Continuous Harmonic Spaces

Clifton Callender

Abstract This article generalizes Ian Quinn's recent harmonic characterization of pitch-class sets in equal tempered spaces to chords drawn from continuous pitch and pitch-class spaces. Using the Fourier transform, chords of any real-valued pitches or pitch classes are represented by their *spectra* and located in a harmonic space of all possible chord spectra. Euclidean and angular distance metrics defined on chord spectra correlate strongly with common interval-based similarity measures such as IcVSIM and ANGLE. Thus, we can approximate these common measures of harmonic similarity in continuous environments, applying the corresponding harmonic intuitions to all possible chords of pitches and pitch classes in all possible tuning systems. This Fourier-based approach to harmony is used to compare the properties of twelve-note chords in Witold Lutosławski and Elliot Carter, to analyze the opening section of Gérard Grisey's *Partiels*, and to investigate the structural properties underlying the Z-relation (part of ongoing research with Rachel Hall).

1. Introduction

Suppose we wish to compare sonorities in the just intonation systems of Harry Partch and Lou Harrison, the various equal tempered systems of Easley Blackwood, the spectral music of Tristan Murail and Gérard Grisey, or the more freely composed nontempered work of György Ligeti and others. In other words, we wish to compare chords drawn from continuous pitch or pitch-class space rather than some form of discrete tuning. Traditional similarity measures based on interval vectors or subset embedding are of little help for the simple reason that continuous spaces contain an infinite number of intervals. For example, "C-major triads" drawn from twelve-tone equal temperament, {0, 4, 7}, and just intonation, $\{0, 12\log_2\frac{5}{4}, 12\log_2\frac{3}{2}\} \approx \{0, 3.86, 7.02\}$, would be judged by traditional measures as maximally dissimilar, because they contain no shared intervals.[1] Naturally, we would like a means of comparing these sonorities that corresponds to our perception of these chords as being quite similar and, in a sense, equally deserving of the label "C-major triad."

1 The most natural unit of distance in continuous pitch-class space is the octave, so the interval of an octave is 1, a perfect fifth is 7/12, and so forth. However, given readers' familiarity with intervals measured as multiples of a semitone, throughout this article the unit of distance in pitch and pitch-class spaces will be the semitone. Thus, the pitch class $12\log_2\frac{5}{4}$ is approximately 3.86 semitones "higher" than C♮ or 14 cents "lower" than E♮, the pitch class $12\log_2\frac{3}{2}$ is approximately 2 cents "higher" than G♮, and so forth.

Journal of Music Theory 51:2, Fall 2007
DOI 10.1215/00222909-2009-004

One way to compare the two major triads is to measure the smoothness of the voice leading between them, according to some voice-leading metric.[2] The smoothest voice leading consists of one voice (the third) moving 14/100 of a semitone from 4 to 3.86 and another (the fifth) moving by only 2/100 of a semitone from 7 to 7.02, while the root is held as a common tone. By any reasonable metric, this voice leading is extremely smooth, and the two triads should be very close within any voice-leading space.

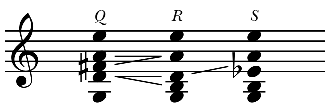

Figure 1. ***Q*** **and** ***R*** **are closer in harmonic space than** ***R*** **and** ***S*****;** ***R*** **and** ***S*** **are closer in voice-leading space than** ***Q*** **and** ***R***

Obviously, such minuscule perturbations of a chord will have only a limited effect on its harmonic content. That is, chords that are very close in a voice-leading space will tend to be close in a harmonic space, as well. (Indeed, in the extreme case of infinitesimal voice leading between two chords, the distance between these two chords in either a voice-leading or harmonic space must also be infinitesimal.) However, the converse is not necessarily true, as demonstrated by the three chords shown in Figure 1. Although *Q* and *R* are not related by transposition or inversion, the two chords sound quite similar. This similarity is easy to explain: both chords are stacks of perfect fifths with one of the fifths divided into a major and a minor third. However, the voice leading between *Q* and *R*, indicated by lines in Figure 1, is not at all smooth. Even allowing for splitting and fusing (in the sense of Straus 2003 and Lewin 1998) {D4} → {B3, D4} and {F♯4, A4} → {A4}, the voice leading requires a total of six semitones of motion. Thus, while *Q* and *R* should be close in harmonic space under any reasonable metric of harmonic distance, these two chords are not particularly close in a voice-leading space. Compare the relation between chords *R* and *S* in which a single voice moves by semitone. Despite the smooth voice leading, the two chords sound quite different—certainly more different than *Q* and *R*.[3] In one case, a relatively large voice leading yields harmonically

2 Roeder (1987), Lewin (1997), Cohn (1998), Straus (2003, 2005), Callender (2004), Tymoczko (2006), Hall and Tymoczko (2007), and Callender, Quinn, and Tymoczko (2008) all propose comparing sets on the basis of their distance from one another in a voice-leading space.

3 *S* does not contain an extended stack of fifths, is not a subset of a diatonic collection, substitutes an embedded augmented triad, {G, B, E♭}, for a major triad, and contains two dissonant intervals not in *R*—interval classes 1 and 6. In contrast, *Q* preserves the stack of perfect fifths, substitutes one major triad for another, is a subset of a diatonic collection, and does not possess a tritone. As pitch-class sets, *Q* and *R* are considered to be significantly more similar than either *R* and *S* or *Q* and *S* by standard similarity measures, including IcVSIM (Isaacson 1990), ISIM2 (Isaacson 1996), similar chords, and in the other a small voice leading yields relatively dissimilar harmonic objects. Thus, there is no simple correlation between distances in voice-leading and harmonic spaces.[4] Though related, these two types of spaces are distinct.

In the remainder of this article, I focus exclusively on approaching continuous harmonic spaces via the Fourier transform—a mathematical technique for discovering periodicities in a wide range of phenomena.[5] Lewin (1959) was the first to note the connection between the Fourier transform and a chord's harmonic content, a connection to which he returned very briefly in 1987 (103–4) and more fully in 2001. Since David Lewin's first article, other theorists have incorporated this mathematical technique into their work, including Vuza (1993), Quinn (2006 and 2007), and Amiot (2007). The following extension of this work into a continuous domain draws substantially upon Ian Quinn's work, which links Lewin's work to a variety of approaches to chord quality, yielding a well-developed theory of chord prototypes and fuzzy harmonic categories. Using the metaphor of "Fourier balances," Quinn employs these mathematical tools in an intuitive manner, such that even those with little mathematical background can follow this Fourier-based approach to music theory. (Indeed, it is quite possible for one to fully digest his work without being aware of the sophisticated mathematical concepts lying just beneath the surface.) While Lewin and Quinn are careful to keep the underlying mechanics of the Fourier transform in the background, it will be necessary to confront some of the details here, given the inherently more complicated nature of continuous versus discrete spaces. However, we will do so only to the extent absolutely necessary, and concrete, clarifying examples appear frequently throughout.[6]

2. The Fourier transform

2.1 Interval cycles

Underlying Quinn's work with the Fourier transform is a particular means of comparing chords with various interval cycles, which provides a kind of harmonic blueprint by which we may characterize a chord and compare it to others. For example, of the three chords discussed in the introduction, *Q* and

---

ANGLE (Scott and Isaacson 1998), SIM (Morris 1979), and SATSIM2 (Buchler 2000). As pitch sets, *Q* and *R* are also considered to be more similar by PM (Morris 1995) and PSATSIM (Buchler 1997), two of the few measures of pitch-set similarity.

4 Tymoczko 2008 explores the relation between voice-leading spaces and harmonic spaces based on the Fourier transform. See note 17 for further discussion.

5 Most musicians come into contact with the Fourier transform in the context of analyzing or modifying audio. Complex sound waves, represented as time-varying changes of intensity, can be transformed into frequency spectra, which can then be analyzed and/or modified and converted back into sound. A good reference for audio applications of the Fourier transform is Smith 2008.

6 My thanks to Michael Buchler, Rachel Hall, Evan Jones, Ian Quinn, two anonymous readers for *JMT*, and especially Dmitri Tymoczko, whose insightful and insistent questions motivated me to think more deeply about numerous sections, especially §6.2, and the more general matter of the relationship between voice-leading and Fourier spaces.

R each contain four members belonging to the cycle of perfect fifths { . . . , G3, D4, A4, E5, . . . }, while S contains only three members belonging to this cycle. Using set intersection as a crude measure of similarity, we could claim that Q and R are more similar to this perfect fifths cycle than is S. But what about the set U = {G3, F♮4, F♯4, A4, E5}? Though U also has three members belonging to the interval cycle in question, is it as similar to this cycle as S, even though the latter just misses having four members in the cycle by a single semitone (substituting E♭4 for D4)? One of Quinn's central insights is that the Fourier transform provides a means of "fuzzification" of membership in (or similarity to) an interval cycle. In the remainder of this section, we get a feel for how this works and the musical relevance of such an approach, before delving into greater specifics in §3.

By an l-cycle we will mean any set of the form $\{\ldots, x - l, x, x + l, \ldots\}$. The l-cycle that contains 0 will be designated $\zeta_l = \{\ldots, -l, 0, l, \ldots,\}$. Consider a set consisting of a single pitch, p, and the interval cycle $\zeta_1 = \{\ldots, -1, 0, 1, \ldots\}$. How well does p "fit" or "correlate" with this semitone cycle? We might interpret the question as asking how "in tune" p is with the twelve-tone equal tempered (12-tet) scale containing pitch 0. The closer p is to any integer, the more in tune it is, or the better it fits, with the 12-tet scale. In the case that p is an integer, then p correlates maximally with ζ_1. The farther p is from any integer, the less in tune it is, or the less well it fits, with the 12-tet scale. In the case that p is exactly halfway between two consecutive integers (1.5, 2.5, 3.5, etc.), then p correlates minimally with ζ_1.

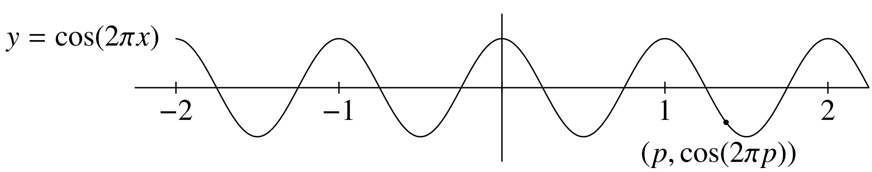

Figure 2. The correlation of p with $\zeta_1 = \{\ldots, -1, 0, 1, \ldots\}$ is $\cos(2\pi p)$

The situation is represented in Figure 2 by a cosine wave with a period of 1 representing the correlation of p with ζ_1. The peaks of the cosine wave correspond to the members of the interval cycle (with a correlation of 1), and the troughs correspond to values that lie exactly halfway between the members of the cycle (with a correlation of −1). For any value of p, the correlation of p with ζ_1 is given by $\cos(2\pi p)$.[7] It is important not to confuse this cosine wave with the interval cycle itself. The interval cycle is a discrete set, while the cosine wave

7 In principle, we could measure this correlation with a number of different periodic functions. For example, we could use the triangle wave, which connects the appropriate maxima and minima in a linear fashion. However, as we shall see, using the Fourier transform yields deep connections with existing methods of measuring similarity based on interval content.

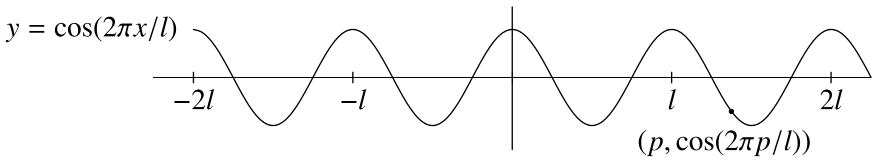

Figure 3. The correlation of *p* with ζ_l = { . . . , –*l*, 0, *l*, . . . } is cos(2π*p*/*l*)

is a continuous function of the correlation of pitches with an interval cycle of the same period and phase.

We can generalize the situation a bit by considering how well the pitch p correlates with the l-cycle containing 0, ζ_l. The situation is represented in Figure 3 by a cosine with a period of l representing the correlation of p with ζ_l. For any value of p, the correlation of p with ζ_l is given by $\cos(2\pi p/l)$. The closer p is to any multiple of l, the better it fits or correlates with ζ_l; conversely, the farther p is from any multiple of l (or the closer it is to $-l/2$, $l/2$, $3l/2$, etc.), the worse it fits or correlates with ζ_l.

Generalizing the situation further, we can consider how well a set of multiple pitches fits with a given interval cycle by simply adding up the correlations of the individual members with this interval cycle. For example, consider the set $S = \{-1, p, 2\}$ and the interval cycle of semitones ζ_1. (Again, we can interpret the "fit" of S with ζ_1 as an indication of how in tune S is with the 12-tet scale.) The pitches −1 and 2 each have a correlation of 1 with ζ_1, and, as we saw above, the correlation of p with ζ_1 is $\cos(2\pi p)$. So, the total "fit" or correlation of S with ζ_1 is $2 + \cos(2\pi p)$. The closer p is to an integer, the closer this fit for set S is to 3; the farther p is from an integer, the closer this fit is to 1. (See Figure 4.)

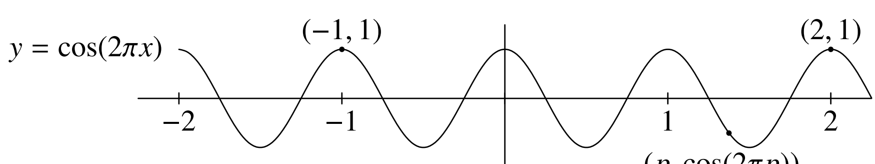

Figure 4. The correlation of *S* = {–1, *p*, 2} with ζ_1 is 1 + 1 + cos(2π*p*)

We can also compare a set not only with a given interval cycle, but also with all possible transpositions of this cycle. For example, consider the set $P = \{-1.5, -.5, 1.5\}$. The correlation between P and ζ_1 is as low as possible, −3, since every pitch is as "out of tune" as possible. However, it is clear that P is maximally in tune with the 12-tet scale $T_{.5}(\zeta_1) = \{\ldots, -.5, .5, 1.5, \ldots\}$, the transposition of ζ_1 up (or down) by a quarter tone. The situation is represented in Figure 5, where the cosine wave has been shifted to the right by .5.

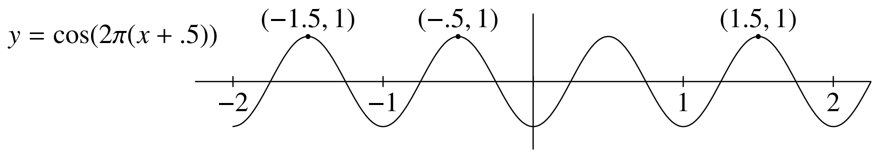

Figure 5. The correlation of $P = \{-1.5, -.5, 1.5\}$ with $T_h(\zeta_1)$ is maximized when $h = .5$ (mod 1); the correlation for individual members of P is given by $\cos(2\pi(p + h))$

Thus, the maximum correlation between P and some transposition of ζ_1 is 3; we will say that P has a *magnitude* of 3 with respect to a 1-cycle. More generally, the magnitude of any set with respect to an l-cycle is the maximum correlation between the set and $T_x(\zeta_l)$ as x varies continuously. Intuitively, we simply slide the cosine wave to the right until the sum of the correlations of individual pitches with the corresponding interval cycle is maximized. (See §2.3 below.)

2.2 Chord spectra

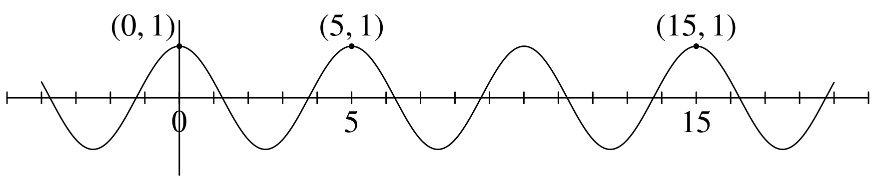

Figure 6. The magnitude of P = {0, 5, 15} with respect to a 5-cycle is 1 + 1 + 1 = 3

Now consider the pitch set $P = \{0, 5, 15\}$ or {C4, F4, E♭5}, which obviously belongs to the 5-cycle ζ_5.[8] The situation is depicted graphically in Figure 6, where the correlation of individual members of P is represented by a cosine wave with a period of 5. Since the values of p corresponding to the members of P coincide with the crests of the cosine wave, by the reckoning of the preceding section, P has a magnitude of 3 with respect to a 5-cycle. In comparison, P correlates less well with 4- and 6-cycles, shown in Figure 7. In the first case, the correlations of individual pitches of P with ζ_4 are 1, 0, and 0, yielding a total magnitude of only 1. (No other transposition of ζ_4 yields a higher correlation with P.) In the second case, the correlations of individual pitches with the 6-cycle in Figure 7 are .5, 1, and –.5, which also yield a total magnitude of 1. (In this case, the sum of the correlations is maximized when the cosine wave of period 6 is shifted to the left by 1/6 of a period. See §2.3 below for details.)

8 Throughout this article, *pitch* sets will be labeled *P*, *Q*, *R*, and so forth. The corresponding *pitch-class* sets generally will be labeled P_O, Q_O, R_O, and so forth whenever it is unclear if a set is of pitches or pitch classes. (The subscript "O" indicates the set modulo octave equivalence.) Additionally, the term *sets* as used in this article is synonymous with the term *multisets*. Multisets are unordered sets in which multiple occurrences of an element are counted separately. For example, the multiset {*a*, *a*, *b*} is not the same as {*a*, *b*}.

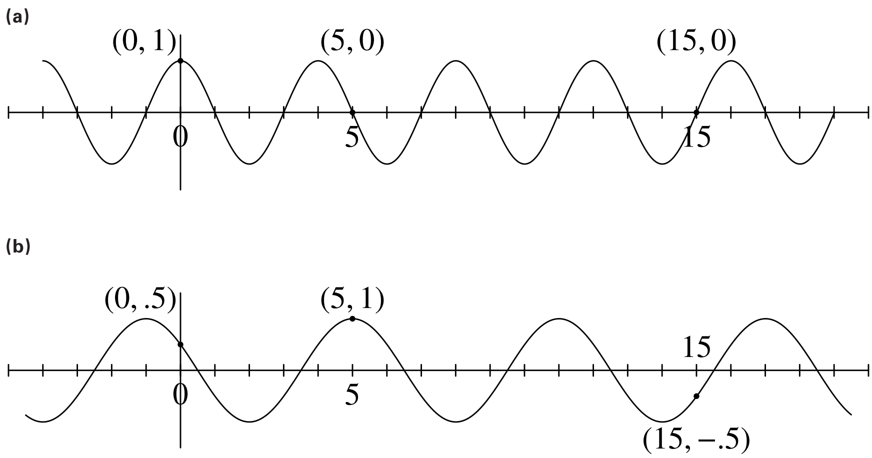

Figure 7. The magnitude of *P* = {0, 5, 15} (a) with respect to a 4-cycle is 1 + 0 + 0 = 1; (b) with respect to a 6-cycle, the magnitude is .5 + 1 – .5 = 1

Since we are working in a continuous space, there is no need to limit our investigation to cycles of an integer number of semitones. Using the intuitive method outlined above, we can find the magnitude of *P* with respect to cycles of 5.1 semitones, 2π semitones, or x semitones for any value of x. Figure 8 graphs the magnitude of *P* with respect to x-cycles as x varies continuously. The magnitudes of *P* with respect to x-cycles for all possible values of x form the *spectra* of *P*. (For more on chord spectra, including the derivation of graphs such as Figure 8, see §3.2.)

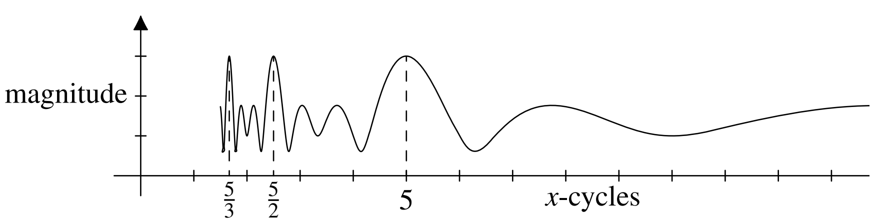

Figure 8. Magnitude of *P* = {0, 5, 15} with respect to *x*-cycles

As expected, since *P* has maximal magnitude with respect to a 5-cycle, there is a peak at $x = 5$ where the graph is equal to 3. (Since there are three members of the set and the crest of the cosine wave is 1, no other peak can be higher than 3, and this value is obtained only when the crests of the cosine wave are perfectly aligned with values corresponding to the members of the set. This is the sense in which a set has maximal magnitude with respect to an interval cycle.) There are, however, a number of other peaks, located at $x = \frac{5}{2}$ and $x = \frac{5}{3}$, that are equally high. Since $\zeta_5 = \{\ldots, -5, 0, 5, \ldots\}$ is a subset of the $\frac{5}{2}$-cycle $\{\ldots, -5, -\frac{5}{2}, 0, \frac{5}{2}, 5, \ldots\}$, *P* obviously has maximal magnitude with respect to a $\frac{5}{2}$-cycle. Likewise, since ζ_5 is a subset of the $\frac{5}{3}$-cycle $\{\ldots, -5, -\frac{10}{3}, -\frac{5}{3}, 0, \frac{5}{3}, \frac{10}{3}, 5, \ldots\}$, P will also have maximal magnitude with respect to a $\frac{5}{3}$-cycle. More generally, if a set has maximal magnitude with respect to an x-cycle, it will also have maximal magnitude with respect to an $\frac{x}{k}$-cycle, where k is any integer. Thus, continuing the graph of Figure 8 to smaller values of x, there will be identical peaks at $x = \frac{5}{4}$, $x = \frac{5}{5} = 1$, and so on.

