GENERALIZED CHORD SPACES

Clifton Callender, Florida State University
Ian Quinn, Yale University
Dmitri Tymoczko, Princeton University

Draft of 20 June 2005. This is a work in progress.

Music theorists have been using geometrical models to explore voice-leading for almost three hundred years. These models come in a wide variety of seemingly unrelated forms. Some model the voice-leading possibilities between individual chords: the familiar circle of fifths (Example 0.1[a]), the Oettingen/Riemann *Tonnetz* (Example 0.1[b]), Jack Douthett and Peter Steinbach's "Cube Dance" (Example 0.1[c]) and "Power Towers" (Example 0.1[d]), and Dmitri Tymoczko's "scale lattice" (Example 0.1[e]). Others model voice-leading possibilities between *set-classes:* Joe Straus's discrete model of trichord space (Example 0.2[a]); Clifton Callender's triangular model of *continuous* trichord-space (Example 0.2[b]); Straus's 2D model of discrete tetrachordal set-class space (Example 0.2[c]); Richard Cohn's 3D tetrahedral model of tetrachordal set-class space (Example 0.2[d]); and the familiar model of diatonic triad-classes shown in Example 0.2(e). Still other geometrical models depict voice-leading possibilities between *ordered* sets and set-classes: most notably, John Roeder's "ordered interval space" (Example 0.3[a]), which has been further explored and developed by Clifton Callender (Example 0.3[b]). **[Examples 0.1-0.3 will be found in a separate file.]**

The great diversity of these geometrical models raises a host of difficult theoretical questions. First, how do these models relate to one another? Are they like maps of different continents, or do they offer alternate perspectives on the same underlying terrain? Second, can we *generalize* these models, so that they are as complete and comprehensive as possible? Many of the above models deal with only a restricted range of sets and set-classes: for instance, the models in Example 0.1 treat only a small set of familiar scales and chords; while those in Example 0.2(a) and (c) deal only with the *equal tempered* set-classes—themselves a small subset of the available harmonic resources. Is it possible to generalize these maps so that they include all possible n-note chords or chord-types? Finally, what are the music-theoretical consequences of thinking geometrically about harmonic objects? Theorists are accustomed to describing music in discrete terms, using combinatorics and finite groups. Yet the notion of geometry essentially involves mathematical concepts that go beyond those that are used in discrete mathematics—notions such as continuity, "straight line," and distance. Does a mature geometrical perspective require music theorists to master a new set of tools, and do these new tools shed new light on familiar musical concepts such as "interval" and "chord"?

In this paper we attempt to answer all these questions. We show that the various structures in Examples 0.1-0.3 are in fact deeply interrelated—different perspectives on, or regions of, the same underlying geometrical space. We show how to extend these geometrical models so that they describe *all* the voice-leading possibilities between *all* possible sets and set-classes, both unordered and ordered. We show that understanding these voice-leading spaces requires music theorists to absorb new mathematical tools—
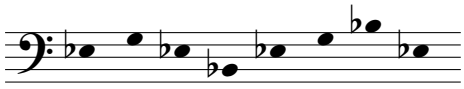
notions from geometry and topology, including the concept of a "quotient space." And we will argue that these new tools provide interesting new perspectives on some basic music theoretical concepts, including the notion of "interval," "interval class," and "transformation." Our ultimate goal is to attempt to move geometrical ideas to the center of music theory, displacing—or at least complementing—the finite, combinatorial paradigm that has dominated the field for the last few decades.

More specifically, we will argue that the spaces in Examples 0.1-0.3 belong to a larger family of geometrical spaces that share a few essential features. A *point* in these spaces corresponds to a harmonic object such as a "chord" or "set-class." A straight *path* between two points corresponds to a voice-leading between them, with the *length* of the path equal to the "size" of the voice-leading it represents. Consequently, the *distance* between two points corresponds to the size of the *minimal voice-leading* between the two harmonic objects. These features together imply that the graphs in Examples 0.1-0.3 can be treated in a unified way: they are all *quotients* of the same fundamental type of space, or spaces that result from "gluing together" points in a larger "parent space." We will begin, in §§1-2, by introducing showing how to translate traditional music theoretical ideas into the language of geometry. In §3-6 we use geometrical concepts to demonstrate the essential unity of the graphs in Examples 0.1-0.3. In §§7-9 we begin to take up some of the larger theoretical issues raised by the discussion.

§1. Sixteen ways of looking at a warhorse

1.0. Example 1.1 presents the beginning of the main theme of the first movement of Beethoven's *Eroica* symphony, suppressing all information about rhythm, tempo, timbre, and instrumentation. What remains is an abstract notation encoding various types of harmonic information—information about the *order, register, absolute transpositional level, absolute direction of intervallic motion,* and *cardinality* of the notes in this passage. Depending on our musical interests, we may wish to disregard some of or all of this information. We can, for example, describe Example 1.1 as an ordered series of pitches (E♭3, G3, E♭3, B♭2, E♭3, G3, B♭3, E♭3)[1], as an *unordered* collection of *pitch-classes* {E♭, E♭, E♭, E♭, G, G, B♭, B♭}, as an "E♭ major triad," as a "major triad," or—perhaps most generally or all—a "consonant triad" or member of set-class 3-11[037]. These descriptions are all equally correct, though they involve different degrees of precision. They therefore suggest different analytical consequences: under some descriptions the harmonic object shown Example 1.1 is "the same" as the chord stated at the opening of the movement; but under others it is not. Likewise, under some descriptions it is "the same" as the chord that opens the *Symphony of Psalms,* while under others it is not.

[1] In this paper, regular parentheses denote *ordered* series, so that (a, b, c) ≠ (b, c, a). Curly braces denote unordered sets. Thus {a, b, c} = {b, c, a}.



Example 1.1. The opening of the first theme of Beethoven's Third Symphony

When we disregard such harmonic information, we create an *equivalence class*—a collection of objects that differ only in terms of the information we are disregarding. Thus, if we disregard register, the ordered series (E♭3, G3, B♭4) is equivalent to (E♭2, G5, B♭7). If we disregard order, then the series (E♭3, G3, B♭4) is equivalent to (G3, E♭3, B♭4). If we disregard *both* order and octave then (G3, E♭3, B♭4) is equivalent to (E♭2, G5, B♭7). The more information we disregard, the larger our equivalence classes. This allows us to make analytical statements that are less precise, but more general. Disregarding musical information therefore testifies not to conceptual or analytical laziness, but rather to a legitimate interest in abstracting away from the musical surface of the piece. Conversely, the less information we discard, the more precise—but less general—our analytical statements becomes. The art of analysis consists, in part, in the skillful balance between the conflicting virtues of precision and generality.

Let us now consider four types of information conveyed by Example 1.1, and the different sorts of equivalence classes that result when we disregard this information.

1. Octave (**O**). We can use the boldface letter **O** to refer to octave equivalence, writing (E♭3, G3, B♭4) $\sim_O$ (E♭2, G5, B♭7). (**NB:** this notation indicates that two *distinct* objects are equivalent when we disregard octave information.) We can denote the **O**-*equivalence class* containing (E♭3, G3, B♭4) in a number of different ways. For example, we can use a subscript "O" to indicate that one is to disregard registral information; thus $(E♭3, G3, B♭4)_O = (E♭2, G5, B♭7)_O$. (**NB:** this notation indicates that a single object—an O-equivalence-class—is identical to itself.) Alternatively, we can simply eliminate the Arabic numerals that indicate register, as in $(E♭3, G3, B♭4)_O = (E♭, G, B♭)$. Finally, if we are using numbers to indicate pitches and pitch-classes, we can simply stipulate that numbers in the range $0 \leq x < 12$ to refer to pitch-classes rather than pitches.[2] Thus the **O**-equivalence class of Example 1.1 is $(E♭3, G3, E♭3, B♭2, E♭3, G3, B♭3, E♭3)_O = (E♭, G, E♭, B♭, E♭, G, B♭, E♭) = (3, 7, 3, 10, 3, 7, 10, 3)$. As the notation indicates, we can consider an **O**-equivalence class to be an ordered series of pitch-classes. Indeed, a "pitch-class" itself is simply the **O**-equivalence class of a one-element series.[3]

2. Permutation (**P**). The pitches in Example 1.1 are ordered. If we disregard their ordering, we obtain an equivalence class consisting of all 840 distinct orderings of these eight elements: (E♭3, G3, E♭3, B♭2, E♭3, G3, B♭3, E♭3), (E♭3, B♭3, G3, E♭3, B♭2, E♭3, G3, E♭3), (B♭2, E♭3, E♭3, E♭3, E♭3, G3, G3, B♭3), and so on.[4] All of these ordered sequences are *permutations* of each other. We will use the boldface **P** to refer to permutation, and the symbol $\sim_P$ to refer to equivalence-to-within permutation: (E♭3, G3, E♭3, B♭2, E♭3, G3, B♭3, E♭3) $\sim_P$ (E♭3, B♭3, G3, E♭3, B♭2, E♭3, G3, E♭3), We can refer

[2] We discuss mapping pitches to numbers in §3, below.

[3] See §3.

[4] Were our set to contain no duplications, there would be 40320 (8!) of these. Since the set contains duplications, however, there are only 8!/(2!4!) = 840.

to the P-equivalence class containing an ordered set either with the subscript "P" or (somewhat more traditionally) using curly braces. Thus the equivalent notations {B♭2, E♭3, E♭3, E♭3, E♭3, G3, G3, B♭3} and (E♭3, G3, E♭3, B♭2, E♭3, G3, B♭3, E♭3)$_P$ refer to the class containing all 840 orderings of the 8 elements B♭2, E♭3, E♭3, E♭3, E♭3, G3, G3, and B♭3.

3. Transposition (**T**). We can also disregard the specific *notes* in Example 1.1, considering only the intervallic content of the passage. That is, we disregard the *transpositional level* at which the pattern of Example 1.1 is stated. We will use a boldface **T** to refer to equivalence-under-transposition, writing: (E♭3, G3, E♭3, B♭2, E♭3, G3, B♭3, E♭3) ~$_T$ (D4, F♯4, D4. A3. D4, F♯4, A4, D4). Similarly, we can use the subscript "**T**" to indicate equivalence-classes under transposition. Thus (E♭3, G3, B♭3)$_T$ refers to the **T**-equivalence class containing (E♭3, G3, B♭3).

4. Inversion (**I**). Finally, we can abstract from the absolute *direction* of the intervals in Example 1.1. That is, we could consider the sequence (E♭3, G3, E♭3, B♭2, E♭3, G3, B♭3, E♭3) to be equivalent to one of its inversions in pitch-space, such as (E♭3, C♭2, E♭3, A♭3, E♭3, C♭2, A♭2, E♭3). Here, *ascending* intervals have been replaced by *descending* intervals and vice versa. **I**-equivalence, described in this way, is a good deal less familiar than **O**, **P**, and **T**, equivalence. This is largely because music theorists are accustomed to combining **I**-equivalence with **T**-equivalence.[5] The details here are somewhat subtle, and will be explored in depth below. For now, however, we will simply declare that an *inversion-class* consists of two sequences of pitches (a_1, a_2, …, a_n) and ($\mathbf{I}_x(a_1)$, $\mathbf{I}_x(a_2)$, …, $\mathbf{I}_x(a_n)$), where $\mathbf{I}_x$ is some absolute (non-contextual) inversion in pitch-space.

Naturally enough, the equivalence-classes generated by **O**, **P**, **T**, and **I** can be combined to form larger equivalence classes. Thus, for example, we could disregard both octave and order information, considering all objects that are **OP**-equivalent to the one shown in Example 1.1. This **OP**-equivalence class is the set of all possible orderings of the pitch-classes {E♭, E♭, E♭, E♭, G, G, B♭, B♭}. This equivalence-class is more familiarly described as a *multiset* of pitch-classes—a collection in which cardinality but not order is significant. In the same way, we could consider the equivalence-class of all ordered pitch-series equivalent to Example 1.1 under transposition or inversion. We might call this equivalence-class the *ordered* **TI**-*class of pitches* containing Example 1.1. There are 16 (= $2^{1.1}$) possible ways to combine of the equivalence-classes generated by **O**, **P**, **T**, and **I**, many of which have familiar music-theoretical names.

We will now briefly describe these 16 ways of looking at the harmonic content of Example 1.1. Each corresponds to a different level of "abstraction" from the musical surface of the piece, and hence suggests a different level of analytical generality. Some of these ways of abstracting correspond to very familiar music-theoretical objects, while others are much less familiar.

1. (No equivalences) Without any equivalences, Example 1.1 is an ordered series of pitches: (E♭3, G3, E♭3, B♭2, E♭3, G3, B♭3, E♭3).

[5] To make matters worse, there seem to be two different sorts of inversion: "absolute" inversions and "contextual inversions." We discuss these later in the paper.

2. **O**. An ordered series of pitch-classes: (E♭, G, E♭, B♭, E♭, G, B♭, E♭) or (E♭3, G3, E♭3, B♭2, E♭3, G3, B♭3, E♭3)$_{O}$.
3. **P.** An unordered collection of pitches in which cardinality is important, or a *multiset of pitches:* {B♭2, E♭3, E♭3, E♭3, E♭3, G3, G3, B♭3}.
4. **T**. An ordered *transposition-class* of pitches (E♭3, G3, E♭3, B♭2, E♭3, G3, B♭3, E♭3)$_{T}$ = (D4, F♯4, D4. A3. D4, F♯4, A4, D4)$_{T}$.
5. **OP.** A multiset of *pitch-classes:* {E♭, E♭, E♭, E♭, G, G, B♭, B♭}.
6. **OPT.** A transpositional multiset-class. {0, 0, 0, 0, 4, 4, 7, 7}$_{T}$, in "normal form."
7. **OPTI**. A transpositional and inversional multiset-class. In "normal form," {0, 0, 0, 0, 3, 3, 7, 7}$_{TI}$.
8. **OT**. An ordered transposition-class of pitch-classes, consisting of all sequences of pitch-classes transpositionally equivalent to (E♭, G, E♭, B♭, E♭, G, B♭, E♭).
9. **OTI**. An ordered transposition-and-inversion class of pitch-classes, consisting of all sequences of pitch-classes transpositionally or inversionally equivalent to (E♭, G, E♭, B♭, E♭, G, B♭, E♭).
10. **PT**. An unordered transposition-class of pitches, consisting of every different ordering of every collection of pitches that is transpositionally equivalent to {B♭2, E♭3, E♭3, E♭3, E♭3, G3, G3, B♭3}.
11. **PTI.** An unordered transposition-and-inversion class of pitches, consisting of every different ordering of every collection of pitches that is transpositionally or inversionally equivalent to {B♭2, E♭3, E♭3, E♭3, E♭3, G3, G3, B♭3}.
12. **TI**. An ordered transposition-and-inversion class of pitches, consisting of all sequences of pitches transpositionally or inversionally equivalent to (E♭3, G3, E♭3, B♭2, E♭3, G3, B♭3, E♭3).
13. **I**. An ordered *inversion-class* of pitches, consisting of (E♭3, G3, E♭3, B♭2, E♭3, G3, B♭3, E♭3) and one ordered sequence of pitches equivalent to it under inversion, such as (E♭3, C♭2, E♭3, A♭3, E♭3, C♭2, A♭2, E♭3).
14. **OI**. An ordered *inversion class* of *pitch-classes,* consisting of (E♭, G, E♭, B♭, E♭, G, B♭, E♭) and one ordered sequence of pitch-classes equivalent to it under inversion, such as (E♭, C♭, E♭, A♭, E♭, C♭, A♭, E♭).
15. **PI**. An unordered inversion-class of pitches, consisting of {B♭2, E♭3, E♭3, E♭3, E♭3, G3, G3, B♭3} and one unordered sequence of pitches equivalent to it under inversion, such as {A4, E4, E4, E4, E4, C4, C4, A5},
16. **OPI**. An unordered inversion-class of pitch-classes, consisting of {E♭, E♭, E♭, E♭, G, G, B♭, B♭} and one unordered sequence of pitches equivalent to it under inversion, such as {A, A, C, C, E, E, E, E},

Consider now the different geometrical models shown in Examples 0.1-0.3. Each of these models can be interpreted as displaying the voice-leading relationships among equivalence classes generated by some combination of **O**, **P**, **T**, and **I**. Those in Example 0.1 depict voice-leading relationships among "chords," or—in our terminology—**OP**-equivalence classes. Those in Example 0.2 depict voice-leading relationships among "set classes," or in our terminology **OPTI**-equivalence classes. Finally, those in Example 0.3 depict voice-leading relationships among ordered transposition-classes, or **T**-equivalence

classes.[6] The essential idea of the present paper is to provide a geometrical characterization of the notions of **O**, **P**, **T**, and **I** equivalence, thereby specifying the relation between the models shown in Examples 0.1-0.3. We show that in disregarding information about octave, order, transposition level, or the "direction" of motion in pitch-class space, we are identifying, or "gluing together," points in a geometrical space. This intuitive notion of "gluing together" is captured by the mathematical notion of a *quotient-space*. Thus the different geometrical models in Example 0.1-0.3 represent different quotient-spaces that can be formed from the same underlying space.

However, the geometrical approach works best if we restrict the range of acceptable equivalence classes to those that give rise to consistent geometrical interpretations. This requirement leads to several important (and somewhat surprising) departures from music-theoretical tradition. Perhaps the least surprising of these departures is the suggestion that **P**-equivalence is very closely analogous to **O**-, **T**-, and **I**-equivalence: from the geometrical perspective, *unordered* harmonic objects are simply equivalence classes of *ordered* harmonic objects, just as multisets of pitch-classes are equivalence-classes of multisets of pitches. By contrast, the traditional approach tend to make a much stronger distinction between unordered and ordered objects.[7] Rather more surprising, however, is the fact that the geometrical approach leads us to ignore three types of equivalence-class that are very important to traditional music-theory: equivalence to within pitch-class duplication, equivalence under intervallic multiplication, and "Z-equivalence," or equivalence to within total intervallic content. We end this section by briefly considering these three types of equivalence class, and sketching how they conflict with the geometrical approach of the present paper.

5. Cardinality (**C**). Equivalence to within duplication of pitches and pitch-classes is central to the traditional music-theoretical notions of "set" and "set-class." According to the traditional view, the unordered collections of pitch-classes {E♭, G, E♭, B♭} and {E♭, G, B♭} both instantiate the *same* pitch-class set, namely {E♭, G, B♭}. Our geometrical approach suggests that we depart from tradition in this regard. For us, the unordered collections {E♭, G, E♭, B♭} and {E♭, G, B♭} are *not equivalent*, because the former has *two* instances of pitch-class E♭, while {E♭, G, B♭} has only one. We will explore the reasons for this in detail below; in essence, however, the problem is that in our geometrical models, the number of notes in a chord corresponds to the *dimension* of the underlying space. In identifying, say, the three-note set {0, 0, 4} with the two-note set {0, 4} we are attempting to glue together a point in 3D-space with a point in 2D-space. As we will see, this sort of identification prevents us from conceiving of the size of a voice-leading as a "distance."[8] It is interesting to reflect that such a central aspect of

---

[6] More specifically, Roeder's spaces are **OT**-spaces, while Callender's are **O**-spaces. Both authors further investigate regions of their spaces containing no chords equivalent under **P** and **I**. This amounts to identifying a *fundamental domain* of the appropriate group action. For an explanation of these concepts, see below.

[7] Examples?

[8] To get a feel for the problem, note that we would intuitively like the distance between the ordered chords (0, 4, 4) and (0, 4, 5) to be 1, and similarly for the distance between (0, 0, 4) and (11, 0, 4). However, we would presumably like the distance between (0, 4, 5) and (11, 0, 4) to be greater than 2, since the minimal voice-leading between these two

music-theoretical thinking—treating cardinality as inessential—is in conflict with the geometrical approach to voice-leading. We hypothesize that this conflict with **C**-equivalence is what prevented the earlier emergence of a comprehensive geometrical approach to voice-leading spaces.

6. Multiplication (**M**). Another traditional equivalence concerns *equivalence under multiplication* by x semitones or $\mathbf{M}_x$-equivalence. The pitch-class sets {0, 1, 2} and {0, 5, 10} are M_5-equivalent because {0 * 5, 1 * 5, 2 * 5} = {0, 5, 10}. We will find that $\mathbf{M}_x$-equivalence is again incompatible with the geometrical approach to voice-leading spaces. Here the issue is that $\mathbf{M}_x$-equivalence produces counter-intuitive results when combined with **O**-equivalence in the context of *continuous* geometrical spaces. Take, for example, the $\mathbf{OM}_5$-class containing the pitch-class C (= 0) and its "just fifth," the pitch-class $12\log_2 3$ - 12. Since $12\log_2 3 - 12$ is irrational, its M5-equivalence class is infinitely large; furthermore, it contains elements *arbitrarily close* to every other pitch-class interval containing C.[9] Thus $\mathbf{OM}_5$-equivalence implies that *every pitch-class interval is almost equal to the acoustically pure fifth!* Though this result does not involve outright mathematical inconsistency, it is extremely counter-intuitive, leading to radically nonstandard geometries of dubious musical utility. For this reason, we will not consider $\mathbf{M}_x$-equivalence in this paper.

7. Interval-class equivalence (**Z**). Two set-classes are **Z**-related (or **Z**-*equivalent*) if they share the same total interval-class content. Thus the set-classes $\{0, 1, 4, 6\}_{\mathbf{OPTI}}$ and $\{0, 1, 3, 7\}_{\mathbf{OPTI}}$ are **Z**-related since they both contain exactly one interval from each interval-class. (These are in fact the only four-element, **Z**-related set-classes.) However, **Z**-related set-classes do not, in general, have the same voice-leading capabilities: a chord belonging to $\{0, 1, 4, 6\}_{\mathbf{OPTI}}$ can be connected by single-semitone voice-leading to a chord belonging to $\{0, 2, 4, 6\}_{\mathbf{OPTI}}$, while a chord belonging to $\{0, 1, 3, 7\}_{\mathbf{OPTI}}$ cannot. As we will see, this means that the most reasonable definitions of "voice-leading size" are inconsistent with the notion of **Z**-equivalence: we can generate geometrical models in which points represent **Z**-equivalence classes, and we can generate geometrical models in which distances correspond to the size of minimal voice-leadings, but not both.[10] We will explore this matter further, below.

---

chords involves at least six semitones of motion. But this is extremely problematic if the distance between (0, 0, 4) and (0, 4, 4) is 0, as it needs to be if these two chords are represented by a single point.

[9] This notion of "closeness" implies a concept of "distance" or a metric. For more on this, see below.

[10] Suppose—as seems reasonable—the voice-leading distance between $\{0, 1, 4, 6\}_{\mathbf{OPTI}}$ and $\{0, 2, 4, 6\}_{\mathbf{OPTI}}$ is 1, the distance between $\{0, 1, 2, 7\}_{\mathbf{OPTI}}$ and $\{0, 1, 3, 7\}_{\mathbf{OPTI}}$ is 1, and the distance between $\{0, 1, 2, 7\}_{\mathbf{OPTI}}$ and $\{0, 2, 4, 6\}_{\mathbf{OPTI}}$ is greater than 2. We can represent **Z**-equivalence by letting the distance between $\{0, 1, 3, 7\}_{\mathbf{OPTI}}$ and $\{0, 1, 4, 6\}_{\mathbf{OPTI}}$ be 0. But then the resulting assignments violate the triangle inequality, and cannot represent true "distances."

§2. Quotient spaces, metrics, and group actions

We will now start to translate these familiar music-theoretical ideas into geometrical terms. We begin by introducing four important mathematical concepts: *space*, *metric*, *quotient space*, and *group action.*

A *space* is simply be a collection of objects without any further structure. Thus we can consider the space of all integers, or of all ordered pairs of real numbers, or of pitch-classes, or indeed of almost anything else. The important point is that in defining a space, we simply list the set of objects it contains, without saying anything at all about how they relate. Mathematicians typically use the term "space" to refer to *topological* spaces, which have slightly more structure than our minimalist "spaces."[11] For our purposes, we can think of topology as providing a very general sense in which two distinct spaces can equivalent (or *homeomorphic*).[12] As we will see, some familiar musical spaces are topologically—but not geometrically—equivalent.

A *metric* gives us a way to measure "distances" in a space. Formally, a *metric* for a space S is a function **d**(x,y), from points in S to real numbers, that possesses four properties:

1. (Non-negativity) **d**(x,y) ≥ 0.
2. (Identity of indiscernables) if **d**(x, y) = 0, then x = y.
3. (Symmetry) **d**(x,y) = **d**(y,x).
4. (Triangle inequality) **d**(x,z) ≤ **d**(x,y) + **d**(y,z), for all x, y, z in S.

A *metric space* consists of a space, along with a metric for that space.

A *quotient space,* intuitively, is the result of "dividing" a space by an equivalence relation: given a space S and an equivalence relation ~ over S, the *quotient space* S/~ (pronounced "S modulo ~" or "S mod ~") is the space consisting of the *equivalence classes* of ~. We can think of the quotient space S/~ as mapping all points in S to their equivalence classes under ~. Equivalently, we can think of the quotient space S/~ as creating a new space by *identifying* or "gluing together" all those points of S that are equivalent under the relation ~.

Note that a metric for a space will not necessarily give rise to a metric for the quotients of that space. Let S be a space, with metric **d**, and let ~ be an equivalence relation over S. There is a natural way to extend **d** to the elements of S/~: define **D**(A, B), for all A, B in S/~, as the smallest value **d**(a, b) such that a ∈ A, b ∈ B. The function **D** assigns to each pair of equivalence classes a number corresponding to the minimum of the distances between their elements. However, **D** does not always define a metric for S/~. (We will provide a practical illustration at the end of §3, below.) A metric on S generates a metric on S/~ only when the metric and the equivalence relation are well-matched; in such cases, we say that the quotient space *inherits* a metric from its parent

---

[11] A *topological* space is a space S, along with a set O of subsets of S. The elements of O are called the *open sets* of S, and must satisfy four axioms: the empty set ∅ is in O; S itself is in O; the union of an arbitrary number of elements of O is in O; and the intersection of a finite number of elements of O is in O. A set O satisfying these axioms is called a *topology* for S, and can be used to define the notion of continuity.

[12] Two topological spaces A and B are topologically equivalent, or homeomorphic, if there is a bijective function A→B that sends only open sets in A to open sets in B.

space. As we will see, it is highly desirable to work with such spaces, since they allow us to use a single notion of "voice-leading size" to measure distances between a variety of related harmonic objects, including ordered chords, unordered chords, and multiset-classes.

Next, we consider quotient spaces that are formed by the action of a group on a set. We will say that the group Γ *acts on* space S, if, for every group element $g \in \Gamma$, and every point $x \in S$, we can assign a unique point $gx \in S$, such that the following conditions obtain:

1. $ex = x$, for all $x \in S$, where e is the identity element of Γ; and
2. $g(h(x)) = (gh)x$, for all $x \in S$, and all $g, h \in \Gamma$.

A *group action* of Γ on S is therefore a function $\Gamma \times S \rightarrow S$.[13] The *orbit* of a point x under the group action is the set of points $\{gx\}$, for all $g \in \Gamma$. It is readily checked that group orbits are equivalence classes. Consequently, any group action of Γ on S defines an equivalence relation: $x \sim y$ only if there exists some $g \in \Gamma$ such that $gx = y$. We can therefore take the quotient of S by a group action, notated S/Γ. Note that since a group Γ can act on a space S in a variety of ways, it is important, when discussing a quotient S/Γ, to specify the intended action.

