

Chapter 1

A brief history and critique of interval/subset class vectors and similarity functions.

1.1 Basic Definitions

There are a few terms, fundamental to musical atonal theory, that are used frequently in this study. For the reader who is not familiar with them, we present a quick primer.

Pitch class (pc) is used to denote a pitch name without the octave designation. C4, for example, refers to a pitch in a particular octave; C is more generic, referring to any and/or all Cs in the sonic spectrum. We assume equal temperament and enharmonic equivalence, so the distance between two adjacent pcs is always the same and $\text{pc C} = \text{pc B}\sharp$. Pcs are also commonly labeled with an integer. In such cases, we will adopt the standard of $0 = \text{C}$, $1 = \text{C}/\text{D}$, ... $9 = \text{A}$, $a = \text{A}\sharp/\text{B}\flat$, and $b = \text{B}$ ('a' and 'b' are the duodecimal equivalents of decimal 10 and 11).

Interval class (ic) represents the smallest possible pitch interval, counted in number of semitones, between the realization of any two pcs. For example, the interval class of C and G— $\text{ic}(\text{C}, \text{G})$ —is 5 because in their closest possible spacing, some C and G (C down to G *or* G up to C) would be separated by 5 semitones. There are only six interval classes (1 through 6) because intervals larger than the tritone (ic6) can be reduced to a smaller number (interval 7 = ic5, interval 8 = ic4, etc.).

A *pcset* is an unordered set of pitch classes that contains at most one of each pc. Using numeric notation, the set {C, D, E \flat , F \sharp } would be represented as {0,2,3,6}.

The *aggregate* (U) is the complete set of all twelve pcs.

The *interval-class vector* (ICV) is an enumeration of the interval classes found between members of a pcset. Let $X = \text{pcset } \{0,2,3,6\}$. There are six distinct pairs of notes in any tetrachord; the ics within these pairs are as follows:

$$\{0,2\} = \text{ic}2$$

$$\{0,3\} = \text{ic}3$$

$$\{0,6\} = \text{ic}6$$

$$\{2,3\} = \text{ic}1$$

$$\{2,6\} = \text{ic}4$$

$$\{3,6\} = \text{ic}3$$

We now place the numbers of each ic present into a six-argument vector, with each place in the vector representing the corresponding ic (the first argument in the vector represents the number of ic1s in the pcset, etc.). In other words, the number of ic n s is placed in argument n of the ICV; there are two ic3s in pcset X , so the third argument in $\text{ICV}(X)$ is 2. Therefore, $\text{ICV}(\{0,2,3,6\}) = \langle 112101 \rangle$. From this, we can quickly see that there are no instances of ic5, two instances ic3, and one of each of the remaining ics.

Even though there is one and only one interval present in a set of cardinality 2, a dyad (or dyad class) is not the same as an interval. An interval represents a distance between notes; a two-pc (cardinality 2, or #2) set is a group with two members. Despite this important distinction, it is possible to

think of the interval-class vector as a #2 subset-class vector—that is a listing of the number and type of #2 subsets embedded in a set.¹ While such a vector would be identical in appearance to the ICV, conceptualizing the elements are subsets rather than intervals will allow us to generalize outwards and produce other subset-class vectors (e.g., a trichord-class vector, tetrachord-class vector, etc.).

Before further discussing other subset-class vectors, it will be useful to define abstract pcset inclusion. Given two set classes (SCs) X (denoted $/X/$) and $/Y/$ where $/X/$ has the same number or fewer elements than $/Y/$ ($\#X/ \leq \#Y/$),² $AS(/X/, /Y/)$ is a Boolean function (relation) which returns “true” if at least one X of $/X/$ is embedded in some Y of $/Y/$.³ $AS(/X/, /Y/)$ returns “false” if no form of $/X/$ is embedded in some Y of $/Y/$. If $AS(/X/, /Y/) = \text{true}$ then $/X/$ is said to be an abstract subset of $/Y/$.⁴ For example, let us consider three set classes: $/A/ = [014]$, $/B/ = [0145]$, $/C/ = [0167]$. $AS(/A/, /B/) = \text{true}$; $AS(/A/, /C/) = \text{false}$.⁵

David Lewin’s $EMB(/X/, /Y/)$ function⁶ provides more specific information than does our AS relation. $EMB(/X/, /Y/)$ returns the number of distinct forms of $/X/$ which are embedded in some Y of $/Y/$.⁷ If, for example, $/X/ = SC(3-1) [012]$ and $Y = \{1,2,3,6,7,8,a\}$, $EMB(/X/, /Y/) = 2$; that is, both $\{1,2,3\}$

¹There is a one-to-one relationship between the number of ic n in a pcset and the number of SC $2-n$ in the same pcset.

² $\#X/ =$ cardinality of (number of elements in) pcset X .

³We will, for the time, assume that forms of X are all possible transpositions ($T_n(X)$) and inversions ($T_nI(X)$) of X .

⁴Because the use of set class is inherent to the notion of “abstract subset,” $AS(/X/, /Y/)$ is functionally the same as $AS(X, Y)$. If $X \subseteq Y$, $/X/$ is an abstract subset of $/Y/$.

⁵As a numerical values, “true” and “false” are normally represented by the digits 1 and 0, respectively.

⁶Lewin 1977 and Lewin 1979-80a. In the latter article, Lewin uses the form $EMB(/X/, Y)$, but comments that $EMB(/X/, /Y/)$ provides the same information. (Lewin 1979-80a, 433.) Consistent with our definition of AS , we prefer the latter because it deals entirely with set classes.

⁷We will, for now, use the common canonical operations transposition (T_n) and transposition and inversion (T_nI) to calculate the distinct forms of $/X/$.

and {6,7,8} are forms of [012] and are included in {1,2,3,6,7,8,a}. Given that {1,2,3} (or any form of [012]) is an inversionally symmetrical set, it can map onto itself through two distinct operations: the identity operation ($T_0\{1,2,3\} = \{1,2,3\}$); and under transposition and inversion ($T_3I\{1,2,3\} = \{1,2,3\}$). These two transformations yield the same pcset and will be considered a single *distinct* form of 3-1.

Abstract inclusion vectors, including n -class vectors (or n CVs⁸), such as the #2 subset-class vector (which, again, is equivalent in appearance to the ICV) can be derived by performing $EMB(/X/, /Y/)$ for each distinct $SC(X)$ where $\#X = n$. When $\#X = 2$, $EMB(/X/, /Y/)$ returns an argument of the #2 subset-class vector of $SC /Y/$ ($ICV(Y)$). Formally, we can define each argument in the 2CV as follows: $2CV(X)_i = EMB(i, /X/)$.⁹

1.2 Evolution of the interval class vector

The ICV has played an integral part in the history of atonal theory, serving as the prototype for most other abstract subset inclusion vectors and, more specifically, to the new vectors and relations which will be introduced in chapter 2 of this study. In this section, we will chronicle its evolution and variations.

⁸This abbreviation is adopted from Castrén (Castrén 1995, 3) who, in turn, adopted it from Lewin (Lewin 1987, 106-7).

⁹We could also have defined the ICV using Lewin's function $COV(/X/, /Y/)$. This function, also presented in Lewin 1979-80, 434, returns the number of distinct forms of $SC Y$ that cover—or include—a member of $/X/$. Where $\#X = 2$ (that is, when X is a dyad), $COV(/X/, /Y/) = EMB(/X/, /Y/)$. COV and EMB will also return the same value in all cases where pcset X is inversionally symmetrical. When pcset X is not inversionally symmetrical, these two functions return different values. Morris 1987, 90 provides an example where these two functions serve different purposes. Since we will be dealing with abstract inclusion of all pcset classes in the course of this study, $EMB(/X/, /Y/)$ will serve as a more useful model than $COV(/X/, /Y/)$.

1.2.1 Hanson, 1960

Howard Hanson's chord or scale "interval analysis," presented in his 1960 book *Harmonic Materials of Modern Music: Resources of the Tempered Scale*, is, to our knowledge, the first enumeration of the interval-class vector and thus its first use in comparing and relating different pcsets.¹⁰ The actual term "interval vector" was coined several years later by Allen Forte. Hanson represented each of the six interval classes with a different letter: *p* ("perfect") represented ic5, *m* ("major third") represented ic4, *n* ("minor third") represented ic3, *s* ("major second") represented ic2, *d* ("minor second") represented ic1, and *t* ("tritone") represented ic6. These letters were used to delineate SC and always fell in the order "*pmnsdt*."¹¹ Each letter was only used in a SC label if its respective ic was embedded in the set. E.g., SC(3-11) [037] = *pmn* in Hanson's terminology. This is the same as saying that it has one ic5, ic4, and ic3, and none of the other ics. When there was more than one embedding of a particular ic in a SC, superscripts carrying the EMB(*ic*,/*X*/) number followed each letter. For example, SC(3-9) [027] = *p²s*.

Hanson's *pmnsdt* interval analysis was, in part, meant to show composers how they could project certain intervallic structures using particular SCs. His unique labeling scheme introduced a number of concepts that were reintroduced more formally by later theorists. For example, while Hanson labeled SCs using his mnemonic version of the ICV, he understood that ICV

¹⁰Even if that comparison is a bit less systematic than more recent attempts.

¹¹This ordering was from most-to-least consonant intervals. Hanson's tonally-analogous names and the order in which they occur represent his manner of thinking about these abstract structures.

equivalence was not the same as transpositional or inversive equivalence and that, in some cases, SCs with the same ICV might be neither transpositions or inversions of each other (Forte later labeled such SCs, Z-related). Hanson first mentioned this within the context of what he called “isomeric twins.”¹²

Furthermore, Hanson described one aspect of what Morris calls the ZC relation. Hanson remarked that “Every *isometric* six-tone scale formed by the simultaneous projection of two intervals has an *isomeric* “twin” having the identical intervallic analysis.”¹³ Using Morris’s definition, “complementary pcsets A and \bar{A} are ZC related if there is no TTO [twelve-tone operator] that can map one into a subset of the other.”¹⁴ Hanson initially only deals with the case of hexachords; complementary hexachords are either T_n/T_nI related or they are ZC related.

We have already noted ... that every six-tone scale has a *complementary* scale consisting in each case of the *remaining six tones* of the twelve-tone scale. We have also noted that these complementary scales vary in their formation. In certain cases, ... the complementary scale is simply a *transposition* of the original scale. In other cases, ... the complementary scale is the *involution* [inversion] of the original scale. However, in fifteen cases the complementary scale has an entirely different order, although the same intervallic analysis.¹⁵

Later, he described a “maverick” sonority—a five-note collection which “is *not itself a part of its own complementary scale*. It is, instead, a part of the “twin” [i.e., Z-relation] of its own complementary scale.”¹⁶

¹²Hanson’s first description of isomeric twins comes very early in his book (p. 22), where he notes that: “There are a few sonorities which *have the same components but which are not involutions one of the other*, although each has its own involution. Examples are the tetrads C-E-F[♯]-G and C-F[♯]-G-B^b. Each contains one perfect fifth, one major third, one minor third, one major second, one minor second, and one tritone (*pmnsdt*), but one is *not* the involution [inversion] of the other—although each has *its own* involution.

We shall describe such sonorities ... as *isomeric* sonorities.”

¹³Hanson 1960, 196 (Hanson’s italics).

¹⁴Morris 1987, 74.

¹⁵Hanson, 254 (Hanson’s italics).

¹⁶Hanson, 331 (Hanson’s italics).

Hanson noticed that all complementary hexachords (hexads) contained the same intervallic analysis (i.e., same ICV). He also noticed a trend that will be central to this present study: that all complementary sets are similarly saturated with the same interval classes, and, further that composers could use complementary sets (in a non-twelve-tone serial environment) to “project” the same intervallic profile. Hanson devoted one chapter of his book to examining this phenomenon.

An analysis of all the sonorities of the twelve-tone scale will reveal the fact that *every* sonority has a complementary sonority composed of the *remaining tones* of the twelve-tone scale and that the complementary scale will always have the same *type* of intervallic analysis, that is, the predominance of the same interval or intervals. In other words, every two-tone interval has a complementary ten-tone scale, every triad has a complementary nine-tone scale, ...¹⁷

The major triad C-E-G has a complementary nine-tone scale consisting of the *remaining nine tones of the chromatic scale*, the tones C#-D-D#-F-F#-G#-A-A# and B. We shall observe in analyzing this scale that it has seven perfect fifths, seven major thirds, and seven minor thirds, but only six major seconds, six minor seconds, and three tritones—*that it predominates in the same three intervals which form the major triad*. ... This nine-tone scale we shall call the *projection* of the major triad, since it is in fact the *expansion* or projection of the triad to the nine-tone order.¹⁸

The importance of this principle to the composer can hardly be overstated, since it allows the composer to *expand any tonal relation with complete consistency*.¹⁹

Hanson discussed both how a composer could project a particular trichord class and, more fundamentally, each interval class. Through the former, he anticipated Martino’s discussion of how source trichords can be combined into combinatorial hexachords; through the latter, he generated the

¹⁷Hanson, 263 (Hanson’s italics).

¹⁸Hanson, 264 (Hanson’s italics).

¹⁹Hanson, 263 (Hanson’s italics).

cyclic set classes²⁰ and briefly discussed some of their properties. While Hanson's descriptions of cyclic sets, interval- or set-class projection, and complementary equivalence are not formally stated, they are there and remain true. Moreover, they were the first published coherent investigations of relations later restated and developed in atonal music theory.

1.2.2 Lewin, 1960

David Lewin did not present the ICV as a single six-part array as did Hanson (and later Martino and Forte), but he did create all the tools necessary to derive its values and, more importantly, posited theorems that generalize some properties of pcset complementation. En route to these theorems, Lewin introduced the interval function, which measures the number and size of intervals that occur either between elements of two sets (Lewin 1959) or among elements of a single set (Lewin 1960). $p(i)$ represents the number of interval i that occur between any two elements of set P and a copy of P (this can be seen as a precursor to Lewin's $EMB(/X/,/Y/)$ or $i(X, Y)$ functions). Using Lewin's symbols, p is an element of set P and p' is an element of \bar{P} (the complement of P). The variable x is the cardinality of set P .

Formula I: $p'(i) = 12 - 2x + p(i)$ for every i .

... it is obvious that transposing or inverting a collection of notes [pcs] does not alter its intervallic content. Thus, if P is a collection, any transposition or inversion of P has the same intervallic content as does P . Second, from formula I, we see that if P is a collection of six notes, where $x = 6$, then $12 - 2x = 0$, and $p'(i) = p(i)$ for all i . Hence any six-note collection whatsoever has the same intervallic content as its

²⁰Chapter 41 (pp. 274-284) of Hanson ("Projection of the Six Basic Series with Their Complementary Sonorities") is an appendix of the cyclic set classes discussed in chapter two of this present study.

complement, and therefore the same intervallic content as any transposition or inversion of its complement.²¹

In addition to this rule of hexachordal complementation, Lewin, like Hanson, also recognized that there are set pairs that share the same ic content, but are not transpositionally and/or inversionally related. He listed the exceptional pairs, but could not offer a formal explanation of how they are related as transformations under a pc operator.²²

1.2.3 Martino, 1961

Donald Martino's article, "The Source Set and its Aggregate Formations," appeared in *Journal of Music Theory* in 1961. The primary focus of this study was to create a catalog of all the different SCs in pc space and describe their invariant and aggregate-forming properties in the context of twelve-tone composition. As a means of ordering the SCs in his tables (and to illustrate a method of determining transpositional invariance), Martino "counts the intervals" (to use his words) in each pcset's normal form. His six-argument interval count was ordered from smallest-to-largest rather than Hanson's tonally biased "consonant-to-dissonant" ordering. It is identical to the structure Allen Forte's interval vector of 1964.

Martino's set class list assumes equivalence based upon identical ICVs, not upon transformation. Therefore, like Hanson (and unlike Lewin), Martino offers only one name class for each unique ICV, thereby placing Z-related pcsets into the same equivalence class.

²¹Lewin 1960, 99.

²²No other theorist since 1960 has been able to make this connection either.

1.2.4 Forte, 1964 and 1973

Allen Forte first coined the term “interval vector” in his 1964 article “A Theory of Set Complexes for Music.”²³ In that important article, he reformulated, if independently, Hanson’s concepts and also provided a series of similarity functions which relate sets using the ICV and a theory of set complexes based upon mutual inclusion of larger SCs.²⁴ These theoretic tools were also described, if in different guises, in his 1973 book, *The Structure of Atonal Music*. Most of our references will therefore come from that later work.

The interval vector is an ordered array; that is, the first (left-most) number is always the number of occurrences of ic1, the next number to the right is always the number of occurrences of ic2, and so on. If some interval class is not represented, the entry 0 (zero) should be in the appropriate position in the vector.²⁵

While Lewin (1960), Hanson, and Martino acknowledge the existence of sets which have the same ICV, yet are not related under transposition or inversion, Forte was the first theorist to create a labeling system which distinguishes these set class pairs. Forte calls these special SC pairs Z-related.²⁶ While in Hanson, Martino, and Forte (1964) Z-related pairs had been treated as a single SC, Forte (1973) defines each SC as equivalent under T_n/I (transposition and/or inversion); each so-defined SC has its own label because “unless a distinction is made on the basis of fundamental pcset

²³Forte 1964, 141. We refer to it as the “interval-class vector” because the vector’s six arguments are more precisely said to represent *interval class* content.

²⁴In Forte’s writings, he encloses the ICV in square brackets ([]). Throughout this study, ICVs will be enclosed in angle brackets (< >). We prefer this notation because it allows us to reserve square brackets for prime forms and angle brackets for ordered arrays, such as the ICV. This is the notation that John Rahn uses in his text *Basic Atonal Theory* (Rahn 1980).

²⁵Forte 1973, 15.

²⁶Forte 1973, 21.

structure, with the identical vector characteristic as secondary, it is not possible to make the necessary and (as will be shown) fruitful differentiation between pitch-class structures and interval-class structures over the full range of the available resources.”²⁷ “Moreover,” Forte added, “the members of a Z-related pair are also distinct with respect to subset structure if the cardinal number of the subset is greater than 2.”²⁸

1.2.5 Morris, 1982

The ways in which pcsets are considered “equivalent” are explored in great detail by Robert Morris in his article “Set Groups, Complementation, and Mappings among Pitch-Class Sets.” Morris clarified two basic types of pcset equivalence: two sets which can be related under some canonical twelve-tone operator (TTO) are considered *operational equivalences*; two sets which share the same interval-class content are considered *vector equivalences*.²⁹ The latter is a superset of the former; for example, the standard interval-class vector yields invariant values for transpositionally- and/or inversionally-related (T_n/I) sets, and also for the twenty-three SC pairs that Allen Forte calls Z-related.

Morris (1982) defines a series of set groups (SG) that distinguish equivalence under different pc transformations. Under SG(1), only transpositionally-related sets are members of the same SC. Therefore {1,2,4} and {2,3,5} are SG(1) equivalent, but neither is equivalent to {1,3,4} (which is related by T_5I to the former and by T_6I to the latter). {1,2,4} and {1,3,4} are,

²⁷Forte 1973, 21.

²⁸Forte 1973, 21n.

²⁹For example, Hanson defined SCs using vector equivalence; Forte 1973 defined SCs using operational equivalence (allowing T_n and T_nI as canonical operations).

however, Z-related in SG(1) since they contain the same interval content. SG(2) expands the operational canon to include inversionally-related sets. Therefore, under SG(2), each of these three sets is equivalent to the other two. SG(3) includes transposition, inversion, and/or multiplication by 5 (M).³⁰ Thus, {0,1,4} and {0,3,7} are SG(3) related (since the former is related by T₇M to the latter).

Just as the sets which have the same interval-class vector include SG(1) and SG(2) equivalent sets, one can create an interval-class vector that includes SG(3) equivalences as well.³¹ Since the operations T_nM and T_nMI effectively swap the values in the ic1 and ic5 arguments,³² the SG(3)V (set group 3 vector) combines the ic1 and ic5 values in the ICV into a single argument (the first) of a five-argument vector. For example, ICV(6-1 [012345]) = <543210>; SG(3)V(6-1) = <64320>.³³

This alternative vector creates an atmosphere where a set and its M-relation yield identical ICVs (and are therefore considered SG(3) Z-relations). For example, let us consider the set 6-32 [024579] relative to 6-1:

$$\begin{array}{ll} \text{ICV}(6-32) = \langle 143250 \rangle & \text{SG}(3)\text{V}(6-32) = \langle 64320 \rangle \\ \text{T}_n\text{M}(6-32) \quad 6-1 & \text{SG}(3)\text{V}(6-1) = \text{SG}(3)\text{V}(6-32) = \langle 64320 \rangle. \end{array}$$

Morris refers to these sets (and all other SG(3)-related pairs) as a single SC. In this case, the composite SG(3) SC is 6-1/32.

SG(3) vector equivalence is broader than SG(3) transformationally-related sets. Just as the ICV yields some equivalences that we refer to as SG(1)

³⁰By closure, SG(3) also includes M and I (MI, or multiplication by 7).

³¹Starr 1978 deals at length with T_nM and T_nMI equivalence (see pp. 7ff).

³²For example, T₀M{0,1} {0,5} and T₀M{0,5} {0,1}. All pc pairs separated by ic1 map to pc pairs separated by ic5 and vice versa, while all other ic relations are held invariant.

³³Morris 1982, 102-9.

or SG(2) Z-relations, this new vector also yields a number of SG(3) Z-relations.³⁴ For example let us now consider 6-8 [023457] in relation to the aforementioned SG(3) 6-1/32 SC:

$$\begin{aligned} \text{ICV}(6-8) &= \langle 343230 \rangle & \text{SG}(3)\text{V}(6-8) &= \langle 64320 \rangle \\ \text{SG}(3)\text{V}(6-8) &= \text{SG}(3)\text{V}(6-1/32) & \text{SG}(3) \text{ 6-8} & \text{SG}(3) \text{ 6-1/32} \\ & & \text{6-8 is SG}(3) \text{ Z-related to 6-1/32} & \end{aligned}$$

Just as SG(3) SG(2) SG(1), SG(3)V SG(2)V (i.e., ICV).

Morris also added another SG which includes all the relations of SG(3). SG() includes the transformation where one whole-tone collection is held invariant while the other is transposed by T_n (where n is always even—that way an odd-numbered pc will not get transposed onto an even one, or vice versa).³⁵ This operation was also posited in Mead 1989 as the O operation.³⁶ O^z is an operation which transforms only the odd-numbered pcs (where 0 = C), while maintaining the even-numbered pcs invariant. Its pc cycles are (0) (2) (4) (6) (8) (a) (1, 3, 5, 7, 9, b).³⁷ For example, $O^1\{0, 1, 2, 3, 4, 5\}$ $\{0, 2, 3, 4, 5, 7\}$. O^z equivalence, by itself, subsumes $T_0\text{MI}$ equivalence since that operation maintains all the even pcs in a set while transposing the others by T_6 (i.e., $O^3 = T_0\text{MI}$). When we add O to our canon of T and I (SG(2)), we have a superset of SG(3) equivalence: $T_n O^3 = T_n \text{MI}$ and $T_n I O^3 = T_n \text{M}$. Under

³⁴SG(2) Z-relations are the same as Forte's original conception of the Z-relation since Forte's ICV = Morris's SG(2)V.

³⁵Morris 1982, 115-6.

³⁶Mead 1989, 224ff. Mead's O operation is not the same as Morris's . Morris's operation includes Mead's. In this study, we will slightly reformulate O.

³⁷In Mead 1989, O is followed by a subscript (z), that represents the transposition level applied to the odd numbered pcs. Thus, O_2 transforms the odd numbered pcs by T_2 . In this study, we are slightly reformulating O so that O^z where $z = 1$ represents a single shift (to the right) in the pc cycles of O. Our $O^1 = \text{Mead's } O_2$, our $O^2 = \text{Mead's } O_4$, and so forth. $O^0 = O^6$, the identity operation.

SG() our SG(3) Z-related pair (6-8 and 6-1/32) would map onto each other,

For example:

$$A = \{0,1,2,3,4,5\}$$

$$B = \{0,2,3,4,5,7\}$$

$$C = \{0,2,4,5,7,9\}$$

$$T_0O^1(A) = B \quad T_0O^1(B) = C \quad T_0O^2(A) = C$$

Like all groups of operators, O is also transitive (this property is illustrated in the last line of the above example). These three SG(2) SCs (A, B, and C) fold into a single SG() SC, which we will denote 6-1/8/32.³⁸

The SG()V is created by combining all the odd ics in the ICV together into a single argument—the first argument of the vector, followed by the number of ic2, ic4, and ic6, respectively.³⁹ For example, SG()V(6-1/8/32) = <9420> (i.e., there are four ic2s, two ic4s, no ic6s, and a combined total of nine ic1s, ic3s, and ic5s). Z-relations also occur among SG() classes, but there are only five SG() Z-related pairs⁴⁰—fewer than found in the other SG(n)s discussed herein. These extensions to the ICV posited by Morris⁴¹ suggest that those new functions which will be presented in chapter 2 can be applied *mutatis mutandis* to any ICVs.

³⁸Presumably to avoid such awkward labelings, Morris renumbers the set classes in his SG() in an appendix to his article. (Morris 1982, 137-8.)

³⁹SG() itself is determined not by vector equivalence, but by transformation under T_n , I, M, and O_z . If $SC(X) = O_z SC(Y)$ then $SC(Y), SC(X) \in SG()$.

