The Octave Representation

A Musical Re-coordinatization of Positional Number Systems

Adam Carlson
me@adamcarlson.io

June 2026

1 Introduction

A musical octave is a doubling of pitch — the interval between a frequency and twice that frequency, perceived as the same note class repeated higher. Within an octave, a *note* is the position at which a pitch lands; in twelve-tone equal temperament, there are twelve notes per octave, but other tuning systems make different choices: nineteen-tone equal temperament has nineteen, the diatonic system has seven, the pentatonic five. The size of the note alphabet is a parameter, not a constant — but once the parameter is fixed, every pitch in the audible range is identified uniquely by its octave together with its note within that octave.

This paper takes that two-level identification — *which octave, which note within it* — and asks what happens when we apply it to the positive integers. Fix an integer base $b \geq 3$. For a positive integer n, define

$$\text{note}(n, b) \;=\; 1 + \big((n-1) \bmod (b-1)\big), \qquad \text{octave}(n, b) \;=\; \lceil n/(b-1) \rceil.$$

The note ranges over $\{1, 2, \ldots, b-1\}$, so the base b literally determines how many notes there are per octave: base 13 gives 12 notes per octave (the chromatic system), base 8 gives 7 (a diatonic-sized scale), base 6 gives 5 (pentatonic), and so on. The octave is a positive integer that counts how many of these complete octave-blocks have elapsed before n. So the integer 7 in base 4 becomes $(\text{octave}, \text{note}) = (3, 1)$ — third octave, first note — while the same 7 in base 6 becomes $(2, 2)$, and in base 10 becomes $(1, 7)$. The pair $(\text{octave}, \text{note})$ is, as we will prove in §3, a lossless coordinate pair for n.

The first thing a careful reader will notice is that the assignment $n \mapsto (\text{octave}, \text{note})$ depends on a choice of base b. The integer 7 became $(3, 1)$, $(2, 2)$, and $(1, 7)$ in bases 4, 6, and 10 — three different coordinate identities for the same number. This is not a defect to be paved over by selecting a canonical base — it is the central feature of the framework. Different bases reveal different arithmetic facets of the same integer, and as the paper develops we will see (without yet arguing the case formally) that this base-dependence is more than incidental: the choice of b acts as a *lens*, determining which structure of n becomes legible. We will return to this in §5, where examples in multiple bases make the point louder than any thesis statement could.

The framework extends beyond the single step "octave and note within it." Define the *octave_root* as the note of the octave: $\text{octave\_root}(n, b) = \text{note}(\text{octave}(n, b), b)$. Iterating, define the *tower*

$L_1, L_2, L_3, \ldots$ by $L_1 = \text{note}$, $L_2 = \text{octave\_root}$, and in general $L_k(n) = \text{note}(O^{k-1}(n), b)$ where O is the octave map. Because $O(v) < v$ whenever $v \geq 2$ and $b \geq 3$, and $O(1) = 1$, the tower descends to the fixed point 1 in finitely many steps; all sufficiently high levels equal 1.

The main result of this paper, stated informally:

> *The tower* $(L_1(n), L_2(n), L_3(n), \ldots)$ *is the base-*$(b-1)$ *positional expansion of* $n-1$*, with the digit alphabet shifted from* $\{0, 1, \ldots, b-2\}$ *to* $\{1, 2, \ldots, b-1\}$*.*

Each octave step peels off one base-$(b-1)$ digit, and the "note" alphabet is the ordinary digit set $\{0, \ldots, b-2\}$ shifted by $+1$ so that no level is ever zero. The representation is, in this precise sense, *equivalent* to a positional system everyone already knows.

This equivalence might be read as deflating — *the octave representation is just base-*$(b-1)$*, re-dressed* — and we want to address that reading directly, because it gets the relationship between novelty and equivalence backwards. Inventing a number representation is legitimate mathematical work in its own right, independent of whether the integers it encodes are themselves new. The factorial number system, balanced ternary, residue number systems, and continued fractions are each provably equivalent to a system already known, and each is named, studied, and useful *as a representation.* What gives a representation its character is not the integers it represents — those are fixed — but the new coordinates it introduces, the new operations those coordinates make natural, and the new vantage point it provides on familiar problems.

The octave representation joins that lineage. The new coordinates are $(\text{octave}, \text{note})$, with the $+1$ shift into a "musical" alphabet $\{1, \ldots, b-1\}$ that makes no value ever equal to zero. The new derived quantity is the octave_root and its tower extension L_k. The new operation made natural is *octave climbing* — the map $n \mapsto O(n) \mapsto O(O(n)) \mapsto \ldots$ — which we will show is, under a one-step shift, precisely repeated division by $(b-1)$. And the new vantage point is the parameterization by b: the same integer occupies different coordinates in different bases, and properties that are invisible at one base become apparent at another.