We are now able to introduce the basic geometrical model of harmonic objects that we will use for the rest of the paper. To begin, we associate pitches with real numbers **R** by stipulating that middle C corresponds to the number 60 and that ascending by semitone corresponds to addition by 1. Thus B3 = 59, C4 = 60, C♯4 = 61, D = 62, D "quarter-tone-sharp" = 62.5 and so on. Note that since the real numbers are *continuous*, there is a real number corresponding to every conceivable pitch: every microtone and every tone in every tuning system. (For example, the pitch a just fifth above middle C corresponds to the number $48 + 12\log_2 3$.) We will interpret an *ordered n-note series of pitches* as a point in the space $\mathbf{R}^n$, the n-fold Cartesian product of real numbers. Given any ordered pair of points in $\mathbf{R}^n$, $(x_1, x_2, \ldots, x_n)$ and $(y_1, y_2, \ldots, y_n)$ we can associate a unique voice-leading between them, one that maps x_n to y_n, for all n. A *metric* for $\mathbf{R}^n$ can be interpreted as measuring the size of these voice-leadings. As we will see, music theorists have developed a variety of methods of measuring voice-leading size, almost all of which can be interpreted as metrics.

The central thesis of this paper is that the graphs in Examples 0.1-0.3 all belong to a family of quotients of the same space under different equivalence relations, and inheriting a single metric from the parent space. The underlying spaces are simply $\mathbf{R}^n$, containing all possible ordered n-note series of pitches. The equivalence relations are the sixteen possible combinations of **T**, **P**, **I**, and **O** discussed in the previous section. A metric for $\mathbf{R}^n$ corresponds to a method of measuring voice-leading size. What is special about **T**, **P**, **I**, and **O** is that they can be represented by groups of *isometries* of $\mathbf{R}^n$: that is, groups that act so as to preserve the distances between any two points, according to any "reasonable" metric of voice-leading size. (We discuss the notion of a "reasonable" metric in Section 000, below.) As we will see, this means that these quotient spaces inherit metrics from their parent spaces. They therefore comprise a powerful and mutually consistent set of geometrical models of voice-leading, in which one and the

[13] The notation $A \times B$ refers to the *Cartesian product* of A and B, or the set of all ordered pairs (a, b), such that $a \in A$, and $b \in B$.

same notion of "voice-leading size" can be used to describe the relation between a variety of harmonic objects. (Some of these models are shown in Examples 0.1-0.3; others will be described below.) By contrast **C**, $\mathbf{M_x}$, and **Z** do not correspond to isometries of $\mathbf{R^n}$, and the quotient spaces they generate do not inherit the standard metrics of voice-leading size. Consequently, these operations do not generate well-behaved geometrical models such as those shown in Examples 0.1-0.3.

§3. One-note spaces

To illustrate these ideas, let us consider the various spaces associated with one-note "chords." Although one-note spaces are not intrinsically very interesting, our discussion will allow us to introduce some additional geometrical concepts that will be important in the remainder of the paper.

We begin by defining *pitch-space*, the space of all possible pitches. As explained in §2, this is simply **R**, the space of all real numbers. Next, we choose a metric for **R** that identifies the distance between pitches. Equivalently, our metric represents a method of measuring the size of the (one-voice) "voice-leadings" connecting one pitch to another. In general, it is quite difficult to choose a metric for the higher-dimensional spaces $\mathbf{R^n}$.[14] However, there is an extremely natural notion of "distance" in one-dimensional pitch-space: we can simply take the distance between two pitches to be determined by the function $\mathbf{d}(x,y) = |x - y|$.[15] Musically, this metric asserts that the distance between two pitches corresponds to the number of semitones between them.[16] The metric thus captures the intuitively appealing idea that a "semitone" is a unit of distance.[17] It is easily checked that the metric $|x - y|$ satisfies the four axioms in §2.

We will now describe the spaces that result from taking the quotient of the real line **R**, with the natural metric, by the equivalence-classes **T**, **P**, **I**, and **O**.

1. One-note **O**-space. Octave equivalence is modeled by the identification

$$\mathbf{O} : x \sim_O x + 12$$

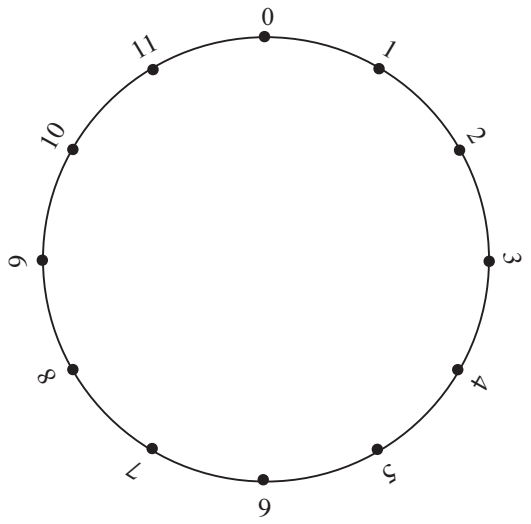
[14] The difficulty is exacerbated if we want to find a metric that accurately reflects our perceptual judgments. Indeed, it is not even clear that *any* metric is consistent with our intuitive perceptual judgments about voice-leading. However, there are numerous alternative criteria for measuring voice-leading size, including logical consistency. See §000 for more discussion.

[15] This represents the "natural" metric for **R**. The most commonly-used mathematical metrics agree with this definition of one-dimensional distances. A key consideration here is that, given some other metric $\mathbf{d_1}$ for **R**, we can often relabel the elements of **R** so as to satisfy $\mathbf{d_1}(x,y) = |x - y|$. In effect, this sort of "relabeling" is what occurs when we use the logarithm of a pitch's fundamental frequency to generate the standard numerical model of pitches.

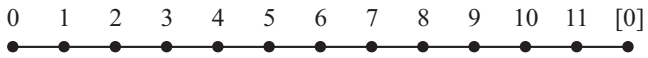
[16] Note that this number need not be an integer, since we are considering continuous pitch-space.

[17] Indeed, we are virtually required to use this metric, or some trivial variant thereof, if we want to consider a "semitone" to be a musical distance.

We can represent this equivalence relation as an action of the *group* 12**Z** of integer multiples of 12 on the *space* **R** of real numbers. The group 12**Z** acts on **R** by ordinary addition. Thus, the equivalence classes of **R**/**O** consist of the sets {…, x – 36, x – 24, x – 12, x – 0, x +12, x + 24, x + 36, …}. These are simply the pitch-classes. A metric for **R**/**O** (or *pitch-class space*) can be defined by taking the distance between two pitch-classes A and B to be the smallest of the distances between elements of A and elements of B. Thus **R**/**O** inherits a metric from **R**. This metric corresponds to the familiar music-theoretical notion of the "interval class" between two pitch-classes.



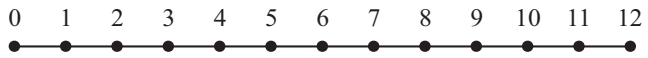
Example 3.1(a). Pitch-class space, represented as a circle



Example 3.1(b). Pitch-class space, represented as a line segment with endpoints identified

Examples 3.1(a) and 3.1(b) depict the quotient space **R**/**O** in two different ways. Example 3.1(a) shows **R**/**O** as a familiar circle. This circle appears to be curved and two-dimensional. These, however, are *extrinsic* features of the representation, products of the way it **R**/**O** has been embedded in higher-dimensional Euclidean space, rather than features of the space itself. If one considers the circle **R**/**O** as a space unto itself, and not embedded in any higher-dimensional space, then it is flat (uncurved) and one-dimensional. (It is one-dimensional, roughly, because at any point there are only two directions to move; it is flat, roughly, because it inherits the "natural" metric from **R**.) Example 3.1(b) removes any temptation to regard **R**/**O** as curved and two-dimensional, depicting the quotient space as a line segment whose endpoints are identified. Unfortunately, Example 3.1(b) is misleading in a different way: although represented by distinct points in Euclidean space, the two endpoints of Example 3.1(b)

correspond to the same point in **R**/**O**. Thus Examples 3.1(a-b) have complementary flaws: Example 3.1(a) represents a flat, one-dimensional space as a curved, two-dimensional object, but represents each point in **R**/**O** only once. Example 3.1(b) is clearly flat and one-dimensional, but appears to contain a "duplication" since it doubly represents a single point in the underlying space.



Example 3.1(c). A region of the real line, whose endpoints are *not* identified

It is important not to confuse the space shown in Example 3.1(b) with another, superficially similar geometrical object. Example 3.1(c) shows a region of **R**, the line segment $0 \leq x \leq 12$. Unlike Example 3.1(b), the endpoints of Example 3.1(c) are not identified. They represent distinct points of **R**, 0 and 12. These points are identified by the action of 12**Z** on **R**, but in the space **R** itself they remain distinct. Example 3.1(c) is what is called a *fundamental domain* (or *fundamental region*) of the action of 12**Z** in **R**. A fundamental domain of a group Γ on space S is a region that "tiles" S under the action of the group. That is: S is the union of the regions $\{g\mathrm{T}\}$, for all $g \in \Gamma$, and any two regions $g\Gamma$ and $h\Gamma$ intersect only at their boundaries. A fundamental domain can therefore contain "duplications" only at its boundaries. A fundamental domain can be transformed into the quotient space S/Γ by identifying, or gluing, the appropriate boundary points. Note that fundamental domains are not unique: in the present case, we could use any interval [x, x+12] in **R** as a fundamental domain of 12**Z** in **R**.

2. One-note **P**-space. No two one-note chords are **P**-equivalent, since a one-note chord has only one ordering. Thus "one note P-equivalence" can be modeled by the trivial equivalence class:

$$\mathbf{P} : x \sim_{\mathrm{P}} x$$

P-equivalence results from the action of the "trivial group" *E* on the space **R**. The group *E* contains only the identity element. It "acts" on **R** by leaving **R** unchanged: thus the quotient **R**/**P** is simply **R** itself.

3. One-note **T**-space. Two pitches a and b are T-equivalent if there is some real number x such that a = b + x. Thus T-equivalence can be modeled by the relation:

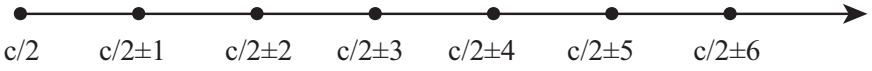
$$\mathbf{T} : x \sim_{\mathrm{T}} x + \mathrm{c}, \text{ for all real numbers c.}$$

T-equivalence can be represented as an action of the *group* **R** on the *space* **R**. The group is by ordinary addition, and generates a single equivalence class, since for any two points $x, y \in \mathbf{R}$, there is always some real number n, such that x + n = y. Thus the quotient space **R**/**T** consists in a single point. This is the geometrical expression of the fact that there is only one one-note set-class.

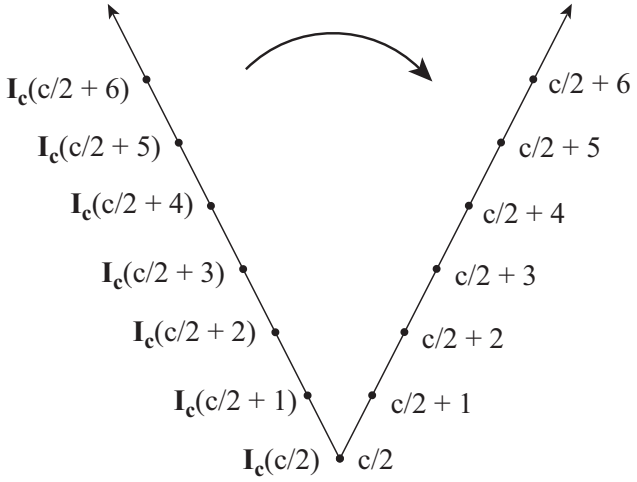
4. One-note **I**-space. Two pitches a and b are equivalent under $\mathbf{I}_c$ if a + b = c. Thus

$$\mathbf{I}_c : x \sim_{\mathrm{I}} \mathrm{c} - x \text{ , where c is a constant.}$$

The number c corresponds to the "index number" of the inversion. We can describe inversion-with-index-c using the function $\mathbf{I}_c(a) = c - a$. To find a group containing this function $\mathbf{I}_c$, observe that $\mathbf{I}_c(\mathbf{I}_c(a)) = \mathbf{I}_c(c - a) = c - (c - a) = a$. Thus $\mathbf{I}_c$ is an involution: it composes with itself to form the identity operation. $\mathbf{I}_c$ and the identity together comprise the two-element group $\mathbf{Z}_2$. Thus, one-note **I**-space is the quotient $\mathbf{R}/\mathbf{Z}_2$, with $\mathbf{Z}_2$ acting by subtraction from a constant value. The geometrical term for subtraction from a constant value is *reflection.*



Example 3.2(a). The quotient space **R**/**I** is a half-line.

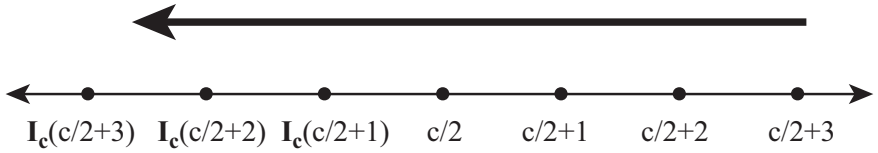


Example 3.2(b). We can fold the real line **R** in half to form quotient space **R**/**I**.

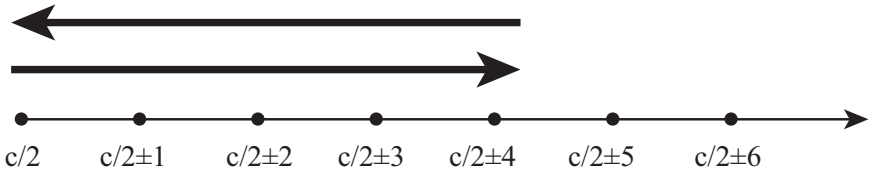
Example 3.2(a) shows that one-note **I**-space is a "half-line": a ray with a single endpoint at c/2. As Example 3.2(b) shows, we arrive at this geometrical structure by "folding" the real line **R** at the point c/2, thereby identifying all $\mathbf{I}_c$-related pairs. Algebraically, the vertex of the fold is a *fixed point* of the function $\mathbf{I}_c$, since $\mathbf{I}_c(c/2) = c/2$. The fact that a non-trivial element of $\mathbf{Z}_2$ "fixes" a point in **R** means that the quotient space $\mathbf{R}/\mathbf{Z}_2$ will have a *singularity*: an "exceptional point" at which the space inherits an unusual geometry from its parent. None of the other three quotient spaces we have examined thus far, **R**/**O**, **R**/**P**, and **R**/**T**, have had singularities, since no *non-trivial* element of their respective groups fixes any point in the space **R**.[18]

[18] Of course, the *identity* element in these groups trivially fixes *all* points of the space; these do not count as true fixed points.

How should we think about the singularity c/2 in the space **R**/**I**? Intuitively, it might appear to be a straightforward "boundary" or "edge" of the space. We might imagine that a very tiny (0-dimensional!) creature, walking along the space **R**/**I** would "bump into a wall" when it reached the point c/2: it would simply "run out of space" and be unable to continue moving in the same direction. Natural as this perspective may be, it is not the most useful way to think about the point c/2. To see why, consider Example 3.3. Example 3.3(a) represents the path of our 0-dimensional creature as it moves along **R** toward c/2 from the positive direction. The example shows the two fundamental domains of the $\mathbf{Z}_2$ action in **R**, consisting of the half-lines $x \geq c/2$ and $x \leq c/2$. (The two fundamental domains overlap at the single point c/2.) The 0-dimensional creature will not notice anything special when it reaches the point c/2: it will not "bump into a wall" or find that it can no longer keep moving "in the same direction." Instead, it simply passes seamlessly from one fundamental domain to another. As Example 3.3(a) shows, upon crossing c/2, it begins to encounter pitches $\mathbf{I}_c$-related to those it already encountered, in the reverse order.



Example 3.3(a). A straight-line path in **R**.



Example 3.3(b). The image of this straight-line path in the quotient space **R**/**I**.

Example 3.3(b) presents the image of the path shown in Example 3.3(a) in the quotient space **R**/**I**, as depicted in Example 3.2(a). Motion toward c/2 appears as leftward motion on the example; motion *away* from c/2 appears as rightward motion. Our creature's straight-line path in **R** through c/2 gets represented in **R**/**I** as a path that moves toward c/2 from the right, "bounces" or "reflects off" c/2, and moves back in the direction it came from. Crucially, however, this "change of direction" *results from the structure of the quotient space* rather than from a genuine change of direction in parent space **R**.[19] To capture this sense of an "effortless" change of direction, we need to say

[19] One important caveat: the path shown in Example 3.3(b) could also arise from a genuine change in direction in **R**—our one-dimensional creature could move leftwards to the fixed point c/2, only to "turn around" and begin moving rightwards in **R**. (It is possible to walk right up to the mirror and turn around, rather than passing through it.)

that there is something about the *structure* of the space **R**/**I** that causes straight lines to retrace their steps when they reach the point c/2. The metaphor of a mirror or reflecting boundary allows us to do just that. Most, but not all of the singularities we will encounter in this paper act like mirrors.[20] Singularities result when we import the geometrical structure of the parent space—including its "angles" and "straight lines"—into the quotient space. The mathematical term *orbifold* describes a quotient space that inherits a geometry from its parent space, including its angles, straight-lines, and distances. Singularities are a characteristic feature of orbifolds.

These concepts have very natural musical interpretations. We could reframe our narrative about the "0-dimensional creature" as a narrative about a musical voice that glissandos continuously in register through the point c/2. From a phenomenological perspective, we would not notice a dramatic "change in direction" as the voice passed through the fixed point c/2. *Analytically,* however, we might wish to assert that the musical voice retraces familiar harmonic territory once it passes through c/2. Here, the "harmonic objects" are given by the equivalence classes of $\mathbf{I}_c$—the points of the quotient space **R**/**I**. Our perceptual notion of "moving in a straight line," however, derives from the parent space **R**. These two notions unite when we describe **R**/**I** as an orbifold that inherits its geometry from **R**.

Finally, note that Example 3.3 shows that the quotient **R**/**I** inherits a metric from **R**. As before, we define the "distance" between two equivalence classes in **R**/**I** as the minimum of the distances between any element of the first class and any element of the second. This definition simply reproduces the standard metric $|x - y|$ on the half-line $x \geq c/2$. Note that it is a consequence of this definition that the shortest path between two points in **R**/**I** never passes through the mirror point c/2. This is a straightforward consequence of the triangle inequality. As we will see, this idea is closely connected to the traditional contrapuntal principle that one should avoid "voice-crossings" between successive chords.

We conclude this section by considering two more spaces: the combined **OP** space, and a "nonstandard" quotient space that does not inherit a metric from **R**.

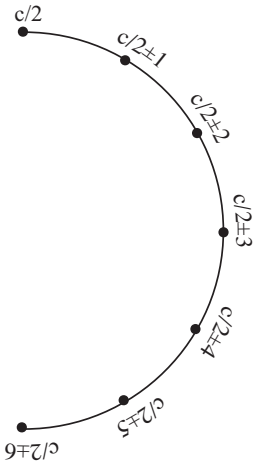
5. We have seen that one-dimensional **P** space is the real line **R**, while one-dimensional **T** space is a single point. One of these spaces is "too big," in the sense that it is equivalent to the parent space **R**; the other is "too small" in the sense that consists only of a single point. Consequently, we do not expect **P** or **T** to combine with the other equivalence classes to form interesting spaces. The two equivalence relations **O** and **I**, however, can be combined. We define $\mathbf{OI}_c$-equivalence as follows:

Thus, two paths in **R**, one "kinked" and one "straight" are mapped to the same "kinked" path in **R**/**I**.

[20] As we will see, the singularities in our spaces manifest themselves in two ways: by the fact that straight-lines passing through them "change direction," and by the fact that infinitesimal circles surrounding them have an area less than πr^2. Singularities having the first property invariably have the second property as well. Singularities having the second property need not have the first. These singularities manifest themselves solely in the fact that circles surrounding them do not have the "correct" area—that is, that some of the space around the singularity is missing. We will encounter an example of this when we consider trichordal set-class space in §5, below

$\mathbf{OI_c}$: $x \sim_{OI} x + 12 \sim_{OI} c - x$, where c is a constant.

OI-equivalence can be represented as a group action of $12\mathbf{Z} \times \mathbf{Z}_2$ on **R**, with $12\mathbf{Z}$ and $\mathbf{Z}_2$ acting as before. To understand the space, consider the effect of reflection around point c/2 on the circular space **R**/**O**, shown in Example 3.1(a). This reflection "folds" one half of the circle onto the second half—just as we did in deriving Example 3.2(b). As Example 3.4 shows "folding" a circle in half creates two fixed points: the expected fixed point c/2, and a second point halfway around the circle at c/2 + 6. This second point is also a mirror point in **R**/**OI**. (**NB:** the fact that Example 3.4 is "curved" has no intrinsic significance, and merely illustrates the relation between Example 3.4 and Example 3.1; intrinsically, **R**/**OI** is a line-segment.)



Example 3.4. The quotient space **R**/**OI** is a line-segment with two mirror boundaries.

Where does this second mirror point come from? Consider ordinary pitch inversion around middle C. In pitch-space, this inversion has only one fixed point—the axis of inversion, corresponding to the vertex of the "fold" shown in Example 3.2. Now consider the effect of this inversion on some pitches near middle C: E4, a major third above middle C, is mapped onto G♯3, a major third below middle C; F♯4, a tritone above middle C is mapped onto F♯3, a tritone below middle C. While E4 and G♯4 belong to different pitch-classes, F♯4 and F♯3 belong to the *same* pitch-class. In other words, in **O**-space, F♯3 and F♯4 are identified. Thus $\mathbf{I_c}$ has only one fixed point in pitch-space, but *two* fixed points in pitch-class space: one corresponding to the fixed point of the pitch-space inversion, the other resulting from the fact that inversion in pitch-space preserves the *pitch-class* of pitches one or more tritones away from the inversional axis.

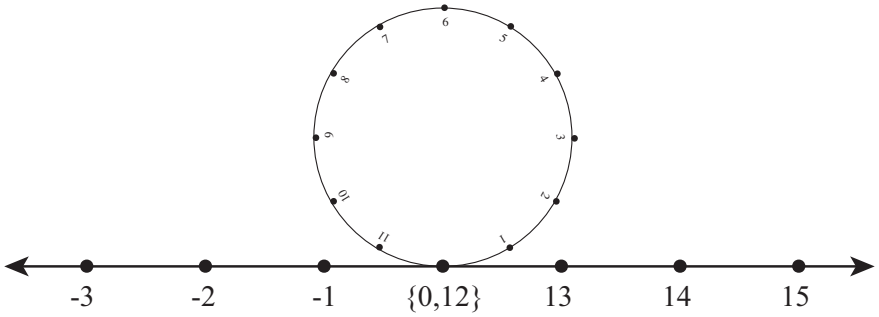
As we will see, the one-note **OI** space shown in Example 3.4 is isomorphic to two-note **OPT**-space, more familiarly known as the voice-leading space of two-note *set-classes.*[21] This represents our first encounter with an interesting phenomenon: a single space that has two different musical applications.

[21] Strictly speaking, multiset classes.

6. Finally, by way of contrast, let us consider a non-standard equivalence relation $\sim_?$. We will define $\sim_?$ as the equivalence relation that arbitrarily identifies two points in **R**, say 0 and 12. Consider the function *f* on **R**, defined as follows:

$f(0) = 12$, $f(12) = 0$, $f(x) = x$, for $x \neq 0, 12$.

f, like $\mathbf{I}_c$, is an involution. Thus the nonstandard equivalence relation $\sim_?$ can be represented as an action of $\mathbf{Z}_2$ on **R**.[22] Example 3.5 depicts the quotient space $\mathbf{R}/\sim_?$; it is a circle perched atop a line. The intersection of these two shapes is the equivalence class {0, 12}, the "arbitrary" identification encoded by our equivalence relation $\sim_?$.



Example 3.5. Identifying 0 and 12 in **R** creates a circle attached to a line.

Clearly, the space $\mathbf{R}/\sim_?$ does *not* inherit the "natural" metric from **R**. Suppose we define a function **D** on pairs of elements of $\mathbf{R}/\sim_?$, as described in §2, above: **D**(A, B) = the smallest distance between any pitches $a \in A$, $b \in B$. According to this definition, **D(**{0, 12}, {3}) = 3, **D**({0, 12}, {15}) = 3, and **D(**{15}, {3}) = 12. Thus **D** cannot be a metric, since it violates the triangle inequality: according to **D**, the distance from A to B is *larger* than the distance from A to C, plus the distance from C to B! To be sure, it would be entirely possible to provide an alternate metric for $\mathbf{R}/\sim_?$. There is nothing *inherently* non-metrical about the quotient space $\mathbf{R}/\sim_?$. But it is *not* possible to define a metric on $\mathbf{R}/\sim_?$ that is consistent with the natural metric we have defined for the parent space **R**: there will always be some distances between A, $B \in \mathbf{R}/\sim_?$ that are not equal to the distances in **R** between any $a \in A$, $b \in B$. In musical terms, this means we have to consider voice-leading in $\mathbf{R}/\sim_?$ to be fundamentally different from voice-leading in **R**: we cannot draw conclusions about the size of voice-leadings in **R** on the basis of information about voice-leadings in $\mathbf{R}/\sim_?$, nor vice versa.

[22] Note that these two $\mathbf{Z}_2$ actions generate very different quotient spaces. Thus we see that it is very important to specify how the group Γ acts to produce a quotient space S/Γ.

§4. Classification of chord spaces I: $\mathbf{R}^n$, **O**- and **P**-equivalence.

4.0. We now turn to higher-dimensional spaces. In §4 we will consider spaces involving **O-** and **P**-equivalence. §5 considers **T-**equivalence, while §6 considers **I**-equivalence. Each of these sections concludes with a space central to music-theoretical thinking: §4 concludes with **OP**-space, or the space of sets; §5 concludes with **OPT**-space, or the space of transpositional set-classes; and §6 concludes with **OPTI**-space, or the space of **TI**-set-classes. These three types of space subsume the familiar music-theoretical graphs shown in Examples 0.1-0.3.

We begin by considering ordered chords of pitches. As noted above, we can represent such chords using $\mathbf{R}^n$, the space of ordered n-tuples of real numbers. $\mathbf{R}^1$ is the real number line itself, where each "number" is a point with a single coordinate. $\mathbf{R}^2$ is the space of ordered pairs, with each element of the space having two coordinates (x_1, x_2). Note that this space is *not* the familiar "Cartesian plane" since we do not presuppose a Euclidean (or any other) metric; for this reason, we do not have the resources to say that $\mathbf{R}^2$ has "coordinate-axes" that are "orthogonal." Similarly, $\mathbf{R}^3$ is not to be confused with the three-dimensional space of ordinary Euclidean geometry. Our spaces $\mathbf{R}^n$ are simply spaces of ordered n-note series of real numbers $(x_1, x_2, \cdots, x_n)$.

The elements of $\mathbf{R}^n$ describe ordered musical objects. This ordering can occur in time, as discussed in conjunction with Example 4 of §1 (the opening theme of the Eroica symphony). It can also occur in register, as when we consider a chord to be ordered from "bottom to top."[23] Even more abstractly, we can use the ordering to identify instruments or musical voices. Here, while the *consistent* assignment of notes to voices, and of voices to dimensions, is of significance, the *particular* assignment is not; there is no geometrical reason to choose an analysis of a four-voice chorale texture that consistently makes the assignment (S,A,T,B) over one that consistently makes the assignment (S,T,A,B). To be sure, the former assignment ensures that order positions will typically also correspond to registral position. But this is a mere curiosity: we can chose to use order to represent register, and we can chose to use order to represent instrumental voices, but we cannot do both at the same time—at least if "voice-crossings" are a musical possibility.

Finally, it is worth reemphasizing that our use of real numbers to represent pitches is to some degree arbitrary. The decision to associate the distance 1 with the equal-tempered semitone is arbitrary, and proponents of other tuning systems might associate the distance 1 with the octave, or the pure twelfth, or anything else. Similarly, the decision to associate middle C with the number 60 is arbitrary. These arbitrary choices, however, do not fundamentally affect the geometry of the models we will be considering here: the spaces remain for all intents and purposes the same, no matter whether we measure distances in semitones, octaves, or otherwise, and no matter whether we associate middle C with the number 60, 0, or any other number. Thus it is purely for the sake of familiarity that we adopt the convention that 60 means middle C (C4), 62 means D4, 59 means B3, and so on.

