Submitted exclusively to the *Journal of Mathematics and Music*
Last compiled on April 27, 2026

NOTE

Prime-power indexed multiscale graph diagnostics for symbolic temporal data: methodological exploration and delimitation via BWV 1007

Jesús Rogelio Pérez Buendía[a*]

[a] *SECIHTI – Centro de Investigación en Matemáticas (CIMAT), Unidad Mérida, Mérida, Mexico*

We introduce a prime-power indexed multiscale framework for symbolic temporal data. Windows of length $N = p^n$ define a tower of pattern sets from a MIDI-derived time series; a symmetrized kNN graph at each level yields the connectivity count β_0 as primary diagnostic, supplemented by giant-component fraction, Fiedler eigenvalue, and clustering as auxiliary observables. Comparing $p = 2$ against $p = 3$, the two prime factors of the chromatic subdivision $12 = 2^2 \cdot 3$, we track the contrast $\Delta_{23}(n) = \beta_0(G_{2,n}) - \beta_0(G_{3,n})$ across levels, with pre-registered control primes $p \in \{5, 7\}$ and two null models. The main finding is a methodological delimitation of β_0 as a per-level diagnostic. Across the tested protocols, $R_{\text{ctrl}} \leq 1$ under matched level ranges and beats-axis contrasts are eliminated by local time-shuffle, an outcome more consistent, under the tested null models, with short-range temporal correlations in onset patterns than with harmonic information alone; auxiliary connectivity observables (giant-component fraction, Fiedler eigenvalue, average clustering) display sporadic but non-systematic structure on the same grid. This delimitation identifies a bottleneck of per-level connectivity and motivates the next natural step: an inter-level compatibility diagnostic between successive towers, developed in a separate companion paper (Pérez-Buendía 2026). Experiments cover synthetic toys and all six movements of Bach's Cello Suite No. 1 (BWV 1007), with two time axes and systematic parameter sweeps.

Keywords: prime-power multiscale analysis; kNN graph invariants; symbolic music analysis; graph connectivity; hierarchical structure; Bach cello suites

2020 Mathematics Subject Classification: 00A65; 05C82; 11S82
2012 Computing Classification Scheme: applied computing

1. Introduction

Multiscale analysis of temporal data typically relies on continuous filtrations indexed by a geometric scale parameter $\varepsilon > 0$ (Carlsson 2009; Edelsbrunner and Harer 2010); in symbolic music this approach has been pursued through persistent homology of intervallic transition graphs (Mijangos, Bravetti, and Padilla-Longoria 2024) and through several distance choices on music networks (Heo et al. 2025), while related work explores higher-order network topology of Bach's solo violin works (Mrad and Najem 2025) and explicit hierarchical structure synthesis from symbolic corpora (Shapiro et al. 2025). We propose an alternative: a *prime-power indexed* framework in which the observation window has length $N = p^n$ for a fixed prime p and varying $n \geq 1$. From a MIDI-derived multivariate time series (Müller 2015; Raffel and Ellis 2014), one obtains at each level a pattern set

*Corresponding author. Email: rogelio.perez@cimat.mx.

$D_{p,n}$; a symmetrized k-nearest-neighbor (kNN) graph built on $D_{p,n}$ then yields graph-theoretic invariants whose behavior across levels constitutes the multiscale diagnostic.

The sequence p^n is arithmetically distinguished: the rings $\mathbb{Z}/p^n\mathbb{Z}$ together with the natural reduction maps $\mathbb{Z}/p^{n+1}\mathbb{Z} \twoheadrightarrow \mathbb{Z}/p^n\mathbb{Z}$ form the standard inverse system whose limit recovers the ring of p-adic integers $\mathbb{Z}_p$ (Serre 1979; Gouvêa 2020), the canonical projective tower attached to the prime p. This canonicity is internal to the formalism; whether it captures a given empirical phenomenon is an experimental question depending on corpus, protocol, and observable.