⁴⁰Morris 1982, 137-8.

⁴¹For the sake of space, not all of Morris's set groups and related vectors have been summarized herein.

1.3 Alternate representations of the interval-class vector

In the second chapter of this study, we will introduce several interpretations on the interval-class vector which facilitate comparisons among SCs of different cardinalities. Two other theorists have also used this basic approach in constructing means for relating sets: Tore Ericksson (1986) and Marcus Castrén (1994). We will briefly explore their variations in sections 1.3.1 and 1.3.2 respectively.

1.3.1 Tore Ericksson's mm vector

In his effort to weight the different ICV arguments, Tore Ericksson (1986) created the mm vector (maximum to minimum vector), a reordering of the ICV that eliminates the specific arguments and replaces them with an ordering of the numbers one through six, depending on which ics have greatest prominence. Because of the transpositional symmetry of ic6, its ICV value is doubled for the sake of the mm-vector.⁴² Thus, the mm-vector of SC [036] (3-10) (ICV = <002001>) is <36, 1245>. This means that the SC has the same amounts of ic3 and ic6 and the same amounts of ic1, ic2, ic4, and ic5, and that it has more of the first group than of the second (though it is not specified how much more).

While the mm-vector lacks the specificity of the ICV (i.e., we do not know how many of each ic are included in a given SC), it is useful in establishing families of related SCs. In particular, Ericksson mentions that the mm-vectors of all cyclic set classes (see chapter 2 of this study for a detailed

⁴²Because the interval of the tritone divides the aggregate in half, there are half as many distinct members of SC[06] as there are members of any other dyad class (6 as opposed to 12). In creating a vector which shows the relative content of each ic, Ericksson doubles the number of ic6s to help level the playing field.

explanation of this term) of a particular T_n cycle share the same qualities.⁴³ In figure 1.1, we provide the ICVs and mm-vectors for four such pcsets, each of which has the maximum amount of ic3 content for a set of its cardinality:

Despite their different cardinalities, sets *A*, *B*, and *D* all share the same mm-vector, creating a Z-relation of sorts among SCs of different cardinality; the mm-vector of *C* differs only slightly from the others (since Ericksson doubles the amount of ic6 for the sake of parity, it is impossible to have an odd amount of ic6 and it is therefore impossible for there to be the same amount of ic3 and ic6 in this situation). Complementary SCs produce identical mm-vectors, reflecting the fact that complementary sets are saturated with the same interval classes to precisely the same ratio (this observation was initially made by Howard Hanson⁴⁴ and will be further developed in chapter 2 of this dissertation). Using the mm-vector as a means for vector equivalence, Ericksson refers to complementary sets in tandem (e.g., 5/7-33).

1.3.2 Castrén's $nC\%$ -vector

Marcus Castrén (1994), in a similar effort, has created the $nC\%$ -vector. As above, “ nC ” stands for the cardinality of subset which it lists. A $2C\%V$, therefore lists dyad (interval) classes, a $3C\%V$ lists trichord classes, etc. Like the mm-vector, the $2C\%V$ is a relational reworking of the ICV. Unlike the mm-vector, however, the relative numbers of embedded ics are not entirely lost. The $2C\%V$ is derived by dividing the number in each argument of an ICV by the cardinality of the ICV ($\#ICV$) then taking the quotient and multiplying it

⁴³Ericksson groups the mm-vectors into families of partially-ordered mm-vectors. For example, the four SCs in example 1.c all belong to the mm-vector family $\langle 36 | 1245 \rangle$. That is, each member of the family will have more ic3 and ic6 than any other ic. But any particular member might have more ic3 than ic6 or more ic6 than ic3 and more ic2 than ic1, etc. See Ericksson, 100-104.

⁴⁴Hanson, 263 (see quote in section 1.2.1 of this chapter).

by 100 (and discarding any remainder). The 2C%V of set class 3-10 [036] would be derived as follows:

$$\begin{array}{rcl}
 \text{ICV}(3-10) & = & \langle 0 \quad 0 \quad 2 \quad 0 \quad 0 \quad 1 \quad \rangle \\
 \#\text{ICV}(3-10) & = & 3 \\
 \frac{\text{ICV}(3-10)_i}{\#\text{ICV}(3-10)} & = & \langle 0.00 \quad 0.00 \quad 0.67 \quad 0.00 \quad 0.00 \quad 0.33 \rangle \\
 \cdot 100 & = & \\
 2\text{C}\%V(3-10) & = & \langle 0 \quad 0 \quad 67 \quad 0 \quad 0 \quad 33 \rangle
 \end{array}$$

More formally: $2\text{C}\%V_n(X) = \frac{\text{ICV}_n(X)}{\#\text{ICV}(X)} \cdot 100$, for all n .

Castrén's $n\text{C}\%V$ has the advantage of having the same vector cardinality (100),⁴⁵ no matter what cardinality set class it is representing. This makes comparing two set classes of any cardinality very simple and assures that comparison values will always fall within a constant range.

Unfortunately, there is a problem that comes with using the $n\text{C}\%V$ in a similarity function: its range of possible values decreases both when the cardinality of the set increases and also, to a lesser degree, when the number of distinct subset classes of cardinality n increases. The reason behind this is not difficult to fathom; the larger the set or the more arguments in the vector (i.e., there is one vector argument for each subset class of cardinality n), the smaller the proportion of the whole any one can occupy.

Figure 1.2 lists the minimal and maximal 2C%V values for any argument in the vector, given pcsets of cardinality 2 through 12. One can readily see

⁴⁵ Actually, the cardinality is 100 ± 2 since the values are rounded. For example, both 4-z15 [0146] and 4-z29 [0137] have the ICV $\langle 111111 \rangle$. This means that each argument in their 2C%V is derived by dividing 1 by 6 (the cardinality of the ICV). $\frac{1}{6} = 0.1\overline{66}$, which is rounded to 0.17 and multiplied by 100. 17 times 6 = 102, the cardinality of 2C%V(4-z15).

that the $2C\%V$ maximal values decrease as set cardinality increases. This effect will naturally be even more pronounced when dealing with vectors with more arguments ($nC\%V$, where $2 < n < 10$), and particularly when calculating the $nC\%V$ of hexachords and larger sets. While the minimal value for any argument stays relatively constant for all cardinalities, the maximal values differ sharply, as do the range of values between the extremes.

1.4 Basic definitions regarding pcset similarity functions.

For the present purposes, we will consider a similarity function to be one which is passed two sets⁴⁶ as input and returns a single number which represents whether and/or to what degree the two sets are related using some determined criteria. There are two distinct types of similarity functions: Boolean and real number. A Boolean similarity function returns either 1 or 0 as output (this can be understood as true or false), indicating that the two given sets either are or are not related under the similarity function. This type of function is properly called a *similarity relation*. A real number similarity function returns a real number (frequently between 0 and 1) which represents the degree to which the two sets are related under the similarity function. We will call this type of function a *similarity index*.

Let R be a similarity relation, N be a similarity index and X, Y , and Z be pcsets. The following statements are possible:

$$R(X, Y) = \text{true}; R(Y, Z) = \text{true}; R(X, Z) = \text{false}.$$

$$N(X, Y) = 0.00; N(Y, Z) = 0.57; N(X, Z) = 0.75.$$

The following statements are not possible:

⁴⁶While this study is dedicated to the examination of pcset functions, this definition is not necessarily limited to pcsets. Similarity functions can certainly weigh two sets based upon any well-defined parameter, not just pitch class.

$$R(X, Y) = 0.35$$

$$N(X, Y) = \text{True}.$$

All similarity functions are symmetrical and non-transitive. By symmetrical, we mean that $R(X, Y) = R(Y, X)$ and $N(X, Y) = N(Y, X)$. The lack of transitivity, is illustrated in the first line of the above possible statements; a relation of X to Y and Y to Z does not necessitate or imply a like relation from X to Z . The two statements which are not possible return incorrect value types (i.e., the relation is returning a real number and the index is returning a Boolean value).

The criteria for determining resemblance that pcset similarity functions employ can include pc content, interval-class content, set class content (this can include ic content if one views the 6 ics as the group of #2 SCs), and transformations within and/or between the pcsets. Each of the functions examined in this chapter measure similarity by examining some nCV . The interval-class vector (ICV) is an example of such an array (see section 1.2).⁴⁷

1.5 What criteria should similarity functions fulfill?

Theorists who have posited similarity indices for relating pcsets have, in summarizing the work of others, frequently defined a number of criteria that an “effective” index should meet. These criteria were, quite naturally, reflections of their own ideas regarding how two pcsets might be related. Moreover, some use their criteria to evaluate how other similarity functions compare to their own. In these cases, one could argue that the faults and praises that are sung of prior similarity functions reflect the biases of the evaluators as much as the characteristics of the evaluated functions.

⁴⁷Block and Douthett discuss abstract inclusion vectors of subsets larger than cardinality 2. They label functions which utilize such vectors “flexible similarity” functions. (Block and Douthett 1994, 22.)

Naturally, if one considers one's own similarity function to be the ideal, other functions are certain to fall short of it. Nevertheless, the technique of weighing all other functions against one's own is fruitful in contrasting the different types of similarity functions.

In this dissertation, a number of new indices will be defined, but no new criteria for evaluating previous indices will precede them. The reasons for this are two-fold: (1) the criteria that other theorists have listed (particularly Castrén and Hermann) are quite adequate for our purposes. Rather than create new criteria for evaluation, we will instead review the criteria of others and, in some cases, discuss their motivation. (2) Each of the new indices do not operate under a single philosophy of what should and should not be considered "equivalent" or "similar." We will therefore examine the criteria of others in an effort to explore the motivation of the evaluators.

1.5.1 Isaacson

Eric Isaacson (1990) lists only three criteria for evaluating similarity functions. He believes that similarity functions should:

- 1) provide a distinct value for every pair of sets
- 2) be useful (not just usable) for sets of any size
- 3) provide a wide range of discrete values⁴⁸

The first criterion is satisfied by most similarity functions since, by their very nature, functions return values. Even (Boolean) similarity relations (such as Forte's K or K_h set complexes, Ericksson's ic max-point structure, or Kaplan's K_i subcomplex), where any pair of pcsets is either in or out of the relation, satisfy this condition. We will introduce two similarity indices in chapter 2

⁴⁸Isaacson 1990, 2.

(PSATSIM and PSAT%SIM) which, in fact, do not satisfy this criterion. This is because these similarity functions require dividing by zero in a handful of cases. Since any number divided by zero produces an undefined quotient, there is no distinct value in these cases. A defense of the exceptions to this criterion is presented in section 2.5.

Isaacson's second criterion—that a similarity function must “be useful (not just usable) for sets of any size”—dismisses all of Forte's R_n relations as well as the similarity indices by Teitelbaum, Lord, and Morris's $SIM(X, Y)$ since these do not have a built-in adjustments to account for differences in set cardinality. The indices that we will present in chapter 2 do indeed meet this criterion,⁴⁹ and it is easy to see why one would want to relate sets of different cardinality since it is rare to find a piece of music in which the structural groups are always the same size.

Nevertheless, there are valid arguments against this criterion. Sets of different size naturally have different attributes. There are no SCs larger than a hexachord that exclude any single interval class, for example, yet every trichord contains only three intervals and must therefore exclude at least half of the ics. If one is comparing a cardinality 3 pcset to one of cardinality 7, there are a number of difficulties to be faced. If one favors a similarity function which compares the total (literal or abstract) subset content of the two pcsets, then one is faced with having subsets of cardinal 0, 1, and 2 as the only potential commonalities. In this case, such a function is not appreciably different from one which only compares the ic content of the two pcsets.⁵⁰ If

⁴⁹Except, as we mentioned, for a small handful of exceptions in the cases of PSATSIM and PSAT%SIM. These indices cannot be used to compare sets of cardinality 0, 1, 11, or 12 because division by zero is necessitated.

⁵⁰This is not to suggest that functions which only use ic content are in any way inferior to those which examine the total subset content; it is only meant to indicate that one must make allowances in

one favors a function which examines the number of ways that either one pcset can be transformed onto the other (or the ways in which the elements of each pcset are related internally to the other members of the same pcset), then once again the number of comparisons which come into play is determined exclusively by the size of the smaller pcset.

Additionally, one could argue that there are four set classes which are difficult to relate to others: the empty SC (\emptyset , or 0-0[]), the one-element SC 1-1[0], and their complements, the aggregate (U , or 12-1[0123456789ab]) and the eleven-element SC (11-1[0123456789a]). Each of these SCs are the sole representatives of their cardinality; each therefore has both the maximal and minimal amounts of each ic for any pcset of the same cardinality, making their interval content rather entropic. One might well consider that any pcset of cardinality 2 through 10 should be related equally to any of these four difficult pcsets. How, after all, can any set X (where $1 < \#X < 11$) really be more or less similar or dissimilar to the one-element SC than is any other X ?

Another case against Isaacson's second criterion can be built around the claim that sets of different cardinality are inherently different because they have different ranges of possibility in almost every domain. There are, for example, 41% more tetrachords than trichords and 76% more pentachords than tetrachords. Tetrachords have twice the number of embedded intervals as trichords, etc. How, then, can we reasonably say that sets of different size must be relatable under some common index? Most means of "correcting" the "cardinality problem"⁵¹ (i.e., the problem of comparing pcsets of different

a "total" index for the variable numbers of subsets which are brought into the comparison of different cardinality pcsets.

⁵¹This is the term Richard Hermann uses in "A General Measurement for Similarity Relations: a Heuristic for Constructing or Evaluating Aspects of Possible Musical Grammars." (Hermann 1993)

cardinality) involve simpler adjustments which blur together individual similarities and differences among different object types.⁵²

The last of Isaacson's criteria is, in a sense, too vague to challenge. After all, how many values constitute a "wide range?" He makes it clear in his summary of Forte's R_n relations and K/Kh set complexes that, for him, Boolean relations produce too narrow a range of discrete values. Morris's ASIM index is also found to produce too few values.⁵³ But Isaacson specifies no preferable range of values.

1.5.2 Castrén

Marcus Castrén also provides a listing of "similarity relation evaluation criteria." In his recent dissertation, *RECREL: A Similarity Measure for Set-Classes*, he says that a similarity function should:

- C1) allow comparisons between SCs of different cardinalities
- C2) provide a distinct value for every pair of SCs
- C3) provide a comprehensible scale of values, so that
 - C3.1) all values are commensurable
 - C3.2) the end points are not just some extreme values, but can be meaningfully associated with maximal similarity and dissimilarity.
 - C3.3) the values are integers or other easily manageable numbers
 - C3.4) the degree of discrimination is not too coarse or unrealistically fine
- C4) produce a uniform value for all comparable cases
- C5) observe mutually embeddable subset-classes of all meaningful cardinalities
- C6) observe also the mutually embeddable subset-classes *not* in common between the SCs being compared.⁵⁴

⁵²Either through dividing the differences between the two sets (however they are measured) by the number of objects compared (whether they be subsets, intervals, or elements) or by taking the standard deviation of the set of differences between the two sets.

⁵³Isaacson 1990, 5. Isaacson gives an incorrect definition of Morris's ASIM(X , Y) index, dividing the SIM(X , Y) value by $\#X + \#Y$ rather than $\#ICV(X) + \#ICV(Y)$. While the actual values that Isaacson lists for Morris's ASIM index are therefore incorrect, the number of values produced by ASIM is correct.

⁵⁴Castrén 1994, 18.

Isaacson's criteria 1 and 2 are included in Castrén's longer list as C2 and C1, respectively; Castrén's C3 is an expanded (and better defined) version of Isaacson's criterion 3.⁵⁵

Castrén felt that Isaacson's third criterion was not as meaningful as his first two and believed that it should be more of a recommendation than a condition. Perhaps in an effort to make this "recommendation" more substantial, he offers four elaborations on it.⁵⁶ C3.1 indicates that comparisons across all cardinalities should reflect the same range of values (e.g., a range of possible values from 0 to 100 for all cardinality SC pairs). Castrén explains his rationale for C3.2 by saying that "it is easier to relate 0.75 to 1 or 75 to 100 than, say, 2.73 to 3.64, even though the first figure of each pair is three quarters of the latter in each case."⁵⁷ 3.64, he points out, is the maxima of Isaacson's IcVSIM index.⁵⁸ By "easily manageable values" in C3.3, Castrén is simply expressing his preference for integer values over real numbers. Finally, in C3.4, Castrén says that where real number values are

⁵⁵Interestingly, while Castrén adopts these three criteria, he does not do so without some reproach. In his critique of Isaacson's criteria he states:

In the literature, authors often point to unsatisfactory features of previously presented similarity relations, and point to favorable ones in those they are about to introduce. It is rare to find detailed analysis of minimum conditions which a relation must meet in order to be valid, however.

...

The first criterion rejects a number of relations for not being similarity *measures*. They provide an insufficient degree of discrimination, the yes-or-no type of outcome indicating only whether two SCs enjoy the relation or not. We assume that the second criterion means that an index must produce meaningful results from comparisons between SCs of all cardinalities in order to be useful. What is more, this criterion discourages the use of some measures, as they are designed for SCs of the same cardinality only. (Castrén, 17.)

⁵⁶Castrén mentions that the four parts of criterion 3 are "listed in order of decreasing importance." (Castrén, 18)

⁵⁷Castrén, 20.

⁵⁸Hermann also registers this criticism, but offers a solution, showing how Isaacson's values could be brought into a range of 0.00 to 1.00. (Hermann 1993, 112.)

used, they should be rounded at a reasonable place.⁵⁹ Just as Castrén believes that Isaacson's third criterion are more suggestions than necessities, it seems reasonable to assert that Castrén's four-part third criterion are also meant to be preference rules rather than structural conditions.

Criterion C4 is an interesting critique of indices that use standard deviation, particularly Isaacson's IcVSIM. Castrén correctly points out that two SC pairs can have the same subgroups in common, but the pair that has the more extreme difference in the elements which are *not* in common will be considered more distantly related when using standard deviation.⁶⁰ Castrén's criticism of IcVSIM is apt, particularly in his summary of Isaacson's index, where he points out that using standard deviation leads to an instance of maximal similarity between SCs 3-10 [036] and 6-30 [013679].⁶¹ While, by

⁵⁹This is presumably a criticism of Isaacson, who, in certain cases, rounds the values returned to five significant digits to the right of the decimal point. (See Isaacson 1990, 14-15)

⁶⁰In further explanation, Isaacson's IcVSIM may be defined as follows: X and Y are pcsets; $\text{diffV}(X, Y)$ is a six-argument vector which derives each of its values by $|\text{ICV}(X)_i - \text{ICV}(Y)_i|$. $\text{IcVSIM} = (\text{diffV}(X), \text{diffV}(Y))$. Where σ = the standard deviation. More formally, the standard deviation is derived as follows:

$$\text{IcVSIM}(X, Y) = \sqrt{\frac{\sum_{n=1}^6 (\text{diffV}(X, Y)_n - \text{AVG}(\text{diffV}(X, Y)))^2}{6}}$$

$$\text{AVG}(\text{diffV}(X, Y)) = \frac{\sum_{n=1}^6 (\text{diffV}_n(X, Y))}{6}$$

$$\text{where } \text{AVG}(\text{diffV}(X, Y)) = \frac{\sum_{n=1}^6 (\text{diffV}_n(X, Y))}{6}$$

⁶¹In other words, of all trichord/hexachord pairs, 3-10/6-39 is the most "similar." In fact, there are no SC pairs of any cardinality which are closer, including any SC compared to itself. $\text{IcVSIM}(3-10, 6-30)$ is calculated as follows:

$$\begin{array}{rcl} \text{ICV}(3-10) & = & \langle 0 \ 0 \ 2 \ 0 \ 0 \ 1 \ \rangle \\ \text{ICV}(6-30) & = & \langle 2 \ 2 \ 4 \ 2 \ 2 \ 3 \ \rangle \\ \text{diffV}(3-10, 6-30) & < & \langle 2 \ 2 \ 2 \ 2 \ 2 \ 2 \ \rangle \\ \text{AVG}(\text{diffV}(3-10, 6-30)) & = & 2 \end{array}$$

$$= \sqrt{\frac{(2-2)^2 + (2-2)^2 + (2-2)^2 + (2-2)^2 + (2-2)^2 + (2-2)^2}{6}} = \sqrt{\frac{0}{6}} = 0.00$$

See Castrén, 64-67.

virtue of their maximal inclusion of ic6 (for their particular cardinalities), these two sets are no doubt similar, the hexachord does not have the maximal amount of ic3 or the minimal amount of ic1, as does the trichord. It is difficult, therefore, to imagine a rationale for considering them maximally similar.

Criterion C5 expresses Castrén's notion that more than just interval classes should be compared. He is careful to state that "the expression 'all meaningful cardinalities' is an intentionally ambiguous one, since participating subset-class cardinalities may be selected differently in different indices."⁶² Through this criterion, he is privileging indices such as Rahn's TMEMB or ATMEMB and Lewin's REL, which examine the entire subset content of each set in the comparison. Castrén states that "the more points of reference we have between the SCs to be compared, the more accurately we can hope to demonstrate structural similarities or differences between them."⁶³ He then mentions that it is necessary to compare subset classes larger than cardinal 2 because "we want to discriminate between inversionally related and Z-related SCs. They are identical from the point of view of ic content-based similarity indices."⁶⁴

While indices that consider all subset-classes certainly take more information into account than those that compare only interval classes, I believe that this is not always helpful in obtaining a satisfactory comparison. Because indices that consider all applicable subset-classes (< or to the smaller cardinality set being compared) deal with a variable number of arguments, comparisons among different sized pairs of set classes will

⁶²Castrén, 25.

⁶³Castrén, 25.

⁶⁴Castrén, 25. Castrén's own index, RECREL, works within SG(1).

necessarily yield different gradations of values. Indices that compare only interval classes, by contrast, always deal with a sort of “lowest common denominator” and are therefore much more likely to yield commensurable values in comparisons between sets of different cardinalities.

I do not suggest a wholesale dismissal of such indices (and Castrén’s method of utilizing all applicable subsets in his RECREL index is quite useful), but neither do I believe that they are necessarily “more accurate” or even more likely to produce “intuitive” results.⁶⁵ Similarity functions that examine the total subset content do not make a distinction between cardinality 2 SCs and all other cardinalities of SCs. Perhaps, however, this is a valuable distinction. After all, pcsets (including SCs) are entities which can be found embedded in other pcsets. Ics, however, are distances which separate the elements of a pcset, not actual embedded entities. If one holds this distinction firm, it does not necessarily follow that one can generalize from ics to trichords and larger set classes.

While it is true that examination of subset-classes larger than cardinal 2 allows one to discern between SG(1) inversionally related SCs and Z-related SCs in SG(1) and SG(2),⁶⁶ there is no reason to believe that working within one set group is inherently superior to working within another. To Castrén’s credit, he acknowledges this, stating the validity of a particular set group depends “only on how well its generality or particularity fulfills the needs of a given task.”⁶⁷ He nevertheless defends his choice of using only

⁶⁵ Three comparable indices, one which uses only ics and the others which consider all mutually embedded SCs, are presented in sections 2.3 and 2.10 of this dissertation.

⁶⁶ Assuming, of course, that the functions one employs use SG(1) SCs as data (Castrén does do this).

⁶⁷ Castrén, 31.

transpositional equivalence in both his own index, RECREL, and in evaluating other people's indices:

We want inversionally related SCs to be independent objects for the simple reason that we experience them to be different. More important still, we experience that the degree of similarity between pairs of inversionally related classes is not a constant. Some inversionally related SCs may give the impression of being extremely close to each other, while others can seem more distant.⁶⁸

The last of Castrén's criteria forms the foundation of his index. While we will summarize his index later in this chapter, the basis of C6 is that all subset-classes of the set classes being compared should also be compared to each other to determine their degree of similarity. This basically sets up a recursive process under which all non-mutually embedded SCs of the same cardinality are compared to each other and their degree of similarity is taken into consideration when calculating the value that the similarity index ultimately returns.

1.5.3 Hermann

Richard Hermann does not present an exhaustive list of criteria that, as Castrén writes, "point to unsatisfactory features of previously presented similarity relations."⁶⁹ Rather, he defines several categories of "resemblance" theories, based upon mathematically rigorous concepts.⁷⁰ These categories allow Hermann to group the theories unambiguously into a number of

⁶⁸Castrén, 32. It seems a bit peculiar that Castrén would, after acknowledging that a variety of SG-types are useful in different situations, defend a single SG as preferable.

⁶⁹Castrén, 17.

⁷⁰Hermann 1993. His classification scheme is presented in pp. 1-15 and his summaries of "published theories of resemblance" follow on pp. 15-119.

different categories, according to type. In every case, the functions that he summarizes either fit or don't fit into each category.

