hexagon:2611.01089v1

Rational Numerals and Carry Geometry in Iteris

Patricio M. Vera
patricio@simbionte.org

25 November 2026

Abstract

Iteris is a positional notation in base 10^4. It writes each four-digit decimal block as a Cistercian-inspired glyph and arranges the glyphs in a horizontal row. We specify a canonical representation of every rational number and determine the lengths of its fractional prefix and repeating block cycle. The graph of legal single-stroke edits in a quadrant is an isometric subgraph of a five-dimensional binary cube. The Cartesian product of these quadrant graphs makes node distances additive. We derive the exact distribution of decimal carry costs on words of m decimal digits, including leading zeros as needed to fill the width. Stroke parity rules out a one-stroke ordering of all 10^m words. A construction attains the minimum total edit cost $10^m + 2^m - 2$ for an open order. For the four positions of one node the open minimum is 10,014 and numerical order costs 19,986.

1 Introduction

Cistercian numerals combine several decimal digits around a common stem. The historical forms varied in their orientation, stroke shapes, and placement of decimal positions. Documented uses include foliation, indexes, and dates [1, 2]. Historical sources also describe extensions beyond 9,999, and the modern `xistercian` package already renders Cistercian numerals in LaTeX [2, 3]. Iteris adopts one vertical convention [4] and places its glyphs in a positional sequence extending in both directions from the radix point.

We first describe how place value, the four quadrants, and the horizontal row work together. We then ask how to write a rational whose decimal repeat starts inside a four-digit group. The same set of glyphs defines a graph by single-stroke edits. We use it to measure the cost of an increment and to find the least costly order in which to visit all glyphs of a given width.

2 The Iteris notation

2.1 Base and place value

Iteris groups decimal digits in fours. Each group is one digit of the positional system in base

$$B = 10^4 = 10{,}000.$$

We call the glyph for a base-B digit a *node*. Write $D = \{0, \dots, B-1\}$. A *well-formed Iteris numeral* is a tuple $w = (s, \mathbf{a}, \mathbf{b}, \mathbf{e})$, where

$$s \in \{-1, +1\}, \qquad \mathbf{a} = (a_k, \dots, a_0) \in D^+, \qquad \mathbf{b} = (b_1, \dots, b_t) \in D^*, \qquad \mathbf{e} = (e_1, \dots, e_\ell) \in D^*.$$

Here D^+ consists of nonempty finite words, and D^* also includes the empty word. The words store, in reading order, the integer blocks, the fractional prefix, and the repeating blocks. A minus sign is drawn exactly when $s = -1$. The radix point appears exactly when $t + \ell > 0$, followed by **b** and then **e**. An overbar appears exactly when $\ell > 0$ and spans **e**. A numeral with empty **e** is terminating and has decoded value

$$x = s \left(\sum_{i=0}^{k} a_i B^i + \sum_{j=1}^{t} b_j B^{-j} \right), \qquad a_i, b_j \in \{0, \ldots, 9999\}. \tag{1}$$

The radix point marks the boundary. The last integer node has weight B^0, and the first fractional node has weight B^{-1}. To convert a terminating decimal, group its integer and fractional digits in fours starting at the radix point and moving outward. Pad the outer ends of incomplete groups with zeros. For example, the blocks `0001 0000 0005` represent $1 \cdot B^2 + 0 \cdot B + 5 = 100{,}000{,}005$, and the fractional block `0700` contributes $700/B = 0.07$.

2.2 The node and its four quadrants

A node's four quadrants encode, in order of increasing place value, units (upper right), tens (upper left), hundreds (lower right), and thousands (lower left). A stem is always drawn. Each quadrant contains a subset of five available segments. In the units quadrant, p_1 and p_2 are the upper and lower horizontal segments, p_3 joins the upper stem point to the lower outer point, p_4 joins the lower stem point to the upper outer point, and p_6 joins the outer points. Reflections in the stem and in the horizontal midline generate the other quadrants. Table 1 gives the common stroke table. The five primitive labels follow the digits formed by single segments. Digit 5 combines p_1 and p_4. An all-zero node consists of the stem alone.