The comparison between $p = 2$ and $p = 3$ is motivated by the arithmetic of the chromatic scale, but its use here is explicitly methodological rather than ontological. Western equal temperament divides the octave into $12 = 2^2 \cdot 3$ semitones. Since $\gcd(4, 3) = 1$, the Chinese Remainder Theorem gives $\mathbb{Z}/12\mathbb{Z} \cong \mathbb{Z}/4\mathbb{Z} \times \mathbb{Z}/3\mathbb{Z}$, decomposing the chromatic group into independent binary and ternary components (Balzano 1980; Tymoczko 2011). Comparing $p = 2$ against $p = 3$ interrogates these two arithmetically independent factors at the level of *temporal* window lengths; the parallel transfer of the chromatic CRT to a metric/temporal axis is a modeling decision, declared explicitly here, not a canonical correspondence in music theory.

Within that scope, the paper offers three contributions, organized as *object*, *protocol*, and *result*:

(i) **Object.** An arithmetically natural tower of pattern sets $\{D_{p,n}\}_n$ with $N = p^n$, motivated by the standard projective tower $\{\mathbb{Z}/p^n\mathbb{Z}\}_n$ associated with the prime p and by the CRT decomposition $12 = 2^2 \cdot 3$, together with the connectivity count β_0 as primary diagnostic.

(ii) **Protocol.** A reproducible experimental architecture: prime contrasts $\Delta_{23}(n)$, a control ratio R_{ctrl} against $p \in \{5, 7\}$, null models (pitch-shuffle and time-shuffle), and systematic sweeps over k, cap, and time axis.

(iii) **Result.** A precise delimitation: per-level graph connectivity does not produce robust prime-specific signal in the tested corpus ($R_{\mathrm{ctrl}} \le 1$ under matched level ranges; null-model separation fails in most conditions); this identifies β_0 as a natural first diagnostic but one too coarse to be conclusive.

The experimental scope is deliberately finite. Experiments cover synthetic binary/ternary toys and all six movements of Bach's Cello Suite No. 1 (BWV 1007), with two time axes (seconds and beat-synchronous binning), systematic sweeps over cap and k, and two external control pieces used only for null-model evaluation (BWV 1049, BWV 1050). This delimitation is methodologically informative: it identifies the exact level at which arithmetic indexing alone ceases to provide robust discrimination, and it justifies the move to inter-level compatibility diagnostics as the natural next instrument.

2. Prime-power towers and graph diagnostics

Fix a prime p and a tower level $n \ge 1$; set $N = p^n$. The input is a multivariate time series $X(t) = \big(H(t), a(t), f(t), g_W(t)\big)$ extracted from a MIDI file, where $H(t) \in [0,1]^{12}$ is the chroma vector, $a(t) \ge 0$ the onset density, $f(t) \ge 0$ the spectral flux, and $g_W(t)$ a local inter-onset-interval (IOI) histogram, all at frame t. Two time axes are considered: (i) seconds-based, with a fixed bin width $\Delta = 0.05\,\mathrm{s}$, and (ii) beat-synchronous, with bin width $\Delta_b = 1/12$ of a beat.

Definition 2.1 (Pattern set) Given a time series X of length T and a window length

$N = p^n$, the *pattern set at level* n is

$$D_{p,n} \;=\; \big\{\, X[t : t+N) \;\big|\; t = 0,\, s,\, 2s,\, \ldots \big\},$$

where $s \geq 1$ is the stride and each element is a matrix in $\mathbb{R}^{N\times d}$ (d is the number of channels). If $|D_{p,n}|$ exceeds a prescribed cap C, only the first C windows are retained.

The squared distance between two windows $w, w' \in D_{p,n}$ is

$$d(w,w')^2 \;=\; \overline{\|H_w - H_{w'}\|^2} \;+\; \alpha^2\,\overline{(a_w - a_{w'})^2} \;+\; \beta^2\,\overline{(f_w - f_{w'})^2} \;+\; \gamma^2\,\|g_W - g_{W'}\|^2,$$

where overlines denote frame-wise means. Three weight configurations are tested: (A) $(\alpha, \beta, \gamma) = (1, 0, 0)$, using chroma and onset density; (B) $(1, 1, 0)$, adding spectral flux; and (C) $(1, 1, 1)$, adding the IOI histogram. Config A is the primary configuration reported in the results; B and C serve as sensitivity checks.

Definition 2.2 (Pattern tower) The *pattern tower* for prime p is the family $\{D_{p,n}\}_{n\geq 1}$. As n increases, the window length p^n grows and the number of available windows decreases; each successive level captures coarser temporal structure.