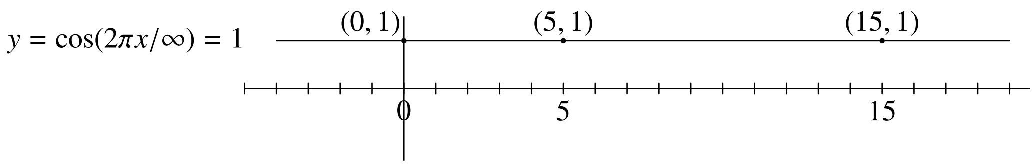

Figure 9. Magnitude of *P* = {0, 5, 15} with respect to the ∞-cycle

There is one other cycle with respect to which *any* set has maximal magnitude—the ∞-cycle. A cosine with an infinitely long period is simply a horizontal line at $y = 1$, shown in Figure 9. It is clear that the magnitude of a set with respect to the ∞-cycle will be equal to the cardinality of the set.

There are other (lesser) peaks in the graph of P at $x \approx 3$, $x \approx 3\frac{2}{3}$, and numerous other values for x. The peak near $x = 3$ is easy to understand by raising the middle pitch from F4 to F♯4 and noting that the resulting chord, $P' = \{0, 6, 15\}$, belongs to a minor-third cycle. The closeness of P and P' in voice-leading space accounts for the relatively strong magnitude of P with respect to a 3-cycle as well as the magnitude of P' with respect to a 5-cycle. (See Figure 10 for the graph of P'.)

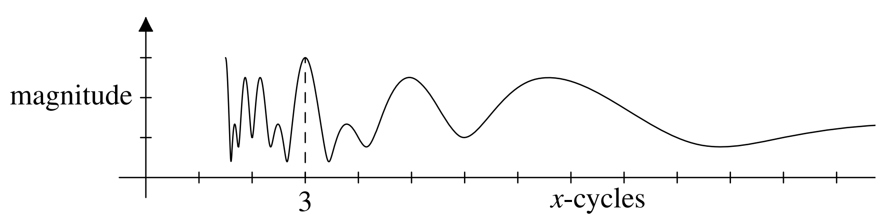

Figure 10. Magnitude of *P′* = {0, 6, 15} with respect to *x*-cycles

The peak at $x \approx 3\frac{2}{3}$ in Figure 8 is only slightly more complicated. First, note that taking every third pitch in a $3\frac{2}{3}$- or $\frac{11}{3}$-cycle yields a cycle of major sevenths: $(\ldots, 0, 3\frac{2}{3}, 7\frac{1}{3}, 11, \ldots)$. Lowering the middle pitch of P from F4 to E4 forms a major seventh with E♭5 and a major third with C4, which is *nearly* equal to $3\frac{2}{3}$ semitones. Raising C4 by $\frac{1}{3}$ of a semitone yields $P'' = \{\frac{1}{3}, 4, 15\}$, which belongs to the $3\frac{2}{3}$-cycle $\{\ldots, \frac{1}{3}, 4, 7\frac{2}{3}, \ldots\}$ and is graphed in Figure 11.

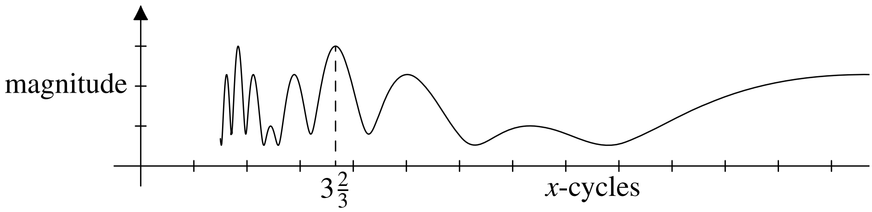

Figure 11. Magnitude of *P*″ = {1/3, 4, 15} with respect to *x*-cycles

Again, the closeness of P and P'' in voice-leading space accounts for the relatively strong peaks near $x = 3\frac{2}{3}$ in the graph of P and $x = 5$ in the graph of P''.

Care must be taken when interpreting these graphs to distinguish between intervals within a chord and the magnitude of this chord with respect to various interval cycles. For example, since there is one minor seventh in P and the value of the graph in Figure 8 at $x = 10$ is 1, we might be tempted to assert an exact correspondence between the two. However, this is not the case, as shown by the graph of {0, 5, 10, 15} in Figure 12. The addition of pitch 10 to chord P creates a second minor seventh, but the value of the corresponding graph at $x = 10$ is now zero. This is because the minor sevenths are separated by a perfect fourth—an interval that bisects a minor seventh. Accordingly, if we shift a cosine wave with a period of 10 so that its crests fall on one of the minor sevenths, its troughs will fall on the other minor seventh, and the two will cancel one another. Some readers may find this lack of a direct correspondence between intervals and magnitudes with respect to interval cycles disconcerting, but it is precisely this property that allows us to employ an approach based on interval content in continuous spaces. As we shall see (particularly in §6.3), there is a very strong connection between interval content and the Fourier transform.

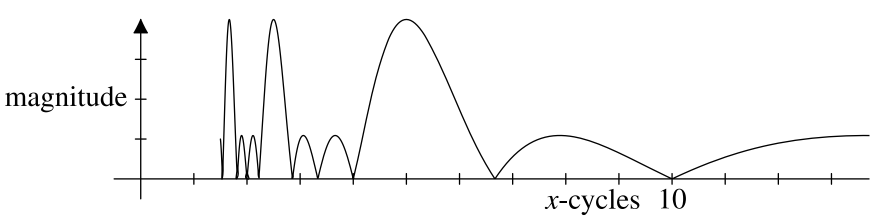

Figure 12. Magnitude of {0, 5, 10, 15} with respect to *x*-cycles

2.3 Transposition and phase

In §2.1 we measured the correlation of a pitch with ζ_l by a cosine wave with a period of l. Suppose we wish to similarly measure the correlation of a pitch with the transposition of this interval cycle by x semitones, $T_x(\zeta_l)$. In this case, the correlation is given by the same cosine wave *shifted to the right* by x. In other words, transposition of an interval cycle corresponds to a *phase shift* of the cosine wave indicating the correlation of a pitch with the interval cycle. For example, consider the correlation of a pitch with the interval cycle $T_1(\zeta_5) = \{\ldots, 1, 6, 11, \ldots\}$. In this case, correlation with the interval cycle is measured by a cosine wave with a period of 5 shifted by 1/5 of a cycle (the transposition factor divided by the generating interval of an interval cycle), which yields a phase of $\frac{360°}{5} = 72°$ or $2\pi/5$ radians (Figure 13a). (2π radians is equivalent to 360°. I will use radians instead of degrees for the remainder of the article.) Transposing ζ_5 by x semitones corresponds to a shift of the relevant cosine wave by $x/5$ cycles or a phase of $2\pi x/5$. More generally, transposing ζ_l by x corresponds to a shift of the relevant cosine wave by x/l cycles or a phase of $2\pi x/l$ (Figure 13b).

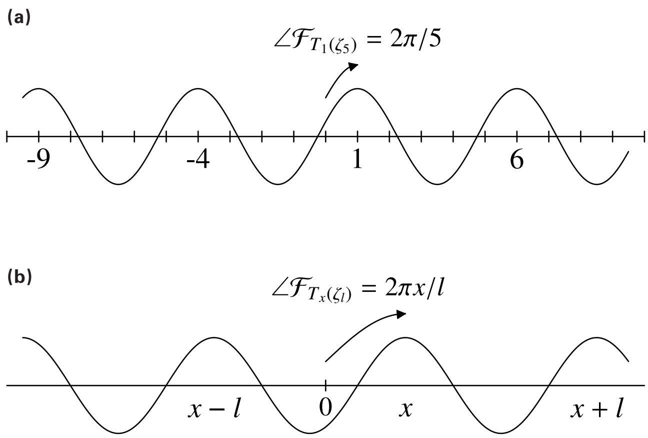

Figure 13. Transposition corresponds to a phase shift: (a) $T_1(\zeta_5)$, (b) $T_x(\zeta_l)$

Dividing by the generating interval, l, of an interval cycle as in the previous paragraph is the same as multiplying by $1/l$, which is the interval cycle *frequency*. The frequency is the number of cycles per semitone. For example, a semitone spans 1/5 of a 5-cycle and two complete cycles of a quartertone cycle; thus, these cycles have frequencies of 1/5 and 2, respectively. Using $z = 1/l$ for the frequency, we can rewrite the phase of the cosine wave measuring correlation with $T_x(\zeta_l)$ as simply $2\pi xz$.

We can combine the magnitude and phase of a set with respect to a given frequency by vectors in a two-dimensional plane, where the magnitude corresponds to the length of the vector and the phase corresponds to the angle of a clockwise rotation from due north.[9] For example, with respect to a 5-cycle, $P = \{0, 5, 15\}$ has a magnitude of 3 and a phase of 0. We can represent the

9 The association in this article of zero phase with due north and a *positive* change of phase with clockwise rotation (or rightward shift) runs contrary to mathematical convention. By convention, zero phase is associated with the positive x-axis and a positive change of phase with *counterclockwise* rotation (or leftward shift). The contrary association adopted in this article is in agreement with the clock-face diagrams in Quinn 2007.

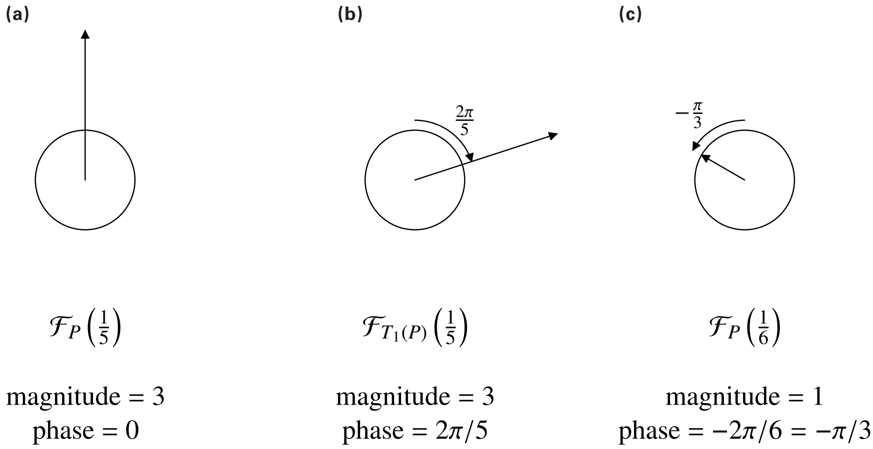

Figure 14. Representing magnitudes and phases of (a) *P* = {0, 5, 15} with respect to a 5-cycle, (b) $T_1(P)$ with respect to a 5-cycle, and (c) *P* with respect to a 6-cycle

situation by a vector with a length of three pointing due north (see Figure 14a). The vector corresponding to $T_1(P)$ with respect to a 5-cycle is shown in Figure 14b, where the vector in Figure 14a has been rotated clockwise by 1/5. Recall from Figure 7b that with respect to a 6-cycle P has a magnitude of 1. Also note that the cosine wave that gives the highest correlation is shifted to the left by one semitone or one-sixth of a cycle, yielding a phase of $-2\pi/6 = -\pi/3$. The corresponding vector, with a length of 1 and rotated counterclockwise by 1/6, is shown in Figure 14c.

Combining magnitudes and phases yields the *Fourier transform* of set P with respect to frequency $z = 1/l$, written $\mathcal{F}_P(z)$ or, in abbreviated form, $\mathcal{F}_P$.[10] The magnitude is indicated by $|\mathcal{F}_P|$ and the phase by $\angle\mathcal{F}_P$.[11] For example, the Fourier transform of set P with respect to a 6-cycle is $\mathcal{F}_P(1/6)$, where the magnitude is $|\mathcal{F}_P(1/6)| = 1$ and the phase is $\angle\mathcal{F}_P(1/6) = -\pi/3$.

The Fourier transform takes a set from the domain of pitches or pitch classes into the domain of interval frequencies where the underlying periodicities of a set are explicitly represented. It is in the interval-frequency domain that the infinitely many sets in continuous pitch or pitch-class space can be analyzed, compared, and transformed on equal footing. First, we must consider how to calculate the Fourier transform.[12]

10 Strictly speaking, we take the Fourier transform of the *characteristic* function of set P. (See §8.) Nonetheless, for our purposes the slight abuse of mathematical language is convenient and not particularly problematic.

11 The magnitude and phase can be combined into the single complex number, $\mathcal{F}_P(z) = re^{i\theta}$, where r is the magnitude, θ is the phase, and $i = \sqrt{-1}$.

12 Quinn 2007 (part 3) demonstrates the relevance of the Fourier series with respect to equal tempered pitch-class spaces. Sections 3 and 4 of the present article summarize Quinn's work (especially in the comments on Figure 21), tease out the underlying mathematics, and both extend it to continuous pitch-class spaces and connect it with pitch space and the continuous Fourier transform.

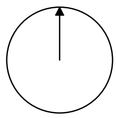

Figure 15. The Fourier transform $\mathcal{F}_{\{0\}}(z)$ for all values of z

3. Pitch sets

3.1 Fourier transform of pitch sets

We begin with the Fourier transform of the singleton pitch set {0}. Since a singleton is (trivially) a subset of all interval cycles (up to transposition), the magnitude of its transform will be 1 for all values of z. Moreover, since the interval cycle to which {0} belongs will obviously contain pitch 0, the phase of the transform will be 0 for all values of z. Figure 15 shows $\mathcal{F}_{\{0\}}(z)$ for all values of z: a unit vector pointing due north.

Next we consider the Fourier transform of the arbitrary singleton pitch set $\{p\}$. While the magnitude of the transform will again be 1 for all values of z, the phase will vary according to the specific values of p and z. Since p is simply $T_p(0)$, we know from section 2.3 that the phase will be $2\pi pz$. For example, Figure 16 shows the Fourier transforms of {0}, {5}, and {15} for $z = 1/6$ (or a 6-cycle) with magnitudes of 1 and phases of 0, $2\pi \cdot \frac{5}{6} = \frac{5}{3}\pi = -\pi/3$, and $2\pi \cdot \frac{15}{6} = 5\pi = \pi$, respectively.

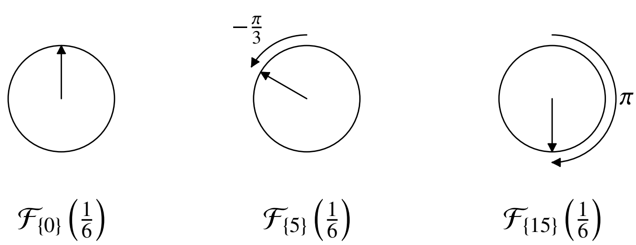

Figure 16. The Fourier transforms of {0}, {5}, and {15} with respect to a 6-cycle

The Fourier transform of a pitch set containing multiple pitches is found by adding together the transforms of the individual pitches.[13] Adding two vectors together is accomplished by simply placing the tail of one vector on the head of the other. The resulting vector that extends from the tail of the first to the head of the last is the sum of the vectors. Since arrow addition is commutative, we can add the arrows in any order. For example, to find the Fourier transform of $P = \{0, 5, 15\}$ with respect to a 6-cycle, we add the vectors

13 This follows from the general principle that $\mathcal{F}_{P \cup Q} = \mathcal{F}_P + \mathcal{F}_Q$. (See §8.)

for {0}, {5}, and {15} shown in Figure 16. The vector resulting from this summation, shown in Figure 17a, has a magnitude of 1 and a phase of $-\pi/3$, confirming the informal results shown in Figure 7b. In this particular case, there is a means to simplify the addition of vectors. Note that the vectors associated with the transforms of {0} and {15} with respect to a 6-cycle point in opposite directions with the same magnitude. Thus, when these two vectors are added, they will cancel each other, yielding the zero vector (a point) shown in Figure 17b. Since adding any vector to the zero vector will yield the original vector, the transform of {0, 5, 15} with respect to a 6-cycle will be the same as that for {5}, shown in Figure 17c. This canceling of vectors is an important phenomenon that will become particularly useful in §7.

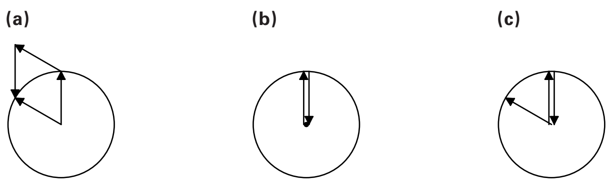

Figure 17. The Fourier transform of {0, 5, 15} with respect to a 6-cycle

In order to accurately measure the magnitude and phase of the Fourier transform of pitch sets, it is helpful to first convert the corresponding vectors to Cartesian coordinates. Again, we begin with singleton pitch sets. Given a set $\{p\}$ and a frequency z, the vector associated with the Fourier transform $\mathcal{F}_{\{p\}}(z)$ extends from the origin to the point (x, y), where $x = \sin(2\pi pz)$ and $y = \cos(2\pi pz)$. For example, letting $p = 5$ and $z = 1/8$ (an 8-cycle), the vector associated with the Fourier transform $\mathcal{F}_{\{5\}}(1/8)$ extends from the origin to the point $\mathbf{x}_1 = (-\sqrt{2}/2, -\sqrt{2}/2)$, since $\sin(2\pi \cdot \frac{5}{8}) = \cos(2\pi \cdot \frac{5}{8}) = -\sqrt{2}/2$. Likewise, we can calculate that the vectors associated with the Fourier transforms for {0} and {15} with respect to an 8-cycle extend from the origin to the points $\mathbf{x}_2 = (0, 1)$ and $\mathbf{x}_3 = (-\sqrt{2}/2, \sqrt{2}/2)$, respectively.

As before, we can find the vector associated with the Fourier transform of $P = \{0, 5, 15\}$ with respect to an 8-cycle by adding the vectors for the individual pitches. Having converted the magnitude and phase information to Cartesian coordinates, this is equivalent to adding the x and y coordinates of each of the three points, $\mathbf{x}_1$, $\mathbf{x}_2$, and $\mathbf{x}_3$, independently. Summing the x coordinates, we have $-\sqrt{2}/2 + 0 - \sqrt{2}/2 = -\sqrt{2}$; summing the y coordinates, we have $-\sqrt{2}/2 + 1 + \sqrt{2}/2 = 1$. Thus, the Fourier transform of P with respect to an 8-cycle, $\mathcal{F}_P(1/8)$, is represented by the vector extending from the origin to the point $(-\sqrt{2}, 1)$.

More generally, for an arbitrary pitch set Q and frequency z, the Fourier transform $\mathcal{F}_Q(z)$ is represented by the vector extending from the origin to the

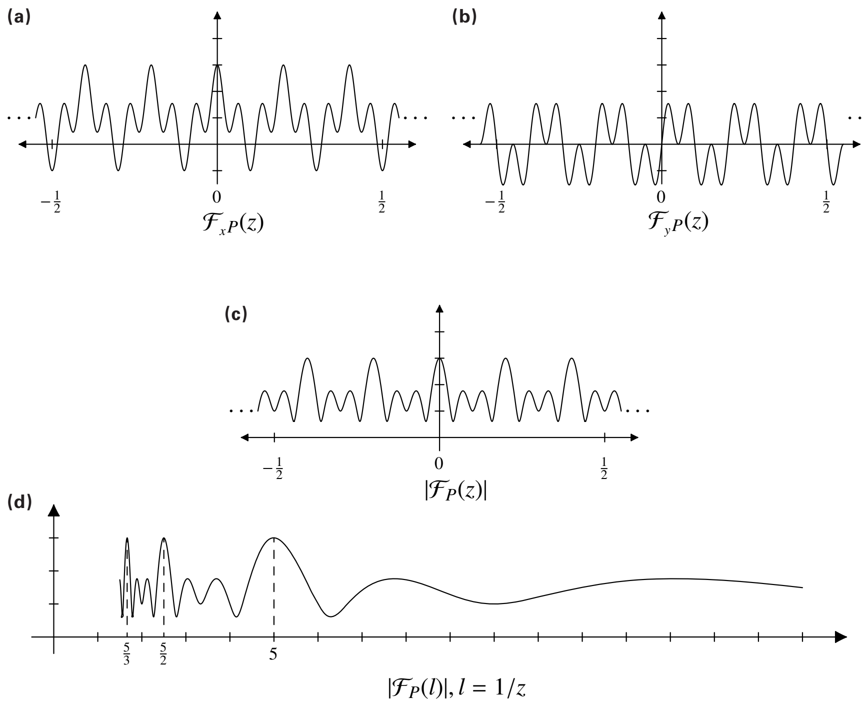

Figure 18. Derivation of the spectrum $|\mathcal{F}_{\{0,5,15\}}|$ graphed in Figure 8

point (x, y), where $x = \sum_{q \in Q} \sin(2\pi qz)$ and $y = \sum_{q \in Q} \cos(2\pi qz)$. We will call x the *sine* component of the transform and y the *cosine* component, respectively, written $\mathcal{F}_{xQ}$ and $\mathcal{F}_{yQ}$.

We can calculate the length of this vector, and thus the magnitude of the associated Fourier transform, by using the Pythagorean theorem. The distance from the origin to the point (x, y) is given by $\sqrt{x^2 + y^2}$. For example, recalling that the vector for the Fourier transform of P with respect to an 8-cycle extends from the origin to the point $(-\sqrt{2}, 1)$, the length of the vector is $\sqrt{(-\sqrt{2})^2 + 1} = \sqrt{2 + 1} = \sqrt{3}$. Thus, the magnitude of the transform is $|\mathcal{F}_P(1/8)| = \sqrt{3}$. This value does not necessarily tell us much without comparisons with other triads, which is why graphs such as Figure 8 are particularly helpful.

More generally, the magnitude of the Fourier transform of Q with frequency z is

$$|\mathcal{F}_Q(z)| = \sqrt{\left(\sum_{q \in Q} \sin(2\pi qz)\right)^2 + \left(\sum_{q \in Q} \cos(2\pi qz)\right)^2}. \quad (1)$$

We will refer to this function as the *spectrum* of Q and use the symbols $\mathcal{Q}$, $\mathcal{R}$, $\mathcal{S}$, and so on, to refer to the spectra of Q, R, and S. We will see later that the *squared* magnitude of the Fourier transform ($|\mathcal{F}_Q|^2$ or $\mathcal{Q}^2$), often called the *power spectrum*, will be even more useful for our purposes.