There are six categories that Hermann uses for classifying theories of resemblance: resemblance criteria, measurement types, collection structures, element types, spatial types, and relation types. They are elaborated at the end of the first section of chapter 1 of his dissertation (“A General Measurement for Similarity Relations: a Heuristic for Constructing or Evaluating Aspects of Possible Musical Grammars.”) in a chart that is reproduced as our figure 1.3. He uses these six categories when summarizing each of the indices that he cites in his dissertation.⁷¹

Hermann's descriptions of other similarity functions are encyclopedic; his language is clear and his summaries are accurate and relatively value neutral.⁷² In the first chapter of his dissertation, Hermann lists only three “problems” that occur among similarity functions. They are: “the cardinality problem,” which occurs when a similarity function is only designed to compare SCs of the same cardinality;⁷³ “the scaling problem,” which occurs when the range of values output by the similarity function is not constant among different sized SC pairs;⁷⁴ and “the z SC problem,” which occurs when the similarity function cannot discern z-related SCs.⁷⁵ As is the case with Castrén's and Isaacson's evaluation criteria, Hermann's “problems” are a matter of

⁷¹For example, Hermann's classification of Forte's R_p relation reads as follows: “The resemblance criterion is that of inclusion; the element type is the pc (object) for strong R_p and ic (interval) for the others; the measurement type is by difference; the spatial type is linear finite for strong R_p and linear finite cyclic for the others and weak R_p , the collection structure is the pcset (unordered) for strong R_p and the SC (unordered) for the others and the relation type is similarity.” (Hermann 1993, 23.)

⁷²Castrén's summaries are also very clear, but his descriptions also involve an evaluation of each index or relation using his own criteria (necessitating some degree of bias).

⁷³This is essentially the same as Isaacson's criterion 2 and Castrén C1.

⁷⁴This is essentially the same as Castrén's C3.1.

⁷⁵This is the motivation behind Castrén's C5.

perspective; they are not necessarily universal goals that all similarity functions do or should strive to meet.

In addition to providing an evaluation of each similarity function, Hermann also traces their influence, where applicable, on later theories and he also provides citations of any published reviews of the function at hand. The inclusion of additional sources and Hermann's concise and rigorous descriptions make his summaries very valuable and I recommend them to the reader.

1.6 Summary of some referential similarity relations/indices

As mentioned in section 1.5, the recent dissertations of Marcus Castrén and Richard Hermann do an admirable job of summarizing the various similarity functions that have been published in the last thirty-seven years. While we will provide descriptions of the two similarity functions that most clearly influence our own indices, including Castrén's RECREL (particularly since that has not previously been summarized), we will forgo a complete description of all prior resemblance theories (to borrow Hermann's more inclusive term). Such an extended glossary would largely recreate Castrén's and Hermann's efforts.

Instead, in figure 1.4, we provide a list of indices that have been evaluated by Hermann and Castrén and page references for the reader who wants to examine their summaries. A shorter summary of similarity functions that use the ICV as data can be found in Isaacson, 1990. However, that article is much shorter, less complete, and is not entirely free of error.

Figure 1.5 places the similarity functions listed in the previous table into three categories. The first category is set complexes and subcomplexes. Each of these complexes represents a similarity relation (Boolean) that includes a finite number of pcsets that are drawn together under some criteria. The second category is ICV-based similarity functions. These are either indices or relations (the function type is given in the third column) that use the interval-class vector of a pcset as their only data for determining resemblance. The third category is subset-based similarity functions. These are either indices or relations that compare the subset content (larger than cardinal 2) of two sets as data for determining resemblance.

1.6.1 Morris's SIM and ASIM indices

$SIM(X, Y)$, a similarity index on pcsets, totals the absolute values of the differences in each place of the two pcsets ic vectors. Formally, we can define $SIM(X, Y)$ as follows:

$$SIM(X, Y) = \sum_{n=1}^6 |ICV(X)_n - ICV(Y)_n|$$

where $ICV(X)_n$ represents the n -th place in the interval-class vector of pcset X .

Morris's $ASIM(X, Y)$ index was the first pcset similarity function that allowed one to compare pcsets of different cardinality in a meaningful manner.

$ASIM(X, Y)$ is an adjustment on $SIM(X, Y)$, in which the SIM total is divided by the sum of the arguments in both ic vectors, returning a value between zero and one. The smaller the returned value, the more similar two set classes are said to be.⁷⁶ Formally, we can define $ASIM(X, Y)$ as follows:

⁷⁶Because two set classes are said to be less similar as the returned value approaches 1, $ASIM$ is more properly said to be a "dissimilarity index" (as are most of the indices discussed herein).

$$\text{ASIM}(X, Y) = \frac{\text{SIM}(X, Y)}{\#\text{ICV}(X) + \#\text{ICV}(Y)}$$

where $\#\text{ICV}(X)$ —the cardinality of the interval-class vector of X is the sum of the $\text{ICV}(X)$ entries.

1.6.2 Castrén’s RECREL index

Marcus Castrén’s RECREL (or **recursive relation**) is also a similarity index on pcsets. RECREL returns a number between 0 and 100; the lower the number, the more similar the two pcsets are said to be.⁷⁷ Basically, $\text{RECREL}(X, Y)$ is a weighted composite index, derived from repeated iterations of an internal similarity function. These iterations compare not only X to Y , but also all of the abstract subsets of X and Y to each other. In order to define RECREL, we must first define the internal similarity function $\%REL_n$.

1.6.2.1 The difference vector (DV)

The difference vector (DV) of two pcsets is an array which contains the absolute values of the differences in the corresponding places of two vectors of the same size and type. For example, the difference vector comparing the ICVs of 5-35[02479] and 6-20[014589] is shown in figure 1.6. We could, for example, revise our definition of Morris’s SIM as follows:

$$\text{SIM}(X, Y) = \frac{1}{n} \sum_{n=1}^6 \text{DV}(\text{ICV}(X)_n, \text{ICV}(Y)_n)$$

or even:

⁷⁷This is comparable to ASIM, which returns a value between 0 and 1.

$$\text{SIM}(X, Y) = \#DV (\text{ICV}(X), \text{ICV}(Y))$$

where #DV is the cardinality of the difference vector (i.e., the sum of its arguments).

By contrast, Castrén's difference vector is non-symmetrical. He does not take the absolute values of the differences in each vector argument, but rather subtracts the arguments of vector A from those of B and shows only those differences greater than zero. He then repeats the process, subtracting the arguments of vector B from those of vector A. The sum of all the arguments in both Castrén DVs is the same as the sum of the single DV used by Isaacson. Castrén takes this longer route in order to show which vector elements are exclusive to set A and which are exclusive to set B. Figure 1.7 shows Castrén's DV(ICV(5-35), ICV(6-20))

1.6.2.2 %REL_n

Castrén's similarity index %REL_n(X, Y) compares the %nCV of two pcsets X and Y.⁷⁸ It is defined as follows:

$$\%REL_n(X, Y) = \frac{\#DV (\%nCV(X), \%nCV(Y))}{2}$$

In essence it is a version of Morris's ASIM which uses the %nCV instead of the ICV as data. The numerator is a SIM-like comparison which sums the absolute values of the differences in the corresponding %nCV arguments. Since the arguments in each %nCV total 100,⁷⁹ the denominator represents a scaled version of the combined cardinality of both %nCVs (reinforcing the

⁷⁸See section 1.2.2 for a review of Castrén's %nCV.

⁷⁹Because Castrén's multiplies each value by 100.

comparison to Morris's ASIM). Figure 1.8 provides an example of a %REL₂ comparison between 5-35[02479] and 6-20[014589]. The value 50 returned by %REL₂ indicates that 50% of the %2CVs of the two SCs are different.

1.6.2.3 RECREL in summary

Given two hypothetical pcsets 6-A and 6-B, RECREL uses %REL_{*n*} to measure the subset content of 6-A and 6-B as well as the subset content of each of 6-A's and 6-B's subsets. RECREL begins by calculating %REL₅(6-A, 6-B). If 6-A embeds 5-A, 5-B, 5-C, and 5-D; and 6-B embeds 5-C, 5-D, 5-E, and 5-F, RECREL then compares each of the subsets which are *not* mutually embedded to each other using %REL₄(i.e., 5-A and 5-B are both compared to 5-E and 5-F). The non-mutually embedded tetrachordal SCs are then compared using %REL₃. Each of the non-mutually embedded trichordal SCs are then compared using %REL₂ (dyad classes are the smallest SCs employed by %REL_{*n*}).

The series of comparisons described above comprises only a single branch of the RECREL function. It is called branch 5 because it is initiated by a single %REL₅ comparison performed on the two hexachords being compared by RECREL. Branch 4 is a similar string of comparisons, initiated by calculating %REL₄(6-A, 6-B) followed by %REL₃ on the tetrachordal subsets of 6-A and 6-B, then %REL₂ on the trichordal subsets of the tetrachordal subsets of 6-A and 6-B. Similarly, branch 3 is a string of comparisons initiated by calculating %REL₃(6-A, 6-B) then %REL₂ on the trichordal subsets of 6-A and 6-B. Finally, branch 2—the lowest branch—is performed by calculating %REL₂(6-A, 6-B).

Each of the $\%REL_n$ values is weighted using what Castrén calls a proportioned weight, derived from a weighted difference vector (WDV). To derive a weighted difference vector, one divides the arguments of a difference vector by the sum of its arguments (the vector cardinality) then that value is multiplied by 100 to maintain an integer. More formally stated:

$$WDV_i(V(X), V(Y)) = \frac{DV_i(V(X), V(Y))}{\#DV(V(X), V(Y))}$$

The DV and WDV of the $\%2CV$ s of 5-35, 6-20 are given in figure 1.9.

When two subsets of 6-A and 6-B are compared using $\%REL_5$ in branch 5 of RECREL, the difference vectors used in the $\%REL_5$ comparison are converted to weighted difference vectors, the arguments of which are used to weight the various $\%REL_4$ comparisons that utilize the #5 subsets of 6-A and 6-B. The weighting of a $\%REL_n$ comparison is arrived at by multiplying the product of the two relevant arguments in the weighted difference vectors of the prior $\%REL_{n+1}$ comparison by the $\%REL_n$ value. More formally:

$$\text{weight} = \frac{WDV_X * WDV_Y}{100}$$

where WDV_X is the X -th place in a $WDV(A, B)$ comparison

$$\text{weighted } \%REL_n(X, Y) \text{ value} = \frac{\text{weight} * \%REL_n(X, Y)}{100}$$

Each of the weighted $\%REL_n$ values in a single level of the hierarchy are added together (we will call this family of comparisons TEST in the formalization below), then their sum is multiplied by the weighted $\%REL_n$ value of the next level up. This process is repeated recursively until finally a single value for a branch is determined. More formally, the value of a single RECREL branch can be described as follows:

$$\text{RECRELbranch}_n(X, Y) = \frac{n}{a=2 \text{ test TEST}} \left(\text{weight}_{X,Y} \% \text{REL}_a(X, Y) \right)$$

This process is repeated for each RECREL branch. When all $\% \text{REL}_n$ comparisons have been performed, weighted, added together, and multiplied by the weighted $\% \text{REL}_n$ values which precede them hierarchically, the RECREL branch values are averaged together to reach a final RECREL value for the two pcsets being compared. So, finally, we can say that:

$$\text{RECREL}(X, Y) = \text{average}(\text{RECRELbranch}_2 \dots \text{RECRELbranch}_n).$$

Marcus Castrén gives a very good RECREL tutorial in chapter 4 of his dissertation,⁸⁰ “RECREL: A Similarity Measure for Set Classes.” In that chapter, he goes through two RECREL comparisons step by step, facilitating the sometimes difficult concepts involved. His exposition of the index goes into far more depth than we can reasonably provide in our summary.

⁸⁰Castrén, 101-125.

Chapter 2

Saturation vectors and associated similarity indices

In this chapter, we will define a body of analytic tools for use in measuring how closely a pair of pcset classes resemble each other. We recognize that there is no single best way to measure resemblance between two pcsets; accordingly, we will introduce indices which utilize a variety of approaches, including those which compare the interval-class content, subset-class content, and the presence of cyclic adjacencies in each pcset. We will begin by explaining the concepts of abstract pcset inclusion and minimal and maximal saturation of interval classes and pcset classes. We will then define a number of “saturation vectors” which compare the quantity and types of embedded interval- or subset-classes to what is possible in any pcset of the same cardinality. Each saturation vector will, in turn, be used as our means for comparing two pcsets to each other using a similarity index. A summary of these new vector types and similarity indices appears at the end of this chapter.

2.1 Relative abstract pcset inclusion

In section 1.1 we studied the concept of relative abstract pcset inclusion using two functions, $AS(/X/, /Y/)$ and $EMB(/X/, /Y/)$. In this chapter we will introduce a new function, $SATEMB(/X/, /Y/)$, which returns two arguments that reflect a comparison between $EMB(/X/, /Y/)$ and the largest and smallest $EMB(/X/, /Y/)$ values for all sets of $\#Y$. We will call $SATEMB$'s output the degree to which $/Y/$ is saturated with $/X/$. The function itself shall be explained

in greater detail later in this chapter.⁸¹ Just as abstract inclusion (n -class) vectors such as the interval-class vector can be derived by performing $\text{EMB}(/X/, /Y/)$ for each distinct $\text{SC}(X)$ where $\#X = n$, we will similarly create n -class “saturation vectors” by concatenating $\text{SATEMB}(/X/, /Y/)$ values for each $/X/$ where $\#X = n$. Wherever the places in a vector $V(X)$ show the number of SCs A_1 through A_n that are embedded in pcset X , a saturation vector $\text{SATV}(X)$ can be constructed to show how the quantity of subsets A_1 through A_n fall into the range of what is possible given any pcset of $\#X$. Several n -class (largely interval-class) saturation vectors will be introduced in this chapter.

One can construct saturation vectors which measure the degree to which different collections of SCs ($X_1 \dots X_n$) are maximally and/or minimally embedded in pcset Y in a variety of musical pitch spaces, using different means of defining equivalence. A SC collection might contain all SCs of a particular cardinality (e.g., interval classes, trichord classes,...) or might (or might not) share some other property (e.g., all transpositionally-symmetrical SCs). Certain SCs might even be weighted in the vector (this concept shall be demonstrated later in this chapter). While the principal focus will be on interval-class vectors, it should be implicit that the same saturation functions that we use to determine relative ic embedding can also determine relative embedding of any other size subset class. Additionally, while T_nI -based n -class vectors ($nCVs$) will form the basis for the vectors demonstrated herein, one can certainly use any of Morris’s $SG(n)$ vectors for the same purpose.

⁸¹The function SATEMB can also be seen as a special case of the membership function (μ) used in “fuzzy” set theory. Where inclusion in classical (“crisp”) set theory deals exclusively with whether or not element x is included in set A , fuzzy set theory allows a statement such as element x (or $\text{SC}(X)$) is only partially a member of set A . While SATEMB does not model partial membership, it does model partial “saturation,” that is, the degree to which $\text{SC}(X)$ is maximally and/or minimally embedded in pcset A , given what is possible in any pcset of $\#A$.

2.2 Minimal and Maximal Saturation of Interval Classes

In section 2.3, we will introduce a variation on the interval-class vector that compares the amount of each ic that separates members of a particular set class with the minimum and maximum such values (ICV arguments) found in any pcset of the same cardinality. Before we can discuss this vector, however, it will be useful to examine how one calculates the minimal and maximal amounts of each ic in any pcset of a particular cardinality. One could, of course, simply examine the ICVs of all sets of a particular cardinality and, by inspection, keep track of the values at each extreme. But one can derive this information more directly, and with a greater appreciation of the inherent boundaries of pitch-class space, through an examination of the properties of interval cycles. The relationship between our systematic discussion of interval cycles in pc space and minimal and maximal saturation of interval classes shall become apparent as we progress.

An i -cycle (where i is a variable which represents any interval in standard 12-pc space) is defined as a closed and finite ordered collection of pcs whose adjacent elements map onto each other with a constant operation.⁸² The members of an i -cycle are defined as $(x+i, x+i^2, x+i^3, \dots, x+i^p = x)$ where p is the periodicity of the i -cycle. Because each of the i -cycles evenly partitions the aggregate (U), the number of i -cycles for any ic i (we call this m) = $12/p$. The complete i -cycles are shown in figure 2.1.

We will define an i -set X as a pcset which meets any of the following three conditions: (1) it is the union of one or more completed i -cycle(s), (2) it is a set of i -cycle adjacencies where cardinality prohibits completing the cycle

⁸²In abstract algebra, this is called a “ring.”

(we will call this an i -cycle fragment), or (3) it is a combination of the above where only one i -cycle fragment is permitted. An i -cyclic fragment Y is defined as a group of pcs where each $y \in Y$ is an adjacency within a single incomplete i -cycle.⁸³ Examples of the first condition above are $\{0, 3, 6, 9\}$ and $\{0, 1, 3, 4, 6, 7, 9, a\}$. The first pcset is comprised of a completed 3 cycle; the second is a concatenation of two completed 3 cycles. Examples of the second condition are $\{0, 3, 6\}$ and $\{0, 2, 4, 5, 7, 9\}$. The first pcset is too small to contain a completed 3 cycle, but it does contain exclusively 3-cyclic adjacencies; the second pcset is a cyclically-adjacent subset of a completed 5 cycle— $\{4, 9, 2, 7, 0, 5\}$. Examples of the third condition are $\{0, 1, 3, 4, 6, 9\}$ and $\{0, 1, 2, 4, 6, 8, a\}$. The first pcset is entirely comprised of a completed 3 cycle $\{0, 3, 6, 9\}$ and a cyclically-adjacent subset of a 3 cycle $\{1, 4\}$; the second pcset is entirely comprised of a completed 2 cycle $\{0, 2, 4, 6, 8, a\}$ and a cyclically-adjacent subset of a 2 cycle $\{1\}$. It does not matter that the last cyclic subset contains only one member. Because the original pcset is a septachord and a completed 2 cycle contains six distinct members, it is impossible to add another ic2 with only one additional member since it takes two elements to form an interval. $\{0, 1, 2, 3, 4, 6, 8\}$ is not a 2-cyclic set because it is comprised of two 2-cycle fragments: $\{0, 2, 4, 6, 8\}$ and $\{1, 3\}$ (only one is ever allowed in a cyclic set).

The relation between the i -cycles and maximal and minimal saturation will stem from the following theorem:

Theorem 2.2.1: All instances of ic i are found between adjacent pcs within i -cycles.

⁸³For example, $\{0, 2, 4, 6\}$ is a 2-cyclic fragment and $\{0, 2, 6, 8\}$ is not. While each element in both sets belongs to a single 2-cycle, all the elements in the second set do not form a single unbroken fragment of a 2-cycle (pcs 2 and 6 are not adjacent).

Corollary 2.2.1: Pitch classes from different i -cycles are not related by i .

From theorem 2.2.1, we can conclude that any i -set X will always contain the maximal possible amount of i in any set of cardinal X (which we shall abbreviate as c). Because the various i -cycles that partition U have different periodicities (p), the minimal and maximal numbers of each i are not uniform within a particular cardinality. For example, $\{0, 4, 8\}$ and $\{0, 3, 6\}$ are both i -sets (where $i = 4$ and 3 respectively) and therefore contain the maximal amount of their respective cyclic generators. Despite this similarity, the former contains 3 i 4s, while the latter contains only two i 3s. This is because a closed i -cycle of cardinal c (for all i except 6) contains two more i s than an i -set of $c-1$. In the cases of $\{0, 4, 8\}$ and $\{0, 3, 6\}$, the former pcset is a closed 4 cycle, while the latter is a 3-cycle fragment. Both have the maximal possible amount of their generating interval for their cardinality: the difference is that their shared cardinality has a different range of possible values for i 3 and i 4.

The value equal to the maximal amount of i in any pcset of c shall be called $\max(c, i)$. We can now formally describe the algorithm for deriving this value. All variables will carry integers and, accordingly, all division will be integer division, which truncates the remainder, if any (e.g., $\frac{7}{2} = 3$). Similarly, “mod” is a binary operation which returns the remainder of integer division (e.g., $7 \bmod 2 = 1$).⁸⁴

⁸⁴To review, the variables that we are using are defined as follows:

Let s be a temporary variable that we shall use *en route* to deriving $\min(c, i)$ and $\max(c, i)$.

Let c = the cardinality of pcset X .

Let p = the periodicity of interval cycle i ($p = 12/m$).

Let m = the number of distinct cycles of i in U ($m = 12/p$).

Derivation of $\max(c, i)$:⁸⁵

$$(c \bmod p = 0) \quad s = \frac{c}{p} p$$

$$\sim (c \bmod p \neq 0) \quad s = \frac{c}{p} p + (c \bmod p) - 1$$

$$(i = 6) \quad s = \frac{s}{2}$$

$$\max(c, i) = s$$

In less formal language: If the periodicity (p) of an i -cycle divides evenly into c , then s = the product of p and the integer quotient of c and p . This first statement is a test of i -set condition #1. If the left side of the statement is true, then the completed i -cycle or combination of completed i -cycles will yield the maximum number of i in any pcset of cardinality c . That number will equal the product of the periodicity of the i -cycle and the number of completed i -cycles that fit into c (this is represented by $\frac{c}{p}$), except in the case where $i = 6$.

In the last case, s must be divided in half to account for the fact that there are only half as many distinct i 6s in U as any of the other i cs. This “fix” is provided on line 3 of the above equation. Line 2 of the equation describes conditions 2 and 3 where completed i -cycles cannot be evenly partitioned into a pcset of c . In this case, we take the number of i from any potentially-completed cycles $\frac{c}{p} p$ and add to it the number of i s found in a

⁸⁵For those readers who are more comfortable examining this structure in C code, it can be encoded as follows:

```

int s, c, p;           /* see above definitions      */
s = (c / p) * p;
if (c % p != 0)        /* if c does not divide evenly into p */
    s = s + (c % p) - 1;
if (n == 6)
    s = s / 2;          /* s is the max(c, i) value      */

```

fragment the size of the remainder of $\frac{c}{p}$. The number of *ic is* in any *i*-cycle fragment equals its cardinality minus 1 (since it is not a closed cycle, the “wraparound” interval must be subtracted).

The minimal values of *i* in any pcset of *c* ($\min(c, i)$) can also be formally described using these parameters. $\min(c, i)$ is not always 0 but is variable, dependent on the periodicity of the *i*-cycle and the cardinality of the pcset. Before we present the entire equation for $\min(c, i)$ let us state the conditions where $\min(c, i) = 0$:

$$c \geq \frac{p}{2} m \quad \min(c, i) = 0$$

As we stated in the corollary to theorem 2.2.1, “there are no *i*-cyclic non-adjacencies or elements in different *i*-cycles that are separated by *ic i*.” This is the motivation for the first half of the above formula. If we remove every other element of a particular *i*-cycle (thus removing all instances of *ic i*) the length of what remains is, quite obviously, $\frac{p}{2}$. We then multiply the length of our half-cycle by the number of *i*-cycles to determine the maximum cardinality with no embedded *ic is*. The values of $\frac{p}{2} m$ for all *p* and *m* are given in figure 2.2.

With this condition in place, we can now compose an algorithm to determine $\min(c, i)$ values for all *c* and *i*:

Derivation of $\min(c, i)$:⁸⁶

$$s = c - \frac{p}{2} m$$

$$\frac{p}{2} > 1 \quad \frac{p}{2} m - 2 < c \quad s = (s - 2) - 12 - \frac{p}{2} m - 2$$

$$(s < 0) \quad s = 0$$

$$\min(c, i) = s$$

In some cases, $c - \frac{p}{2} m$ will yield a negative number. Since it is impossible for members of any pcset X to be separated by a negative number of i , the third line of the formula maps all negative values to 0. In most cases where $c = \frac{p}{2} m$, the pcset X which has the minimum number of i will be optimally fragmented with respect to the i -cycle. This means that it is impossible to add a single pc to pcset X without adding two i s. As we have already illustrated, $\{0, 2, 4, 6, 8, a\}$ is optimally fragmented with respect to a 1-cycle. Any (non-redundant) pc that we add to that pcset will add two i s (yet the resultant seven-note pcset will still have the fewest i s in any septachord). Therefore, $\min(c, i)$ is incremented by 2 for every c higher than $\frac{p}{2} m$.

The only exceptions to this rule arise when deriving the minimal amounts of ic_6 and ic_4 where $\min(c, i) = 0$. Because it is impossible to add one note to

⁸⁶Once again, for those readers who are more comfortable examining this structure in C code, it can be encoded as follows:

```

int s, c, p, m;           /* see above definitions */
s = c - (p / 2) * m;
if (((p / 2) > 1) || (((p / 2) * m) * 2) < c) /* "||" is logical or */
    s = (s * 2) - (12 - ((p / 2) * m) * 2);
if (s < 0)
    s = 0;                 /* s is the min(c, i) value */

```

a set and consequently add two ic6s, we only increment $\min(c, 6)$ by 1 for every c higher than $\frac{P}{2} m$. The case of $\min(c, 4)$ is rather complex. Since all completed 4-cycles have periodicity 3, it is impossible to have two non-adjacent elements. Thus, a set's ic4 content can only be 0 when all the elements of the pcset are members of the four different 4-cycles. Unlike maximally fragmented 1-, 2-, 3-, and 5-cycles, any pc added to a maximally fragmented 4-cycle will only add one ic4 to the set's ic content. This is true for four additional pcs ($c = 5 \dots 8$). An eight-note pcset that minimally includes ic4 will have two pcs from each of the four 4-cycles. Any note added to it will complete a 4-cycle and therefore add two more ic4s. Thus, $\min(c, i)$ is incremented by two when $c = 9 \dots 12$. The conditions required by ic4 and ic6 are handled in line 2 of the above formula.