Remark 2.3 (Relation to ultrametric data analysis) The use of p-adic indexing of pattern sets is, in spirit, related to the line of work on p-adic encoding and ultrametric clustering of hierarchical data initiated by Murtagh (Murtagh 2016). Our construction differs in that the p-adic structure is imposed on *window lengths* (a coarse-graining of the temporal axis), not on the metric of the data points themselves. The two approaches are complementary; an ultrametric embedding of the pattern sets $D_{p,n}$ at each level remains an open methodological direction.

At each level n, we build an undirected graph $G^{(k)}_{p,n}$ on the vertex set $D_{p,n}$ by symmetrizing the k-nearest-neighbor relation: vertices w and w' are connected whenever w is among the k nearest neighbors of w' or vice versa. Four invariants are recorded. The connectivity count β_0 is the primary analytic observable: it is the coarsest probe of pattern-cloud fragmentation under arithmetic refinement, and determining whether it already discriminates between primes is a natural first step before deploying finer invariants. The remaining three serve as auxiliary diagnostics; their role is to test whether the delimitation identified through β_0 persists under alternative graph summaries.

(a) **Connected components** (β_0). The zeroth Betti number of $G^{(k)}_{p,n}$, counting the number of connected components.
(b) **Giant-component fraction** (g). The fraction of vertices in the largest connected component: $g = |C_{\max}|/|D_{p,n}|$.
(c) **Fiedler eigenvalue** (λ_2). The second-smallest eigenvalue of the normalized graph Laplacian of the giant component, measuring algebraic connectivity.
(d) **Average clustering coefficient**. The mean local clustering coefficient over all vertices of degree ≥ 2.

The prime contrasts are then defined at the same level of the tower. For two primes p_a, p_b and a fixed piece π with condition Ξ (axis, bin, cap, and configuration), the contrast at level n is

$$\Delta_{ab}(\pi; \Xi, n) \;:=\; \beta_0^{(p_a)}(\pi; \Xi, n) - \beta_0^{(p_b)}(\pi; \Xi, n).$$

We track Δ_{23} (main contrast, $p = 2$ vs. $p = 3$) together with Δ_{25} and Δ_{27} (control contrasts). Given a set of cells $\mathcal{I}$ (ranging over levels $n \geq 2$ and neighborhoods k, for a fixed configuration), the mean absolute contrast is

$$\mathrm{MeanAbs}\Delta_{ab} := \frac{1}{|\mathcal{I}|} \sum_{(n,k)\in\mathcal{I}} |\Delta_{ab}(n,k)|.$$

The control ratio is $R_{\mathrm{ctrl}} := \mathrm{MeanAbs}\Delta_{23} / \max(\mathrm{MeanAbs}\Delta_{25}, \mathrm{MeanAbs}\Delta_{27})$. A value $R_{\mathrm{ctrl}} > 1$ would indicate that the 2–3 contrast exceeds the control baselines.

3. Experimental protocol

The primary corpus consists of two groups. *Synthetic toys*: a binary toy (period 2) and a ternary toy (period 3), generated as deterministic chroma sequences with known ground-truth periodicity. *Real music*: the six movements of Bach's Cello Suite No. 1 (BWV 1007), namely Prelude, Allemande, Courante, Sarabande, Menuets I–II, and Gigue, encoded as MIDI. An external corpus, used exclusively as null-model control cases (not as part of the primary analysis indicated in the title), comprises BWV 1049 (Brandenburg Concerto No. 4, movement 1, capped at 300 patterns) and BWV 1050 (Brandenburg Concerto No. 5, movement 2, capped at 300 patterns).

The parameter grid is kept small enough to make the negative result interpretable. Two time axes are tested: seconds (bin width $\Delta = 0.05\,\mathrm{s}$) and beats ($\Delta_b = 1/12$ of a beat). The kNN neighborhood size takes values $k \in \{8, 10, 12\}$; the pattern cap $C \in \{200, 300, 500\}$; and three weight configurations (A, B, C) control the relative contribution of chroma and onset density. Primes $p \in \{2, 3\}$ are the main experimental factors; $p \in \{5, 7\}$ serve as pre-registered negative controls.