3.2 Graphing chord spectra

A couple of examples should help to make the details of §3.1 clear. We begin by re-creating the spectrum of $P = \{0, 5, 15\}$, which was graphed in Figure 8, according to the following steps. (See Figure 18.)

(a) Find the sine component of the Fourier transform, $\mathcal{F}_{xP}$:
$\mathcal{F}_{xP}(z) = \sin(0) + \sin(5 \times 2\pi z) + \sin(15 \times 2\pi z)$

(b) Find the cosine component of the Fourier transform, $\mathcal{F}_{yP}$:
$\mathcal{F}_{yP}(z) = \cos(0) + \cos(5 \times 2\pi z) + \cos(15 \times 2\pi z)$

(c) Find the magnitude of the Fourier transform:

$$|\mathcal{F}_P(z)| = \sqrt{\left[\mathcal{F}_{xP}(z)\right]^2 + \left[\mathcal{F}_{yP}(z)\right]^2}$$

(d) Since the graph in Figure 8 is plotted as a function of interval size, l, which is the reciprocal of interval frequency z, the final step is to substitute $l = 1/z$ for z and graph

$$\left|\mathcal{F}_P(l)\right| = \sqrt{(1 + \cos(10\pi/l) + \cos(30\pi/l))^2 + (\sin(10\pi/l) + \sin(30\pi/l))^2}$$

For a second example, we compare the spectra of the three pitch sets from the introduction: $Q = \{$G3, D4, F♯4, A4, E4$\}$, $R = \{$G3, B3, D4, A4, E4$\}$, and $S = \{$G3, B3, E♭4, A4, E4$\}$. The spectra of Q and R are superimposed in Figure 19a, and those of R and S are superimposed in Figure 19b. (In the graphs, the horizontal axis is $l = 1/z$, the period of each sinusoid or, equivalently, the size of the intervals in each interval cycle.) It is clear from visual inspection of the two spectra in Figure 19a that Q and R are quite similar harmonically. Both contain strong and virtually identical peaks near $l = 7$ and all of its integer divisors, $l = 7/2$, $l = 7/3$, and so forth, but no other peaks of significant magnitude. In addition, all of the differences between the two spectra involve relatively weak components. By comparison, the spectra of R and S differ significantly. Their peaks do not align except for the somewhat weaker peak near $x = 7$, and those portions that are similar (e.g., the region between $x = 4$ and $x = 4.5$) involve weak components. The intuitions about harmonic similarity discussed in the introduction are made manifest in these graphs of chord spectra. (We will see in §6 how to make these intuitions more concrete by measuring the distance between these spectra.)

4. Pitch-class sets

4.1 Fourier transform of pitch-class sets

For the Fourier transform of pitch sets, it is necessary to consider all possible interval cycles. However, for pitch-class sets, it is necessary only to consider those interval cycles that divide the octave into an integral number of parts, such as the octave itself, the tritone, the augmented triad, the diminished

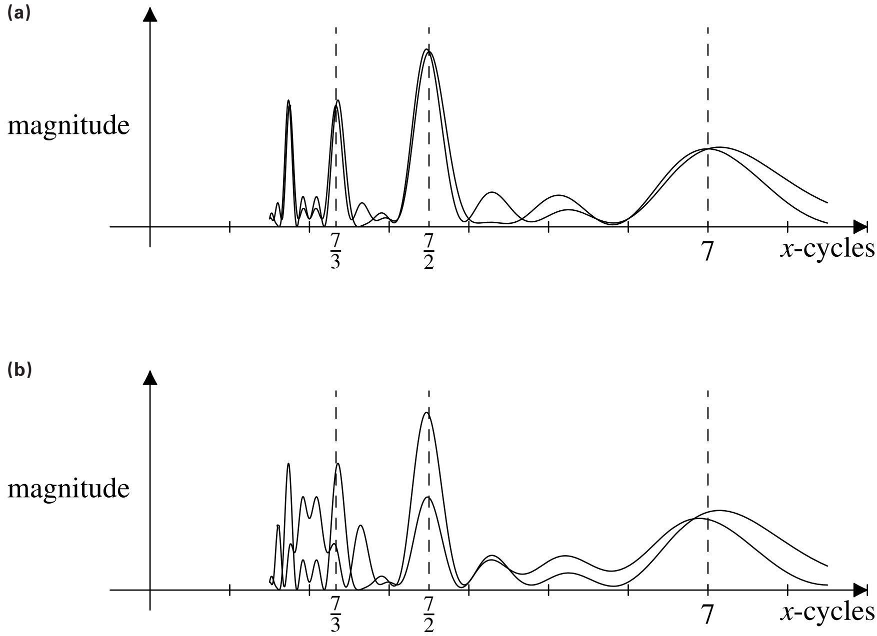

Figure 19. Comparison of the spectra of the three pitch sets from Figure 1: (a) *Q* and *R*, (b) *R* and *S*

seventh chord, the five-note set-class {0, 2.4, 4.8, 7.2, 9.6}, and so forth.[14] Dividing the octave into $k \in \mathbb{Z}$ parts yields a $12/k$-cycle, which has a corresponding frequency of $z = k/12$. Thus, while the Fourier transform of a pitch set is (generally) a continuous function of interval frequency, the transform of a pitch-class set is a discrete series (at all frequencies $k/12$) called the Fourier series.

Knowing the spectrum of any pitch set, P, we can easily derive the spectrum of its corresponding pitch-class set, notated P_O, in a manner demonstrated by Figure 20. Figure 20a shows the spectrum of the pitch set $Q = \{0, 5, 15\}$, with the values at frequencies of the form $k/12$ marked by vertical lines. These vertical lines are the spectrum of the pitch-class set Q_O, shown in Figure 20b. In other words, the Fourier transform of Q_O is a *sampled* version of the transform of Q, where the sampling occurs at all frequencies of the form $k/12$. We will refer to these frequencies as the *harmonics* of the octave, with the frequency $k/12$ being the kth harmonic. Formally, the sine and cosine

14 If I is rational, then any I-cycle in pitch space is equivalent to some equal division of the octave in pitch-class space. For example, consider an infinite cycle of perfect fifths, a 7-cycle. As is well known, an infinite cycle of perfect fifths is equivalent to a 12-fold division of the octave, disregarding order. More generally, any I-cycle in pitch space where I is a rational number of the form a/b is equivalent to a $\frac{12b}{\gcd(a,12)}$-fold division of the octave in pitch-class space. Irrational interval cycles in pitch space such as $(\ldots, -\sqrt{2}, 0, \sqrt{2}, \ldots)$ do not yield closed interval cycles in pitch-class space. (See §5.)

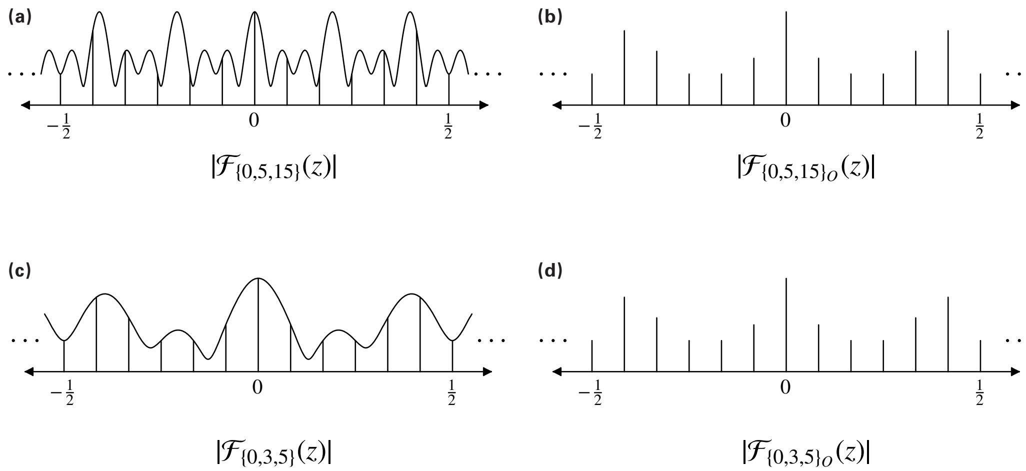

Figure 20. Derivation of the transform of $\{0, 3, 5\}_O$ from (a and b) {0, 5, 15} and (c and d) {0, 3, 5}

components of the Fourier transform of any pitch-class set, P_O, are

$$\mathcal{F}_{xP}(k/12) = \sum_{p \in P} \sin(2\pi pk/12),\ \mathcal{F}_{yP}(k/12) = \sum_{p \in P} \cos(2\pi pk/12),\ k \in \mathbb{Z}. \qquad (2)$$

Note that it does not matter which representative of each pitch class we use for Equation 2. For example, suppose pitch 15 in Q is replaced by 3 in $Q' = \{0, 3, 5\}$. Even though the spectra of Q and Q' are different (Figure 20c), the spectra of Q_O and Q'_O are the same (Figure 20d). (This, of course, had better be the case, since Q_O and Q'_O are the same pitch-class set.)

4.2 Examples of pitch-class spectra

For reasons that will become clear in §5, it is only necessary to consider a finite number of harmonics for the spectra of equal tempered pitch-class sets. Specifically, for a pitch-class set drawn from n-tone equal temperament, it is only necessary to consider harmonics 0 through $\lfloor n/2 \rfloor$, the integral part of $n/2$. For example, if n is 12 or 13, the entire infinite spectrum of a pitch-class set can be reconstructed from harmonics 0 through 6 (since 6 is the integral part of 13/2).

Figure 21 graphs the spectra of the twelve classical pitch-class trichordal set-classes (those without pitch-class duplications) in twelve-tone equal temperament. The numbered comments below correspond to each of the harmonic components of these twelve set-classes:[15]

15 For more on the interpretation of the magnitudes for various harmonics, see Quinn 2007. Analogous comments apply to the spectra of set-classes with any number of pitch classes.

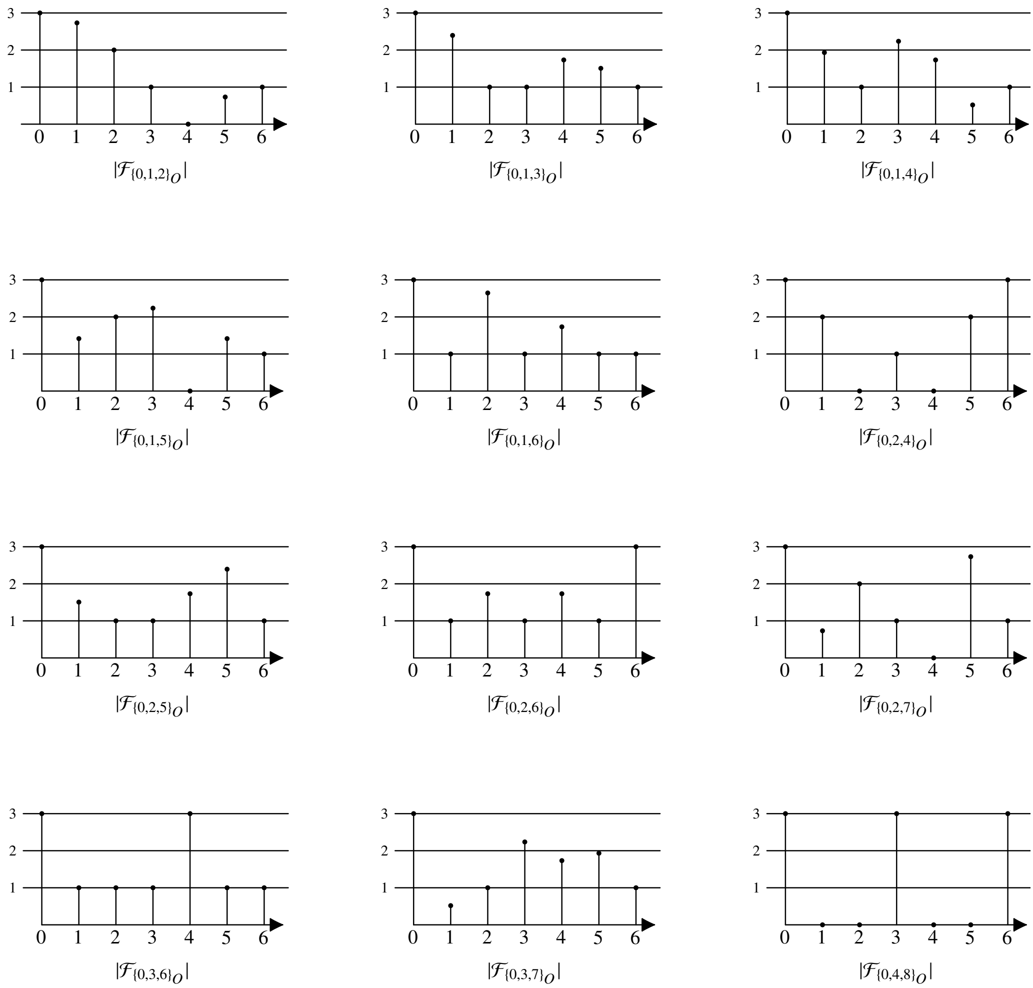

Figure 21. Spectra of twelve-tone equal tempered trichordal pitch-class sets

(0) As noted before, harmonic 0 corresponds to the cardinality of a given pitch or pitch-class set. Thus, each of the spectra graphed in Figure 21 has a magnitude of 3 for the zeroth harmonic.

(1) Harmonic 1 is the magnitude of a pitch-class set with respect to an octave or unison. The magnitude of this harmonic is thus a measure of how unevenly a set divides the octave. Highly uneven sets such as $\{0, 1, 2\}_O$ have a high magnitude for this harmonic, while sets that divide the octave into equal divisions such as $\{0, 4, 8\}_O$ have zero magnitude.

(2) Harmonic 2 is the magnitude of a set with respect to a twofold division of the octave, or a tritone. Not surprisingly, the magnitude

of the second harmonic is quite high for $\{0, 1, 6\}_O$, while the magnitude vanishes for $\{0, 2, 4\}_O$, since this set-class divides the tritone into three equal parts.

(3) Harmonic 3 is the magnitude of a set with respect to a threefold division of the octave, or an augmented triad. Obviously, $\{0, 4, 8\}_O$ has maximal magnitude for this harmonic, and, not surprisingly, the magnitudes for $\{0, 1, 4\}_O$ and $\{0, 1, 5\}_O$ are also quite high.

(4) The magnitude of the fourth harmonic is maximal for $\{0, 3, 6\}_O$, since this set-class belongs to a single fourfold division of the octave, or a diminished seventh chord. The corresponding magnitude vanishes for $\{0, 1, 2\}_O$, $\{0, 1, 5\}_O$, $\{0, 2, 4\}_O$, $\{0, 2, 7\}_O$, and $\{0, 4, 8\}_O$, since the members of these set-classes are drawn equally from the three unique diminished seventh chords.

(5) The fifth harmonic is the magnitude of a set with respect to a fivefold division of the octave, or the set-class $\{0, 2.4, 4.8, 7.2, 9.6\}_O$. The generating interval of this cycle, 2.4, is nearly halfway in between an equal tempered major second and minor third, but twice this interval, 4.8, is very close to an equal tempered perfect fourth. Thus, those set-classes that belong to a small segment of a 5-cycle, such as $\{0, 2, 7\}_O$ and $\{0, 2, 5\}_O$, have significant magnitudes for the fifth harmonic.

(6) Finally, harmonic 6 is the magnitude of a set with respect to a sixfold division of the octave, or a whole-tone collection. The three trichordal set-classes that are subsets of a single whole-tone collection, $\{0, 2, 4\}_O$, $\{0, 2, 6\}_O$, and $\{0, 4, 8\}_O$, have maximal magnitude for this harmonic. Since the remaining trichordal set-classes all have two members that belong to one whole-tone collection while the remaining member belongs to the other collection, each of these set-classes has a magnitude of 1 for this harmonic.

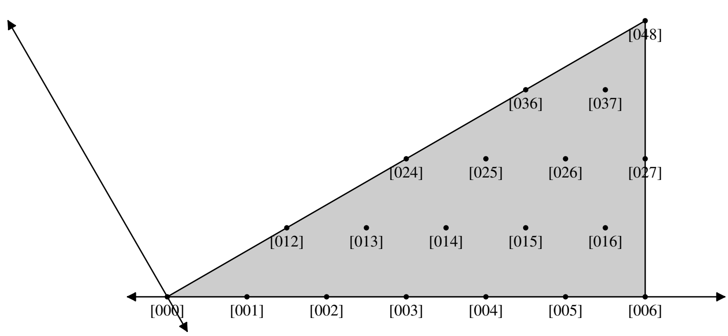

Figure 22. Fundamental region for trichordal set-classes in continuous pitch-class space

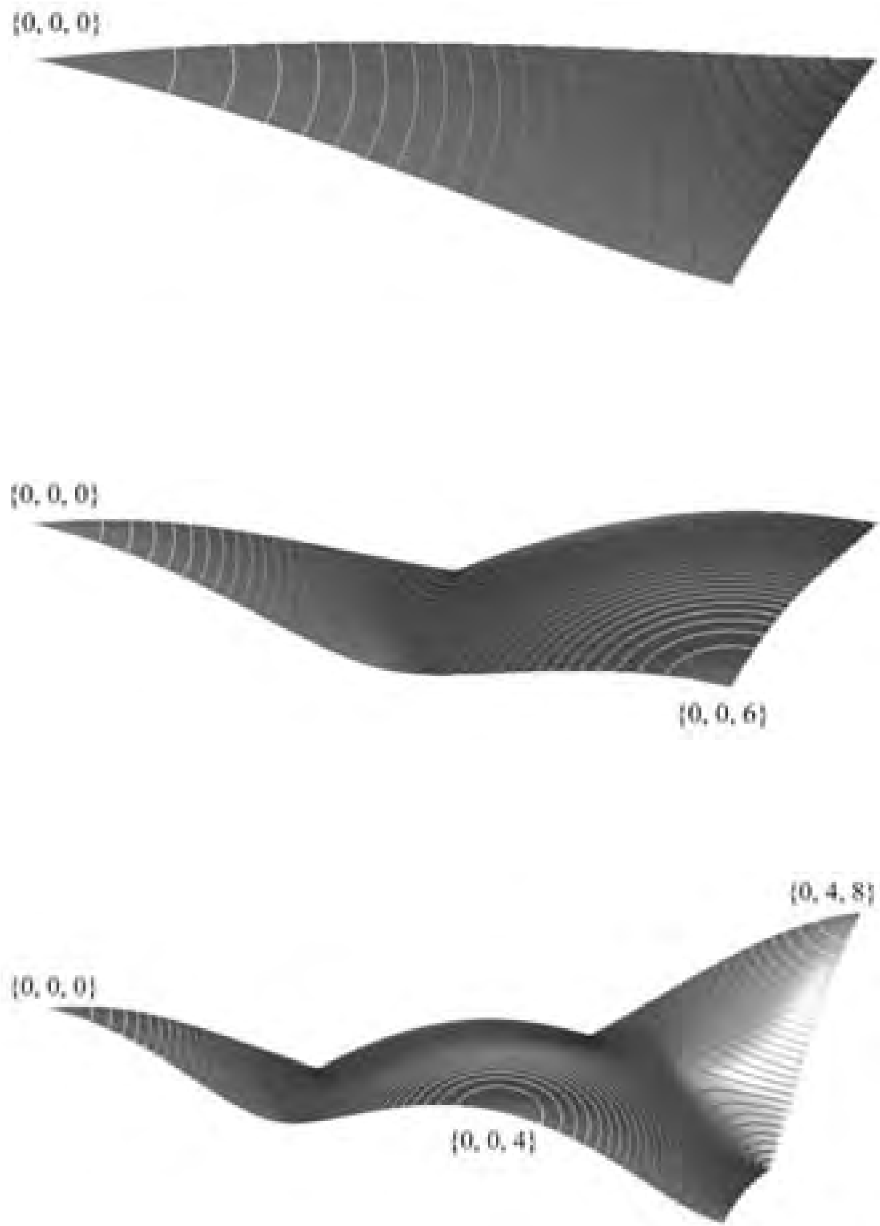

Figure 23. Magnitudes of all trichordal set-classes in twelve-tone equal temperament for the first six harmonics of the octave

While the snapshots provided in Figure 21 provide a harmonic blueprint for each of the set-classes, we can gain a more complete understanding of the underlying topology by considering the spectra of trichords in continuous pitch-class space. We begin by plotting trichord set-classes in two dimensions

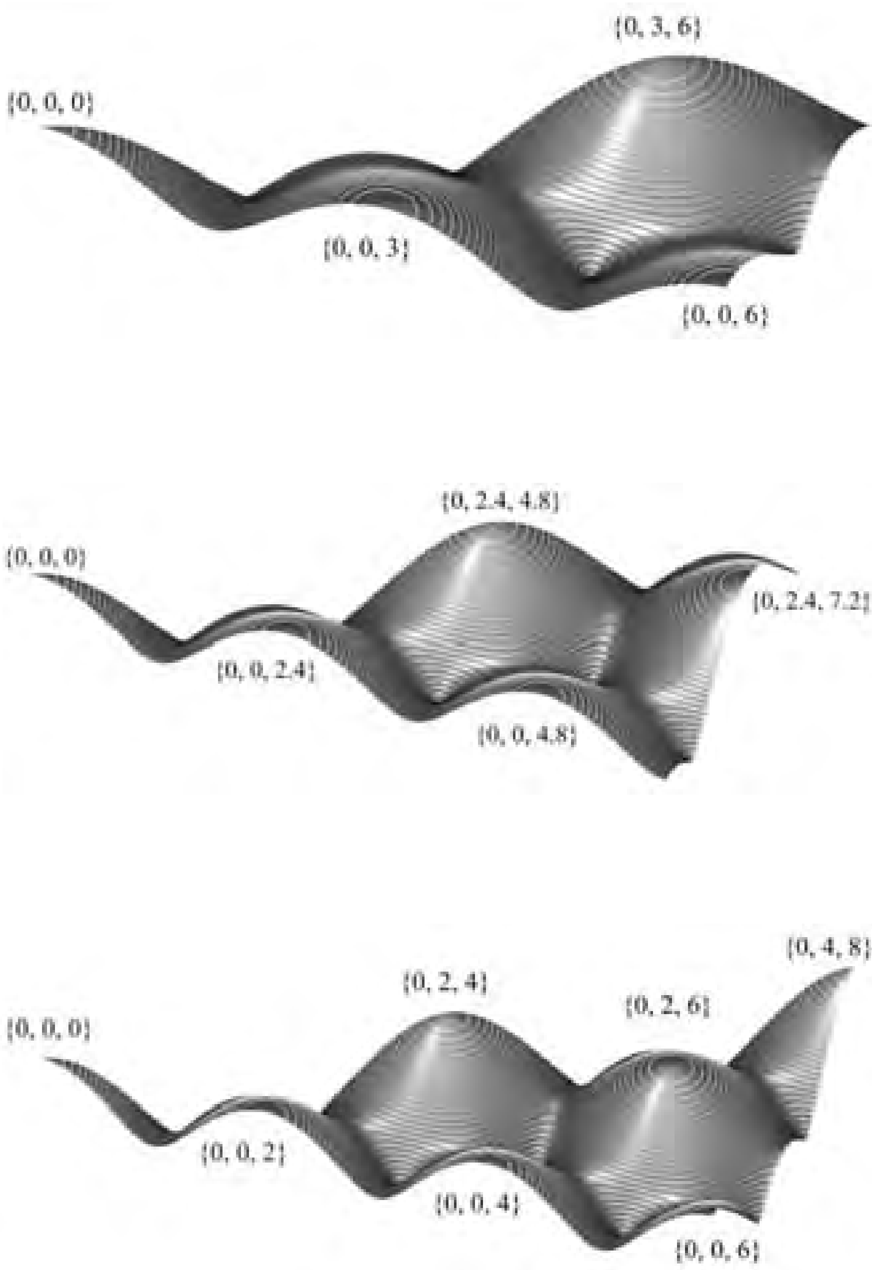

Figure 23. (cont.)