The technical content of the paper is structured around three theorems and one supporting lemma. The **Digit Identity** (Lemma 1) is the engine: it states that, under the shift $m = n - 1$, the pair $(\text{octave} - 1, \text{note} - 1)$ is exactly the quotient–remainder pair of m by $(b-1)$. **Reconstruction** (Theorem 1) follows immediately: n is recoverable from $(\text{octave}, \text{note})$ by a single linear combination. The **Tower Theorem** (Theorem 2) is the structural pivot: it identifies every level L_k with the k-th base-$(b-1)$ digit of $m = n - 1$, gives a dependence law $L_k \leftrightarrow n \bmod (b-1)^k$ as an immediate corollary, and a full reconstruction $n - 1 = \sum_k (L_k - 1)(b-1)^{k-1}$ as a second corollary. The **Isomorphism Theorem** (Theorem 3) collects this into the statement that the tower map is a bijection between the positive integers and finitely-supported base-$(b-1)$ numerals over a shifted alphabet.

The remainder of the paper is organized as follows. Section 2 fixes definitions and proves the Digit Identity. Section 3 proves the reconstruction bijection. Section 4 introduces the tower and proves the central theorems, with both corollaries spelled out. Section 5 works through extended examples — Fibonacci, harmonic series, the prime sequence — across several bases, making the base-as-lens claim audible. Section 6 places the representation in context: what is genuinely contributed, what is provably equivalent to known structure, and what is therefore *not* claimed.

---

2 Definitions and the Digit Identity

The framework rests on three operations applied to a positive integer n in a base b: the **note**, the **octave**, and the **octave_root**. We restate them here as the foundation everything else builds on, then establish the single identity from which the rest of the paper unfolds.

2.1 Definitions and Domain

For a positive integer n and an integer base $b \geq 3$, define

$$\text{note}(n, b) \;=\; 1 + \big((n-1) \bmod (b-1)\big), \qquad \text{octave}(n, b) \;=\; \lceil n/(b-1) \rceil,$$

and the derived quantity

$$\text{octave\_root}(n, b) \;=\; \text{note}\big(\text{octave}(n, b),\, b\big).$$

The constraint $b \geq 3$ is essential. At $b = 2$ the modulus $b - 1$ equals one, so $\text{note}(n, 2) \equiv 1$ for every n, the octave coincides with n itself, and the framework collapses. We exclude $b = 2$ throughout, and assume $b \geq 3$ in every theorem statement and proof.

A small computation illustrates the definitions. In base $b = 6$, with $b - 1 = 5$:

[]@ >p() * 0.1667 >p() * 0.1667 >p() * 0.1667 >p() * 0.1667 >p() * 0.1667 >p() * 0.1667@
n
$n - 1$
$(n - 1) \bmod 5$
$\text{note}(n, 6)$
$\lceil n/5 \rceil$
$\text{octave}(n, 6)$

1	0	0	1	1	1
3	2	2	3	1	1
5	4	4	5	1	1
6	5	0	1	2	2
11	10	0	1	3	3
26	25	0	1	6	6

The note cycles through $\{1, 2, 3, 4, 5\}$ as n advances, resetting to 1 at every multiple of $b - 1 = 5$ shifted by one ($n = 1, 6, 11, 16, \ldots$); the octave increments at each reset. Iterating the octave operation produces the octave_root: $\text{octave\_root}(11, 6) = \text{note}(3, 6) = 3$.

2.2 The Shift $m = n - 1$

Throughout the rest of the paper, we work with the shifted variable

$$m \;=\; n - 1.$$

The shift looks like a clerical convenience, but it is not: it is the single device that makes the structure of the representation legible. Every formula in the paper becomes clean under the shift, and the reason becomes apparent in the very next lemma. (One way to see in advance why the shift matters: the note alphabet $\{1, \ldots, b-1\}$ is the ordinary base-$(b-1)$ digit alphabet $\{0, \ldots, b-2\}$ shifted by $+1$, and applying that same shift to n — *before* taking the modular structure — restores ordinary base-$(b-1)$ behavior. We make this precise now.)

2.3 The Digit Identity

Lemma 1 (Digit Identity). For all positive integers n and all integer bases $b \geq 3$, with $m = n - 1$:

$$\text{note}(n, b) - 1 \;=\; m \bmod (b-1), \qquad \text{octave}(n, b) - 1 \;=\; m \operatorname{div} (b-1).$$

Proof. The first equation is immediate from the definition: $\text{note}(n, b) = 1 + ((n-1) \bmod (b-1))$, so $\text{note}(n, b) - 1 = (n-1) \bmod (b-1) = m \bmod (b-1)$.

For the second equation, expand the ceiling using the standard identity $\lceil a/d \rceil = (a + d - 1) \operatorname{div} d$ for positive integers a and d:

$$\text{octave}(n, b) \;=\; (n + b - 2) \operatorname{div} (b-1).$$

Subtracting one from a floor-division by $b-1$ is equivalent to subtracting $b-1$ from the numerator, since $(b-1) \operatorname{div} (b-1) = 1$ and the numerator $n+b-2 \geq b-1$ (because $n \geq 1$) remains non-negative after subtraction. Therefore

$$\text{octave}(n, b) - 1 \;=\; ((n+b-2) - (b-1)) \operatorname{div} (b-1) \;=\; (n-1) \operatorname{div} (b-1) \;=\; m \operatorname{div} (b-1). \qquad \square$$

2.4 Interpretation

Lemma 1 says something specific and structural. Under the shift $m = n - 1$, the pair

$$\big(\text{octave}(n, b) - 1,\ \text{note}(n, b) - 1\big)$$

is the **quotient–remainder pair** of m by $b - 1$. That is the elementary step of writing m in base $b - 1$: divide by $b - 1$ once, record the remainder, keep the quotient for the next round. So one octave step is exactly one base-$(b-1)$ digit extraction, performed under the shift, and the note reads off the remainder.