Digit	Stroke set	Digit	Stroke set
0	$\varnothing$	5	$\{p_1, p_4\}$
1	$\{p_1\}$	6	$\{p_6\}$
2	$\{p_2\}$	7	$\{p_1, p_6\}$
3	$\{p_3\}$	8	$\{p_2, p_6\}$
4	$\{p_4\}$	9	$\{p_1, p_2, p_6\}$

Table 1: The stroke set for each digit in the adopted convention. The common stem does not enter the stroke set.

2.3 Horizontal reading order

Each stem stays vertical within its node. The nodes occupy separate slots on a single line, read left to right. The row begins with a minus sign if needed, followed by the integer nodes from most to least significant. A radix point introduces the fractional nodes, and a continuous overbar marks exactly the nodes that repeat. This layout lets a multi-node numeral occupy the same position in a formula as a conventional numeral. Arithmetic operators, grouping symbols, and exponents retain their usual roles. For example, 2^{10000} can be written $\restriction^{\restriction\restriction}$, with the entire two-node exponent reduced to superscript size.

The reference layout sets the quadrant width to $W = H/5$, where H is the stem height. Adjacent node centers are $3W$ apart. The radix point occupies a separate slot of width W. The bar extends from one quadrant width before the first repeating stem to one quadrant width after the last.

Value	Blocks in reading order	Iteris numeral
9,999	9999	
10,000	0001 0000	
100,000,005.07	0001 0000 0005 . 0700	

Figure 1: Crossing 9999 requires a second, more significant node. A zero inside a numeral occupies a full slot and retains its stem. In the last row the radix point separates three integer nodes from one fractional node.

Figure 1 illustrates the node boundaries and how a bare stem carries place value. The full horizontal geometry is specified in Iteris Core [11].

2.4 Recurring fractions

Write $A = \sum_{i=0}^{k} a_i B^i$ for the value of the integer word. If $\ell > 0$, put $C = \sum_{i=1}^{\ell} e_i B^{\ell-i}$. The decoded value is

$$V(w) = s\left(A + \sum_{j=1}^{t} b_j B^{-j} + \frac{C}{B^t(B^\ell - 1)}\right). \tag{2}$$

When $\ell = 0$, omit the last term, giving (1). The period length ℓ counts every block under the bar, including initial zeros. Well-formed numerals can decode to the same value. For example, a zero integer block followed by a repeating 9999 block decodes to one.

3 The canonical rational expansion

Block division specializes the classical theory of rational radix expansions to $B = 10^4$ [7, Chapter IX]. The prefix and period formulas below are the corresponding radix-B formulas. The Iteris specification determines how to display this expansion as a canonical sequence of nodes, with the overbar aligned to node boundaries.

Given $x \in \mathbb{Q}$, write $|x| = p/q$ in lowest terms with $p \geq 0$, $q > 0$. The block division rule $N(x)$ sets $s = -1$ exactly when $x < 0$. It writes $A = \lfloor p/q \rfloor$ in base B with no leading zero unless the integer word is (0). Starting at $r_0 = p \bmod q$, emit b_{j+1} at step $j + 1$ by Euclidean division

$$Br_j = qb_{j+1} + r_{j+1}, \qquad 0 \leq r_{j+1} < q.$$

Stop at the first zero remainder or first repeated nonzero remainder. At a zero remainder the fractional word consists of all blocks emitted so far and the cycle is empty. When $r_{t+\ell}$ is the first repeated remainder and r_t is its earlier occurrence, the first t emitted blocks form the prefix and the next ℓ blocks form the cycle. Each remainder lies in $\{0, \ldots, q-1\}$, so the rule stops after finitely many steps.

