The Shape of Musical Possibility

The Shape of Musical Possibility

Cyclic Content, Ordered Traversal, and the Geometry of Twelve-Tone Space

Jason St George

GAMUT — Geometric-Algebraic Music Theory

March 2026

This volume collects and extends the three-part essay series *The Shape of Musical Possibility*, first published online at https://shapeofmusicalpossibility.org.

The companion formal papers, "Cyclic Autocorrelation, Permutation Fibers, and a Layered Symplectic Model for Pitch-Class Space," are available separately.

The interactive Layered Bundle Explorer is available at https://shapeofmusicalpossibility.org/visualizations/layered-bundle-explorer/

Contents

Preface

This volume gathers the three essays of *The Shape of Musical Possibility* into a single arc. The argument moves from a first question—why the twelve-tone universe resists inherited categories—through a layered construction that separates pitch content from ordered traversal, and finally into the phase-space language of symplectic geometry, where the atlas of musical objects acquires a natural law of motion.

The text is addressed to a mathematically curious reader who may or may not hold a degree in either music theory or mathematics. Where formalism is needed it is presented, but the emphasis throughout is on ideas, interpretations, and the reasons each mathematical turn was forced rather than merely chosen.

Alongside the essays, a companion suite of formal papers develops the full proof-oriented treatment, including rigorous statements of the layered Fourier embedding, the stratified symplectic quotient, and the Hamiltonian action of the dihedral and cyclic symmetry groups. These are available as separate downloads.

An interactive visualization—the *Layered Bundle Explorer*—lets the reader traverse the 223-node content graph, inspect local permutation fibers, and watch the guided tour animate a path through common Western harmonic structures. It is freely available on the project website.

J. St G., March 2026

List of Figures

Chapter 1

Why Music Wants a Geometry

1.1 The impossible atlas

A piano keyboard invites a flattering illusion. It suggests that musical possibility is laid out before us in a neat sequence, visible and finite, like a well-ordered street. There are twelve pitch classes in the chromatic system, and after a lifetime spent hearing scales, chords, cadences, arpeggios, and riffs, it is tempting to think that the space of musically meaningful objects is therefore large but manageable. One imagines a shelf of familiar categories: major and minor, dominant and diminished, modes, scales, rows, progressions. The imagination is wrong.

The moment one stops asking for the familiar and instead asks for the possible, the floor drops away. If we care only about pitch content and ignore octave placement, then every musical object is some subset of the twelve pitch classes. There are $2^{12} = 4096$ such subsets. That number is already large enough that no ordinary theoretical vocabulary can hold it comfortably. But content is only the first layer. If order matters, the numbers become much worse. There are 12! distinct ordered arrangements of all twelve pitch classes, which is 479,001,600 possible twelve-tone rows before one has said anything at all about rhythm, register, repetition, or emphasis. If one allows ordered collections of any cardinality from one note up to twelve distinct notes, the count rises past 1.3 billion. None of this is paradoxical. It is simply the cost of taking musical possibility seriously.

And yet human hearing does not meet that enormity as a spreadsheet. It meets it as shape. We do not merely catalog sounds; we feel neighborhoods, analogies, returns, tensions, shortcuts, symmetries, and singular landmarks. We hear that a major triad belongs somehow near a minor triad and far from a chromatic cluster. We feel that the whole-tone collection is less locally directional than the diatonic scale. We recognize that a diminished seventh chord has a peculiar self-similarity, as though one could rotate it and lose track of where it began. These are geometric intuitions before they are theoretical conclusions. They suggest that what music needs is not just a nomenclature, but an atlas.

The trouble is that inherited theory, brilliant as it is within its own historical corridor, was not designed to chart the whole territory. Functional harmony is a theory of directed motion in a particular repertorial world. Modal theory is a theory of scalar behavior under another set of stylistic assumptions. Serial theory widened the lens, but even there the practical focus often remained on special classes of ordered objects rather than on the total shape of the space they inhabit. Each of these theories sees something real. None of them, by itself, gives a satisfying answer to the simplest global question: what would it mean to picture the universe of twelve-tone pitch objects in a way a human mind could actually navigate?

To ask that question is already to move from grammar to geometry.

1.2 Folding the keyboard into a clock

The first act of compression is both severe and liberating. We stop treating pitches as isolated events in acoustic space and instead identify notes that are separated by octaves. C4, C5, and C6 cease to be distinct destinations and become three registrations of a single pitch class. The keyboard folds into a circle of twelve positions. This is not yet a theory of music; it is a decision about what counts as the same for the purpose of a first map.

Once that decision is made, the natural ambient space is $\mathbb{Z}_{12}$: twelve positions arranged not in a line but on a loop. That loop matters. A line has edges; a circle does not. On the line, B and C look far apart because the notational alphabet breaks there. On the circle of pitch classes, they are neighbors. The step from the keyboard to the clock is therefore not just a convenience. It is the first recognition that chromatic pitch organization is inherently cyclic.

Within this clock, a musical object can be represented simply by the positions it occupies. A major triad becomes the set $\{0, 4, 7\}$ after choosing some temporary zero. A whole-tone collection becomes $\{0, 2, 4, 6, 8, 10\}$. A diminished seventh chord becomes $\{0, 3, 6, 9\}$. One may then go a step further and identify sets related by transposition or inversion. A C-major triad and an E-flat-major triad become instances of the same abstract content. Under transposition/inversion equivalence, even larger families collapse together. This is not an evasion of difference; it is a principled folding along symmetries that the first approximation does not need to distinguish.

That reduction can feel like loss, but it is better understood as a quotient. In mathematics, a quotient is what happens when we identify points that differ only by a symmetry we no longer wish to track. Rotating the pitch-class clock does not change the intervallic shape of a set. Reflecting it does not change its inversional profile. So the space is folded along those motions. What emerges is not a smooth picture in the ordinary sense, but a compressed one. Many superficially different musical objects collapse to a single abstract type.

This already explains something that ordinary musical language often leaves mysterious. Why do all major triads feel like instances of one thing? Why do major and minor triads feel related at a deeper level than either does to a tone cluster? Why does the diminished seventh chord seem unusually slippery and self-similar? The answer is that some objects have more symmetry than others. Most pitch-class sets have the full expected collection of distinct transpositions and inversions, but symmetrical sets have fewer distinct forms because nontrivial transpositions or inversions map them back onto themselves. Symmetry is not an ornamental property of the space. It is one of the forces that shapes it.

At this stage, one might think the problem is nearly solved. There are only finitely many set classes. They can be named, listed, and grouped by cardinality. Twentieth-century music theory did exactly this with admirable rigor. Allen Forte's classification assigns each transposition-inversion equivalence class a canonical label of the form k-n, where k is the cardinality and n is the ordinal within that cardinality group. Under this system the 4096 subsets of the chromatic aggregate collapse to exactly 223 distinct set classes. Many of these classes correspond to objects that every musician already knows by ear, even if not by number: 3-11 is the major/minor triad, 4-27 is the dominant/half-diminished seventh chord, 7-35 is the diatonic scale. Others—5-Z12, 6-Z44, 8-22—lack familiar common names but occupy no less definite a position in the pitch-class universe. The Forte catalog is, in effect, the periodic table of twelve-tone content: compact, exhaustive, and indispensable as a naming scheme.

But a list is not yet a geometry. A classification tells us which objects are equivalent. It does not yet tell us what should count as near.

1.3 Why naming is not enough

Imagine being handed a periodic table without any indication of adjacency, valence, or structure. It would still be useful. One could name things. One could sort them. But one could not yet think spatially about chemical behavior. Something similar happens in musical set theory. Knowing that the major/minor triad is Forte 3-11 and that the whole-tone scale is Forte

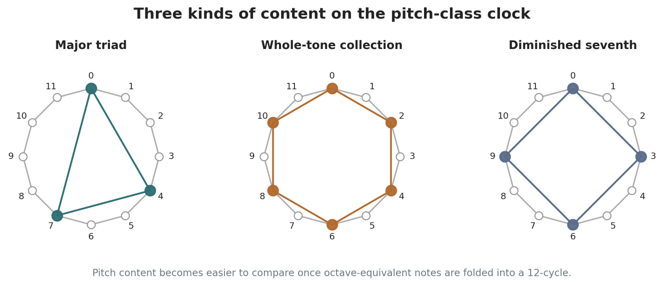

Figure 1.1: Pitch-class clock examples showing several collections around a twelve-point circle.

6-35 does not, by itself, tell us how the two are related, how far apart they sit, or what bridges connect them. Naming a set class is not the same as understanding its position in the larger terrain.

There are several reasons for this. The first is psychological. Human beings do not think well in exhaustive tables once the table becomes large. We think better when we can recognize regions, bridges, bottlenecks, and special cases. The second is musical. The reason theorists care about content classes at all is not because naming is satisfying for its own sake, but because similar content tends to support similar sonic affordances. The third is conceptual. The moment we begin talking about similarity, we have entered the domain of metrics, neighborhoods, and embeddings whether we say so or not.