Figure 2.3 lists all the $\min(c, i)$ and $\max(c, i)$ values. These values will be referred to throughout this dissertation and particularly in this chapter. For example, $\min_{6,4}$ —the minimal amount of ic4 in any pcset of #6—carries the value 2 and $\max_{6,4}$ carries the value 6.

In some cases, it is possible to have more than one i -set (for a given i) in a particular cardinality. This is, in part, because we have not specified an intervallic distance that must separate the multiple completed cycles (under condition 2) or completed cycle(s) and cyclic fragment (under condition 3). The following hexachords are all 6-cyclic sets:

$$\{0, 1, 2, 6, 7, 8\}$$

$$\{0, 1, 3, 6, 7, 9\}$$

$$\{0, 2, 4, 6, 8, a\}$$

All three contain three ic6s, but they have very different properties in other

respects. The first and third are all-combinatorial sets; the second is not. The third pcset is also a completed 2 cycle and the intersection of two completed 4 cycles. Thus, we see that it is possible for a pcset to satisfy a condition of cyclic set membership for more than a single interval class. Such a pcset will be called a “multiply-cyclic set.” There are only three varieties of multiply cyclic sets: 1) those which are products of all six i -cycles, 2) those which are products of i -cycles 2, 4, and 6, and 3) those which are products of i -cycles 3 and 6. All multiply cyclic sets are therefore 6-cycle sets (the converse is not true).

Figure 2.4 contains all i -sets in 12-pc space. The first column provides the cardinality of the cyclic set class, the second column lists the Forte number for each set class, the third column lists the prime form of each set class (under T_nI equivalence), the fourth column, labeled “type,” states the condition under which each pcset qualifies as a cyclic set (“C” = completed cycle or concatenation of completed cycles; “F” = cyclic fragment; “CF” = completed cycle or cycles plus one cyclic fragment), the fifth column identifies which sets are multiply cyclic (“1” = cyclic set for all interval class cycles; “2” = ic2, ic4, and ic6 cyclic set; “3” = ic3 and ic6 cyclic set), and the remainder of the columns show the interval class content (interval-class vector).⁸⁷

We have already demonstrated that every i -set X embeds the maximal number of ic i in any pcset of $\#X$. However, an interesting by-product of this maximal saturation is that *most* cyclic sets also include the minimal amount for some other ic or ics (that is to say that a pcset which has the maximum number of i -cyclic adjacencies will usually necessitate that a different i -cycle be

⁸⁷The collection of multiply-cyclic set classes also appears in Ericksson 1986 as “maxgroup 2” (2/4/6-cyclic sets) and “maxgroup 3” (3/6-cyclic sets), though Ericksson does not discuss their cyclic properties. (Ericksson 1986, 97-99.)

maximally fragmented). The italicized “most” becomes a more definitive “all” if a single condition is added (we will call this the *maximal cyclic fragmentation condition*): an i -set will not only feature maximal values for ic_i , but also minimal values for some other ic_j if, after one cycle has been completed in the generation of the pcset, a new cycle is begun at $T_1, T_B(T_{-1}), T_5$, or $T_7(T_{-5})$ of some pc in the completed cycle. All 1-, 2-, 3-, and 5-cyclic sets (those i -cycles where $\frac{p}{2} > 1$) will meet this condition. In the case of the 1- and 5-cycles, completion is only attained with the entire aggregate of pcs; in the case of 2- and 3-cycles, there are no pcs which do not lie within a semitone of some element of a completed 2- or 3-cycle. In effect, this condition sifts out type 2 and type 3 multiply-cyclic sets from the group of 4- and 6-cyclic sets, assigning them sole status as either 2- or 3-cyclic sets.

The above condition is true of all i -sets, regardless of their cardinality (e.g., any 2-cycle pcset X will have the minimal amount of ic_1, ic_3 , and ic_5 for any pcset of $\#X$). With the maximal cyclic fragmentation condition in place, we can now generalize which ics are maximally fragmented (minimally included) in the six distinct i -cycles. Figure 2.5 presents the information in the form of a table. Where, for example, “Min” is found in a location in the table, all i -sets (where $i = \text{row}$) will generate the $\min(c, i)$ amount of ic_i (where $i = \text{column}$). Places where “Min” is in parentheses are true only when the maximal cyclic fragmentation condition is satisfied. Similarly, where “Max” appears in the table, all i -sets will generate the $\max(c, i)$ amount of ic_i . It should be noted that while all sets that are *maximally* saturated with a particular interval class are cyclic sets generable from the same i -cycle, not all sets that are *minimally* saturated with a particular interval class are i -sets (e.g.,

$\{0, 1, 3, 4, 5, 8\}$ and $\{0, 2, 3, 4, 5, 7\}$ are both minimally saturated with ic_6 , yet neither is a cyclic set).

Figure 2.5 yields some interesting patterns which will not only impact how we interpret the data from the interval-class-based vector proposed in the next section, but should also effect how we interpret data from any ICV-based similarity relation. We can, for example, make the following observations about particular families of pcsets: any time that one finds a maximal amount of ic_1 or ic_5 , a minimal amount of ic_6 and ic_4 is guaranteed as a by-product; it is impossible to maximize ic_3 content without also maximizing ic_6 content—but inverse of this statement is not true; ic_4 and ic_6 are either maximized or minimized in every cyclic set, but while the six ic cycles yield minimal ic_6 content equally as often as maximal ic_6 content, there is twice the likelihood of finding minimal ic_4 content as maximal ic_4 content.

As mentioned earlier, it might have been simpler to survey all set classes and make note of which in each cardinality had the largest and smallest numbers of each ic . But the approach we have taken has been more revealing about the interplay among ics , the cardinalities of pcsets, and the properties associated with minimal and maximal interval-class saturation. The generalizations that we have drawn will be integral in our exposition of the cyclic saturation vector later in this chapter.

2.3 The Interval-Class Saturation Vector (SATV)

Having just explained the idea of relative interval-class saturation, we can now define a new vector-class which compares the arguments of an interval-class vector of a set class X to the minimal and maximal possible

values in each place of the vector for any pcset of $\#X$. $SATV(X)$ is a dual six-argument vector: each place in the first (A) row of the vector carries the designation “min+ n ” or “max- n ” (n represents the difference between the corresponding ic vector place i (that is, ic i) and the minima ($\min(c, i)$) or maxima ($\max(c, i)$) of its cardinality), depending on which extreme is closer to the ic vector value. In the case of a tie, the ic vector place is compared to the maximal value. Each place in the second (B) row of the vector shows the opposite (furthest) comparison (if the A row compares the ic vector place with $\min(c, i)$ the B row compares the ic vector place with $\max(c, i)$, and vice versa). For example, the SATV for all-combinatorial hexachord [012678] (ic vector $\langle 420243 \rangle$) is shown below.

$$\begin{aligned} SATV_A[012678]: & \quad \langle \text{max-1, min+2, min+0, min+0, max-1, max-0} \rangle \\ SATV_B[012678]: & \quad \langle \text{min+4, max-4, max-5, max-4, min+4, min+3} \rangle \end{aligned}$$

For the sake of shorthand, and in order to use vectors that are entirely numerical, we will henceforth omit the keyword “min” or “max” from the vectors and retain only the sign (+ or -) and the difference value. “Positive” values will imply a comparison to the appropriate $\min(c, i)$ value and “negative” values will imply a comparison to the appropriate $\max(c, i)$ value.⁸⁸ Thus the SATV for [012678] will appear as follows:

$$\begin{aligned} SATV_A[012678]: & \quad \langle -1, +2, +0, +0, -1, -0 \rangle \\ SATV_B[012678]: & \quad \langle +4, -4, -5, -4, +4, +3 \rangle \end{aligned}$$

Both rows of the vector are necessary because the range of possible values for ic i in $\#X$ is not consistent for all i values. In other words, the difference between $\min(c, i)$ and $\max(c, i)$ is different in many cases (e.g., when

⁸⁸The words “positive” and “negative” are in quotes because the values are not actually positive or negative. The sign indicates whether the ICV value is being compared to a smaller or larger ($\min(c, i)$ or $\max(c, i)$) value. +0 and -0 are opposite ICV(c, i) extremes: the first indicates minimal saturation of ic i in cardinality c ; the second indicates maximal saturation of ic i in cardinality c .

$c = 4$ and $i = 2$, the distance between $\max(c, i)$ and $\min(c, i)$ is 3, but when $c = 4$ and $i = 3$, the distance between the two extremes is 4). The reason we scatter min- and max-related values between the two rows of the SATV rather than segregating them into a row of min-related values and a row of max-related values will become apparent in the next section of this chapter when we use SATV for the purpose of relating set classes.

The ranges for each interval class in each cardinality (i.e., the distance between $\min(c, i)$ and $\max(c, i)$ for all c and i) are shown in figure 2.6. As one can see, complementary cardinalities yield the same sized ranges of possible values (bandwidths) for each interval-class vector place. In other words, the difference values $\max(c, i) - \min(c, i)$ and $\max(d, i) - \min(d, i)$ are the same for all i where c and d are inverses, mod 12.

The above observation combined with Ericksson's observation that when a set class maximally embeds a particular interval class or interval classes, its complement will also maximally embed the same interval class(es)⁸⁹ leads to a very interesting property of interval-class saturation vectors. Because the ic saturation vector places each interval-class vector argument into the context of its range of possible values and, because complementary set classes saturate each interval class to the same degree (within that range), ic saturation vectors of complementary set classes yield the same values. For example, 3-1 [012] and 9-1 [012345678] both have the maximal possible ic1 content, the maximal-1 possible ic2 content, and the minimum amount of all the other interval classes.

⁸⁹Ericksson, 96.

The saturation vector can contribute significantly to analyses of pieces in which complementary relationships play a structural role. However, the equivalence of complementary SCs under SATV should not limit its application *only* to such pieces since fundamentally all compositions written in twelve-tone equal temperament work within the boundaries of pc-space.⁹⁰ The SATV simply represents one method of examining these boundaries. Complementary SATV equivalence may also reinforce aggregate-based relations such as Kh⁹¹ and ZC⁹² even in pieces that do not actually exhibit pcset or SC complementation.

2.4 The Interval-Class Saturation Similarity Index—SATSIM(X, Y)

We will now describe one method of utilizing SATV to examine how closely two set classes resemble each other. The ic saturation similarity index—SATSIM(X, Y)—is a function that compares ic saturation vectors of two sets, returning a real number between 0 and 1 that serves as an indicator of the two sets' relative similarity, following the model set by Morris's ASIM(X, Y). Since a high SATSIM(X, Y) value indicates a lack of similarity among pcsets X and Y , one might more properly call this a “dissimilarity index.”⁹³ The principal difference between the construction of ASIM(X, Y) and SATSIM(X, Y) is that the former deals with one-part ic vectors while the latter uses values in a two-part vector.

⁹⁰Of course, this is to say nothing of the innumerable possibilities for orchestration into p-space and the structures which might exist on that much larger plane.

⁹¹Forte 1973, 96-100.

⁹²Morris 1982, 103-9 and Morris 1987, 74.

⁹³This could easily be transformed into a similarity index by subtracting SATSIM(X, Y) values from 1.

When relating two set classes X and Y using their ic saturation vectors, it is necessary to allow for the possibility (in fact, the likelihood) that the respective rows (A and B) of the two vectors might have a different pattern of max- and min-related values. Since row A of the SATV always contains the “closest” comparison of the $ICV(X)_i$ or $ICV(Y)_i$ value to either $\min(c, i)$ or $\max(c, i)$, we are most interested in those values. To compare two SATVs, we first compare row A of pcset X 's saturation vector to the corresponding min- or max-related value in either row A or B of pcset Y 's saturation vector. We must then compare row A of pcset Y 's saturation vector to the corresponding min- or max-related value in either row A or B of pcset X 's saturation vector. Because the comparison of pcset X to pcset Y frequently yields different values from the comparison of pcset Y to pcset X , it is necessary to perform both in order to insure reflexivity.

To compare two vectors using the SATSIM index, the absolute values of the numerical differences found in the above comparison are added, then this sum is divided by the combined cardinality of the two vectors. Cardinality of saturation vectors is obtained by adding together the distances between the numerical values in the respective arguments of both lines of the vector. If, for example, a particular place in $SATV_A = +4$ and the same place in $SATV_B = -1$, the distance between $+4$ and $-1 = 5$. SATV cardinality is formally defined in figure 2.7. The combined values in saturation vectors will always total the same number for sets of the same cardinality, just as they do in ic vectors. Unlike the ic vector, however, saturation vectors of complementary sets will also add up to the same combined total, since complementary sets yield precisely the same saturation vectors and have the same sized ranges of

possible values. A list of saturation vector cardinalities is provided in figure 2.8 and similarity index SATSIM is formally defined in figure 2.9.

We now provide a demonstration of how SATSIM values are attained. In figure 2.10, we see that pcset X , [012678], has the value (max)-1 in the ic1 column. Pcset Y , [0369], on the other hand, has +0, which means it is minimally saturated with ic1. Because sets X and Y are different cardinalities (and because it will always be possible that $\#X \neq \#Y$ for any sets X and Y), $\min(c, i)$ and $\max(c, i)$ will represent different extremes for each i . It is therefore impossible to compare a min-related value directly with a max-related value; in this case, we must look to line B of pcset Y 's ic saturation vector, which shows that [0369] contains the (max)-3 amount of ic1 (function *row*, described formally in figure 2.9, is a mechanism for determining which row in $\text{SATV}(Y)$ should be compared to $\text{SATV}_A(X)$ for each of the six arguments). The absolute value of the difference between -1 and -3 (i.e., 2) is the value returned for the ic1 column. In the ic2 column, $\text{SATV}_A(X)$ has the value (min)+2, while $\text{SATV}_A(Y)$ row has the value (min)+0, also yielding a difference of 2. In this case, one need not check the value in $\text{SATV}_B(Y)$ since row A had the necessary min-related value. This procedure (step 1 in figure 2.10) is repeated for each place in $\text{SATV}_A(X)$. One then compares each argument in $\text{SATV}_A(Y)$ to either row A or B of pcset X 's saturation vector, creating a two part difference vector. Because only the A rows of one saturation vector are compared to whichever row has the matching argument in the other vector, not all the max- and min-related values are necessarily employed in the comparison. In fact, when both sets have, for example, a max-related value in some ic column of row A, the corresponding min-related

values in the B rows are never actually compared. While an index that does not always consider all available arguments might be viewed as incomplete, by comparing only the *closest* arguments in the SATVs we are greatly reducing the effect of cardinality. If, for example, a trichord and a hexachord both have the $\max(c, i)$ amount of ic1 (that is to say that the first argument in SATV_A of both sets would contain the value -0), they would have SATV_B ic1 values of +2 and +5, respectively. If SATSIM employed all these values in its comparison, then we would see that the two sets are $\frac{|0 - 0| + |2 - 5|}{2 + 5} = \frac{3}{7} = 42.8\%$ different in their ic1 content (we have added together the difference between the min-related values and max-related values and divided that sum by the sum of the distance between SATV_A and SATV_B in the ic1 place for each pcset). Considering that these two sets are both maximally similar in ic1 content (for a trichord and hexachord), this difference seems rather extreme and it occurs solely as a product of their difference in cardinality, a factor that we have striven to temper in our index.

When the differences between row A of pcset X 's saturation vector and the corresponding min- or max-related values in either row of pcset Y 's vector are added together, they are not necessarily the same as the equivalent comparison of pcset Y to pcset X (i.e., $\text{SATV}_A(X) : \text{SATV}_{\text{row}}(Y) \quad \text{SATV}_A(Y) : \text{SATV}_{\text{row}}(X)$ for all X and Y). This is the case of the two sets compared in figure 2.10. Therefore, in order to obtain the same value from a comparison of pcset X to pcset Y and pcset Y to pcset X , it is necessary to add all the difference values together to obtain a composite that reflects both comparisons (step 2 in figure 2.10).

We could have stopped here and had a perfectly acceptable context-free similarity index, one that has, in large part, solved the problem of comparing sets with different cardinalities. However, an even better index is attained by dividing the sum of the differences by the combined cardinality of the two vectors (step 4 in figure 2.10). As we mentioned, when dealing with saturation vectors, the cardinality is obtained by summing the numerical distances in all places of both lines of the vector (step 3 in figure 2.10). This cardinality adjustment better allows us to compare $\text{SATSIM}(X, Y)$ and $\text{SATSIM}(S, T)$ where $\#S$ or $\#T$ are not necessarily equal to $\#X$ or $\#Y$.⁹⁴

The comparison value of 0.54 that SATSIM yields for [012678] and [0369] (step 4 in figure 2.10) represents the very great differences in their ic1, ic2, ic3, and ic5 content. It also represents the congruence of values in the ic4 and ic6 columns, returning a value which indicates that more than 50% of the interval classes in the two sets are not mutually saturated. The larger the number, the more dissimilar two sets are said to be using this formula. The number zero indicates an equivalence relation, while the number one indicates maximal dissimilarity. Maximal dissimilarity, however, is impossible to attain in pitch-class space. Even in the case of two trichords with no interval classes in common (for example, [012] and [048]), there are several shared mutual exclusions.

By definition, any similarity index or equivalence relation that uses the interval-class vector as sole data will consider members of the same set class and Z-related sets as either maximally similar or equivalent.⁹⁵ Because of the

⁹⁴The construction of SATSIM is similar to Morris's ASIM index. (Morris, 1979-80.)

⁹⁵One could argue that this is only a semantic difference in the case of maximally similar SATSIM relations (i.e., $\text{SATSIM}(X, Y) = 0$). While similarity measures normally have no transitivity features, there is transitivity among maximally similar SATSIM relations. If

way SATV interprets the data from the interval-class vector, maximal similarity is found in more places using this index than with most because, in addition to members of the same set class and Z-related sets, complementary sets also return a value of zero using SATSIM, reflecting the fact that they are equally saturated with the same interval classes.

We have already shown (in section 2.3) that the SATVs of complementary pcsets are equivalent. Because SATV provides the basis for SATSIM, complementary pcsets will also be SATSIM-equivalent. In particular, this means that when we are looking through the lens of SATSIM, any mention of SC 4-28 [0369] will also implicitly refer to SC 8-28 [0134679a]. In general, any mention of the SATV or SATSIM properties of #4 SCs will also be taken to include #8 SCs. Like Forte, we will therefore consider a broader definition of cardinality that equates sets and their complements.

To avoid confusion, we will adopt a standard of referring to SATSIM-equivalent sets together at all times. SCs 4-28 and 8-28 will simply be referred to as SC 4/8-28 or [0369] / [0134679a]. We will call this coupling of complementary cardinalities “cardinality pairs” (or c.p.). For the sake of comparison, sets of c.p. #4/#8 will be considered larger than sets of c.p. #3/#9; more formally stated: the larger the difference in cardinality between a pcset and its complement, the smaller the combined set class size. When a single cardinality is required for the sake of a formalization, the lower number is chosen (thus, # $W = 4$ where $W = \text{SC}[0369] / [0134679a]$).

Z-related SCs are also SATSIM-equivalents (as is, by transitivity, the complement of a Z-related SC). Therefore, Z and ZC-related set classes will

SATSIM(X, Y) = 0 and SATSIM(Y, Z) = 0 then SATSIM(X, Z) = 0. Maximal SATSIM similarity therefore meets the criteria for equivalence; that is, it is reflexive, symmetrical, and transitive.

also be referred to jointly as 6-3/36 [012356] / [012347] or 5/7-18/38 [01457] / [01258] / [0145679] / [0124578]. Let us define a SATSIM group as an equivalence class which contains all SCs which are SATSIM equivalent.⁹⁶ There is only one distinct SC in the 6-14 [014358] SATSIM group; there are two distinct SCs in the 6-3/36 or 5/7-1 SATSIM groups; there are four distinct SCs in the 5/7-18/38 SATSIM group.

In addition to the SATSIM equivalence of complementary and Z-related SCs, there are also six SATSIM groups that include SCs which are neither complements nor Z-relations (see figure 2.11). We will consider these special equivalences to be SATV_A Z-related sets since each pair shares the same SATV_A; SATSIM(X, Y) will yield the number 0.00 since it only evaluates line A when SATV_A(X) and SATV_A(Y) contain the same pattern of min- and max-related (“positive” and “negative”) values. The SATV_B values for sets X and Y in the second through sixth groups are, in fact, different, reflecting the differences in cardinality between the two SCs, and the resultant differences in ic ranges of possibility.⁹⁷ (Both SATV_A and SATV_B are given for each complementary SC pair in figure 2.11.)

Each SATSIM group in figure 2.11 shares some common traits. Unlike the common definition of Z-related sets (using the ic vector), none of these special SATV_A Z-pairs are the same cardinality. Each is a pair of i -sets, using the same i -cycle(s) as generator. Each X/Y or \bar{X}/\bar{Y} pair differs in size either by a single member (i.e., pcset X is either one more larger or smaller than pcset Y).⁹⁸

⁹⁶This is, in essence, an equivalence group of SATV Z-relations.

⁹⁷The SATV_B vector is not different for pcset X and pcset Y in the first SATSIM group in figure 2.11 because the empty set/aggregate and the one-member pcset/eleven-member pcset both maximize and minimize all possible interval classes for their cardinalities. This is because there is only one set class that belongs to each of those cardinalities (#0/#12 and #1/#11).

⁹⁸SATSIM groups 2 through 6, shown in figure 2.11, are all segments of K_i subcomplexes about 10-3, 10-6, 10-2, and/or 10-4 (see Kaplan 1990). The set classes in SATSIM group 1 would

In each case, pcset Y (and \bar{Y}) is a comprised of a completed cycle (or concatenation of completed cycles) and $\bar{X} \setminus (\bar{Y} \cap X)$ (the complement of X is a superset of the complement of Y (if the complement of Y is distinct from Y) is a superset of Y is a superset of X).

While there are more occurrences of maximal similarity found using $\text{SATSIM}(X, Y)$ than with other similarity indices⁹⁹, there are no instances of total dissimilarity (SATSIM value of 1). The highest value returned by SATSIM in all comparisons of sets of cardinality 2/10 through 6 is 0.74. This value is found between SCs 2-4/10-4 [04] / [012345689a] and either 6-27 [013469] or 6-30 [013679]. Figure 2.13 shows this comparison. For the sake of space, we will herein shorten our demonstrations of SATSIM from the illustration in figure 2.10. Rather than showing each $\text{SATV}(X)_A : \text{SATV}(Y)_{row}$ and $\text{SATV}(Y)_A : \text{SATV}(X)_{row}$ comparison, we will use a difference vector ($\text{Diff SATV}(X, Y)$) that, in each place, has the former comparison followed by the latter (the two comparisons are separated by commas). The values in each difference vector are summed at the end of the same line. For the sake of comparison, figure 2.10 is duplicated in shorter form as figure 2.12.

The smallest non-zero SATSIM value is 0.02 and is found in two separate SC pairs. The first is between SC 5/7-5 [01236] / [0123467] and SC 6-2 [012346]; the second is between SC 5/7-29 [01368] / [0124679] and 6-33 [023579].¹⁰⁰ These two comparisons are illustrated in figure 2.14. Notice that

(almost trivially) be the progenitors/end-points of all Ki subcomplexes had Kaplan elected to extend the boundaries of Ki beyond #2 and #10 sets.

⁹⁹There are only 111 distinct SATV_A s among the 4096 possible pcsets; by contrast, there are 224 different set classes (T_n/T_nI equivalence groups) among the 4096 possible pcsets (including the empty set).

¹⁰⁰This relationship would be further simplified by allowing M/MI equivalence into our canon. These two pcset pairs would fold into a single pair since SC 5/7-4 is M/MI related to SC 5/7-29 and SC 6-2 is M/MI related to SC 6-33.

the $SATV_A$ of each pair of sets is identical except in the ic6 columns, accounting for the very close relation.

A summary of the possible SATSIM values is given in figure 2.15 in the form of a “value group matrix” (after Castrén 1994).¹⁰¹ Each cell represents a statistical summary of the values possible using $SATSIM(X, Y)$ where X is a pcset of the X -axis cardinality and Y is a pcset of the Y -axis cardinality (or vice versa). The upper left corner of each cell is the lowest SATSIM value possible in the value group;¹⁰² the upper right corner is the highest SATSIM value possible in the value group; the middle left value is the lowest non-zero SATSIM value; the lower left corner contains the average of all the values in the group; and the lower right corner contains the number of distinct SATSIM values in the value group.

In examining the SATSIM value group matrix, we notice some patterns. Almost without exception, the average SATSIM value increases with the difference in cardinality. (Again, for our purposes, sets larger than #6 count as the cardinality of their complement; therefore the comparison of a #4 pcset to a #9 pcset is equivalent to the comparison of a #4 pcset to a #3 pcset, a difference of one, not five.) This seems reasonable when one considers the number of variants possible in a #3/#9 pcset compared to the number possible in a #6 pcset. In the #3/#9 pcset, not only are there many fewer set classes (12 tri/nonachord classes compared to 50 hexachord classes), but there is a much smaller range of possible values in each ICV place. Naturally, the greater the

¹⁰¹Castrén defines a value group as follows: ‘The value group # X / Y contains the values that a given similarity index returns to the SC pairs in the comparison group # X / Y .’ (Castrén, 5.) A comparison group # X / Y ‘contains all SC pairs $\{X, Y\}$ such that X belongs # X and Y belongs to # Y .’ (Castrén, 5.)