Finally, two null models separate harmonic marginal information from temporal ordering. Let a source MIDI file be encoded as a finite sequence of note events $\mathcal{N} = \{(t_i, d_i, \mathrm{pitch}_i)\}_{i=1}^{M}$, with onset time t_i, duration d_i, and pitch $\mathrm{pitch}_i \in \{0, \ldots, 127\}$. Two null variants of each source are generated:

(i) *Pitch-shuffle.* A uniformly random permutation σ of $\{1, \ldots, M\}$ is drawn, and the pitches are reassigned by $\mathrm{pitch}_i \mapsto \mathrm{pitch}_{\sigma(i)}$, leaving $\{(t_i, d_i)\}$ unchanged. This destroys harmonic and melodic correlations while preserving the rhythmic envelope and the pitch multiset.

(ii) *Local time-shuffle.* The event sequence is partitioned into consecutive non-overlapping windows of fixed length W in beats; within each window, the inter-onset interval sequence is permuted uniformly while the pitch sequence is left intact. This disrupts sub-beat ordering while preserving beat-scale rhythmic density and full harmonic content. We fix $W = 1$ beat throughout, independently of the analysis time axis (seconds or beats); this choice matches the granularity of beat-synchronous binning. The implementation flag is `--window-beats 1.0` in `scripts/nullmodel_midi.py`, and sensitivity to $W \in \{0.5, 1, 2\}$ beats is documented in the companion repository.

Each null variant is generated with a fixed random seed (recorded in the manifest), and the full pipeline is rerun on the null variant exactly as on the original. The pitch-shuffle preserves the chroma marginal distribution; the local time-shuffle preserves the pitch sequence. Their joint use tests the contribution of harmonic vs. temporal correlations to the diagnostic. Full results are reported in Section 4.

4. Results: synthetic controls and BWV 1007 suite

The binary and ternary toys produce only small contrasts ($|\Delta_{23}| \leq 5$): both towers converge to $\beta_0 = 1$ at high n, and the Δ_{23} pattern lacks a consistent sign across all levels. This indicates that β_0 is too coarse to discriminate even purpose-built periodic structures in the prime-power tower.

The same limitation persists when the main primes are compared with the pre-registered controls. Using the pre-registered control primes $p \in \{5, 7\}$ (Section 3), we compute R_{ctrl} as defined in Section 2. For the Prelude (the only movement with $p = 5, 7$ data at present), we compute R_{ctrl} under seconds, config A, cap 300. Because the control towers reach fewer levels than the main primes ($n_{\max} = 3$ for $p = 7$, $n_{\max} = 4$ for $p = 5$, versus $n_{\max} = 5$ for $p = 2, 3$), we report two variants: using per-pair ranges, $R_{\text{ctrl}} = 0.125$ ($k = 10$) and $R_{\text{ctrl}} = 1.125$ ($k = 8$); restricting all contrasts to the common range $n \in \{2, 3\}$, both values drop below 1 ($R_{\text{ctrl}} = 0.25$ and 0.17, respectively). The marginally above-1 value at $k = 8$ thus reflects unequal level coverage rather than a genuine excess of the 2–3 signal.

The complete BWV 1007 suite gives the main empirical test. Tables 1 and 2 report $\Delta_{23}(n)$ for all six movements under config A, cap 300, for both time axes.

Under the *seconds* axis (Table 1), the contrasts concentrate at level $n = 1$: 94% of movement–k pairs yield $\Delta_{23} \neq 0$ there, and the sign is consistently positive, indicating that the $p = 2$ tower fragments more at the coarsest resolution. At $n = 3$ the contrast vanishes entirely (0/18 nonzero), confirming convergence. However, at $n \geq 4$ the Gigue and Courante develop large *negative* deltas that are k-invariant: $\Delta_{23} = -14$ for the Gigue at $n = 4$ (all k) and $\Delta_{23} = -9$ for the Courante at $n = 5$ (all k). The Prelude shows a weaker, k-dependent effect ($\Delta_{23} = -3, -5$ at $n = 4, 5$ only for $k = 8$). These negative contrasts arise in the movements with the highest rhythmic density and indicate that the $p = 3$ tower fragments at levels where the window length 3^n becomes large relative to 2^n (the ratio $(3/2)^n \geq 5$ at $n \geq 4$).

Under the *beats* axis (Table 2), non-zero contrasts persist across all levels. At $n = 4$, all 18 movement–k pairs are nonzero (100%), and the sign is uniformly negative: the $p = 3$ tower consistently shows more connected components when beat-synchronous binning is used. This systematic $n = 4$ effect is the strongest pattern in the entire grid, though its magnitude (typically $|\Delta_{23}| \in \{1, 2\}$) remains modest.