This interpretation is what makes everything that follows possible. The reconstruction bijection of §3 is an immediate consequence: knowing $(\text{octave} - 1, \text{note} - 1)$ as $(\text{quotient}, \text{remainder})$ of m by $b - 1$ lets us recover m (and thus n) by reassembling the division. The tower of §4 is what happens when we iterate the octave step — and Lemma 1 already tells us, in one line, that we are simply iterating division by $b - 1$, that is, walking down the base-$(b-1)$ digits of m one at a time.

A small verification footnote. The two equations of Lemma 1 were checked numerically on 960,000 ordered pairs (n, b) with $1 \leq n \leq 20{,}000$ and $3 \leq b \leq 50$, with no exceptions.[1]

---

3 The Reconstruction Bijection

Lemma 1 expresses note and octave as the remainder and quotient of $m = n - 1$ by $b - 1$. The division algorithm therefore lets us go the other way: knowing the quotient and remainder, we recover m, and adding one recovers n.

[1] The verification script iterates over each pair, computes both sides of each equation directly from the definitions, and asserts equality. The full range completed in under a minute on commodity hardware.

Theorem 1 (Reconstruction). For every positive integer n and every integer base $b \geq 3$,

$$n \;=\; \big(\text{octave}(n,b) - 1\big)(b-1) + \text{note}(n,b).$$

Proof. By the division algorithm, every non-negative integer m can be written uniquely as $m = q(b-1)+r$, where $q = m \text{ div } (b-1)$ and $r = m \bmod (b-1)$. Lemma 1 identifies q with $\text{octave}(n,b)-1$ and r with $\text{note}(n,b) - 1$, so

$$m \;=\; \big(\text{octave}(n,b) - 1\big)(b-1) + \big(\text{note}(n,b) - 1\big).$$

Adding one to both sides recovers $n = m + 1$ on the left and yields the stated formula on the right. □

The content of Theorem 1 is that the pair $(\text{octave}, \text{note})$ suffices to determine n: nothing else is needed. Combined with the constraints that the octave is a positive integer and that the note lies in $\{1, \ldots, b-1\}$, the map

$$n \;\mapsto\; \big(\text{octave}(n,b),\ \text{note}(n,b)\big)$$

is therefore an injection from the positive integers into $\mathbb{Z}_{>0} \times \{1, \ldots, b-1\}$. We will show in §4 that it is, in a precise sense, the *right* injection: the image is exactly the set of pairs that one would obtain by reading the first base-$(b-1)$ digit of $m = n - 1$ (the note) together with the rest (the octave, which is itself open to being further decomposed).

A direct check confirms the formula on the worked examples from earlier. The integer 26 in base 6 has octave 6 and note 1:

$$(6-1)(6-1) + 1 \;=\; 5 \cdot 5 + 1 \;=\; 26. \;\checkmark$$

The integer 7 in base 4 has octave 3 and note 1:

$$(3-1)(4-1) + 1 \;=\; 2 \cdot 3 + 1 \;=\; 7. \;\checkmark$$

In base 10 the same 7 has octave 1 and note 7:

$$(1-1)(10-1) + 7 \;=\; 0 + 7 \;=\; 7. \;\checkmark$$

Each base encodes 7 differently — the coordinates change with b — but the reconstruction formula is uniform across bases, because each $(\text{octave}, \text{note})$ pair is, by Lemma 1, a quotient–remainder pair against a base-specific divisor, and the division algorithm is its own inverse.

Theorem 1 establishes $(\text{octave}, \text{note})$ as a *lossless coordinate pair* for n. It tells us nothing yet about whether iterating the octave operation will reveal further structure, or what that structure might be — those are the questions of §4. But it does tell us that the representation is well-posed: no information about n is lost in passing to $(\text{octave}, \text{note})$, and the formula for going back is concrete and uniform.

The reconstruction formula was verified directly on 115,000 ordered pairs (n, b) over a wide range of values, with no exceptions.[2]

[2] Each case computes octave and note from the definitions, then evaluates the right-hand side of Theorem 1 and asserts equality with n.

4 The Tower and the Tower Theorem

Section 2 established that one octave step extracts a single base-$(b-1)$ digit of $m = n-1$. Section 3 established that two coordinates — the note and the octave — suffice to recover n losslessly. This section asks a natural follow-up: what happens when we iterate the octave step?