For the division rule, extending the output by zero blocks after termination or by repetition in the recurring case gives an infinite word $d_1 d_2 \cdots$. With the corresponding continuation of remainders, Euclidean division gives, for every $j \geq 0$,

$$\frac{p}{q} = A + \sum_{i=1}^{j} d_i B^{-i} + \frac{r_j}{qB^j}, \qquad 0 \leq r_j < q.$$

The bound on r_j proves $V(N(x)) = x$ by taking the limit, using the geometric series in (2). The same argument beginning at step j gives the value of the remaining fractional tail as r_j/q at its own scale.

A well-formed numeral w is *canonical* when it satisfies the following conditions on its sign and block words.

1. The integer sequence begins with a nonzero block or consists of one block. The sign is positive when every block in the three words is zero.
2. With no bar, the fractional prefix is empty or ends in a nonzero block.
3. With a bar, the cycle is primitive, meaning it is not a repetition of a shorter word.
4. With a bar, the cycle is neither the singleton `0000` nor the singleton `9999`.
5. With a bar and a nonempty prefix, $b_t \neq e_\ell$.

Proposition 1 (Canonical form and block division)**.** *A well-formed numeral w is canonical if and only if $w = N(V(w))$.*

Proof. First suppose $w = N(V(w))$. The rule gives condition (1). If it terminates after emitting a block, that block is nonzero because at the final step $Br_{t-1} = qb_t$ and $r_{t-1} > 0$. This gives condition (2). If it repeats, the value of the fractional tail after each position j is r_j/q. A cycle that is a proper power of a shorter word of length $h < \ell$ would give equal tails after t and $t + h$, hence an earlier repeated remainder. A singleton cycle of 0000 or 9999 would give $r_t/q = 0$ or 1, respectively, so both are excluded. If $t > 0$ and $b_t = e_\ell$, the cycle could start one block earlier with a right rotation. Then the equal tails after $t - 1$ and $t + \ell - 1$ would also give an earlier repeat. This proves conditions (3)–(5).

Conversely, suppose conditions (1)–(5) hold. Extend the fractional blocks by zeros or by indefinite repetition of the cycle. Conditions (3) and (4) exclude tails of all $B - 1$ blocks, so each tail has value in $[0, 1)$. Thus $A = \lfloor |V(w)| \rfloor$, and successive multiplication of the tails by B followed by taking the floor recovers the displayed blocks. This is the digit-selection step of division. The sign and integer word match by (1). In the terminating case, (2) fixes the first zero remainder, including the empty fractional word when $t = 0$. No nonzero remainder repeats before step t: equal remainders $r_i = r_j$ with $i < j < t$ would give equal tails after i and j, so the tail after i would be periodic and nonzero, contradicting termination.

In the recurring case, (3) and (4) exclude eventual zeros. Primitivity gives least period ℓ, and any eventual period h is a multiple of ℓ. Otherwise a shift by h would give a repeated word of length $\gcd(h, \ell) < \ell$. If the repeat started before t, the blocks at positions t and $t + h$ would agree, since both would lie inside the periodic tail. The block at $t + h$ is e_ℓ, which violates (5). Since the tails determine their remainders and vice versa, division must return the displayed prefix and cycle. Hence $w = N(V(w))$. □

The tests use only the displayed blocks. The check for primitivity can be done in linear time by computing the prefix function of the cycle word, which gives the length of its longest proper prefix that is also a suffix. The cycle is a proper power exactly when the candidate period given by subtracting that length is smaller than the word length and divides it.

Theorem 1 (Existence, uniqueness, and lengths)**.** *Every rational number has a unique canonical Iteris numeral. Write its reduced denominator as* $q = 2^a 5^b M$*, with* $\gcd(M, 10) = 1$*, and set*

$$t = \left\lceil \frac{\max(a, b)}{4} \right\rceil .$$

If $M = 1$*, there are exactly* t *fractional blocks and no period. If* $M > 1$*, there are exactly* t *fractional blocks before the bar and the period has length*

$$\ell = \operatorname{ord}_M(B) = \frac{\operatorname{ord}_M(10)}{\gcd(\operatorname{ord}_M(10), 4)}$$

in blocks, where ord_M *denotes multiplicative order modulo* M*.*

Proof. By Proposition 1, the normalized block conditions characterize the well-formed numerals w satisfying $w = N(V(w))$. The identity $V(N(x)) = x$ then shows that $N(x)$ is canonical and that every canonical numeral of value x equals $N(x)$.