Traditional pitch-class set theory offers one important bridge from naming to comparison: interval content. One may summarize a set by counting how many intervals of each size occur within it. This gives a compact signature of the set's internal makeup, and it often captures a great deal of what listeners mean when they speak of a collection's color or profile. But even here one senses the deeper question pressing in. A summary is not yet a location. It tells us how many intervals occur, but not how the set is arranged around the cycle. It is useful precisely because it forgets something. The real question is whether there is a more fundamental operation beneath this summary: something that does not merely tabulate intervals, but arises naturally from the cyclic structure of the pitch-class clock itself.

This is where cyclic autocorrelation enters the scene.

1.4 Duncan's fingerprint

Andrew Duncan's work on combinatorial music theory is especially suggestive because it does not begin with the language of harmonic function or stylistic usage. It begins with a periodic object and asks how that object overlaps with

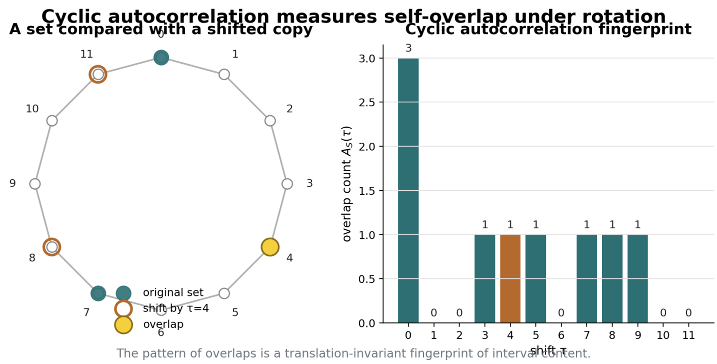

Figure 1.2: Cyclic autocorrelation: a pitch-class set is rotated against itself and the resulting overlap profile serves as a content fingerprint.

itself under cyclic shift. The move is mathematically simple and musically deep. One treats a pitch-class set as a pattern on a twelve-point circle and measures what survives when the pattern is rotated against itself.

Take a pitch-class set $S \subset \mathbb{Z}_{12}$. Define its indicator function $\chi_S(n)$ to be 1 when $n \in S$ and 0 otherwise. Then define the cyclic autocorrelation of S by

$$A_S(\tau) \;=\; \sum_{n \in \mathbb{Z}_{12}} \chi_S(n)\, \chi_S(n+\tau).$$

In plain English, this asks: if I shift the set around the pitch-class circle by τ semitones, how many points does it share with its original position? The answer is a twelve-entry cyclic fingerprint. At lag zero, the set overlaps with itself completely, so $A_S(0)$ is just the number of notes in the set. At other lags, the entries count how often intervals of that size occur inside the set. Because the space is cyclic, this overlap-counting procedure is perfectly adapted to the object being studied. And because it counts overlap after rotation, it is automatically insensitive to where the set begins. That is exactly what a good content fingerprint should do.

One can see immediately why this is more than a fancy reformulation of interval counting. The operation is geometric in spirit. It does not ask first for a label. It asks how a shape behaves when moved around its native space. A major triad (3-11), a whole-tone collection (6-35), and a diminished seventh chord (4-28) all leave different overlap traces when shifted around the twelve-point cycle. Those traces are not arbitrary statistics pasted onto the object after the fact. They are responses of the object to motion.

This is the first truly persuasive reason to say that music wants a geometry. The relevant invariants arise not from a taxonomy imposed from outside, but from an operation internal to the space itself.

1.5 The Fourier shadow

There is a second reason this idea matters, and it is easy to miss if one focuses only on the counting interpretation. Cyclic autocorrelation has a spectral twin. By the finite Wiener–Khinchin correspondence, autocorrelation and the squared magnitude of the discrete Fourier transform carry the same information. One may read the structure either in the domain of overlap or in the domain of modes.

For music, this means that the same pitch-class set can be understood in two coordinate languages. In one language, we count how the set overlaps with itself under cyclic shift. In the other, we measure how strongly different Fourier modes are present in the set's characteristic function. What sounds like two theories is really one theory seen from opposite sides.

That equivalence matters because it explains why interval content and spectral shape keep reappearing together in modern music theory. It also plants a seed for everything that comes later. If the cyclic fingerprint of content already has a Fourier shadow, then musical geometry is not merely combinatorial. It is already brushing against a phase space. The richer mathematics to come are not imported artificially; they are latent from the start.

At the level of this first essay, the point is simpler. Cyclic autocorrelation gives us the first serious coordinate system for content. It is musically meaningful, mathematically natural, and adapted to the cyclic world in which pitch classes live.

And yet it is not enough.

1.6 Slonimsky and the pressure of order

Before explaining the incompleteness of content-based fingerprints, it helps to notice another pressure that has been building all along. Musicians do not care only about membership. They care about traversal.

This is one reason Nicolas Slonimsky's *Thesaurus of Scales and Melodic Patterns* remains so revealing. However eccentric the book may appear at first glance, its governing intuition is not eccentric at all. It organizes patterns by interval cycles and then generates richer patterns by interpolation, insertion, and ornamentation. Slonimsky is not merely cataloging pitch collections. He is cataloging ways of moving through them.

That historical fact matters because it uncovers a structural distinction. If the pitch-class set answers the question "what notes are present?," the ordered pattern answers the question "how are they traversed?" These are not rival descriptions. They are different layers. A scale may be treated as content when we abstract away its order, but as pattern when we care about its steps. A chord may be treated as content when we ask about interval makeup, but as arpeggiated motion when we care about voice-leading or melodic realization. Slonimsky's enduring value lies in the fact that he forces the second question back into view.

The consequence is unavoidable. Any geometry of musical possibility that captures content but not traversal will be useful, but partial. It will tell us what is there, not how it behaves from within.

1.7 Where the first fingerprint fails

The incompleteness of content-based fingerprints is not merely philosophical. It is mathematically explicit. Autocorrelation records interval content. Interval vectors record interval content. The magnitudes of the discrete Fourier coefficients of a pitch-class set also record, in another guise, interval content. But there exist pitch-class sets that share this content without being equivalent under transposition or inversion. In music theory these are called Z-related sets—the "Z" in Forte numbers like 5-Z12 and 5-Z36 signals exactly this phenomenon—and in broader mathematical language they are homometric sets.

The significance of these examples is easy to understate. Two genuinely different sets can cast the same intervallic shadow. Or, in spectral language, two different objects can have the same Fourier magnitudes. Something positional has been lost. When we count pairwise spacings, we remember how many intervals occur but forget enough of their arrangement that distinct objects can collapse onto the same signature.

This is not a technical nuisance. It is the turning point of the whole project. It tells us that interval content, though powerful, cannot be the whole story. The first layer of the map succeeds and fails at once. It succeeds because it gives us a natural, transposition-invariant coordinate system for content. It fails because musical objects are not exhausted by their interval counts. Content has shape beyond content, and pattern has shape beyond membership.

1.8 Why the object must be layered

The lesson of the last two sections is subtle but decisive. Duncan shows that cyclic autocorrelation gives a mathematically natural fingerprint for pitch content. Slonimsky shows, from another direction, that ordered traversal is structurally real and musically irreducible. Z-relatedness shows that even the best content fingerprint is incomplete. Taken together, these facts imply that the space we want cannot be flat.

There is a base layer consisting of pitch content up to symmetry, and there are further layers that preserve information discarded by that base. Order is one such layer. Positional or phase-sensitive information is another. The right global object, then, will not be a single uniform cloud of points. It will be a layered construction in which one level records what is present and another records how that presence is internally organized or traversed.

This layered view is not only mathematically cleaner; it is closer to musical experience. When we recognize a sonority, we recognize its content. When we recognize an arpeggiation, a scalar figure, or a melodic cell, we recognize an ordering of that content. When we distinguish two homometric objects that "contain the same intervals" yet do not sound or behave the same, we are responding to structure that survives beyond unordered interval counts. In each case, perception itself is already stratified.

A geometry adequate to music must therefore do two things at once. It must compress by symmetry, because otherwise the raw combinatorics are unthinkable. But it must also preserve enough layered distinction that compression does not become erasure. This is why the project does not end with a set-class table or an interval vector. Those are indispensable instruments of reduction, but they are not the final image of the space.

1.9 Threshold

We can now say what this first essay has accomplished. It has argued that the universe of twelve-tone musical objects is too large to be understood as a mere list. It has shown that the first necessary simplification is to move from pitches to pitch classes and then to quotient by transposition and inversion. It has identified cyclic autocorrelation as the first genuinely geometric fingerprint of pitch content: a way of reading a set by its self-overlap under cyclic motion. It has noted the Fourier shadow of that fingerprint and the deeper spectral language it implies. It has used Slonimsky to show that order is not an optional afterthought but an independent axis of structure. And it has used homometric and Z-related examples to prove that content alone, however elegantly summarized, cannot be the whole map.

What we possess, then, is not yet the geometry of musical possibility, but its first coordinate chart. We know how to speak about content as a cyclic object. We know why symmetry forces us to fold the space. We know why the result is informative but incomplete.