[23] Cf. John Roeder's work.

4.1. We will refer to the order positions of points in $\mathbf{R}^n$ as "voices"—mindful of the fact that these "voices" can be manifested in a very wide range of musical ways. Between any two points in $\mathbf{R}^n$, $(x_1, x_2, \ldots, x_n)$ and $(y_1, y_2, \ldots, y_n)$, we can associate a unique "voice-leading" between them, that sends x_1 to y_1, x_2 to y_2, and so on, for all n. The voice-leading therefore sends voice n (or order position n, or instrument n, or …) in chord X to voice n in chord Y. A metric for $\mathbf{R}^n$ measures the distance between any two points, and can be interpreted as measuring the "size" of the voice-leading between them.

Two common metrics are the Euclidean metric

$$\delta_e(x,y) = \sqrt{(x_1 - y_1)^2 + (x_1 - y_1)^2 + \cdots + (x_n - y_n)^2}$$

and the taxicab metric

$$\delta_t(x,y) = |x_1 - y_1| + |x_2 - y_2| + \cdots + |x_n - y_n|.$$

Roughly speaking, the Euclidean metric measures distance "as the crow flies" (by means of the generalized Pythagorean theorem), and the taxicab metric measures distance in each dimension separately, summing the results. It is called the taxicab metric because it resembles the distance traveled by a taxicab taking the shortest possible route in a city with a grid of two-way streets.

Different choices of metric model different intuitions about voice leading. One of the differences between Euclidean geometry and taxicab geometry concerns the locus of points a fixed distance from a given point—a "generalized sphere," relative to the metric. In Euclidean geometry, the locus is the familiar sphere, whereas in taxicab geometry, it has edges and corners. (In Euclidean space, it appears to be a square, octahedron, or generally something called an "orthoplex.") The following figure shows an example in $\mathbf{R}^3$ of the locus of chords two semitones away from (60,64,67), a particular instance of a C-major triad.

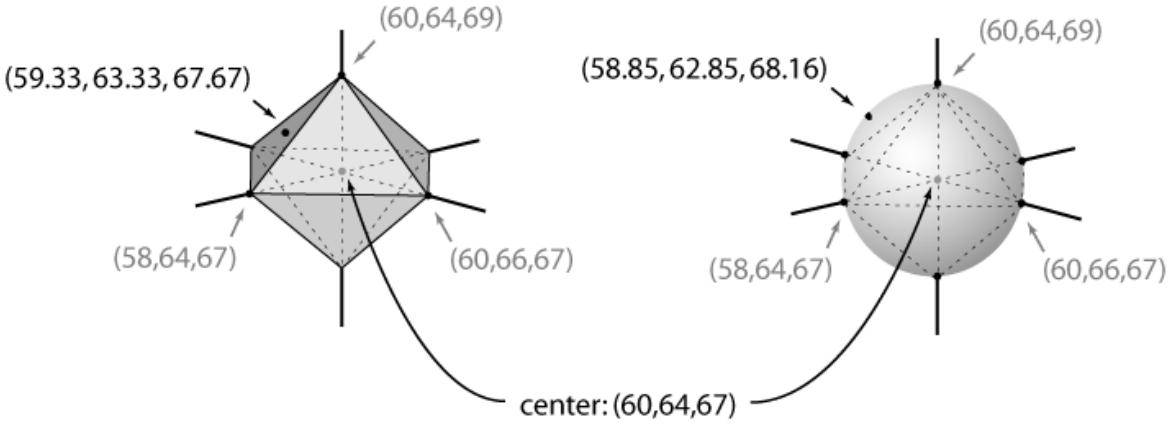


Figure 4.1

In taxicab geometry, this includes all chords that can be reached by moving one voice two semitones, two voices one semitone each, one voice 1.5 semitones and another 0.5 semitones, and so on--the idea is something like "among the three voices, two semitones' worth of motion occurs." These chords are found on the surface of a regular octahedron with side 4 (in taxicab geometry) and with (60,64,67) at its center. In Euclidean geometry, on the other hand, the locus of chords two semitones away forms a sphere of radius 2 (in Euclidean geometry), which is the smallest sphere that encloses the octahedron just discussed. The two geometries agree about the six cases in which only one voice moves, located at the vertices of the octahedron, but in all other cases Euclidean geometry seems to allow "more" motion in most cases, at least relative to taxicab geometry. The discrepancy is maximal for cases in which all voices move an equal distance, an example of which is shown in the figure. Musically, the Euclidean metric asserts that distances moved by individual voices do not add linearly: the more voices are moving, the less each motion contributes to the calculation of overall distance. This feature has a natural perceptual interpretation: the more voices that are moving, the less each additional voice contributes to the overall sensation of motion.[24]

Different theorists have favored different ways of measuring voice-leading size. The issue is a complicated and subtle one. We will return to it in greater detail below.

4.2. Having explored the raw space $\mathbf{R}^n$, let us take up the question of octave equivalence. We have already seen that $\mathbf{R}^1$ has the structure of a line and that $\mathbf{R}^1/\mathbf{O}$ has the structure of a circle whose circumference under our conventions, is 12; we have also seen that $\mathbf{R}^1/\mathbf{O}$ is the quotient of an action of the group $\mathbf{Z}$ on $\mathbf{R}^1$ generated by the identification $x \sim_{\mathbf{O}} x+12$.

Consider now the effect of octave equivalence on the plane $\mathbf{R}^2 = \mathbf{R}^1 \times \mathbf{R}$. Two identifications generate this equivalence:

$$\begin{aligned}(x_1, x_2) &\sim_{\mathbf{O}} (x_1+12, x_2) \\ &\sim_{\mathbf{O}} (x_1, x_2+12)\end{aligned}$$

These identifications correspond to an action of the group $\mathbf{Z}\times\mathbf{Z}$ (or $\mathbf{Z}^2$) on the space. The effect is to wrap the space around itself in each of its two dimensions. In three-dimensional space, we can represent a torus as the (curved) surface of a donut. Note, however, that the curvature of this space is an *extrinsic* feature. It is perfectly consistent to speak of a "flat" torus with Euclidean metric—even though these surfaces cannot be embedded in 3-dimensional Euclidean space.

At the beginning of the last paragraph we used a curious notation meant to highlight a particular instance of the general fact that $\mathbf{R}^n = \mathbf{R}^{n-1} \times \mathbf{R}$. A useful way to grasp the meaning of this and other product spaces is to read it as "$\mathbf{R}^n$ is a $\mathbf{R}^{n-1}$ of lines." For instance, the plane ($\mathbf{R}^2$) is a line of lines, or the result of replacing every point in a line with a line. Algebraically, this means that we add to $\mathbf{R}^{n-1}$ a single linear coordinate. Thus, to form $\mathbf{R}^3$, we add a single linear coordinate to $\mathbf{R}^2$, expressing the fact that three dimensional space is a plane of lines.

[24] For more on these issues, see Callender (unpublished).

Likewise, we can imagine the torus $\mathbf{R}^2/\mathbf{O} = \mathbf{R}^1/\mathbf{O} \times \mathbf{R}/12\mathbf{Z}$ as a "circle of circles," or the result of replacing every point in a circle with a circle. (Here we add a single *circular* coordinate to the original coordinate, which itself is circle.) Since a circle is the same thing as a line-segment-with-identified-ends, a circle of circles becomes a line-segment-with-identified-ends of line-segments-with-identified-ends, or, in plainer language, a square with opposite sides identified. This way of visualizing a torus is quite familiar to music theorists. In the following figure, a square torus representing $\mathbf{R}^2/\mathbf{O}$ is shown at (a), with a few points in the space specified.

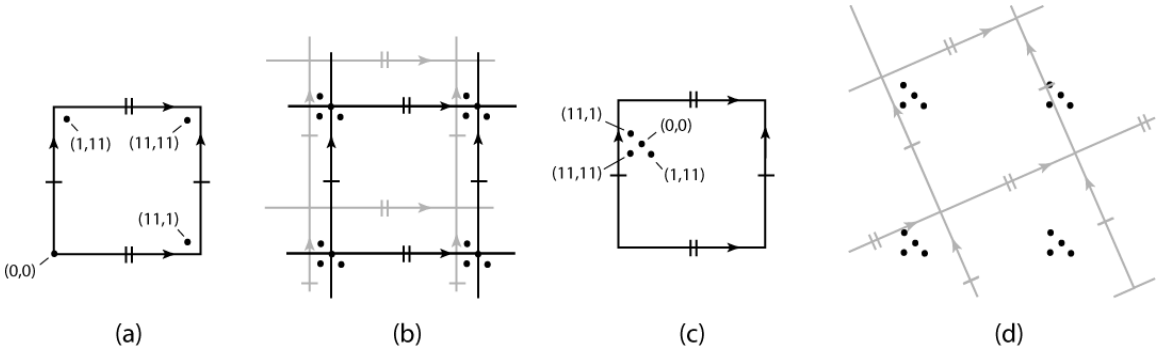


Figure 4.2

The default way of drawing the square torus is with its lower left-hand corner at (0,0). It should be noted, however, that there is nothing topologically (or geometrically) special about the origin; toroidal space is just as uniform as $\mathbf{R}^n$. Likewise, there is nothing special about the edges that are identified; it's just that they have to be drawn somewhere. Toroidal space simply has the property that it circles back on itself in all of its dimensions.

In the parent space $\mathbf{R}^2$, the fundamental domains of $\mathbf{R}^2/\mathbf{O}$ repeat periodically. At (b) the configuration of points from (a) is repeated briefly by realizing the identification of edges; this makes it easier to see that the square representation is a window that can be translated indiscriminately over the space. The diagram at (c) shows a different view of exactly the same space. The metaphor of the window has its limits, however. First of all, it does not mean to imply that the window excerpts a large, repeating space; the window contains the entire space, and indeed, with its two "extra" edges, just a bit more. Second, while the window can be moved indiscriminately over the space, it cannot be rotated; see (d), which shows that imagining a large, repeating space with an incorrectly oriented window leads to a contradiction between views. Thus while toroidal space is uniform in the sense that every point has the same neighborhood, it has a natural "grain" to it: walking in certain directions will bring you back to your starting point faster than others, and there are also directions in which you can walk in a straight line without ever covering exactly the same ground.

We are ready to move to higher dimensions. The space $\mathbf{R}^3$ can be construed as a plane of lines, replacing points in a plane with lines like we did to get $\mathbf{R}^2$ from $\mathbf{R}^1$, and likewise we can construct the space $\mathbf{R}^3/\mathbf{O}$, a "3-torus," as a torus of circles. A 3-torus is difficult to visualize in three dimensions except by analogy with the square representation of a 2-torus; a 3-torus is a cube with opposite faces identified:

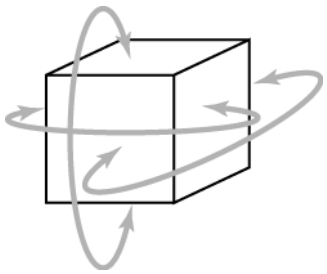


Figure 4.3

The group acting on $\mathbf{R}^3$ is $\mathbf{Z}\times\mathbf{Z}\times\mathbf{Z}$, and the action **O** has the three generating identifications

$$\begin{aligned}(x_1,x_2,x_3) &\sim_{\mathbf{O}} (x_1+12,x_2,x_3)\\ &\sim_{\mathbf{O}} (x_1,x_2+12,x_3)\\ &\sim_{\mathbf{O}} (x_1,x_2,x_3+12)\end{aligned}$$

The general case should be clear enough by now: while $\mathbf{R}^n$ is the usual infinite (i.e., topologically "noncompact") n-dimensional space, $\mathbf{R}^n/\mathbf{O}$ is finite and without-boundary (i.e., topologically "compact" or "closed") version of that space having the form of an n-torus, which can be imagined either as a circle of circles of ... of circles or as an n-hypercube with opposite "faces" identified. The quotient space derives from an action of $\mathbf{Z}\times\mathbf{Z}\times\cdots\times\mathbf{Z}$ generated by the identifications

$$\begin{aligned}(x_1,x_2,\cdots,x_n) &\sim_{\mathbf{O}} (x_1+12,x_2,\cdots,x_n)\\ &\sim_{\mathbf{O}} (x_1,x_2+12,\cdots,x_n)\\ &\vdots\\ &\sim_{\mathbf{O}} (x_1,x_2,\cdots,x_n+12)\end{aligned}$$

4.3. We now consider the effect of permutations--reassignments of notes to voices, in the manner of a generalized voice exchange, on the spaces $\mathbf{R}^n$. Permutational equivalence (**P**) is different from octave equivalence (**O**) by virtue of having fixed points; that is, there are points in $\mathbf{R}^n$ that are unaffected by one or more group elements (other than the identity) in the action of **P**. These fixed points are exactly those corresponding to a chord with pitch doublings. For example, the Alberti-bass point (60,67,64,67) in $\mathbf{R}^4$ is fixed under one group element in the action of **P**, specifically the exchange of the second and fourth voices. The nearby point (60,67,67,67), on the other hand, is fixed by the six group elements that do not affect the first voice. Points corresponding to quadrupled pitches, such as (67,67,67,67), are fixed by the entire action. (In general, we will use the term "**P**-symmetries" to refer to doublings, triplings, quadruplings, etc.)

Since there are $n!$ ways to assign n notes to n voices, every orbit (set of identified points) of $\mathbf{R}^n$ under the action of **P** that does not involve fixed points consists of $n!$ distinct points. The fixed points (corresponding to **P**-symmetric chords) form boundaries between $n!$ connected regions of $\mathbf{R}^n$, each of which maps into the others under the action of **P** and each of which can stand as a fundamental domain of the quotient space $\mathbf{R}^n/\mathbf{P}$.

Loosely speaking, then we can understand $\mathbf{R}^n/\mathbf{P}$ as comprising $1/n!$ as much space as the covering space $\mathbf{R}^n$. (There are ways to make this statement mathematically reputable, though it would take us too far afield to discuss them here.) Any fundamental domain of $\mathbf{P}$ can be characterized in terms of the contour of the ordered sets it contains; a particularly convenient fundamental domain is the locus of ordered sets $(x_1, x_2, \cdots, x_n)$ whose notes simply appear in ascending order $(x_1 \le x_2 \le \cdots \le x_n)$. This fundamental domain is bounded by the $n-1$ hyperplanes having equations of the form $x_i = x_{i+1}$.

Before considering the quotient spaces modulo $\mathbf{P}$ in detail, we need to make some preliminary definitions. A *line of transposition* is defined as any of the infinite number of parallel diagonal lines in $\mathbf{R}^n$ comprising voicewise pitch transpositions, and the *line of unisons* as the particular line of transposition that passes through the origin (so called because it consists of exactly those points corresponding to "chords" in which every voice has the same pitch). The system of equations defining the line of unisons in $\mathbf{R}^n$ is simply $x_1 = x_2 = \cdots = x_n$.

Now consider the case $n = 2$; our formula says that $\mathbf{R}^2/\mathbf{P}$ consists of $1/n! = 1/2$ as much space as $\mathbf{R}^2$. The relevant group of permutations is the symmetric group $\mathbf{S}_2$, which is isomorphic to the cyclic group $\mathbf{Z}_2$ and every other group of order 2. Since $\mathbf{R}^2$ is a plane, the line of unisons (like any line) divides it into two half-planes. These two half-planes are the fundamental domains of the $\mathbf{S}_2$ action on $\mathbf{R}^2$. That is, the nonidentity element of $\mathbf{S}_2$ under the action $\mathbf{P}$ moves every point in one half-plane into the other. Specifically, the motion carried out by that permutation is equivalent to a reflection in the line of unisons. We may therefore construct a model of the space $\mathbf{R}^2/\mathbf{P}$ by throwing away either of the half-planes into which the line of unisons divides $\mathbf{R}^2$ (it doesn't matter which), and placing a mirror at the line of unisons.

Our next step is the case $n = 3$; now our formula tells us that $\mathbf{R}^3/\mathbf{P}$ consists of $1/n! = 1/6$ as much space as $\mathbf{R}^3$. The figure below shows that the canonical fundamental region, bounded by two (hyper)planes intersecting at the line of unisons, takes the form of a wedge when viewed along the line of unisons. Points close to the thin edge of the wedge correspond to chords with notes packed closely together.

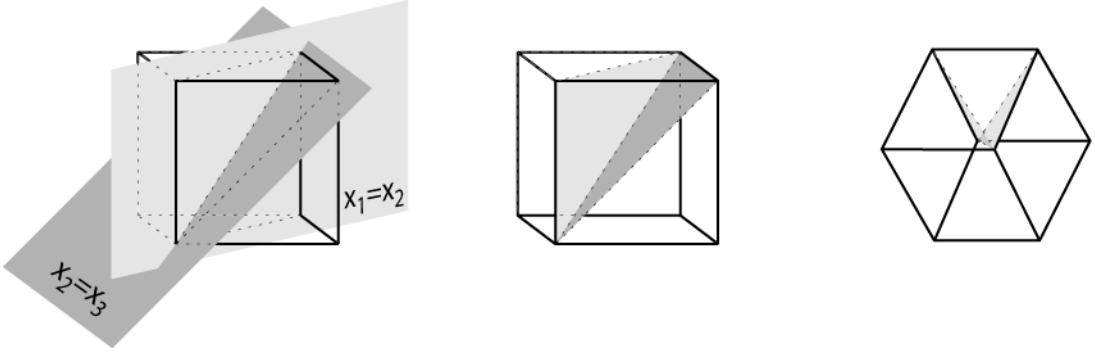


Figure 4.4

The relevant group of permutations is the symmetric group $\mathbf{S}_3$, which is of order 6 but is not cyclic; thus the action of the group on $\mathbf{R}^3$ corresponds not to sixfold rotation about the line of unisons, but reflection in each of the hyperplane mirrors. One of these

reflections exchanges a point's first and second co-ordinates; the other exchanges its second and third coordinates. The two reflections combine to form the full set of permutations of three objects, including the two circular permutations, which correspond to 120-degree rotations around the line $x_1 = x_2 = x_3$. Thus, from the standpoint of **P**-equivalence, $\mathbf{R}^3$ has a threefold rotational symmetry, each section of which has bilateral symmetry.

The general case follows a pattern that readily suggests itself. The quotient space $\mathbf{R}^n/\mathbf{P}$ has the form of an *n*-dimensional wedge bounded by $n-1$ hyperplane mirrors intersecting in *n* *n*-flats (1 line, 2 planes, and so on). The wedge's thin edge runs along the line of unisons, and points close to this edge correspond to cluster-type chords.

We have seen the effects of **O**- and **P**-equivalence individually on $\mathbf{R}^n$; now let us consider the question of how they interact to form equivalence classes of unordered pitch-class sets. The resulting space is very important from a music-theoretical perspective, as it represents the space of multisets, or unordered chords of pitch-classes. The geometrical models shown in Example 1, including the familiar circle of fifths, the tonnetz, and many others, all represent subsets of these spaces.

As shown in Tymoczko (unpublished), the space $\mathbf{R}^2/\mathbf{OP}$ is a Möbius strip--a two-dimensional figure usually described in three dimensions as a circular band with a half-twist that has just one side and one edge:

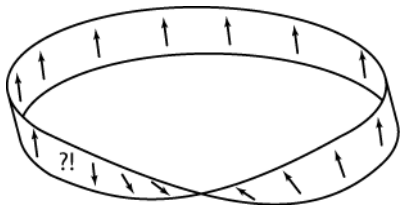


Figure 4.5

We can arrive at this figure by cutting a torus in half, and identifying two of the remaining edges. The next figure, which derives from Tymoczko (unpublished), shows how to do this.

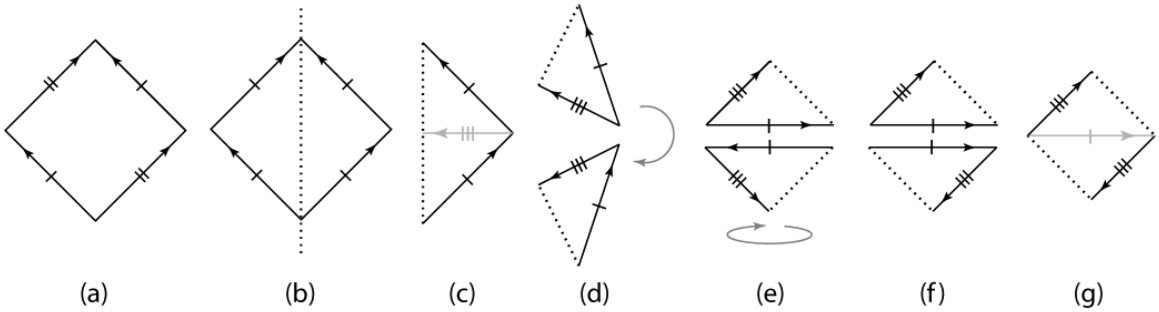


Figure 4. 6

We begin with a square (a) that has its opposite sides (marked I and II) identified in the familiar way that connects the square up like a torus. We identify all points related by

reflection through the line of unisons, as shown in (b). The resulting figure is a triangle triangle (c) whose base is a mirror and whose legs (I) are identified. To see clearly that this structure is a Möbius strip requires some additional surgery: slice the triangle into congruent halves along its altitude (d), compensating with an identification of the two new edges marked III, then reunite the original identified edges labeled I, which necessitates flipping one of the triangles over (e, f). The two triangles have now fit together into a square (g); two of its opposed edges are mirrors, and the other two are identified. To bring the identified edges together in Euclidean space, a half-twist needs to be executed in the third dimension--the characteristic half-twist of the Möbius strip, which also reunites the mirror edges (the line of unisons) into a single edge.

This mirror is the line of unisons; the line parallel to the mirror and furthest from it is the locus of tritones. The transpositional symmetry of the tritone under **O**-equivalence is reflected in the peculiar geometry of the Möbius strip: the line of tritones is exactly half as long as any other line of transposition in $\mathbf{R}^2/\mathbf{OP}$. While that may be counterintuitive at first, it reflects the musical fact—familiar in the discrete case—that there are "half as many tritones" as there are other intervals. We generalize this fact to the continuous case (where there is an uncountably infinite number of every type of interval) using the notion of *length*. The result can be expressed in terms of a readily graspable musical situation. Imagine first two Shepard tones executing a constant upward glissando, in lockstep a whole-step apart; then imagine the same situation but with a tritone separating the Shepard tones instead of a whole step. In the case of the tritone, the glissando will appear to cycle twice as fast.

A significant feature of the Möbius strip is that it is a nonorientable surface. What this means is translation along certain closed paths is equivalent to reflection. Thus, if you lived in a nonorientable space, it would be possible for you to take a trip such that on your return, everything seemed backwards to you, including writing, whirlpools, handedness, clockfaces, and so on; at the same time, everybody who did not go on such a trip would insist that you're the one who came back backwards. (The situation could be remedied by retracing your journey.) In the picture of the Möbius strip on the previous page, arrows drawn on the strip show that a trip through the space in the direction of transposition causes such a reversal. On the Möbius strip, although "north" and "south" can be fixed -- in the present case, we might say that north is the direction in which one transposes chords up by a small interval -- there is no way to fix east and west. (**NB**: this is also true of some *orientable* spaces, such as the space $\mathbf{R}^3/\mathbf{OP}$.) Again, there is a ready musical interpretation of this oddity, which has to do with the question of direction in pitch-class space. While it is clearly possible to orient one's experience of pitch-class in terms of ascent and descent under most circumstances, what is not possible is an absolute orientation with respect to relative motion between lines. That is, in pitch-class space there cannot be a way to settle the question whether two lines are converging or diverging solely on the basis of how much and in which direction each line is moving, as can be done in pitch space. Instead, the question is typically settled with reference to the change in interval class (shortest distance) between two lines; on a Möbius strip with fixed north and south, this is something like defining east and west in terms of whether one is

moving from the middle of the strip to its edge or vice versa. Many quotient spaces modulo octave equivalence turn out to be nonorientable.[25]

Turning our attention now to the quotient of $\mathbf{R}^3$ modulo **O** and **P**, we recall that the 3-torus $\mathbf{R}^3/\mathbf{O}$ can be modeled as a cube with identifications (Figure 4.3). To take the quotient modulo **P** we can take the image of any wedge-shaped fundamental domain of **P** in the torus. As Figure 4.4(b) suggests, this is a tetrahedron having two pairs of congruent faces. By reasoning analogous to that we undertook in the two-note case, the mirrors that bound this tetrahedron reflect the identified walls of the cube onto each other so that both pairs of faces of the tetrahedron are identified. The result, unlike the two-dimensional case, is an *orientable* surface,[26] which again has a single circular edge consisting of the line of unisons. It also has a single face, on which all trichordal pitch-class sets with one degree of **P**-symmetry (one doubling) are located.

It is one thing to know that $\mathbf{R}^3/\mathbf{OP}$ is a tetrahedron with identifications; it is quite another to visualize the identification. The industrious reader who makes a number of models of the tetrahedron and glues them together according to the indicated procedure will find that the result is a prism whose section is an equilateral triangle. In fact, the tetrahedron itself can be sliced into three pieces and rearranged so as to produce a piece of this prism, as Figure 4.7 shows.

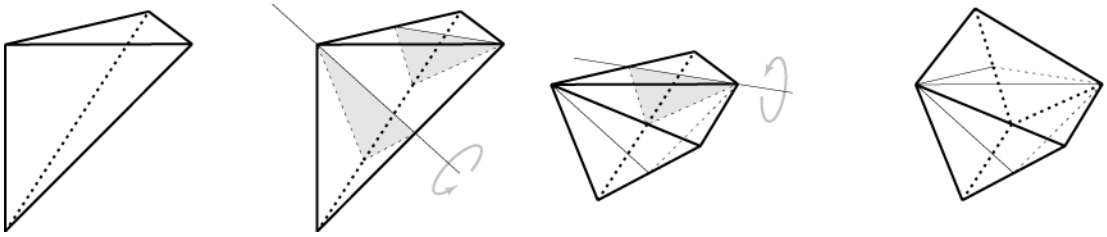


Figure 4.7

(Alternatively, Tymoczko [unpublished] provides an algebraic proof.) Fanciful though it may be, this exercise demonstrates, in connection with our earlier discussion of the Möbius strip, a general feature of the $\mathbf{R}^n/\mathbf{OP}$ family of spaces. Figure 4.6 showed a triangle with identifications reassembled into a square with identifications; here we have a tetrahedron with identifications reassembled into a prism with identifications. As the identifications in the Möbius strip require the half-twist that give it a peculiar geometry, so do the identifications on our prism require a one-third twist of the triangular faces before they can be joined. This twist is what gives the final product a single face and a single edge: a path that traverses one edge of the prism will move to another edge following the twisted identification, and to the third after the next twisted identification. Similar observations can be made about paths along the face of this space.

As we prepare to generalize, recall that an *n*-simplex is an *n*-dimensional figure with $n+1$ vertices in which each pair of vertices determine an edge, each triplet of vertices determine a face, and so on. A 1-simplex is a line segment; a 2-simplex is a triangle; a 3-

[25] In this connection see, e.g., Peck 2003 ("Klein-Bottle Tonnetze").

[26] More generally, the spaces $\mathbf{R}^n/\mathbf{OP}$ are orientable only when n is odd.

simplex is a tetrahedron, and any *n*-simplex can be characterized inductively as the result of pulling an $(n-1)$-simplex out into a new dimension while contracting it to a point.

We may now make the general statement that a fundamental domain for the space $\mathbf{R}^n/\mathbf{OP}$ is a "right *n*-simplex" whose vertices are the points

$$(0,0,\cdots,0,0)$$
$$(0,0,\cdots,0,12)$$
$$(0,0,\cdots,12,12)$$
$$\vdots$$
$$(0,12,\cdots,12,12)$$
$$(12,12,\cdots,12,12)$$

By identifying the appropriate faces, we obtain a shape that can be described as the product of an n-1 simplex with a circle, modulo the action that cyclically permutes the simplex's vertices while rotating the circle through $1/n$ of a turn ($360/n$ degrees). For more information, see Tymoczko (unpublished).