¹⁰²If the value in the upper left corner is italicized if it is 0.000 and that number only represents the trivial case of one SC compared with itself (in cases where # $X = Y$). If the upper left number is 0.000 and there is some SATSIM Z-relation in the value group the value is not italicized.

range of ic values, the more potential there will be for relatively similar constructions. Therefore, it should not be surprising that, as a general trend, the larger the c.p. of the set classes in the value group, the smaller the SATSIM values will tend to be. Thus, the smallest average SATSIM values occur when comparing hexachords to hexachords. The largest average SATSIM values occur when comparing #2/#10 sets to #6; this is, not coincidentally, also the largest cardinality difference in our value group.¹⁰³

The above statistical statements should not imply that either SATV or SATSIM is weighted in such a way to reward SC comparisons which have a small difference in cardinality. Such mention of statistical likelihood is more a reflection of the range of possible ICV values in different cardinalities of pcsets. As we have seen when considering individual comparisons, SATSIM's primary determinant of SC similarity is the degree to which each SC is saturated with the same ics. The fact that SATSIM can return the value 0.000 when comparing SCs of different size should dispel any notion that such sets are automatically considered less similar. Despite the statistical information presented above, one could actually make what might seem like a counterintuitive, or even contradictory, claim: that sets of different size are likely to be more similar than sets of the same size. Let us sort out the instances of maximal similarity yielded by SATSIM and examine the smallest

¹⁰³We can now refine our observation as follows: the primary determinant of average SATSIM values in any value group is the difference in cardinality between the two set classes. Thus, if $| \#W - \#X | > | \#Y - \#Z |$, $\text{SATSIM}(W, X)$ will, on average, yield a larger value than $\text{SATSIM}(Y, Z)$. If, however, the difference between the cardinalities of W and X and Y and Z is the same, the secondary determinant of average SATSIM values is the size of the set classes being compared. If $| \#W - \#X | = | \#Y - \#Z |$ and $(\#W + \#X) > (\#Y + \#Z)$, $\text{SATSIM}(W, X)$ will, on average, yield a smaller value than $\text{SATSIM}(Y, Z)$. This generalization does not always hold true: for example, the average $\text{SATSIM}(X, Y)$ value where $\#X = 2$ and $\#Y = 3$ is smaller than the average $\text{SATSIM}(X, Y)$ value where $\#X = 3$ and $\#Y = 4$. The difference between these two averages, however, is small (0.027) and does not discourage us from making the above claim.

non-zero values returned by SATSIM in each value group (the middle left value in each matrix cell). This removes comparisons of each SC to itself and to its SATV Z-relations and leaves the nearest relations possible where the vectors are not the same.

We can see from these values that given any SATSIM(X, Y) comparison, the closest possible non-equivalence for any $\#X$ or $\#Y$ will almost always occur when $\#X \neq \#Y$.¹⁰⁴ That is to say that comparisons where $\#X = \#Y$ yield a higher minimal value than comparing pcset X to the “most similar” set class of pcset Y where $\#Y \neq \#X$. Between sets of non-equivalent cardinality, however, we can see that the greater the difference in cardinality, the higher the smallest non-zero SATSIM value. One can say with some assurance that any given set class will be very similar to at least one set class that is one pc larger or smaller than it (its superset(s) and subset(s), for instance). One cannot say the same about any given set class and some other set class of the same cardinality. Cyclic set classes, for example, are structurally unique in their cardinality. While there may be a SC of another size that shares the same structural features¹⁰⁵, there will not be any other SC of the same size that does. It is therefore reasonable to assert that the most similar (non-equivalent) SC pairs will be different sizes. Naturally, however, when $\#X = \#Y$, the larger the cardinality, the larger the comparison group, and the greater the likelihood that some pcset Y will share many of pcset X 's degrees of intervallic saturation. This is reflected in the gradually smaller numbers that occur from left to right along the top diagonal.

¹⁰⁴The comparison SATSIM(X, Y) where $\#X$ or $\#Y = 2/10$ is frequently larger than SATSIM(X, Y) where $\#X = \#Y = 2/10$

¹⁰⁵E.g., those from the same Ki subcomplex (Kaplan 1990) or ic maxpoint structure (Ericksson 1986).

Appendix B contains a SC-specific statistical summary of SATSIM values. It is a table which shows the average $\text{SATSIM}(X, Y)$ value for each X compared to all Y in the comparison group of all SCs of c.p. 2/10 through 6. Additionally, it shows the lowest and highest possible $\text{SATSIM}(X, Y)$ value for each X compared to all Y in the comparison group. Appendix B also contains the *average* average, lowest, and highest values among each cardinality of SC X compared to the entire range of SC Y s. Both figure 2.15 and appendix B are useful in contextualizing SATSIM values, thus allowing us to make more meaningful analytical statements. In particular, appendix B illustrates that some SCs are, on average, more distantly related to all other sets. For example, we can see from the table that i -cyclic set classes and SCs which are similar to i -cyclic SCs are more distinctive—or *singular*—than average (i.e., their average SATSIM value is much higher than the average of the average SATSIM values for other SCs of the same cardinality).

2.5 PSATV: A one-part alternative to SATV

In this section we will introduce a simpler saturation vector—one that conveys largely the same information in a single series of real numbers rather than a pair of min- and max-related integer values. The proportional saturation vector (PSATV) is a single-part six-argument vector that reflects the saturation level of each interval class as a percentage of what is maximally possible. Simply dividing each value in the interval-class vector by the maximal possible ic value for that cardinality ($\max(c, i)$) would effectively eliminate the comparison with the $\min(c, i)$ values. To avoid this situation, we first create a “min-adjusted ic vector” (MAV) by subtracting the $\min(c, i)$

values from the ICV values. PSATV is then derived by dividing the MAV values from the respective “min-adjusted” $\max(c, i)$ values (i.e., $\max(c, i) - \min(c, i)$). PSATV is demonstrated in two steps in figure 2.16.

PSATVs of complementary set classes are exactly the same, as was the case with SATV (we will therefore continue to consider complementary SCs and complementary cardinalities of SCs in tandem). Using PSATV, however, we can see a bit more easily why complementary sets yield the same interval class saturation vectors (of either sort). The proof lies in the relationship between each interval-class vector argument for a given pcset X and the associated $\min(c, i)$ value. This relationship, which is the foundation for the MAV, always yields identical values for complementary set classes. Figure 2.17 contains an illustration using complementary set classes 5-z36 [01247] and 7-z36 [0123568]. By the same logic, $\max(c, i) - \min(c, i) = \max(\bar{c}, i) - \min(\bar{c}, i)$ for all c and i . In other words, the min-adjusted $\max(c, i)$ values are the same for complementary pcset cardinalities. The equivalence of complementary set classes is obviously maintained when dividing the values in the MAV by a set of min-adjusted $\max(c, i)$ values (which are also the same in complementary cardinalities).

PSATV, like SATV_A , also yields certain special Z-relations beyond those of complementary set classes and ic vector Z-relations (we will call these PSATSIM groups¹⁰⁶). There are only three special PSATSIM groups, and these are a subset of the SATSIM groups (shown in figure 2.11). The PSATSIM groups are shown in figure 2.18. The special PSATSIM groups share the same traits as the special SATSIM groups, and they are also bound

¹⁰⁶While we haven't yet introduced the $\text{PSATSIM}(X, Y)$ index, it should be clear that any index that uses PSATV will return maximal similarity when $\text{PSATV}(X) = \text{PSATV}(Y)$.

by one additional constraint: all PSATV values must be either 0 or 1, reflecting minimal or maximal saturation of each ic . This is a necessary constraint for these special non-complementary PSATV Z-pairs of different cardinality because the range of ic values in each cardinality is different and the real number values of PSATV will reflect those (sometimes subtle) differences, making it nearly impossible for PSATVs of two different cardinality sets to have the same values.

One might expect that SATSIM group #1 from figure 2.11 (0/12-1 compared to 1/11-1) would also appear as a PSATSIM group. This would seem logical given that each ic is both minimally and maximally saturated in each SC. However, because the ICV content of these sets is, in each case, equal to the $\min(c, i)$ place, the MAV yields entirely zeros. The min-adjusted max values are also therefore zeros. Since one cannot divide zero by zero, it is impossible to create PSATVs for these set classes. One could include them as PSATV Z-relations, however, if one is willing to allow for PSATVs with undefined values. Perhaps PSATV is only of limited use because it cannot represent these two SC pairs in a manner consistent with the way it represents all the other SCs. Perhaps, however, “undefined” is the best possible value to describe their intervallic content. After all, they each maximally *and* minimally saturate each ic . Therefore, neither the value 0.00 nor 1.00 is entirely representative—and certainly neither is anything in between.

2.6 PSATSIM: The PSATV Similarity Index

We will now examine the potential for comparing set classes using PSATV. We have already shown that, by-in-large, PSATV conveys the same

information as SATV, but without the necessity of a two-part vector. Using PSATV as the basis for a similarity index has one potential downfall: it is that the vector does not produce a constant cardinality for each #X.¹⁰⁷ For example, the values in PSATV[036] add up to 2.00 while the values in PSATV[048] add up to only 1.00 (see figure 2.19).

The problem with having no constant cardinality is that it is not immediately obvious how one might use this vector in an ASIM-style index; that is, an index which compares the arguments in two vectors then divides the sum of the difference vector values by the combined cardinality of the saturation vectors. Arguably, one doesn't need this type of index where the differences in cardinality have, in large part, been adjusted for in the vector. This is a valid premise—even more so than in the case of comparing SATVs—but it is useful to bring the measure's values within the constant range of 0 and 1 for the sake of comparison.

The largest value in any argument of a difference vector (diffV) that compares two PSATVs is 1.00. This is the case when one SC has the maximal amount (1.00) of some ic and the other has the minimal amount (0.00) of the same ic. Given that we are currently working with a six-place vector, maximal dissimilarity would therefore be represented by the number 6.00, which is the number we will divide the sum of the PSATV diffV values by to create an index which is comparable to SATSIM. Figure 2.20 gives a step-by-step demonstration of PSATSIM; figure 2.21 gives a formalization of the index.

¹⁰⁷Unlike SATV, the cardinality of PSATV is derived like the ICV:

$$\#PSATV = \sum_{n=1}^6 (PSATV(X)_n)$$

Appendix C contains a SC-specific statistical summary of PSATSIM values. A more general summary of the possible PSATSIM values is given in figure 2.22 in the form of a value group matrix (after Castrén, 1994). As before, each cell represents a statistical summary of the values possible using $PSATSIM(X, Y)$ where X is a pcset of the X-axis cardinality and Y is a pcset of the Y-axis cardinality (or vice versa). The upper left corner of each cell is the smallest value possible in the value group; the upper right corner is the largest PSATSIM value possible in the value group; the middle left value is the smallest possible non-zero value; the lower left corner contains the average of the values; and the lower right corner contains the number of distinct PSATSIM values in the value group. This last value — the number of distinct PSATSIM values — might seem extraordinarily large in some instances (e.g., there are 238 different $PSATSIM(\#4/\#8, \#6)$ values). To reflect this degree of precision, all similarity index values have been rounded to three significant digits (digits to the right of the decimal point) in the value group matrices. While the analytic usefulness of this degree of precision is perhaps questionable (this will be discussed in chapter 3), the figures have been rounded from the 8-bit real numbers used by the computer when calculating the measurement values. We believe this presents a sufficiently accurate account of the measures' precision.

The same patterns that we noticed when examining the SATSIM value group matrix in figure 2.15 remain invariant in the PSATSIM value group comparisons. We commented, in the former case, that “almost without exception, the average value increases with the cardinality difference.” When discussing PSATSIM, we can remove the word “almost” from that

observation and make it more systematic. Without exception, the primary determinant of average PSATSIM values in any value group is the difference in cardinality between the two set classes; once again the secondary determinant is the size of the set classes being compared.

The overall range of possible values using PSATSIM is a bit smaller than those yielded by SATSIM. SATSIM yields values ranging from 0.00 to 0.74; PSATSIM yields values ranging from 0.00 to 0.69. Unlike SATSIM, however, the #6 : #6 group gives us not only the widest range of values (the #5 : #6 group also has the same range), but also the largest possible single comparison (most dissimilar set class pair).

The smallest non-zero number yielded by PSATSIM is 0.03, between SCs 5/7-7 [01267] / [0123678] and 6-7 [012678] (see figure 2.23).

The largest PSATSIM relation is found between any of the first-order all-combinatorial hexachords (SCs 6-1 [012345], 6-8 [023457], or 6-32 [024579]) and either the whole-tone collection (6-35 [02468a]) or 5/7-33 [02468a] / [012468a]—the five/seven-note sub/superset of the whole-tone collection (recall that 6-32 and 5/7-33 form one of the special PSATV Z-relations). While the latter pair clearly has the maximal amounts of each of the even ics, the first-order all-combinatorial hexachords have the maximal total amount of the odd ics and the minimal amounts of ic4 and ic6. These comparisons are shown in figure 2.24.

As we did with the SATSIM values, let us now examine the smallest non-zero values returned by PSATSIM in each value group. All the observations we derived regarding the smallest non-zero SATSIM values hold true with PSATSIM as well. In summary, the values in the middle row of each cell of the

PSATSIM value group matrix increase in size from the top to the bottom of each column, except for the first value, which is the highest. Each diagonal decreases in size from left to right. The top cell in each column indicates the PSATSIM(X, Y) values where $\#X = \#Y$ and $X \neq Y$.

2.7 PSAT%V: An relational adjustment to the PSATV

As a way to normalize the values of PSATV into a vector which maintains a constant cardinality, we will now create a %-vector¹⁰⁸ using the values of PSATV(X). This new vector, which, for the sake of consistency with Castrén, we will call PSAT%V(X), is created by dividing each argument in PSATV(X) by the $\#PSATV(X)$. The result is a vector that's elements always add up to 1.00, like Castrén's nC%-vector.¹⁰⁹ This vector is demonstrated in figure 2.25 and formalized in figure 2.26.

Since PSAT%V is an elaboration on PSATV, it yields precisely the same equivalences and Z-relations as PSATV, including the “special” Z-relations listed in figure 2.18 above (i.e., PSATSIM group $_n =$ PSAT%SIM group $_n$ for all n). Unlike the other saturation vectors, which, by design, can have the full range of values (minimal to maximal saturation) for each ic argument in every cardinality SC #2/#10 through #6, this one is limited by the constraints of pcset cardinality. SATV_A and PSATV both represent the amount to which each ic is or isn't saturated in a set class; but they do so at the cost of an inconsistent vector cardinality. That is to say that the value in the n th place of a PSATV has no direct bearing on any other place in the vector (other than

¹⁰⁸As described in Castrén 1994.

¹⁰⁹We will not follow Castrén's convention of multiplying each value by 100 for the sake of working with integers.

those described in figure 2.5). PSAT%V takes the values in PSATV and weighs them against each other. Maximal saturation of a particular interval class will always appear as the largest number in any single PSAT%V, but a pcset which carries the maximal amount of, say, two ics, may appear to have relatively fewer of both than a pcset which is maximally saturated with only one ic. Figure 2.27 presents PSAT%V[036] and PSAT%V[048]. The former is a multiply cyclic set; the latter is a singularly cyclic set. We mentioned this difficulty when describing Castrén's RECREL index in the first chapter and we will revisit the issue in greater detail in the next section of this chapter (2.8).

Figure 2.28 shows the highest and lowest possible values of any single argument of PSAT%V for sets of all cardinalities. While minimal saturation, reflected by the value 0.00 is always possible for sets of #2/#10 through #6, maximal saturation is reflected by different numbers depending on the number of other interval classes in a pcset and their degree of saturation. Therefore, listing 0.44 as the largest *possible* single value in any PSAT%V(#5), does not infer that every time an ic is maximally saturated in a pentachord, it will carry that particular value (e.g., SC [02468], which maximally saturates all the even ics, has the PSAT%V: <0.00, 0.33, 0.00, 0.33, 0.00, 0.33>). Naturally, like PSATV, PSAT%V cannot be used for the empty set, one-element pcset, eleven-element pcset, or aggregate because the values would all be undefined.

2.8 PSAT%SIM: The PSAT%V Similarity Index

Of the three vectors we have examined, PSAT%V(X) most readily facilitates both SIM- and ASIM-like comparisons. Since the vector cardinality

is always set at 1.00, the sum of two vectors, or 2.00, is the constant denominator for each ASIM-like PSAT%V comparison. This new relation, PSAT%SIM, is demonstrated in figure 2.29 and formally defined in figure 2.30. Having defined a “simplification” of SATV (PSATV), an adjustment of it (PSAT%V) that ranks the relative prominence of its saturation values, and similarity indices for comparing either of them, we will now examine the differences between these two elaborations on SATV and the similarities and differences in the numbers the generated by PSATV and PSAT%V.

The value group matrix for the PSAT%SIM index (figure 2.31) and the SC-specific summary of PSAT%SIM values (appendix D) show markedly different results than the analogous figures and appendices which used the SATSIM and PSATSIM indices. PSAT%SIM yields the same number of maximal similarities as does PSATSIM. This is not surprising since the PSAT%V Z-relations are the same as the PSATV Z-relations. Unlike the other indices, however, there are instances of maximal dissimilarity ($PSAT\%SIM(X, Y) = 1.00$), and quite a number of them at that. PSAT%SIM is therefore the only saturation vector-based index that returns the widest range of possible values from zero to one.

This expanded range of values and the possibility of both maximal and minimal similarity distinctions might be considered a desirable feature. On the other hand, one might consider it undesirable to allow maximal dissimilarity in pc space for a reason cited in section 2.4: there are no SC pairs X and Y where $ICV_i(X) + ICV_i(Y) = 0$ and $ICV_i(X) + ICV_i(Y) = 1$ for all i . The same holds true using either SATV or PSATV. In other words, there are no pairs of pcsets where their ICVs (or SATVs or PSATVs) are in complementary distribution.

There therefore cannot be maximal dissimilarity when using SATSIM or PSATSIM.

Since PSATSIM does return maximal dissimilarity in some cases, it is clear that the index does not effectively measure mutual pc exclusion (non-saturation). Castrén's nC -vectors and our extension of them to the realm of saturation vectors, force each vector to have the same cardinality, regardless of the set class cardinality. This tends to overplay the presence of each SC in nC cardinalities with relatively few SCs (e.g., dyad and trichord classes) and it downplays the presence of each element in nC cardinalities with large numbers of SCs (e.g., hexachord classes). A listing of PSATSIM maximal dissimilarities is provided in figure 2.32 (only the smaller SCs in each c.p. SC pair are provided for the sake of space).

Maximal PSATSIM(X, Y) dissimilarity occurs when the intersection of non-zero values in the respective arguments of PSATV(X) and PSATV(Y) is zero.¹¹⁰ In the #2 : #2 group, every set class pair that is not a PSATSIM equivalence (found, in this case, when comparing any particular SC to itself) is maximally dissimilar. For example, [01] is maximally dissimilar to [02]. This produces the same result as non-saturation-based indices such as ASIM: it only takes into account the intervals that are mutually saturated and not those that are mutually minimized (in the case of [01] and [02], the mutually minimized ics are [03], [04], [05], and [06]—two thirds of the ics).

We can see in figure 2.32 that, for the most part, set class pairs that are maximally dissimilar have no ics in common. This is true, in fact, of all comparisons except those between [04] or [048] and pentachords or

¹¹⁰More formally, PSATSIM(X, Y) = 1.00 iff PSATV(X) _{n} ∩ PSATV(Y) _{n} = 0 for all n N .

hexachords. While there are very few maximal dissimilarities between most comparison groups, there are quite a number of maximal dissimilarities to SCs [04] and [048]. Remarkably, both of these SCs are maximally dissimilar to over half the number of hexachords. The reason behind this seeming glut of maximal dissimilarities to these SCs is two-fold. They are both equally related to the same SCs because both [04] and [048] have the same PSATV and PSAT%Vs (both are maximally saturated with ic4 and embed no other ics). By contrast, the listed hexachords and pentachords are all minimally saturated with ic4 (recall that $\min_{5,4} = 1$ and $\min_{6,4} = 2$). This would seem to indicate that while minimal saturation of most ics is a unique and distinguishing feature, minimal saturation of ic4 for #5 and #6 pcsets is rather commonplace. While this feature might call the particular usefulness of PSAT%SIM into question, it is also valuable to keep in mind when drawing analytic conclusions from saturation vector data in general.

2.9 The Generalized Saturation Vector (SATV n)

For ease of exposition, we have, up to this point, restricted all our saturation vectors to interval classes. This was a useful limitation to demonstrate the construction of each vector (using an easily manageable number of arguments) and its associated similarity index. One can, however, easily create saturation vectors (of any of the three varieties that we have already demonstrated) that show relative content of any subset size. As we mentioned at the beginning of chapter 1 (section 1.1), an interval class is not the same as a #2 subset: the former represents a distance between two elements; the latter represents a group of two elements (which are necessarily

separated by some interval). That said, one can produce an equivalent structure to the $ICV(Y)$ by performing $EMB(/X/, /Y/)$ six times, where $/X/$ cycles through all SCs of #2 (2-1 through 2-6) and displaying the results in a 6-argument array.

By extension, rather than the interval class (or #2 subset class) saturation vector (SATV), we could have a trichord class SATV for SCs larger than #3, or a hexachord class SATV for SCs larger than #6, etc. In this section, we will generalize SATV to allow any or all subset-class cardinalities. The generalization, $SATV_n(X)$, is constructed precisely like our earlier $SATV(X)$, but carries an extra variable, n , which signifies the cardinality of subsets. The vector SATV, defined in section 2.3, will now be called $SATV_2(X)$ and the relation SATSIM, defined in section 2.4 will now be called $SATSIM_2(X)$ in keeping with our new convention.

Let us return to our function, $SATEMB(/X/, /Y/)$, first suggested in section 2.1. $SATEMB(/X/, /Y/)$ returns the degree that SC X is saturated in SC Y compared to maximal and minimal $EMB(/X/, /Y/)$ values for any SC of # Y . If, for example, $X = \text{SC } 3\text{-}1 [012]$ and $Y = \text{SC } 6\text{-}7 [012678]$, $EMB(/X/, /Y/) = 2$. There are, however, a maximum of four embedded $[012]$ s in any hexachord (there are four in SC 6-1 $[012345]$) and the minimum number is zero (there are none in SC 6-35 $[02468a]$, among others). Therefore, $SATEMB(/X/, /Y/) = \langle (\max)-2, (\min)+2 \rangle$

We can now refine our definition of $SATV_n(Y)$ as a complete listing of all $SATEMB(/X/, /Y/)$ values where each $/X/$ is a SC of cardinality c . The $SATV_3(6-7)$ is shown in figure 2.33. In section 2.2, we showed how $\min(c, i)$ and $\max(c, i)$ values where $c = 2$ (interval classes) can be derived from

examining the cyclic set classes. Naturally, we could do the same sort of thing, examining imbricated chains of all the trichord-classes, tetrachord-classes, etc. to generate set classes which maximally saturate each SC of $\#c$. The number of SCs involved and the length of the process make this exercise prohibitively long, however. Appendix A contains a complete list of the $\max_{c,i}$ and $\min_{c,i}$ values where $c = \#2$ through $\#12$ and i is the complete range of SCs smaller than or equal to the superset cardinality. These values were generated by a computer program which simply calculated all possible $\text{EMB}(X, Y)$ values where $\#X$ and $\#Y$ fall between 2 and 12. The program then saved the largest and smallest values for each SC Y .

Before introducing SATSIM_n , an index for comparing SCs using SATV_n , we will present an interesting corollary to the SATV_2 equivalence of complementary SCs. We recall that all complementary SCs generate identical SATV_2 s, PSATV_2 s, and $\text{PSAT}\%V_2$ s. The same does *not* hold true with complementary SCs using SATV_n where $n > 2$. Figure 2.34 shows the SATV_3 s of complementary SCs 5-1 [01234] and 7-1 [0123456]. While there are distinct similarities between the two SATV_{3A} s, they are not precisely the same and their SATV_{3B} s are markedly different, reflecting the strong differences in the range of possible $\text{SATEMB}(\#3, \#5)$ and $\text{SATEMB}(\#3, \#7)$ values.

Clearly, then, complementary SCs would not be considered equivalences under an index that uses SATV_n s where $n > 2$. However, just as SATV_2 yielded an interesting complementary equivalence, a similar property is found when examining all the SATV_n s of $\#10$ SCs. Figure 2.35 shows all the SATV_n s (for $n = 2$ through 9) of SC 10-3. Notice that all complementary set

classes are saturated to precisely the same degree. In other words, $\text{SATV}_3(10-3) = \text{SATV}_9(10-3)$, $\text{SATV}_4(10-3) = \text{SATV}_8(10-3)$, etc. Similarly, complementary hexachord classes are also saturated to the same degree (e.g., $\text{SATEMB}(6-z_3, 10-3) = \text{SATEMB}(6-z_{36}, 10-3) = \langle (\max)-2, (\min)+2 \rangle$). More formally, we can declare that when $\#Y = 2$ (and $\#\bar{Y} = 10$), $\text{SATEMB}(Y, X) = \text{SATEMB}(Y, \bar{X})$ and $\text{SATEMB}(X, \bar{Y}) = \text{SATEMB}(\bar{X}, \bar{Y})$. That is to say that all complementary set class pairs (larger than #2) saturate the same interval classes (#2 SCs) to the same degree and the complements of the interval classes (#10 SCs) saturate complementary set class pairs (larger than #2) to the same degree. This is true of no other complementary SC cardinalities.