The next step is to attach the missing structure rather than merely lament its absence. Content gives us the answer to the question of what notes are present. What remains is to describe the local worlds of order that hang above each content class, and to show how those worlds can be assembled into a single intelligible object.

Content gives us a map of what is present, but not yet of how a musical object internally unfolds.

Source notes and further reading

1. Allen Forte, *The Structure of Atonal Music* (Yale University Press, 1973).
2. John Rahn, *Basic Atonal Theory* (Longman, 1980).
3. Andrew Duncan, "Combinatorial Music Theory," *Journal of the Audio Engineering Society* 39, no. 6 (1991): 427–448.
4. Ian Quinn, "General Equal-Tempered Harmony," *Journal of Music Theory* 50/2 (2006).
5. Emmanuel Amiot, *Music Through Fourier Space* (Springer, 2016).
6. Jason Yust, *Organized Time: Rhythm, Tonality, and Form* (Oxford University Press, 2018).
7. Nicolas Slonimsky, *Thesaurus of Scales and Melodic Patterns* (Schirmer, 1947).

Chapter 2

Attaching Order to Content

2.1 The second question

Suppose two musicians are given exactly the same five notes. One arranges them as a smooth ascending scalar figure. The other leaps through them in alternating thirds, breaking the line into a lattice of skips and returns. In the most literal sense, the two passages contain the same material. In the most immediate musical sense, they do not inhabit the same world.

This is the second question that emerges the moment the first one is answered. Part I asked what notes are present and showed that pitch content can be organized by cyclic fingerprints on the pitch-class circle. But once content has been reduced, another form of structure appears in relief. A set tells us membership. It does not tell us itinerary. It tells us what belongs to the object. It does not yet tell us how the object moves through itself.

The distinction is familiar in practice even when it is not named. A scale is not only a collection of scale degrees; it is also a way of passing from one degree to another. A chord is not only an unordered sonority; it can also be a broken figure, a voicing pattern, an arpeggiated route, a melodic cell distributed over time. Even the same twelve-tone row can become palpably different when its order is segmented, cycled, inverted in part, or rethreaded through a different registral or rhythmic design. Once one notices this, the inadequacy of a content-only geometry becomes impossible to ignore.

The problem is not that content theory is wrong. The problem is that it is only one layer. To add order honestly, however, one must avoid a common mistake. One must not simply throw all orderings back into the same pot and call the result a richer space. That move would destroy the very compression that made the first map intelligible. The real task is subtler: content must remain the base layer, while order is attached locally as a further structure above each content point.

2.2 Why order does not belong in the base layer

It is tempting to think that the easiest solution is simply to replace each pitch-class set with the list of its notes in order. But that would confuse two different kinds of identity. Content answers a symmetry-reduced question: which pitch classes are present up to transposition and inversion? Order answers a different question: in what sequence or cyclic traversal is that content realized? If both are collapsed into a single undifferentiated coordinate system, the first layer loses its explanatory power. We no longer see what different orderings have in common, and we no longer know which differences belong to content and which belong to traversal.

A useful analogy is a city map. One layer gives the streets and blocks. Another gives the routes travelers may take. No one mistakes the street grid for the collection of all possible itineraries; nor would anyone understand traffic by erasing the grid and keeping only route histories. The route space must be attached to the map, not substituted for it.

Musical order behaves in the same way. There is first a content point: some abstract pitch-class object after the relevant symmetries have been folded out. Above that point sit many possible ordered realizations. Some are almost trivial neighbors, differing only by a local swap. Others are remote from one another and require a much larger rethreading of the pattern. The right question is therefore not “What is the order of this set?” but “What is the local order-space associated with this content class?” That wording already hints at the geometry to come.

2.3 Rooting the cycle

Before building that local order-space, one must remove a final nuisance symmetry. Many ordered patterns are cyclic rather than linear: one can start reading them at different points on the cycle and still encounter the same circular object. If this is ignored, one and the same pattern appears many times for no musically substantial reason.

The standard cure is to root the cycle. Choose a distinguished starting point—often the smallest pitch class after normalization, or some other chosen anchor—and read the cycle from there. Once a root is fixed, cyclic rotation is no longer treated as a meaningful difference. The gain is not merely clerical. Rooting turns a redundant family of circular readings into a manageable set of orderings whose remaining differences are musically substantial.

This is the first place where the geometry of the fiber becomes visible. We are not cataloging arbitrary strings of notes. We are studying rooted cyclic traversals of a fixed content object. A rooted traversal remembers enough order to matter while modding out the trivial equivalence of “the

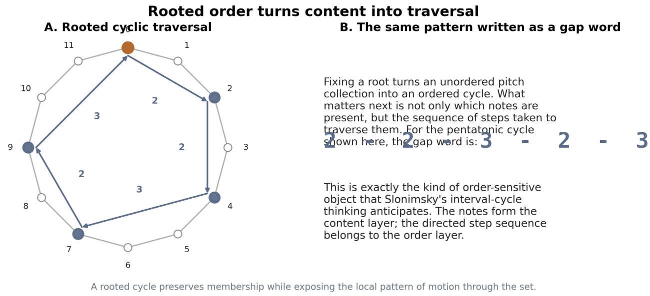

Figure 2.1: A rooted cyclic ordering of five pitch classes and the resulting gap word of successive intervals.

same cycle read from a different starting point." In practical terms, the fiber over a content class becomes smaller, cleaner, and more interpretable.

2.4 Gap words and interval cycles

Once a cyclic ordering is rooted, it acquires an especially revealing summary: the cyclic gap word. Read around the ordered pattern and record the successive pitch-class intervals from each note to the next, taken modulo 12. The result is a circular word of step sizes. If the pattern is $[p_0, p_1, \ldots, p_{k-1}]$, then its gap word records $p_1 - p_0$, $p_2 - p_1$, and so on, returning at the end to $p_0 - p_{k-1}$ modulo 12.

This is where Slonimsky suddenly becomes mathematically crisp. His interval cycles are not vague stylistic habits but periodic gap structures. A pattern built on repeated minor thirds corresponds to a recurrent gap word dominated by the value 3. A whole-tone cycle gives repeated 2s. Interpolated patterns, infra-polations, and other ornamented constructions appear as deformations of a simpler cyclic backbone: the cycle remains legible, but the gap word acquires extra local variation.

The advantage of this description is that it preserves order without abandoning cyclic structure. A set says what notes are occupied on the clock. A gap word says how one travels from occupied point to occupied point. The first is static membership. The second is organized motion. And because the gap word is itself a cyclic object, it is ready for the same kinds of combinatorial and Fourier analysis that proved fruitful for content.

At this point one can already sense the two-layer picture sharpening. Content classes form the lower geography. Gap structures, rooted traversals, and their deformations form a local upper geography attached to each lower point.

2.5 The local geometry of permutations

Orderings do not merely differ; they differ by more or less. That simple observation is what turns the local order-space from a bag into a geometry. If one ordering can be reached from another by a single adjacent swap, they feel close. If it takes many such swaps to pass from one to the other, they feel more distant. This is the intuition behind permutation metrics such as Kendall tau, which counts how many adjacent exchanges are needed to transform one ordering into another.

The importance of such a metric is philosophical as much as technical. It says that local differences in order are not all equal. A tiny rerouting of a pattern should not be treated as remote from the original. Conversely, a radical rethreading should not be plotted as a near neighbor merely because the underlying content is unchanged. The order-space needs its own notion of small displacement.

Combinatorially, the picture is elegant. The set of all orderings of a fixed list can be organized by adjacent swaps; each ordering is a vertex, and an edge is drawn when one adjacent transposition converts one ordering into another. In the background lurks the permutohedral idea: permutations are not only countable; they have a natural neighborhood graph. For musical purposes, one does not need to insist on the whole polytope in full generality. What matters is the insight it offers. Orderings can be navigated locally.

This gives the fiber its first honest geometry. It is no longer merely the set of all possible traversals. It is a space in which one ordering sits near another for concrete, musically intelligible reasons.

2.6 The fiber over a content class

We can now state the structural move that Part I was pointing toward. Let c be a content class—some symmetry-reduced pitch object in the base space. The fiber over c is the set of admissible rooted cyclic orderings of that content. In schematic form one may write

$$F_c = \{\text{rooted cyclic orderings of the content class } c\}.$$

The sentence matters more than the symbolism. Each point in the base does not stand alone. It carries above it a local fringe of internal possibilities: ways of reading, traversing, and articulating that same underlying content. Two content classes may be near in the base because they share similar interval fingerprints. Two orderings in the same fiber may be near because one is a small rerouting of the other. These are different kinds of nearness, and the total space must preserve both.

This is why the emerging object begins to resemble a bundle. In a fiber bundle, one has a base space together with a fiber attached to each base point, and the total object is built by assembling those local fibers over the whole base. The musical construction is not a fiber bundle in every strict differential-topological detail from the outset, because singular symmetries and changing cardinalities complicate the story. But as a guiding architecture, the bundle picture is exactly right. Content lies below. Order hangs above. The total space is the union of these local attachments.