Each quotient space $\mathbf{R}^n/\mathbf{OP}$ is bounded by a single hyperplane mirror that (when $n>2$) intersects with itself in such a way as to produce a single 1-dimensional singularity on the line of unisons, and various combinations of other mirrors. There is at least one *i*-dimensional singularity for each $0<i<n$, corresponding to multisets having just *i* distinct notes. In the 2- and 3-note cases, there is only one permutational singularity for each *i*. In the 4-note case, however, we find two intersecting but distinct plane singularities $(i=2)$ arising from **P**. One plane contains chords of the form $(x,x,x,y)_{\mathbf{OP}}$ and has the same form as $\mathbf{R}^2/\mathbf{O}$, a 2-torus. The other contains chords of the form $(x,x,y,y)_{\mathbf{OP}}$ and has the same form as $\mathbf{R}^2/\mathbf{OP}$, a Möbius strip. A detailed study of why these have different forms and what the general principles are would derail us unduly at this point, but the reader is encouraged to give some thought to the matter as an exercise.

The curious "short circuit" of the tritone dyad in $\mathbf{R}^2/\mathbf{OP}$ has an analogue for any chord having transpositional invariance. In $\mathbf{R}^3/\mathbf{OP}$, the augmented triads are located on a line that runs through the center of the twisted prism and that is one-third the length of all other lines of transposition in the space. If and only if *n* is prime, there is exactly one class of chords (up to **T**) having transpositional invariance (the equal divisions of the octave), and therefore exactly one short circuit among the lines of transposition, having $1/n$ the length of normal lines of transposition. For composite *n*, the situation is different. In $\mathbf{R}^4/\mathbf{OP}$, for example, the diminished triads (the equal divisions of the octave) have the expected quarter-length circuit, but there is also a band of chords having the form $(x,y,6+x,6+y)_{\mathbf{OP}}$ whose lines of transposition have half the normal length. We will take up this issue further in connection with the spaces $\mathbf{R}^n/\mathbf{OPT}$, once we have had a chance to work with **T**-equivalence, to which we now turn.

Musically, the spaces $\mathbf{R}^n/\mathbf{OP}$ are important because they allow us to represent any possible voice-leading between unordered pitch-classes. They therefore provide a powerful analytical tool for representing a wide range of musical phenomena. Furthermore, the structure of these spaces allows us to see how chord-structure, and in

particular the symmetry or near-symmetry of a chord, produces smooth voice-leading. For further details, see Tymoczko (unpublished).

§5. Classification of chord spaces II: **T**-equivalence.

5.0. The next four classes of space we will study are the same as the first four, but with the additional complications of transpositional equivalence (**T**). We saw that **O** and **P** represent structurally rather different types of actions--**P** actions have fixed points and manifest themselves as mirror boundaries, while **O** actions have no fixed points and manifest themselves by wrapping spaces in on themselves. The actions of **T** are structurally different from both **O** and **P** actions in still another respect: since the basic structure of $\mathbf{R}^n$ involves a continuous model of pitch space, the group of transpositions is *itself* continuous: for any two transpositions $\mathbf{T}_x$ and $\mathbf{T}_y$ there is another transposition that lies "between" them. By contrast **O**, **P**, and **I** correspond to *discrete* (and in the case of **P** and **I**, *finite*) groups.

The continuity of **T** means that whereas $\mathbf{R}^n$ modulo **O** and/or **P** are n-dimensional spaces, any quotient space of $\mathbf{R}^n$ modulo **T** is at most $(n-1)$-dimensional. To see why this is the case, recall our so-called lines of transposition, which are the lines parallel to the diagonal line of unisons $x_1 = x_2 = \cdots = x_n$. These lines fill the space $\mathbf{R}^n$, and they are the equivalence classes of **T**. The quotient space $\mathbf{R}^n/\mathbf{T}$, therefore, is such that (in the terminology of section 000) $\mathbf{R}^n$ is a "$\mathbf{R}^n/\mathbf{T}$ of lines." Since we also know that $\mathbf{R}^n$ is $\mathbf{R}\times\mathbf{R}^{(n-1)}$, or a "$\mathbf{R}^{(n-1)}$ of lines ," it must be the case that $\mathbf{R}^n/\mathbf{T}$ has the same topological structure as $\mathbf{R}^{(n-1)}$.

A more intuitive way to state the same argument begins with the premise that the **T**-equivalence class of any chord x includes its $(-x_1)$-transpose, whose coordinates are $(0, x_2 - x_1, \cdots, x_n - x_1)$. Let this latter chord be called the *FB-representative* of the equivalence class, since it can be interepreted as a figured-bass-like collection of intervals relative to the referential "first" note. [27] Clearly every **T**-class has exactly one FB-representative. The set of all FB-representatives in $\mathbf{R}^n$ constitute a single *hyperplane*, by which we mean an instance of $\mathbf{R}^{n-1}$ embedded "flatly" in $\mathbf{R}^n$ (i.e., a point in $\mathbf{R}^1$, a line in $\mathbf{R}^2$, a plane in $\mathbf{R}^3$, and so on). In other words, the mapping of points in $\mathbf{R}^n$ to the representatives of their **T**-classes constitutes a projection onto this hyperplane; this projection discards one dimension's worth of information about the structure of $\mathbf{R}^n$. The discarded dimension contains information about the specific "bass note" of our figured-bass structures.

While the connection with figured bass makes this particular way of projecting $\mathbf{R}^n$ onto a hyperplane conceptually attractive, it has an unfortunate geometrical side effect that bears exploration. The following figure shows the chord (60,64,67) and the six chords that are 2 semitones away under both Euclidean and taxicab geometry (b).

[27] Compare Morris (JMTP) and Roeder 1987 (JMT).

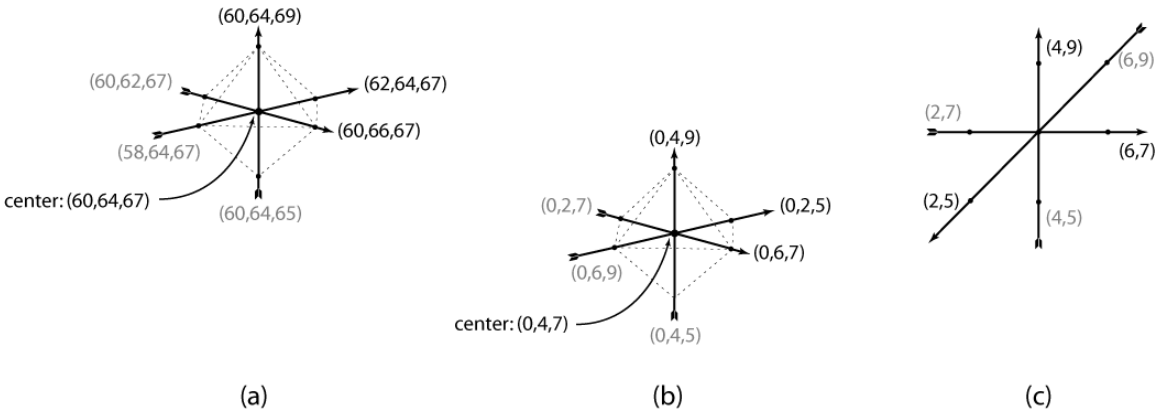


Figure 5.1

At (b) these chords are replaced with their FB-representatives, and at (c) they are plotted the FB-representatives on their (hyper)plane of projection after throwing away the vestigial first coordinate. The figure clearly shows that the projection has introduced a distortion: in the (hyper)plane of projection, the FB-representatives of the **T**-classes $(0,2,5)_{\mathbf{T}} = (62,64,67)_{\mathbf{T}}$ and $(0,6,9)_{\mathbf{T}} = (58,64,67)_{\mathbf{T}}$ are further away from the FB-representative of the "central" **T**-class $(0,4,7)_{\mathbf{T}} = (60,64,67)_{\mathbf{T}}$ than the other four, which are all still 2 semitones away. Under Euclidean geometry, the distorted distances increase by a factor of $\sqrt{2}$, and under taxicab geometry they double! What makes this distortion particularly worrisome is that it is voice-specific, affecting all and only motions involving the referential "bass" voice, since a move in that voice is manifested on the hyperplane of FB-representatives as a move of the same size in *each* of the other voices.

For this reason we will prefer a different canonical projection of $\mathbf{R}^n$ into a hyperplane of **T**-class representatives, one guaranteed to treat all voices equally.

5.1. $\mathbf{R}^n$ can be described as a "line" (or "stack") of *sum-class hyperplanes,* each described by a system of equations having the form $x_1 + x_2 + \cdots + x_n = c$, where c varies over the real numbers. In Euclidean space, these hyperplanes are *orthogonal* (or perpendicular) to the line of unisons. We will focus on the case $c = 0$, which is the unique sum-class hyperplane that includes the origin. As a convenient point of reference, this *sum-0 hyperplane* is the counterpart to the line of unisons, which it intersects at the origin and to which it is orthogonal. We can visualize it as follows: Imagine the two points $(1,1,\cdots,1)$ and $(-1,-1,\cdots,-1)$; these are two points lying on the line of unisons, on either side of the origin. (They are opposite vertices on the hypercube $(\pm1,\pm1,\cdots,\pm1)$, which is centered on the origin.) The sum-class hyperplane is (for either Euclidean or taxicab metric) the locus of points equidistant from these two points. In $\mathbf{R}^2$ the sum-0 hyperplane is a diagonal line traversing the origin from northwest to southeast; it intersects the two points on the square $(\pm1,\pm1)$ that are not on the line of unisons. In $\mathbf{R}^3$ the sum-0 hyperplane is a plane that comes close to each of the six vertices of the cube $(\pm1,\pm1,\pm1)$ that are not on the line of unisons, but does not intersect any of them.

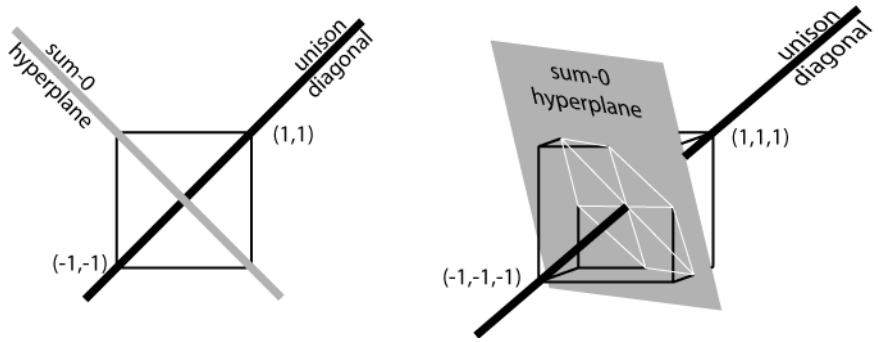


Figure 5.2

The projection of $\mathbf{R}^n$ onto any sum-class hyperplane provides us with a special kind of "coordinate system" referred to as *simplicial* or *synergetic*.[28] (We use quotations around "coordinate system" because, unlike a true coordinate system, in which every point has a unique representation, simplicial coordinates have an extra, or redundant term. We explore this matter below.) This projection, like any other, collapses lines of transposition into points in a hyperplane; these points are **T**-equivalence classes. But it has a special feature that no other projection has: the n unit vectors that form the Cartesian basis of $\mathbf{R}^n$, when projected onto an $(n-1)$-flat in this manner, describe the vertices of a regular $(n-1)$-simplex. In Figures 4.4 and 5.2 we see that a cube becomes a hexagon under this projection. Two of the cube's vertices are collapsed to the center of the hexagon. The other six vertices, projected as the vertices of the hexagon, form two equilateral triangles that are complementary -- i.e., they have a common center but opposite orientations. Each of these complementary triangles' vertices were, prior to the projection of the cube, connected to one or the other of the two "collapsed" vertices of the cube. Thus motion from the center of the hexagon to one of its vertices represents either upward motion or downward motion (relative to the lines of transposition), depending on which of the complementary triangles contains the vertex. (See Figure 5.3.)

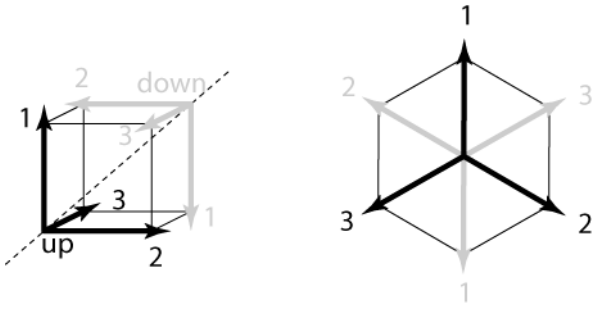


Figure 5.3

[28] The latter term is due to R. Buckminster Fuller, who felt it to be superior to Cartesian geometry. Callender (2003) describes these coordinates in the trichordal case.

The advantage of this type of projection is that it preserves many distance relationships, although there are quirks that result from the factoring-out of transposition. For example, equal motion in the same direction in all three voices is the same as no motion at all, and equal motion in the same direction in two voices is equivalent to the same amount of motion in the opposite direction in the third voice. The advantage of the resulting coordinate system is that we can use the n coordinates of $\mathbf{R}^n$ to label points in the $(n-1)$-dimensional $\mathbf{R}^n/\mathbf{T}$ without working out the trigonometric details. The fact that there is a redundant coordinate can be dealt with by simply ignoring one of them. For example, we could label a trichordal T-equivalence class (0, x, x + y) using any of the coordinates (x, x + y), (y, x + y), or (x, y). Figure 5.4 shows a small portion of $\mathbf{R}^3/\mathbf{T}$, labeling points by the latter method.

5.2. We are now in a position to dispense relatively quickly with the quotient spaces of $\mathbf{R}^n$ modulo **T** and either **O** or **P**. The case of $\mathbf{R}^n/\mathbf{OT}$ is particularly easy, as it is topologically (but not geometrically) identical to $\mathbf{R}^{n-1}/\mathbf{O}$. This space takes the form of a $(n-1)$-torus, like $\mathbf{R}^{n-1}/\mathbf{O}$, but it inherits the simplicial geometry of $\mathbf{R}^n/\mathbf{T}$.

$\mathbf{R}^n/\mathbf{PT}$ is only slightly more complicated. Recall that **P** is an action of $\mathbf{S}_n$ on $\mathbf{R}^n$ that effectively slices it into a wedge bounded by $n-1$ hyperplane mirrors intersecting the line of unisons. In the case of $\mathbf{R}^n/\mathbf{PT}$, **P** is still an action of $\mathbf{S}_n$, but this time its domain is effectively $\mathbf{R}^n/\mathbf{T} \cong \mathbf{R}^{n-1}$, which it slices into a wedge bounded by $n-1$ hyperplane mirrors intersecting the line of unisons. $\mathbf{R}^3/\mathbf{PT}$, for example, is a 60-degree slice of the plane $\mathbf{R}^3/\mathbf{T}$. We could think of it as a triangle with two infinite sides; likewise, $\mathbf{R}^4/\mathbf{PT}$ is a tetrahedron with three infinite faces, and $\mathbf{R}^n/\mathbf{PT}$ is an infinitely extended $(n-1)$-simplex. and The following table summarizes the relationships between $\mathbf{R}^n/\mathbf{P}$ and $\mathbf{R}^n/\mathbf{PT}$.

	$\mathbf{R}^n/\mathbf{P}$			$\mathbf{R}^n/\mathbf{PT}$		
n	Base space	Boundaries		Base space	Boundaries	
		no.	type		no.	type
1	$\mathbf{R}^1$ (line)	0	n/a	$\mathbf{R}^0$ (point)	0	n/a
2	$\mathbf{R}^2$ (plane)	1	$\mathbf{R}^1$ (line)	$\mathbf{R}^1$ (line)	1	$\mathbf{R}^0$ (point)
3	$\mathbf{R}^3$ (space)	2	$\mathbf{R}^2$ (plane)	$\mathbf{R}^2$ (plane)	2	$\mathbf{R}^1$ (line)
4	$\mathbf{R}^4$	3	$\mathbf{R}^3$ (space)	$\mathbf{R}^3$ (space)	3	$\mathbf{R}^2$ (plane)

We now come to a second very important music-theoretical space: $\mathbf{R}^n/\mathbf{OPT}$. What makes this space important is the interpretation of its contents: equivalence classes of chords up to permutation of voices, octave equivalence, and transposition. In other words, we have finally reached something like the pcset-class, although we have not yet tackled the question of inversion.

Let us begin with some generalities. There are two kinds of singularities possible in the quotient spaces $\mathbf{R}^n/\mathbf{OPT}$. The first kind arises from doublings among voices, or permutational symmetries. These lie on hyperplane mirrors and intersections of these, which are familiar as the boundaries of the wedge-shaped spaces $\mathbf{R}^n/\mathbf{P}$, $\mathbf{R}^n/\mathbf{OP}$, and $\mathbf{R}^n/\mathbf{PT}$.

The second kind arises from transpositional invariance of chords, a property that is dependent upon octave equivalence.[29] The formations these singularities take are highly cardinality-dependent, the general rule being as follows: For every prime factor p of n other than 1 there is one singularity in the form of a $(n/p-1)$-flat containing all and only chords that can be partitioned into instances of a p-fold equal division of the octave; all are transpositionally invariant.[30] The following table shows all the possibilities for small cardinalities:

n	p	$(n/p-1)$	TINV chords
2	2	0	$(0,6)_{\mathbf{OPT}}$
3	3	0	$(0,4,8)_{\mathbf{OPT}}$
4	2	1	$(0,x,6,6+x)_{\mathbf{OPT}}$
5	5	0	$(0,2.4,4.8.,7.2,9.6)_{\mathbf{OPT}}$
6	2	2	$(0,x,y,6,6+x,6+y)_{\mathbf{OPT}}$
	3	1	$(0,x,4,4+x,8,8+x)_{\mathbf{OPT}}$

A brief inspection of the table should suffice to show that the dimensionality of the singularity corresponding to each type of transpositionally invariant chord is tied to the number of degrees of freedom that type has, which is in turn tied back to the numerical properties of n. The two types of transpositionally invariant hexachord, for example, are (a) those containing two copies $(0,4,8)_{\mathbf{OPT}}$ and (b) those containing three copies of $(0,6)_{\mathbf{OPT}}$. The singularity in $\mathbf{R}^6/\mathbf{OPT}$ consisting of type-(a) TINV hexachords has precisely the same structure as $\mathbf{R}^2/\mathbf{OPT}$, each hexachord $(0,x,4,4+x,8,8+x)_{\mathbf{OPT}}$ corresponding to the dyad $(0,3x)_{\mathbf{OPT}}$. Likewise, the singularity consisting of type-(b) TINV hexachords has precisely the same structure as $\mathbf{R}^3/\mathbf{OPT}$, each hexachord $(0,x,y,6,6+x,6+y)_{\mathbf{OPT}}$ corresponding to the trichord $(0,2x,2y)_{\mathbf{OPT}}$. These two singularities intersect at exactly one point, $(0,2,4,6,8,10)_{\mathbf{OPT}}$, which is not only transpositionally invariant, but also the maximally even set of cardinality 6.

Generalizing the foregoing observations about hexachords, we observe for any prime factor p of an integer n, there is a singularity in $\mathbf{R}^n/\mathbf{OPT}$ containing those transpositionally invariant chords that can be characterized as the "Boulez product" of the p-fold equal division of the octave with some chord of cardinality n/p. This singularity has the same global structure as $\mathbf{R}^{n/p}/\mathbf{OPT}$, possibly possessing singularities of its own but otherwise locally resembling a $(n/p-1)$-flat. The intersection of any two TINV singularities corresponding to prime factors p and q is a $(n/pq-1)$-flat.

Let us now see how these generalities play out in the case of trichords. The space $\mathbf{R}^3/\mathbf{PT}$, as we have seen, can be described as a 60-degree wedge of the plane bounded by mirrors. This is equivalent to observing that the space is a plane with three singularities: two hyperplanes (lines in this case) and one corner at which they intersect. Each hyperplane contains mostly dyads-with-doublings. On one hyperplane we find dyads with the lower note doubled, and on the other we find dyads with the higher note

[29] The one point in the space $\mathbf{R}^0$ is transpositionally invariant in a trivial sense, but otherwise transpositional invariance requires **O**-equivalence.

[30] Cohn, "Properties and Generability...," Quinn diss.

doubled. They intersect at a point corresponding to the tripled singleton. Except for these singularities, which bound the space on two sides, the space locally resembles a plane at every point.

Suppose we now identify octave-equivalent points. Locally the space will continue to resemble a plane (just as, for example, a torus does), but its global structure changes significantly. These changes concern the two classes of singularity: permutational and transpositional. The permutational singularities, the two hyperplanes (lines) of doubled notes, are themselves identified, since chords with their lower notes doubled are identified with chords having doubled higher notes. For example:

$$(0,0,3)_{\mathbf{OPT}} \sim (12,12,3)_{\mathbf{OPT}} \sim (0,9,9)_{\mathbf{OPT}}$$

Thus in $\mathbf{R}^3/\mathbf{OPT}$ the images of all dyads-with-doublings are located on a single line, with their specific locations determined by the ordered pc interval from the doubled note to the other note.

So these lines, which meet at a 60-degree angle at the origin (a corner singularity), are the same line, which has finite length, and which is the only boundary of the space. It has the structure of a circle with one corner "pinched" into it, the shape of a teardrop.

So much for the permutational singularities; let us consider the transpositional singularities. According to our earlier observations, since the only prime factor of 3 greater than 1 is 3, there is just one transpositional singularity, and since $3/3-1=0$, it takes the shape of a 0-flat: a point. This is, of course, the point $(0,4,8)_{\mathbf{OPT}}$. Because of this chord's threefold symmetry, the neighborhood of the point singularity contains just one-third of the usual amount of space; thus it is a cone point that gives triple vision to those standing before it.

The space $\mathbf{R}^3/\mathbf{OPT}$, then, is a 2-dimensional surface with a single teardrop-shaped boundary and a cone point. The following figure shows that Callender's (2003) kite-shaped trichord space, subjected to an identification of its short lets, rolls up into just such a shape. When looking at the figure it will be useful to keep the following three identities in mind:

$$(0,5,5)_{\mathbf{OPT}} \sim (0,0,7)_{\mathbf{OPT}}$$
$$(0,6,6)_{\mathbf{OPT}} \sim (0,0,6)_{\mathbf{OPT}}$$
$$(0,7,7)_{\mathbf{OPT}} \sim (0,0,5)_{\mathbf{OPT}}$$

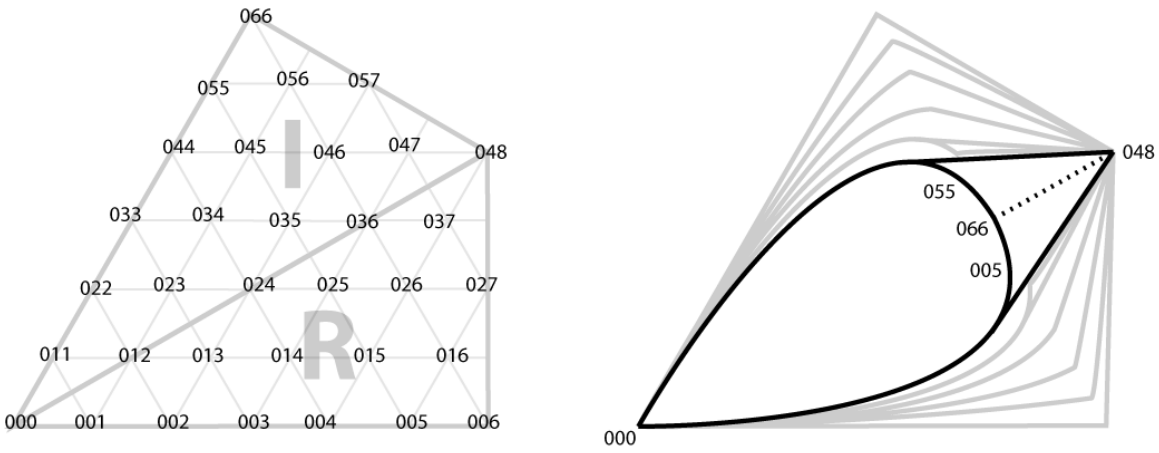


Figure 5.4

Notice that while Callender's kite seems to have four edges, the space it represents has just one, the teardrop-shaped locus of dyads-with-doublings. The lines that connect the two tritone points to the augmented triad are identified and therefore the same; in fact, neither the tritone nor the other points on the line that runs to the augmented triad in the kite representation, including $(0,2,7)_{\mathbf{OPT}}$, are special at all, and do not represent a boundary or any kind of singularity in the space. In fact, a flat representation of the space can be made from the three-dimensional representation can be made by making a cut from any point on the boundary to the singularity at $(0,4,8)_{\mathbf{OPT}}$, identifying the cut edges, and flattening the resulting figure:

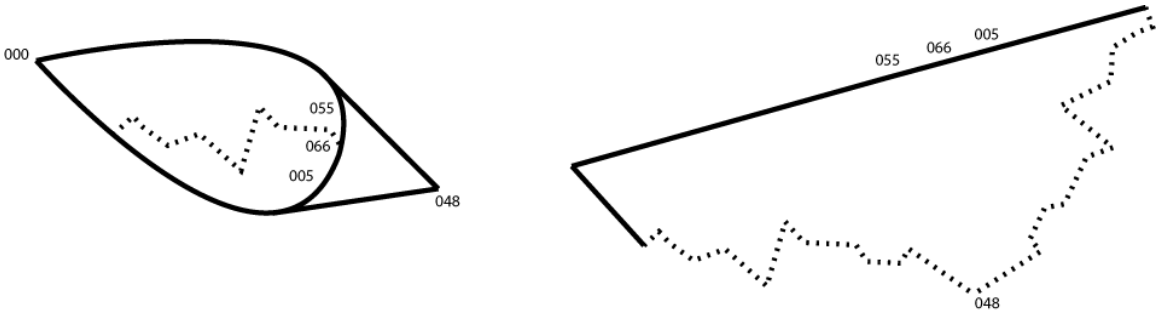


Figure 5.5

Such a representation is geometrically just as accurate as the kite, or the wide triangle that results from making the cut from the origin to $(0,4,8)_{\mathbf{OPT}}$. There is, however, a handy (if superficial) feature that only the kite-shaped and triangular representations have, which is a bilateral symmetry corresponding to the inversional relationship between one half of the space (marked R in an earlier figure) and the other (marked I). This inversional relationship is of no consequence to the structure of any of the **OPT** spaces, but will become quite important when we take up **I**-equivalence in part 6 of this paper.

Another interesting perspective on smooth voice leading possibilities among trichord types is afforded by surveying the view from the singularity $(0,4,8)_{\mathbf{OPT}}$. Someone standing at that point might imagine that the space looks like Figure 5.6, which is derived

from the space's extrinsically conic manifestation by viewing it from a point outside the space above the cone point. In this perspective, $(0,4,8)_{\mathbf{OPT}}$ is at the center of the action, and not sticking out on sime kind of point at the periphery of the space. Because $(0,4,8)_{\mathbf{OPT}}$ is a singularity having only 120 degrees of space around it, we have to bend the lines to get them to join up as they are supposed to. But since the actions that generate this space all preseve local distances, there is not any "real" curvature in the neighborhood of any point outside the singularity. Instead, an observer standing near a singularity like $(0,4,8)_{\mathbf{OPT}}$ will see multiple copies of herself and of everything in her neighborhood. (In this space, she will ordinarily see two additional copies of herself such that all three of her are standing at the vertices of an equilateral triangle whose centroid is the singularity.)