The division rule gives $r_j \equiv pB^j \pmod q$. When $M = 1$, the first zero occurs at the least j for which $q \mid B^j$, since $\gcd(p, q) = 1$. This gives $j = t$. When $M > 1$, no zero remainder occurs. Equality of remainders $r_j = r_{j+h}$ for $h \geq 1$ is equivalent to

$$q \mid B^j(B^h - 1).$$

Since $B^h - 1$ is coprime to 10, this is equivalent to $j \geq t$ and $\operatorname{ord}_M(B) \mid h$. If $L = \operatorname{ord}_M(10)$, then $B^h \equiv 1 \pmod M$ is equivalent to $L \mid 4h$, giving the stated length. Any block expansion with a prefix of length j and a cycle of length h has a reduced denominator dividing $B^j(B^h - 1)$ by (2), so these lengths are minimal. □

For 1/7, the ordinary period has six digits. The Iteris period contains three blocks, or twelve digits, so the bar ends on a block boundary. The 1/30000 example illustrates how a zero block before the bar retains its place value.

Value	Fractional blocks	Iteris numeral
1/6	1666 $\overline{6666}$	
1/7	$\overline{1428\ 5714\ 2857}$	
1/30000	0000 $\overline{3333}$	
−3.14159	1415 9000	

Corollary 1 (Block period lengths)**.** *Assume* $M > 1$ *in Theorem 1 and put* $L = \operatorname{ord}_M(10)$*. One block cycle spans* $4\ell = \operatorname{lcm}(4, L)$ *decimal digits. It occupies one node exactly when* $M \mid 9999$*, and at most two nodes exactly when* $M \mid 10^8 - 1$*. For a prime* $p \nmid 10$*, the block period of* $1/p$*, measured in base-B blocks, divides* $(p-1)/\gcd(p-1, 4) < p-1$*. Relative to the least decimal period* $L = \operatorname{ord}_p(10)$*, the block count is* $L/\gcd(L, 4)$*. In particular, the two counts are equal when* L *is odd.*

Proof. The length relation follows from Theorem 1. The one- and two-node criteria are $B \equiv 1 \pmod M$ and $B^2 \equiv 1 \pmod M$, respectively. For the prime claim, put $d = \gcd(p - 1, 4)$. Then $B^{(p-1)/d}$ is a power of $10^{p-1} \equiv 1 \pmod p$. Since p is odd, $d \geq 2$. □

For 1/101, the period 0099 fits in one node. For 1/17, 10 has order 16 modulo 17, since $10^8 \equiv -1 \pmod{17}$. The cycle spans sixteen decimal digits grouped into four nodes.

$$\frac{1}{17} = 0.\overline{0588\ 2352\ 9411\ 7647} = \ |\cdot$$

4 The graph of valid digit edits

Let $S(d)$ be the stroke set of digit d. Represent $S(d)$ by its incidence vector on $(p_1, p_2, p_3, p_4, p_6)$. The ten vectors are vertices of the binary cube Q_5. Let G be the graph induced on these vertices. Isometric subgraphs of hypercubes were characterized by Djoković [5]. We give a direct isometry proof for this particular graph. Its edges join valid digits that differ by exactly one added or deleted stroke. A move along an edge therefore always ends at a valid digit. Figure 2 gives all thirteen edges.