This bundle language does more than provide a pretty metaphor. It disciplines the theory. It prevents us from confusing similarity of content with similarity of order. It lets us talk about motion inside a fiber separately from motion along the base. And it prepares the way for a later dynamics in which both kinds of displacement matter but need not be measured by the same ruler.

2.7 Measuring likeness within the fiber

A local order-space still needs coordinates robust enough to support visualization. Adjacent-swap distance provides one piece of the metric, but it is not the whole story. Two orderings might lie the same number of swaps away from a reference ordering while producing very different step patterns. The fiber therefore benefits from an order-sensitive feature map built from the pattern itself.

The simplest such features are summaries of the gap word: a histogram of step sizes, counts of successive step pairs, perhaps an entropy-like measure of how regular or irregular the traversal is. These do not replace the ordering; they describe aspects of it that remain musically legible. A pattern with repeated equal steps, a pattern that alternates long and short intervals, and a pattern that meanders irregularly should not collapse into one local neighborhood merely because they contain the same notes.

A practical local metric therefore blends two ingredients: the cost of permuting one ordering into another and the difference between their order-sensitive summaries. In schematic form one may write

$$d_F(\pi, \rho)^2 \; = \; \alpha \, d_K(\pi, \rho)^2 \; + \; \beta \, \|f(\pi) - f(\rho)\|^2.$$

Here d_K is a permutation distance such as Kendall tau, and f records whatever order-sensitive features one has chosen to preserve. The specific weights α and β are not sacred. What matters is the principle. The geometry of the fiber should preserve musically meaningful nearness rather than merely counting reorderings in the abstract.

This is one of the decisive interpretive gains of the layered model. It gives an explicit answer to the question of why some orderings of the same content feel like variants of one another while others feel like genuinely different pattern-worlds. They inhabit the same fiber but not the same neighborhood.

2.8 From exact combinatorics to visible geometry

At this point the exact object is already present in principle: a base of content classes, fibers of rooted orderings, and metrics on both levels. But human beings do not think directly in high-dimensional feature spaces or large neighborhood graphs. To make the space visible, one needs a controlled projection.

This is where methods such as multidimensional scaling and diffusion maps enter. Their role is not to define the geometry but to render it. One first builds the exact or approximate distance data that expresses musical nearness. One then asks for a low-dimensional configuration of points whose visible distances preserve that deeper structure as faithfully as possible. The 3D and 4D bundle renderings arise at precisely this stage: they are shadows of a richer exact object, not the object itself.

This distinction between model and rendering is essential. A plot can be seductive and misleading at once. A cluster that looks important may be only a projection artifact. A loop that appears smooth in two dimensions may hide combinatorial singularities. The right discipline is therefore to treat the graph, metric, and quotient relations as primary, and the visible embedding as explanatory. The rendering is a map of the geometry, not a substitute for it.

Even so, the visual gain is enormous. A 3D tangent-plane attachment lets the eye see each content point surrounded by a small local halo of orderings. A 4D direct-sum construction, when rotated and projected, separates base motion from fiber motion even more cleanly. The first is often easier to read in a static image. The second is often closer to the mathematics. Together they teach the same lesson: content and order belong together, but they do not live in the same coordinates.

2.9 Where symmetry pinches the bundle

No honest account of the total object can ignore its singular points. Some content classes have exceptional symmetry. The diminished seventh chord (4-28), for example, folds onto itself under nontrivial transpositions; the whole-tone collection (6-35) does something similar in its own way. These are not generic points in the base. They are hubs where the quotient has extra isotropy, and the local geometry behaves differently.

The same phenomenon reappears in the fibers. Certain orderings are stabilized by residual symmetries or by repeated local structures that make

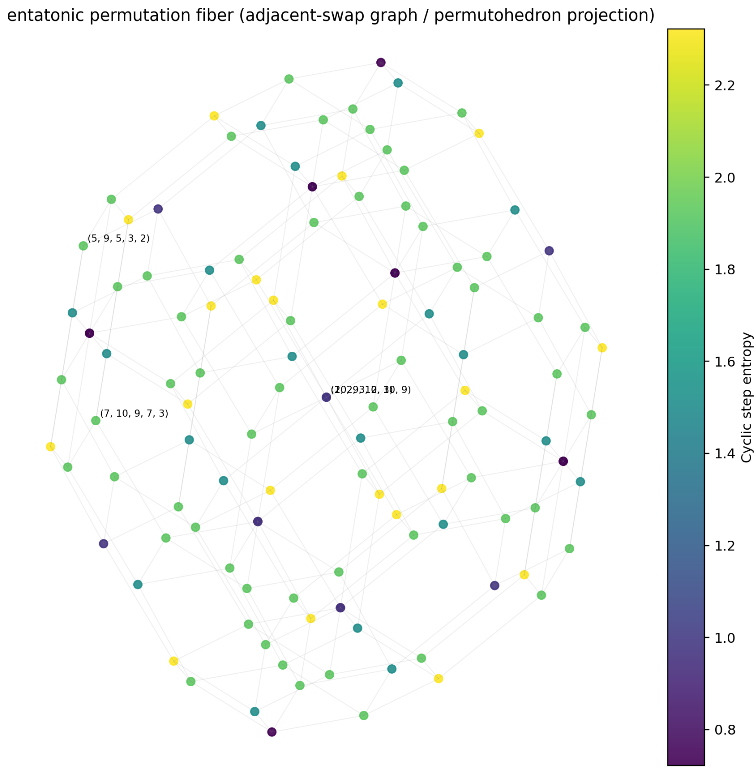

Figure 2.2: The pentatonic permutation fiber: each vertex is a distinct rooted cyclic ordering and each edge connects orderings differing by a single adjacent swap.

distinct-looking readings partially equivalent. In such places the fiber does not behave like a generic smooth fringe. It pinches, folds, or acquires reduced local dimension. What looked like a bundle in the broad sense reveals itself, more precisely, as a stratified bundle-like object with singular orbit types.

This is not a flaw in the model. It is one of the things the model is trying to explain. Symmetry-rich musical objects really do behave like landmarks. They gather pathways. They collapse distinctions. They serve as points of unusual equivalence in the space. Any geometry that forced them to look generic would be less truthful than one that lets the singularities remain visible.

2.10 A space of positions is not yet a space of motion

By the end of this second stage, the map is far richer than the one we began with. We have a content base shaped by cyclic fingerprints and quotient symmetries. We have local fibers of rooted orderings organized by permutation and pattern metrics. We have a total object that can be approximated visually as a bundle-like 3D or 4D construction. And we have a principled account of why singular landmarks arise where symmetry becomes unusually strong.

But however satisfying that architecture may be, it is still static. It tells us where musical objects live and how they cluster. It does not yet tell us what lawful motion through the space should mean. A shortest path in a metric space is not the same thing as a musically privileged flow. A local rerouting inside a fiber is not yet a dynamics. The map is now much more than a catalog. It is not yet a mechanics.

That final turn requires a new language. The moment the layered object is lifted into Fourier coordinates, phase begins to matter alongside magnitude, and the geometry begins to resemble a true phase space rather than a mere arrangement of points. Only then does the theory acquire the resources to speak not just about similarity and adjacency, but about conserved structure, Hamiltonian-type symmetries, and lawful motion on a stratified musical space.

Once content and order are separated and reattached, the remaining question is no longer where musical objects live, but how they move.

2.11 Computational anatomy of the fiber space

The fiber construction described above is not merely theoretical. The accompanying Layered Bundle Explorer computes exact or sampled fibers for every content class with three or more pitch classes under transposition-inversion equivalence. The following table summarizes the combinatorial landscape.

Three observations deserve emphasis. First, the combinatorial explosion is dramatic: from 2 orderings at cardinality 3 to tens of millions at cardinality 12. This is why naively plotting all orderings flat would produce an illegible cloud. Second, every content class of cardinality k has the same rooted-fiber count $(k-1)!$, because fixing a root eliminates translational symmetry. The diminished seventh chord (cardinality 4, symmetry order 4) thus has the same $6 = 3!$ rooted orderings as any other tetrachord. Symmetry enters only at a further level of quotient: fewer of those orderings give globally inequivalent patterns for a highly symmetric set. The whole-tone collection (cardinality 6,

Cardinality	TI-classes	Orderings per fiber	Sampling	Total orderings
1	1	1	—	—
2	6	1	—	—
3	12	2	exhaustive	24
4	29	6	exhaustive	174
5	38	24	exhaustive	912
6	50	96	sampled from 120	4,800
7	38	96	sampled from 720	3,648
8	29	96	sampled from 5,040	2,784
9	12	96	sampled from 40,320	1,152
10	6	96	sampled from 362,880	576
11	1	96	sampled from 3,628,800	96
12	1	96	sampled from 39,916,800	96

Table 2.1: Fiber counts by cardinality. For every content class of cardinality k, the exact number of rooted cyclic orderings is $(k-1)!$, independent of the symmetry order of the T/I-class (fixing a root eliminates translational symmetry).

symmetry order 6) has all $120 = 5!$ rooted orderings, but only 20 remain distinct modulo its sixfold translational symmetry. Third, even the sampled fibers preserve meaningful structure: adjacent-swap edges connect orderings that differ by a single neighbor transposition, so the local topology of the permutohedron is visible at every scale.