Close to the singularity are the triads $(0,3,7)_{\mathbf{OPT}}$ and $(0,4,7)_{\mathbf{OPT}}$, whose proximity to the nearby cone point is well known to be an important source of the properties studied by neo-Riemannian theory. The line of "inversional symmetry" runs horizontally through the center of the figure. T-equivalence classes can be linked to themselves by smooth voice-leading to the extent that they are close to either the maximally even set $(0,4,8)_{\mathbf{OPT}}$ or the mirror boundary containing duplications. (For example, a C-major triad can be smoothly voice-led to an E-major or A-flat-major triad, and any chord containing a semitone can be smoothly voice-led to another instance of itself by exchanging the voices that are a semitone apart.) T-equivalence classes close to the horizontal line of inversional symmetry can be linked to their inversions by smooth voice-leading.[31]

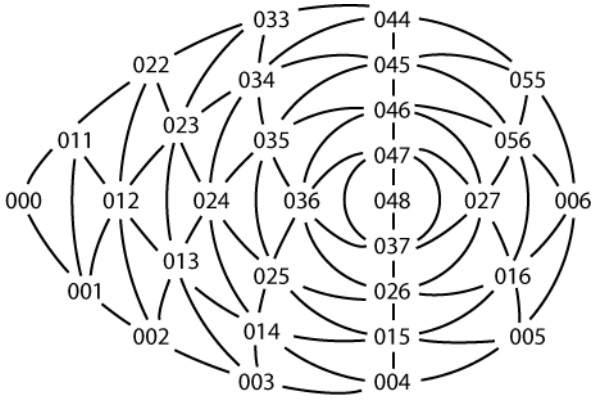


Figure 5.6

The analogous space for tetrachords, $\mathbf{R}^4/\mathbf{OPT}$, is as closely related to Cohn's (2003) tetrahedral model as the trichord space is to Callender's triangular model. Let us review the structure we expect to find in this space. Locally, it ought to resemble 3-dimensional space with a simplicial geometry. It should be bounded by a single plane containing tetrachords-with-doublings, and this plane should wrap in on itself so as to intersect with itself at the origin, where it looks like a corner of three planes. What look like three lines of intersection at the origin are actually two, one for twice-doubled dyads of the form $(x,x,y,y)_{\mathbf{OPT}}$ and one for tripled dyads of the form $(x,x,x,y)_{\mathbf{OPT}}$. The line of twice-

[31] For a generalized algebraic proof of these statements, see Tymoczko (unpublished).

doubled dyads ends in a mirror and has the structure of $\mathbf{R}^2/\mathbf{OPT}$, and the line of tripled dyads loops back on itself like the edge of $\mathbf{R}^3/\mathbf{OPT}$. The remaining permutational singularity is the simplest one: the quadrupled singleton at the origin.

The situation we have just described, unfortunately, cannot actually be represented in three dimensions without recourse to identifications, much like the 2-dimensional space $\mathbf{R}^3/\mathbf{OPT}$ requires a third dimension in which to wrap around and resolve the identifications. (Or, more familiarly, the way the humble 1-dimensional circle requires a second dimension in Euclidean space.) Unfortunately, even if a fourth dimension were easily available, it would not help us understand even the 2-dimensional boundary of $\mathbf{R}^4/\mathbf{OPT}$. As the edge of $\mathbf{R}^3/\mathbf{OPT}$ was a circle pinched once at the origin to form a corner, the edge of $\mathbf{R}^4/\mathbf{OPT}$ is an oddly behaved nonorientable 2-dimensional surface called a *projective plane*. (A projective plane, unlike a Möbius strip cannot be represented in three dimensions; therefore even a fourth Euclidean dimension would not help us depict $\mathbf{R}^4/\mathbf{OPT}$.) The figure below shows a triangular surface with the necessary identifications. All three corners of the triangle are identified; they represent the origin. The two halves of the rightmost edge of the triangle are identified by folding them together. Finally, the surface is closed by identifying the two remaining edges in opposite directions: this is the part that cannot be done in "normal" space. One way to imagine this operation as folding the leftmost vertex of the triangle into the newly folded corner on the right, but with a vertical twist that is necessary to line up the points properly. Even though it is difficult to imagine how the operation is completed, the exercise is sufficient to demonstrate that from the point of view of the origin, $\mathbf{R}^4/\mathbf{OPT}$ looks like a 3-sided wedge, just like $\mathbf{R}^4/\mathbf{PT}$, and that this wedge has one 2-dimensional boundary and two 1-dimensional boundaries. (The dashed lines in the figure are for reference only; they do not indicate additional creases or singularities.)

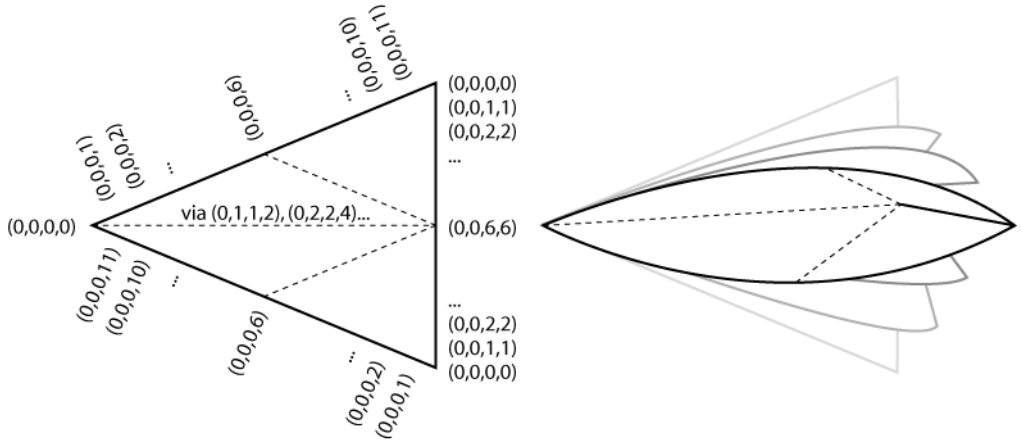


Figure 5.7

Despite the strange way in which this triangular boundary is twisted, its interior is nothing other than the double tetrahedron discussed by Cohn (2003) in connection with noninvertible tetrachord-classes. The duplications Cohn notices in his tetrahedral model

were a source of some concern; by this point in our work it should be clear that they are simply a manifestation of Cohn's attempt to depict the structure in three-dimensional Euclidean space.

As for the transpositional singularities, we expect to find here a line singularity containing tetrachords of the form $(0,x,6,6+x)_{\mathbf{OPT}}$ and having the same structure as $\mathbf{R}^2/\mathbf{OPT}$, a line segment with mirrors at its two ends, $(0,0,6,6)_{\mathbf{OPT}}$ and $(0,3,6,9)_{\mathbf{OPT}}$. Note that $(0,0,6,6)_{\mathbf{OPT}}$ is at the end of two distinct line singularities: a corner singularity of permutational provenance, and a radial singularity resulting from transpositional symmetry. Most points on the corner singularity have order-4 symmetry, and most points on the radial singularity have order-2 symmetry; their common point $(0,0,6,6)_{\mathbf{OPT}}$ has order-8 symmetry.

We also observe that in the geometric conception we have been building up here, the transpositional singularity takes the form of a line segment beginning at the point $(0,0,6,6)_{\mathbf{OPT}}$ on an exterior corner of the space and pointing into the interior. Thus while $(0,3,6,9)_{\mathbf{OPT}}$ appeared on one of the edges of Cohn's tetrahedron, it seems to be buried in the middle of our space. But just as $(0,4,8)_{\mathbf{OPT}}$ is both "in the middle" of $\mathbf{R}^3/\mathbf{OPT}$ with respect to that space's boundary and "on the outside" thanks to the way its symmetry seems to bend its space, and there are ways of representing the space that correspond to each point of view, so is it possible -- even desirable -- to conceive of $(0,3,6,9)_{\mathbf{OPT}}$ on the outside; indeed, on Cohn's double tetrahedron, it lies on an edge formed by two planes meeting at the 90-degree angle characteristic of the chord's fourfold symmetry.

To sum up our rather bewildering exploration of $\mathbf{R}^4/\mathbf{OPT}$, we would like to reiterate that the space can be completely described in terms of its "parent space" $\mathbf{R}^4/\mathbf{PT}$ and the singularities that arise under **O**-equivalence. The parent space has a fairly simple structure: a structurally homogeneous interior of 3-dimensional space with a simplicial geometry, bounded by three mirror (hyper)planes intersecting at the origin and containing all 4-voice chord types that have doublings. Things get decidedly more complicated when we add **O**-equivalence. The boundary wraps in on itself completely, so that the three planes become a single surface; in addition, the three lines that provided the edges of the wedge in $\mathbf{R}^4/\mathbf{PT}$ reduce to two. Finally, the closure of the space under **O** induces singularities on the interior thanks to transpositional invariance. We noted some similarities in structure between, on the one hand, boundary (permutational) singularities and internal (transpositional) singularities areas of $\mathbf{R}^4/\mathbf{OPT}$ and, on the other hand, other **PT**- and **OPT**-spaces.

Alternatively, we can describe the spaces $\mathbf{R}^n/\mathbf{OPT}$ as the quotients of n-1 dimensional simplexes by the action that cyclically permutes their vertices. This description shows that they represent projections of the spaces $\mathbf{R}^n/\mathbf{OP}$ (the space of sets) onto a hyperplane, just as $\mathbf{R}^n/\mathbf{T}$ is a projection of $\mathbf{R}^n$ onto a hyperplane. However, we will defer further generalization about **OPT**-spaces until after our study of **I**-equivalence, to which we now turn.

§6. Classification of chord spaces III: **I**-equivalence.

6.0. Having discussed **O**, **P**, and **T**, we now turn our attention to the eight spaces that incorporate **I**. The first five spaces to be discussed in this section have been selected for their potential to explain the geometric nature of inversional equivalence and to set up the most abstract yet familiar of our spaces, $\mathbf{R}^n/\mathbf{OPTI}$, the space of pcset-classes. The remaining three spaces are rather esoteric and are reserved for the end of this section.

As we have seen in §3.4, the space of one-note chords under $\mathbf{I}_c$-equivalence, $\mathbf{R}^1/\mathbf{I}_c$, is a half line bounded by a single mirror point at $c/2$. This space is the quotient of $\mathbf{R}^1$ modulo the action of $\mathbf{Z}_2$ generated by the identification $x \sim_{\mathbf{I}_c} c-x$. As we have also seen, this action is the *function* $\mathbf{I}_c(x) = c-x$.

From this one-dimensional space we can observe a property of $\mathbf{I}_c$ that generalizes to higher-dimensional cases: the quotient of $\mathbf{R}^1$ modulo $\mathbf{I}_c$ is a half space. A half-space is obtained by discarding the space that lies to one side of a hyperplane. In the case of the action of $\mathbf{I}_c$ on $\mathbf{R}^1$, the hyperplane is simply the point $c/2$, which divides the infinite line into two rays, or half lines. Discarding one of these half-spaces yields the space $\mathbf{R}^1/\mathbf{I}_c$.

This highlights an important difference between **I** and the other equivalences. Every point of $\mathbf{R}^n/\mathbf{O}$ or $\mathbf{R}^n/\mathbf{T}$ is an equivalence class that contains an infinite number of members. One consequence of this is that given a fundamental region of either space, there are an infinite number of copies of this fundamental region in $\mathbf{R}^n$. For example, $\mathbf{R}^2$ can be tiled by an infinite number of 12 x 12 squares each of which can serve as a fundamental domain of **O** in the plane. Likewise, $\mathbf{R}^n$ contains an infinite number of sum-class hyperplanes. In contrast, the typical **P**-equivalence class contains $1/n!$ members, corresponding to the number of permutations of n objects. $\mathbf{R}^n$ is packed by $n!$ copies of the fundamental region of $\mathbf{R}^n/\mathbf{P}$, which agrees with our intuition that $\mathbf{R}^n/\mathbf{P}$ contains $1/n!$ as much space as $\mathbf{R}^n$ (with the necessary caveat that this is, technically speaking, merely a useful fiction). The typical $\mathbf{I}_c$-equivalence class contains exactly two elements: the original point, x, and it's inversion, $\mathbf{I}_c(x)$. In $\mathbf{R}^1/\mathbf{I}_c$, every point $c/2+x$ is an equivalence class containing the pitches $c/2+x$ and $c/2-x$; the sole exception is the fixed point at $c/2$, which contains only a single member. Thus in the general case, we expect that the quotient $\mathbf{R}^n$ modulo $\mathbf{I}_c$ will be in some sense half of $\mathbf{R}^n$.

This hypothesis is applicable to the other spaces of §4-5 modulo $\mathbf{I}_c$. Let **E** be an equivalence class of some combination of **O** and/or **P** and/or **T**. The equivalence class $\mathbf{x}_\mathbf{E}$ (where **x** is any chord) is $\mathbf{I}_c$-equivalent to one and only one **E**-equivalence class, namely, $(\mathbf{I}_c(\mathbf{x}))_\mathbf{E}$. Since the vast majority of **E**-classes will not remain fixed by $\mathbf{I}_c$, we again expect that the quotient $\mathbf{R}^n/\mathbf{E}$ modulo $\mathbf{I}_c$ will be in some sense half of $\mathbf{R}^n/\mathbf{E}$.

6.1. $\mathbf{I}_c$ acts on dyads by a combination of its action on the individual voices of the dyad:

$$\mathbf{I}_c(x_1, x_2) = (\mathbf{I}_c(x_1), \mathbf{I}_c(x_2)) = (c-x_1, c-x_2).$$

Geometrically we can interpret this action as a reflection in the plane $x_1 = c/2$, which sends voice x_1 to $c-x_1$, followed by a reflection in the plane $x_2 = c/2$, which sends the voice x_2 to $c-x_2$. (See Figure 6.1.1(a).) However, the combination of these two reflections lines in the plane is equivalent to a *reflection in the point* at which the lines

intersect. In this case, we can see that $\mathbf{I}_c$ acting on $\mathbf{R}^2$ corresponds to a reflection in the point $(c/2, c/2)$, which we will refer to as the *inversion point*. The inversion point for any inversion will lie on the line of unisons.

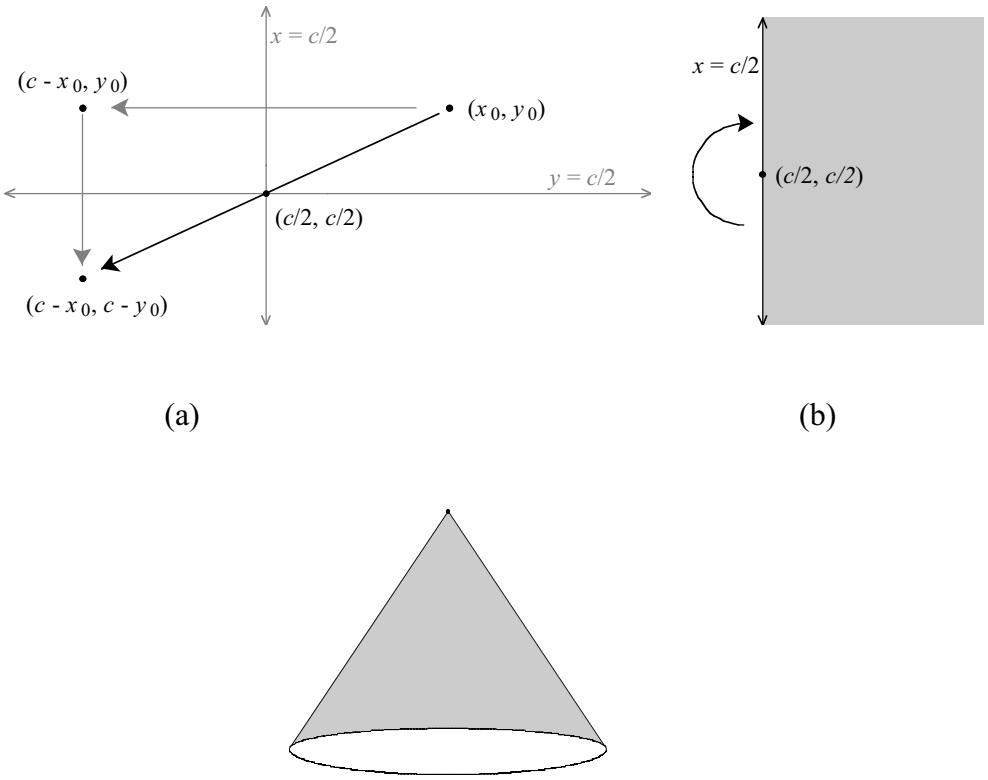
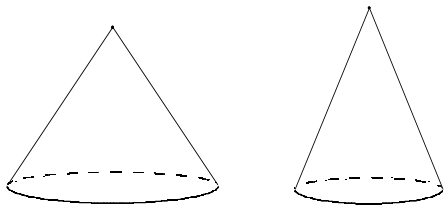


Figure 6.1.1

Any half plane bounded by a line passing through the inversion point is a fundamental domain of the action of $\mathbf{I}_c$ on $\mathbf{R}^2$. For example, take the half plane bounded by the line $x = c/2$, shown in Figure 6.1.1(b). Clearly any point on the right of this line has an $\mathbf{I}_c$-equivalent mate on the left, since this was one of our two mirrors in Figure 6.1.1(a). To obtain the quotient space $\mathbf{R}^2/\mathbf{I}_c$ from this fundamental domain, we must identify points on this line with their reflection in the inversion point, as indicated in Figure 6.1.1(b). Joining the two halves of the boundary line yields an infinite cone with a cone point at the inversion point, the single fixed point for $\mathbf{I}_c$. (See Figure 6.1.1(c).) Note that there is nothing special about the line $x = c/2$. This line is transformed into a half line bounded by the inversion point just as every other line in $\mathbf{R}^2$ that passes through $(c/2, c/2)$. We simply choose one of these lines to serve *temporarily* as a boundary of a half plane. Once $\mathbf{I}_c$-equivalent points on this line are identified, the temporary half space boundary vanishes. Try as we might, we cannot tell by examining Figure 6.1.1(c) exactly which line served this temporary, though useful, purpose.

That $\mathbf{R}^2/\mathbf{I}_c$ is a cone should not be surprising, since in two dimensions reflection in a point is equivalent to a two-fold rotation about this point, and we have already seen that $\mathbf{R}^3/\mathbf{OPT}$, which possesses three-fold rotational symmetry about the equivalence class of maximally-even sets, $(0,4,8)_{\mathbf{OPT}}$, is also a cone. The cone point of $\mathbf{R}^2/\mathbf{I}_c$ differs in two important ways from that of $\mathbf{R}^3/\mathbf{OPT}$. First, since the former arises from a two-fold rotation and the latter a three-fold rotation, we intuitively suspect that the neighborhood about the former is in some sense "larger" than the latter. We can buttress this intuition with some precision: Take a disk of radius r centered on p, which has an area of A.[32] Transforming this disk into a cone by two-fold rotation reduces the area by a factor of two, $A/2$. (See Figure 6.1.2(a) and (b).) A similar transformation by three-fold rotation reduces the area by a factor of three, and so on. It is in this sense that the neighborhood about the cone point of $\mathbf{R}^2/\mathbf{I}_c$ is 3/2 as large as the neighborhood about the cone point for $\mathbf{R}^3/\mathbf{OPT}$ and half as large as the neighborhood about any other point in $\mathbf{R}^2/\mathbf{I}_c$. In any dimension, neighborhoods about singularities arising from the action of **I** on a space always contain "half as much space," or content, as these same points in the parent space. This corresponds to our understanding that a chord either possesses inversional symmetry or it does not; unlike transpositional and permutational symmetry, we do not speak of degrees of inversional symmetry.[33]



(a) (b)

Figure 6.1.2

Second, the cone point of $\mathbf{R}^2/\mathbf{I}_c$ is a 0-dimensional mirror, as can be verified by comparing the path of any line that intersects this point with any path that does not pass through the inversion point. (See Figure 6.1.3, which gives a "bird's eye view" of the cone $\mathbf{R}^2/\mathbf{I}_0$ from directly above the apex.) Any straight path that does not include the inversion point, no matter how close it comes to this point, will continue in a straight line. For example, observe the path (3,1) → (1,0) → (-1,-1) $\sim_{\mathbf{I}_0}$ (1,1) → (3,2) as it curls around the inversion point. (Despite the curvature as depicted on the plane of the page,

[32] Using Euclidean distance, $A = \pi r^2$. However, since we are eschewing any particular metric for the time being, we simply say that the area of the circle under some metric is A.

[33] Of course, in pc-space there are many sets that remain invariant under multiple inversions. For example, (0,4,8) is invariant under $\mathbf{I}_x$, where x is even. However, this is due to the sets *transpositional* symmetry. If a set that possesses **I**-symmetry does not also possess **T**-symmetry, it will *not* remain invariant under more than one inversion.

this is a straight-line path on the surface of the cone.) Any straight path that does intersect this point reflects back in the opposite direction in the same manner as any path in $\mathbf{R}^1/\mathbf{I}_c$ that intersects $c/2$. For example, observe the path (2,4) → (1,2) → (0,0) → (-1,-2) $\sim_{\mathbf{I}_0}$ (1,2) → (2,4) and so forth, which retraces its steps after reflection off of the mirror point. Observing the behavior of this mirror point in combination with other equivalences is essential in understanding these more abstract spaces and explains much about why the **I**-spaces are generally less intuitive and more complicated than those spaces without **I**-equivalence.

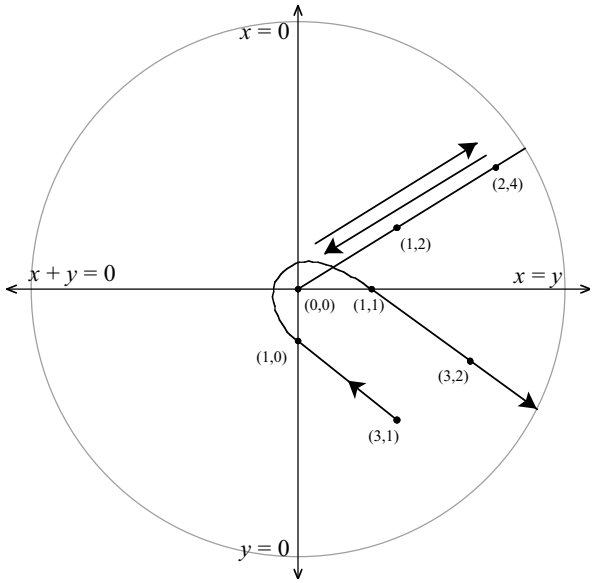


Figure 6.1.3

Turning to trichords, any point (x_1, x_2, x_3) in $\mathbf{R}^3$ is $\mathbf{I}_c$-equivalent to $(\mathbf{I}_c(x_1), \mathbf{I}_c(x_2), \mathbf{I}_c(x_3))$ or

$$\mathbf{I}_c(x_1, x_2, x_3) = (c - x_1, c - x_2, c - x_3).$$

We can interpret this as reflections in three different mirrors: $x_1 = c/2$, $x_2 = c/2$, and $x_3 = c/2$, which send the voices x_1, x_2, and x_3 to $c - x_1$, $c - x_2$, and $c - x_3$, respectively. Similar to the case for $\mathbf{R}^2$, the combination of reflections in any these three planes in $\mathbf{R}^3$ is equivalent to reflection in the point at which all three planes intersect. Thus, $\mathbf{I}_c$ sends any point in $\mathbf{R}^3$ to its reflection in the point $(c/2, c/2, c/2)$, the inversion point for trichords.

Any half space of $\mathbf{R}^3$ bounded by a plane passing through the inversion point can serve as a fundamental domain of $\mathbf{R}^3$ under $\mathbf{I}_c$-equivalence. Identifying points on the bounding plane with their reflection in the inversion point yields $\mathbf{R}^3/\mathbf{I}_c$. It is important to remember that once we identify points on this plane, the boundary it represented vanishes. There are no boundaries in $\mathbf{R}^3/\mathbf{I}_c$; there is only a single, exceptional point at the fixed point of the inversion.

Generalizing the above for n-note chords, $\mathbf{I}_c$ sends any point $\mathbf{x}=(x_1,\ldots,x_n)$ to $\mathbf{I}_c(\mathbf{x})=(\mathbf{I}_c(x_1),\ldots,\mathbf{I}_c(x_n))$ or

$$\mathbf{I}_c(\mathbf{x})=\mathbf{c}-\mathbf{x}, \text{ where } \mathbf{c} \text{ is } (c,\ldots,c).$$

This action corresponds to reflections in the hyperplanes $x_i=c/2,\ 1\le i\le n$, the combination of which is equivalent to reflection in the single point of their intersection, the inversion point $(c/2,\ldots,c/2)$.[34] We have just seen that $\mathbf{R}^3/\mathbf{I}_c$ consists of an infinite number of $\mathbf{R}^2/\mathbf{I}_c$, which in turn consists of an infinite number of $\mathbf{R}^1/\mathbf{I}_c$, a line folded in half at a single point. The quotient space $\mathbf{R}^n/\mathbf{I}_c$ follows in a similarly recursive fashion: $\mathbf{R}^n/\mathbf{I}_c$ is a half space of $\mathbf{R}^n$ bounded by a hyperplane containing the inversion point and identified in the manner of $\mathbf{R}^{n-1}/\mathbf{I}_c$, which in turn may be described as a half space of $\mathbf{R}^{n-1}$ with its boundary identified in the manner of $\mathbf{R}^{n-2}/\mathbf{I}_c$, and so forth. In every instance, there is a sole singularity at the inversion point which contains half as much space about it as every other point. The oddities of $\mathbf{R}^n/\mathbf{I}_c$ arise largely from this deficit of space about the inversion point.

6.2. We next consider the combination of inversional and permutational equivalence by observing the action of $\mathbf{I}_c$ on $\mathbf{R}^n/\mathbf{P}$. This action gives rise to reflection in a particular kind of flat, which, as we shall see, is an important part of the geometry of $\mathbf{R}^n/\mathbf{OPTI}$, since it coincides with inversionally-invariant set classes. It is therefore useful to study these flats in the simplest space in which they arise.

Recall that the space of unordered dyads, $\mathbf{R}^2/\mathbf{P}$, is a half plane bounded by a mirror at $x_1=x_2$, the unison diagonal. The action of $\mathbf{I}_c$ on $\mathbf{R}^2/\mathbf{P}$ glues together points according to the identification

$$(x_1,x_2)_{\mathbf{P}} \sim_{\mathbf{I}_c} (c-x_1,c-x_2)_{\mathbf{P}}.$$

That is, any member of the first **P**-class is $\mathbf{I}_c$-equivalent to some member of the second **P**-class and vice versa. The geometric interpretation of this action is shown in Figure 6.2.1: (a) As we have seen, $\mathbf{I}_c$ acts on any *ordered* dyad by reflection in the inversion point, $(c/2,c/2)$, taking the point (x_1,x_2) to $(c-x_1,c-x_2)$. (b) However, in $\mathbf{R}^2/\mathbf{P}$ this path reflects off of the mirror at $x_1=x_2$, which has the effect of permuting the voices of the dyad. This reflection yields the "V"-shaped path leading to $(c-x_2,c-x_1)$, the sole permutation of $(c-x_1,c-x_2)$. (c) This reflected path is equivalent to the direct (straight) path between these two points, which corresponds to reflection in the line $x_1=c-x_2$ or $x_2=c-x_1$ (we simply substitute $c-x_2$ for x_1 and $c-x_1$ for x_2), which we will call the *inversion flat*. (For dyads, the inversion flat is also the sum-c hyperplane, $x_1+x_2=c$.) Since any point in $\mathbf{R}^2/\mathbf{P}$ is $\mathbf{I}_c$-equivalent to its reflection in this line, any point *on* this line must remain fixed under this reflection, and the inversion flat consists of all and only those points in $\mathbf{R}^2/\mathbf{P}$ that are invariant under $\mathbf{I}_c$-equivalence. These are precisely those

[34] Reflection in a point that does *not* lie on the unison diagonal corresponds to a split transformation, as discussed by O'Donnell. For example, reflecting (1,2,4,7,9,10) about the point (2,2,2,9,9,9) is equivalent to inverting the first three voices about 2 and the last three voices about 9 yielding (0,2,3,8,9,11).

dyad classes of the form $(x, c-x)_{\mathbf{P}}$, since $\mathbf{I}_c(x) = c - x$ and $\mathbf{I}_c(c-x) = x$. Thus, the inversion flat divides $\mathbf{R}^2/\mathbf{P}$ into two halves. Discarding one of the halves yields the quotient space $\mathbf{R}^2/\mathbf{PI}_c$, a half space of $\mathbf{R}^2/\mathbf{P}$ or the quarter plane bounded by mirrors at $x_1 = x_2$ and $x_1 = c - x_2$.