such that [0, *y*, *x*] is located at the point (*x*, *y*). For example, augmented triads are located at (8, 4), chromatic trichords at (2, 1), and tripled unisons at (0, 0). For reasons that lie beyond the scope of this article, we will also use an oblique coordinate system in which the *x*- and *y*-axes are in a 120° relation rather than 90°. Plotting pitch-class sets in this manner, all possible trichord set-classes lie within the shaded region of Figure 22. This is the *fundamental region* of trichord set-classes, since all possible three-note set-classes lie in this region and there are no duplications—no two points within the region represent the same set-class. Set-classes (including those with pitch-class duplications) from twelve-tone equal temperament are labeled only as a guide for the reader; all points within the region represent unique set-classes, and none are to be given conceptual priority.[16]

Next, we add a third dimension that corresponds to the magnitude of each set-class with respect to a given harmonic, shown in Figure 23. For harmonic *k*, there will be global maxima at those set-classes that are subsets of a single *k*-fold division of the octave. Thus, for the first harmonic, there is a single maximum at the tripled unison, $\{0, 0, 0\}_O$, since this is the only trichordal set-class that belongs to a onefold division of the octave. In fact, since a unison (with any number of pitch-class duplications) is trivially a member of any interval cycle (up to transposition), there will be a maximum at this set-class for all harmonics. For the second harmonic, there is an additional maximum at $\{0, 0, 6\}_O$; for the third harmonic, there are additional maxima at $\{0, 0, 4\}_O$ and $\{0, 4, 8\}_O$ and similarly for the remaining harmonics. Since the division of the octave into five parts does not yield a twelve-tone equal tempered set, the maxima for the fifth harmonic do not lie on familiar set-classes. Instead, there are maxima at $\{0, 0, 0\}_O$, $\{0, 0, 2.4\}_O$, $\{0, 0, 4.8\}_O$, $\{0, 2.4, 4.8\}_O$, and $\{0, 2.4, 7.2\}_O$.

Set-classes lying near maxima will have relatively high magnitudes for a given harmonic. For example, $\{0, 1, 6\}_O$ lies near $\{0, 0, 6\}_O$ and thus has a relatively high magnitude for the second harmonic, while the magnitude of $\{0, 2, 7\}_O$ with respect to the fifth harmonic is relatively high due to its proximity to the maximum at $\{0, 2.4, 7.2\}_O$.[17] Also note that the higher the harmonic, the greater the number of maxima and the steeper the slope away from these high points. Thus, the topology for the first harmonic consists of a gradual descending slope from a single maximum at the minimally even set-class, while the landscape of the sixth harmonic is studded with seven peaks and correspondingly steep gradients.

Analogous *n*-dimensional structures obtain for the spectra of *n*-note sets in continuous pitch-class space. Despite the difficulties in visualizing the cases for sets of more than three notes, the spaces are easily described: for the *k*th

16 For a more detailed explanation of this region and its derivation, see Callender 2004 and Callender, Quinn, and Tymoczko 2008.

17 Tymoczko 2008 makes the same point in a much more rigorous manner, demonstrating a strong inverse relationship between the magnitude of the *k*th harmonic in a chord's spectrum and the voice-leading distance from that chord to the nearest subset of a *k*-fold division of the octave. This means that the location of a chord in voice-leading space determines the chord's spectrum to a high degree of accuracy. Conversely, two chords will be close in Fourier space if the set of their voice-leading distances to the nearest subset of equal divisions of the octave are similar. Of course, two chords can be similarly distant from these subsets without necessarily being close to one another in voice-leading space. While Tymoczko's work exposes a deep connection between voice-leading and Fourier spaces, it also helps to explain the situation encountered in the introduction in which chords are close in one space but relatively more distant in the other.

harmonic, peaks are located at all set-classes of the form $\{a_i\}_{i=1}^{n}$, where $a_i \equiv 0$ mod $12/k$. The peaks (and troughs) form equal-spaced lattices corresponding to triangles for three-note chords, tetrahedra for four-note chords, and so forth.[18]

5. Properties of chord spectra

It is appropriate at this point to collect a number of general properties of chord spectra, focusing on those that are most relevant for continuous pitch and pitch-class spaces.

(1) Chords with identical interval content have identical spectra. This is evident by rewriting Equation 1 in the alternate form

$$|\mathcal{F}_P(z)| = \sqrt{\sum_{i,j} \cos(2\pi(p_i - p_j)z)}. \qquad (3)$$

In this alternate form the spectrum is defined in terms of the directed intervals from a chord to itself: $p_i - p_j$. Thus, transpositionally related, inversionally related, and Z-related sets have identical spectra.

(2) Multiplying a chord has the effect of dilating its spectrum; specifically, if $P = M_x(Q)$, then $\mathcal{F}_P(z) = \mathcal{F}_Q(xz)$ or, equivalently, $\mathcal{F}_P(z/x) = \mathcal{F}_Q(z)$. Quinn (2007) labels this property the *multiplication principle*. For example, since {0, 2, 4} is related to {0, 1, 2} by M_2, the spectrum of {0, 2, 4} with respect to frequency z is equal to the spectrum of {0, 1, 2} with respect to frequency $2z$.

Putting this property together with property (1) above, if the spectra of $M_x(P)$ and $M_x(Q)$ are equal at frequency z, then the spectra of P and Q will be equal at frequency xz. As an example of the usefulness of this property, consider the pitch-class sets $P_O = \{0, 1, 3\}_O$ and $Q_O = \{0, 1, 4\}_O$. Multiplying both sets by 2 yields $M_2(P_O) = \{0, 2, 6\}_O$ and $M_2(Q_O) = \{0, 2, 8\}_O$. Since these two sets, $M_2(P_O)$ and $M_2(Q_O)$, are related by inversion, their spectra are identical for all harmonics. By the multiplication principle, this implies that the spectra of P_O and Q_O are equal for all harmonics of the form $2k$ (where k is any integer), which can be verified by examining the even harmonics of the relevant spectra in Figure 21. More generally, given any pitch-class sets P_O and Q_O, if $M_x(P_O)$ and $M_x(Q_O)$ are transpositionally related, inversionally related, or Z-related, then the spectra of P_O and Q_O will be equal at all harmonics of the form xz. This property will play an important role in the discussion of Z-related sets in §7.

18 For more on higher-dimensional analogues of Figure 22 and on voice-leading spaces in general, see Callender, Quinn, and Tymoczko 2008.

(3) Given a pitch set P, if every member of P belongs to the same l-cycle, then the spectrum of P is periodic with a period of $1/l$:

$$P \subseteq T_y(\zeta_l) \Rightarrow |\mathcal{F}_P(z)| = |\mathcal{F}_P(z + 1/l)|$$

For example, the pitch set {0, 5, 15} belongs to a 5-cycle, so its spectrum, already shown in Figure 18c, has a period of 1/5. For another example, the pitch set {0, 3/4, 5/3} belongs to a (1/12)-cycle, since 1/12 is the greatest common divisor of the set's intervals. Thus, this set's spectrum has a period of 12.

The same principle applies for sets with irrational intervals, as long as the *ratio* of these intervals is rational. For example, the set $\{0, \sqrt{2}, 3\sqrt{2}\}$ belongs to a cycle of $\sqrt{2}$ semitones, $\{\ldots, 0, \sqrt{2}, 2\sqrt{2}, 3\sqrt{2}, \ldots\}$, so its spectrum has a period of $1/\sqrt{2}$. However, the set $\{0, \sqrt{2}, \pi\}$ cannot belong to an interval cycle since the ratios of its intervals, such as $\sqrt{2}/\pi$, are not rational. Thus, the spectrum of this set is aperiodic.

(4) The spectrum of a pitch-class set P_O is periodic if the spectrum of the corresponding pitch set P has a period of the form $k/12$, $k \in \mathbb{Z}$. Another way to state this is that if a pitch-class set belongs to some k-tone equal tempered system, then the spectrum of the pitch-class set is periodic with a period of $k/12$. For example, while the pitch set $P = \{0, 5, 15\}$ belongs to a 5-cycle and its spectrum has a periodicity of 1/5, the spectrum of the pitch-class set P_O cannot have a periodicity of 1/5, since a 5-cycle does not evenly divide the octave. However, P also belongs to a 1-cycle, which divides the octaves into the familiar twelve-tone equal tempered collection, so the spectrum of P_O has a period of 1. Similarly, $\{0, 3/4, 5/3\}_O$ belongs to a 144-fold division of the octave, so its spectrum has a period of 12. However, while the pitch set $\{0, \sqrt{2}, 3\sqrt{2}\}$ has a periodic spectrum, this set can never be embedded in a single equal tempered system. Thus, the spectrum of the pitch-class set $\{0, \sqrt{2}, 3\sqrt{2}\}_O$ is aperiodic.

(5) Chord spectra are symmetric about $z = 0$: $|\mathcal{F}_P(z)| = |\mathcal{F}_P(-z)|$. It is clear that an interval cycle generated by l, $\{\ldots, -l, 0, l, \ldots\}$, will be the same as a cycle generated by $-l$, $\{\ldots, l, 0, -l, \ldots\}$. Thus the magnitude of a given set with respect to $z = 1/l$ will be the same as its magnitude with respect to $-z = -1/l$. (This is also evident from Equation 3, since $\cos\theta = \cos -\theta$.)

(6) Putting properties (3), (4), and (5) together, if the spectrum of a pitch or pitch-class set is periodic with a period of $1/l$, then the entire spectrum can be generated by the magnitudes for $0 \leq z \leq 1/2l$. By properties (3) and (4), we can see that the entire spectrum can be generated by any interval spanning $1/l$. In particular, we can generate the entire spectrum by the magnitudes for

$-1/2l \leq z \leq 1/2l$. By property (5), we can generate the negative portion of this range from the positive frequencies. For a pitch-class set drawn from k-tone equal temperament, it is thus only necessary to consider those frequencies in the range $0 \leq z \leq k/(2 \cdot 12)$, which corresponds to harmonics 0 through $\lfloor k/2 \rfloor$. Conversely, for a pitch-class set that is not drawn from any k-tone equal temperament, it is necessary to consider the infinite number of harmonics, since the spectra of these sets are aperiodic.

6. Metrics for continuous harmonic space

In this section we will consider a number of metrics defined on chord spectra that may be interpreted as dissimilarity measures in a continuous harmonic space. However, the primary aim is not to put forth any metric in particular, but to demonstrate both how to construct metrics on chord spectra and that these metrics can correlate very strongly with similarity measures based on interval vectors. As discussed in Scott and Isaacson 1998, if two similarity measures correlate very well, "they probably will stand together or fall together when subjected to empirical tests." In this context, if a metric on chord spectra, SPECTRA, and a similarity measure on interval vectors, XSIM, correlate strongly with one another, we can think of SPECTRA as approximating XSIM in a continuous environment. The advantage is that, to the extent XSIM captures intuitions concerning harmonic similarity, SPECTRA applies these intuitions to all possible chords in all possible tuning systems (to within some level of accuracy). In order to facilitate such comparisons, we will begin with the spectra of pitch-class sets drawn from twelve-tone equal temperament, then move to continuous pitch-class space, and finish with metrics defined on continuous pitch space.

6.1 Pitch-class sets in twelve-tone equal temperament

Recall that the spectra of pitch-class sets in twelve-tone equal temperament are determined by the magnitudes of the first six harmonics of the octave (in addition to harmonic zero, which corresponds to cardinality). Let harmonics one through six define a six-dimensional Fourier transform space, $\mathbb{F}^6$. Using the magnitudes of these six harmonics, we can locate the spectrum of any equal tempered pitch-class set: $\mathcal{P} = (p_1, \ldots, p_6)$, where p_k is the magnitude of the kth harmonic. For example, the spectrum of $P = \{0, 4, 8\}$, which has a magnitude of 3 for the third and sixth harmonics and zero magnitude for the other harmonics, is located in $\mathbb{F}^6$ at $\mathcal{P} = (0, 0, 3, 0, 0, 3)$. A number of familiar similarity measures are defined on a different six-dimensional space defined by the interval vector, including Teitelbaum's (1965) s.i. (for "similarity index"), Isaacson's (1990) IcVSIM, Rogers's (1999) cos Θ, and Scott and Isaacson's (1998) ANGLE. We now consider two metrics on chord spectra that are analogous to these four similarity measures.

6.1.1 Euclidean metric in $\mathbb{F}^6$ A natural way to measure the distance between two spectra is to take the Euclidean distance between their corresponding points in $\mathbb{F}^6$. The Euclidean distance between the spectra of two pitch-class sets, *P* and *Q*, is

$$d(\mathcal{P},\mathcal{Q}) = \sqrt{\sum (p_i - q_i)^2}, \tag{4}$$

or simply the length of the differences of the two vectors,

$$d(\mathcal{P},\mathcal{Q}) = |\mathcal{P} - \mathcal{Q}|. \tag{5}$$

For example, the distance between the spectra of {0, 4, 8}, $\mathcal{P} = (0, 0, 3, 0, 0, 3)$, and {0, 3, 6}, $\mathcal{Q} = (1, 1, 1, 3, 1, 1)$, is $d(\mathcal{P},\mathcal{Q}) = \sqrt{1 + 1 + 4 + 9 + 1 + 4} = \sqrt{20} \approx 4.47$. This is quite similar to Teitelbaum's (1965) s.i., which measures the Euclidean distance between two points in the space defined by the interval vector. Both *d* and s.i. use the same metric; the difference is the space on which this metric is defined.

Richard Teitelbaum uses s.i. only to compare sets of the same size, since the metric tends to give unintuitive results when comparing sets of different size. (As Isaacson [1990] notes, s.i. judges all seven-note sets to be more similar to {0, 1, 2} than are all eight-note sets, which are all more similar to the chromatic trichord than are the nine-note sets.) As might be expected, *d* and s.i. correlate well; comparing pitch-class sets of equal cardinality (ranging from three to nine notes), the correlation coefficient for *d* and s.i. is $r = .87$.[19]

While this is a fairly high correlation, a much higher value results from a variant of *d*. Recall that the square of the magnitude spectrum is the power spectrum, which emphasizes relatively strong components and deemphasizes relatively weak ones. Since it is the strong components of a spectrum that most fully characterize the content of a set, it makes sense to define the Euclidean metric on the power rather than the magnitude spectrum. For instance, what matters most in distinguishing {0, 4, 8} and {0, 3, 6} is that the peak magnitudes of 3 occur in different harmonics; the fact that $\mathcal{Q}$ has magnitudes of 1 in harmonics where $\mathcal{P}$ has a magnitude of 0 is of relatively little consequence. Denoting Euclidean metrics defined on the power spectra by d_{pow}, we have

$$d_{\text{pow}}(\mathcal{P},\mathcal{Q}) = |\mathcal{P}^2 - \mathcal{Q}^2|. \tag{6}$$

The correlation between d_{pow} and s.i. is quite high: comparing sets of cardinality 3 through 9 (limited to sets of equal cardinality), $r = .96$, meaning that 92% of the variance in the metrics is related ($r^2 = .92$).

In order to overcome the limitation of s.i. with respect to cardinality, Isaacson (1990) proposes measuring the standard deviation of the differences

19 All correlations given in this article are statistically significant, with $p \leq .0001$.

of two interval vectors, a measure he termed IcVSIM. For two interval vectors, $a = (a_1, \ldots, a_6)$ and $b = (b_1, \ldots, b_6)$, let the differences of the two vectors be $c = (a_1 - b_1, \ldots, a_6 - b_6) = (c_1, \ldots, c_6)$. IcVSIM is defined as the standard deviation of c:

$$\sigma(c) = \sqrt{\tfrac{1}{6}\sum(c_i - \bar{c})^2}, \tag{7}$$

where $\bar{c}$ is the mean of the values of c. As Eric Isaacson notes, when comparing sets of equal cardinality, IcVSIM is the same as s.i. (up to a constant scaling factor: IcVSIM = s.i. $/\sqrt{6}$). Thus, it should not be surprising that d_{pow} and IcVSIM correlate strongly; comparing all sets of cardinalities 3 through 6 (not limited to those of equal cardinality), $r = .95$, while sets of cardinalities 3 through 9 yield a slightly lower correlation of $r = .89$.

While d_{pow} correlates strongly with IcVSIM, the former does not possess some of the more idiosyncratic features of the latter. In particular, sets for which the interval vectors differ by a constant are measured as maximally similar by IcVSIM and thus are located at the same point in the corresponding harmonic space. For instance, the interval vectors for $A = \{0, 3, 6\}$ and $B = \{0, 1, 3, 6, 7, 9\}$ are ⟨002001⟩ and ⟨224223⟩, respectively, differing by a constant value of 2. While these two sets are clearly similar in sound, it is hard to justify this equivalence, especially given that other, seemingly *more* similar pairs of sets, such as {0, 3, 6} and {0, 3, 6, 9}, are considered less similar. (Both Castrén [1994] and Buchler [1997] also question this equivalence.) In contrast, d_{pow} finds the two sets to be similar, but not equivalent: $d_{\text{pow}}(\mathcal{A},\mathcal{B}) \approx 5.48$, which is a smaller distance than that between 91% of all possible pairings.[20]

6.1.2 Angular distance in $\mathbb{F}^6$ Another natural way to measure the distance between two points in a space is to take the angle between the vectors extending from the origin to each of the points. This is a very common metric, one that is used instinctively whenever we compare the locations of objects in the sky. We cannot judge the actual distance between two stars by the naked eye, but we can approximate their angular distance fairly easily. Figure 24 gives examples of how angular distance applies to interval vectors. In Figure 24a, {0, 2, 4} and {0, 4, 8} are plotted on a Cartesian plane with the x-axis representing the number of interval class 2 and the y-axis representing the number of interval class 4. The angular distance between the two vectors is $\theta \approx 63.5°$ or $\approx .45$ radians. Smaller values for θ indicate interval vectors that are more similar with $\theta = 0$, indicating maximal similarity. For example, Figure 24b plots the interval vectors for {0, 3, 6} and {0, 3, 6, 9} on a Cartesian plane with the x-axis representing the number of

20 One idiosyncratic feature of measuring distance in $\mathbb{F}^6$ is that complements have identical magnitudes for each of these harmonics. Thus, complements are located at the same point in $\mathbb{F}^6$ and are judged to be maximally similar. One solution would be to incorporate the zeroth harmonic, which is equal to cardinality, so that chords of different cardinalities would not be considered equivalent.

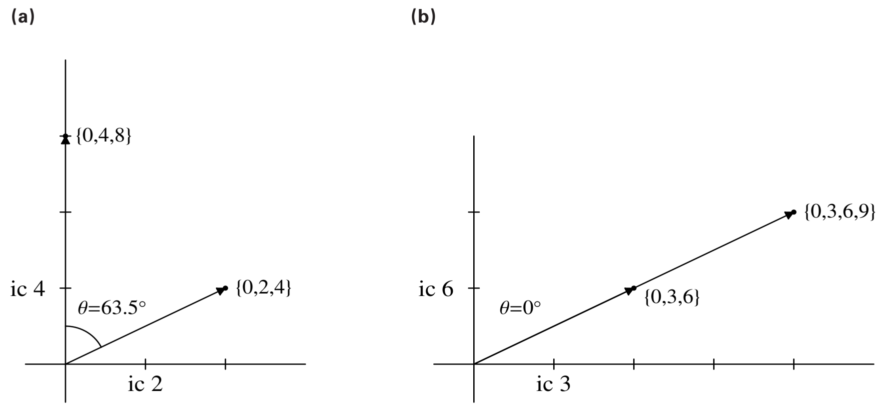

Figure 24. Angular distance between interval vectors for (a) {0,2,4} and {0,4,8} and (b) {0,3,6} and {0,3,6,9}

interval class 3 and the y-axis representing the number of interval class 6. Since the interval vectors of these pitch-class sets are exactly proportional, their vectors point in the same direction, so the angle between them is 0. Maximal dissimilarity, indicated by $\theta = 90°$ or $\pi/2$ radians, occurs precisely when two sets have no intervallic content in common.[21]

Scott and Isaacson's (1998) ANGLE measures the angle between vectors in a six-dimensional interval-vector space, while Rogers's (1999) cos Θ measures the cosine of this angle. The cosine of the angle θ between two vectors, $x = (x_1, \ldots, x_n)$ and $y = (y_1, \ldots, y_n)$, is given by

$$\cos(\theta) = \frac{x \cdot y}{|x||y|}, \quad (8)$$

where $|x|$ is the magnitude of x (or the length from the origin to x) and $x \cdot y$ is the dot product of x and y: $x \cdot y = \sum x_i y_i$. The values of cos Θ range from 0, indicating maximal dissimilarity, to 1, indicating maximal similarity. We can convert this measure to a metric in a harmonic space with the widely used cosine distance, $1 - \cos \Theta$, or by taking the arccos of Equation 8 yielding θ, which is the same as ANGLE. While each of these measures yields different sets of numbers, they are all measuring precisely the same thing—angular distance—so we can take ANGLE as the representative for all three.