The full dataset contains 223 content classes, 1,051 add/remove edges, and 216 computed fibers (the seven classes of cardinality $k \leq 2$ have trivial single-element fibers and are excluded from nontrivial fiber computation) with a total of 14,166 orderings and their associated swap-edge graphs. Fiber node colors in the visualization encode step-word entropy, providing an immediate visual cue about whether a given ordering distributes its intervals evenly (smooth, blue) or irregularly (jagged, amber).

2.11.1 Bridging abstract and familiar: Forte numbers and common names

Each of the 223 content classes is identified by its Forte number—the standard label introduced in Part I (e.g. 3-11, 7-35). But Forte numbers, while precise, can feel opaque to musicians trained in tonal or jazz vocabulary. The explorer therefore supplements each Forte number with a common Western name where one exists. Selecting a node reveals, for instance, that 3-11 is the major/minor triad, 4-27 is the dominant/half-diminished seventh chord, 5-35 is the pentatonic scale, and 7-35 is the diatonic collection. Where no single familiar name applies—as with many of the more exotic pentachords and hexachords—the Forte number stands alone, serving as a neutral coordinate rather than a cultural claim.

This double labeling bridges two modes of understanding that the geometry deliberately holds in tension. The Forte number locates a set class within the abstract lattice of transposition-inversion equivalence: a position in the 223-node graph defined purely by pitch-class content and single-note add/remove adjacency. The common name, where it exists, anchors that same point in the lived practice of Western harmony. Neither description reduces to the other.

The mapping also makes visible a striking asymmetry. Of the 223 set classes, only about fifty carry widely recognized names—the triads, seventh chords, common scales, augmented-sixth sonorities, and a handful of iconic aggregates. The remaining majority are nameless in traditional pedagogy, not because they are musically inert, but because the tonal and post-tonal traditions that generated our vocabulary explored only a fraction of the combinatorial space. The graph makes this gap visible: named nodes cluster in a dense core of cardinalities 3 through 7, while the periphery at high and low cardinalities is populated almost entirely by Forte-number-only classes. This is itself a kind of data about the history of Western musical thought.

Source notes and further reading

1. Nicolas Slonimsky, *Thesaurus of Scales and Melodic Patterns* (Schirmer, 1947).

2. Robert Morris, *Composition with Pitch-Classes* (Yale University Press, 1987).

3. Dmitri Tymoczko, *A Geometry of Music* (Oxford University Press, 2011).

4. Clifton Callender, Ian Quinn, and Dmitri Tymoczko, "Generalized Voice-Leading Spaces," *Science* 320, no. 5874 (2008): 346–348.

5. Art Samplaski, "Mapping the Geometries of Pitch-Class Set Similarity Measures via Multidimensional Scaling," *Music Theory Online* 11, no. 2 (2005).

Chapter 3

Why the Space Wants a Symplectic Form

3.1 A map is not yet a mechanics

By the end of Part II, the space we had built was already far richer than the one with which we began. Content no longer floated as a mere list of set classes. It occupied a base shaped by cyclic fingerprints and quotient symmetries. Order no longer hid anonymously inside familiar musical categories. It appeared as a local fiber of rooted cyclic traversals attached to each content point. The whole object could be rendered, at least approximately, as a 3D tangent-plane bundle or a 4D direct-sum construction. One could see families, corridors, neighborhoods, and singular landmarks.

And yet something essential was still missing.

> *A map is not yet a mechanics. One may have a beautiful atlas of musical objects and still no law of motion.*

A map is not yet a mechanics. One may have a beautiful atlas of locations without any account of how motion through that atlas should work. A city map does not, by itself, tell us traffic laws, currents of use, or the forces that make some paths stable and others unstable. Likewise, a geometric arrangement of musical objects does not yet tell us what should count as lawful transformation. It can tell us that two objects are near, but not whether there is a principled way of moving from one to the other. It can tell us that a region is densely connected, but not whether a flow through that region preserves anything worth calling musical structure.

This distinction matters because music is not merely a population of objects. It is a traffic of transformations. A major triad bends toward a minor triad. A diatonic collection leans toward a whole-tone haze. A pattern built from an equal division of the octave acquires interpolated notes and returns altered but recognizable. A row is rethreaded, a cycle inflected, a region revisited. To hear music is not only to hear what is present, but to hear how presence changes.

Metric geometry takes us part of the way. It gives us distance, neighborhood, and local comparison. That is already a great advance over bare classification. But a metric answers the question "how far?" more readily than it answers the question "under what rule?" A shortest path in a similarity space is not automatically a musically privileged path. A visible cluster in a low-dimensional embedding is not, by itself, a dynamics. The next step therefore requires a change in language. We need a geometry in which motion is not an afterthought appended to a static map, but something built into the structure from the start.

The surprise is that this next language has been quietly waiting for us from the beginning. Part I introduced cyclic autocorrelation as a fingerprint of pitch content. Part II separated content from order and then reattached them as base and fiber. Both steps already pointed toward Fourier analysis. And Fourier analysis, once written honestly in complex coordinates, leads very naturally toward phase space.

That is the real subject of Part III: not the importation of an alien mathematics into music, but the recognition that the layered musical object we have been building already wants to be written in a language of amplitudes and phases, and that this language carries with it a natural symplectic form.

3.2 From overlap counts to phase

The hinge of the argument is easy to state.

Cyclic autocorrelation tells us how often a pitch-class set overlaps with itself under cyclic shift. That was Duncan's crucial insight in the first essay. For a set of pitch classes, the resulting overlap profile is a highly economical summary of interval content. One can hear why it matters almost immediately: a diminished seventh (Forte 4-28), a whole-tone collection (6-35), and a major triad (3-11) all leave different traces when they are rotated against themselves. The trace is not an arbitrary statistic glued onto the object from outside. It is the object's response to motion in the cyclic space where it already lives.

But autocorrelation also has a second life. It is equivalent, by the finite Wiener–Khinchin principle, to the squared magnitudes of the discrete Fourier coefficients of the same object. The first description speaks in the language of overlap. The second speaks in the language of spectral content. These are not rival theories. They are two coordinate systems for the same invariant information.

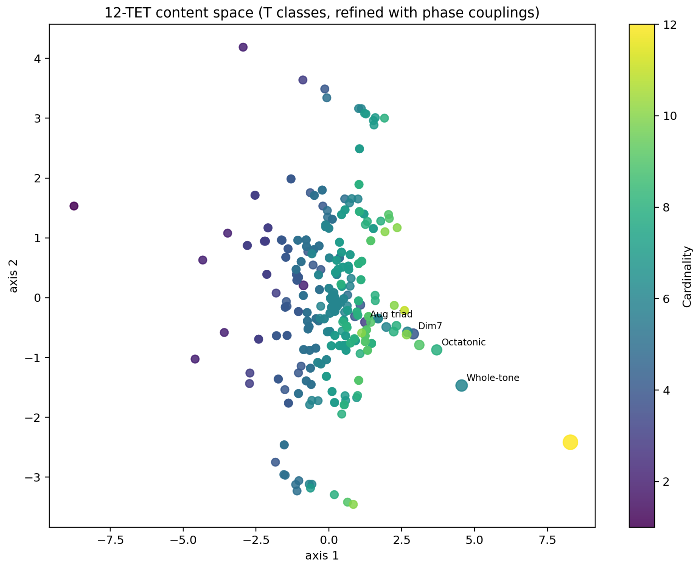

Figure 3.1: Refined spectral content space—a low-dimensional rendering of the content layer after phase-sensitive refinement.

That equivalence is more than a technical convenience. It tells us what the first fingerprint remembers and what it forgets. Fourier magnitudes preserve how strongly a set responds to different periodicities around the twelve-point cycle. A strong response in one mode indicates that the set aligns conspicuously with a particular equal division or cyclic pattern. This is why the Fourier language clarifies the special status of collections such as the whole-tone set (6-35), the diminished seventh (4-28), and the diatonic scale (7-35). They are not merely historical categories. They are sets with distinctive spectral saliences, and their Forte numbers serve as discrete addresses inside this spectral landscape.

Yet magnitude is only half of the story. Two different objects can share the same magnitudes while differing in arrangement. This is one way to restate the phenomenon of homometry and Z-relatedness that forced the layered turn at the end of Part I. Magnitude remembers how much of a given periodicity is present. It does not, by itself, fully remember how those periodicities are arranged relative to one another. What is missing is phase.

Magnitude remembers how much of a periodicity is present.
Phase remembers how that periodicity is placed.