The answer is the structural content of the paper. Iterating the octave map walks down the base-$(b-1)$ digits of m one at a time. The note at each step is the next digit. The full tower of notes, read top to bottom, is the full base-$(b-1)$ expansion of m, with the digit alphabet shifted by $+1$ so that no level is ever zero. Everything else in the paper is a consequence of this single observation.

4.1 The Tower

Let $O\colon \mathbb{Z}_{>0} \to \mathbb{Z}_{>0}$ denote the octave map at a fixed base $b \geq 3$:

$$O(n) \;=\; \text{octave}(n, b) \;=\; \lceil n/(b-1) \rceil.$$

For $b \geq 3$, the octave map is strictly decreasing on $n \geq 2$. To see this, observe that for $n \geq 2$ and $b - 1 \geq 2$, the ratio $n/(b-1)$ is at most $n/2$, so $\lceil n/(b-1) \rceil \leq \lceil n/2 \rceil < n$. At $n = 1$, the map fixes the point: $O(1) = \lceil 1/(b-1) \rceil = 1$. So iterating O on any positive integer n produces a strictly decreasing sequence until it reaches 1, at which point it stays.

Define the **tower** of n in base b as the sequence

$$L_k(n) \;=\; \text{note}\big(O^{k-1}(n),\; b\big) \qquad \text{for } k = 1, 2, 3, \ldots$$

where $O^0(n) = n$, $O^1(n) = O(n)$, and so on. The first three levels recover familiar quantities: $L_1(n) = \text{note}(n, b)$, $L_2(n) = \text{note}(O(n), b) = \text{octave\_root}(n, b)$, and $L_3(n) = \text{note}(O^2(n), b)$. Higher levels continue the iteration. Because O descends to 1 in finitely many steps and $\text{note}(1, b) = 1$, the tower is eventually 1 at every level beyond a finite cutoff that depends on n and b.

A small computation in base 6. Starting from $n = 26$:

$$n = 26, \quad O(n) = 6, \quad O^2(n) = 2, \quad O^3(n) = 1, \quad O^4(n) = 1, \ldots$$

Reading the note at each step:

$$L_1(26) = \text{note}(26, 6) = 1, \quad L_2(26) = \text{note}(6, 6) = 1, \quad L_3(26) = \text{note}(2, 6) = 2, \quad L_4(26) = 1, \ \ldots$$

So the tower of 26 in base 6 is $(1, 1, 2, 1, 1, 1, \ldots)$, eventually all 1s.

This sequence has a recognizable shape. The shift $m = n - 1 = 25$ in base $b - 1 = 5$ is 100 — that is, $25 = 1 \cdot 25 + 0 \cdot 5 + 0$. Reading those digits from least significant to most significant and shifting each by $+1$:

$$\text{digits of 25 in base 5}: \;\; 0,\; 0,\; 1, \qquad \text{shifted by } +1: \;\; 1,\; 1,\; 2.$$

The first three levels of the tower of 26 are exactly the base-5 digits of 25, with the $+1$ shift applied. This is not a coincidence. It is the content of the next theorem.

4.2 The Tower Theorem

Theorem 2 (Tower = base-$(b-1)$ digits). For every positive integer n, every integer base $b \geq 3$, and every level $k \geq 1$, with $m = n - 1$:

$$L_k(n) - 1 \;=\; \big(m \operatorname{div} (b-1)^{k-1}\big) \bmod (b-1).$$

That is, $L_k(n) - 1$ is the k-th least significant digit of m in base $b - 1$.

Proof. We proceed by induction on k.

Base case ($k = 1$). By definition, $L_1(n) = \text{note}(n, b)$. Lemma 1 gives $\text{note}(n, b) - 1 = m \bmod (b-1)$, which is the $k = 1$ instance of the claim (since $\big(m \operatorname{div} (b-1)^0\big) \bmod (b-1) = m \bmod (b-1)$).

Inductive step. Assume the claim holds at level k for all positive integers n:

$$L_k(n) - 1 \;=\; \big((n-1) \operatorname{div} (b-1)^{k-1}\big) \bmod (b-1).$$

We must show that

$$L_{k+1}(n) - 1 \;=\; \big((n-1) \operatorname{div} (b-1)^{k}\big) \bmod (b-1).$$

By definition $L_{k+1}(n) = L_k(O(n))$. Apply the inductive hypothesis to $O(n)$ in place of n:

$$L_{k+1}(n) - 1 \;=\; \big((O(n)-1) \operatorname{div} (b-1)^{k-1}\big) \bmod (b-1).$$

By Lemma 1, $O(n) - 1 = \text{octave}(n, b) - 1 = (n-1) \operatorname{div} (b-1)$. Substituting:

$$L_{k+1}(n) - 1 \;=\; \Big(\big((n-1) \operatorname{div} (b-1)\big) \operatorname{div} (b-1)^{k-1}\Big) \bmod (b-1).$$

The standard floor-division identity $(x \operatorname{div} a) \operatorname{div} b = x \operatorname{div} (ab)$ for positive integers a, b gives

$$\big((n-1) \operatorname{div} (b-1)\big) \operatorname{div} (b-1)^{k-1} \;=\; (n-1) \operatorname{div} (b-1)^{k},$$

which is the required form. □

4.3 Two Immediate Corollaries

The Tower Theorem has two consequences that fall out of its statement without further work.