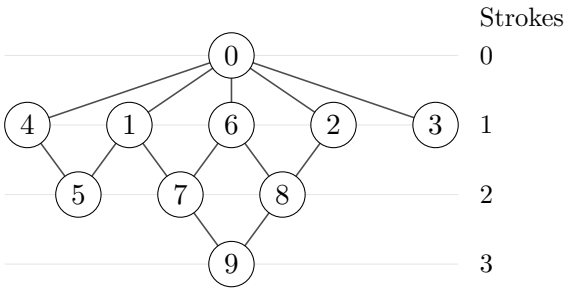

Figure 2: The graph G of one-stroke edits in one quadrant. Vertices are placed by the number of active strokes. The stem is excluded.

Theorem 2 (Stroke edits and Hamming distance)**.** *For any two digits d, e, the shortest path length in G is*

$$d_G(d, e) = d_{\mathrm{H}}(S(d), S(e)) = |S(d) \mathbin{\triangle} S(e)|.$$

The diameter of G is 4. The node graph, in which one move edits a single stroke and preserves a valid node, is the fourfold Cartesian product $G^{\square 4}$. It has 10^4 vertices, 52,000 edges, and diameter 16.

Proof. Every edge changes one stroke, so Hamming distance is a lower bound. For a family of stroke sets closed under taking subsets, this bound is attained by deleting strokes down to the intersection of the two sets and then adding strokes up to the second set.

Our family comprises all subsets of $S(3), S(5), S(9)$ except $F = \{p_1, p_2\}$. The intersection of two valid sets is valid: only $S(9)$ among them contains F, and its intersection with itself is $S(9)$. A deletion path from any set other than $S(9)$ stays valid. When deleting from $S(9)$ to a proper valid subset, that subset omits p_1 or p_2, so delete one of those omitted strokes first. The rest of the deletions stay valid. Applying this to both sets and reversing the second path shows that the intersection construction works here as well.

Only digit 9 has three strokes. Its distance can reach 4 only with a disjoint singleton, giving pairs $\{3, 9\}$ and $\{4, 9\}$. Each two-stroke set intersects $S(9)$, and the only disjoint two-stroke pair is $\{5, 8\}$. The other pairs have distance at most 3. Thus the diameter is 4.

A node move changes exactly one quadrant, giving the Cartesian product. Projecting a path to each quadrant gives a lower bound equal to the sum of quadrant distances. Following a shortest path in each quadrant in turn attains the sum. Each choice of a varying quadrant and three fixed digits contributes 13 edges, giving $4 \cdot 10^3 \cdot 13$ edges. Four pairs of quadrants each at distance 4 attain the product diameter. □

Downward closure guarantees cube isometry, and this table exhibits isometry with downward closure broken by the missing set F.

5 The geometry of decimal carry

Here m counts decimal positions. The vertices of $G^{\square m}$ are words of exactly m decimal digits, padded with zeros on the left when needed. The padding fixes which positions contribute to the edit cost. Canonicity in Section 3 applies to the displayed sequence of base-B nodes. In particular, for $m = 4$ the whole word labels one node and its leading zero digits occupy quadrants within that node.

The following two jumps share a destination and show how numeric and stroke distances differ.

Decimal transition	Iteris transition	Numeric increase	Stroke edits
0000 → 1000	▏ → ▗	1000	1
0999 → 1000	▛ → ▗	1	10

The carry in the second jump clears three strokes from each of three quadrants and adds one stroke in the thousands quadrant.

In the usual radix-addition algorithm, incrementing a decimal integer propagates a carry through its trailing run of nines [8, Section 4.3.1]. Here the cost of the increment is the graph distance between the stroke patterns of the two fixed-width words. The following theorem gives the exact distribution of that cost.

Theorem 3 (Carry cost)**.** *Let n be a nonnegative integer written in m decimal positions, padded on the left with zeros, and suppose $n < 10^m - 1$. Use one copy of G per position and let $\Delta_m(n)$ be the shortest graph distance between n and $n+1$ in that product. If n ends in k nines and $a < 9$ is the digit immediately preceding them, then*

$$\Delta_m(n) = 3k + c_a, \qquad (c_0, \ldots, c_8) = (1, 2, 2, 2, 1, 3, 1, 2, 1).$$

For each $k = 0, \ldots, m-1$, the costs $3k+1, 3k+2, 3k+3$ occur $4 \cdot 10^{m-k-1}$, $4 \cdot 10^{m-k-1}$, and 10^{m-k-1} times, respectively.