Once that point is seen, the logic of the next step becomes difficult to resist. If content already has a natural spectral encoding, and if the incompleteness of that encoding is tied to the loss of phase information, then the honest continuous thickening of the content space should not live in an anonymous Euclidean cloud. It should live in a space of complex coordinates, where amplitude and phase are both present and where symmetries act in the simplest possible way: by rotating phases.

This conclusion becomes even stronger when order enters the picture. Part II introduced rooted cyclic orderings, gap words, interval cycles, and local permutation fibers. Those objects are also cyclic by nature. A rooted pattern is traversed around a loop. Its step sequence has periodicity, regularity, and modulation. Slonimsky's interval-cycle imagination can be re-described very precisely as a theory of patterns whose order signal concentrates strongly in a small family of periodic modes. Once again, Fourier language is not a decoration added later. It is an exact way of speaking about the cyclic structure already present in the musical object.

So the question shifts. We are no longer asking only which pitch classes are present, or which ordering has been chosen. We are asking how to write both layers in a common spectral language.

3.3 Content modes and order modes

For content, the move is straightforward. A pitch-class set can be represented by its indicator function on the twelve-point cycle. Taking the discrete Fourier transform produces a finite list of complex coefficients:

$$X_j(S) \;=\; \sum_n \chi_S(n)\, e^{-2\pi i\, jn/12}.$$

The formula says only this: for each cyclic mode j, measure how strongly the set S resonates with a j-fold periodicity around the chromatic circle. The coefficient is complex because every periodic component has not only a strength but a phase. Its magnitude tells us how much of that periodicity is present. Its angle tells us how that periodic component is situated on the circle.

For order, the basic idea is the same. A rooted cyclic ordering $p = (p_0, \ldots, p_{k-1})$ can be treated as a traversal of points on the pitch-class circle. One may encode the visited pitch classes as complex points $z_m = e^{2\pi i\, p_m/12}$ on the unit circle and then take a Fourier transform along the order index m. This produces coefficients of the form

$$Y_l(p) \;=\; \sum_m z_m\, e^{-2\pi i\, lm/k}.$$

Here the mode l measures periodicity in the traversal itself. A pattern that returns every third step, every fourth step, or alternates between a small

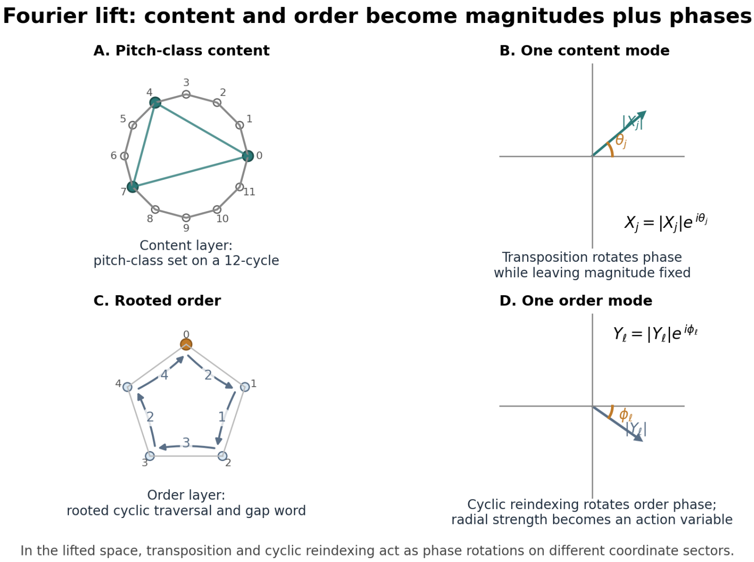

Figure 3.2: The layered Fourier lift: content and order coordinates are written as complex amplitudes and phases in a common ambient space.

number of step types will reveal that structure spectrally. If one prefers to work with the gap word rather than the visited points, one can define related coefficients from the successive intervals instead. The exact choice matters less than the principle. The order layer, like the content layer, has natural cyclic modes, and those modes are most honestly written as complex quantities carrying both amplitude and phase.

Once this is done, the base-and-fiber picture from Part II is enriched rather than erased. The content modes X_j describe what the pitch collection is, stripped to its cyclic essence. The order modes Y_l describe how that collection is traversed. They are not the same variables. But they can live side by side in a common ambient space. The layered object can now be lifted from a purely combinatorial description into a complex coordinate system in which both content and order are simultaneously present.

This is the decisive transition. The object is no longer merely a graph of equivalence classes with locally attached permutation clouds. It is an embedded family of points in a product of complex planes.

3.4 Why the natural completion is symplectic

Once one arrives at a product of complex planes, a certain geometry comes for free.

Every complex coordinate $z = x + iy$ carries a canonical oriented area form. In ordinary real coordinates, that form is just $dx \wedge dy$. In complex notation, it may be written as $\frac{i}{2}\, dz \wedge d\bar{z}$. This is the standard local building block of symplectic geometry, the geometry that underlies Hamiltonian phase spaces in classical mechanics. One need not yet know much symplectic theory to feel why it belongs here. If our chosen coordinates already come in pairs of intensity and phase, or more abstractly in canonically conjugate directions, then a geometry that preserves those pairings is exactly what we want.

For the layered musical construction, the most natural ambient form is the sum of the standard forms on all content modes and all order modes:

$$\omega \;=\; \frac{i}{2}\sum_j dX_j \wedge d\bar{X}_j \;+\; \lambda\,\frac{i}{2}\sum_l dY_l \wedge d\bar{Y}_l.$$

The constant λ is a balancing parameter. It controls how strongly the order layer is weighted relative to the content layer. Conceptually, nothing mysterious is happening. The content coordinates contribute one symplectic block; the order coordinates contribute another; the total form is the direct sum of the two.

What does this buy us? First, it turns the spectral lift into something more than a convenient parametrization. The ambient space is no longer just a place to put points. It has a built-in geometry of motion. Second, it makes precise a fact that has been implicit ever since Part I: amplitudes without phases are incomplete, and phases are not merely labels. They are dynamical coordinates.

The significance becomes even clearer when the complex variables are rewritten in polar form. A nonzero complex coordinate can be written as a magnitude times a phase factor. In the present setting, that means each content mode and each order mode can be described by an intensity variable and an angular variable. On the open locus where no chosen mode vanishes, the symplectic form becomes

$$\omega \;=\; \sum_j dI_j \wedge d\theta_j \;+\; \lambda \sum_l dJ_l \wedge d\varphi_l.$$

This is the action-angle form familiar from Hamiltonian geometry. Read in plain language, it says that each spectral mode splits into a quantity measuring "how much" and a quantity measuring "where in the cycle," and the symplectic form pairs those two kinds of data. Intensity and phase are treated as conjugate variables.

> *The symplectic form is not ornamental mathematics pasted onto music. It is the bookkeeping device that preserves amplitude-phase pairings in the spectral lift.*

This is the moment at which the musical interpretation sharpens. The content layer ceases to be merely a bag of spectral saliences. It becomes a family of phase-bearing modes. The order layer ceases to be merely a set of local features. It too becomes a family of phase-bearing modes. The full object behaves like a layered phase space: one part remembers what periodic structures are present, the other remembers how traversal organizes them, and the geometry knows how to keep both kinds of information in coordinated relation.

One should resist a crude physical analogy here. The claim is not that a fugue or a jazz solo literally obeys Newton's equations. The claim is subtler and stronger. Once musical content and musical order are encoded spectrally, the correct structure-preserving continuous geometry is symplectic. This is a statement about how best to organize the coordinates, not about reducing art to mechanics.

3.5 Symmetry becomes internal motion

The case for the symplectic picture would already be strong if it only provided a clean ambient geometry. But the argument goes further. The musically basic symmetries act especially well in these coordinates.

Begin with transposition. In ordinary pitch-class terms, transposition adds the same interval to every note. In content spectral coordinates, this does not scramble the modes. It simply rotates the phase of each mode by an amount proportional to the transposition and to the mode number. What looked, in pitch space, like a rigid displacement becomes, in the spectral space, a collection of coordinated phase rotations. Because these rotations preserve the standard symplectic form, the transposition action is not merely geometric. It is Hamiltonian in the usual sense: it is generated by a conserved quantity, or more precisely by a moment-map coordinate built from the mode intensities.

Now consider cyclic reindexing in the order layer. Part II identified rooted cyclic order as the right local quotient. If one shifts the starting point of a rooted traversal around its cycle, the order modes again transform by simple phase factors. So cyclic reindexing acts on the order coordinates exactly as transposition acts on the content coordinates: as a phase rotation preserving the symplectic form. The parallel is beautiful and revealing. What key change is for content, root shift is for order.

Inversion behaves differently, and that difference matters. In the pitch-class world, inversion reflects the cycle. Spectrally, that reflection is represented by complex conjugation together with an index reversal appropriate to the chosen encoding. Complex conjugation reverses the sign of the standard symplectic form. So inversion is not symplectic but anti-symplectic.