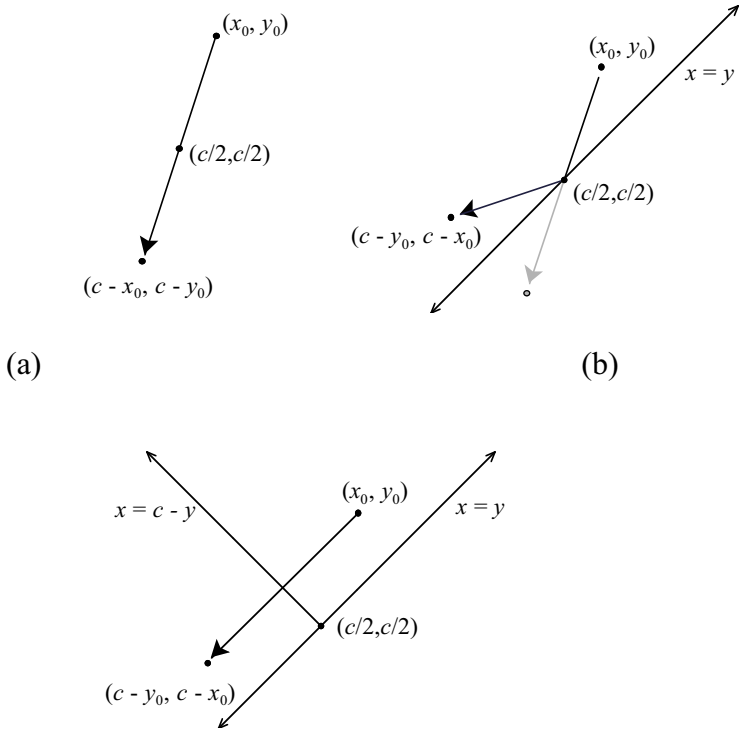


Figure 6.2.1

That $\mathbf{I}_c$ acts on $\mathbf{R}^2/\mathbf{P}$ by reflection in the line of $\mathbf{I}_c$-invariant **P**-classes, or the inversion flat, points to a general principle of the action of $\mathbf{I}_c$ on *any* of the spaces from sections four and five: $\mathbf{I}_c$ corresponds to a reflection in the flat of chords or equivalence classes of chords that remain invariant under $\mathbf{I}_c$. Therefore, the operative question when considering $\mathbf{I}_c$-equivalence on any space is "Where are the $\mathbf{I}_c$-invariant points?" Once we know the answer to this question, we know the precise action of $\mathbf{I}_c$ on the original space. We have already seen this principle at work in the quotient spaces $\mathbf{R}^n/\mathbf{I}_c$: there is precisely one fixed point, *A*, such that $\mathbf{I}_c(A) = A$, and $\mathbf{R}^n$ is transformed into $\mathbf{R}^n/\mathbf{I}_c$ by gluing together points with their reflection in *A*.

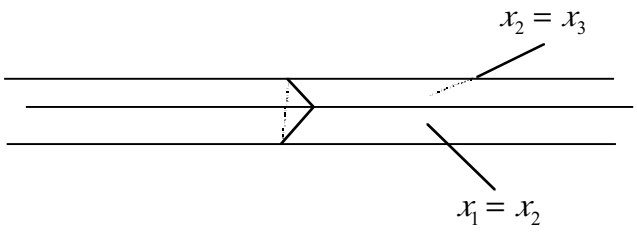
We now proceed to find the inversion flat consisting of points in $\mathbf{R}^3/\mathbf{P}$, trichords under **P**-equivalence, that remain invariant by the function

$$\mathbf{I}_c(x_1, x_2, x_3)_{\mathbf{P}} = (c - x_1, c - x_2, c - x_3)_{\mathbf{P}}.$$

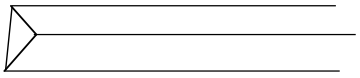
Suppose an unordered trichord contains the pitch x. In order for this trichord to remain invariant under $\mathbf{I}_c$, it must also contain $\mathbf{I}_c(x)$, or $c-x$. So any unordered trichord on the inversion flat must be of the form $\{x, c-x, ?\}$. Since we have paired off two $\mathbf{I}_c$-equivalent members of the trichord, the remaining pitch, "?", must be its own $\mathbf{I}_c$ inversion, and we have already seen that $c/2$ is the pitch that satisfies this requirement. Thus the inversion flat consists of all points in $\mathbf{R}^3/\mathbf{P}$ of the form $(x, c-x, c/2)_{\mathbf{P}}$. In this case, the "flat" of inversion is simply a line, since these points can be specified by a single variable.

To obtain the quotient space $\mathbf{R}^3/\mathbf{PI}_c$, take a half space of $\mathbf{R}^3/\mathbf{P}$ (an infinite wedge) bounded by any plane containing the inversion flat and identify points on this boundary with their reflection in the flat. The result is a half space of $\mathbf{R}^3/\mathbf{P}$ with its boundary folded along the inversion flat, like a book with $(x, c-x, c/2)_{\mathbf{P}}$ at its spine.

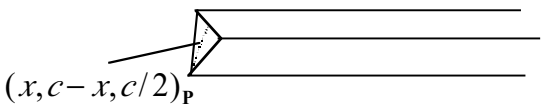
Since this is the first time we've encountered a reflection in something other than a hyperplane or a point, it is worthwhile to take a closer look at the transformation from the wedge of $\mathbf{R}^3/\mathbf{P}$ to the "half" wedge of $\mathbf{R}^3/\mathbf{PI}_c$, shown in Figure 6.2.2. We begin with the infinite wedge of $\mathbf{R}^3/\mathbf{P}$ bounded by mirrors at $x_1 = x_2$ and $x_2 = x_3$. (a) The cross-section indicated in the middle is where the boundary for the half-space will be placed. If this boundary were a mirror, the resulting space would be the half wedge in (b). Instead, this boundary this boundary will be folded about the inversion flat, shown by a dashed line in (c). As this boundary is folded (d), the permutation mirrors gradually get closer until they meet when the identification is complete (e).



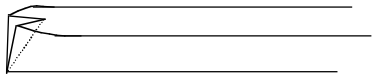
(a)



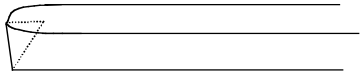
(b)



(c)



(d)



(e)

Figure 6.2.2

In constructing our "folded" half wedge, we have introduced a fair amount of distortion in the graphical representation. In actuality, the mirrors at $x_1 = x_2$ and $x_2 = x_3$ are not bent together in three dimensions; we could represent this identification without distortion in a fourth dimension. This is another example of the difference between the extrinsic and intrinsic features of a space, first encountered in §3.1. However, our representation does clearly show that the wedge $\mathbf{R}^3/\mathbf{P}$ is transformed into a "folded" half wedge $\mathbf{R}^3/\mathbf{P}\mathbf{I}_c$ without adding an additional boundary. In fact, the transformation reduces the number of boundaries, since the two mirrors are fused into a single boundary for the entire space.[35]

Generalizing the above for arbitrary cardinalities, $\mathbf{I}_c$ sends the **P**-class of **x** to the **P**-class of $\mathbf{c}-\mathbf{x}$ by reflection in the inversion flat, which consists of those **P**-classes that remain invariant under $\mathbf{I}_c$. Unordered sets satisfy this condition if for every member, y_i, $\mathbf{I}_c(y_i)$ is also a member of the set. For even cardinalities, $\mathbf{I}_c$-invariant unordered sets can thus be partitioned into $n/2$ $\mathbf{I}_c$-equivalent pairs, y_1 and $c-y_1$, y_2 and $y-x_2$, and so forth, so that the inversion flat consists of those points of the form:

$$(y_1, c-y_1, y_2, c-y_2, \ldots, y_{n/2}, c-y_{n/2})_{\mathbf{P}}. \tag{1}$$

For odd cardinalities, $\mathbf{I}_c$-invariant unordered sets can thus be partitioned into $(n-1)/2$ $\mathbf{I}_c$-equivalent pairs plus a single $\mathbf{I}_c$-invariant pitch, $c/2$, so that the inversion flat consists of those points of the form:

$$(y_1, c-y_1, y_2, c-y_2, \ldots, y_{n/2}, c-y_{n/2}, c/2)_{\mathbf{P}}. \tag{2}$$

Since for either even or odd cardinalities $\mathbf{I}_c$-invariant unordered sets can be specified with $\lfloor n/2 \rfloor$ variables, these sets comprise an $\lfloor n/2 \rfloor$-dimensional flat.

Any half space of $\mathbf{R}^n/\mathbf{P}$, a generalized wedge, bounded by a hyperplane containing the inversion flat can serve as a fundamental domain of the action of $\mathbf{I}_c$ on $\mathbf{R}^n/\mathbf{P}$.

[35] The reflection of (x_1, x_2, x_3) in the inversion flat is $(c-x_3, c-x_2, c-x_1)$. If the first two voices are equal, (x, x, y), then this point lies on the mirror $x_1 = x_2$, and its reflection in the inversion flat, $(c-y, c-x, c-x)$, lies on the mirror $x_2 = x_3$. Since $\mathbf{I}_c$ acts on any point of one mirror by sending it to a point on the other mirror, these two mirror boundaries fuse.

Identifying points on this bounding hyperplane with their reflection in the inversion flat yields the desired quotient space, $\mathbf{R}^n/\mathbf{PI}_c$. The best intuitive description of this space is a generalized half wedge where, for $n>2$, the bounding hyperplanes of $\mathbf{R}^n/\mathbf{P}$ are "twisted" into coincidence about the inversion flat.[36]

6.3. Having studied $\mathbf{R}^n/\mathbf{I}_c$ and $\mathbf{R}^n/\mathbf{PI}_c$ in detail, we already know most of what is necessary to understand how these spaces combine with **T** to yield $\mathbf{R}^n/\mathbf{TI}$ and $\mathbf{R}^n/\mathbf{PTI}$. As we have seen in section §5.0, $\mathbf{R}^n/\mathbf{T}$ is topologically (but not geometrically) equivalent to $\mathbf{R}^{n-1}$. Thus the action of $\mathbf{I}_c$ on $\mathbf{R}^n/\mathbf{T}$ is essentially the same as the action of $\mathbf{I}_c$ on $\mathbf{R}^{n-1}$, which we have already investigated: $\mathbf{R}^{n-1}/\mathbf{I}_c$ results from identifying points of $\mathbf{R}^{n-1}$ with their reflections in the inversion point. However, recall that the inversion point always lies on the unison diagonal in $\mathbf{R}^n$, which is collapsed to a single point, the origin, upon identifying **T**-equivalent points. Thus, for any value of *c*, identifying $\mathbf{I}_c$-equivalent points in $\mathbf{R}^n/\mathbf{T}$ corresponds to reflection in the origin of $\mathbf{R}^{n-1}$. (This is simply a geometric expression of the well-known fact that all inversions are related by some transposition.) Since this corresponds to the action of $\mathbf{I}_0$ on $\mathbf{R}^{n-1}$, $\mathbf{R}^n/\mathbf{TI}$ is topologically equivalent to $\mathbf{R}^{n-1}/\mathbf{I}_0$ but possesses the simplicial geometry inherited from $\mathbf{R}^n/\mathbf{T}$.

In the case of dyads, this means that $\mathbf{R}^2/\mathbf{TI}$ is equivalent to $\mathbf{R}^1/\mathbf{I}_0$—a half line bounded by a mirror point at the origin, which is also identical to the quotient space $\mathbf{R}^2/\mathbf{PT}$. For dyads (and dyads only) under transpositional equivalence, **P** and **I** are indistinguishable. That is, it is impossible to tell if the transformation from $\mathbf{R}^2/\mathbf{T}$ to the space of Figure 3.1.4(a) is achieved by identifying **P**- or **I**-equivalent points, since pairs of points that are equivalent under one are also equivalent under the other.[37] The common assertion that all dyads are invariant under inversion relies on an assumption of both **P**- and **T**-equivalence. One of the benefits of exploring all sixteen combinations of **O**, **P**, **T**, and **I** is to tease out these types of assumptions. In this case, since $\mathbf{R}^2/\mathbf{PT}$ and $\mathbf{R}^2/\mathbf{TI}$ are the same spaces, we could just as accurately assert that all **TI** set-classes of ordered pitch dyads are invariant under permutation.

The relation of the quotient spaces $\mathbf{R}^n/\mathbf{PI}_c$ and $\mathbf{R}^n/\mathbf{PTI}$ is similarly straightforward. Recall that collapsing the wedge of $\mathbf{R}^n/\mathbf{P}$ along lines of transposition yields the wedge of $\mathbf{R}^n/\mathbf{PT}$ in $\mathbf{R}^{n-1}$. **I** (or $\mathbf{I}_c$) acts on both types of wedges in exactly the same manner—by reflection in the inversion flat. In $\mathbf{R}^n/\mathbf{PT}$ these flats are identical to those defined in (1) and (2) except that they are flats of **I**-invariant **PT**-classes rather than $\mathbf{I}_c$-invariant **P**-

---

[36] The precise manner in which the bounding hyperplanes of $\mathbf{R}^n/\mathbf{P}$ are brought into coincidence is somewhat fussy and complicated. For instance, in $\mathbf{R}^4/\mathbf{P}$ the mirrors at $x_1=x_2$ and $x_3=x_4$ are made to coincide by reflection in the inversion flat, but the mirror at $x_2=x_3$ is folded onto itself.

[37] Another way to see that **P** and **I** are indistinguishable for dyadic **T**-classes is to recall that
I is an action of $\mathbf{Z}_2$ corresponding to reflection in a point, while **P** is an action of $\mathbf{S}_n$ corresponding to reflections in hyperplanes. In the impoverished geometry of $\mathbf{R}^1$, each point *is* a hyperplane.

classes.[38] Thus $\mathbf{R}^n/\mathbf{PTI}$ is a half wedge that possesses the same relation to $\mathbf{R}^n/\mathbf{PI}_c$ as $\mathbf{R}^n/\mathbf{PT}$ possesses to $\mathbf{R}^n/\mathbf{P}$.

6.4. We are now ready to consider **OPTI**-equivalence and those spaces corresponding to the familiar **TI** unordered pcset-classes, referred to in the remainder of this section simply as set classes. Proceeding as before, we can obtain the quotient space $\mathbf{R}^n/\mathbf{OPTI}$ by taking a half space of $\mathbf{R}^n/\mathbf{OPT}$ and identifying points under the action of **I**. Accurate as this may be, we have reached a sufficient degree of complexity in the spaces $\mathbf{R}^n/\mathbf{OPT}$ that this approach is not always helpful in understanding the space of set classes. For example, $\mathbf{R}^4/\mathbf{OPT}$, the join of a point and a "pinched" version of the real projective plane, is difficult enough to comprehend without trying to *transform* the space by **I**! Instead, our approach will be to zoom out, so to speak, from the quotient space itself to a larger region that is more directly accessible to our Euclidean intuitions. From this larger region we will survey the various flats corresponding to chords that are invariant under permutation, transposition, or inversion before zooming back in to see how this helps us understand $\mathbf{R}^n/\mathbf{OPTI}$.

Figure 6.4.1 is an example of the region we have in mind, which might called the normal order region of $\mathbf{R}^3/\mathbf{OT}$. Students are taught to find a normal order of a pc set by listing the pcs in "increasing order within the octave."[39] For trichords, this corresponds to a region of $\mathbf{R}^3$ in which the pitches are listed in non-descending order, $x_1 \le x_2 \le x_3$, and all pitches are within 12 semitones of one another, $x_3 \le x_1 + 12$. Projecting this region onto any sum-class plane in the manner discussed in §5.1 yields the region shown in Figure 6..4.1. Thus, any normal ordering of pcs is located somewhere in this region of $\mathbf{R}^3/\mathbf{T}$. Since all normal orderings of a given pcset are related by rotation, this region contains all of the cyclic permutations of any point in the region, but none of the other permutations. For example, Figure 6.4.1 contains $(0,4,7)_{\mathbf{OT}}$, $(4,7,0)_{\mathbf{OT}}$, and $(7,0,4)_{\mathbf{OT}}$, the three cyclic permutations of a $(0,4,7)_{\mathbf{T}}$ in "close" position, or, interpreting any one of the voices as the *bass* voice of the texture, the three chordal inversions of a major triad. For this reason we will refer to the region in Figure 6..4.1 and its higher-dimensional analogues as the *cyclic region*.

[38] For simplicity of notation, we can substitute 0 for *c* in (1) and (2) when specifying the inversion flat in $\mathbf{R}^n/\mathbf{PTI}$. To see that this is the case, transpose any **P**-classes of the form of (2) by $-c/2$. Each $\mathbf{I}_c$-equivalent pair is now $x_i - c/2$ and $c/2 - x_i$. Substituting x_i' for $x_i - c/2$, this pair becomes x_i' and $-x_i'$. The remaining $\mathbf{I}_c$-invariant pitch is now 0, which it must be, since $\mathbf{I}_c$ acts on any space incorporating **T** by reflection in the origin.

[39] Rahn, p. 33.

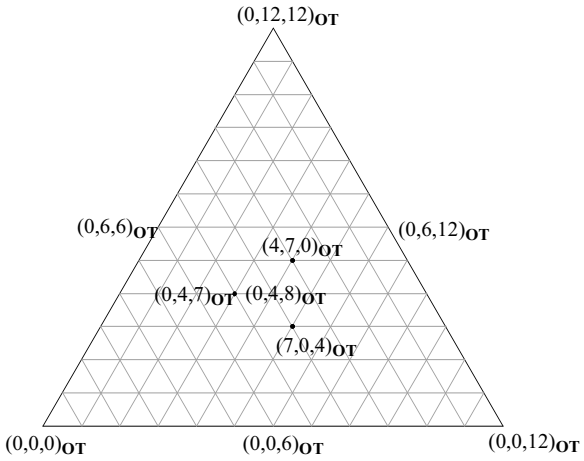


Figure 6.4.1.

Due to the manner in which we constructed the cyclic region, there are only three points that are **O**-equivalent, which are the three vertices of the triangle. Gluing these points together gives the shape of this region on the $\mathbf{R}^3/\mathbf{OT}$ torus. Indeed, beginning with the $\mathbf{R}^3/\mathbf{OT}$ torus and placing mirrors at $x_1 = x_2$, $x_2 = x_3$, and $x_3 = x_1$ would give precisely this region. Points on these mirrors contain pc duplications and invariant under some (non-cyclic) permutation. We can also see that the centroid of the cyclic region is the transposition class of maximally-even sets, which remains fixed under cyclic permutation.[40] This point is not yet a singularity, since we have not yet glued together points related by cyclic permutation. For now, $(0,4,8)_{\mathbf{OT}}$ is merely a latent singularity.

Figure 6.4.2 shows the location of **I**-invariant sets within the cyclic region, which form the line $(-x,0,x)_{\mathbf{OT}}$ and its cyclic permutations $(0,x,-x)_{\mathbf{OT}}$ and $(x,-x,0)_{\mathbf{OT}}$. These lines correspond to sets in which two of the voices are positioned symmetrically about an inversional axis that falls on the remaining voice. Notice that each line extends from the minimally even set class through the maximally even set to the member of $(0,0,6)_{\mathbf{OPTI}}$ that lies on the opposite face of the region. Notice also that these three lines carve the cyclic region into six congruent 30^o-60^o-90^o triangles. Each of the triangular regions is the quotient space $\mathbf{R}^3/\mathbf{OPTI}$: each triangle contains all possible trichordal set classes. The cyclic region must be tiled by six copies of $\mathbf{R}^3/\mathbf{OPTI}$, since the region contains all three cyclic permutations for both a chord and its inversion. Figure 0.2(b) plots the

[40] Clough and Douthett (1991) and many other sources use the term "maximally even" to describe sets like (0,2,4,7,9), which are not strictly even because of the assumed discrete division of the octave. Since we are operating in continuous pitch-class space, which can be seen as a limit case of equal temperament as the number of octave divisions, our "maximally even" sets always belong to the special class that Carey & Clampitt call "degenerate well-formed scales" and Quinn (2004) calls "Class IIa ME sets." The reader should be mindful of this distinction; for example, the centroid of any cyclic region in a quotient space of $\mathbf{R}^n$ is always a maximally even set of this special type, and thus is only a possible chord in c-tone equal temperaments where c is a multiple of n.

twelve familiar trichordal set classes (and seven multiset classes) from twelve-tone equal temperament, which form a lattice on $\mathbf{R}^3/\mathbf{OPTI}$.

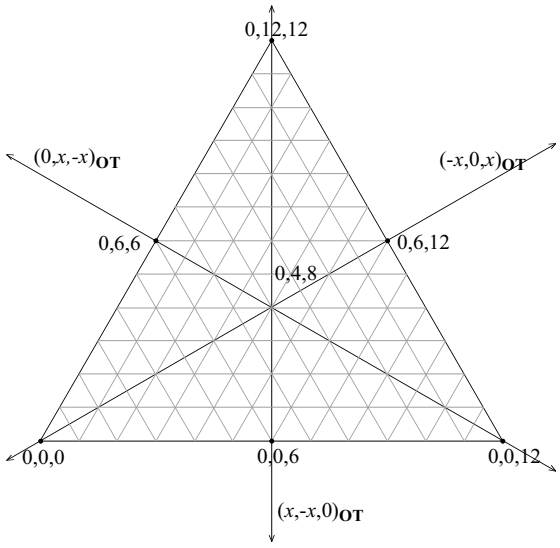


Figure 6.4.2.

The three boundaries of $\mathbf{R}^3/\mathbf{OPTI}$ are all mirrors. (If this was not the case, then one or more boundaries would contain duplications and the triangle would be a fundamental domain but not a quotient space.) The boundary extending from $(0,0,0)_{\mathbf{OPTI}}$ to $(0,0,6)_{\mathbf{OPTI}}$ is a permutation mirror, while the other two mirror boundaries correspond to inversion. Flipping the triangle about either of these two inversional boundaries yields one of the two fundamental domains of OPT for trichords discussed in §5.3. Recall that the Callender kite resulted from cutting the $\mathbf{R}^3/\mathbf{OPT}$ cone from $(0,0,6)_{\mathbf{OPT}}$ to $(0,4,8)_{\mathbf{OPT}}$ and the other fundamental domain resulted from cutting this cone from the origin to the cone point.[41] Combining these two cuts and proceeding in a straight line on the cone from the origin through the cone point to $(0,0,6)_{\mathbf{OPT}}$ on the "back" side of the cone yields the triangular quotient space $\mathbf{R}^3/\mathbf{OPTI}$. Thus, while the triangle has three *boundaries*, it only possesses two *mirrors*: a single permutation mirror along the bottom and a single mirror on the inversion flat that is "bent" at the maximally-even set class under its three-fold rotational or T_4-symmetry.

The various singularities in $\mathbf{R}^3/\mathbf{OPTI}$ are summarized in the table below. The columns labeled *p*, *t*, and *i* correspond to the degrees of permutational, transpositional,

[41] In a later section devoted to $\mathbf{R}^n/\mathbf{OPT}$ and $\mathbf{R}^n/\mathbf{OPTI}$, we will consider a generalization of the latter fundamental domain that may be described as follows: The cyclic region for *n*-note chords is the (*n*-1)-simplicial region of $\mathbf{R}^n/\mathbf{T}$ with *n* vertices at $A_i = \left(0^{n-i},12^n\right)_{\mathbf{T}}$, $0 \le i < n$. Replacing A_{n-1} with a vertex at the maximally-even set yields a fundamental domain of **OPT** in $\mathbf{R}^n$. Starting with this fundamental domain, replacing A_{n-2} with a vertex at $(0,0,6,\ldots,6)_{\mathbf{T}}$ yields a fundamental domain of **OPTI** in $\mathbf{R}^n$. This construction is due to Noam Elkies.

and inversional symmetries of a given chord. The column 1/pti indicates the fraction of space about the given object in comparison with a typical point in the space. For instance, $(0,0,3)_{\mathbf{OPTI}}$ lies on the permutation boundary and contains only half as much space as $(0,3,7)_{\mathbf{OPTI}}$. We can fit a sufficiently small disk centered on $(0,3,7)_{\mathbf{OPTI}}$ with the quotient space, but for a disk centered on $(0,0,3)_{\mathbf{OPTI}}$, half a the disk will be reflected onto the other half. Notice how this deficit of space is reflected in the geometry of $\mathbf{R}^3/\mathbf{OPT}$: $(0,0,6)_{\mathbf{OPTI}}$ lies at the $360^o/4 = 90^o$ angle, $(0,4,8)_{\mathbf{OPTI}}$ lies at the $360^o/6 = 60^o$ angle, and $(0,0,0)_{\mathbf{OPTI}}$ lies at the $360^o/12 = 30^o$ angle.

Object	p	t	i	1/*pti*
$(0,0,x)_{\mathbf{OPTI}}, 0 < x < 6$	2	1	1	1/2
$(-x,0,x)_{\mathbf{OPTI}}, 0 < x < 6, x \neq 4$	1	1	2	1/2
$(0,0,6)_{\mathbf{OPTI}}$	2	1	2	1/4
$(0,4,8)_{\mathbf{OPTI}}$	1	3	2	1/6
$(0,0,0)_{\mathbf{OPTI}}$	6	1	2	1/12

It is tempting but terribly misleading to generalize the very simple space of trichordal set classes to spaces of larger cardinalities. However, in light of the many examples of complicated quotient spaces we have already encountered, it is clear that any naïve hopes of geometric simplicity for the space of set classes must be set aside in order to understand the actually spaces these abstract chord types inhabit. Accordingly, we will proceed with a detailed look at the space of tetrachordal set classes, $\mathbf{R}^4/\mathbf{OPTI}$, followed by quicker glances at the cases for pentachords, hexachords, and heptachords before generalizing the observed principles.

Figure 6.4.3 shows the tetrachordal cyclic region, a tetrahedron with vertices at $(0,0,0,0)_{\mathbf{OT}}$, $(0,0,0,12)_{\mathbf{OT}}$, $(0,0,12,12)_{\mathbf{OT}}$, and $(0,12,12,12)_{\mathbf{OT}}$. As with the cyclic region for trichords, the boundaries are mirrors corresponding to **P**-equivalence other than cyclic permutations. There are two planes cutting through the interior of this tetrahedral region. These are the inversion flat $(-x,-y,y,x)_{\mathbf{OT}}$, and its sole unique cyclic permutation $(-y,y,x,-x)_{\mathbf{OT}}$. There are only two unique cyclic permutations of $(-x,-y,y,x)_{\mathbf{OT}}$, since the third cyclic permutation, $(-x,-y,y,x)_{\mathbf{OT}}$, in which the "outer" and "inner" voices form **I**-equivalent pairs, contains the same sets as the first. Thus, since there are only two unique cyclic permutations of the inversion flat, there are only two planes dividing the cyclic region into four subregions. Following the previous discussion of trichords, for most points in the cyclic region, there are seven other points to which a given point is **OPTI**-equivalent—three related by cyclic permutation and four related by inversion.[42] For example, there are eight points in the cyclic region that belong to $(0,2,5,8)_{\mathbf{OPTI}}$, namely, $(0,2,5,8)_{\mathbf{OT}}$, $(0,2,6,9)_{\mathbf{OT}}$, $(11,2,5,9)_{\mathbf{OT}}$, $(11,3,6,9)_{\mathbf{OT}}$, $(11,2,6,8)_{\mathbf{OT}}$, $(0,3,6,8)_{\mathbf{OT}}$, $(11,3,5,8)_{\mathbf{T}}$, and $(0,3,5,9)_{\mathbf{T}}$, the four chordal inversion of half-diminished and dominant seventh chords. This means that each of the four subregions in Figure 6.4.3 must contain points that are **OPTI**-equivalent, which in turn implies the presence of singularities within each of these subregions. Where are these singularities and how do they arise?