Proof. The increment replaces the last k nines by zeros, increments the preceding digit a, and preserves the rest. By Theorem 2 and the same product distance argument for m positions, the total is the sum of those edits. The $9 \to 0$ distance is 3. Reading the two stroke sets for a and $a+1$ in Table 1 gives c_a. Among the nine possibilities for a, four have $c_a = 1$, four have $c_a = 2$, and one has $c_a = 3$. Each choice leaves $m - k - 1$ unrestricted digits, giving the counts. □

For one node, $m = 4$. The theorem counts all 9,999 successor steps within a node. Since $c_5 = 3$ and all other c_a are smaller, the cost 12 occurs only for $5999 \to 6000$. The transition $9999 \to 10000$ changes the number of nodes and lies outside this fixed-width calculation.

Corollary 2 (Total cost, mean, and variance)**.** *For $m \geq 1$, the total cost of counting from 0 to $10^m - 1$ is*

$$\sum_{n=0}^{10^m-2} \Delta_m(n) = 2(10^m - 1) - 3m.$$

For a uniformly chosen increment among those $10^m - 1$ steps,

$$\mathbb{E}[\Delta_m] = 2 - \frac{3m}{10^m - 1},$$
$$\mathrm{Var}(\Delta_m) = \frac{14}{9} - \frac{9m^2 10^m}{(10^m - 1)^2}.$$

The mean tends to 2 and the variance to 14/9.

Proof. The digit tour $0 \to 1 \to \cdots \to 9 \to 0$ costs 18. Position i, with $i = 0$ for the units place, goes through this tour 10^{m-i-1} times with the final $9 \to 0$ omitted. Its contribution is $18 \cdot 10^{m-i-1} - 3$. Summing over $i = 0, \ldots, m-1$ gives the total. Dividing by $10^m - 1$ gives the mean.

Theorem 3 makes the digit a preceding the trailing nines uniform on $\{0, \ldots, 8\}$ and independent of the count K of trailing nines. Thus c_a takes values $1, 2, 3$ with probabilities $4/9, 4/9, 1/9$, respectively, and has variance $4/9$. The law of K is proportional to r^k for $0 \le k < m$, with $r = 1/10$. For $Z(r) = \sum_{k=0}^{m-1} r^k$, $r\, d(\log Z)/dr = \mathbb{E}[K]$, and applying $r\, d/dr$ again gives $\mathrm{Var}(K)$. Using $Z(r) = (1 - r^m)/(1 - r)$ gives

$$\mathbb{E}[K] = \frac{r}{1-r} - \frac{mr^m}{1-r^m} = \frac{1}{9} - \frac{m}{10^m - 1},$$

so that $\mathbb{E}[\Delta_m] = 3\,\mathbb{E}[K] + 5/3$ agrees with the mean above, and

$$\mathrm{Var}(K) = \frac{10}{81} - \frac{m^2 10^m}{(10^m - 1)^2}.$$

Independence gives $\mathrm{Var}(\Delta_m) = 9\,\mathrm{Var}(K) + 4/9$. □

6 Optimal stroke orderings

A combinatorial Gray code lists objects with a prescribed small change between consecutive entries [6, 10]. Here a one-stroke code would visit every fixed-width word once along a Hamiltonian path in the edit graph. We charge graph distance between consecutive listed words, including the return to the first word for a cyclic order. An edit path between two listed words can pass through another valid word without visiting that word in the order.