Equation 8 can also be used to measure the angular distance between points in Fourier transform space. As with the Euclidean metrics of section 6.1.1, defining angular distance on power spectra yields a stronger correlation with ANGLE than the same metric defining on magnitude spectra.

21 Since the sets in Figure 24a contain only interval classes 2 and 4 and the sets in Figure 24b contain only interval classes 3 and 6, all nonzero dimensions are taken into account, and the graphs give complete information.

Designating this metric as d_θ, we have

$$d_\theta(\mathcal{P}^2,\mathcal{Q}^2) = \arccos\frac{\mathcal{P}^2 \cdot \mathcal{Q}^2}{|\mathcal{P}^2||\mathcal{Q}^2|}. \tag{9}$$

For example, if $\mathcal{P} = \{0, 3, 6\}$ and $\mathcal{Q} = \{0, 3, 6, 9\}$, then

$\mathcal{P}^2 = (1, 1, 1, 9, 1, 1)$ and $\mathcal{Q}^2 = (0, 0, 0, 16, 0, 0)$,

$\mathcal{P}^2 \cdot \mathcal{Q}^2 = 9 \cdot 16$,

$|\mathcal{P}^2| = (\sqrt{1 + 1 + 1 + 9^2 + 1 + 1} = \sqrt{86}$, and

$|\mathbb{Q}^2| = 16$.

Thus,

$$d_\theta(\mathcal{P}^2,\mathcal{Q}^2) = \arccos\frac{9 \cdot 16}{\sqrt{86} \cdot 16} \approx \arccos .97 \approx 13.95° \text{ or } \approx .24 \text{ radians},$$

indicating that according to d_θ the two sets are quite similar (closer than 99 percent of set-class pairs), but not equivalent as is the case for ANGLE.

The correlation between d_θ in $\mathbb{F}^6$ and ANGLE is not particularly high, with $r = .7$ for cardinalities 3 through 6 and $r = .52$ for cardinalities 3 through 9. A much higher correlation obtains between ANGLE and d_θ in $\mathbb{F}^{12}$, the twelve-dimensional space defined by the magnitudes for harmonics 1 through 12. The correlation between d_θ in $\mathbb{F}^{12}$ and ANGLE is $r = .94$ and $r = .93$ for cardinalities 3 through 6 and 3 through 9, respectively. d_θ in $\mathbb{F}^{12}$ also correlates strongly with ISIM2 (Scott and Isaacson 1998), $r = .88$, and RECREL (Castrén 1994), $.81 \leq r \leq .99$, for cardinalities 3 through 9. We will consider the reason for this significantly higher correlation in §6.3.

6.2 Pitch-class sets in continuous space

As we saw in §5, it is necessary to consider an infinite number of harmonics of the octave in order to fully capture irrational pitch-class sets such as $\{0, \sqrt{2}, \pi\}$ that do not belong to any equal tempered system. The spectrum of a given pitch-class set is a single point within an infinite dimensional Fourier space, $\mathbb{F}^\infty$, where each dimension corresponds to a unique harmonic of the octave. In theory, then, it should be necessary for a metric on the Fourier transform of continuous pitch-class space to be defined on $\mathbb{F}^\infty$, but there are reasons why we should not and, in fact, cannot do so, at least not if we wish for the results to be meaningful.

Let's consider three harmonics of four pitch-class sets: $P = \{0, 0.46, 0.95, 1.41\}$, $Q = \{0, 1.01, 6, 7.01\}$, $R = \{0, 3, 6, 9\}$, and a slight variation of R, $R' = \{0, 3.005, 6, 9.005\}$. The magnitudes of the first harmonic are $p_1 \approx 3.83$, $q_1 \approx 0$, $r_1 = 0$, and $r'_1 \approx 0$. This harmonic divides the four pitch-class sets into two categories: the unbalanced set, P, and the balanced sets Q, R, and R'. The magnitudes of the second harmonic are $p_2 \approx 3.41$, $q_2 \approx 3.41$, $r_2 = 0$, and $r'_2 \approx 0.01$. This harmonic divides the four pitch-class sets into two different categories: those that correlate strongly with a twofold division of the octave, P and Q, and those that do not, R and R'.

These two harmonics provide different levels of resolution with which to investigate these pitch-class sets and thus give different types of information about their harmonic content. While the first harmonic successfully distinguishes P from Q, its resolution is too low to discriminate between two different ways of obtaining balance within a pitch-class set: a maximally even distribution of all pitch classes in R, and a maximally even distribution of a pair of uneven dyads in Q. Likewise, while the second harmonic successfully distinguishes between Q and R, its resolution is too high to discriminate two different ways of correlating strongly with a tritone, by clustering around a single point, P, or by having an equal number of pitch classes clustered around two points a tritone apart, Q. Obviously, neither harmonic is successful at distinguishing R and R'. In order to make this type of fine distinction, it is necessary to use a very high frequency of resolution. Jumping (far) ahead to the 1200th harmonic, the magnitudes are $p_{1200} = q_{1200} = r_{1200} = 4$ and $r'_{1200} = 0$. The maximal magnitudes at $z = 1200$ for P, Q, and R result from the fact that in each case all four pitch classes belong to the equal division of the octave into 1200 parts, or a cent scale. Since R' contains two pitch classes from one cent scale and another two from the opposing cent scale, it has a magnitude of zero. (Opposed cent scales are separated by a transposition of one-half cent.) This harmonic successfully distinguishes R' from R but provides far too high a resolution to distinguish P, Q, and R or any of the unique pitch-class sets in 1200-tone equal temperament.

This example is purposely extreme in order to drive home the main point that different harmonics provide different levels of information about harmonic content. It is not that some harmonics are inherently superior to others, but that, depending on the application at hand, some harmonics will be more *useful* than others. If, for whatever reason, one wishes to investigate the *harmonic* difference between R and R', then it is necessary to use an extremely high harmonic and correspondingly high frequency of resolution.

The particular application at hand in this section is to measure the distance between pitch-class sets in a continuous harmonic space. If distances are taken between points in $\mathbb{F}^\infty$ and all harmonics are weighted equally, then the difference between r_{1200} and r'_{1200}, which is 4, will have a slightly greater effect on the measured distance than the difference between p_1 and r_1, which is approximately 3.83. However, the implicit assumption is that we wish for the measured distance to reflect our intuitions of perceived harmonic similarity. Since R and R' are perceptually identical, the large difference between r_{1200} and r'_{1200} is meaningless as a factor in measuring perceived distance. In other words, harmonic 1200 provides far too high a resolution for this particular application. At the other extreme, it is clear that harmonics 1 and 2 together do not provide enough information to serve as the sole basis for measurements of perceived distance—augmented triads and diminished seventh chords (which are certainly easily distinguishable sonorities!) both have zero magnitude for harmonics 1 and 2. The trick is to find a set of harmonics that are both necessary and sufficient to serve as a basis for a reasonable metric of harmonic distance in continuous pitch-class space.[22]

There are two common solutions to this problem. The first is to weight dimensions in a descending manner. For example, suppose we scale the *n*th harmonic by $1/n$: p_n/n. Then differences in the first harmonic will be 1200 times more important in determining distance than differences in the 1200th harmonic. For instance, according to this method of weighting dimensions, the weighted difference between r_{1200} and r'_{1200} is only $4/1200 \approx 0.0033$, which would have the desired negligible impact on the overall distance. Another solution is to simply take the first *n* harmonics as the basis for the Fourier space F^n, selecting *n* as appropriate. We will adopt the second approach, since it corresponds more closely with our approach to pitch-class sets in twelve-tone equal temperament.

In order to get a feel for metrics in continuous pitch-class space, we will compare distances between two pitch-class sets, *P* and *Q*, where *P* is some given trichord and *Q* varies among all possible trichords. Recall from §4.2 that the fundamental region of trichord set-classes is the triangular region graphed in Figure 22. To this two-dimensional region we add a third dimension that corresponds to *normalized similarity*, which varies from 0 to 1, indicating minimal and maximal similarity, respectively. The higher a point within the region is, the more similar its associated set-class, $/Q/$, is to the given comparison set-class, $/P/$. Conversely, lower points correspond to set-classes that are relatively dissimilar to $/P/$. For example, Figure 25a plots the normalized similarity between the augmented triad and all trichord set-classes using the metric d_{pow} defined on the first six harmonics of the octave, or the space $\mathbb{F}^6$. Note that the plot contains a global maximum at the augmented triad, indicating maximal similarity to {0, 4, 8}, and a global minimum at the tripled unison, indicating minimal similarity. Note also the local peaks at or near {0, 0, 4}, {0, 2, 6}, and {0, 2, 4}, and the slight ridge extending from {0, 0, 4} to both {0, 4, 8} and {0, 2, 4}. The center of this ridge corresponds to all trichordal set-classes that contain a major third—those set-classes of the form $/\{0, x, x + 4\}/$. Thus, we can see the influence that the major third, the defining interval of the augmented triad, exerts over the similarity contour for the continuous space of trichords.

However, not all trichords containing a major third are judged to be equally similar to the augmented triad. The similarity contour takes into

22 There is also a strictly mathematical problem. Consider the pitch-class sets $P = \{0, 4, 8\}$, $Q = \{0, 4 + \varepsilon, 8 + 3\varepsilon\}$, and $R = \{0, \varepsilon, 3\varepsilon\}$, where $\varepsilon = 1/p$, $p \in \mathbb{Z}$. Taking *p* to be arbitrarily large, and thus ε to be arbitrarily small, in $\mathbb{F}^{12p}$ $d_{\text{pow}}(\mathcal{P},\mathcal{Q}) = d_{\text{pow}}(\mathcal{P},\mathcal{R})$. At the very least, we should want a set that is indistinguishable from {0, 4, 8} (*Q*) to be closer to the maximally even set than a set that is indistinguishable from the maximally uneven set (*R*)! This is because all three sets belong to 12*p*-tone equal temperament and their respective interval functions are maximally dissimilar. See Equations 11 and 36 with their accompanying discussions for more details. (The same problem exists for the angular distance in $\mathbb{F}^{\infty}$.)

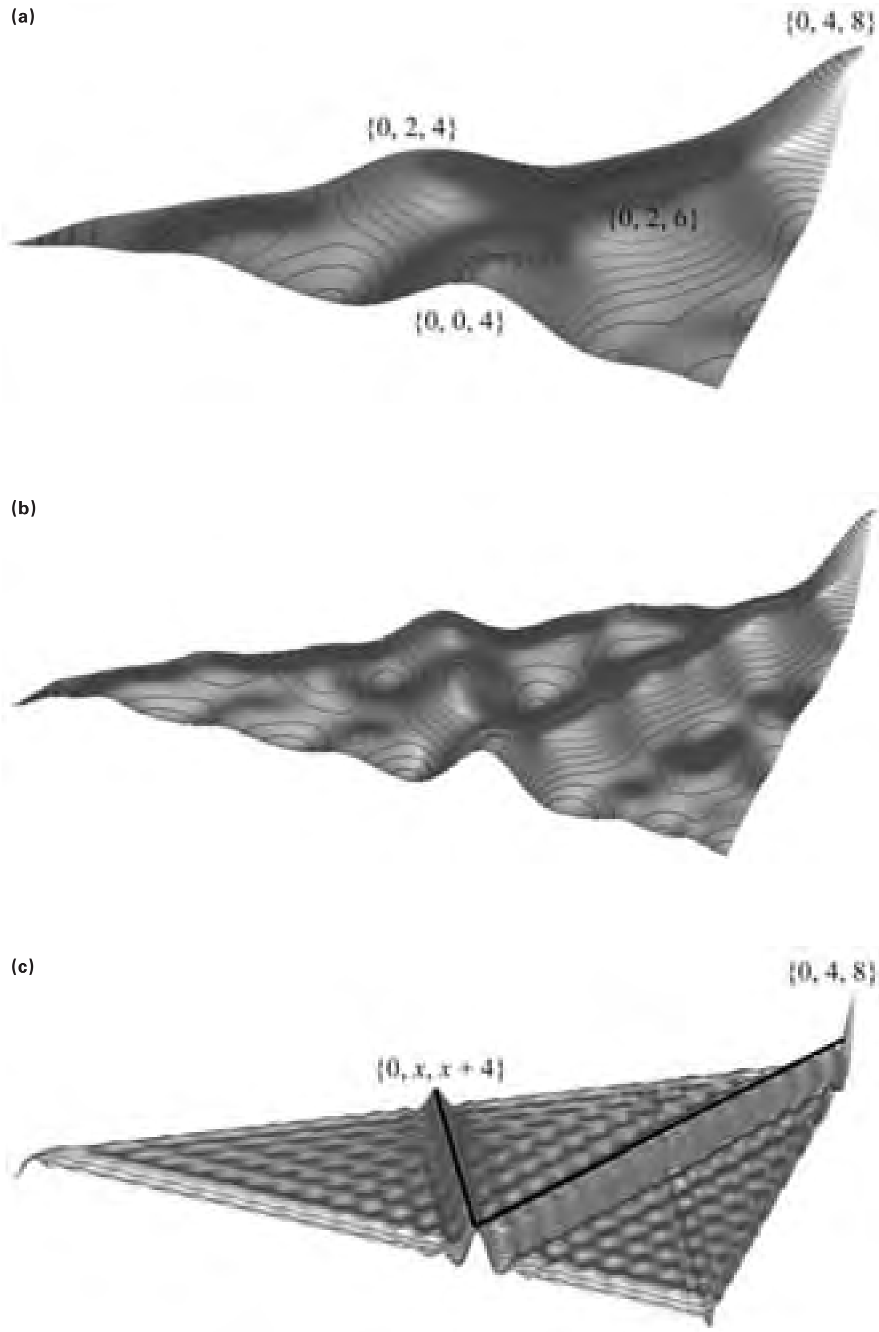

Figure 25. Plot of normalized similarity between {0, 4, 8} and all other trichordal set-classes based on d_{pow} in (a) $\mathbb{F}^6$, (b) $\mathbb{F}^{12}$, and (c) $\mathbb{F}^{48}$

account other important features of {0, 4, 8}, including its maximal fit with whole-tone and major-third cycles and its maximally even distribution of pitch classes. For instance, the peak at {0, 0, 4}, which also belongs to both whole-tone and major-third cycles, is higher than the other local peaks at or near {0, 2, 4} and {0, 2, 6}. Of the two smaller peaks, the one near {0, 2, 6} is higher, since this set-class has a more even distribution of pitch classes than {0, 2, 4}. Measures of harmonic similarity that are based on the Fourier transform thus capture many features of a given chord.

Also note that the similarity contour is smooth, such that set-classes lying near one another in the fundamental domain are judged to be roughly equal in their similarity to the augmented triad. This conforms to our expectation that similarity judgments should not be greatly affected by very small displacements of chord members. This feature is directly related to the number of harmonics taken into account. For example, Figure 25b is the contour map that obtains from distances based on the first *twelve* harmonics, $\mathbb{F}^{12}$, rather than just the first six. This contour is noticeably more jagged than Figure 25a, while the ridge corresponding to set-classes containing a major third has become less differentiated. These tendencies are even more pronounced in Figure 25c, which is based on distances in the Fourier space defined by the first 48 harmonics, $\mathbb{F}^{48}$. In the limit case, the normalized similarity based on d_{pow} in $\mathbb{F}^n$, where n goes to infinity, would yield a contour of similarity between {0, 4, 8} and all other trichords that consists of a single impulse at the augmented triad and an almost entirely undifferentiated, infinitesimally narrow ridge along set-classes containing a major third.[23] In other words, the higher the value of n (and the greater number of harmonics taken into consideration), the more distances based on chord spectra become like those based solely on interval content, losing their unique and advantageous features along the way.

Figure 26 is a contour map of similarities based on distances in $\mathbb{F}^6$ measured by d_θ rather than d_{pow}. A comparison of the graphs in 25a and 26 shows how similar the topographies resulting from d_{pow} and d_θ are. Both contain the same ridge and three local peaks discussed above. At least when limited to comparing sets of the same cardinalities, d_{pow} and d_θ (and their interval-vector based cousins, IcVSIM and ANGLE) correlate quite strongly. The sole exception occurs at the tripled unison, which is a global minimum according to the Euclidean metric but a rather strong peak according to angular distance.[24]

23 More precisely, the normalized similarity between {0, 4, 8} and all other trichords based on d_{pow} in $\mathbb{F}^\infty$ yields a contour map consisting of a single maximal spike of 1 at {0, 4, 8}, a ridge of .78 at all sets of the form {0, x, x + 4} with the exception of a slight spike of .89 at {0, 0, 4} and slight dips of .74 and .7 at {0, 2, 6} and {0, 2, 4}, and a broad plain of .54 for all other sets except for cliffs of .48 and .4 at sets of the form {0, x, 2x} and {0, 0, x} and a downward spike to 0 at {0, 0, 0}.

24 The spectrum of the maximally uneven set {0, . . . , 0} contains maximal magnitude in every harmonic. Thus, the vector associated with this spectrum points directly in the middle of that part of $\mathbb{F}^n$ defined by nonnegative values for each harmonic. This means that the maximum angular distance between the maximally uneven set and any other pitch-class set is 45° or π/4 radians.

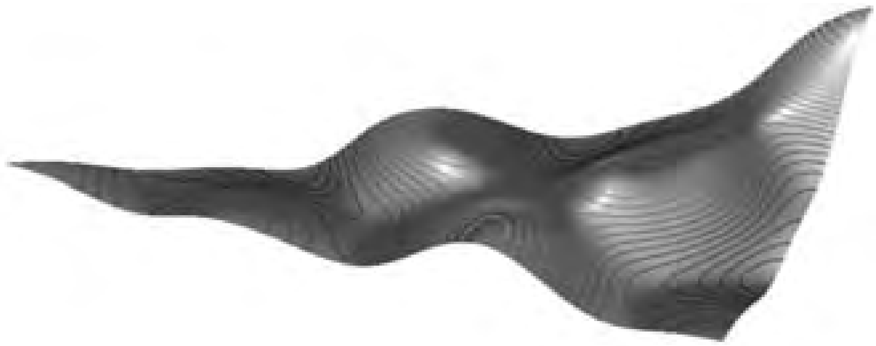

Figure 26. Plot of normalized similarity between {0, 4, 8} and other trichordal set-classes based on d_θ in $\mathbb{F}^6$

As discussed above, part of the strength and attraction of Fourier spaces is the flexibility to consider only those harmonics that are relevant to a specific context. For instance, most readers have been conditioned to hear anything belonging to twelve-tone equal temperament as "in tune" and (most) everything else as more or less "out of tune." This suggests that, all other things being equal, twelve-tone equal tempered chords will be perceived as more similar to one another than the neighboring "out of tune" chords. Since the twelfth harmonic measures the extent to which a set belongs to a twelvefold division of the octave, we can use this harmonic to help model this phenomenon, measuring distances in a seven-dimensional Fourier space defined by harmonics 1–6 and 12. (Note that this space does *not* contain harmonics 7–11.) Figure 27 once again shows the normalized similarity between trichords and the augmented triad, but here distances are measured with respect to this seven-dimensional Fourier space. The resulting contour map contains local maxima at each set-class associated with twelve-tone equal tempered chords, including those with pitch-class duplications. (Thus, the peaks in Figure 27 correspond to the points labeled in Figure 22.)

6.3 More on Euclidean and angular distance on chord spectra and interval content

We have seen two metrics on chord spectra that correlate strongly with metrics on interval content. It was suggested that this strong correlation allows us to apply the intuitions about harmonic similarity that are implicit in prominent similarity measures to continuous pitch-class spaces. But what are these intuitions? What exactly is the relationship between these metrics on chord spectra and interval content?

To understand the connection between these metrics, we must consider Lewin's (1987) interval function (what mathematicians call the Patterson function). For a given set P, the interval function, which we will write as ΔP, is the set of all directed intervals between members of P: $\Delta P = \{p_i - p_j \bmod 12\}$,

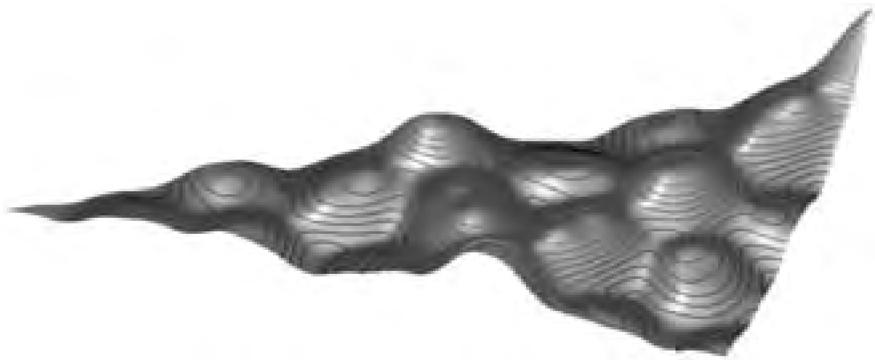

Figure 27. Plot of normalized similarity between {0, 4, 8} and all other trichordal set-classes based on d_{pow} in the Fourier space defined by the magnitudes of harmonics 1–6 and 12

for all $p_i, p_j \in P$.[25] For example, in the set $Q = \{0, 1, 3\}$ there are three ways of moving by 0 (from any member of Q to itself) and one way of moving by intervals of 1, 2, 3, 9, 10, and 11 semitones. Thus, the interval function of Q is $\Delta Q = \{0, 0, 0, 1, 2, 3, 9, 10, 11\}$. We could also write this set as a vector where the ith entry (beginning with $i = 0$) corresponds to the multiplicity of i in the interval function; for example, $\Delta Q = (3, 1, 1, 1, 0, 0, 0, 0, 0, 1, 1, 1)$.