The two axes thus exhibit qualitatively different behavior: seconds-binned data yields a "convergence then fragmentation" profile, while beat-binned data shows distributed contrasts with a characteristic negative drift at high n.

Aggregating these tables by level sharpens the contrast between axes. Table 3 reports the fraction of movement–k pairs (out of $18 = 6$ movements $\times$ 3 values of k) that yield $\Delta_{23} \neq 0$ at each level, by axis.

The pattern is stable across weight configurations. Configs B and C (Section 2) yield qualitatively similar patterns: under seconds, nonzero Δ_{23} at $n = 1$ drops from 94% (A) to 67% (B) and 65% (C); under beats, the $n = 4$ nonzero rate remains at or above 88% for all three configs. Adding spectral flux and IOI channels leaves the delimitation essentially intact.

The auxiliary graph invariants refine the picture but do not remove the limitation of the per-level diagnostic. They were recorded at every level for all movements and conditions. Under the seconds axis, the clustering coefficient shows the highest frequency of inter-prime differences (79% of movement–k–n cells exceed $|\delta_{\text{clust}}| > 0.05$), but the sign reverses between low and high n, preventing any consistent directionality. The Fiedler eigenvalue λ_2 yields negligible contrasts ($|\delta_{\lambda_2}| < 0.05$ in almost all cells). Under beats,

Table 1. $\Delta_{23}(n) = \beta_0^{(2)} - \beta_0^{(3)}$ for BWV 1007, seconds axis, config A, cap 300. Rows grouped by movement; three k values per movement.

Movement	k	$n=1$	$n=2$	$n=3$	$n=4$	$n=5$
Prelude	8	+2	+1	0	−3	−5
	10	+7	+1	0	0	0
	12	+4	+1	0	0	0
Allemande	8	+2	+2	0	0	0
	10	+1	+2	0	0	0
	12	0	+1	0	0	0
Courante	8	+3	0	0	0	−9
	10	+3	0	0	0	−9
	12	+2	0	0	0	−9
Sarabande	8	+6	0	0	0	0
	10	+6	0	0	0	0
	12	+6	0	0	0	0
Menuets	8	+4	0	0	0	0
	10	+2	0	0	0	0
	12	+1	0	0	0	0
Gigue	8	+3	0	0	−14	−13
	10	+4	0	0	−14	−13
	12	+3	0	0	−14	−14

Table 2. $\Delta_{23}(n) = \beta_0^{(2)} - \beta_0^{(3)}$ for BWV 1007, beats axis, config A, cap 300.

Movement	k	$n=1$	$n=2$	$n=3$	$n=4$	$n=5$
Prelude	8	+2	+5	0	−1	−2
	10	+4	−1	+1	−1	−2
	12	+2	−1	+1	−1	−2
Allemande	8	+2	+1	−2	−2	0
	10	+1	+1	−2	−2	0
	12	+1	0	−2	−2	0
Courante	8	+2	0	0	−2	−2
	10	+2	0	0	−2	−2
	12	+2	0	0	−2	−2
Sarabande	8	+2	0	+2	−2	−2
	10	+3	0	+1	−2	−2
	12	+3	+1	0	−2	−2
Menuets	8	+1	+4	0	−2	−2
	10	+1	+2	0	−2	−2
	12	0	+1	0	−2	−2
Gigue	8	+2	+2	−1	−1	0
	10	+1	+1	0	−1	0
	12	0	0	−1	−1	0

Table 3. Fraction of 18 movement–k pairs with $\Delta_{23}(n) \neq 0$, config A, cap 300.

Axis	$n=1$	$n=2$	$n=3$	$n=4$	$n=5$
Seconds	17/18 (94%)	6/18 (33%)	0/18 (0%)	4/18 (22%)	7/18 (39%)
Beats	16/18 (89%)	11/18 (61%)	9/18 (50%)	18/18 (100%)	12/18 (67%)

the giant-component fraction at $n = 4$ stands out: $\delta_g = g^{(2)} - g^{(3)}$ is positive for all 18 movement–k pairs, perfectly k-invariant, and ranges from $+0.31$ to $+0.67$, confirming that the $p = 3$ tower fragments while the $p = 2$ tower remains connected. This is the most k-stable auxiliary signal in the entire grid, but it is concentrated at a single level under a single axis, and so does not resolve the delimitation identified through β_0.