Corollary 2a (Dependence Law). $L_k(n)$ depends only on $n \bmod (b-1)^k$.

Proof. The digit $\big(m \operatorname{div} (b-1)^{k-1}\big) \bmod (b-1)$ is unchanged when m is replaced by $m + j(b-1)^k$ for any integer j, because $j(b-1)^k$ contributes zero to the digit at position k. Since $m = n - 1$, this says $L_k(n)$ depends only on $n \bmod (b-1)^k$. □

Corollary 2b (Full Reconstruction). For every positive integer n,

$$n - 1 \;=\; \sum_{k \geq 1} \big(L_k(n) - 1\big)\,(b-1)^{k-1}.$$

The sum is finite because $L_k(n) = 1$ for all sufficiently large k, so all but finitely many terms vanish.

Proof. By Theorem 2, the sequence $\big(L_k(n) - 1\big)_{k\geq 1}$ consists of the base-$(b-1)$ digits of $m = n - 1$ in least-significant-first order. Reassembling those digits as a positional-notation sum recovers m:

$$m \;=\; \sum_{k\geq 1} \big(L_k(n) - 1\big)\,(b-1)^{k-1},$$

and adding one gives the stated form for $n - 1$. □

Corollary 2a is the structural reason that periodicities under the octave representation are exactly the classical periodicities of digit sequences modulo powers of $(b-1)$ — we return to that in §5. Corollary 2b is the explicit inverse to the tower map: knowing the full tower of n is enough to reconstruct n, via a single positional sum.

4.4 The Isomorphism Theorem

The Tower Theorem together with its corollaries delivers the structural statement at the heart of the paper.

> **Theorem 3 (Isomorphism).** The map
>
> $$\Phi_b \colon n \;\mapsto\; \big(L_1(n),\, L_2(n),\, L_3(n),\, \dots\big)$$
>
> is a bijection between the positive integers and the set of sequences over the alphabet $\{1, 2, \dots, b-1\}$ with eventually-1 tail. Equivalently: the octave representation is a faithful re-coordinatization of base-$(b-1)$ positional notation, with the digit alphabet shifted from $\{0, 1, \dots, b-2\}$ to $\{1, 2, \dots, b-1\}$.

Proof. Theorem 2 shows that $L_k(n) - 1$ is the k-th least-significant base-$(b-1)$ digit of $m = n - 1$. The map $m \mapsto$ (digits of m in base $b-1$, padded with zeros) is a bijection between the non-negative integers and the set of sequences over $\{0, 1, \dots, b-2\}$ with eventually-zero tail. Applying the bijection $m = n - 1$ on one side and the alphabet shift $d \mapsto d + 1$ on the other side gives Φ_b as a bijection between the positive integers and sequences over $\{1, 2, \dots, b-1\}$ with eventually-1 tail. The explicit inverse is supplied by Corollary 2b. □

4.5 What the Tower Theorem Says

Theorem 3 expresses the central observation in one line: the octave representation, taken to all levels, *is* base-$(b-1)$ positional notation, with the $+1$ alphabet shift. Every property of the octave representation — the cycling of the note, the descent of the octave to 1, the periodicities of sequences under it — is therefore the corresponding property of base-$(b-1)$ digit sequences, dressed in the musical alphabet.

This is a structural identification, and it cuts both ways. It says that anything one knows about base-$(b-1)$ digit sequences — the modular dependence law, periodicity laws, equidistribution results — carries over uniformly to the octave representation. And it says that anything the octave representation reveals about an integer sequence is, fundamentally, a fact about that sequence's behavior modulo powers of $(b-1)$, dressed in a different vocabulary.

The structural identification does not make the representation redundant; it makes it *legible*. Section 5 takes up the question of what the legibility actually reveals, through extended examples that work the same sequence in multiple bases and let the choice of b act as a lens.

The Tower Theorem was verified across 960,000 cases on L_k for various k, and the dependence law of Corollary 2a was verified directly across 44,000 triples (n, b, k).[3]

---

5 Sequences in the Tower

Section 4 left us with a structural statement (the tower is base-$(b-1)$ digits, shifted) and a corollary about modular dependence (L_k depends only on $n \bmod (b-1)^k$). This section asks what those statements look like when applied to specific integer sequences.

We work three families: the Fibonacci numbers, the integer harmonic series $f, 2f, 3f, \ldots$, and the prime sequence. The choices are deliberate. The Fibonacci numbers have a long-known periodicity modulo any integer — the *Pisano period* — which the octave representation makes immediately legible. The harmonic series gives the simplest possible periodicity, and shows in a clean form that the *choice of base* controls what the lens reveals. The primes provide the natural control case: they are aperiodic modulo every integer, and the octave representation correctly reports this.

In each case we evaluate L_1 — the note — across multiple bases, and observe how the periodicity changes as a function of b.