[42] Points that possess some degree of symmetry will have fewer **OPTI**-equivalent siblings in the cyclic region.

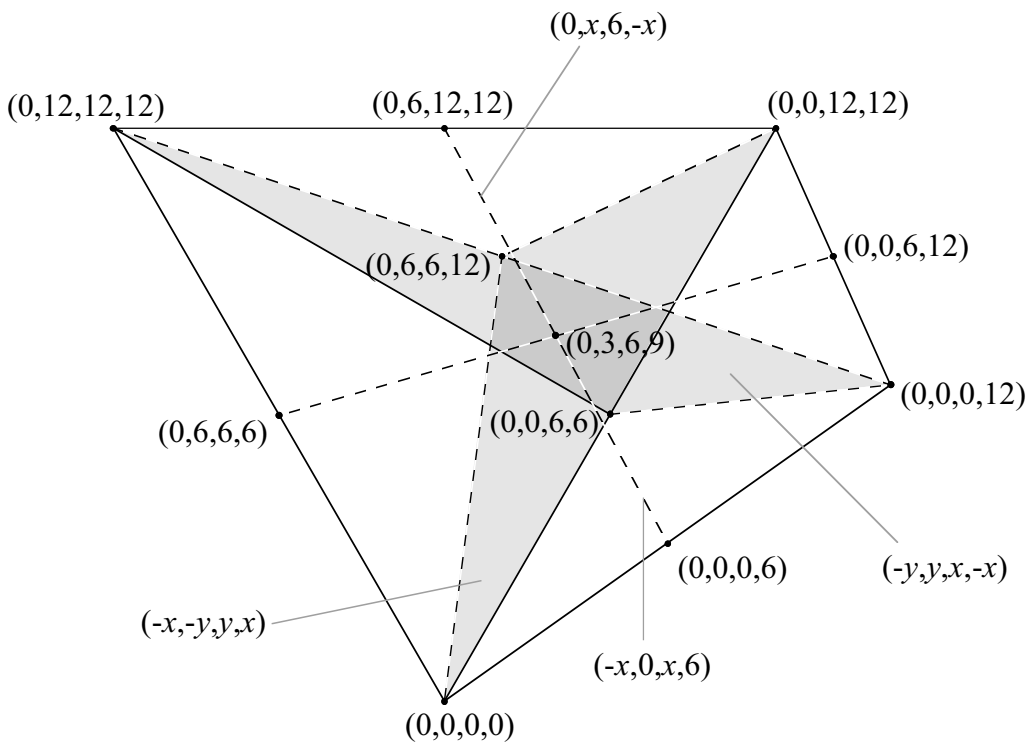


Figure 6.4.3. The cyclic region in $\mathbf{R}^4/\mathbf{OT}$ with inversion planes and lines.

Consider trichordal **OT**-classes of the form $(-x,0,x)_{\mathbf{OT}}$, which are invariant under **PI**-equivalence. We can add a fourth voice a tritone away from the second without disturbing the inversional symmetry, since both pcs remain fixed by the same inversion. (This is a direct consequence of the discussion in §3.5, in which tritone-related mirror points arose from the combination of **O** and $\mathbf{I}_c$ in the context of one-note chords.) In particular, tetrachordal **OT**-classes of the form $(-x,0,x,6)_{\mathbf{OT}}$ are also invariant under **PI**-equivalence. These **OT**-classes form two lines in the tetrachordal cyclic region of Figure 6.4.3. One line corresponds to $(-x,0,x,6)_{\mathbf{OT}}$ and contains the equal-tempered **OT**-classes $(0,0,0,6)_{\mathbf{OT}}$, $(11,0,1,6)_{\mathbf{OT}} = (0,1,2,7)_{\mathbf{OT}}$, $(10,0,2,6)_{\mathbf{OT}} = (0,2,4,8)_{\mathbf{OT}}$, and so forth. The other line is the sole unique cyclic permutation of the first, containing **OT**-classes of the form $(0,x,6,-x)_{\mathbf{OT}}$.[43] These lines correspond to sets that are **I**-invariant in *pitch-class* space but not *pitch* space. For instance, (0,1,2,7) and (0,2,4,8) are **I**-invariant under some inversion in pc-space, but this invariance cannot be *instantiated* in pitch space. This contrasts with sets lying on the planes of inversional invariance in Figure 6.4.3, such as (0,1,6,7), where this invariance *can* be realized in pitch space.

We will now construct a fundamental domain of **OPTI** for tetrachords, making sure that the boundaries of this domain contain the appropriate inversion flats, as they must. We begin by taking any of the four subregions of the cyclic region bounded by the two

[43] The other cyclic permutations of $(-x,0,x,6)_{\mathbf{OT}}$ are identical to one of these two lines for the same reason that there are only two unique inversion planes in the cyclic region.

(mirror) inversion planes. For convenience, we take the subregion with vertices at $(0,0,0,0)_{\mathbf{OT}}$ and $(0,0,0,12)_{\mathbf{OT}}$, a tetrahedron whose remaining two vertices are $(0,0,6,6)_{\mathbf{OT}}$ and $(0,6,6,12)_{\mathbf{OT}}$. Since this region contains **I**-equivalent points related by reflection in the inversion line $(-x,0,x,6)_{\mathbf{OT}}$, a fundamental domain of **OPTI** must divide this subregion into two halves bounded by a plane containing this inversion line. The obvious choice is the plane that also includes the intersection of the two inversion planes. Taking the half that contains the origin yields the fundamental domain of **OPTI** shown in Figure 6.4.4.[44]

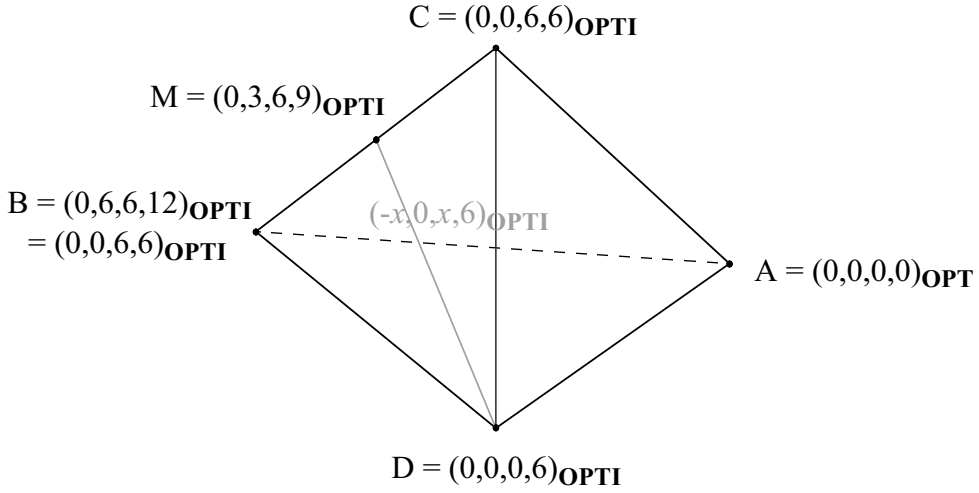


Figure 6.4.4. Fundamental domain of OPTI for tetrachords.

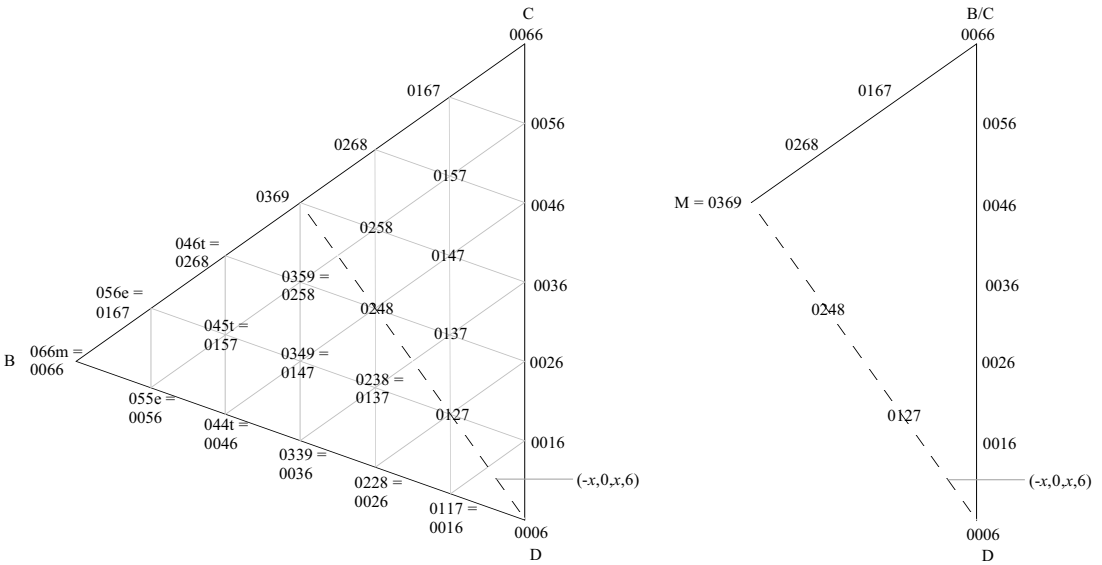
We have just reconstructed Cohn's tetrahedron, expanded to include multisets. His vertex at the minimally-even set class of the 12 pc-universe, [0123], has been replaced by the corresponding point in a continuous space, the origin. Similarly, his two vertices at [0167], a maximally-even distribution of the smallest interval in the 12 pc-universe, have been placed at the corresponding points in a continuous space, $(0,0,6,6)_{\mathbf{OT}}$ and $(0,6,6,12)_{\mathbf{OT}}$, and his vertex at [0127] is replaced by $(0,0,0,6)_{\mathbf{OT}}$. The ABD and ACD faces lie on the boundary of the cyclic region and are thus mirrors. The ABC face lies on the inversion plane and is also a mirror. However, the BCD face is not a mirror, since it contains duplicate set classes arranged symmetrically about the inversion line. These duplications are the source of much discussion in the final section of Cohn (2003), where he suggests "creasing the BCD face along its central column and folding the corresponding points together."[45] We will now carry out this identification transforming, the fundamental domain into the quotient space $\mathbf{R}^4/\mathbf{OPTI}$.[46]

[44] This is only one of many possible fundamental domains of OPTI for tetrachords.

[45] Since Cohn focuses on the 12 pc-universe, his tetrahedron is a lattice contained with the tetrahedron of Figure 6.4.4. The number of duplicate set-classes on this lattice are

The BCD face folds in half along the inversion line in the manner already discussed in connection with $\mathbf{R}^3 / \mathbf{PI}_c$ and shown in Figure 6.4.5(a-b). (In Figure 6.4.5, t = 10, e = 11, and m = 12.) The vertices B and C are placed at a single point, B/C = $(0,0,6,6)_{\mathbf{OPTI}}$, but the fold creates a new vertex at the maximally-even set class, $(0,3,6,9)_{\mathbf{OPTI}}$, labeled M. Points on the interior of this face and the boundary that they form vanish into the interior of the space. Since the remaining faces of the tetrahedron share an edge with this now-collapsed face, all are affected by this transformation. Folding the BC edge of the ABC face along its midpoint (c) yields a cone (d) with a cone point at M. On this cone there are separate straight-line paths from $(0,0,0,0)_{\mathbf{OPTI}}$ to $(0,0,6,6)_{\mathbf{OPTI}}$; one in which two voices move in parallel, e.g., $(0,0,0,0)_{\mathbf{OPTI}} \rightarrow (0,0,1,1)_{\mathbf{OPTI}} \rightarrow \ldots (0,0,6,6)_{\mathbf{OPTI}}$, and another in which two voices move in contrary motion, e.g., $(0,0,0,0)_{\mathbf{OPTI}} \rightarrow (11,0,0,1)_{\mathbf{OPTI}} \rightarrow \ldots (6,0,0,6)_{\mathbf{OPTI}} = (0,0,6,6)_{\mathbf{OPTI}}$. (Readers can verify from Figures 6.4.5(a-d) that the A(B/C)M cone and the DM line contain all **I**-invariant tetrachordal set classes.)

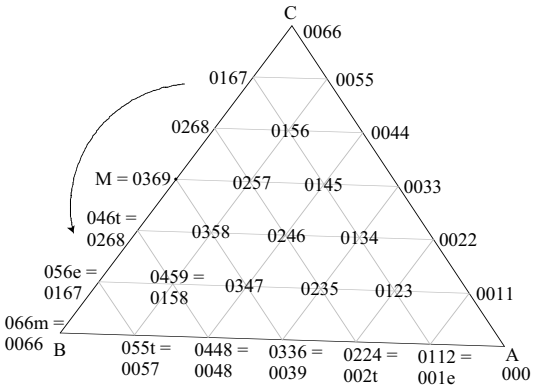
The remaining two faces, ABD and ACD, are united along their BC and CD edges (e), yielding a second cone (f) with the same two distinct straight paths from A to B/C and a cone point at D. Putting this all together, the cones in (d) and (f) are united along their shared A to B/C paths, with the inversion line running through the interior attached to the two cone points (g). (As before, it would be necessary to represent this three-dimensional space in four-dimensional Euclidean space to "see" it without distortions.)



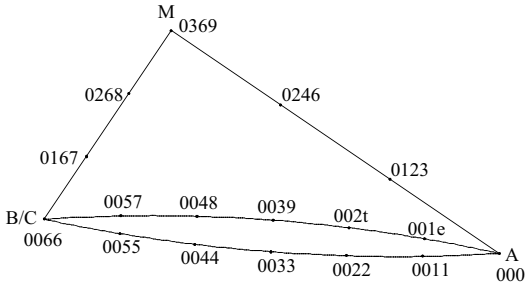
(a) (b)

quite limited; in the corresponding continuous space, there are any infinite number of duplications.

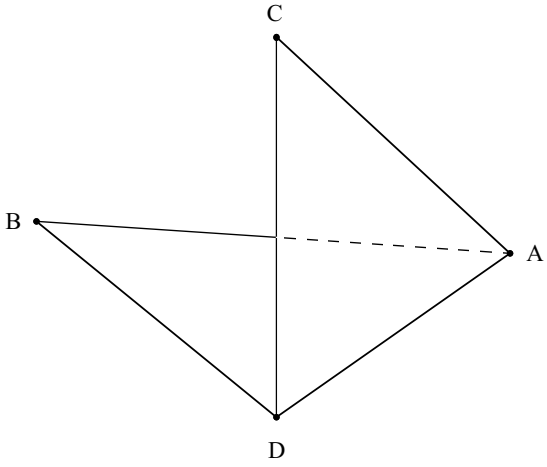
[46] Indeed, Cohn's intuition is writ large in the present paper.


C
0066
0167
0055
0268
0156
0044
M = 0369
0257
0145
0033
046t =
0268
0358
0246
0134
0022
056e =
0167
0459 =
0158
0347
0235
0123
0011
066m =
0066
B
055t =
0057
0448 =
0048
0336 =
0039
0224 =
002t
0112 =
001e
A
000

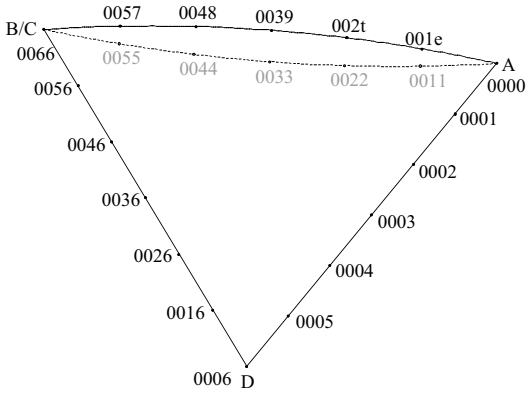
(c)


M
0369
0268
0246
0167
0123
0057
0048
0039
002t
001e
B/C
0066
0055
0044
0033
0022
0011
A
000

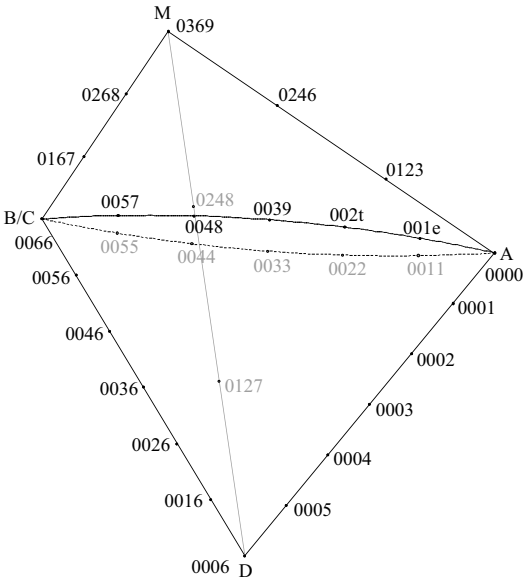
(d)


C
B
A
D

(e)


0057
0048
0039
002t
001e
B/C
0066
0055
0044
0033
0022
0011
A
0000
0056
0001
0046
0002
0036
0003
0026
0004
0016
0005
0006
D

(f)



(g)

Figure 6.4.5. Transformating the fundamental to the quotient space R4/OPTI.

Figure 6.4.5(g) is the quotient space $\mathbf{R}^4/\mathbf{OPTI}$, which is bounded by two mirrors wrapped into cones, with a single permutation mirror on the bottom, a single inversion mirror on the top, and a line mirror connecting the two cone points. Most of the points on these two boundaries are singularities whose neighborhood possess half the usual space, but there are numerous singularities that possess a greater deficiency of space about them. The table below summarizes the location and degree of spatial contraction about these singularities:

Object	Sets of the form	*p*	*t*	*i*	1/*pti* of ball
AD	$(0,0,0,x)_{\mathbf{OPTI}}, 0<x<6$	6	1	1	1/6
AB/C	$(0,0,x,x)_{\mathbf{OPTI}}, 0<x<6$	4	1	2	1/8
AB/C	$(-x,0,0,x)_{\mathbf{OPTI}}, 0<x<6$	2	1	2	1/4
B/CM	$(-x,0,x,6)_{\mathbf{OPTI}}, 0<x<3$	1	2	2	1/4
M	$(0,3,6,9)_{\mathbf{OPTI}}$	0	4	2	1/8
D	$(0,0,0,6)_{\mathbf{OPTI}}$	6	1	2	1/12
B/C	$(0,0,6,6)_{\mathbf{OPTI}}$	4	2	2	1/16
A	$(0,0,0,0)_{\mathbf{OPTI}}$	24	1	2	1/48

Now we turn to pentachords. For the remainder of this section we will focus on the inversion flat(s) and their relation to the permutation boundary. Since the space about

any points in $\mathbf{R}^n/\mathbf{OPTI}$ that are not singularities is identical to $\mathbf{R}^n$, focusing on these singularities allows us to study the more interesting areas in the quotient space. Additionally, since the inversion flat for up to seven-note chords is only three dimensions, these flats allow us to apply our spatial intuitions to cardinalities as large of heptachords.

Figure 6.4.6(a) shows the cyclic region for pentachords projected onto the plane as a pentagon. (As before, m = 12.) The inversion flat $(-x,-y,0,y,x)_{\mathbf{T}}$ includes the origin and passes through the center of the pentagon.[47] This flat intersects the edges connecting vertices arranged symmetrically about the origin at their midpoint. That is, $(-x,-y,0,y,x)_{\mathbf{T}}$ is a two-dimensional (planar) flat that intersects the cyclic region at three points: the origin, the midpoint of BE, $(0,6,6,6,12)_{\mathbf{T}}$, and the midpoint of CD, $(0,0,6,12,12)_{\mathbf{T}}$. The inversion flat in $\mathbf{R}^5/\mathbf{OPTI}$ is thus a triangle with vertices at the origin, the set class of quintuple unisons, and the two unique set classes that correspond to a tritone with pc duplications, $(0,0,0,0,6,)_{\mathbf{OPTI}}$ and $(0,0,6,6,6)_{\mathbf{OPTI}}$.[48] (See Figure 6.4.6(b), where the reader can verify that this triangle contains all pentachordal set classes that are **I**-invariant.) Singularities in $\mathbf{R}^5/\mathbf{OPTI}$ are found in two places: on the single four-dimensional permutation boundary that contains all set classes with pc duplications, and the inversion flat, which attaches to the permutation boundary along the edges of the triangle. We can imagine placing this triangle directly against the permutation boundary, then grabbing the maximally even set-class, $(0,{}^{12}\!/_{5},2\cdot{}^{12}\!/_{5},3\cdot{}^{12}\!/_{5},4\cdot{}^{12}\!/_{5})_{\mathbf{OPTI}}$ or $(0,2.6,5.2,7.8,9.4)_{\mathbf{OPTI}}$, and pulling it along with the rest of the interior of the triangle off of the permutation boundary. The inversion flat is thus "bent" around the maximally-even set class under the influence of its five-fold rotational or $T_{12/5=2.6}$-symmetry, just as was the case for the inversion line $\mathbf{R}^3/\mathbf{OPTI}$. There are two classes of singularities in $\mathbf{R}^5/\mathbf{OPTI}$ that stand out from the rest: the maximally-even set class, which possesses both **T**- and **I**-symmetry, and the sets on the edges of the triangular inversion flat, which possess both **P**- and **I**-symmetry.

[47] There are a total of five cyclic permutations of this flat that pass through the cyclic region, each containing a different vertex of the region. We choose this particular flat for the obvious reason that it contains the origin.

[48] Note that these two set classes are *not* equivalent under the four equivalences we are considering. See §2 for a discussion as to why we are not consider "cardinality" equivalence.

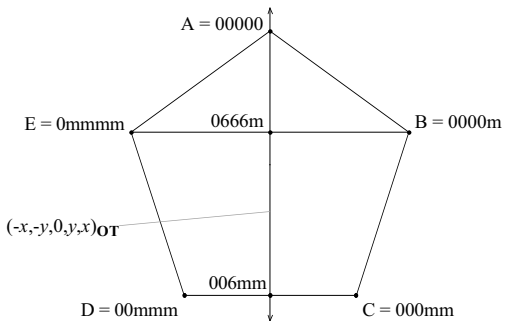


Figure 6.4.6 (a)

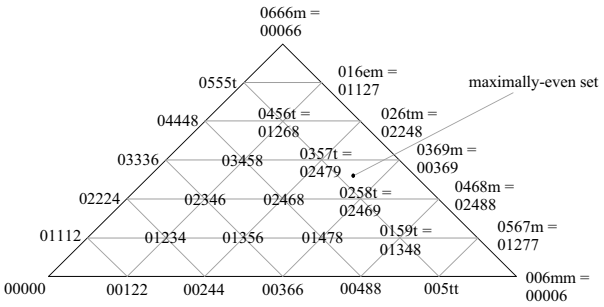


Figure 6.4.6 (b)

It is important to remember that the inversion flat is *not* a boundary of $\mathbf{R}^5/\mathbf{OPTI}$. Though it is a flat that attaches to the permutation boundary, it no more bounds the four-dimensional space of $\mathbf{R}^5/\mathbf{OPTI}$ than an arc attached to a plane bounds a three-dimensional object. The flat belongs to the category of "radial singularities," discussed in §5.3.

Jumping ahead to the next odd cardinality, Figure 6.4.7 shows the projection of the cyclic region for seven-note chords onto the plane as a heptagon. As before, the inversion flat that passes through the origin, $(-x,-y,-z,0,z,y,x)_{\mathbf{OT}}$, intersects the cyclic region at the midpoints of edges connecting vertices arranged symmetrically about the origin. The inversion flat in $\mathbf{R}^7/\mathbf{OPTI}$ is thus a tetrahedron with vertices at the origin and those three unique set-classes that are equivalent to a tritone with pc duplications: $(0,0,0,0,0,0,7)_{\mathbf{OPTI}}$, $(0,0,0,0,0,6,6)_{\mathbf{OPTI}}$, and $(0,0,0,0,6,6,6)_{\mathbf{OPTI}}$. This flat attaches to the single permutation boundary at the faces of the tetrahedron, which all contain pc duplications. The tetrahedron will be "bent" at the singularity of the maximally-even set class $(0,{}^{12}\!/_7,2\cdot{}^{12}\!/_7,3\cdot{}^{12}\!/_7,4\cdot{}^{12}\!/_7,5\cdot{}^{12}\!/_7,6\cdot{}^{12}\!/_7)_{\mathbf{OPTI}}$.

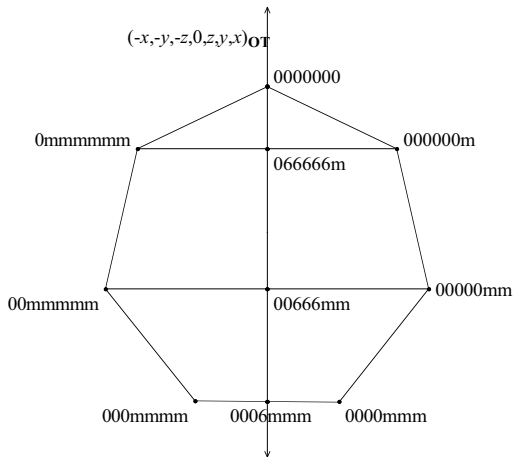


Figure 6.4.7. Cyclic region and inversion flat in R7/T.

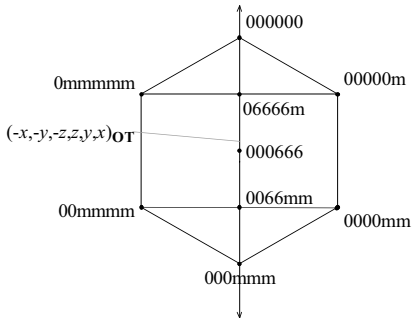
Generalizing for $\mathbf{R}^n/\mathbf{OPTI}$, when n is odd, the inversion flat is an $\lfloor n/2 \rfloor$-simplex with vertices at the origin and the $\lfloor n/2 \rfloor$ unique set-classes that are equivalent to a tritone with pc duplications. This simplex is "bent" under the influence of the n-fold rotational or $\mathcal{I}_{12/n}$-symmetry of the maximally-even set class, and attaches to the permutation boundary at its facets.[49] For prime cardinalities there will be four types of singularities: 1) those that lie solely on the permutation boundary; 2) those that lie solely on the inversion flat; 3) those that lie at the intersection of the permutation boundary and the inversion flat; and 4) the maximally-even set. For odd but not prime cardinalities there will be additional flats corresponding to set classes with **T**-symmetry, as described in §5.3.

The situation for even cardinalities is much more complicated, as we have seen with $\mathbf{R}^4/\mathbf{OPTI}$. We now turn to a brief look at the case for hexachords.