Theorem 4 (Optimal open and cyclic orders)**.** *For $m \ge 1$, the even and odd stroke-parity classes in $G^{\square m}$ have respective sizes $(10^m + (-2)^m)/2$ and $(10^m - (-2)^m)/2$. The even class is larger for even m, and the odd class is larger for odd m. The graph has no one-stroke Gray code visiting all 10^m words. The minimum total stroke distance is $10^m + 2^m - 2$ for an open order and $10^m + 2^m$ for a cyclic order. Both minima are attained by orders with steps of one or two strokes.*

Proof. A single stroke flips parity. The even digit class is $\{0, 5, 7, 8\}$, and the remaining six digits are odd. The difference between the numbers of even and odd words is the product of the signed parity sums of the coordinates, namely $(4 - 6)^m = (-2)^m$, with absolute difference 2^m.

Let $A > B$ denote the sizes of the majority and minority classes. In an open order, the minority vertices separate the majority vertices into at most $B + 1$ runs. At least $A - B - 1 = 2^m - 1$ consecutive pairs therefore have both words in the majority class, so each of those pairs costs at least two strokes. Each other step costs at least one stroke. This gives the open bound and rules out a one-stroke order. The same argument gives at most B majority runs in a cyclic order, so at least 2^m of its steps cost at least two strokes.

For a matching construction, use the one-stroke path

$$P = (1, 5, 4, 0, 2, 8, 9, 7, 6).$$

It visits every digit except 3, which is at distance two from both endpoints. Partition the words into groups indexed by the set T of positions with digit 3. The free coordinates in each group are copies of P.

Within each group, we use a reflected ordering of the free coordinates, as in the standard generation of tuples [9, Section 7.2.1.1].

Each such grid has a one-stroke traversal from the corner with all free digits 1 to the corner with all free digits 6. For no free coordinates the traversal visits only one word. Otherwise, inductively traverse the nine slices of the first free coordinate in path order, alternating between the traversal of the remaining coordinates and its reverse. Consecutive slices meet at the same corner in those coordinates, because by induction the traversal of the remaining coordinates ends at their all-6 corner and its reverse starts there. Since P has nine vertices, an odd number, the traversal ends at the all-6 corner.

List the groups in binary reflected Gray order of the masks for T, starting with $T = \varnothing$. The order is obtained recursively by prefixing 0 to the previous list and 1 to its reversal. Consecutive masks, including the closing pair, differ in one bit. Index the groups by $r = 0, \ldots, 2^m - 1$ in this Gray order. Traverse even-indexed groups from the all-1 corner of their free coordinates to the all-6 corner, and odd-indexed groups in reverse. For the singleton group $T = \{1, \ldots, m\}$, both corner descriptions denote the same word $(3, \ldots, 3)$.

At a join from an even index to an odd index, the outgoing and incoming corners both use 6 on their free coordinates. At a join from an odd index to an even index, both use 1. The closing join is of the latter kind because 2^m is even. The masks differ in exactly one coordinate, where a 3 is exchanged with the corresponding endpoint digit. This costs two strokes since $d_G(3, 1) = d_G(3, 6) = 2$. At either join incident to the singleton, there are no common free coordinates to match, so the same cost calculation applies.

The open order has $10^m - 2^m$ internal steps of cost one and $2^m - 1$ joins of cost two. The cyclic order adds the closing join of cost two. These costs attain the respective lower bounds. □

For one node the optimal open order uses 9,984 steps of cost one and 15 of cost two, a total of 10,014. The optimal cyclic order costs 10,016. Numerical order uses 19,986 strokes from 0000 through 9999.

Reassigning the glyphs to decimal digits. Keeping the bare stem at zero, the same ten stroke sets admit 9! assignments to the digits. By Theorem 4 with $m = 1$, the minimum closed digit tour cost is 12. The adopted assignment costs 18. Of the 362,880 assignments, 18,192, or about 5.01%, cost less than 18. To reproduce the count, enumerate permutations π of the digits with $\pi_0 = 0$ and sum the ten distances $|S(\pi_i) \triangle S(\pi_{(i+1) \bmod 10})|$ from Table 1. The two orientations are counted separately. The tour cost T governs the counting cost only to leading order. By the argument of Corollary 2, position i contributes $T \cdot 10^{m-i-1} - d_G(\pi_9, \pi_0)$, so the exact total also depends on the closing distance.