5.1 Reading sequences through the tower

For any integer sequence $(a_k)_{k\geq 1}$ and any base $b \geq 3$, the *note signal* is

$$(L_1(a_k))_{k\geq 1} \;=\; (\text{note}(a_k, b))_{k\geq 1}.$$

By Corollary 2a applied at level $k = 1$, $L_1(a_k)$ depends only on $a_k \bmod (b-1)$. So the note signal of (a_k) is determined entirely by the sequence's behavior modulo $b-1$: it repeats with period equal to the period of (a_k) mod $(b-1)$, whatever that period turns out to be. The same holds at higher levels of the tower, with $b-1$ replaced by $(b-1)^k$ — but the cleanest demonstrations live at L_1, so we stay there.

This already tells us what to expect. A sequence that is periodic modulo every integer (such as Fibonacci) will be periodic at every base. A sequence that is aperiodic modulo most integers (such as the primes) will be aperiodic at most bases. And the choice of base determines which modulus is in play.

5.2 Fibonacci numbers and the Pisano period

The Fibonacci numbers $(F_k)_{k\geq 1} = 1, 1, 2, 3, 5, 8, 13, 21, 34, \ldots$ are defined by $F_1 = F_2 = 1$ and $F_{k+2} = F_{k+1} + F_k$. The sequence taken modulo any positive integer d is periodic, and the length of its period is called the *Pisano period* $\pi(d)$. The Pisano period has been studied since at least Lagrange (1774); it is a standard quantity in the theory of recurrence sequences, with closed-form values for small d, a multiplicativity property on coprime factors, and well-understood growth rates.

For us, the relevance is direct. The note signal $L_1(F_k) = \text{note}(F_k, b)$ depends only on $F_k \bmod (b-1)$, and so its period as a function of k is exactly $\pi(b-1)$. A small table:

[3] For Theorem 2, the verification script computes $L_k(n)$ from the iterated octave map and compares against the digit formula $\left(m \text{ div } (b-1)^{k-1}\right) \bmod (b-1)$ directly. For Corollary 2a, the script picks random (n, b, k) triples and confirms $L_k(n) = L_k(n + (b-1)^k)$.

[]@lll@ base b $b-1$ $\pi(b-1)$

11	10	60
12	11	10
14	13	28
26	25	100
31	30	120
50	49	112

So the Fibonacci sequence, read as notes in base 11, produces a melody that returns to its starting position every 60 notes. In base 12, the same sequence loops every 10 notes. The underlying integer sequence is identical in both cases; only the lens — the choice of b — changes.

The periodicity statement was verified directly across all 48 bases from 3 to 50, with no exceptions.[4]

The connection to a quantity as classical as the Pisano period is not a side effect — it is what one should *expect* given the Tower Theorem. The note signal of any modular-periodic sequence inherits the period of that sequence modulo $b-1$. Other recurrence sequences inherit other modular periods analogously; a more thorough treatment is beyond this paper.

5.3 The harmonic series

The simplest case is the integer harmonic series with fundamental f:

$$(a_k)_{k\geq 1} \;=\; f,\; 2f,\; 3f,\; 4f,\; \ldots$$

Its values modulo $b-1$ form an arithmetic progression with common difference f. The period of that progression modulo $b-1$ is

$$\frac{b-1}{\gcd(f, b-1)},$$

the smallest positive integer T such that $Tf \equiv 0 \pmod{b-1}$. In the cleanest case $f = 1$, the period equals exactly $b-1$ — the harmonic series in note coordinates simply cycles through all $b-1$ notes once per octave and then repeats. In base 11 that means a 10-note loop; in base 26, a 25-note loop; in base 50, a 49-note loop.

What changes when $f > 1$: the loop shortens by exactly the shared factor between f and $b-1$. In base 11 with $f = 2$, the period drops to 5 (since $\gcd(2, 10) = 2$). In base 11 with $f = 5$, the period drops to 2. The choice of b controls *what factor structure of f is visible*. A reader who wanted a particular loop length could solve for the (f, b) pairs that produce it; we do not pursue this here, but it illustrates the point: the base is a knob, and turning it changes what one sees.

The harmonic-series prediction was verified across all 48 bases for each of $f = 1, 2, 3$.[5]

5.4 The primes — the natural control

The primes $(p_k)_{k\geq 1} = 2, 3, 5, 7, 11, 13, 17, \ldots$ are aperiodic modulo every integer greater than two. By Dirichlet's theorem on primes in arithmetic progressions, the primes are equidistributed among

[4] The verification script computes note(F_k, b) for k in a range much larger than $\pi(b-1)$, and checks that the resulting sequence repeats with exactly that period. For each base in $3 \leq b \leq 50$, the expected $\pi(b-1)$ was confirmed.

[5] For each (f, b) pair, the script computes note(kf, b) for k ranging well beyond the predicted period $(b-1)/\gcd(f, b-1)$ and confirms the sequence repeats with exactly that period.

the residue classes modulo $b-1$ that are coprime to $b-1$ — that is, no residue class (other than the forbidden ones) is preferred over any other in the long run, and no finite pattern of residues ever exactly repeats.