The relevance of the interval function is that, when measuring distance between twelve-tone equal tempered sets in $\mathbb{F}^{12}$ rather than $\mathbb{F}^6$, d_θ and d_{pow} can each be rewritten in terms of the interval function. (See §8.) For example, we can rewrite Equation 9 as

$$d_\theta(\mathcal{P},\mathcal{Q}) = \arccos\frac{\Delta P \cdot \Delta Q}{|\Delta P||\Delta Q|}. \quad (10)$$

In other words, measuring angular distance in $\mathbb{F}^{12}$ is identical to measuring angular distance in the twelve-dimensional space defined by the interval function.[26] Returning to a previous example, if $P = \{0, 3, 6\}$ and $Q = \{0, 3, 6, 9\}$, then

$$\Delta P = (3, 0, 0, 2, 0, 0, 2, 0, 0, 2, 0, 0),$$
$$\Delta Q = (4, 0, 0, 4, 0, 0, 4, 0, 0, 4, 0, 0),$$
$$\Delta P \cdot \Delta Q = 3 \cdot 4 + 2 \cdot 4 + 2 \cdot 4 + 2 \cdot 4 = 36,$$
$$|\Delta P| = \sqrt{3^2 + 3 \cdot 2^2} = \sqrt{21}, \text{and } |\Delta Q| = \sqrt{4 \cdot 4^2} = 8.$$

Thus, measuring angular distance in twelve-dimensional interval function space, we have

$$\text{ANGLE}_{12}(P,Q) = \arccos\frac{36}{\sqrt{21 \cdot 8}} \approx \arccos .982 \approx 10.9° \text{ or} \approx .19 \text{ radians.}$$

25 Lewin's interval function is actually defined on two sets, P and Q, and is the set of all directed intervals from members of P to members of Q: $\text{IFUNC}(P,Q) = \{p - q \bmod 12\}$ for all $p \in P$ and $q \in Q$. We will write ΔP as an abbreviation for IFUNC(P,P).

26 More generally, for n-tone equal tempered sets, d_θ in $\mathbb{F}^n$ is identical to measuring angular distance in n-dimensional interval-function space.

Likewise, measuring angular distance in $\mathbb{F}^{12}$, we have

$$|\mathcal{F}_{\Delta P}| = (1, 1, 1, 9, 1, 1, 1, 9, 1, 1, 1, 9),$$
$$|\mathcal{F}_{\Delta Q}| = (0, 0, 0, 16, 0, 0, 0, 16, 0, 0, 0, 16),$$
$$|\mathcal{F}_{\Delta P}| \cdot |\mathcal{F}_{\Delta Q}| = 3 \cdot 9 \cdot 16,$$
$$\|\mathcal{F}_{\Delta P}\| = \sqrt{3 \cdot 9^2 + 9} = 16\sqrt{7}, \text{ and}$$
$$\|\mathcal{F}_{\Delta Q}\| = \sqrt{3 \cdot 16^2} = 16\sqrt{3},$$

which yields the same result:

$$d_\theta(|\mathcal{F}_{\Delta P}|,|\mathcal{F}_{\Delta P}|) = \arccos\frac{3 \cdot 9 \cdot 16}{6\sqrt{7} \cdot 16\sqrt{3}} \approx 10.9° \text{ or } \approx .19 \text{ radians.}$$

It is this identity between angular distance measured on power spectra and the twelve-dimensional interval function that explains the high correlation noted above between d_θ, measured in $\mathbb{F}^{12}$, and ANGLE, based on the six-dimensional interval vector, because ANGLE correlates strongly with angular distance on interval functions.

Similarly, Euclidean-based metrics on chord spectra can be expressed in terms of the interval function. Consider d_{pow} defined in $\mathbb{F}^{12}$ rather than $\mathbb{F}^6$. As measured in the twelve-dimensional space of power spectra, when applied to twelve-tone equal tempered sets Equation 6 may be simplified as

$$d_{\text{pow}}(\mathcal{P},\mathcal{Q}) = \sqrt{12}\,|\Delta P - \Delta Q|. \tag{11}$$

That is, Euclidean distance between power spectra in $\mathbb{F}^{12}$ is identical (up to a constant scaling factor) to Euclidean distance between interval functions.[27]

Defined in $\mathbb{F}^{12}$, d_{pow} correlates very strongly with Teitelbaum's s.i (or the Euclidean distance between interval vectors), with $r = .997$. (Indeed, the two measures are nearly identical.) To the extent that this way of measuring distance captures something of our sense of harmonic similarity, d_{pow} (in $\mathbb{F}^{12}$) allows us to apply this intuition in continuous pitch-class space.

The correlation between s.i. and IcVSIM is fairly weak ($r = .41$ for cardinalities 3 through 9), so the strong correlation between d_{pow} in $\mathbb{F}^{12}$ and s.i. cannot be the reason for the strong correlation between d_{pow} in $\mathbb{F}^6$ and IcVSIM.[28] The interpretation of d_{pow} in either $\mathbb{F}^6$ or $\mathbb{F}^{12}$ is very similar, with the former differing from the latter in giving additional weight to the degree to which the relevant sets embed within a single whole-tone collection. Let O_X and E_X be the number of odd and even intervals in the interval function of X. The quantity $O_X E_X$ represents the degree to which the members of X belong to opposing whole-tone collections.[29] Equation 11 for d_{pow} in $\mathbb{F}^{12}$ can be modified in the

27 Again more generally, for n-tone equal tempered sets, d_{pow} in $\mathbb{F}^n$ is a scaled version of s.i. defined on the n-dimensional interval function rather than the $\lfloor n/2 \rfloor$-dimensional interval vector.

28 The correlation between d_{pow} in $\mathbb{F}^6$ and a variant of IcVSIM defined on the interval function rather than the interval vector is similarly strong ($r = .93$ for cardinalities 3 through 9).

29 For example, consider three tetrachords: in X all four members belong to the same whole-tone collection, in Y there are three members in one collection and one member in the other, and in Z there are two members in each collection. Thus, we have $O_X E_X = 0 \times 16 = 0$, $O_Y E_Y = 6 \times 10 = 60$, and $O_Z E_Z = 8 = 8 = 64$, with the values increasing as the pitch classes are distributed more evenly between the two whole-tone collections.

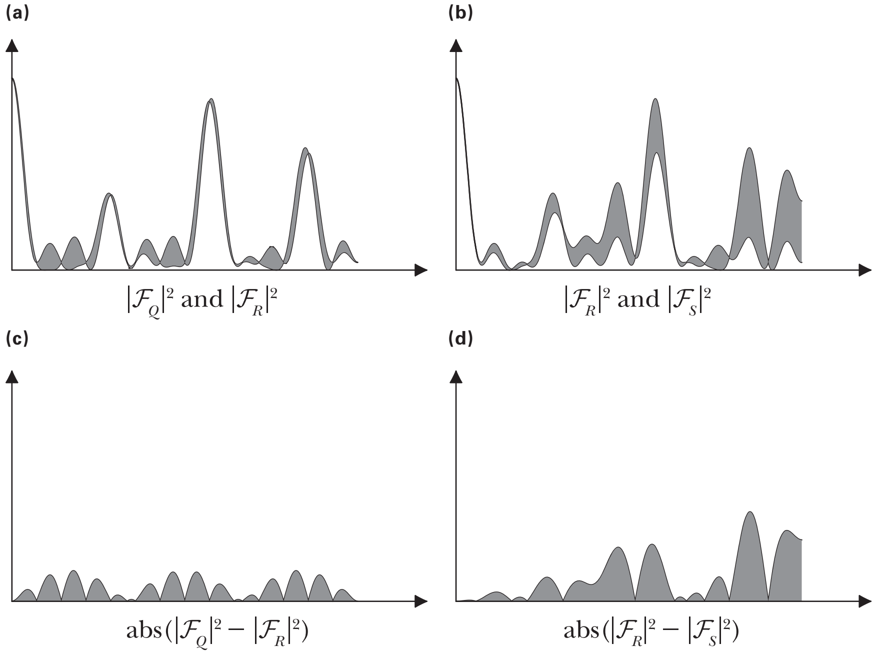

Figure 28. Defining distance between pitch spectra in terms of area

following way to yield the equation for d_{pow} in $\mathbb{F}^6$:

$$d_{\text{pow}}(\mathcal{P},\mathcal{Q}) = \sqrt{6}\,|\Delta P - \Delta Q| + \sqrt{2(O_P E_P + O_Q E_Q - (O_P E_Q + O_Q E_P))}. \quad (12)$$

Essentially, the expression $O_P E_P + O_Q E_Q - (O_P E_Q + O_Q E_P)$ represents the degree to which P and Q are similarly divided between whole-tone collections. If the two sets are divided between whole-tone collections to the same degree—that is, $O_P = O_Q$ and $E_P = E_Q$—then the value of the expression is 0. The maximum value for the expression occurs when one set is drawn from a single whole-tone collection while the other is as evenly divided between whole-tone collections as possible. While the difference between s.i. and IcVSIM is due to the manner in which each measure deals with comparisons between sets of different cardinalities, it is suggestive that the corresponding metrics on chord spectra, d_{pow} in $\mathbb{F}^{12}$ and $\mathbb{F}^6$, differ not in terms of cardinality directly but in terms of whole-tone embedding.

6.4 Pitch sets

Defining distance between spectra is much more complicated for pitch sets than for pitch-class sets, because the spectra of pitch sets are continuous. To better understand the difficulties, let's return to the three sets from the introduction: Q = {G3, D4, F♯4, A4, E4}, R = {G3, B3, D4, A4, E4}, and S = {G3, B3, E♭4, A4, E4}. Musical intuition suggests that the spectra of R should be more similar to that of Q than to that of S, implying that Q and R should lie closer together in harmonic space than R and S. But how can we measure this distance? Since the spectra of pitch sets are continuous, we cannot use discrete summation as with pitch-class sets. Instead, we must rely on the analogous geometric concept of area. The spectra of Q and R are superimposed in Figure 28a, while the spectra of R and S are superimposed in Figure 28b. In both cases, the region bounded by the two spectra and vertical boundaries at $x = 0$ and $x = 1/2$ is shaded. Equivalently, we can take the absolute value of the difference of two spectra and shade the region bounded by the resulting curve, the x-axis, and the same vertical boundaries, as shown in Figure 28, (c) and (d). The shaded regions in (a) and (c) are equal, as are the regions in (b) and (d). If we were to measure the area of the shaded regions, we would observe that the area of the region in (c) is less than that in (d). We can now make our intuitions of distance more concrete: one way to measure the distance between two spectra is to measure the *area under the curve* of their absolute difference, or squared difference, or any other means of measuring the difference between the two. Since the shaded region in (d) is larger than that in (c), the harmonic distance between R and S is greater than that between Q and R.

The "area under the curve" of a function, $g(x)$, between vertical boundaries at $x = a$ and $x = b$ is denoted by the *definite integral*

$$\int_a^b g(x)\,dx. \tag{13}$$

Readers are encouraged not to be put off by the mathematical symbolism; the definite integral is simply the familiar notion of area from geometry applied to the more general setting of continuous functions. Fortunately, it is not necessary to understand calculus to *apply* Equation 13. There are many elegant means of estimating definite integrals that can be easily implemented in a computer program. (See §8 for a simple algorithm to evaluate Equation 13.)

Recall that the squared Euclidean distance between pitch-class sets is the sum of the squared differences between their discrete spectra. The analogous squared Euclidean distance between pitch sets is the *area* of the squared differences between their continuous spectra. As with pitch-class sets, it is necessary to decide which frequencies to include in the measurement of distance and whether to scale the differences in the magnitudes of these frequencies:

$$d_{\text{pow}}(\mathcal{P},\mathcal{Q}) = \sqrt{\int_a^b f(z)\left(\mathcal{P}(z)^2 - \mathcal{Q}(z)^2\right)^2 dz}, \tag{14}$$

where the frequencies included are $a \leq z \leq b$, and $f(z)$ is the scaling factor.

As we did in Figure 25 for pitch-class sets, we can get a feel for metrics in continuous pitch space by comparing distances between a pitch set and all possible three-note pitch sets. The fundamental region for trichordal pitch sets is similar to that for pitch-class sets given in Figure 21, but without the vertical boundary on the right. That is, the space of all possible three-note pitch sets is an infinite wedge.

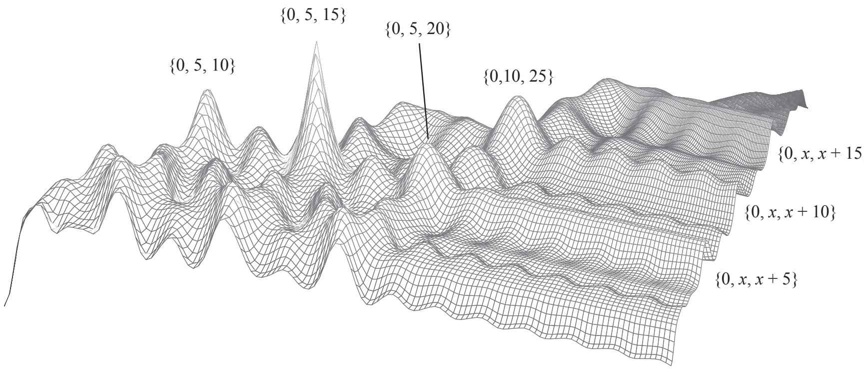

Figure 29. Plot of normalized similarity to {0, 5, 15} in a region of continuous space for three-note pitch set-classes based on d_{pow} in the Fourier space defined by all frequencies from 0 to 1/2 with no scaling factor.

Figure 29 plots the normalized similarity based on the Euclidian distance between {0, 5, 15} and all possible three-note pitch sets within a portion of this infinite region. As before, similarity is indicated by height. For this example, we assume a range of frequencies from 0 to 1/2 and no scaling factor, $f(z) = 1$. This range of frequencies covers all interval cycles generated by a whole-tone or larger interval and is analogous to defining a distance measure on pitch-class sets on the first six harmonics. (Recall that the sixth harmonic has a frequency of 6/12 = 1/2.) There are four primary peaks in the resulting similarity contour graph: a global maximum at {0, 5, 15} and three local maxima at {0, 5, 10}, {0, 5, 20}, and {0, 10, 25}. In addition, there are noticeable ridges corresponding to pitch sets that contain at least one interval of 5, 10, or 15 semitones. This is the same phenomenon as discussed with reference to Figure 25 and demonstrates the influence of interval content on Fourier-based harmonic spaces. The reason for the local maxima at {0, 5, 10}, {0, 5, 20}, and {0, 10, 25} is that these pitch set-classes lie at the intersection of two of these ridges. The global maximum obviously lies at the intersection of all three.

Figure 29 is fairly jagged, with pronounced ridges, rather than smooth like Figure 25a, due to the lack of a scaling factor. All frequencies from 0 to 1/2 have equal weighting. However, the range of frequencies from 1/4 to 1/2 corresponds to all interval cycles between major seconds and major thirds, while the range from 0 to 1/4 corresponds to all interval cycles larger than (or equal to) a major third. These smaller interval cycles thus may be thought to have a disproportionate effect on the measured distance between pitch sets in a manner similar to the inclusion of higher frequencies in Figure 25, (b) and (c). By using a scaling factor, we can weight frequencies such that the effect of smaller interval cycles is lessened relative to larger interval cycles. One such scaling factor is $f(z) = (1 - 2z)^2$, which is equal to 1 at $z = 0$

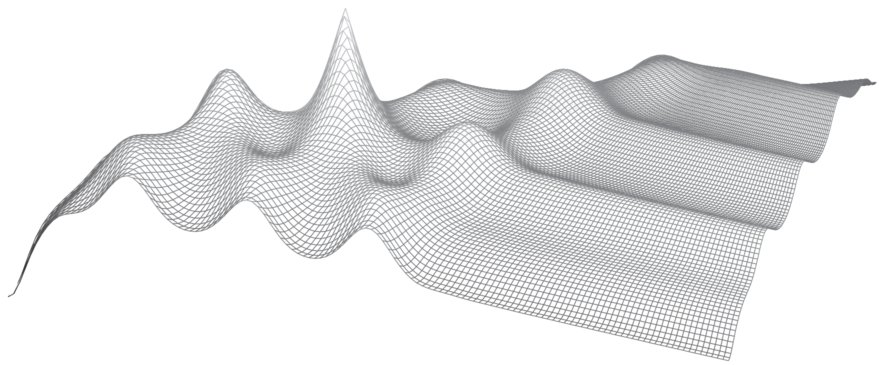

Figure 30. Plot of normalized similarity to {0, 5, 15} in a region of continuous space for three-note pitch set-classes based on d_{pow} in the Fourier space defined by all frequencies from $0 \le z \le 1/2$ scaled by the function $f(z) = (1 - 2z)^2$.

and decreases to 0 at $z = 1/2$. Figure 30 plots the similarity graph using this scaling factor. The basic profile of the graph remains the same, but the contour has been smoothed, and the ridges corresponding to interval content are less pronounced.

For pitch sets drawn from twelve-tone equal temperament, Equation 14 may be simplified in a manner nearly identical to the simplification of Euclidean distance for pitch-class sets. Assuming a range of frequencies from 0 to 1/2 and no scaling factor, for twelve-tone equal tempered pitch sets Equation 14 is equivalent to

$$d_{\text{pow}}(\mathcal{P},\mathcal{Q}) = \frac{|\Delta P - \Delta Q|}{\sqrt{2}}. \tag{15}$$

For example, applying Equation 15 to the three-pitch set from the introduction yields a distance of $\sqrt{2}$ between Q and R and a distance of $\sqrt{10}$ between R and S, which is in agreement with our intuitions about the relative harmonic distances between these chords. Furthermore, given the closeness of Q and R in harmonic space, we would expect Q and S to be about as distant as R and S. This is indeed the case, with the measure yielding a distance of $\sqrt{12}$ between Q and S.[30]

We can also construct measures of angular distance between pitch spectra that are analogous to ANGLE and d_θ. As with Euclidean distance between pitch spectra, it is necessary to replace the discrete summations of Equation 9 with definite integrals:

$$d_\theta(\mathcal{P}^2,\mathcal{Q}^2) = \arccos \frac{\int_a^b |\mathcal{F}_{P_O}|^2 \cdot |\mathcal{F}_{Q_O}|^2}{\|\mathcal{F}_{P_O}\|_a^b \, \|\mathcal{F}_{Q_O}\|_a^b}, \tag{16}$$

30 There are many choices of a, b, and $f(z)$ such that Equation 14 simplifies to Equation 15 (up to a constant scaling factor). See §8.

Table 1. Average correlations between d_{pow}, d_θ, s.i., and ANGLE for vertical aggregates spanning less than six octaves

	d_{pow}	d_θ	*s.i.*	*ANGLE*
d_{pow}	1			
d_θ	.97	1		
s.i.	.97	.99	1	
ANGLE	.96	.87	.88	1

where, for any function $g(x)$, $\|g(x)\|_a^b$ is the "magnitude" of the function $g(x)$ between a and b, defined as $\|g(x)\|_a^b = \sqrt{\int_a^b |g(x)|}$. (We could also add a scaling factor if desired.) As with angular distance for pitch-class power spectra, we begin with pointwise multiplication of the two pitch power spectra. Instead of summing this result, we find the area of the multiplication of the power spectra and divide this area by the maximum possible value. The more similar the two power spectra, the larger the area of their multiplication and the closer the value of the numerator to the denominator; that is, the more similar the two power spectra, the closer the fraction is to 1 and the arccos is to 0.[31]

As with Euclidean distance, angular distance may be simplified for pitch sets drawn from twelve-tone equal temperament. Indeed, assuming a frequency range of 0 to 1/2 and no scaling factor, Equation 16 simplifies to Equation 10 for twelve-tone equal tempered pitch sets. Applying this simplified equation to the three pitch sets from the introduction yields angular distances that are in relative agreement with d_{pow}: .257 between Q and R, .604 between R and S, and .666 between Q and S.

While s.i., IcVSIM, and ANGLE were originally defined on pitch-class sets, we can easily apply them to pitch sets, keeping in mind that interval vectors for pitch sets are infinite-valued vectors of undirected pitch intervals rather than finite-valued vectors of interval classes.[32] This will allow us to compare distance measures for pitch sets that are based on chord spectra with those based on interval vectors as we did with pitch-class sets in §6.1.1 and §6.1.2. It is difficult to compare measures of distance between pitch sets due to the enormous (indeed, infinite) space of all possible pitch sets. However, based on comparisons made within smaller portions of this space, both d_{pow} and d_θ correlate strongly with s.i. and ANGLE. Consider the space of all possible vertical aggregates up to transposition and inversion. (A vertical aggregate is a chord in which every pitch class is represented exactly once.) Furthermore, limit the

31 A program to calculate values of d_{pow} and d_θ for specific values of *a*, *b*, and *f*(*z*) is available at the author's Web site, http://mailer.fsu.edu/~ccallend.

32 One measure specifically proposed for pitch sets is Morris's PM (1995), where pitch sets are compared on the basis of the intersection of their interval vectors and pitch content. Focusing on the intersection of interval vectors and cast as a measure of distance rather than similarity, Morris's approach is to sum the absolute differences of two interval vectors: PM(P,Q) = $\sum_i p_i - q_i$, where p_i and q_i are the multiplicities of interval *i* in *P* and *Q*, respectively. Buchler's (1997) pSATSIM adapts PM to compare pitch sets of different cardinality.