2.10 SATSIM_n: A similarity index to compare SCs using SATV_n

In order to compare two set classes using our expanded SATV_n, we will introduce a generalized version of $\text{SATSIM}(X, Y)$ which is capable of comparing SATV_ns where $n =$ a single cardinality or a range of cardinalities that are smaller than both $\#X$ and $\#Y$. Any $\text{SATSIM}_n(X, Y)$ that uses all applicable SATV_n values (from $n = 2$ to the smaller of $n = \#X-1$ and $n = \#Y-1$) will be referred to as $\text{TSATSIM}(X, Y)$ (the “total” saturation vector similarity index).

Figure 2.36 shows all significant SATV_n vectors for 6-z₃₇ [012348]. Figure 2.37 shows all possible SATV_n vectors for 6-z₄ [012456]. As one can see from a quick glance at the SATV₂s, the two hexachords are SATV₂ Z-related (and also ICV Z-related, since there are no SATV₂ Z-relations among same-sized SCs that are not also ICV Z-relations). Therefore, $\text{SATSIM}_2(6-z_{37}, 6-z_4) = 0.00$. For all other SATV_ns (where $n > 2$), however, the two SCs are

rather different. In fact, there are no ICV or SATV2 Z-related SCs that are also SATV $n>2$ Z-related. Even the two all-interval tetrachords have different SATV3 vectors. Figure 2.38 shows the SATV3 vectors of 6-z37 and 6-z4 and gives a SATSIM3 comparison of the two. SATSIM3 is calculated precisely the same as SATSIM[2], described earlier in section 2.4. Like the SATV2(X), the values in SATV3(X) (and any SATV n (X)) add up to constant cardinalities for constant c and $\#X$ values. The list of SATV n cardinalities is given in figure 2.39.

Rather than rigorously discussing SATSIM2, SATSIM3, SATSIM4, etc. as separate similarity indices, providing value group comparisons and demonstrations of each, we will instead limit our discussion to TSATSIM—the total subset saturation similarity index. TSATSIM is an amalgam of all possible and non-trivial SATSIM n relations. A possible SATSIM n (X, Y) relation is one in which $n \leq \#X$ and $n \leq \#Y$ (i.e., we must be examining pcsets of at least the cardinality of the smaller pcset). A non-trivial relation is one in which $n < \#X$ and $n < \#Y$. A SATSIM n (X, Y) relation where $n = \#X$ or $\#Y$ is considered trivial because the smaller SC will be maximally saturated with itself and minimally saturated with all other SCs of its cardinality.

Let us now work through a TSATSIM comparison using the same ICV Z-related pair of hexachords (6-z37 and 6-z4). We will demonstrate this vector comparison only once and not in as much detail as the other ones because of its rather cumbersome length. Because these are ICV (and SATV2) Z-related, the SATSIM2(6-z37 and 6-z4) comparison = $\frac{0}{56} = 0.00$. As noted above, the SATSIM3 comparison = $\frac{28}{140} = 0.20$. The SATSIM4 comparison = $\frac{40}{166} = 0.24$; and the SATSIM5 comparison = $\frac{24}{160} = 0.15$. The TSATSIM value is

calculated by dividing the sum of the numerators by the sum of the denominators. Therefore:

$$\text{TSATSIM}(6\text{-}z37, 6\text{-}z4) = \frac{0 + 28 + 40 + 24}{56 + 140 + 166 + 160} = \frac{92}{522} = 0.18$$

Depending on how one wanted to weight this comparison, one could also derive a total subset saturation index by averaging the individual SATSIM n values. This would give equal weight to each SATSIM n comparison rather than biasing the TSATSIM comparison toward the SATSIM n comparison(s) where n is/are closest to 6 (i.e., the comparisons with the greatest number of elements). This second way of deriving the comparison is as follows:

$$\begin{aligned} \text{Average of SATSIM}_n = 2..5(6\text{-}z37, 6\text{-}z4) = \\ \frac{0.00 + 0.20 + 0.24 + 0.15}{4} = 0.15 \end{aligned}$$

This value is naturally lower than that TSATSIM value that we first calculated because the SATV2 Z-relation is one quarter of the total weight whereas it is only a bit more than 10% of the weight of the earlier index. One can certainly make a case for calculating the index either way. We will call the first method TSATSIM(X, Y) and the second method AvgSATSIM n (X, Y). The value group matrix for TSATSIM(X, Y) is given in figure 2.40 and the SC-specific summary of values is given in appendix E; the value group matrix for AvgSATSIM n (X, Y) is given in figure 2.41 and the SC-specific summary of values is given in appendix F.

The range of TSATSIM values is flatter than the range of SATV2 values. There are fewer possible low numbers, largely because of the decreased probability of similar embedding patterns of #3 and larger subset classes. The highest TSATSIM values are not as high as any of the other indices we have examined because of the very large SATV n cardinalities when $n > 3$.

Dividing by the combined cardinality of the vectors creates a remarkably large denominator, which invariably leads to a relatively low dissimilarity level.

TSATSIM and AvgSATSIM n both have limitations that SATSIM2 does not. Neither comparison can compare dyad classes to any other set class because both comparisons compare SATV2 through SATV($(\#X < \#Y)-1$). When either $\#X$ or $\#Y = 2$, TSATSIM or AvgSATSIM n would have an impossible range of values to compare. Also, when the smaller of the two SCs, X or Y , is a trichord, TSATSIM or AvgSATSIM n are the same as SATSIM2, and therefore yield the complementary equivalences which are not present in other (larger) TSATSIM or AvgSATSIM n comparisons. This means that when comparing a pcset larger than a hexachord to a trichord, the larger the difference, the *smaller* the average TSATSIM or AvgSATSIM n comparison would be. This is the opposite of what happens when two SCs that are both larger than a trichord are compared using those indices. In both of these cases, the larger the difference, the higher the average comparison value.

Despite these admitted quirks, TSATSIM and AvgSATSIM n are both useful indices where the analyst wishes to use a saturation-based index, but also wishes to differentiate between Z-related set classes and complements. In sections 2.12 and 2.13, we will introduce another vector and measure which differentiates between these SCs, but which produces values which are much closer to the interval-class saturation indices.

While TSATSIM and AvgSATSIM n both differentiate Z-related and complementary SCs (as before, all values of 0.000 in the value group matrices are italicized when the case(s) of maximal is/are trivial—that is when the two sets being compared are members of the same SC). There are two special cases

of TSATSIM and AvgSATSIM n equivalence, however. These are between the set class pairs 5-21[01458] / 6-20[014589] and 5-33[02468] / 6-35[02468a]. Not surprisingly, these two set class pairs are also maximally similar using the SATSIM index.

2.11 PSATV n , PSAT%V n , and associated similarity indices.

Naturally, it is possible to create one-part proportional vectors and %-vectors out of the generalized SATV n . Constructing these vectors and associated similarity indices is accomplished in precisely the same manner as shown in sections 2.5 through 2.8.

2.12 The Cyclic Saturation Vector (CSATV)

The remainder of this chapter will deal with extensions to SATV2(X) and SATSIM2(X) that weight cyclic adjacencies in pc-space. The assumption behind the cyclic saturation vector (CSATV(X)) and its associated index (CSATSIM(X, Y)) is that adjacent pcs within an i -cycle project ic i more strongly than pcs separated by ic i (s) that either fall in different i -cycles or are non-adjacent within the same i -cycle. For example, under this premise we might claim that a four-note quartal or quintal chord projects ic5 better than a chord with three non-adjacent ic5s. Even though we are dealing with pcsets and not their particular orientations in p-space, we believe that it is still reasonable to assert that many realizations of set class [0257] will likely project ic5 to a great extent.¹¹¹ The same will, of course, be true of any of the

¹¹¹Furthermore, there is a direct correlation between the number of shared i -cycle adjacencies in two sets classes $/X/$ and $/Y/$ and their potential degree of intersection.

i -cyclic set classes presented in section 2.2. In essence, CSATV allows one to determine the degree to which any SC resembles a cyclic set class.

Because CSATV(X) is a weighted vector (i.e., an intermediate weighting step must be included in its derivation), it is more complex than SATV $_n(X)$. Cyclic groups for use with CSATV(X) and CSATSIM(X, Y) will initially be drawn from the standard interval-class cycles, but we will also explore the possibility of allowing other one-to-one TTO and non-TTO mappings (such as those discussed in Morris 1987, 170-177).

We will define CSATV(X) in several parts. Because its data is derived fundamentally from cyclic subsets, we will first examine the manner in which pcset X is fragmented into cyclic subsets. Figure 2.43 provides an example of cyclic fragmentation where X is the pcset {0, 1, 2, 3, 4, 8} (SC 6-z37, ICV = <432321>). CycFrag(X, i) is a reorganization of the pcs of pcset X to show adjacencies within the i cycle(s). Parentheses delineate completed ic cycles and adjacent pcs within the brackets (including the wraparound) are i -cyclic adjacencies (producing a single embedded ic i). Dashes (-) indicate vacant places in each i cycle. For the sake of reference, the complete cyclic fragmentation of 12-1 [0123456789ab] (where each cycle is complete and there are no vacant places) is given in figure 2.42.

From the CycFrag(X, i) representations, we can now extract the number of distinct fragments and their cardinality. This information will be collected into a new array — Part $_i(X)$ — which lists the cardinalities of the i partitions of X from largest to smallest. The sum of the Part $_i(X)$ numbers will add up to the cardinality of set X as long as the operation i is one-to-one and onto. In the case where $X = 6\text{-z37}$ and $i = \text{ic1}$, Part $_i(X) = \langle 51 \rangle$ because the set

{0, 1, 2, 3, 4, 8} has one five-note 1-cycle fragment ({0, 1, 2, 3, 4}) and one one-note 1-cycle fragment ({8}). The complete $\text{Part}_i(X)$ values where $i = 1 \dots 6$ are shown in figure 2.44.

We will now construct a new vector, the “ic cycle vector” (ICCycV), that shows how the interval-class content of a set is partitioned among cyclic fragments. $\text{Part}_i(X)$ will be our vehicle for deriving this information. In figures 2.43 and 2.44 we can see that the four ic1s embedded in 6-z37 stem from a single unbroken cyclic fragment and that only a single pc in the pcset is disjunct from that cycle. Similarly, the three ic4s are formed from a single completed ic4 cycle. The three pcs that are not a part of that cycle are therefore all from different ic4 cycles. By contrast, there are three ic2s—two from one cycle and one from the other.

The $\text{ICCycV}(X)$ has six arguments — each of which is a vector of variable arguments, depending on the number of fragments possible in each ic cycle. To derive the $\text{ICCycV}(X)$ from the six $\text{Part}_i(X)$ arrays, simply subtract 1 from each Part_i value (e.g., a one-note partition yields no interval i s, a two-note partition yields a single interval i , etc.) except in cases where a particular Part_i is equal to the periodicity of the i -cycle (e.g., a three note partition of a 4-cycle (periodicity = 3) yields three ic4s). Figure 2.45 shows the ICCycV of SC 6-z37.

The sum of the numbers in each of the six internal vectors is the same as the argument in the parallel place of the ICV of the same SC; more formally, $\text{ICV}_i(X) = \sum_n \text{ICCycV}_i(X)$. The ICV tells us that SC 6-z37 embeds four ic1s, 3 ic2s, etc. The first internal vector ($\text{ICCycV}_1(X) = \langle 4000 \rangle$) tells us that the four ic1s are all from a continuous 1-cycle fragment; the second internal vector

($\text{ICCycV}_2(X) = \langle 2100 \rangle$) that the three ic2s are divided into two from one ic2-cycle and 1 from the other. The internal vectors have variable numbers of places which represent both the number of complete cycles for each ic and the degree to which each cycle can be fragmented. We recall from section 2.1 that the size of a maximally fragmented cycle equals the periodicity (p) integer divided by 2.¹¹² In this case, however, we are looking for the greatest number of disjunct ic i s in each cycle. This number is found by dividing p by 3 (rounding up or down to the nearest integer) and multiplying quotient by the number of distinct i cycles (m). More formally, the number of elements in each $\text{ICCycV}_i = m \frac{p}{3}$. These values are given in figure 2.46.

Ic1, for example, can be broken into no more than four disjunct fragments (e.g., {0, 1, 3, 4, 6, 7, 9, a} has four ic1s, none of which are conjunct and there is no way to add another pc to this pcset without adjoining two of the cyclic fragments). There is, therefore, no need for more than four arguments in ICCycV_1 . While there is only a single unbroken ic1 cycle, there are three possible ic3 cycles, but it is impossible to have two ic3-cyclic fragments from the same cycle (i.e., if there are two ic3s from the same cycle, they must be adjacent within the cycle). Therefore, there can only ever be three non-zero places in the ICCycV_3 vector.

For comparison, in figure 2.47 we examine the $\text{CycFrag}(X, i)$ and $\text{ICCycV}(X)$ of 6-z4 [012456] — the Z-relation and abstract complement of SC 6-z37. Recall that earlier in this chapter we compared the SATVns of these two SCs. While 6-z4 and 6-z37 embed each ic to the same degree, the way that the pcs are partitioned among the ic cycles is different in four of the six ic

¹¹²i.e., divided with the remainder truncated.

cycles. Clearly, then, any index which compares these two vectors should yield some degree of difference as long as the four ic1s represented by the $\text{ICCycV}_1(6-z37) = \langle 4000 \rangle$ were weighed differently than those represented by $\text{ICCycV}_1(6-z4) = \langle 2200 \rangle$. The former indicates that all the ic1 content comes from a single cyclic adjacency; the latter indicates that the ic1 content is fragmented once along the 1-cycle. The ic4 content is even more different. As we mentioned, SC 6-z37 embeds three ic4s, all from a single unbroken ic4 cycle [048] ($\text{ICCycV}_4(6-z37) = \langle 3000 \rangle$); SC 6-z4, however, embeds three ic4s from three separate (disjunct) and fragmented ic4 cycles [04], [15], and [26] ($\text{ICCycV}_4(6-z4) = \langle 1110 \rangle$).

The intuition that we are attempting to model is that non-adjacent i -cyclic fragments project ic i less successfully, or at least differently, than completed cycles or cyclic adjacencies. Therefore, we will want to weight larger numbers higher than groups of small numbers (e.g., $\langle 3000 \rangle$ would indicate a greater cyclic presence than $\langle 1110 \rangle$). Perhaps the easiest way to weight these would be to square all the values, then add them together. This would create vectors with values of $\sum_n \sum_N \langle 9000 \rangle = 9$ and $\sum_n \sum_N \langle 1110 \rangle = 3$.

The difference between the two (6) would be the value that a difference vector would return for the ic4 place of these two SCs. While we want to differentiate between cyclically adjacent and non-adjacent intervals and establish a bias favoring the former, we do not want create an exaggerated comparison by weighting the former *too* heavily. Therefore, we believe that simply squaring the ICCycV_{i_n} places produces too course a weighting system.

Taking the square root of the sum of the squared vector places

$\sqrt{\sum_n \sum_N \text{ICCycV}_{i_n}(X)^2}$ is one way to temper this roughness. This difference

between $\sqrt{9}$ (3) and $\sqrt{3}$ (1.73) is 1.27: a much smaller difference, and one which would still allow for a fairly close relation of these two Z-related hexachords. Even this adjustment may not be enough, though. The former SC would, by this index, still appear to have 73% more ic4 salience. While this seems better than saying that it has 200% more ic4 salience (as would be the case if we did not take the square root), it still seems quite exaggerated; and this discrepancy is worsened in cyclic vectors of larger SCs.

We will therefore adopt a variable weighting system that is capable of more gradual scaling. This system, which we will call WEIGHT, is an additive formula which is constructed as follows: We begin with the number 1, which we multiply by a constant (real number or integer) value. If the ICCYcV_{i_n} value is 1, we stop. If not, we add 1 to our product then multiply the new number by our constant. If the ICCYcV_{i_n} value is 2, we stop. Otherwise we continue as above until the number of iterations equals the ICCYcV_{i_n} value. This process is outlined as a C-program segment in figure 2.48. Let us assign a value to the variable *constant* and trace the function. If *constant* = 1.2 and the particular ICCYcV_{i_n} value is 4, the weighting procedure would be as shown in figure 2.49.¹¹³

We can now return to the problem of weighting the two ICCYcV_{i_n} vectors: <3000> and <1110>. $\text{WEIGHT}(3) = 4.37$ (see the third iteration in figure 2.49 above); $3 * \text{WEIGHT}(1) = 1.2 * 3 = 3.6$ (this value is multiplied by 3 because of the three separate 1s in the second vector). Using this system,

¹¹³While I believe that WEIGHT is easiest to understand as a recursive function, it can also be modeled as a simple formula. Let n represent the number that is being weighted and k represent the weighting constant.

$$\text{WEIGHT}(n) = \frac{k}{k-1} [k^n - 1].$$

My sincere thanks to Panayotis Mavromatis for calculating this equation.

the former weighted ICCycV value is only 21.4% larger than the latter. We believe that this is a more reasonable difference than yielded by the above methods and we shall adopt it in our CSATV(X). The constant 1.2 provides a good range of values, but one can certainly adopt other constants in creating this vector. The constant values of each possible ICCycV value (where constant = 1.2) is given in figure 2.50.

We can now define the new weighted ICCycV (WICCV) that is derived by submitting each ICCycV value to function WEIGHT then (separately) summing the values in each of the six parts of the complex vector ($WICCV_i(X) = \frac{1}{n} \sum_{N} (WEIGHT(ICCycV_{i_n}(X)))$). A demonstration of our new WICCV is provided in figure 2.51.

This new weighted vector can be used in place of the ICV in any ICV-based similarity index (including Teitelbaum's similarity index, Morris's ASIM(X, Y), and Isaacson's IcVSIM(X, Y)). It is rather similar to the standard ICV, but uses weighed real number values instead of integers. Naturally, however, in keeping with the spirit of the indices introduced herein, we will want to transform the WICCV into a saturation vector: that is, one that compares each of the values in the WICCV(X) to what is both minimally and maximally possible in any WICCV($\#X$) (given the same weighting constant). We will call the two extremes $cmin(w, c, i)$ and $cmax(w, c, i)$, where w = the weighting constant (1.2 in the cases above), c = the cardinality of X , and i = the interval class. As was the case with the $min(c, i)$ and $max(c, i)$ values, these weighted minimal and maximal extremes are easily derived from examining the i -cyclic set classes (discussed in section 2.2). The $cmin(1.2, c, i)$ and $cmax(1.2, c, i)$ values are given in figure 2.52.

The cyclic saturation vector $\text{CSATV}(X)$ is a two-part six argument vector formed by comparing the each $\text{WICCV}_i(X)$ argument to the $\text{cmax}(w, c, i)$ and $\text{cmin}(w, c, i)$ values in the same way that the $\text{SATV}(X)$ is formed by comparing each $\text{ICV}_i(X)$ argument to the $\text{max}(c, i)$ and $\text{min}(c, i)$ values. In figure 2.53, we demonstrate a $\text{CSATV}(X)$ where $X = 6\text{-z}37$. As with the SATV , CSATV 's two parts (A and B) are delineated by the closest and farthest relations of each WICCV place to $\text{cmin}(w, c, i)$ and $\text{cmax}(w, c, i)$. The CSATV_A values reflect the nearest relationship of each WICCV place to either extreme; the CSATV_B values reflect the opposite (farthest) relationship. As before, the signed values in the CSATV represent abbreviations of either $\text{cmin}(w, c, i)+n$ or $\text{cmax}(w, c, i)-n$.

There are fewer CSATV equivalences (Z-relations) than yielded with SATV or its variants. Complementary SCs do not produce identical CSATVs and neither, for the most part, do ICV Z-related SCs. In fact, there is only one case where two different SCs yield identical CSATVs : the all-interval tetrachords, 4-z15 [0146] and 4-z29 [0137]. Since they both embed one and only one of each interval class, each i cycle is fragmented to the same degree. There are a small handful of cases of CSATV_A equivalences. Except for a single complementary set pair which is CSATV_A equivalent (CSATSIM group #4), these are a subset of the special SATV equivalences shown in figure 2.11 in section 2.4 and also a subset of the smaller group of special PSATV equivalences shown in figure 2.18 in section 2.5. The group of CSATV_A equivalences is presented in figure 2.54 (the complete CSATV is given in each case).

2.13 The Cyclic Saturation Similarity Index (CSATSIM)

We will relate CSATVs using precisely the same formula as we did to relate SATVs (see section 2.4). Formally, similarity function CSATSIM is defined in figure 2.55. The SC-specific summary of CSATSIM values is given in appendix G. The value group matrix for all CSATSIM values (where WEIGHT = 1.2) of SCs of #2 to #10 is shown in figure 2.56. As we did with the other value group matrices, we will examine some patterns of CSATSIM(X , Y) values among different cardinality pairs. The smallest average CSATSIM value between # X and # Y SCs occurs where # X = # Y . The next smallest average value tends to occur between SCs of # X and # \bar{X} . The average CSATSIM value tends to increase the greater the difference between either $|\#Y - \#X|$ or $|\#Y - \bar{\#X}|$. For example, let us examine the #3 : # Y comparisons on the matrix. The closest average comparison (and also the smallest maximal and minimal CSATSIM values) is between SC pairs X and Y where # X = # Y = 3. The next smallest average value in this case happens to be between #3 and #4 SCs. The third closest average, however is between #3 and #9 SCs. The largest average CSATSIM comparison in the #3 value group is between #3 and #6, which constitutes the largest possible difference between either #3 or #9 and any other size SC.

We can determine from this that while CSATSIM does not produce an equivalence relation between complementary SC pairs, it does consider complementary cardinalities to be relatively close in terms of possible cyclic content. Obviously, if we change the WEIGHT value, the CSATSIM values will change; the distribution of relatively large and small values through all

value groups will, however, remain invariant. For this reason, we do not see the need to demonstrate CSATSIM using different WEIGHT constants.

2.14 Proportional CSATV(X) and Proportional SATSIM(X, Y)

For the sake of completeness, we can now create a one-part cyclic saturation vector using exactly the same protocol as with PSATV. In short, we simply deduct the each $cmin(w, c, i)$ value from the respective $WICCV_i$ value then divide the difference by the $cmax(w, c, i)$ value less the $cmin(w, c, i)$ value. This is expressed more formally in figure 2.56.

This vector (PCSATV) will produce largely the same sort of data as CSATV(X). The same equivalences will be yielded and, as was the case with SATSIM *vis-à-vis* PSATSIM, the values are very similar. Like PSATSIM, it is impossible to compare smaller than #3 SCs using the parallel function PCSATSIM because of the mathematical problem of dividing by zero.

2.15 Summary of new saturation vectors and similarity indices

Having just defined a group of new vectors and indices—particularly a group that uses relatively similar acronyms—we will now briefly revisit each one and recapitulate the type of subset class or interval class that each vector gauges:

Abbreviation:	SATV
Vector name:	Interval Class Saturation Vector
Element type used:	Interval classes
Description:	SATV is a dual six-argument vector in which the arguments represent the difference between the corresponding ICV place and the minima ($\min(c, i)$) or maxima ($\max(c, i)$) for any pcset of the same cardinality.

Abbreviation:	SATV2
Vector name:	#2 Interval Class Saturation Vector
Element type used:	#2 subsets
Description:	There is no difference in appearance between this vector class and SATV. The only difference is that $\text{SATV2}(X)$ is calculated from examining the numbers of embedded #2 SC X and $\text{SATV}(X)$ is calculated by examining the interval classes present between members of X .

Abbreviation:	PSATV
Vector name:	Proportional Interval Class Saturation Vector
Element type used:	Interval classes
Description:	This six-argument vector is derived by subtracting the apropos $\min(c, i)$ value from the corresponding ICV place; that difference is then divided by $\max(c, i)$. The result represents the proportion of each ic found in set X compared to what is possible in any set of $\#X$.

Abbreviation:	PSAT%V
Vector name:	Proportional Interval Class Saturation Percentage (%) Vector
Element type used:	Interval classes
Description:	This six-argument percentage vector (see section 1.3.2 on Castrén's $nC\%$ -vector) which is derived by dividing each argument of $\text{PSATV}(X)$ by the sum of the $\text{PSATV}(X)$ arguments (i.e., the cardinality of $\text{PSATV}(X)$).

Abbreviation:	SATVn
Vector name:	$\#n$ Subset Class Saturation Vector
Element type used:	$\#n$ subsets
Description:	This is a generalization of SATV2 to any cardinality subset class.

Abbreviation:	CSATV
Vector name:	Cyclic Saturation Vector
Element type used:	Cyclic adjacencies
Description:	<p>In this dual six-argument vector, the amount of each ic i is weighted based upon the number of the pcs in the set which are i-cycle adjacencies. If two sets X and Y each have four $ic3$s and, in set X, six pcs participate in forming the four $ic3$s, while in set Y, only four pcs participate in forming the four $ic3$s (that is, the latter set contains an unbroken 3-cycle while the former set contains 3-cycle fragments), set Y's $ic3$ content will be weighted more heavily. Like the other saturation vectors, each weighted argument in the CSATV is compared to the corresponding minimal ($cmin(w, i, c)$) and maximal ($cmax(w, i, c)$) values for any set of the same cardinality.</p>

The similarity indices introduced in chapter 2 are summarized as follows:

Abbreviation:	SATSIM
Index name:	Interval Class Saturation Similarity Index
Vector class used:	SATV
Description:	SATSIM compares two SATVs, first by comparing each argument in $SATV_A(X)$ to $SATV_{row}(Y)$, then by comparing each argument in $SATV_A(Y)$ to $SATV_{row}(X)$ (see figures 2.9 and 2.10).