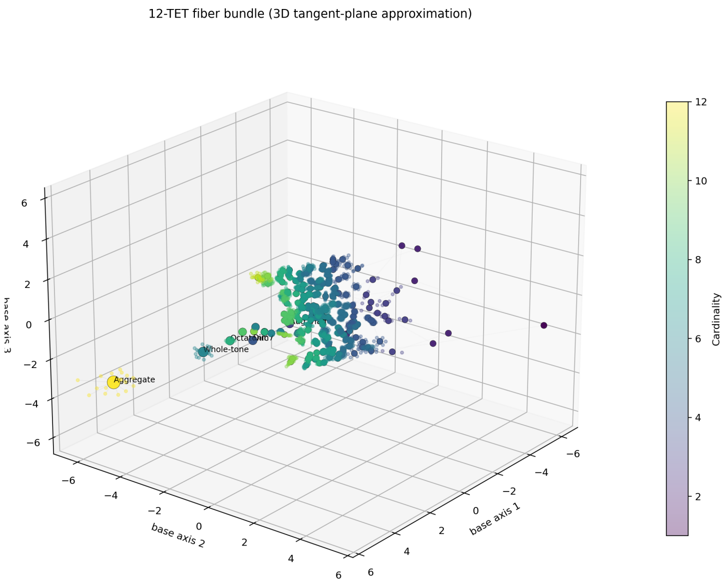

Figure 3.3: Tangent-plane fiber bundle in 3D—the base records content classes, the attached local sheets represent sampled order fibers.

Again, this is not an arbitrary mathematical flourish. It matches the musical intuition that inversion is a reflection rather than a continuous rotation.

Once these actions are visible, a deeper point emerges. The musically fundamental operations are no longer being described from outside as combinatorial rewritings. They appear as internal motions of the ambient structure itself. Transposition is phase rotation. Cyclic reindexing is phase rotation. Inversion is reflection. The coordinate system is not fighting the music; it is letting the music act in its native form.

This is one of the chief reasons the symplectic completion feels like an arrival rather than an embellishment. The same formalism that preserves amplitude-and-phase pairings also makes the core musical symmetries unusually transparent.

3.6 Why the quotient is stratified

The story would be tidier if every point in the space behaved generically. Music does not permit that luxury.

Some pitch collections possess exceptional symmetry. The diminished seventh chord (4-28) is the standard example. Transpose it by a minor third and it maps onto itself. The whole-tone collection (6-35) behaves similarly under its own characteristic shifts. These objects do not sit in the transposition quotient the way a generic triad does. They carry nontrivial stabilizers. In the language of group actions, their orbits are smaller because more transformations fix them.

The same phenomenon can occur in the order layer. Certain cyclic patterns have periodicities or reversal symmetries that make distinct descriptions collapse under the relevant quotient actions. The fiber above such a point is therefore not locally equivalent to the generic fiber. Something pinches, folds, or identifies.

This is why the mathematically honest global object is not a single smooth manifold. It is stratified. One may visualize it as a space built from regions of different regularity stitched together along singular loci. On the generic or free locus, the quotient by transposition or cyclic reindexing behaves smoothly. At symmetry-rich points, one obtains finite quotient singularities of the sort that mathematicians often package as orbifold structure. For an essay, the plainest way to say it is this: the space is smooth where nothing exceptional is being fixed, and folded where symmetry forces several local directions to count as the same.

This is not a pathological nuisance. It is musically correct. Symmetrical objects really are landmarks. They gather pathways and collapse distinctions. They have fewer genuinely different transpositions, fewer genuinely different readings, and more internal self-similarity than generic objects. A model that forced them to look regular would be less faithful than one that allowed the singularities to remain visible.

The symplectic language survives this complication. On each regular stratum, the reduced quotient inherits a genuine symplectic form. Across the singular set, the total object is best understood as a stratified symplectic space: not uniformly smooth, but still organized by compatible symplectic pieces. This is one of the places where the theory acquires real mathematical seriousness. It does not flatten away the exceptional cases; it explains why they are exceptional.

3.7 What the symplectic view buys us

At this point a skeptical reader might ask whether the whole construction has earned more than a refined vocabulary. Why should one care that the spectral thickening is symplectic rather than merely Euclidean, or merely topological, or merely combinatorial?

The first answer is conceptual clarity. A metric space tells us about near and far. A symplectic space tells us about lawful motion. Those are not the same things. The difference between them is precisely the difference between a static similarity map and a phase space. If one wants to talk only about clusters, neighborhoods, and embeddings, a metric may suffice. If one wants to model families of transformations that preserve a deeper structure, one needs more.

The second answer is that symplectic geometry distinguishes two tasks that are often confused in musical discourse. One task is to describe which objects resemble one another. The other is to describe how one object may evolve into another while preserving chosen invariants. These are related, but not interchangeable. A Hamiltonian flow, for example, need not follow a geodesic of any visually convenient metric embedding. It follows the rule generated by a chosen function on phase space. This gives the theory room to express musical preference, tension, or periodic attraction in a principled way rather than by ad hoc nearest-neighbor rules.

It is worth marking the boundary between what the present construction has established and what it makes possible but has not yet demonstrated. The Fourier lift, the exact symplectic form, the action-angle decomposition, the Hamiltonian circle actions of transposition and cyclic reindexing, and the stratified orbifold reduction are all proved. What remains aspirational—and deliberately so—is the next step: the construction of musically motivated Hamiltonians and the demonstration of their dynamical consequences on real repertoire. The framework delivers the phase-space architecture; it does not yet deliver the dynamics.

Suppose, for instance, that one were to write a Hamiltonian favoring strong triadic salience in the content modes and moderate regularity in the order modes. In Forte's catalog this would mean privileging the neighborhood of 3-11 (the major/minor triad) and its immediate harmonic extensions—4-27 (the dominant seventh), 4-26 (the minor seventh), 3-12 (the augmented triad)—all of which cluster in the spectral region where the third and fifth Fourier components are strong. The resulting flow would privilege movements that preserve or exchange those saliences in a controlled manner. A different Hamiltonian might favor equal division of the octave and a periodic gap structure, producing a dynamical bias toward Slonimsky-like interval cycles with limited interpolative drift—a flow that gravitates toward 6-35 (the whole-tone collection), 4-28 (the diminished seventh), and similar high-symmetry landmarks. These are natural next applications of the framework, not claims the present work has already delivered. In either case the model would not be forcing music into one destiny. It would be providing a disciplined way to formulate families of musically meaningful motion.

The third answer is conservation. Whenever a group action is Hamiltonian, one obtains associated conserved quantities. In the present setting, transposition symmetry and cyclic reindexing symmetry naturally suggest moment-like quantities built from spectral intensities. This does not mean that actual pieces of music preserve those quantities globally; compositions

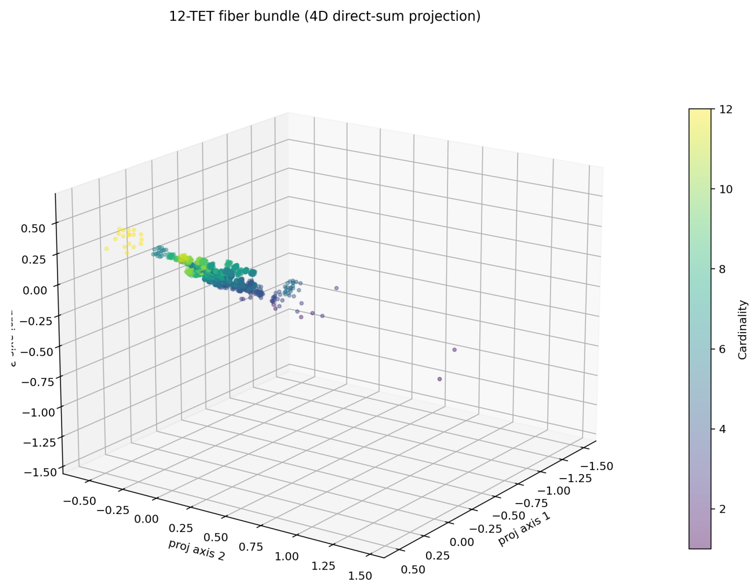

Figure 3.4: A 4D direct-sum projection of the total object—base motion and fiber motion separated more cleanly in the higher-dimensional model.

are not obliged to respect any single model. But it does mean that, inside the model, one can separate changes due to pure symmetry from changes due to genuine deformation. That distinction is already musically valuable. It clarifies the difference between moving the same object to a new key center and altering the object's internal spectral profile.

The fourth answer concerns continuation between discrete and continuous worlds. The musical seed from which we began is discrete: finite sets, finite orderings, finite quotients. Yet many musically illuminating transformations feel continuous: one hears gradual shifts of emphasis, progressive interpolations, rotations through pattern families, or deformations of harmonic quality. A symplectic thickening provides a principled ambient continuum in which the discrete objects sit as distinguished points or lattices. That makes it possible to speak rigorously about continuous families without pretending that the discrete seed has vanished.