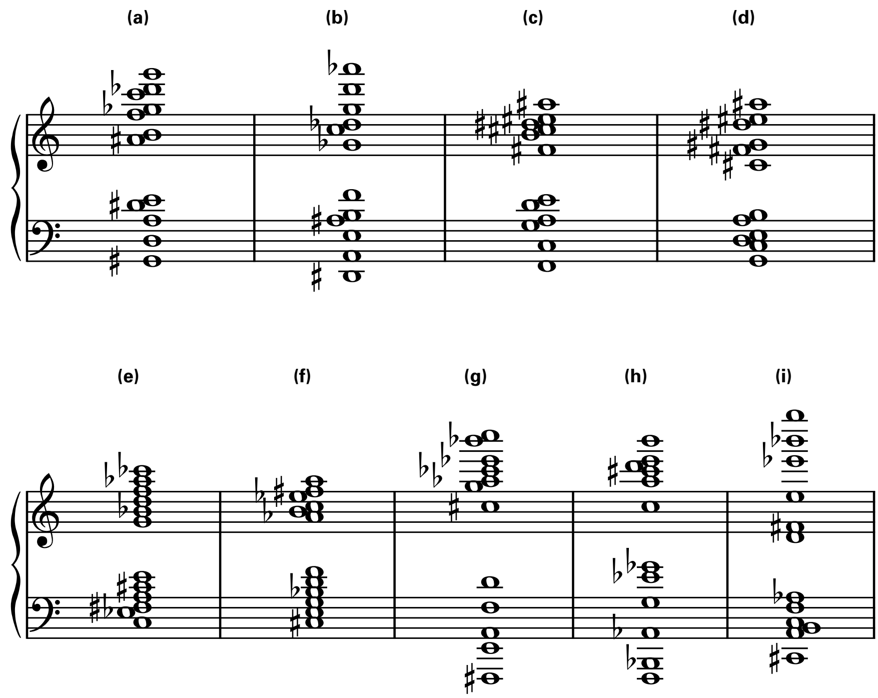

Figure 31. Vertical aggregates in Lutosławski (a–f) and Carter (g–i): (a) Mi-parti; (b) Jeux vénitiens; (c and d) Second Symphony; (e and f) Paroles tissées; (g) Night Fantasies; (h) In Sleep, In Thunder; (i) Fourth String Quartet

space to aggregates that span less than six octaves. Taking random samples of 100 of these pitch sets and comparing all pairs of sets yields strong average correlations among all four measures, shown in Table 1.

Figure 31 shows several vertical aggregates from Witold Lutosławski and Elliott Carter (Stucky 1981, Carter 2002). Since each chord contains all twelve pitch classes, the differentiation between these twelve-note chords is due entirely to their spacing. The two composers preferred to work with very different types of vertical aggregates. Carter often worked with all-interval twelve-note chords in which adjacent pitches (ordered by pitch height) form every interval from a minor second (or ninth) to a major seventh. In contrast, Lutosławski typically used vertical aggregates in which the intervals between adjacent pitches are very limited:

> The fewer kinds of interval that are in your twelve-note chord, the more characteristic is the quality, the physiognomy of the chord, the result. That's one rule because if you add other intervals, it [i.e., the chord] loses its character gradually up to the moment when it's absolutely gray—without any quality at all—if you use all possible intervals in one single chord or in one single melody. It loses the physiognomy. It loses the character. (Rust 1995, 215)

Table 2. Normalized interval function magnitudes for the chords in Figure 31

chord	*normalized magnitude*
E	.80
F	.79
A	.77
B	.76
C	.75
D	.74
I	.35
G	.33
H	.31

We know from Equation 15 that, for equal tempered pitch sets, comparing chords on the basis of the Euclidean distance between their power spectra is essentially the same as comparisons based on the Euclidean distance between their interval functions. Thus, we can compare the Lutosławski and Carter aggregates using Equation 15, which is much easier to calculate than Equation 14. We can also get a more intuitive understanding of the Euclidean distance between power spectra or interval functions by expanding and rewriting Equation 15 in terms of interval function magnitudes:

$$d_{\text{pow}}(\mathcal{P},\mathcal{Q}) = \sqrt{\frac{|\Delta P|^2 + |\Delta Q|^2}{2} - \Delta P \cdot \Delta Q}\,. \tag{17}$$

For pitch set P the magnitude of its interval function is $|\Delta P| = \sqrt{\sum_i p_i^2}$ where p_i is the multiplicity of i in ΔP. This magnitude can be interpreted as a measure of a chord's "distinctiveness." The more a chord's interval content is concentrated in a few intervals, the higher its interval function magnitude and the more characteristic its harmonic quality. Table 2 ranks the chords of Figure 31 by their normalized interval function magnitudes. (In order to make these numbers easier to interpret, the magnitudes are normalized so that the minimum possible value of $\sqrt{276}$ is mapped to 0 and the maximum possible value of 34 is mapped to 1.) As a point of reference, the average normalized interval function magnitude for vertical aggregates spanning less than six octaves is .45. The chords of Figure 31 cluster into two groups: Lutosławski's chords have significantly higher interval function magnitudes than average, while Carter's all-interval aggregates have noticeably lower magnitudes.

In Equation 17, the expression $\sqrt{\frac{|\Delta P|^2 + |\Delta Q|^2}{2}}$ gives the maximum distance between the power spectra of P and Q as the average of their squared interval function magnitudes. The larger this average, the greater the *potential* distance between the chords. From a geometric perspective, the interval function magnitude indicates how far a chord is from the center of a harmonic space based on the interval function or chord spectra. Chords with high interval function magnitudes lie near the extremes of the space and have the potential to be quite distant from one another, while chords with relatively low interval

Table 3a. Euclidean distances between chords from Lutostawski in Figure 31

	A	*B*	*C*	*D*	*E*	*F*
A	0					
B	4.5	0				
C	19.2	19.3	0			
D	19.5	19.6	2.0	0		
E	19.7	19.0	20.5	20.3	0	
F	18.8	18.1	19.8	19.7	4.0	0

Table 3b. Euclidean distances between chords from Carter in Figure 31

	G	*H*	*I*
G	0		
H	7.5	0	
I	8.1	7.9	0

Table 3c. Euclidean distances between chords from Lutostawski and Carter in Figure 31

	A	*B*	*C*	*D*	*E*	*F*
G	14.7	14.7	14.4	14.4	16.6	15.8
H	16.4	16.4	14.8	14.4	15.3	15.1
I	16.1	16.2	14.6	14.4	16.0	15.6

function magnitudes lie near the center of the space. This geometric perspective helps to make Lutosławski's intuitive comments more precise. Due to their location at the extremes of harmonic space, Lutosławski's twelve-note chords have the potential to be highly dissimilar and readily differentiable (due to their "characteristic quality"). In contrast, since all-interval aggregates are clustered near the center of the space, they are likely to be relatively close to one another, yielding chords that are relatively similar (or "gray") and not readily differentiable.

The expression $\Delta P \cdot \Delta Q$ in Equation 17 essentially measures the degree to which P and Q are "similarly distinctive." The greater the extent to which the vectors of their respective interval functions point in the same direction, the closer this dot product is to their average interval function magnitudes, and the smaller the distance between the two chords. The dot product vanishes when the chords have no intervallic content in common, yielding a distance between the two chords that is equal to their average squared interval function magnitudes.

Table 3 shows the pairwise Euclidean distances (according to Equation 15) between chords in Figure 31 in three groups: (a) distances between Lutosławski's aggregates, (b) distances between Carter's all-interval aggregates, and (c) distances between the two different types of chords.[33] As a point of reference, the average Euclidean distance between random vertical aggregates spanning less than six octaves is approximately 10. The extremes of distances are in group (a). Chords that are built on the same limited pitch intervals are very close together: tritones and perfect fifths or minor seconds in chords *A* and *B*, major seconds and perfect fourths or fifths in chords *C* and *D*, and major and minor thirds (with the single exception of a minor second) in chords *E* and *F*. All other distances in group (a) are quite high, confirming our intuitions based on their extreme positions within harmonic space. The all-interval aggregates are somewhat closer than average to themselves (b) and more distant than average to Lutosławski's aggregates (c), but the distances in both of these groups fit within the extremes of group (a). Indeed, the average distance between all-interval aggregates (not limited to the three in Figure 31) and random aggregates of less than six octaves is approximately 10, the same as the average distance between random aggregates. This lends support to Lutosławski's characterization of these chords as "gray" and less "characteristic."

I close this section with an example that approximates continuous pitch space through the extensive use of microtones. The opening chord of Gérard Grisey's *Partiels* is drawn from the overtone series of an E1 fundamental. (See Figure 32a.) Over the course of the first three minutes, the piece moves progressively from pure tones to noise, periodic to aperiodic rhythms, and harmonic to inharmonic chords. The latter is achieved through a series of discrete harmonic changes given in Figure 32b, where open note heads indicate pitches belonging to an E1 harmonic series and filled note heads indicate inharmonic pitches. In order to follow this progression from harmonic to inharmonic sonorities, we can measure the degree to which each chord resembles a harmonic series over an E1 fundamental. Specifically, we can measure the Fourier magnitude of this fundamental frequency for the spectrum of each chord, keeping the following in mind:

(1) Since the relevant cycles are frequency rather than interval cycles, chord members should be represented by their absolute frequencies rather than log-frequency pitches. For example, chord 1 contains E1 $\approx$ 41.2 Hz, B3 $\approx$ 123.47 Hz, and so forth.

(2) While shifting of interval cycles corresponds to transposition, the shifting of frequency cycles corresponds to frequency shifting—a technique common in the music of spectral composers such as Grisey and Murail (Rose 1996). For example, a two-note chord with frequencies of 180 Hz and 380 Hz belongs to a 200-Hz cycle that has been shifted down by 20 Hz, resulting in a phase shift by 1/10

33 These distances correlate well (r = .94) with Morris's PM (and therefore Buchler's pSATSIM). Michael Buchler shared unpublished work with me that relates large chords in Lutosławski and had previously noticed similar relationships between large chords in Lutosławski and Carter using pSATSIM.

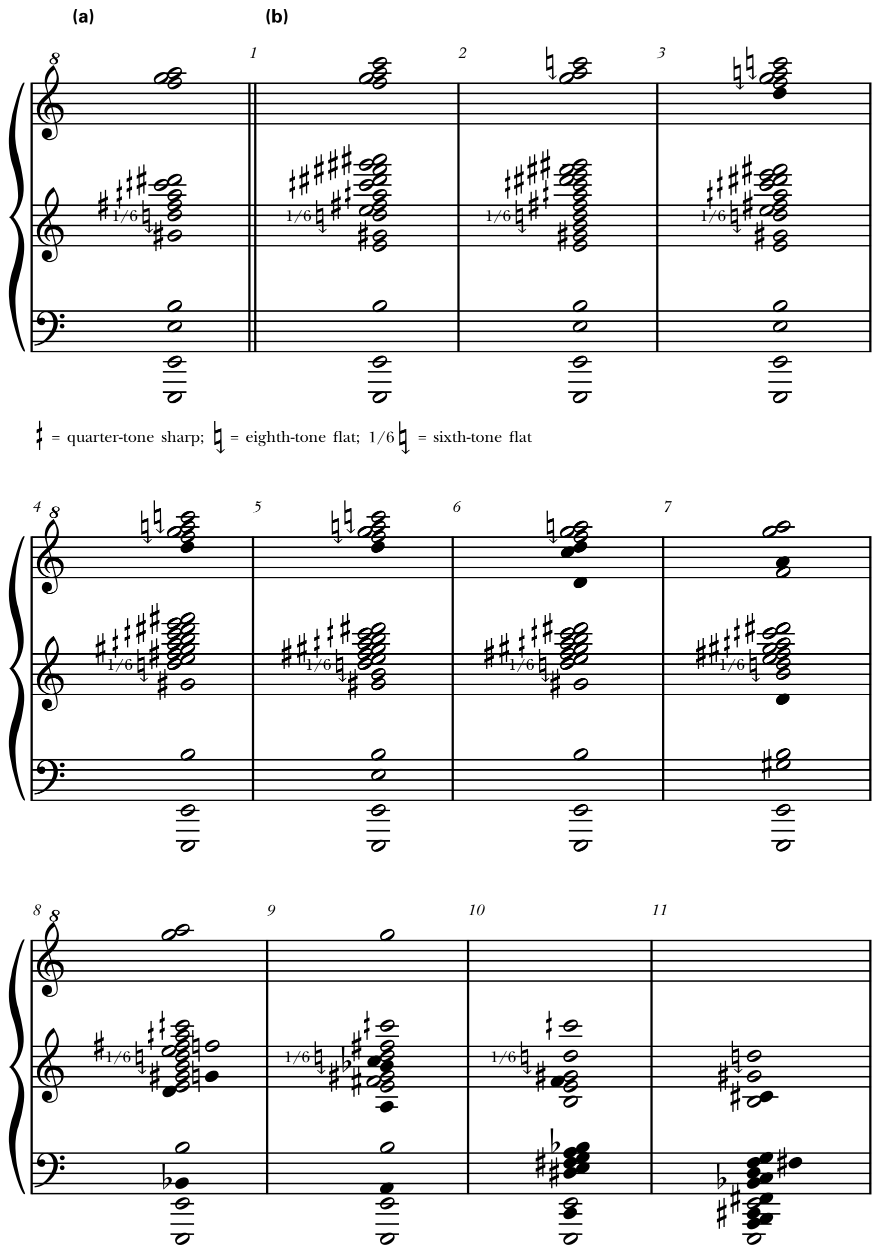

Figure 32. (a) Opening chord of Grisey's *Partiels*. (b) Gradual change from harmonic to inharmonic chords in the opening section of *Partiels*. (Adapted from Rose 1996.)

Table 4. Normalized magnitudes with respect to an E1 fundamental in the opening of *Partiels*

chord	*normalized magnitude* $z \approx 41.2$ *Hz*
1	.87
2	.91
3	.90
4	.87
5	.88
6	.82
7	.77
8	.69
9	.38
10	.08
11	.01

of a cycle. Instead of indicating a *harmonic* relation with a fundamental of 200 Hz, this maximal correlation with a phase-shifted cycle indicates that the chord is a *frequency shifted* version of a harmonic relation. Since we are only interested in the degree to which the sonorities in the opening of *Partiels* correspond to a harmonic series on E1, we are only concerned with cycles that have a phase of 0. This means that it is sufficient to evaluate the cosine portion of Equation 1:

$$\mathcal{F}_{yP}(z) = \sum_{f \in P} \cos(2\pi f z).$$

(3) Lower pitches are more important in establishing a harmonic spectrum than higher ones. For example, in a chord that is otherwise based on the harmonic series, a pitch lying halfway between the 16th and 17th partials is of far less consequence than a pitch halfway between the second and third partials. One way to model this phenomenon is by weighting each frequency, f, according to its ratio with the E1 fundamental, f_0:

$$\sum_{f \in P} \tfrac{f_0}{f} \cos(2\pi f z).$$

(4) Because the number of notes and the cumulative weightings of frequencies vary from chord to chord in the opening of *Partiels*, the magnitude of each chord (with respect to the E1 fundamental) will be normalized to aid in their comparison:

$$\frac{\sum_{f \in P} \tfrac{f_0}{f} \cos(2\pi f z)}{\sum_{f \in P} \tfrac{f_0}{f}}. \qquad (18)$$

Evaluating Equation 18 with z equal to the frequency of E1 ($\approx$41.2 Hz) yields the normalized magnitude of a chord with respect to this fundamental. The higher this magnitude (where the values range from 0 to 1), the greater the correlation between a chord and a harmonic series on E1. The normalized magnitudes with respect to an E1 fundamental for the opening chords in *Partiels* are given in Table 4. After the opening five chords, which are quite harmonic, the sonorities become progressively less harmonic. It is also clear from Table 4 that the *rate* of change from harmonic to inharmonic sonorities increases dramatically toward the end of the section. This is consistent with Grisey's stated preference for accelerating, rather than linear, rates of change (Grisey 1987).

7. The Z-relation problem

Consider the following question: Given a pitch-class set-class, */P/*, in a given tempered space, is *P* Z-related to any other set-classes and, if so, which?[34] At present, answers to this question must in general resort to an exhaustive search of a given pitch-class universe by finding all set-classes and comparing their interval vectors. While this brute-force approach will invariably answer the question posed (eventually! the number of set-classes to check grows exponentially as the number of pitch classes per octave increases), the result does not yield an *understanding* of the deeper structural properties Z-related sets share. The lack of a general explanation of Z-related sets and our inability to *predict* these relations is referred to as the "Z-relation problem."[35] In continuous spaces the problem is worse, since it is impossible to conduct an exhaustive search of an infinite number of sets. For example, how would one determine whether $\{0, \sqrt{2}, \sqrt{2} + 3, 6\}$ possesses a nontrivial Z-related partner? In order to show that continuous spaces not only make the problem more acute but also point the way toward a possible solution, we consider the following three cases, each of which generates an infinite number of Z-relations.

Case 1

Let P and Q be the pitch-class sets $\{0, 6\}_O$ and $\{0, 3\}_O$, respectively, and S_x be the union of P and the transposition of Q by x, $S_x = P \cup T_x(Q)$. We wish to show that S_x and S_{-x} are Z-related by demonstrating that their spectra are equivalent. The spectra of P and Q are $|\mathcal{F}_P| = (2, 0, 2, 0, \ldots)$ and $|\mathcal{F}_Q| = (2, \sqrt{2}, 0, \sqrt{2}, 2, \sqrt{2}, 0, \sqrt{2}, \ldots)$. Where the spectrum for either P or Q is zero, the spectrum of

34 In the common understanding of the Z-relation, two sets, P and Q, are Z-related if $\Delta P = \Delta Q$ *and* P and Q are not related by transposition or inversion. In this section, it will be advantageous to use a more relaxed definition of the Z-relation that includes sets related by transposition and inversion.

35 One attempt at a more general explanation of the Z-relation is the generalized hexachord theorem (Lewin 1987). This theorem states that in 2*k*-tone equal temperament, pitch-class sets of *k* elements are Z-related to their complements. While important, this result is of limited value in continuous spaces. A more recent partial explanation of the Z-relation is Stephen Soderberg's (1995) "Q" inversion, a dual inversion that preserves interval content and thus can be used to generate Z-related sets. Mathematicians refer to Z-related sets as *homometric* sets, which have been studied by Bullough (1961) and Rosenblatt (1984), among others. A more thorough investigation of the Z-relation problem is the focus of ongoing research between Rachel Hall and myself (Callender and Hall 2007) that draws upon existing mathematical work on homometric (or Z-related) sets.

$S_{\pm x}$ is equal to the nonzero magnitude. Thus, where k is odd, $|\mathcal{F}_{S\pm x}(k)| = \sqrt{2}$; where $k \equiv 2 \bmod 4$, $|\mathcal{F}_{S\pm x}(k)| = 2$. The spectrum of $S_{\pm x}$ thus takes the form

$$|\mathcal{F}_{S\pm x}| = (c_0, \sqrt{2}, 2, \sqrt{2}, c_1, \sqrt{2}, 2, \sqrt{2}, \ldots),$$

where c_j is the magnitude of harmonic $4j$. Multiplying S_x and S_{-x} by 4, we have $M_4(S_x) \equiv \{0, 0, 4x, 4x\}$ and $M_4(S_{-x}) \equiv \{0, 0, -4x, -4x\}$, where multiplication is taken mod 12. Since $M_4(S_x)$ and $M_4(S_{-x})$ are related by inversion, their spectra are identical for all harmonics. Recalling the multiplication principle from property (2) of §5, this implies that S_x and S_{-x} have the same spectra for all harmonics of the form $4j$. Thus, for $k \equiv 0 \bmod 4$, $|\mathcal{F}_{S_x}(k)| = |\mathcal{F}_{S_{-x}}(k)|$, which demonstrates that the spectra of S_x and S_{-x} are equal and that these sets are Z-related. More generally, every member of $/S_x/$ is Z-related to every member of $/S_{-x}/$.[36] For a familiar specific example, if $x = 1$, then S_x and S_{-x} are the tetrachords {0, 1, 4, 6} and {11, 0, 2, 6}, which are members of the well-known Z-related set-classes 4-Z15 and 4-Z29, respectively. For a less familiar example, if $x = 1\frac{1}{2}$, then $S_x = \{0, 1\frac{1}{2}, 4\frac{1}{2}, 6\}$ and $S_{-x} = \{10\frac{1}{2}, 0, 1\frac{1}{2}, 6\}$, which are Z-related sets in eight-tone equal temperament.[37] (A chromatic step in eight-tone equal temperament is equal to $1\frac{1}{2}$ semitones.)

Case 2

Let $P = \{0, 4, 8\}_O$, $Q = \{0, 2\}_O$, and (as before) $S_x = P \cup T_x(Q)$. Again, S_x and S_{-x} are Z-related, as we can demonstrate by inspecting the spectra of P and $T_x(Q)$ independently and combined. Taken independently, the spectra of P and Q are

$$|\mathcal{F}_P| = (3, 0, 0, 3, 0, 0, \ldots)$$

and

$$|\mathcal{F}_Q| = (2, \sqrt{3}, 1, 0, 1, \sqrt{3}, 2, \sqrt{3}, 1, 0, 1, \sqrt{3}, \ldots).$$

Notice that one of the two spectra is always zero except for harmonics divisible by 6. Thus, we know that the spectrum of $S_{\pm x}$ takes the form

$$|\mathcal{F}_{S\pm x}| = (c_0, \sqrt{3}, 1, 3, 1, \sqrt{3}, c1, \sqrt{3}, 1, 0, 1, \sqrt{3}, \ldots),$$

where c_j is the magnitude of harmonic $6j$. Similar to case 1, multiplying S_x and S_{-x} by 6 yields $M_6(S_x) \equiv \{0, 0, 0, 6x, 6x, 6x\}$ and $M_6(S_{-x}) \equiv \{0, 0, 0, -6x, -6x, -6x\}$. Since these two sets are related by inversion, by the multiplication principle S_x and S_{-x} we have the same spectra for all harmonics of the form $6j$, which demonstrates that $/S_x/$ and $/S_{-x}/$ are Z-related. A familiar example is given by $x = 1$, which yields the Z-related set-classes SC 5-Z17 = /{0, 1, 3, 4, 8}/ and SC 5-Z37 = /{11, 0, 1, 4, 8}/.