Abbreviation:	PSATSIM
Index name:	Proportional (IC) Saturation Similarity Index
Vector class used:	PSATV
Description:	PSATSIM utilizes a simple ASIM-style index to compare two PSATVs.

Abbreviation:	PSAT%SIM
Index name:	Proportional (IC) Saturation %-Vector Similarity Index
Vector class used:	PSAT%V
Description:	PSAT%SIM also utilizes an ASIM-style index to compare two percent vectors.

Abbreviation:	SATSIMn
Index name:	Cardinality class n Saturation Similarity Index
Vector class used:	SATV n
Description:	SATSIM-style index, generalized to allow any cardinality class of subsets. For example, SATSIM3 compares two SCs through an examination of their trichord-class content.

Abbreviation:	TSATSIM
Index name:	Total Saturation Similarity Index
Vector class used:	SATV n
Description:	TSATSIM uses all SATV n vectors (from $n = 2$ to $n =$ the smaller of $\#X$ and $\#Y$ minus 1) to create a single “total” measure (to use Castrén’s terminology) which evaluates all subsets of all applicable cardinality classes using a SATSIM-style comparison.

Abbreviation:	AvgSATSIMn
Index name:	A verage of SATSIMn comparisons
Vector class used:	SATV n
Description:	Like TSATSIM, AvgSATSIM n uses all SATV n vectors (from $n = 2$ to $n =$ the smaller of $\#X$ and $\#Y$ minus 1) to create a single “total” measure. Unlike TSATSIM, which simply adds together all the differences in the SATV n s and divides that amount by the combined cardinality of each SATV n , AvgSATSIM n is calculated by taking the arithmetic mean (average) of each applicable SATSIM n comparison.

Abbreviation:	CSATSIM
Index name:	C yclic S aturation S imilarity Index
Vector class used:	CSATV
Description:	CSATSIM utilizes a SATSIM-style index to compare two CSATVs.

Chapter 3

A comparative analysis of Stravinsky's *Three Pieces for Clarinet Solo*, first movement (1920) using similarity indices and pcset networks.

In the first two chapters, we reviewed various similarity indices constructed for pcset analysis and provided a detailed account of a number of new saturation-based indices. In the present chapter, we will provide several samples of how similarity relations in general and saturation-based indices in particular might be used analytically. We now attempt to place our new measures into a specific context with those of other theorists. Our vehicle for this comparison is the first of Stravinsky's *Three Pieces for Clarinet Solo*. Using a variety of similarity indices (both new and old), we compare its musical segments. Our aim is to see how the different procedures of comparing pcset classes might lead toward different—even conflicting—readings, and, where conflicts exist, to examine the underlying conceptual differences inherent in the measures' construction.

3.1 Musical segmentation

The first movement of Stravinsky's *Three Pieces for Clarinet Solo* is given in example 3.1. It has been grouped into fifteen non-overlapping segments, based primarily upon Stravinsky's placement of breath marks or rests, and using his placement of slurs as a secondary determinant. There are several places where breath-mark-delimited segments contain more than one slurred segment. In those cases, the slurred groups were either very short (2-3 notes) or, in the case of mm. 16-17, contained exact or near-exact repetitions. There

are few places where slurs rather than breath marks delimit the segmental boundaries (e.g., mm. 3, 22, and 25). Only once in the piece (m. 3) does a segment extend across a breath mark (largely because of the rest in m. 3). In this case, one could place the boundary at the breath mark in m. 2 without altering the results of a pcset-based analysis.¹¹⁴ This is because the $\langle G\#, F\#, G\# \rangle$ sequence in mm. 2-3 duplicates pcs in both the preceding (mm. 1-2) and proceeding (mm. 4-5) slurred groups.

The pcsets formed by the fifteen segments are members of only ten distinct set classes, labeled **A** through **J**. The prime form, Forte label, and interval-class vector of each SC are provided in figure 3.1. There are two pentachord classes, two hexachord classes, four septachord classes, and two octachord classes. Of the ten distinct SCs, two are abstract complements of each other (**A** and **H**—5-23 and 7-23); none of the SC pairs are *Z*-related.

We make no claim that this particular segmentation is uniquely percipient. While it is the segmentation that we prefer, we fully acknowledge that some readers may disagree with our reading of the piece, hearing segmental divisions in different locations, and perhaps even desiring overlapping segments in some cases.¹¹⁵ Such alternate readings may, in fact, also yield interesting relationships. For example, we assert a segmental boundary at m. 22, beat 4, despite the absence of a breath mark. There is, however, the termination of a phrase coincident with a *decrescendo* marking. Two notes

¹¹⁴While an analysis that examines common pcs, common pitches or pitch sets in pitch space (that is, considering register as well as pitch class), contour, and a host of other musical factors would potentially be much more informative, we have, for the sake of both brevity and relative conciseness, chosen to undertake only a pcset-class-based analysis herein.

¹¹⁵For example, one might hear Stravinsky's use of different register as a stronger segmental determinant than his placement of breath marks and slurs. A segmentation based upon this criterion would yield a very different, though likely complementary, reading of this piece.

later, however, there is also a phrase boundary. Had we decided to group from the breath mark in m. 21 to the phrase boundary at the beginning of m. 23, we would have another **H**-type segment (SC 7-23). Similarly, in m. 25, we could have chosen not to draw a segmental boundary (since again there is no breath mark) without significantly altering our results. The **B**-type segment in mm. 24-25 is literal subset (in p-space as well as pc-space) of the **I**-type segment in mm. 25-28. This **B**-type set uses the same pitches in the same register (but in a different order and rhythm) as the initial **B**-type set in mm. 5-6. Had we created a single segment **I** stretching from the breath mark in m. 24 to the one in m. 28, the close relationship between the two **B**-type sets would have been obscured.¹¹⁶

Clearly, further comparison between a variety of segmentations of this piece is analytically fruitful. Such a comparison, however, lies outside the scope of this study. By choosing a single segmentation, we are able to focus our attention on an examination of how a variety of different resemblance criteria group the various segments into clusters of related pcsets. With this in mind, any appropriate and musically-sensitive segmentation would serve equally well for the present purposes. We will further discuss the issue of segmentation in atonal analysis in chapter 4.

3.2 An abstract inclusion network model for analysis

A network of abstract embeddings is provided in figure 3.2. Given two SCs X and Y , where X is an abstract subset of Y ($AS(X, Y) = \text{true}$), X and Y have been

¹¹⁶More importantly, though, a single segment extending from the breath mark in m. 24 to the one in m. 28 was unconvincing to me as I listened repeatedly to the piece. Of course, it is entirely possible that a different performance would lead me toward a different segmentation. The performance on which I based this segmentation was by Walter Boeykens (CD 901356, Harmonia Mundi France, 1991).

connected by an arrow pointing from the smaller set to the larger. Because simple (bivalent) inclusion relations are transitive, any such X and Y are implicitly connected on figure 3.2.¹¹⁷ As one can see from examining this table, **A** is abstractly embedded in (and therefore connected to) six of the eight larger (#6) pcset classes, including its abstract complement. The only two larger pcsets that it is not embedded in are **B** (6-z39) and **J** (7-8), the final SC (**B** and **J** are connected to each other). Additionally, there are two other SCs that are well connected in our structure: **C** (6-8) and **G** (8-22). **C** is the only hexachordal SC which embeds **A** and is embedded by both **H** and **I**. **G** embeds all smaller SCs in the piece except the very last (and arguably most enigmatic) one, **J** (7-8). Both **C** and **G** are in a Kh pcset complex with the complementary pair **A/H**.

From the standpoint of an inclusion-based analysis, **J** appears to be the most remote pcset and **A**, **C**, **H**, and **G** appear to be the most connected pcsets.¹¹⁸ This would suggest that subset-based similarity indices (such as Rahn's ATMEMB, Lewin's REL,¹¹⁹ Castrén's RECREL, and our own AvgSATSIM n) should find **J** to have the fewest segments closely related to it. By contrast, the most connected sets in figure 3.2 should yield the greatest number of close relations when using a subset-based similarity index. Not surprisingly, this prediction is, for the most part, borne out by the data. Before discussing the actual measures, though, we can refine our superset/subset complex to make it a bit more reflective of what the similarity indices are actually measuring. Unlike our AS(X, Y) or Forte's Rp or K/Kh

¹¹⁷More formally: iff AS(X, Z) and AS(Z, Y) then AS(X, Y).

¹¹⁸This viewpoint is strengthened by the musical realization of set **J**: it is the most unique in terms of range, dynamics, and types of adjacent intervals.

¹¹⁹Here the TEST group of subsets is defined liberally to include all SCs smaller than the smaller of the two SCs being compared.

relations, none of the above-mentioned indices deal with abstract inclusion as a binary function.¹²⁰ Rather, *ATMEMB*, *REL*, *RECREL*, and *AvgSATSIM_n* all measure the degree of inclusion of the smaller pcset into the larger. In other words, they use Lewin's *EMB* function (or some equivalent) to determine their data in part. Accordingly, in figure 3.3 we refine our figure 3.2 graph to show how many of each smaller pcset are included in each larger pcset. The reader will notice that there are more connecting lines between the SCs in figure 3.3 than in 3.2. This is because abstract inclusion is transitive. The relation between, for example, **G** and **C** is implicit in figure 3.2 because **H** is embedded in **G** and **C** is embedded in **H**. The extra connecting line between **G** and **C** is necessary in figure 3.3 because the *number of embeddings* found between pairs of SCs is not transitive.¹²¹

We could create a rather similar model using Lewin's *COV(X, Y)* function as well. *COV(X, Y)* returns the number of distinct forms of set class *Y* that embed a particular form of set class *X*.¹²² If *X* and *Y* share the same degree of symmetry, *COV(X, Y)* returns the same value as *EMB(X, Y)*. If, however, the smaller SC is not symmetrical, *EMB* will frequently yield a larger number than *COV*, and vice versa. For example, let us consider the differences between *EMB(A, C)* and *COV(A, C)*. The number of distinct forms of **A**[02357] embedded in any one form of **C**[023457] is two ($\{0,2,3,5,7\}$ $\{0,2,3,4,5,7\}$)

¹²⁰In a binary function that measures inclusion, one SC either is or is not a superset or subset of the other.

¹²¹For example, $EMB(C, H) = 1$, $EMB(H, G) = 1$, and $EMB(C, G) = 1$. But, by contrast $EMB(A, I) = 2$, $EMB(I, G) = 1$, and $EMB(A, G) = 5$. The number of **C** in **G** is not predictable from $EMB(C, H)$ and $EMB(H, G)$; neither is the number of **A** in **G** predictable from $EMB(A, I)$ and $EMB(I, G)$. We know, for example, that there is one **C** in **H** and one **H** in **G**; there is, therefore, *at least* one **C** in **G**, but there may be other **C**s in **G** that are not members of the single embedded **H**. Similarly, if there were multiple members of **H** in **G**, we still could not predict the number of **C**s in **G** because two (or more) of the members of **H** might share the same member of **C** its subset, confounding any simple equation.

¹²²*COV*, like *EMB*, is introduced in Lewin 1977.

and $\{0,2,4,5,7\} \subseteq \{0,2,3,4,5,7\}$). There is, however, only one distinct form of **C** that covers any particular form of **A** ($\{0,2,3,4,5,7\} \subseteq \{0,2,3,5,7\}$). A COV version of figure 3.3 is presented in figure 3.4.

While $\text{COV}(X, Y)$ and $\text{EMB}(X, Y)$ return the same value when pcsets X and Y share the same degree of symmetry, the two functions do measure fundamentally different things. Similarity indices could well be constructed based upon coverings instead of—or even in addition to—embeddings, and this possibility will be investigated in chapter 5.¹²³ In this present chapter, however, we will only be discussing indices that in some way consider embeddings. COV is therefore a less appropriate measure of comparison for our current purposes.¹²⁴

In the spirit of the saturation-based embeddings introduced in chapter 2, we can now contextualize the EMB values in figure 3.3. Figure 3.5 is a reworking of figure 3.3 using proportional SATEMB (PSATEMB) values, expressed as a fraction (where the numerator is the $\text{EMB}(X, Y)$ number and the denominator is the maximum possible number of X included in any SC of cardinality Y —denoted $\max(\#Y, X)$). Because in our particular collection of SCs the minimal number of any SC X included in any SC Y ($\min(\#Y, X)$) is 0, we do not need to adjust for the minimal numbers the way that we did in calculating the proportional SATV (PSATV).

Figure 3.5 contains five cases of maximal embedding: there is the maximal amount of **C** in **H** (1 of 1 possible in any set of $\#\mathbf{H}$), **C** in **I** (1 of 1), **B** in **J** (2 of

¹²³Formally, the relationship between EMB and COV can be described using complementary sets: $\text{COV}(X, Y) = \text{EMB}(\bar{X}, \bar{Y})$. This is detailed in Morris 1990, 181.

¹²⁴In our SATSIM-based analysis, we will consider SCs and their complements in tandem. If every similarity index that we will examine also paired complementary SCs, then comparisons to COV and EMB would be equally appropriate.

2), **A** in **C** (2 of 2), and **A** in **E** (4 of 4). There are two other cases where the embedding number represents a greater than 50% saturation of the smaller pcset into the larger one. We can say that there is an “almost maximal” saturation of **A** in **H** (3 of 4) and **A** in **G** (5 of 6). Figure 3.6 is a superset/subset chart which features only maximal or near-maximal embeddings. In this chart, we can see that the attachment between the **B/J** group and the center group has been severed. All the SCs in the center group are still attached, though not all to each other. It now seems that SC **A** is the nexus of the center group, given that it is the only one which is maximally or near-maximally related to SCs **C**, **E**, **H**, **I**, and **G**. By contrast, it would now seem that **J** is *not* the most detached SC from the superset/subset structure; rather that title falls to the two SCs on the left of our figure (**D** and **F**), which have had all ties to either of them severed. There are two **D**s in **F** ($\max(\#\mathbf{F}, \mathbf{D}) = 4$), but the two connections that the left group had with the center group are shown to be the weakest two links in our pcset network ($\text{PSATEMB}(\mathbf{D}, \mathbf{H}) = \frac{1}{3}$ and $\text{PSATEMB}(\mathbf{A}, \mathbf{F}) = \frac{1}{6}$).

This saturation-based embedding complex could easily be seen as an extension of fuzzy set theory to the realm of abstract subset and superset relations.¹²⁵ While we are not precisely invoking the concept of partial inclusion of an element into a group, we are allowing for various degrees of inclusion dependent on both how many members of SC *X* are embedded in *Y* and how many members of SC *X* could there possibly be in any pcset of #*Y*. The question has evolved from “is SC *X* embedded in SC *Y*?” to “how many

¹²⁵This connection was first suggested to me by Donald Sloan. Ian Quinn has examined the connections of fuzzy set theory and musical analysis in “Fuzzy Transposition of Pitch-Class Sets,” delivered at the Society for Music Theory 1996 annual meeting.

members of SC X are embedded in SC Y ?” and further now to “to what degree is SC X embedded in SC Y ?”. From this example, it would be only a small step to generate “fuzzy” (or saturation-based) K and K_h tables where each K or K_h label would be followed by a fraction or percentage that would show the degree of the saturation of one SC (pair) in the other.¹²⁶

3.3 Resemblance determined by similarity indices which measure interval-class (#2 subset-class) content

With the examples of how the SCs in the first movement of this piece cluster together in both a simple inclusion graph and a saturation-based inclusion graph as our points of departure, we will now examine how the various means of judging resemblance suggest clusters of closely related SCs. All ten SCs have been compared to each other using eight different similarity indices: ASIM, IcVSIM, SATSIM2, ATMEMB, REL, RECREL, AvgSATSIM n , and CSATSIM; the first three are based upon interval-class content, the next four are based upon total shared abstract subset content, and the last is based upon shared content of i -cyclic pcset classes. We will examine the data provided by all these measures and briefly discuss possible similarity index-based analytical narratives. Also, when appropriate, we will consider what factors seem to be at play in determining which SCs are relatively closely related. “Closely related” is, of course, a nebulous term that is difficult to define in *any* particular measure, to say nothing of the problems in defining the concept in a comparable way in eight different measures. For each measure, we could take a number of approaches: we could examine the average values for all sets of

¹²⁶An analogous method has been posited by Robert Morris in his discussion of “multiple complementation” and the degree of abstract complementation (DCOM). (Morris 1990, 196.)

cardinality 5 through 8 (this is the range of cardinalities that we are dealing with in our segmentation of the Stravinsky piece), or the average for all sets that are actually used in the piece. Rather than the average, we decide that the mean value is a better divider between similar and dissimilar pairs.¹²⁷ Alternatively, we might examine the distribution of values and look for naturally-occurring dividing points, or we might simply choose a round number in the range of what (subjectively) seems to us to be similar.¹²⁸ In examining the data provided by each of similarity indices listed above, we will take note of the average values for the group, but will examine the entire range of values and choose relatively large gaps between values as our “naturally-occurring” dividing points to distinguish similar from dissimilar set pairs. For each index, we will examine a number of different dividing points (or filters) in an attempt to understand relative degrees of similarity. While this technique may not allow us to draw very fine distinctions, it will enable us to compare how one might analyze this piece using each of these different measures of resemblance.

3.3.1 Morris’s ASIM

Figure 3.7 gives a comparison matrix of these sets using Morris’s ASIM index. Figure 3.8 provides a list of the values from figure 3.7, sorted from most to least similar. We can see from this list that **E** and **H** are, by far, the most closely

¹²⁷If, for example, a particular index produces a small number of values that are very high, with the rest clustered in a much lower range, a comparison of individual values with the average for the group might make a large percentage of the sets being examined appear similar. Using the mean as the determinant of similar and dissimilar produces a different—but still an arbitrary standard—inferring that half the sets in the group are similar, and the other half are not.

¹²⁸In such cases, an analysis would be colored not only by an index that is based upon an ontologically-subjective interpretation of similarity, but also by an empirically-inter-subjective interpretation of these values.

related pair. In general the values increase rather dramatically both at the top and bottom of the list. The values in the middle of the list generally are separated by smaller increments.¹²⁹ After the gap between the most closely related pair (ASIM = 0.048) and the second-most closely related pair (ASIM = 0.095), the next largest gap is between ASIM values 0.107 and 0.133. If we filter out set pairs that yield ASIM values greater than or equal to 0.133, we are left with the relationships shown in figure 3.9. On the diagram at the right of this figure, sets which meet this standard for relatedness are connected with a thin line segment. The left of figure 3.9 shows the “transitive tuples” formed by this particular standard. A transitive tuple is a group of sets that are all related to each other by some standard of resemblance. In this particular case, all the tuples contain either one or two elements. Larger groups of transitive tuples (particularly three elements—#3—and larger) will be formed as we widen our lens of similarity. Such groups help us detect families—or clusters—of related set classes.

Because so few set pairs fall within this particular similarity cutoff point, it is difficult to make many generalizations about them. We might notice that **E**, **H**, **I**, and **J** are all pitch-class septachords and that **F** and **G** are both octachords. Each of these sets occur (roughly) in the second half of the piece, and none of the sets at the beginning of the piece are related to any others. Both **I** and **J** and **F** and **G** are temporally-adjacent in the music; **E** and **H** and **H** and **I** are not, however. As we mentioned earlier, there are no groups of three or more sets that are all related to each other using ASIM = 0.133 as our cutoff mark. Such groups can be found if we widen our focus a bit and include all those

¹²⁹This is true in most of the eight similarity indices we will examine.

ASIM values that fall below 0.167, our next salient cutoff point. Within this newly-defined range of similarity is one set pair (**B/C**) that is related by the ASIM value 0.133, and nine pairs that are related by the ASIM value 0.143. The transitive tuples are bracketed on the left side of figure 3.10. Again, each pcset within a pair of brackets is closely related to all the other pcsets within the same pair of brackets. On the right of figure 3.10, line segments connect all closely related pcset pairs.

We can see quite clearly that almost all six sets displayed in figure 3.9 are closely related to each other in figure 3.10. The only exceptions are the set pairs **E/F** and **E/J**, which are not closely related using this cutoff value. Sets **E, I, G**, and **H** are all similar to one another, as are sets **F, G, H, I**, and **J**, forming #4 and #5 transitive tuples, respectively. **B** and **C** are also related using this filter, but are distinct from the other sets that we mentioned.

We might now ask what sets **E, I, G**, and **H** have in common. Looking back to figure 3.2, we can see that **G** is the common superset of **E, I**, and **H**. This would seem to be an expected relation. Discovering, the relationship between **F, G, H, I**, and **J** is, however, a bit more difficult. As we just mentioned, **H** and **I** are mutual subsets of **G**. **F** and **G** are both #8 sets, and so we cannot assert an embedding relationship between them. We can, however, remark that they have no mutual subsets other than the pentachord **A** in this group. **J** (7-8) is a subset of neither of the octachords (**F** or **G**), and is connected in no obvious way on the inclusion charts to **H**. It would seem, therefore, that the closest ties one can assert among these groups is that of cardinality. These six sets are all septachords or octachords and there are no septachords or octachords

used in the piece that are *not* in these groups. Similarly, **B** and **C** are the only hexachords in the piece and they too are connected on figure 3.10.

Let us now widen our focus even further and look at the ASIM values 0.167 and 0.2. These are shown in figure 3.11. Looking first at the additional connections provided by moving the filter up to 0.167 (these are shown in medium-thick lines), we see that **B** and **C** are now connected to our larger ASIM structure, forming a transitive tuple with **I** and **H**. Additionally, **B** forms a transitive tuple with set **J**. **H**, **I**, and **J** are three of the four septachords in this piece. Neither **B** nor **C** (again, our two hexachords) are connected to either of the two octachords (**F** and **G**).

The very thick lines in figure 3.11 represent ASIM values greater than 0.167 and less than or equal to 0.2. There are only four set pairs in this range. **E** is now connected to both **F** and **J**, completing a large transitive tuple {**E**, **F**, **G**, **H**, **I**, **J**} that includes all the seven- and eight-note sets in the group (and none of the smaller ones). Additionally, **A**—the first set in the piece—is finally connected to two others, **C** and **D**. However, as in the other ASIM connections, the primary determinant for these two also appears to be cardinality: **A** and **D** are the only two pentachords and **C** is one of the two hexachords.

In figure 3.12, we examine these relations among the sets with regard to Stravinsky's temporal realization. Using Morris's index, some patterns may be seen to emerge: the last five sets (**C3**, **H**, **B2**, **I** and **J**) are related to each other. One can draw a parallel between the end of the piece and mm. 5-9 where sets **B** and **C** first appear. The middle sets, **E** (which is repeated), **F**, and **G** are also connected to each other as part of the largest ASIM transitive tuple shown in

figures 3.10 and 3.11. One might say that **C3** initiates a final section of closely related sets. This observation allows one to draw a rather close parallel to **A1**, the first set of the piece. Both **A1** and **C3** begin with the same pitch classes—and **C3** is a literal superset of **A1** ($C3 = \{A1 + A\# \}$).

With all the connections that *are* asserted using our most inclusive filter (ASIM = 0.2), there are a few seemingly-obvious connections that are surprisingly omitted. Most notably, the connection between **A** and **H**—our two complementary set classes. The ASIM value between these two is 0.355—a number large enough that it would not be included in even the broadest imaginable conception of ASIM similarity. Set classes **A** and **G** are even more dissimilar using ASIM, yielding a value of 0.474, despite the fact that **G** (8-22) embeds an almost maximal number of **A**-type sets (5-23) (see figure 3.5).

To get a slightly different view of the ASIM connections, figure 3.13 provides a chart that individually shows the ASIM = 0.2 connections of each set in the group. From this information, we can easily see that the SCs with the most connections to other SCs in the group are **H** and **I**, with **J** coming in a close second place. The SCs with the fewest connections to others in the group are **D** and **A**. The clusters of related SCs are much different using ASIM than they were using a SC inclusion model. In the latter, we saw that **D**, **F**, and **J** were the sets with fewest connections; in the former, **J** is one of the sets with the most connections. The two systems agree only that **D** is a remotely-connected SC in the group.

One might conclude that a system of measuring SC resemblance based upon ic content simply produces much different results from a system based upon

abstract inclusion. There is, however, a confounding factor at play in this conclusion: ASIM does not neutralize the effect of different cardinalities as strongly as do some other ICV-based indices. Indeed, if we take cardinality into account when examining figures 3.7 and 3.13, we see that **A** and **D**, the two sets with the fewest connections within the group, are the only #5 sets. The two #6 sets, **B** and **C**, have the next-fewest connections, while the most connected three sets, **H**, **I**, and **J** are all #7 SCs. This, it would seem, is their strongest tie to each other and to the others in the group. As we mentioned, any connection between **A** and **H**—the two abstract complementary sets in the group is notably absent. **A** is connected only to other relatively small sets and **H** is connected only to other relatively large sets.¹³⁰

3.3.2 Isaacson's IcVSIM

To see how another interval-class-based index interprets the same data, we will now register the same comparisons using Eric Isaacson's IcVSIM. IcVSIM, which derives its data from the standard deviation of the difference between two ICVs, is a more cardinality-neutral measure than is ASIM, which also examines the differences between two ICVs, but which adjusts for cardinality differences by dividing the sum of the differences by the total number of intervals present in both sets (combined). Figure 3.14 gives a comparison matrix of these same pcsets using IcVSIM. The information in this matrix is sorted into a list of most- to least-closely related set pairs in figure 3.15.