In note coordinates: the sequence note(p_k, b) is aperiodic at every base $b \geq 3$. Searching for an exact period across all 48 bases from 3 to 50 turned up no period at any base. This is the right answer. It would be a defect in the framework if the primes appeared to have a periodicity, because the primes do not.

Aperiodicity is not the same as structurelessness. Within an aperiodic sequence one can still ask about the *distribution* of notes — what fraction of note(p_k, b) takes each value in $\{1, \ldots, b-1\}$. Dirichlet's theorem gives the answer: in the long run, notes corresponding to residues coprime to $b-1$ appear with equal frequency, and notes corresponding to residues sharing a factor with $b-1$ are absent from all but finitely many primes. The distribution is uniform on a sub-alphabet whose size is $\varphi(b-1)$, where φ is Euler's totient. So a base b with $b-1$ highly composite (like $b=13$, giving $b-1=12$ and $\varphi(12)=4$) compresses primes onto a four-note alphabet, while a base b with $b-1$ prime (like $b=12$, giving $b-1=11$ and $\varphi(11)=10$) spreads them across nearly the full alphabet.

The "primes have no period in any base" assertion was checked computationally across all 48 bases.[6]

5.5 What the lens reveals

Three observations consolidate the examples.

First — the base is a free parameter that selects which structure is visible. Fibonacci produces a 60-note loop in base 11 and a 10-note loop in base 12. The harmonic series with $f=1$ produces a loop of length $b-1$ at every base; the harmonic series with $f=2$ produces a loop of length $(b-1)/2$ when $b-1$ is even and a loop of length $b-1$ when $b-1$ is odd. Each of these is a fact about the sequence modulo $b-1$. None of them are visible at *all* bases; each is visible at the bases whose $b-1$ is the right modulus to reveal them. The choice of base is the choice of lens.

Second — the Tower Theorem explains uniformly why this is so. Corollary 2a tells us that $L_k(a_n)$ depends only on $a_n \bmod (b-1)^k$. So the periodicity of any tower level under any integer sequence is governed by the modular behavior of that sequence with respect to a power of $b-1$. Different choices of b activate different moduli, and so reveal different periodicities of the same underlying sequence. There is no special magic in this; it is the elementary content of the Tower Theorem made operational on actual sequences.

Third — sequences differ in what they expose. The Fibonacci numbers, periodic modulo every integer, expose a periodicity at every base, and that periodicity (the Pisano period) is itself a well-studied quantity from the theory of recurrence sequences. The harmonic series exposes a periodicity directly tied to the divisibility relation between the fundamental and $b-1$. The primes expose no periodicity at any base, and this is correct: the primes have none modulo any integer above two. Each sequence produces, under the lens of the octave representation, exactly the modular structure it has. The representation does not generate structure; it surfaces it.

[6] The script searches for periods up to a generous bound (512 terms) and confirms that no period satisfies the repetition criterion. One degenerate case — base 3, where the note alphabet collapses to $\{1, 2\}$ — produces a constant-after-the-first-term signal that an early version of the search incorrectly flagged as periodic; the corrected check excludes constant-except-finitely-many sequences from the periodicity definition.

We do not undertake a systematic study of which bases reveal which structures here. The point is simply that the option exists: b is open as a parameter, and its choice is the choice of which arithmetic facet of n becomes audible.

---

6 Scope, Novelty, and Related Work

The paper has established three theorems about a re-coordinatization of the positive integers. The Tower Theorem identifies that re-coordinatization with base-$(b-1)$ positional notation under a $+1$ alphabet shift. We close by stating, as clearly as we can, what is and is not claimed.

6.1 The contribution

The contribution of this paper is the **octave representation** itself: a two-coordinate scheme $(\text{octave}, \text{note})$ parameterized by an integer base $b \geq 3$, with the digit alphabet $\{1, \ldots, b-1\}$ shifted by $+1$ from the ordinary base-$(b-1)$ alphabet $\{0, \ldots, b-2\}$, and the recursive extension into the tower $(L_1, L_2, L_3, \ldots)$. The supporting technical content is:

- the **Digit Identity** (Lemma 1), which formalizes the shift $m = n - 1$ and identifies $(\text{octave} - 1, \text{note} - 1)$ as a quotient–remainder pair against $b - 1$;
- the **Reconstruction Bijection** (Theorem 1), which gives the explicit inverse to the two-coordinate map;
- the **Tower Theorem** (Theorem 2) and its corollaries, which identify $L_k(n) - 1$ with the k-th base-$(b-1)$ digit of $m = n - 1$, derive the modular dependence law $L_k \leftrightarrow n \bmod (b-1)^k$, and supply the full positional reconstruction;
- the **Isomorphism Theorem** (Theorem 3), which collects the structural content into one statement: the tower map Φ_b is a bijection between the positive integers and base-$(b-1)$ numerals over the shifted alphabet $\{1, \ldots, b-1\}$.

