent `HarmonicPrimeClock`, input validation, a command-line interface, the ternary prefix rule, and independent tests. The clock implements Algorithm 2; a tick returns a certified snapshot rather than merely a residue.

```
def column_pass(p, N, R):
    M = p ** R
    ell, next_power, x = 0, p, 1
    out = [0]
    for j in range(2, N + 1):
        if j == next_power:
            ell += 1
            next_power *= p
            x = p * x % M
        e, m = 0, j
        while m % p == 0:
            m //= p
            e += 1
        x = (x + pow(p, ell - e, M)
               * pow(m, -1, M)) % M
        if x == 0:
            out.append(None)  # unresolved, not infinity
        else:
            t, y = 0, x
            while y % p == 0:
                y //= p
                t += 1
            out.append(t - ell)
    return out

def harmonic_column(p, N):
    R = 1
    while True:
        out = column_pass(p, N, R)
        if all(a is not None for a in out):
            return out, R
        R *= 2  # replay; do not pad old residues
```

For example, `harmonic_column(3, 15)` returns the column

$$(0, 1, -1, -1, -1, 0, 1, 0, -2, -2, -2, -2, -2, -2, -2)$$

and final precision $R = 4$. The minimal precision for this prefix is three digits; the doubling schedule deliberately uses a power of two.

B Operating the persistent clock

The companion implementation supports a saved-state computation with no predetermined final index. Its immutable snapshots contain the prime, row, precision, scale, residue, exact valuation, and cumulative number of replays. The current state starts at H_1; each call to `tick()` advances the output index by one.

```
from harmonic_column_recursion_dynamic import HarmonicPrimeClock

clock = HarmonicPrimeClock(3, initial_digits=3)
row = clock.current
assert (row.n, row.valuation) == (1, 0)

for _ in range(21):
    row = clock.tick()

print(row.n, row.digits, row.scale,
      row.residue, row.valuation, row.replays)
# 22 6 2 513 1 1
```

Rows 1 through 21 are resolved at precision three. At row 22 the clock encounters a zero residue modulo 27, doubles the width to six, and replays only the 3-column through row 22. The new residue is $513 = 3^3 \cdot 19$, so the exact output is $3 - 2 = 1$. Already emitted outputs are unchanged. Subsequent calls continue from the new saved state; they do not replay merely because another row is requested.

The program also verifies that different starting precisions give the same outputs and that each stored residue agrees with an independent exact-rational oracle. These tests exercise both ordinary ticks and precision-recovery ticks. They check an implementation of the proved recurrence, rather than supplying evidence for the general finiteness conjecture.

C Operating the minutes-hand clock without replay

The companion file `harmonic_minute_hand.py` implements Algorithm 3. Each call advances the common clock time by one if desired, but precision refinement uses the minutes address rather than replaying the prefix. The internal cache stores the highest normalized precision reached for each addressed row.

```
from harmonic_minute_hand import MinuteHandHarmonicClock

clock = MinuteHandHarmonicClock(3, initial_digits=3)
for _ in range(22):
    n, valuation, digits = clock.tick()

print(n, valuation, digits)
# 22 1 4
```

At row 22, three normalized digits are insufficient. The implementation requests four digits and refines the minutes chain $22 \to 7 \to 2$; it does not recompute rows 1 through 21.

The same file contains the finite correction evaluator (16). Its self-test compares both the exact valuation and the cached normalized residue with an independent rational computation for 13,000 pairs (p, n).

The replay and minutes-hand programs are both retained because they make two different points. The replay version is the shortest proof-of-concept for a certified sequential scan. The minutes-hand version shows how missing precision can be generated recursively through the minutes address, but its cache deliberately preserves ancestor states so that the time–space tradeoff is visible.

D Operating the carry-compressed clock

The companion file `harmonic_carry_clock.py` implements Algorithm 4. It has no history cache. For a requested width R it reads only the first $\min\{R, L_p(n)\}$ base-p digits, carries a prefix $q < p^R$ and a residue $c < p^R$, and discards all later suffix digits because Theorem 7.5 proves that they cannot affect the state modulo p^R.

```
from harmonic_carry_clock import valuation_by_carry

v, R, prefix = valuation_by_carry(3, 10**100)
print(v, R, prefix)
# The routine never scans H_1,...,H_(10**100).
# It increases R only while the leading carry is zero.
```

For an ordinary finite n the routine begins with $R = 1$. A nonzero carry is an exact certificate and the computation stops. A zero carry means that the current leading prefix belongs to J_p, so one more leading digit/precision level is examined. If n itself is a positive exceptional index, R may exceed the digit length; in that case all index digits have already been read and the extra steps request only more p-adic precision.