The null-model controls ask whether Δ_{23} captures musical structure rather than a generic artifact of the tower. Fix a piece π. Using the contrasts Δ_{ab} defined in Section 2, we aggregate over $n \geq 2$ under config A by

$$\mathrm{MeanAbs}\Delta_{ab}(\pi, \mathrm{axis}) := \frac{1}{|\mathcal{I}|} \sum_{(n,k)\in\mathcal{I}} \left|\Delta_{ab}(\pi; n, k, \mathrm{axis})\right|,$$

$$S_{23}(\pi, \mathrm{axis}) := \frac{1}{|\mathcal{I}|} \sum_{(n,k)\in\mathcal{I}} \mathbf{1}\{\Delta_{23} \neq 0\},$$

where $\mathcal{I}$ ranges over ($n \in \{2, 3, 4, 5\}$, $k \in \{8, 10, 12\}$), giving $|\mathcal{I}| = 12$ cells per piece–axis pair.

To test whether Δ_{23} is sensitive to musical structure, we construct two null variants for each piece π: (i) *pitch-shuffle*: the chroma time series $H(t)$ is randomly permuted across time indices, destroying harmonic temporal correlations while preserving onset density; and (ii) *local time-shuffle*: the onset density $a(t)$ is permuted within non-overlapping windows ($\sim 5\%$ of the series length each), destroying rhythmic micro-structure while preserving global onset distribution and chroma sequence. We compute ($\mathrm{MeanAbs}\Delta_{23}, S_{23}$) for original, pitch-shuffled, and time-shuffled versions under both axes.

Table 4 reports the verified values. Under seconds, all variants yield $\mathrm{MeanAbs}\Delta_{23} = 0$ for both pieces, confirming that the seconds-axis signal at $n \geq 2$ is effectively zero. Under beats, both pieces show nonzero original signal ($\mathrm{MeanAbs}\Delta_{23} = 1.000$ for BWV 1049 and 1.167 for BWV 1050), concentrated at $n = 4$ and 5 where the window-length ratio $(3/2)^n \geq 5$. The time-shuffle eliminates this signal entirely ($\mathrm{MeanAbs}\Delta_{23} = 0$), consistent with the beats-axis signal being carried by local temporal correlations in onset patterns rather than by their global distribution. The pitch-shuffle, by contrast, leaves the signal essentially intact; for BWV 1050 it *increases* it (3.333), because random chroma reassignment amplifies pattern heterogeneity at long window scales. The local time-shuffle is the

Table 4. $\mathrm{MeanAbs}\Delta_{23}$ and S_{23} by piece, variant, and axis (config A, $n \geq 2$, $|\mathcal{I}| = 12$).

Piece	Variant	Axis	$\mathrm{MeanAbs}\Delta_{23}$	S_{23}
BWV 1049 mov. 1	original	seconds	0.000	0.000
BWV 1049 mov. 1	original	beats	1.000	0.500
BWV 1049 mov. 1	pitch-shuffle	seconds	0.000	0.000
BWV 1049 mov. 1	pitch-shuffle	beats	0.583	0.333
BWV 1049 mov. 1	time-shuffle	seconds	0.000	0.000
BWV 1049 mov. 1	time-shuffle	beats	0.000	0.000
BWV 1050 mov. 2	original	seconds	0.000	0.000
BWV 1050 mov. 2	original	beats	1.167	0.667
BWV 1050 mov. 2	pitch-shuffle	seconds	0.000	0.000
BWV 1050 mov. 2	pitch-shuffle	beats	3.333	0.750
BWV 1050 mov. 2	time-shuffle	seconds	0.000	0.000
BWV 1050 mov. 2	time-shuffle	beats	0.000	0.000

more informative null model: it targets the temporal micro-structure that drives pattern similarity, and its effectiveness identifies the beats-axis signal as a consequence of the

window-length asymmetry $(3/2)^n$ rather than of harmonic content. Thus, the strongest surviving signal is ordering-sensitive but not robustly prime-specific.
BWV 1007 movements are not included in Table 4 because they serve as the primary analysis corpus; the external pieces BWV 1049 and BWV 1050 provide independent null-model controls.