The symplectic layer also clarifies why it should not replace the rest of the toolkit. Some analytical or computational tasks will still be better served by multidimensional scaling, diffusion methods, or simplicial chord complexes. The point is not that one formalism abolishes the others. The point is that the symplectic thickening supplies something they do not: a geometry of lawful motion rather than a geometry of adjacency alone.

Finally, the symplectic view unifies the entire series. Part I asked how musical content could be given a geometry and found the first answer in cyclic autocorrelation and spectral fingerprints. Part II asked how order could be attached without destroying that geometry and found the answer in local fibers of rooted cyclic traversal. Part III shows that when both layers are lifted into their natural spectral coordinates, the completed object is not merely a better picture. It is the beginning of a dynamics.

3.8 What this does not claim

A good theory earns trust not only by what it includes, but by what it refuses to claim.

This construction is not "the space of all music." It is not a universal ontology of listening, nor a replacement for style, history, analysis, performance, rhythm, timbre, or embodiment. It is a rigorous model for one particular layer of musical organization: pitch content and ordered traversal in a cyclic twelve-tone setting, together with a natural spectral thickening that supports lawful transformations.

Even within pitch, it is selective. It privileges pitch-class organization over register, timbre, and attack. It works at the level of abstract collections and cyclic patterns rather than at the level of full scores or recordings. Voice-leading, if one wishes to include it, must enter either through additional metrics, further coordinates, or a richer bundle structure. Rhythmic organization would require its own cyclic or non-cyclic layers, possibly coupled to the pitch layers but not automatically reducible to them.

Nor should one confuse mathematical naturalness with psychological necessity. A symplectic form may be the right geometry for the chosen spectral coordinates without being the only geometry relevant to human hearing. Listeners also respond to roughness, expectation, memory, stylistic norm, motor gesture, cultural learning, and the stubborn local grain of sound. The point of the present construction is not to eliminate those realities. It is to isolate a layer that can be treated exactly and then to show how unexpectedly rich that exact treatment becomes.

There is also no need to pretend that every useful musical path will be Hamiltonian. Some analytical or compositional tasks will still be better served by combinatorial graphs, shortest-path metrics, diffusion processes, or simplicial models. The symplectic layer should be understood as an enrichment, not an imperial conquest. It adds a geometry of lawful motion to the existing picture. It does not abolish the other kinds of structure we have already built.

These limits are not weaknesses. They are what keep the theory honest.

3.9 Closing synthesis

The series began with a simple unease: the ordinary categories of music theory, however powerful inside their own domains, do not by themselves chart the full universe of twelve-tone possibility. That unease led first to a compression. We folded pitch into pitch class, pitch class into cyclic content, and content into quotient spaces defined by symmetry. We found in cyclic autocorrelation the first serious fingerprint of musical content, and in its Fourier shadow the first hint that the relevant geometry was already spectral.

The second step was to acknowledge that content is not the whole object. Order matters. Traversal matters. The same notes arranged differently do not inhabit the same local world. So we separated content from order, attached rooted cyclic fibers to the content base, and built a bundle-like object in which local permutation geometry could be seen without being confused with global content similarity.

The third step has now completed the structural ascent. Once content and order are both written in spectral coordinates, the correct continuous thickening is a product of complex mode spaces. And once that space is written honestly, it carries a natural symplectic form. Magnitude and phase become conjugate variables. Transposition and cyclic reindexing become internal phase rotations. Inversion becomes reflection. Symmetry-rich objects become singular strata rather than anomalies. The whole construction passes from the level of atlas to the architecture of a phase space.

To be precise about what has been established: for cyclic pitch content plus cyclic order, the Fourier-based layered formalization admits a compelling and exact symplectic completion with unusually transparent musical symmetries. The stronger claim—that this architecture supports musically validated dynamics on real repertoire—remains an open program rather than a settled result. The framework delivers the kinematics and symmetry structure; the dynamics are the natural next question.

At the beginning of the series, music wanted a geometry. By the end of Part III, that geometry has acquired the structure of a phase space—and with it, the capacity to support lawful motion, once the right Hamiltonians are found.

The result is not the final theory of music. It is something better than that: a tractable, layered, mathematically articulate model in which the universe of pitch content, cyclic order, and their spectral completion can be studied without flattening away the differences that made the problem worth posing in the first place.

The geometry of musical possibility begins as a catalog of patterns, and reaches maturity when it acquires the phase-space structure needed for a geometry of lawful motion.

Notes

1. Andrew Duncan, “Combinatorial Music Theory,” *Journal of the Audio Engineering Society* 39, no. 6 (1991): 427–448.

2. Jason Yust and Emmanuel Amiot, “Non-Spectral Transposition-Invariant Information in Pitch-Class Sets and Distributions,” in *Mathematics and Computation in Music*, LNCS 13267 (2022), 279–291.

3. Nicolas Slonimsky, *Thesaurus of Scales and Melodic Patterns* (New York: Scribner’s, 1947).

4. Clifton Callender, Ian Quinn, and Dmitri Tymoczko, “Generalized Voice-Leading Spaces,” *Science* 320, no. 5874 (2008): 346–348; Dmitri Tymoczko, *A Geometry of Music* (New York: Oxford University Press, 2011).

5. V. I. Arnold, *Mathematical Methods of Classical Mechanics*, 2nd ed. (New York: Springer, 1989); Dusa McDuff and Dietmar Salamon, *Introduction to Symplectic Topology*, 3rd ed. (Oxford: Oxford University Press, 2017).

6. The language of singular stabilizers, orbifold-like quotients, and reduced spaces is motivated by the same general quotient-geometric issues emphasized by Callender, Quinn, and Tymoczko.

7. Art Samplaski, “Mapping the Geometries of Pitch-Class Set Similarity Measures via Multidimensional Scaling,” *Music Theory Online* 11, no. 2 (2005); Louis Bigo et al., “Representation of Musical Structures and Processes in Simplicial Chord Spaces,” *Computer Music Journal* 39, no. 3 (2015): 9–24; Ronald R. Coifman and Stéphane Lafon, “Diffusion Maps,” *Applied and Computational Harmonic Analysis* 21, no. 1 (2006): 5–30.

Consolidated Bibliography

- Amiot, Emmanuel. *Music Through Fourier Space.* Springer, 2016.
- Arnold, V. I. *Mathematical Methods of Classical Mechanics.* 2nd ed. New York: Springer, 1989.
- Bigo, Louis, Daniele Ghisi, Antoine Spicher, and Moreno Andreatta. "Representation of Musical Structures and Processes in Simplicial Chord Spaces." *Computer Music Journal* 39, no. 3 (2015): 9–24.
- Callender, Clifton, Ian Quinn, and Dmitri Tymoczko. "Generalized Voice-Leading Spaces." *Science* 320, no. 5874 (2008): 346–348.
- Coifman, Ronald R., and Stéphane Lafon. "Diffusion Maps." *Applied and Computational Harmonic Analysis* 21, no. 1 (2006): 5–30.
- Duncan, Andrew. "Combinatorial Music Theory." *Journal of the Audio Engineering Society* 39, no. 6 (1991): 427–448.
- Forte, Allen. *The Structure of Atonal Music.* New Haven: Yale University Press, 1973.
- McDuff, Dusa, and Dietmar Salamon. *Introduction to Symplectic Topology.* 3rd ed. Oxford: Oxford University Press, 2017.
- Morris, Robert. *Composition with Pitch-Classes: A Theory of Compositional Design.* New Haven: Yale University Press, 1987.
- Quinn, Ian. "General Equal-Tempered Harmony." *Journal of Music Theory* 50, no. 2 (2006).
- Rahn, John. *Basic Atonal Theory.* New York: Longman, 1980.
- Samplaski, Art. "Mapping the Geometries of Pitch-Class Set Similarity Measures via Multidimensional Scaling." *Music Theory Online* 11, no. 2 (2005).

- Slonimsky, Nicolas. *Thesaurus of Scales and Melodic Patterns.* New York: Schirmer, 1947.
- Tymoczko, Dmitri. *A Geometry of Music: Harmony and Counterpoint in the Extended Common Practice.* New York: Oxford University Press, 2011.
- Yust, Jason. *Organized Time: Rhythm, Tonality, and Form.* Oxford: Oxford University Press, 2018.
- Yust, Jason, and Emmanuel Amiot. "Non-Spectral Transposition-Invariant Information in Pitch-Class Sets and Distributions." In *Mathematics and Computation in Music*, LNCS 13267, 279–291. 2022.

About the Project

GAMUT (Geometric-Algebraic Music Theory) is a research project developing mathematical models of pitch-class space. The project comprises:

- This three-part essay series, presenting the argument at a public-facing level of detail.
- Two companion formal papers: a technical paper developing the layered musical manifold construction, and an AMS-format proof paper formalizing the symplectic geometry.
- The *Layered Bundle Explorer*, an interactive visualization that lets users traverse the 223-node content graph, inspect permutation fibers, and watch a guided tour through common Western harmonic structures.
- A Python library (`musical_manifold`) for computing content spaces, permutation fibers, spectral embeddings, and associated distance data.

All materials are available at:
https://shapeofmusicalpossibility.org