7 Discussion and reproducibility

The arithmetic distance and the stroke distance answer different questions about the same numerals. A large change in value may alter one stroke, and the cost of a unit increment depends on which digits carry. The graph makes this behavior explicit for the adopted set of shapes. For a different stroke table, the isometry and local costs can be tested again.

The project repository contains the core specification, a reference SVG, and a browser demonstration with exact examples of recurring decimals [11]. At the cited revision, the `paper/` directory contains `check_iteris_paper.py` for Sections 4–6 and `iteris.sty` for the worked examples. The script uses the Python standard library and checks the digit graph, the carry and traversal results for widths 1 through 4, and all 9! assignments with zero fixed. Run it with `python paper/check_iteris_paper.py` from the repository root. The companion script

`paper/iteris_extra_checks.py` checks Section 3, verifying Theorem 1 for every denominator $q < 4000$ and the canonical-form conditions by enumeration in small bases. The results in this paper cover exact positional representation and a discrete edit metric. How readily people learn or transcribe the glyphs, and how robustly software reads them from images, are questions for separate evaluations.

AI assistance. The author used AI tools in preparing this work. Gemini 4 Neon was used to explore the notation and its graph interpretation, and it proposed the limiting variance calculation. Claude (Fable 5.2 and Opus 5.5) was used to check the first draft independently; Claude Fable 5.2 proposed the syntactic characterization and the optimal stroke ordering, and Claude Opus 5.5 commented on publication venues and on the Deep Review report. Grok Bot was used to review the second draft, and it suggested clarifications to the definitions, parity classes, and period bounds. ChatGPT 6.2 Sol Ultra was used for proof checks, drafting, and LaTeX preparation. The author takes full responsibility for all content.

References

[1] David A. King. *The Ciphers of the Monks: A Forgotten Number-Notation of the Middle Ages*. Boethius 44. Franz Steiner Verlag, Stuttgart, 2001.

[2] Kirk Miller. Background for Unicode consideration of Cistercian numerals. Unicode Technical Committee document L2/20-290, 10 December 2020. `https://www.unicode.org/L2/L2020/20290-cistercian-digits.pdf`.

[3] Jonathan P. Spratte. `xistercian` Cistercian numerals in LaTeX, version 1.3, 2025. `https://ctan.org/pkg/xistercian`.

[4] Eric W. Weisstein. Cistercian numerals. *MathWorld—A Wolfram Web Resource*. `https://mathworld.wolfram.com/CistercianNumerals.html`. Accessed 25 November 2026.

[5] D. Ž. Djoković. Distance-preserving subgraphs of hypercubes. *Journal of Combinatorial Theory, Series B*, 14(3):263–267, 1973. `https://doi.org/10.1016/0095-8956(73)90010-5`.

[6] Carla Savage. A survey of combinatorial Gray codes. *SIAM Review*, 39(4):605–629, 1997. `https://doi.org/10.1137/S0036144595295272`.

[7] G. H. Hardy and E. M. Wright. *An Introduction to the Theory of Numbers*. 6th ed., revised by D. R. Heath-Brown and J. H. Silverman. Oxford University Press, 2008.

[8] D. E. Knuth. *The Art of Computer Programming, Vol. 2: Seminumerical Algorithms*. 3rd ed., Addison-Wesley, 1997.

[9] D. E. Knuth. *The Art of Computer Programming, Vol. 4A: Combinatorial Algorithms, Part 1*. Addison-Wesley, 2011.

[10] Torsten Mütze. Combinatorial Gray codes—an updated survey. *The Electronic Journal of Combinatorics*, Dynamic Survey DS26, first published 2023, version of 20 September 2024. `https://doi.org/10.37236/11023`.

[11] Patricio M. Vera. Iteris Core v1.0, draft specification and reference implementation. `https://github.com/patriciomvera/Iteris`, revision `65edae6`.