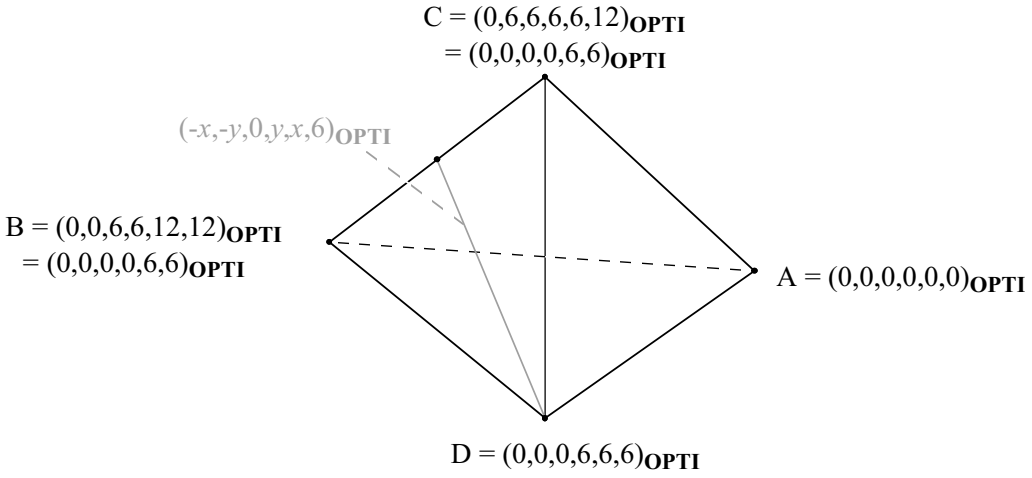
As before, we begin with the cyclic region, shown as a planar projection in Figure 6.4.8(a). The inversion flat $(-x,-y,-z,z,y,x)_{\mathbf{OT}}$ passes through two vertices of the hexagon, the origin and the opposite vertex, $(0,0,0,12,12,12)_{\mathbf{T}}$. This flat also intersects the cyclic region at the midpoints of two edges, $(0,6,6,6,6,12)_{\mathbf{T}}$ and $(0,0,6,6,12,12)_{\mathbf{T}}$. However, notice that the edge connecting A and D itself inversionally symmetric about its midpoint $(0,0,0,6,6,6)_{\mathbf{T}}$; e.g., $(0,0,0,7,7,7) \sim_{\mathbf{OPTI}} (0,0,0,5,5,5)$. In fact, all three cyclic permutations of this flat intersect at the triangular face with vertices at $(0,0,6,6,12,12)_{\mathbf{T}}$, $(0,6,6,6,6,12)_{\mathbf{T}}$, and $(0,0,0,6,6,6)_{\mathbf{T}}$. Every point in the inversion flat on one side of this triangle is I-equivalent to a point on the other side. Taking this face with the origin yields the tetrahedron shown in Figure 6.4.8(b). Since this tetrahedron contains duplications (e.g., two of its vertices belong to the same set class), it is a fundamental domain of the inversion flat. In fact, the entire BCD face is divided into **I**-equivalent halves by the line segment extending from D to the midpoint of BC as shown in Figure 6.4.8(b). This is exactly the same case as the fundamental domain we constructed for **OPTI** in $\mathbf{R}^4$, a

[49] Facets are to higher-dimensional objects like simplexes as edges are to polygons and faces are to polyhedra.

tetrahedron with a single face to be folded in half. Thus, the *inversion flat* in $\mathbf{R}^6/\mathbf{OPTI}$ will be the same basic shape as the *quotient space* $\mathbf{R}^4/\mathbf{OPTI}$, a rather unexpected result at first glance.



(a)



(b)

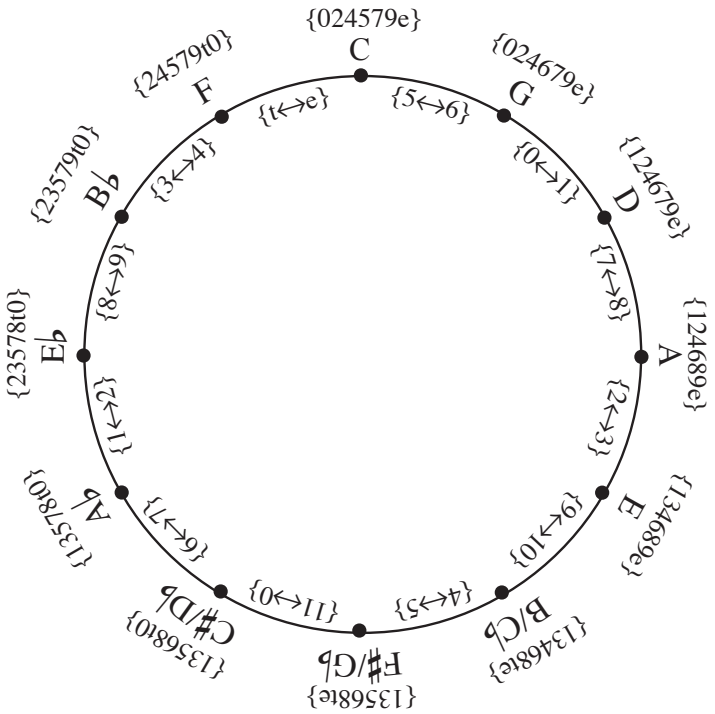
Figure 6.4.8. (a) cyclic region in $\mathbf{R}^6/\mathbf{T}$; (b) fundamental domain of "pitch-and-pc" inversion flat for hexachords.

What is this line along which the BCD face is folded and where does it come from? Recall that sets of even cardinalities can be inversionally invariant in pc-space without being I-invariant in pitch space. Hexachords that display pc-only I-invariance are of the form $(-x,-y,0,y,x,6)_{\mathbf{OT}}$. This two-dimensional flat, which we will call the pc-inversion flat, intersects the tetrahedron of Figure 6.4.8(b) at exactly the line along which the BCD face is folded. In other words, this line is the intersection of the pitch-and-pc-inversion flat and the pc-only-inversion flat.

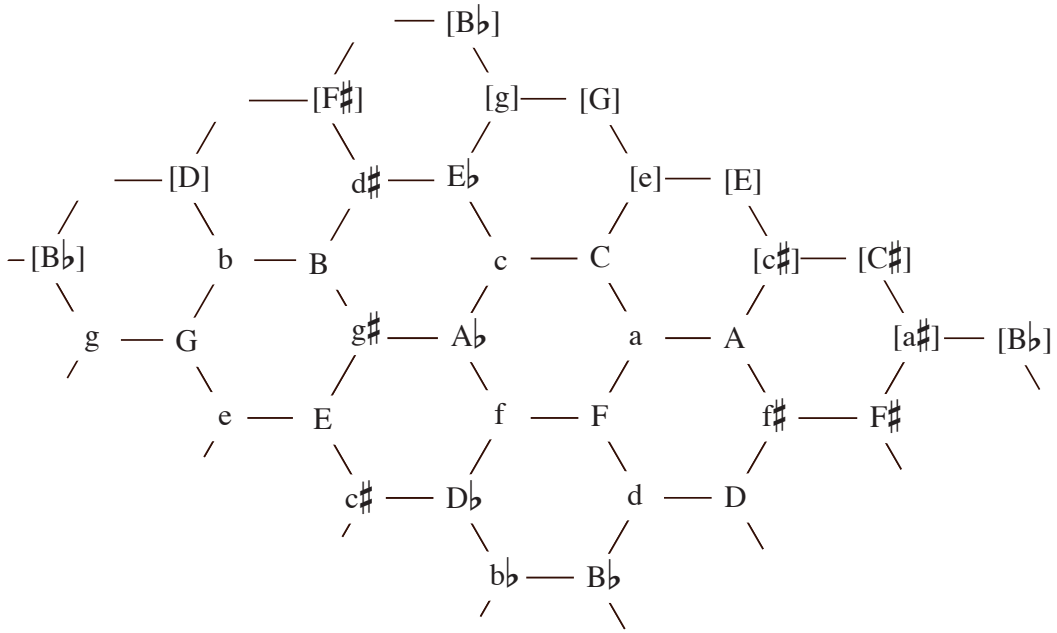
Though we have not yet discussed the exact manner in which the inversion flats attach to the permutation boundary, the manner in which the singularity at the maximally-even set arises, or the other multitude of singularities in this space, our discussion does give a general sense of the complexities that arise for even cardinalities. A more thorough examination is presently deferred for a later section specifically devoted to $\mathbf{R}^n/\mathbf{OPT}$ and $\mathbf{R}^n/\mathbf{OPTI}$.

It is worth taking a moment at the end of our tour the most of the sixteen spaces to summarize our trip from $\mathbf{R}^n$ and the space of ordered pitch sets to the space of *n*-note set-classes. **O**-equivalent points in $\mathbf{R}^n$ are identified by a particular action of $12\mathbf{Z}^n$, yielding an *n*-torus. **P**-equivalent points in $\mathbf{R}^n/\mathbf{O}$ are identified by a particular action of $\mathbf{S}_n$, yielding a (*n*-1)-simplex-based prism with its bases identified by a vertex permuting twist. **T**-equivalent points in $\mathbf{R}^n/\mathbf{OP}$ are identified by a particular action of **R**, a projection onto any sum-class hyperplane; this action on $\mathbf{R}^3/\mathbf{OP}$ yields a cone with a cone point at the **OPT**-class of maximally even sets and a mirror boundary at **P**-invariant **OPT**-classes; on $\mathbf{R}^4/\mathbf{OP}$ this action yields the join of the maximally-even set and the projective plane. **I**-equivalent points in $\mathbf{R}^n/\mathbf{OPT}$ are identified by a particular action of $\mathbf{Z}_2$, taking $\mathbf{x}_{\mathbf{OPT}}$ to $-\mathbf{x}_{\mathbf{OPT}}$, which yields the triangular space $\mathbf{R}^3/\mathbf{OPTI}$ and the "double cone" of $\mathbf{R}^4/\mathbf{OPTI}$.[50] **OPTI**-equivalent points in $\mathbf{R}^n$ are thus identified by a particular action of the group $12\mathbf{Z}^n \times \mathbf{S}_n \times \mathbf{R} \times \mathbf{Z}_2$.

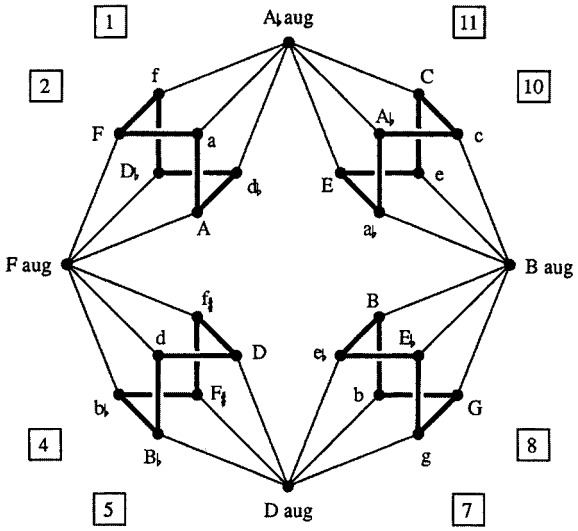
[50] There is nothing unique about this particular trip from $\mathbf{R}^n$ to $\mathbf{R}^n/\mathbf{OPTI}$, since the order in which we take the quotients does not matter.



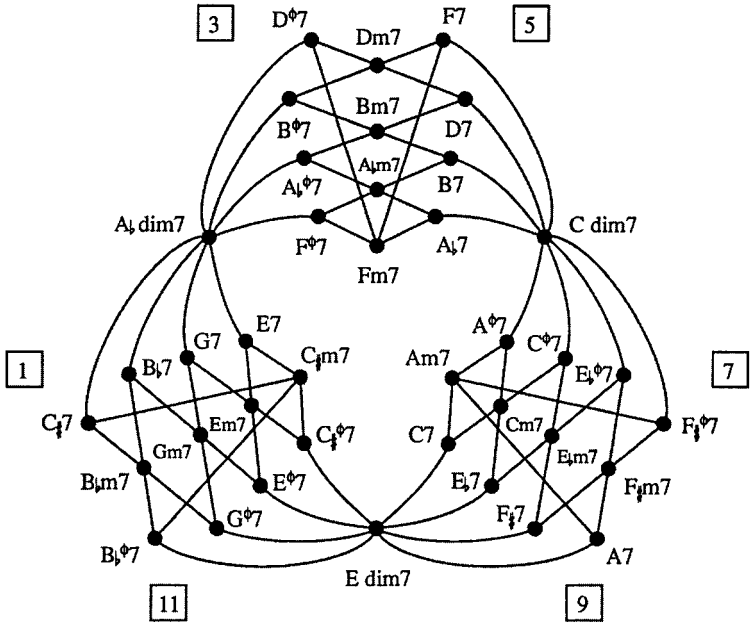
Example 0.1(a). The circle of fifths can be interpreted as depicting minimal voice-leadings between diatonic collections (major scales).



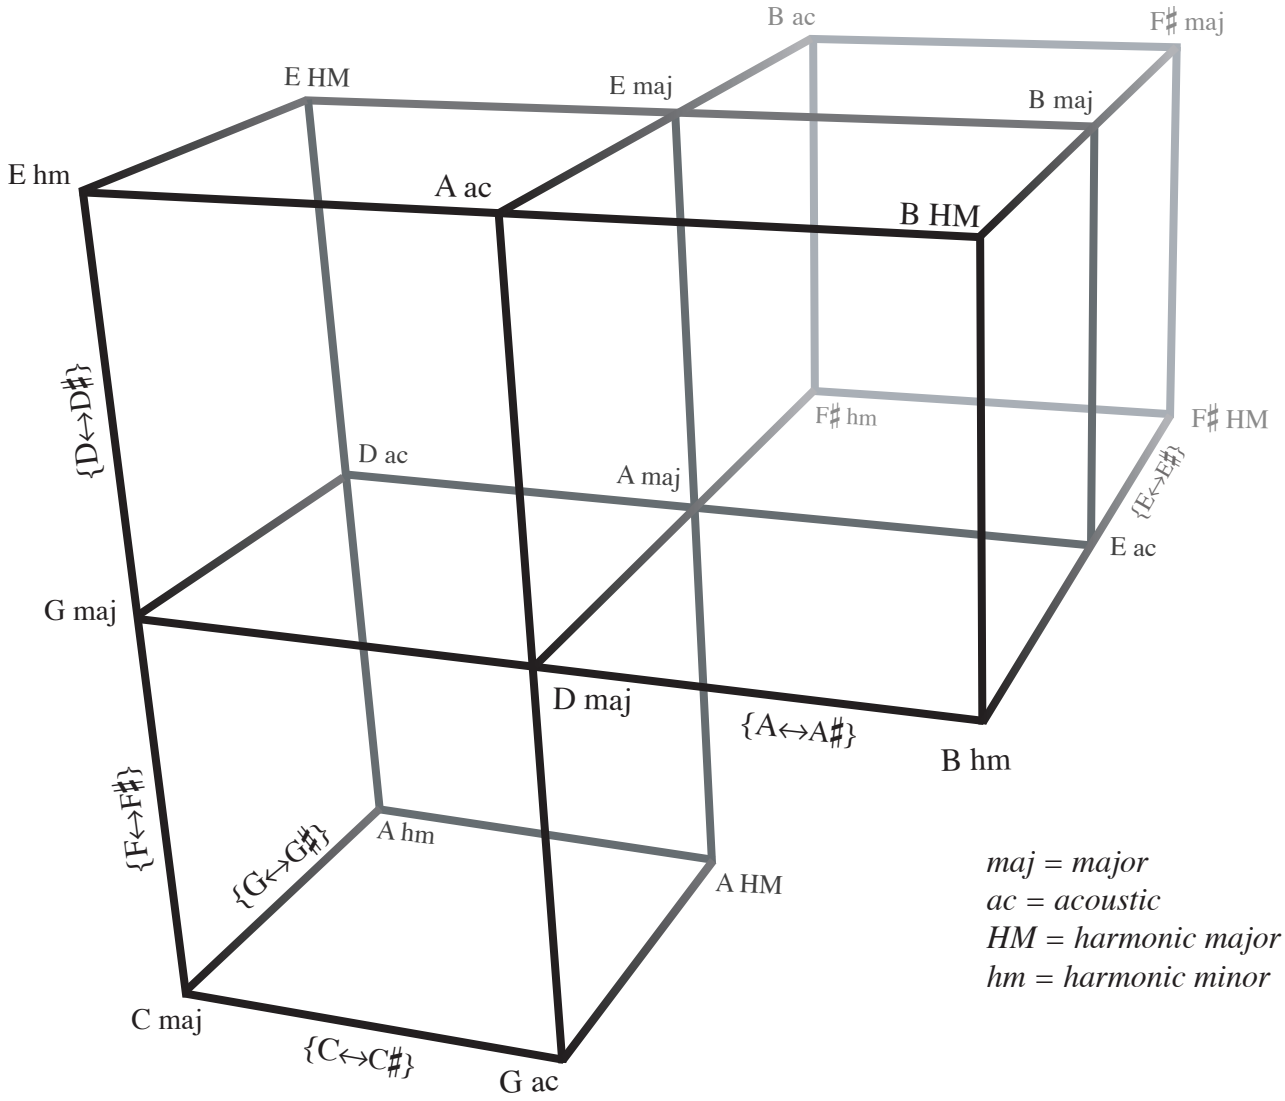
Example 0.1(b). The *Tonnetz,* shown here in its dual form. The graph displays some but not all of the efficient voice-leadings among the 24 familiar major and minor triads. Triads connected by horizontal lines share both "root" and "fifth," and can be connected by voice-leading in which one note moves by one semitone. Triads along the NE/SW diagonal also share two notes and can be connected by single-semitone voice-leading. Triads along a NW/SE diagonal share two notes and can be connected by voice-leading in which one note moves by two semitones.



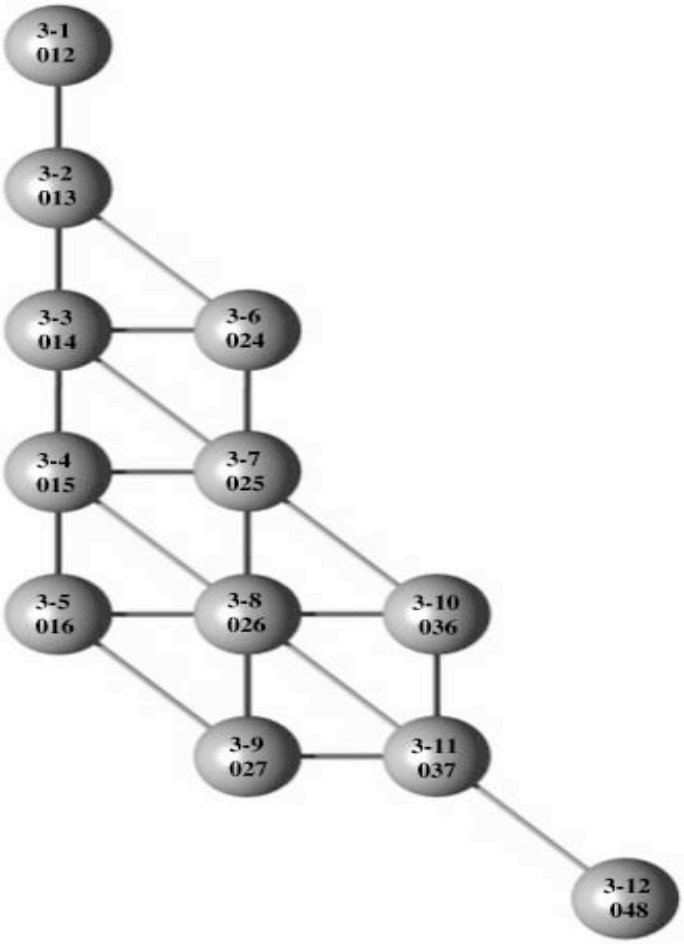
Example 0.1(c). Douthett and Steinbach's "Cube Dance."



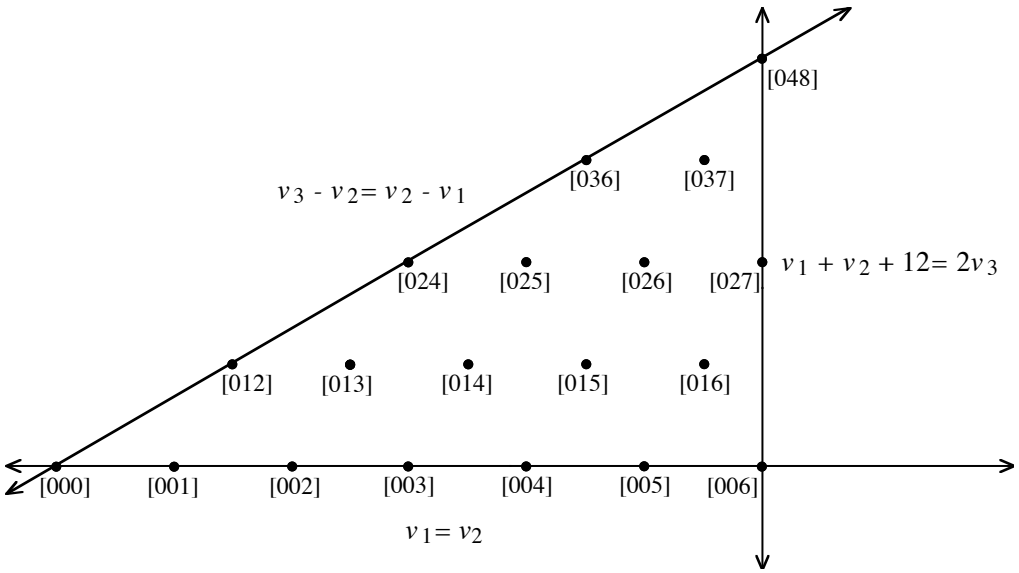
Example 0.1(d). Douthett and Steinbach's "Power Towers."



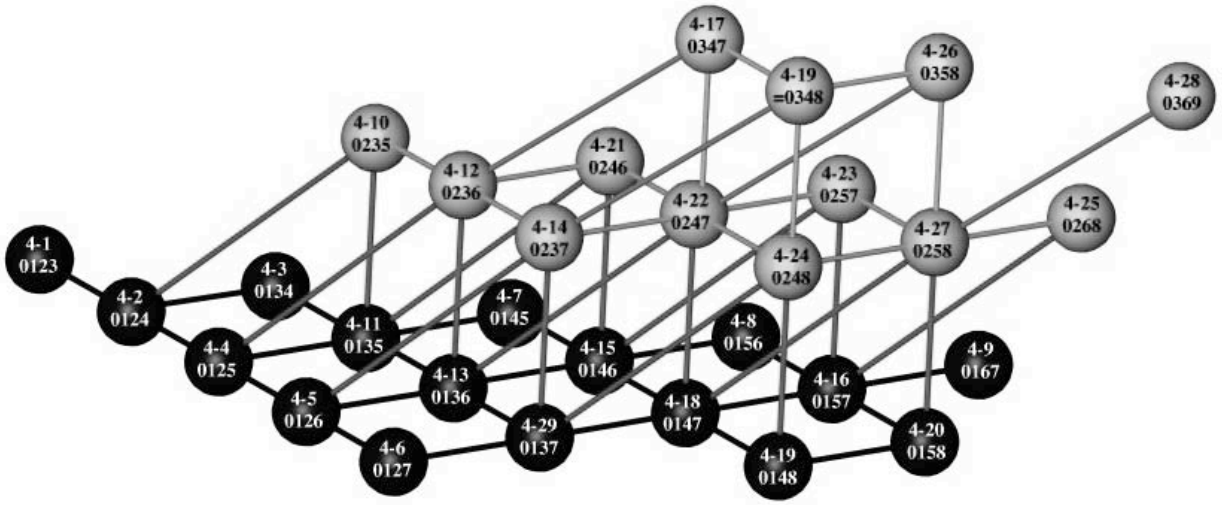
Example 0.1(e). Dmitri Tymoczko's "scale lattice."



Example 0.2(a). Joe Straus's graph of parsimonious connections among trichords

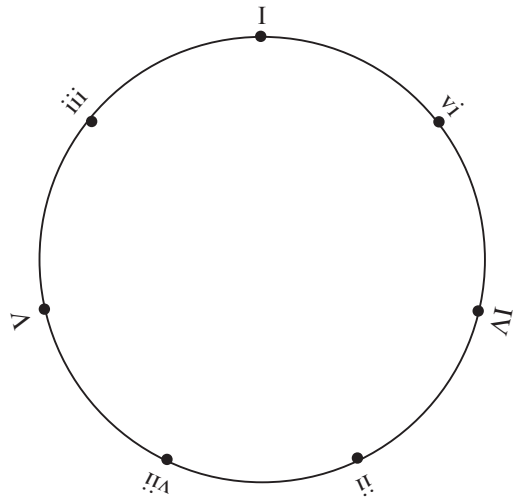


Example 0.2(b). Clifton Callender's continuous "trichord space."

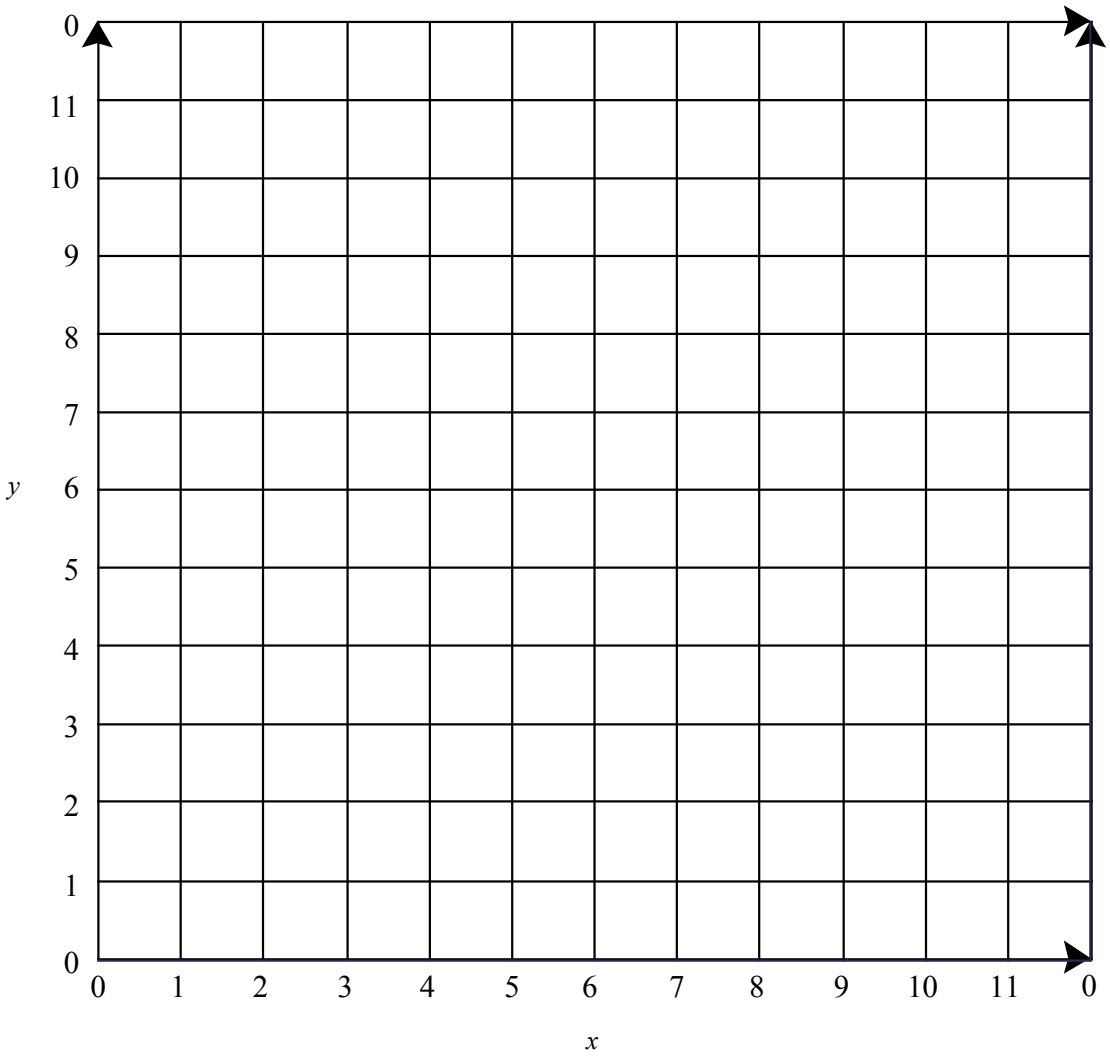


Example 0.2(c). Joe Straus's graph of parsimonious connections among tetrachords

Example 0.2(d) is not available yet.



Example 0.2(e). Parsimonious voice-leading among diatonic triad-classes



Example 0.3(a). John Roeder's "ordered interval" space.

Example 0.3(b) is not available yet.