5. Discussion: delimitation and limitations

The complete BWV 1007 suite (Tables 1–3) reveals two distinct regimes. Under seconds binning, Δ_{23} concentrates at $n = 1$ and vanishes at $n = 3$, with re-emerging negative contrasts at $n \geq 4$ only for rhythmically dense movements (Courante, Gigue). Under beats binning, contrasts persist across all levels and culminate at $n = 4$ (100% nonzero, all negative). The magnitudes remain modest and $R_{\text{ctrl}} \leq 1$ under matched level ranges (Section 4). The null-model analysis (Table 4) shows that both external pieces produce nonzero beats-axis signal under the original protocol, that this signal is eliminated by local time-shuffle, and that it survives pitch-shuffle; this pattern is more consistent, under the tested null models, with short-range temporal correlations in onset patterns than with harmonic information alone.

These observations delimit the method. The invariant β_0 reveals coarse-grained differences between the $p = 2$ and $p = 3$ towers at individual levels, and the auxiliary observables (giant-component fraction at $n = 4$ beats, clustering at low n seconds) show sporadic patterns. However, no per-level summary satisfies all three desiderata of magnitude, sign stability, and k-invariance across the full movement–axis–level grid. This delimitation is the paper's central contribution: enriching the observable set is necessary before any affirmative claim about prime-specific structure can be made. Two directions follow: (a) inter-level compatibility diagnostics measuring whether maps between successive levels form a coherent inverse-system family, and (b) richer per-level observables (spectral gap, persistent homology in k).

The limitations follow directly from this delimitation. The arithmetic distinction of the ladder p^n does not imply optimality for a given musical phenomenon: the property is internal to the inverse-system formalism, while its relevance for a given repertoire is empirical and protocol-dependent. A structural confound is the exponential asymmetry of window lengths: at level n, the ratio $3^n/2^n = (3/2)^n$ grows from 1.5 ($n = 1$) to 7.6 ($n = 5$), so that at high n the $p = 3$ tower samples much longer temporal spans than the $p = 2$ tower. The large negative Δ_{23} values at $n \geq 4$ may therefore reflect window-length heterogeneity rather than intrinsic prime-specific musical structure; disentangling these effects requires future work with matched-length controls. Further limitations: kNN graphs densify with k; cap biases the pattern set; single-piece or single-cap results do not generalize; and the diagnostic probes robustness under a fixed protocol rather than asserting an ontological binary/ternary classification.

6. Conclusion

This paper introduced the prime-power indexed tower $\{D_{p,n}\}_n$, a reproducible experimental protocol, and a precise delimitation (contributions (i)–(iii) of Section 1). The six-movement BWV 1007 analysis shows that β_0, the auxiliary graph invariants, and the control framework yield level-localized differences but no signal that is simultaneously large, sign-stable, and k-invariant across the tested grid. The arithmetically motivated tower remains a viable analytical object and the protocol is internally consistent; the observable set, however, requires enrichment.

The natural next steps are replication across further pieces and performance traditions; richer per-level observables (spectral gap, persistent homology in k); matched-length window controls to disentangle window-size effects from structural asymmetries; and inter-level compatibility diagnostics measuring inverse-system coherence. A companion paper by the present author (Pérez-Buendía 2026) develops the inter-level compatibility direction: it introduces a profinite coherence diagnostic $\mathrm{Coh}_\pi(p,n)$ on the same arithmetic tower, establishes a null-floor identity $\mathrm{Coh}_\pi(p,n) = 1/p$ in the dense-coverage regime when the quantizer branching factor matches the prime, and reports experimental checks on the same BWV 1007 corpus together with extended Bach repertoire.

Acknowledgements

The author acknowledges support from the Supercomputing Center of CIMAT-Mérida "TOOLOK." The author also thanks the members of the P-ADAGIO research group for discussions.

Disclosure statement

The author reports there are no competing interests to declare.

Funding

This work was supported by SECIHTI, Mexico, under grant CF 2019/217367.

Data availability statement

The MIDI files used in this study (BWV 1007, BWV 1049, BWV 1050, synthetic toys) and the Python pipeline (`profinite_echo_midi.py`) are available in the project repository. Reproducibility manifests (Git HEAD, SHA256 hashes, and exact commands) are provided in Appendix A.