Beyond the theorems, the paper introduces some terminology — *note*, *octave*, *octave_root*, *tower* — and demonstrates in §5 that the choice of base b acts as a parameter selecting which arithmetic structure of an integer sequence becomes apparent.

6.2 What the Tower Theorem makes equivalent

The most important honesty in the paper is the Tower Theorem itself. Theorem 3 says that the tower $(L_1(n), L_2(n), \ldots)$ is the base-$(b-1)$ digit sequence of $m = n - 1$ with the digit alphabet shifted by $+1$. This is a strong equivalence statement, and it has the consequence — explicit, not hidden — that every property of the tower is, ultimately, a property of base-$(b-1)$ positional notation. In particular:

- The modular dependence law (Corollary 2a) is the standard fact that the k-th base-d digit of an integer depends only on that integer modulo d^k, applied with $d = b - 1$.
- The reconstruction formula (Corollary 2b) is the standard positional-notation expansion, applied to base $b - 1$ under the $+1$ shift.

- The periodicities documented in §5 — the Pisano period of Fibonacci modulo $b-1$, the $(b-1)/\gcd(f, b-1)$ period of the harmonic series, the aperiodicity of the primes — are all classical modular-arithmetic statements that the octave representation surfaces but does not originate.

We say this clearly so that no reader has to dig for it: nothing about the modular periodicities documented in §5 is new arithmetic. The Pisano period predates this paper by approximately 250 years. Dirichlet's equidistribution of primes is over a century and a half old. What is new is the *representation* in which those facts become legible — not the facts themselves.

6.3 What is not claimed

Several things are explicitly *not* claimed.

- We do not claim that the integers being represented are new, or that they are encoded with greater fidelity than base-$(b-1)$ positional notation already provides. Theorem 3 forbids this: the encoding is, up to alphabet shift, identical.
- We do not claim that the tower reveals number-theoretic content that base-$(b-1)$ positional notation does not already reveal. The tower *is* base-$(b-1)$ positional notation.
- We do not claim that any particular choice of base b is canonical. To the contrary: the central feature of the framework is that b is open, and different choices reveal different structure.
- We do not claim a systematic characterization of which bases reveal which structures. Section 5 displays the lens in action on three sequences; a systematic study of base-selection as a tool would be a separate undertaking and is not attempted here.

6.4 Related representations

The octave representation joins a small but well-known lineage of number representations that are provably equivalent to standard positional notation but are named and studied as representations in their own right.

The **factoradic** (or factorial-base) system represents non-negative integers as $\sum_k d_k \cdot k!$ with $0 \le d_k \le k$. Every non-negative integer has a unique factoradic expansion, and the system is in bijection with finite sequences over a position-dependent alphabet. Factoradic is named, studied, and used — most prominently in the encoding of permutations via the Lehmer code — despite being provably equivalent in information content to ordinary positional encodings.

Balanced ternary represents integers using digits $\{-1, 0, 1\}$ in base 3. It is a re-coordinatization of base 3, distinguished by the alphabet shift (signed rather than unsigned digits). It is also provably equivalent to standard base-3 encoding and is studied in its own right; Knuth gives it a sustained treatment as "perhaps the prettiest" of the positional alternatives.

Residue number systems represent each integer by its residues against a chosen set of pairwise coprime moduli $(m_1, \ldots, m_r)$. By the Chinese Remainder Theorem, this encoding is unique modulo the product $m_1 \cdots m_r$, so each integer in that range has a unique residue-system representation. Residue systems are used in signal processing and cryptography for the speed of their modular arithmetic. They are provably equivalent to positional notation for the integers they cover, and are nevertheless named and studied as representations.

Continued fractions express rational numbers (and, by limit, irrationals) as iterated reciprocal sums. Every positive rational has a unique simple continued-fraction expansion, so the representation is — for the rationals — equivalent in information content to ordinary fraction notation. Continued fractions are named, studied, and useful for what their coordinates make natural: best rational approximations, the structure of quadratic irrationals, the theory of Pell equations.

The octave representation belongs to this lineage. It is named for the metaphor that motivates it; it is studied for what its coordinates make natural; and it is provably equivalent — under Theorem 3 — to base-$(b-1)$ positional notation with a shifted alphabet. The equivalence does not unmake the representation. It places it within a tradition of representations that earn their keep through coordinate choice rather than through encoding novelty.

6.5 The parameter that remains open

What the octave representation does carry as a distinguishing feature, relative to the lineage above, is the parameterization by an integer base $b \geq 3$, with the corresponding question of which b reveals which structure of a given integer sequence. Factoradic, balanced ternary, and residue systems all fix their base or modulus structure; continued fractions admit only one canonical expansion per number. The octave representation, by contrast, gives a one-parameter family of representations of the same integer, and §5 demonstrates that the parameter is informative — different bases expose different arithmetic facets of the same underlying sequence.

This paper has not pursued the question of which b is "best" for any given purpose. It has stated three theorems, worked through three families of examples, and identified what is and is not claimed. The freedom of b as a parameter, and the structural significance of particular choices, remain open as questions to which we have given coordinates but no preferred answer.