From the first two cases, we generalize that if P is a division of the octave into n equal parts and Q is a dyad that divides the equal divisions of P in half,

36 Three of these Z-relations are trivial: if $x = 0$ or 6, then $S_x = S_{-x}$; if $x = 3$, then $S_x = T_6(S_{-x})$.

37 See Soderberg 1995 for how this particular Z-related pair relates to his Q-inversion.

then we can combine P and Q to form an infinite number of Z-related sets. Specifically, let $P = \{0, 12/n, 2 \cdot 12/n, \ldots, (n-1) \cdot 12/n\}$, $Q = \{0, 6/n\}$, and $S_x = P \cup T_x(Q)$. Then $/S_x/$ and $/S_{-x}/$ are Z-related. The only examples of this type of Z-relations in twelve-tone equal temperament are the tetrachords and pentachords already mentioned. However, we can see this principle at work among Z-related hexachords in sixteen-tone equal temperament, where the chromatic step is equal to 3/4 semitones. Temporarily switching to a unit of pitch that is equal to 1/16 of an octave, the fourfold equal division of the octave that includes pitch class 0 is $P = \{0, 4, 8, 12\}_O$. Dividing this equal division in half yields $Q = \{0, 2\}_O$. If $x = 1$, then $S_x = \{0, 1, 3, 4, 8, 12\}_O$ and $S_{-x} = \{15, 0, 1, 4, 8, 12\}_O$. Both $/S_x/$ and $/S_{-x}/$ have the (eight-valued) interval vector ⟨21242022⟩.

Case 3

We begin with a slight variation of case 1, in which pitch class 0 of Q is doubled: let $P = \{0, 6\}_O$, $Q = \{0, 0, 3\}_O$, and S_x be defined as always. Doubling pitch class 0 alters the spectrum of Q such that it never equals zero: $|\mathcal{F}_Q| = (3, \sqrt{5}, 1, \sqrt{5}, 3, \sqrt{5}, 1, \sqrt{5}, \ldots)$. In order to determine the spectrum of $S_{\pm x}$, we must consider those harmonics where the spectrum of P is not zero, which is precisely for all even harmonics. Multiplying S_x and S_{-x} by 2, we have $M_2(S_x) \equiv \{0, 0, 2x, 2x, 2x \pm 6\}$ and $M_2(S_{-x}) \equiv \{0, 0, -2x, -2x, -2x \pm 6\}$. Since $M_2(S_x)$ and $M_2(S_{-x})$ are related by inversion, then by the multiplication principle we know that the spectra of S_x and S_{-x} are equal for all even harmonics. Thus, once again, the spectra of S_x and S_{-x} are equal, and the two sets are Z-related.

Alternatively, instead of simply doubling pitch class 0, we could *split* it into a pair of pitch classes arranged symmetrically about 0: $Q = \{-y, y, 3\}_O$. Once again, we know that the odd harmonics of S_x and S_{-x} are equal, since the spectrum of P vanishes for these harmonics. Likewise, we know that the even harmonics of S_x and S_{-x} are equal, since $M_2(S_x)$ and $M_2(S_{-x})$ are related by inversion. Thus, by the now-familiar chain of reasoning, S_x and S_{-x} are Z-related. Familiar specific Z-related set-classes arise for $x = 2$ and $y = 1$, which yields SC 5-Z12 = /{0, 1, 3, 5, 6}/ and SC 5-Z36 = /{9, 11, 0, 1, 6}/, and for $x = 1$ and $y = 2$, which yields SC 5-Z18 = /{11, 0, 3, 4, 6}/ and SC 5-Z38 = /{9, 0, 1, 2, 6}/.

The same principle holds if we split pitch class 3 instead of pitch class 0: $Q = \{0, 3 - y, 3 + y\}_O$. In fact, the same principle holds if *both* pitch classes 0 and 3 are split. Leaving the details for the interested reader, if $Q = \{-y, y, 3 - z, 3 + z\}_O$, then S_x and S_{-x} are Z-related. Three familiar Z-related hexachord pairs belong to the infinite number of such pairs generated in this way:

(1) for $x = 1/2$, $y = 3/2$, and $z = 1/2$, $S_{\pm x}$ yields 6-Z10/39;
(2) for $x = 1/2$, $y = 5/2$, and $z = 3/2$, $S_{\pm x}$ yields 6-Z24/46; and
(3) for $x = 3/2$, $y = 5/2$, and $z = 5/2$, $S_{\pm x}$ yields 6-Z26/48.

The general principles derived from these and similar cases are sufficient to generate many Z-related sets, including Z-triples, Z-quadruples, and so forth, but such an undertaking is well beyond the scope of the present article.[38]

8. Mathematical appendix

8.1 Fourier transform of pitch sets

In this section I summarize the mathematics underlying much of the foregoing. First, we represent the pitch p by the *delta function* at p, $\delta(x-p)$, where

$$\delta(x-p) = 0 \text{ for } x \neq p, \text{ and } \int_{-\infty}^{\infty}\delta(x-p)dx = 1. \qquad (19)$$

Intuitively, $\delta(x-p)$ can be thought of as an impulse at $x = p$ representing one unit of pitch p.

Next, we represent the pitch set P by its *characteristic function*

$$\chi_P(x) = \sum_{p\in P}\delta(x-p). \qquad (20)$$

That is, the characteristic function of P is a set of unit impulses located at the members of P.

The Fourier transform of the function $f(x)$ is given by

$$\mathcal{F}(z) = \int_{-\infty}^{\infty} f(x)e^{2\pi izx}\,dx. \qquad (21)$$

More specifically for our purposes, since

$$\int_{-\infty}^{\infty}\delta(x)f(x)\,dx = f(0), \qquad (22)$$

the Fourier transform of the delta function at p is

$$\int_{-\infty}^{\infty}\delta(x-p)e^{2\pi izx}\,dx = e^{2\pi izp}. \qquad (23)$$

Since the Fourier transform is linear, the transform of the sum of two functions is the same as the sum of the transforms of the individual functions. That is, if $\mathcal{F}(z)$ and $\mathcal{G}(z)$ are the Fourier transforms of $f(x)$ and $g(x)$, then the transform of $f(x) + g(x)$ is $\mathcal{F}(z) + \mathcal{G}(z)$. Thus, the Fourier transform of the characteristic function of P is the sum of the transforms of its component pitches:

$$\mathcal{F}_P(z) = \sum_{p\in P} e^{2\pi ipz}. \qquad (24)$$

Using Euler's identity, we substitute $\cos\theta + i\sin\theta$ for $e^{i\theta}$ in Equation 24 to yield the Fourier transform of the characteristic function of P expressed in terms of sines and cosines:

$$\mathcal{F}_P(z) = \sum_{p\in P}\cos(2\pi pz) + i\sin(2\pi pz). \qquad (25)$$

The magnitude of any complex number, $a + ib$, is the square root of the number multiplied by its complex conjugate, $a - ib$:

$$|a+ib| = \sqrt{(a+ib)(a-ib)} = \sqrt{(a^2+b^2)}\ . \qquad (26)$$

38 For examples of Z-triples and Z-quadruples, see Lewin (1982) and Soderberg (1995).

Thus, the magnitude, or *modulus*, of $\mathcal{F}_P(z)$ is

$$|\mathcal{F}_P(z)| = \sqrt{\left(\sum_{p\in P}\cos(2\pi pz)\right)^2 + \left(\sum_{p\in P}\sin(2\pi pz)\right)^2}\,. \tag{27}$$

8.2 Simplifying d_{pow} and d_θ

We begin with the mathematical operation of *convolution* to see how the Fourier transform and the interval (or Patterson) function are related. For our purposes, the convolution of two characteristic functions, written $\chi_P(x) * \chi_Q(x)$, is

$$\chi_P(x) * \chi_Q(x) = \sum_{p\in P,\, q\in Q}\delta(x - (p + q)). \tag{28}$$

Another way of stating this is that the convolution of $\chi_P(x)$ and $\chi_Q(x)$ yields the characteristic function of their direct sum, $\chi_P(x) * \chi_Q(x) = \chi_{P\oplus Q}(x)$, where $P \oplus Q = \{p + q \mid p \in P, q \in Q\}$.[39] In the special case of the direct sum of P and $-P$ (the inversion of P about 0), we have the interval function of P:

$$P \oplus -P = \sum_{i,j} p_i - p_j = \Delta P. \tag{29}$$

Now, the spectrum of P can be obtained by multiplying the Fourier transforms of P and $-P$:

$$|\mathcal{F}_P| = \sqrt{\mathcal{F}_P \cdot \mathcal{F}_{-P}}. \tag{30}$$

By the convolution theorem, multiplying the Fourier transforms of two functions, f and g, is the same as taking the Fourier transform of the convolution of these functions:

$$\mathcal{F}_f \cdot \mathcal{F}_g = \mathcal{F}_{f*g}. \tag{31}$$

Putting together Equations 30 and 31, the power spectrum of P is equal to the Fourier transform of the convolution of P and $-P$, which is the interval function of P:

$$|\mathcal{F}_P|^2 = \mathcal{F}_{P*-P} = \mathcal{F}_{\Delta P}. \tag{32}$$

Thus, wherever the power spectrum is used in the metrics discussed in §6, we can substitute the Fourier transform of the interval function. For example, Equation 9 may be rewritten as

$$d_\theta(\mathcal{P},\mathcal{Q}) = \arccos\frac{\mathcal{F}_{\Delta P}\cdot\mathcal{F}_{\Delta Q}}{|\mathcal{F}_{\Delta P}||\mathcal{F}_{\Delta Q}|}. \tag{33}$$

Second, the Fourier transform preserves the dot product (up to a constant scaling factor). (This property is known as Parseval's theorem.) That is,

39 The direct sum is essentially the same as Cohn's transpositional combination and Boulez's pitch multiplication (Cohn 1991, Boulez 1971, Heinemann 1988).

$X \cdot Y$ is proportional to $\mathcal{F}_X \cdot \mathcal{F}_Y$. Thus, in Equation 33 we can substitute ΔX for $\mathcal{F}_{\Delta X}$ and rewrite the equation as

$$d_\theta(\mathcal{P},\mathcal{Q}) = \arccos \frac{\Delta P \cdot \Delta Q}{|\Delta P||\Delta Q|} \tag{34}$$

We can simplify Equation 6 similarly:

$$\begin{aligned} d_{\text{pow}}(\mathcal{P},\mathcal{Q}) &= \left||\mathcal{F}_P|^2 - |\mathcal{F}_Q|^2\right| \\ &= |\mathcal{F}_{\Delta P}| - |\mathcal{F}_{\Delta Q}| \\ &\propto |\Delta P - \Delta Q|. \end{aligned} \tag{35}$$

For n-tone equal tempered pitch-class sets, the above simplifications of d_θ and d_{pow} are only true when evaluated in $\mathbb{F}^{kn}$, where k is any positive integer. Specifically, in this case we have

$$d_{\text{pow}}(\mathcal{P},\mathcal{Q}) = \sqrt{kn}|\Delta P - \Delta Q|. \tag{36}$$

When these two measures are evaluated in Fourier spaces defined by other harmonics, their simplification in terms of the interval function is more complicated. Recall that for twelve-tone equal tempered pitch-class sets, d_{pow} in $\mathbb{F}^6$ simplifies to Equation 12. In order to understand these more complicated situations, we first make use of Equation 3 to rewrite Equation 6 so that $d_{\text{pow}}(\mathcal{P},\mathcal{Q})$ is

$$\left[\frac{\sum[\cos(\theta_{hi} + \theta_{jk}) + \cos(\theta_{hi} - \theta_{jk})] + \sum[\cos(\phi_{hi} + \phi_{jk}) + \cos(\phi_{hi} - \phi_{jk})]}{2} - \sum[\cos(\theta_{hi} + \phi_{jk}) + \cos(\theta_{hi} - \phi_{jk})]\right]^{1/2}. \tag{37}$$

Where $\theta_{ab} = 2\pi(p_a - p_b)z$, $\phi_{ab} = 2\pi(q_a - q_b)z$, and each summation is over all h, i, j, and k, and over $z = \frac{1}{12}, \ldots, \frac{kn/2}{12}$. Where $(p_h - p_i) \pm (p_j - p_k)$, $(q_h - q_i) \pm (q_j - q_k)$, or $(p_h - p_i) \pm (q_j - q_k)$ are not congruent to 0 mod n, the summation of the corresponding cosines for a given h, i, j, and k over all z vanishes. Where $(p_h - p_i) \pm (p_j - p_k)$, $(q_h - q_i) \pm (q_j - q_k)$, or $(p_h - p_i) \pm (q_j - q_k)$ are congruent to 0, the summation of the corresponding cosines for a given h, i, j, and k over all z is equal to kn. The number of ways in which $(p_h - p_i) \pm (p_j - p_k) \equiv 0$ is given by $2|\Delta P|^2$. Likewise, the number of ways in which $(q_h - q_i) \pm (q_j - q_k) \equiv 0$ and $(p_h - p_i) \pm (q_j - q_k) \equiv 0$ is given by $2|\Delta Q|^2$ and $2|\Delta P \cdot \Delta Q|$, respectively. Thus, evaluating n-tone equal tempered pitch-class sets by d_{pow} in $\mathbb{F}^{kn}$, we can rewrite Equation 37 as

$$d_{\text{pow}}(\mathcal{P},\mathcal{Q}) = \sqrt{kn\left(|\Delta P|^2 + |\Delta Q|^2 - 2|\Delta P \cdot \Delta Q|\right)}, \tag{38}$$

which simplifies to Equation 36.

Where n is even and d_{pow} is evaluated in $\mathbb{F}^{kn/2}$, where k is any odd positive integer, this situation is the same with one exception. Where $(p_h - p_i) \pm (p_j - p_k)$, $(q_h - q_i) \pm (q_j - q_k)$ or $(p_h - p_i) \pm (q_j - q_k)$ are odd, the summation of the corresponding cosines for a given h, i, j, and k over all z does *not* vanish, but is equal to –1. The number of ways in which $(p_h - p_i) \pm (p_j - p_k)$ is odd is given by $4O_PE_P$, where O_X and E_X are the number of odd and even intervals, respectively, in the interval function of X. Likewise, the number of ways in which $(q_h - q_i) \pm (q_j - q_k)$ and $(p_h - p_i) \pm (q_j - q_k)$ are odd is given by $4O_QE_Q$ and $2(O_PE_Q + O_QE_P)$, respectively. Since the summations involving $\theta_{hi} \pm \theta_{jk}$ and $\phi_{hi} \pm \phi_{jk}$ are halved in Equation 37, this yields the right side of the expression in the square root of Equation 12, $2(O_PE_P + O_QE_Q - (O_PE_Q + O_QE_P))$. Thus, for n even and k odd, the Euclidean distance between the power spectra in $\mathbb{F}^{kn/2}$ of pitch-class sets in n-tone equal temperament is

$$d_{\text{pow}}(\mathcal{P},\mathcal{Q}) = \sqrt{\tfrac{kn}{2}}\ |\Delta P - \Delta Q| + \sqrt{2\,(O_PE_P + O_QE_Q - (O_PE_Q + O_QE_P))}. \qquad (39)$$

As with pitch-class sets, for n-tone equal tempered pitch sets, the simplification of d_{pow} in Equation 15 is only true for certain values of a, b, and $f(z)$. Specifically, with $a = 0$, $b = kn/24$ $(k \in \mathbb{Z})$, and $f(z) = 1$, we have

$$d_{\text{pow}}(\mathcal{P},\mathcal{Q}) = \sqrt{b}\,|\Delta P - \Delta Q|. \qquad (40)$$

Alternatively, with $a = 0$, $b = kn/12$, and a scaling factor of $f(z) = 1 - (z/b)$, we have

$$d_{\text{pow}}(\mathcal{P},\mathcal{Q}) = \sqrt{\tfrac{b}{2}}\,|\Delta P - \Delta Q|. \qquad (41)$$

8.3 Simpson's rule

One of the most elegant ways of estimating definite integrals (Equation 13) is the composite Simpson's rule. Basically, we divide the interval of integration into many smaller parts and estimate the area of each part. Dividing the interval $[a, b]$ into n parts, Simpson's rule states that

$$\int_a^b g(x)dx \approx \tfrac{h}{3}\left[g(x_0) + 2\sum_{j=1}^{n/2-1} g(x_{2j}) + 4\sum_{j=1}^{n/2} g(x_{2j-1}) + g(x_n)\right], \qquad (42)$$

where $x_k = a + kh$ for $k = 0, 1, \ldots, n$ and $h = (b - a)/n$. By dividing the interval $[a, b]$ into smaller parts and reducing the size of the "step length" h, we can decrease the discrepancy between the estimation and the definite integral. Substituting Equations 14 and 16 for $g(x)$ gives approximations of d_{pow} and d_θ between pitch sets.

Works Cited

Amiot, Emmanuel. 2007. "David Lewin and Maximally Even Sets." *Journal of Mathematics and Music* 1/3: 1–14.

Buchler, Michael. 1997. "Relative Saturation of Subsets and Interval Cycles as a Means for Determining Set-Class Similarity." Ph.D. diss., University of Rochester.

———. 2000. "Broken and Unbroken Interval Cycles and Their Use in Determining Pitch-Class Set Resemblance." *Perspectives of New Music* 38/2: 52–87.

Bullough, R. K. 1961. "On Homometric Sets. I. Some General Theorems." *Acta Crystallographica* 14: 257–68.

Callender, Clifton. 2004. "Continuous Transformations." *Music Theory Online* 10/3.

Callender, Clifton and Rachel Hall. 2007. “Homometric Sets and Z-related Chords.” Paper presented at the American Mathematical Society National Conference, New Orleans.

Callender, Clifton, Ian Quinn, and Dmitri Tymoczko. 2008. “Generalized Voice-Leading Spaces.” *Science* 320: 346–48.

Carter, Elliott. 2002. *Harmony Book*, ed. Nicholas Hopkins and John F. Link. New York: Carl Fischer.

Castrén, Marcus. 1994. “RECREL: A Similarity Measure for Set-Classes.” Ph.D. diss., Sibelius Academy.

Cohn, Richard. 1991. “Properties and Generability of Transpositionally Invariant Sets.” *Journal of Music Theory* 35: 1–32.

———. 1998. “Square Dances with Cubes.” *Journal of Music Theory* 42: 283–95.

Grisey, Gérard. 1987. “Tempus ex Machina: A Composer’s Reflection on Musical Time.” *Contemporary Music Review* 2: 241–77.

Hall, Rachel and Dmitri Tymoczko. 2007. “Poverty and Polyphony: A Connection between Music and Economics.” Preprint available at http://www.sju.edu/~rhall/povertypolyphony.pdf.

Isaacson, Eric. 1990. “Similarity of Interval-Class Content between Pitch-Class Sets: The IcVSIM Relation.” *Journal of Music Theory* 34: 1–28.

———. 1996. “Issues in the Study of Similarity in Atonal Music.” *Music Theory Online* 2/7. *http://societymusictheory.org/mto/issues/mto.96.2.7/mto.96.2.7.isaacson.html.*

Lewin, David. 1959. “Re: Intervallic Relations between Two Collections of Notes.” *Journal of Music Theory* 3: 298–301.

———. 1982. “On Extended Z-Triples.” *Theory and Practice* 7: 38–39.

———. 1987. *Generalized Musical Intervals and Transformations.* New Haven: Yale University Press.

———. 1998. “Some Ideas about Voice Leading between Pcsets.” *Journal of Music Theory* 42: 15–72.

———. 2001. “Special Cases of the Interval Function between Pitch-Class Sets *X* and *Y*.” *Journal of Music Theory* 45: 1–29.

Morris, Robert. 1979. “A Similarity Index for Pitch-Class Sets.” *Perspectives of New Music* 18/1: 445–60.

———. 1995. “Equivalence and Similarity in Pitch and Their Interaction with Pcset Theory.” *Journal of Music Theory* 39: 207–43.

Quinn, Ian. 2006. “General Equal Tempered Harmony (Introduction and Part I).” *Perspectives of New Music* 44/2: 114–58.

———. 2007. “General Equal-Tempered Harmony (Parts II and III).” *Perspectives of New Music* 45/1: 4–63.

Roeder, John. 1987. “A Geometric Representation of Pitch-Class Series.” *Perspectives of New Music* 25/1–2: 362–409.

Rogers, David. 1999. “A Geometric Approach to Pcset Similarity.” *Perspectives of New Music* 37/1: 77–90.

Rose, Francois. 1996. “Introduction to the Pitch Organization of French Spectral Music.” *Perspectives of New Music* 34/2: 6–39.

Rosenblat, Joseph. 1984. “Phase Retrieval.” *Communications in Mathematical Physics* 39: 77–100.

Rust, Douglas. 1995. “Conversation with Witold Lutosławski.” *Musical Quarterly* 79: 207–23.

Scott, Damon and Eric J. Isaacson. 1998. “The Interval Angle: A Similarity Measure for Pitch-Class Sets.” *Perspectives of New Music* 36/2: 107–42.

Smith, Julius O., III. 2008. *Mathematics of the Discrete Fourier Transform (DFT) with Audio Applications*, 2nd ed. *http://ccrma.stanford.edu/~jos/mdft/.*

Soderberg, Stephen. 1995. “Z-Related Sets as Dual Inversions.” *Journal of Music Theory* 39: 77–100.

Straus, Joseph. 2003. “Uniformity, Balance, and Smoothness in Atonal Voice Leading.” *Music Theory Spectrum* 25: 305–52.
Stucky, Steven. 1981. *Lutoslawski and His Music*. Cambridge: Cambridge University Press.
———. 2005. “Voice Leading in Set-Class Space.” *Journal of Music Theory* 49: 45–108.
Teitelbaum, Richard. 1965. “Intervallic Relations in Atonal Music.” *Journal of Music Theory* 9: 72–127.
Tymoczko, Dmitri. 2006. “The Geometry of Musical Chords.” *Science* 313: 72–74.
———. 2008. “Voice Leading and the Fourier Transform.” *Journal of Music Theory* 52, in press.
Vuza, Dan Tudor. 1993. “Supplementary Sets and Regular Complementary Unending Canons (Part Four).” *Perspectives of New Music* 31/1: 270–305.

Clifton Callender is associate professor of composition at Florida State University and Associate Editor of *Perspectives of New Music*. He recently published “Generalized Voice-Leading Spaces,” coauthored with Ian Quinn and Dmitri Tymoczko, in *Science* and completed a commission, *Reasons to Learne to Sing*, for the College Music Society’s 50th Anniversary.