The implementation was checked against exact `Fraction` arithmetic for 2,400 valuation queries (eight primes through 19, $1 \le n \le 300$) and by 8,292 additional carry-freeze comparisons. It is intentionally a reference implementation: the elementary correction evaluator is easy to inspect, while Boyd's convergent series or other specialized arithmetic could replace it in a performance-oriented version.

The three programs therefore form a useful hierarchy for later extraction into a separate algorithmic paper: the Boyd-like/replay scan exhibits the minimal rolling state and the cost of lost precision; the memoized minutes clock exhibits recursive addressability; and the carry reader exhibits the finite prefix quotient, repeated-state compression, and suffix invariance.

References

[1] J. Babb, "Mathematical concepts and proofs from Nicole Oresme: Using the history of calculus to teach mathematics," *Science & Education* **14** (2005), 443–456. doi:10.1007/s11191-004-7937-y.

[2] J. P. Bell, "p-adic valuations and k-regular sequences," *Discrete Mathematics* **307** (2007), 3070–3075. doi:10.1016/j.disc.2007.03.009.

[3] D. W. Boyd, "A p-adic study of the partial sums of the harmonic series," *Experimental Mathematics* **3** (1994), no. 4, 287–302. doi:10.1080/10586458.1994.10504298.

[4] L. Carofiglio, G. Cherubini and A. Gambini, "On Eswarathasan–Levine and Boyd's conjectures for harmonic numbers," *Bulletin of the Australian Mathematical Society* **113** (2026), 15–25. doi:10.1017/S0004972725100154.

[5] A. Eswarathasan and E. Levine, "p-integral harmonic sums," *Discrete Mathematics* **91** (1991), no. 3, 249–257. doi:10.1016/0012-365X(90)90234-9.

[6] M. T. M. Emmerich, "The Prime Clockwork: A Dynamic Representation of Modular and Multiplicative Arithmetic," *arXiv preprint* arXiv:2609.03896 (2026). https://arxiv.org/abs/2609.03896.

[7] K. Kamano, "On 3-adic valuations of generalized harmonic numbers," *Integers* **11** (2011), A69, 8 pp. https://math.colgate.edu/~integers/l69/l69.pdf.

[8] H. J. Keisler, *Elementary Calculus: An Infinitesimal Approach*, 3rd ed., Dover Publications, Mineola, NY, 2012.

[9] T. Lengyel, "On the 2-adic valuation of differences of harmonic numbers," *Integers* **24** (2024), A27, 10 pp. https://math.colgate.edu/~integers/y27/y27.pdf.

[10] A. Robinson, *Non-standard Analysis*, North-Holland, Amsterdam, 1966.

[11] M. D. Rogers, "An extension of Boyd's p-adic algorithm for the harmonic series," *arXiv preprint* arXiv:0708.2439 (2007). https://arxiv.org/abs/0708.2439.

[12] C. Sanna, "On the p-adic valuation of harmonic numbers," *Journal of Number Theory* **166** (2016), 41–46. doi:10.1016/j.jnt.2016.02.020.

[13] NIST Digital Library of Mathematical Functions, Chapter 5, "Gamma Function," especially §5.11 (asymptotic expansions), F. W. J. Olver et al. (eds.). https://dlmf.nist.gov/5.11.

[14] J. C. Lagarias, "Euler's constant: Euler's work and modern developments," *Bulletin of the American Mathematical Society* **50** (2013), no. 4, 527–628. doi:10.1090/S0273-0979-2013-01423-X.

[15] M. Ram Murty and N. Saradha, "Transcendental values of the p-adic digamma function," *Acta Arithmetica* **133** (2008), no. 4, 349–362. doi:10.4064/aa133-4-4.

[16] D. J. Struik (ed.), *A Source Book in Mathematics, 1200–1800*, Harvard University Press, Cambridge, MA, 1969; reprinted by Princeton University Press.