ORCID

Jesús Rogelio Pérez Buendía https://orcid.org/0000-0002-7739-4779

References

Balzano, Gerald J. 1980. "The Group-Theoretic Description of 12-Fold and Microtonal Pitch Systems." *Computer Music Journal* 4 (4): 66–84. https://www.jstor.org/stable/3679467.

Carlsson, Gunnar. 2009. "Topology and data." *Bulletin of the American Mathematical Society* 46 (2): 255–308. https://www.ams.org/journals/bull/2009-46-02/S0273-0979-09-01249-X/.

Edelsbrunner, Herbert, and John Harer. 2010. *Computational Topology: An Introduction*. Vol. 69 of *Mathematical Surveys and Monographs*. Providence, RI: American Mathematical Society. https://www.ams.org/books/mbk/069/.

Gouvêa, Fernando Q. 2020. *p-adic Numbers: An Introduction*. Universitext, 3rd ed. Cham: Springer. https://link.springer.com/book/10.1007/978-3-030-47295-5.

Heo, Eunwoo, Byeongchan Choi, Myung Ock Kim, Mai Lan Tran, and Jae-Hun Jung. 2025. "Persistent Homology of Music Network with Three Different Distances." *Journal of Mathematics and Music* 1–33. https://doi.org/10.1080/17459737.2025.2568844.

Mijangos, Martín, Alessandro Bravetti, and Pablo Padilla-Longoria. 2024. "Musical stylistic analysis: a study of intervallic transition graphs via persistent homology." *Journal of Mathematics and Music* 18 (1): 89–108. https://www.tandfonline.com/doi/abs/10.1080/17459737.2023.2232811.

Mrad, Dima, and Sara Najem. 2025. "Higher-Order Network Representation of J. S. Bach's Solo Violin Sonatas and Partitas: Topological and Geometrical Explorations." *arXiv preprint* https://arxiv.org/abs/2506.08540.

Müller, Meinard. 2015. *Fundamentals of Music Processing: Audio, Analysis, Algorithms, Applications*. Cham: Springer. https://link.springer.com/book/10.1007/978-3-319-21945-5.

Murtagh, Fionn. 2016. "Sparse p-adic data coding for computationally efficient and effective big data analytics." *p-Adic Numbers, Ultrametric Analysis and Applications* 8 (3): 236–247. https://link.springer.com/article/10.1134/S2070046616030055.

Pérez-Buendía, Jesús Rogelio. 2026. "p-adic Towers for Symbolic Music: Hierarchical Coherence and Null Floors." Submitted to *Computational and Applied Mathematics*; companion preprint archived at Zenodo, https://doi.org/10.5281/zenodo.19837367.

Raffel, Colin, and Daniel P. W. Ellis. 2014. "Intuitive Analysis, Creation and Manipulation of MIDI Data With pretty_midi." In *Proceedings of the 15th International Society for Music Information Retrieval Conference, Late-Breaking and Demo Papers (ISMIR-LBD),* Taipei, Taiwan, 84–93. https://colinraffel.com/publications/ismir2014intuitive.pdf.

Serre, Jean-Pierre. 1979. *Local Fields*. Vol. 67 of *Graduate Texts in Mathematics*. New York, NY: Springer. https://link.springer.com/book/10.1007/978-1-4757-5673-9.

Shapiro, Ilana, Ruanqianqian Huang, Zachary Novack, Cheng-i Wang, Hao-Wen Dong, Taylor Berg-Kirkpatrick, Shlomo Dubnov, and Sorin Lerner. 2025. "Synthesizing Composite Hierarchical Structure from Symbolic Music Corpora." In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence (IJCAI-25), Special Track on AI, the Arts and Creativity,* Montreal, Canada, 10153–10161. https://www.ijcai.org/proceedings/2025/1128.

Tymoczko, Dmitri. 2011. *A Geometry of Music: Harmony and Counterpoint in the Extended Common Practice*. Oxford Studies in Music Theory. New York, NY: Oxford University Press. https://global.oup.com/academic/product/a-geometry-of-music-9780195336672.

Appendix A. Reproducibility manifests and hashes

This appendix is included for editorial reproducibility; operational details are external to the argument of the Note. Full reproducibility materials are archived in the companion repository: a manifest records the Git HEAD hash, SHA256 checksums of source files, and the exact commands used to regenerate every table and figure (Table 4 included). No derived outputs are committed; all CSV tables and PNG figures are regenerated deterministically from the documented commands.