¹³⁰In fact, **H** is connected to *all* #6 through #8 SCs in the group and to *neither* of the #5 SCs in the group.

One of the largest gaps in the sorted list of IcVSIM values is between 0.373 (the smallest present IcVSIM value) and 0.577. Despite the small number of sets related by $IcVSIM = 0.373$, this will be where we draw our first distinction. Figure 3.16 shows transitive tuples (in brackets on the left and diagrammed on the right) where $IcVSIM = () 0.373$. There are only three sets in the diagram: **H** is connected to both **A** and **G** (and **A** is not connected to **G**). In contrast to ASIM, IcVSIM finds the only complementary set classes (**A** (5-23) and **H** (7-23)), as well as the Kh-related (**H** and **G** (8-22)) to be the most closely related sets, despite their differences in cardinality.¹³¹

Widening our focus a bit, we see that the next relatively large gap in IcVSIM values (in figure 3.15) is between 0.577 and 0.687. There are six set pairs that are related by $IcVSIM = 0.577$. Figure 3.17 adds these connections to those seen in figure 3.16. In this new graph, we now see that a connection is asserted between **A** and **G**, completing a three-element transitive tuple with **H**. There are also five new connections among set pairs. This is fairly well in accord with what we observed in our inclusion graphs at the beginning of this chapter. Every connection in the graph in 3.17 is also present in figure 3.2, except for the IcVSIM-asserted resemblance between **E** (7-35) and **H** (7-23), which could not be shown in an inclusion graph since they are the same size. In fact, each of the IcVSIM relations among different-sized sets is related by a

¹³¹The relation of **A** to **H** is not surprising given that two complementary sets of cardinal A and B will always differ by the same amount. For example, the difference between the ICVs of any two complementary sets' of #5 and #7 will be $\langle 2,2,2,2,2,1 \rangle$. The degree of deviation between these numbers is obviously quite small, hence the very low IcVSIM value. The greater the difference in the cardinality of two complementary sets, the higher the deviation between the numbers in the difference vector, the less similar IcVSIM will find the set pair. Complementary sets of #4 and #8 are related by $IcVSIM = 0.745$; complementary sets of #3 and #9 are related by $IcVSIM = 1.118$; and complementary sets of #2 and #10 are related by $IcVSIM = 1.491$, a distant relation.

proportional SATEMB value of at least 0.5 (i.e., at least 50% saturated); these can be reviewed in figure 3.5.

If we expand our focus outward a bit, to include IcVSIM ≥ 0.689 relations (this is where the next largest gap in our sorted table occurs), we see more three-element and larger transitive tuples falling into place. These are shown in figure 3.18. The upper half of the graph contains two heavily-overlapping four-element transitive tuples: $\{\mathbf{A}, \mathbf{C}, \mathbf{G}, \mathbf{H}\}$ and $\{\mathbf{A}, \mathbf{E}, \mathbf{G}, \mathbf{H}\}$. There are also several new three-element transitive tuples. \mathbf{B} and \mathbf{J} remain unconnected to \mathbf{I} (a connection that one might expect, looking at the inclusion networks¹³²), and \mathbf{D} remains unconnected to *any* other SC in our segmentation of the piece. Figure 3.19 broadens the focus just a bit farther to include all IcVSIM values ≥ 0.898 (and, coincidentally, all IcVSIM values that are more similar than the average for the group [0.951]). This new filter finally connects \mathbf{D} (5-14) to one other SC (\mathbf{A} (5-23)), but \mathbf{D} is still not considered similar to any of the set classes which cover it (these are \mathbf{F} , \mathbf{G} , and \mathbf{H}). All sets in the piece now belong to least one two-element or larger transitive tuple. \mathbf{D} and \mathbf{J} are, by far, the least connected sets and \mathbf{A} is the most connected set. There is one five-element transitive tuple, formed by sets $\{\mathbf{A}, \mathbf{C}, \mathbf{G}, \mathbf{H}, \text{and } \mathbf{I}\}$, each of which is found in the middle of the inclusion graphs presented earlier in this chapter. A chart showing the temporal arrangement of the SCs, with connections drawn between adjacent segments where IcVSIM ≥ 0.898 is found in figure 3.20. There are fewer connections asserted here than in the parallel ASIM diagram in figure 3.12. Except for the connection between \mathbf{C} and \mathbf{G} , claimed by IcVSIM, all the connections in the present example are also found in the

¹³² $\mathbf{I} - \mathbf{J}$ was one of the closest ASIM relations, though this close relation is likely caused by ASIM's treatment of cardinality.

ASIM example. Figure 3.21 provides a chart that individually shows the IcVSIM 0.898 connections of each set in the group. Using this information, combined with that in figure 3.20 above, we can make a few narrative remarks about the progression of sets in this piece. The piece begins with set **A**, the most connected set in the group and ends with set **J**, one of the two least-connected sets in the group. Set **D**, the other least-connected set is also realized in a rather distinct manner and, like **J**, suggests a sectional division. One might also interpret the IcVSIM data to indicate a reading of the piece that features connections between adjacent sets at its beginning and ending, with a series of distinct (from each other) sets in the middle. The set progression $\langle \mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D} \rangle$ at the beginning of the piece might also be compared to the progression $\langle \mathbf{H}, \mathbf{B}, \mathbf{I}, \mathbf{J} \rangle$ at the end of the piece. In both cases, there is no connection between the first and second set, or between the last and penultimate sets; there is, however a connection between the two interior sets. Adding credibility to this reading is the fact that the initiating sets of each group, **A** and **H**, are abstract complements and the pair which yields the lowest IcVSIM relation in the comparison matrix. **B**, **C**, and **I**—the sets that form the interior groups of these outer sections—are also closely related, comprising a transitive tuple using the IcVSIM 0.898 filter. The comparison breaks down at the end, though: sets **D** and **J** are not closely related, but one could say that they perform a similar role as sets that are unique (in both their interval-class content and dramatic realization) compared to what precedes them. We will discuss this connection a bit more in our SATSIM2-based reading of the piece.

3.3.3 SATSIM2

The final interval-class-based measure that we will examine is SATSIM2, the interval-class (dyad-class) saturation similarity measure, first introduced in chapter 2. Like IcVSIM, SATSIM2 is a relatively cardinality-neutral index. Unlike IcVSIM, it largely adjusts for cardinality before the formal comparison is initiated (through the use of saturation vectors). Figure 3.22 provides a SATSIM2 comparison matrix of the pcsets in this piece. The information in this matrix is sorted into a list of most- to least-closely related set pairs in figure 3.23.

One of the largest gaps in the sorted list of SATSIM2 values is between 0.041 and 0.095. Because both of these numbers are relatively very low (reflecting strong similarity) and because each is derived from only a single comparison, we will choose to make our first (and still very conservative) division between the SATSIM2 values 0.102 and 0.143. Since complementary pcsets are maximally similar using SATSIM2 (yielding a value of 0.000), **A** and **H**, the only such pair in our segmentation, will be considered in tandem in the transitive tuples represented in the following four figures.¹³³ Figure 3.24 shows transitive tuples (in brackets and diagrammed graphically) where $\text{SATSIM} = () 0.102$.

This figure illustrates that **E** (7-35, the diatonic collection) and **A/H** are closely related. Additionally, both sets **C** and **B** form close relations with **I**. These relations are corroborated by the proportional SATEMB network in figure 3.5. From the very first cutoff point in our SATSIM2 analysis, we can already begin asserting the symmetry between the beginning and ending of this piece

¹³³This is because $\text{SATV2}(\mathbf{A}) = \text{SATV2}(\mathbf{H})$.

that was discovered in our IcVSIM analysis. Two realizations of **A** initiate the first section; the last section is initiated by **H**, its complement. Both sets progress to the dissimilar set **B**. At the end of the piece, **B** moves to **I**, a similar set class; at the beginning of the piece, **B** moves to **C**, not a similar relation using SATSIM2 = 0.102 as our cutoff, but the connection between **B** and **C** is established at the next highest SATSIM2 value in the matrix (0.143)—which leads us to our next cutoff point.

Figure 3.25 provides a list and a graphic illustration of the transitive tuples formed by set pairs related by SATSIM2 = 0.163. There are two #3 transitive tuples (counting **A** and **H** as a single member) and two #2 transitive tuples. With the single exception of **D**, every set in the diagram has at least two close relations. **A/H** and **C** are the most well-connected sets in the graph. **F** and **J** are the only two sets that are not deemed similar to any others in the piece. These relations, and most others using this filter, are predictable from the embedding charts at the beginning of this chapter. In fact, every maximal or “near-maximal” proportional SATEMB relation (these are shown in figure 3.6) is present in this reading of the SATSIM2 values as well. This should show, perhaps even more definitively than the analysis of IcVSIM values, that subset-based measures of similarity *can* also be effectively modeled using an interval-class based system of resemblance (naturally, this is not always the case).

We will now move our cutoff point up to include SATSIM2 values that are less than or equal to 0.190. The transitive tuples formed using this filter are shown in figure 3.26. For the first time, a relatively conservative view of similarity within a particular measure has all sets connected by transitive tuples

of at least three elements. There are only five transitive tuples (all of which are at least partially interlocking), facilitating a classification of the sets into rather distinct families.

The next reasonable dividing point is at SATSIM2 = 0.205. These SC pairs and the transitive tuples that they form are shown in figure 3.27. There is not much in figure 3.27 that is not also in figure 3.26. In fact, our charts of the temporally-adjacent sets—found in figure 3.28 (a or b)—are the same using either SATSIM2 = 0.190 or SATSIM2 = 0.205 as our cutoff. Returning to the comparison between the first and last four set types, we can see (on figure 3.28) that there is still no relation between **A** and **B** or **H** and **B**, the first two sets in each group.¹³⁴ However, there are now connections between the last two sets in each section: **C** — **D** and **I** — **J**. Additionally, there are close relations between each penultimate pair as well: **B** — **C** and **B** — **I**. Recall that in our description of the same passage using SATSIM2 = 0.102 as our cutoff value, we were able to assert the connection between **B** and **I**, but not **B** and **C** and that the last two sets in both sections were not closely related. The middle of the piece—from the realization of set class **C** that begins in m. 11 (**C2**) until the realization of **C** in mm. 21-22(**C3**)—can be described using these SATSIM2 values as a palindrome of similar set types. (A chart that individually shows the SATSIM2 = 0.205 connections of each set in the group is provided in figure 3.29.) Working inward from sets **C2** and **C3** in figure 3.28 (a): the two realizations of **E** are quite similar to **G**; **F**, at the center of our palindrome, is similar only to itself when compared to any of the other

¹³⁴Actually, given that there is no SATSIM2 relation between **A** and **B**, it follows trivially that there is no relation between **H** (the complement of **A**) and **B**.

sets in this middle section; and the two realizations of **C2** and **C3** are set-class equivalents and are therefore similar under any pcset based similarity index. We can expand our palindrome if we now work outward on figure 3.28 (a) into the beginning and ending of the piece: **D** is similar to **H**, **C(1)** is similar to **B**, and **B** is similar to **I**. The palindrome stops here, at the penultimate step. **A**, the first set class in the piece, is *not* similar to **J**, the last set class in the piece. This disjunction at either end certainly is aurally salient, emphasized by the extreme contrast of dynamics, timbre, rhythm, and adjacent intervals. The last set in the piece in many ways seems not to “belong” with the rest of the movement. Indeed, we should remember that we are only looking at the first movement of a set of three. The material in the last two measures (SC **J**) seems to point toward the next—very different—movement, as well as refer back to the much milder disruption in textural flow formed by Stravinsky’s realization of SC **D** in mm. 10-11. Perhaps, then, it is more musically convincing to read the linear flow of the outer sections as similar and the inner section as symmetrical. This latter reading is drawn on the bottom of figure 3.28 (b). The first five and last four sets are bracketed and the respective sets in each bracket sections are connected by slurs, indicating SATSIM2 similarity.¹³⁵ The transitive tuples formed by the SATSIM2 relations compare rather well with the inclusion graphs shown in figures 3.2 through 3.6. Almost all connections asserted in the inclusion networks are also asserted in figures 3.26 (SATSIM2 = 0.190) and 3.27 (SATSIM2 = 0.205). There are, however, two notable exceptions: **G** is *not* closely SATSIM2-related to either **I** or **C**. The lack of a close connection between **G** and **C** would appear to be critical

¹³⁵As we did in the ASIM analysis, we could single out the first five and the last five sets for comparison. Rather than treating the two realizations of **A** as a single set, we could compare **A1** to **C3** (SATSIM2 = 0.143), **A2** to **H**, etc.

given that **A**, **C**, **H**, and **G** are all related to each other through a single inclusion network. However, when one sorts out the weaker embeddings, as we have done in figure 3.6, **G** is no longer shown to be closely linked to **I** or **C**.¹³⁶

3.4 Resemblance determined by similarity indices which measure shared subset-class content

The results yielded by Rahn's ATMEMB, Lewin's REL, and Castrén's RECREL are, in many ways, strikingly similar.¹³⁷ This should not be surprising given that each of these indices compares the total abstract subset content of each pcset. The methods used for comparing the respective subset-class content of two pcsets is different with each index, but the data used is the effectively the same. We will examine the results using each of these separately. However, given the number of commonalities among the data provided by each of these three indices, we will posit a single analytic narrative (with some limited variations) using their collective data.

¹³⁶Having compared the inclusion-based networks to those networks formed by ASIM, IcVSIM, and SATSIM2, it is important to note that we are not attempting to prioritize the former as a sort of gold standard which similarity measures should attempt to model (particularly given that each of these three measures is based upon interval-class (or dyad-class) content and not total subset content). The inclusion graphs are used here as a rather simple and largely familiar way of acquainting ourselves with relations among the pcsets used in this piece. The comparisons between the inclusion networks and the networks formed by the total subset measures (ATMEMB, REL, RECREL, and AvgSATSIM n) are, however, much more useful as a sort of yardstick for those measures' success at modeling shared abstract subset content.

¹³⁷The TEST group used to perform the REL comparisons contained all set classes smaller than the smaller of the two sets being compared. This is precisely the method used by Castrén in his dissertation (Castrén, 89-96). Lewin defined REL to be flexible and allow any group of pcsets (TEST) to be used in performing the comparison.

The RECREL values were calculated precisely as Castrén defined the measure. Inversional equivalence was not applied, and therefore, the T_n forms of set classes 5-23 (**A**), 6-z39 (**B**—both realizations), 5-14 (**D**), 8-22 (**G**), and 7-23 (**H**) were used and the T_nI forms of 8-13 (**F**) and 7-11 (**I**) were used. 6-8 (**C**), 7-35 (**E**), and 7-8 (**J**) are all inversionally symmetrical and therefore do not have different T_n and T_nI classes.

3.4.1 Rahn's ATMEMB

Figure 3.30 provides an ATMEMB comparison matrix of the sets in this piece. As before, the information in this matrix is sorted into the list of most- to least-closely related set pairs provided in figure 3.31. Unlike the earlier sorted (ascending) lists of values, these values descend from largest to smallest. This is because larger numbers indicate greater similarity using ATMEMB.¹³⁸ Lewin's REL, which we will examine in the next section, also measures similarity on an ascending scale.

Figures 3.32 through 3.35 list and show sketches of the transitive tuples formed by using four different ATMEMB cutoff points for similar versus dissimilar set class pairs. As with the other indices we examined, the first figure begins with a very conservative cutoff and the figures progress to show transitive tuples formed by increasingly liberal similarity designations until all or most sets are included in at least one transitive tuple.

Using ATMEMB, **B** (6-z39, the second set, which is repeated once toward the end of the movement) and **J** (7-8, the final set of the movement) are by far the most similar set pair (ATMEMB = 0.876). **G** and **I**, the next most similar pair, are also separated by a relatively wide margin, forming a logical cutoff point (ATMEMB = 0.858). Because the first two salient cutoff points together relate only two pairs of sets, our first transitive tuple graph, in figure 3.32, will use ATMEMB values 0.819 as the first cutoff point. There are, as we can see, only #2 transitive tuples in this graph and we are unable to detect any clustering into families of ATMEMB-related SCs. Only one SC is completely

¹³⁸ATMEMB is therefore properly called a similarity index, while the others we have studied (including all the saturation-based indices) are more properly called "dissimilarity indices."

unconnected in this graph: **D**, the least connected SC using ASIM and IcVSIM.

The next graph, in figure 3.33, includes only four additional set pairs, but helps us establish better connections within our graph. While **D** is still not related to any other SC, a few three-set transitive tuples have been formed. Two of these, **{B, I, J}** and **{I, G, H}** are predictable from the inclusion graphs shown at the beginning of this chapter; the third, **{F, G, H}** is a bit surprising given that **H** (7-23) is not a subset of **F** (8-13).

Moving ahead to the graph in figure 3.34, we can see that **D** is *still* unrelated to any other sets and that more relations have been asserted among the sets shown in the more conservative two graphs. In particular, there are two heavily overlapping four-set transitive tuples, **{E, G, H, I}** and **{F, G, H, I}**, which nearly form a single five-element group. These sets encompass all the #7 and #8 SCs except for **J**. Curiously, no relation has yet been asserted between sets **A** (5-23) and **H** (7-23)—the only two complementary SCs found in our segmentation of the piece. This suggests, though by no means proves, that ATMEMB is not quite as cardinality-neutral as it would seem in Castrén’s value group matrix for ATMEMB.¹³⁹

Even in our last graph of transitive tuples (in figure 3.35), there is no connection between **A** and **H**; in fact, **A** is a member of only a single transitive tuple. **D**, as before, has no relations to any of the other sets in this piece.

While ATMEMB provides a much finer gradation of values than does ASIM, the graph in 3.35 is remarkably similar to the ASIM graph in figure 3.11. **{E, F, G, H, I}**, the five-set transitive tuple that was suggested, though not quite

¹³⁹Castrén, 85.

completed using $ATMEMB = 0.733$ is now a complete tuple with $ATMEMB = 0.655$. Additionally, there is one other five-set transitive tuple: $\{\mathbf{F}, \mathbf{G}, \mathbf{H}, \mathbf{I}, \mathbf{J}\}$. The degree of intersection between these two large tuples is 80% (four of five), suggesting an even larger six-element tuple containing *all* #7 and #8 sets in the piece—the same group shown in our final ASIM graph in figure 3.11. Figure 3.36 provides a summary of the SC relations using $ATMEMB = 0.655$ values.

3.4.2 Lewin's REL

Figure 3.37 provides a REL comparison matrix of the sets in this piece. As before, the information in this matrix is sorted into a list of most- to least-closely related set pairs (figure 3.38). Figures 3.39 through 3.42 list and show sketches of the transitive tuples formed by using four different REL cutoff points for similar versus dissimilar set class pairs.

The beginning of the sorted list of REL values (top of figure 3.38) strongly parallels the $ATMEMB$ example (figure 3.31). Both list $\{\mathbf{B}, \mathbf{J}\}$ and $\{\mathbf{G}, \mathbf{I}\}$ as the first- and second-most similar set pairs ($REL = 0.892$ and $REL = 0.870$, respectively). The first cutoff point that we will graph is $REL = 0.830$ (see figure 3.39). A quick comparison between this and the first $ATMEMB$ graph again reveals several parallels (they have been graphed with sets in the same locations to facilitate comparison). Both reveal six two-set transitive tuples, five of which are identical. ($\{\mathbf{B}, \mathbf{I}\}$ is not present in the $ATMEMB$ graph, though it appears at the next cutoff; $\{\mathbf{F}, \mathbf{G}\}$ is not present in the REL graph, though, again, it appears at the next cutoff.)

Our second cutoff point, shown in figure 3.40, is $REL = 0.781$. Again, we can compare this graph to the parallel $ATMEMB$ graph in figure 3.33 ($ATMEMB = 0.790$) and see plenty of commonalities. Both graphs contain the transitive tuples $\{A, C\}$, $\{F, G, H\}$, and $\{G, H, I\}$. Both also have each SC connected to another in the group, with the single exception of **D**.

The next salient dividing point in the list of REL values is at 0.752. We have not graphed it because only two new set pairs would be added to our graph in figure 3.40. These set pairs are worth mentioning, however, before moving on. **J** is connected to **I**, forming the three-set transitive tuple $\{B, I, J\}$, also seen in the $ATMEMB = 0.790$ graph. More interestingly, though, is the connection of **A** and **H**—our two complementary SCs. We recall that even in the most inclusive $ATMEMB$ graph, these two sets were not found to be similar.¹⁴⁰

Our next graph, found in figure 3.41, shows all relations where $REL = 0.731$. Comparing this to the $ATMEMB = 0.733$ graph (figure 3.34), we see that the two four-element transitive tuples are identical ($\{E, G, H, I\}$ and $\{F, G, H, I\}$) and both also include $\{B, I, J\}$ and $\{C, H, I\}$ as well as a still unconnected **D**. In fact, the only major difference between the two graphs is the important inclusion of $\{A, H\}$ by REL .

Our next logical cutoff point might well be 0.707. There are, however, only two new set pairs that are introduced with this cutoff. Our final cutoff point using this measure will therefore be $REL = 0.688$. While there are still no five-element transitive tuples formed under this standard for relatedness, there are, as before, many comparisons to be drawn between the transitive tuples formed

¹⁴⁰And that the lack of an $ATMEMB$ relation between these two sets made us consider the degree to which that measure was in fact cardinality neutral.

by REL and those formed by ATMEMB. **D** is still unconnected to any other set in the group.¹⁴¹ Unlike, ATMEMB, however, the #7 and #8 sets (**E**, **F**, **G**, **H**, **I**, and **J**) do not seem to form a group of interrelated sets. Figure 3.43 provides a summary of REL = 0.688 relations. A comparison between it and the summary of ATMEMB relations (in figure 3.36) shows that both list more relations to sets **H** (7-23) and **I** (7-11) than to any others (these might be seen as forming the nexus of this group using the REL and ATMEMB indices). However, the close relations are spread a bit more evenly among the ten set classes in the REL list than in the ATMEMB list.¹⁴²

3.4.3 Castrén's RECREL

Figure 3.44 provides a RECREL comparison matrix of the sets in this piece. The information in this matrix has been sorted into a list of most- to least-closely related set pairs in figure 3.45.¹⁴³ The top of this list is quite distinct from that of ATMEMB and REL. **E** (7-35) and **H** (7-23) form a single set pair that is by far the closest related (RECREL = 15). The next closest set pairs are related by RECREL = 21; a very large gulf considering that there is at least one set pair related by each RECREL value from 21 up to 37, forming an almost-linear progression from the second-most related set pairs to the least related set pair. This linear scale of values (caused, in part, by the relative lack of precision used in this measure—two rather than three significant digits¹⁴⁴)

¹⁴¹The first relation involving **D** doesn't appear until REL = 0.665 (**D:H**), and the next not until REL = 0.613 (**D:I**).

¹⁴²Again, this could be interpreted as suggesting that REL is more cardinality neutral than ATMEMB, an observation that is confirmed by Castrén's value group matrices (Castrén, 85 and 91).

¹⁴³As with ASIM, IcVSIM, and SATSIM2 (and unlike ATMEMB and REL) smaller values indicate greater similarity. RECREL is therefore also a "dissimilarity index."

¹⁴⁴The precision of the measure was not something that we felt free to alter given Castrén's desire that "the values [be] integers or other easily manageable numbers" and "the degree of discrimination is not too coarse or unrealistically fine." (Castrén, 18.)

makes it rather difficult to decide on dividing points between similar and dissimilar sets. In this case, we will progress by using similarity cutoffs that seem to include relatively equal numbers of additional set pairs and that facilitate comparisons to the ATMEMB and REL graphs.

Our first cutoff point is at RECREL = 21. Sets **A**, **C**, **E**, and **H** form a square of relations, each similar to exactly two others. While this might suggest the eventual formation of a four-set transitive tuple, the sets in this graph only form #2 tuples at this point (in the square, sets are not similar to those located diagonally across from them). This mere implication of larger transitive tuples occurs much earlier using RECREL than with the other two subset-based measures.

Moving up just one notch to RECREL values = 22 (shown in figure 3.47), we see that two adjacent (and overlapping) squares of relations have been formed and that all sets except **D** and **F** are now found to be similar to some other(s) in the group. There are still no transitive tuples with more than two elements, and the two complementary sets are still not similar, but the six sets that form the two squares match those found at the center of the inclusion graphs earlier in this chapter (figures 3.2 through 3.6). To further the comparison, **I**, in both the RECREL and inclusion graphs, is connected to **B**, which, in turn, is connected to **J**.

The next cutoff point is at RECREL = 24 (we are skipping RECREL = 23 because it adds only one set pair to the mix). This is shown in figure 3.48. **D** and **F** remain unconnected to any other SCs, but each of the other SCs now enjoys membership in at least one #3 transitive tuple and the collection of #3 tuples are interlocked (i.e., there are no transitive tuples that do not share at

least one set with another tuple). Graphically, there are no transitive tuples that are detached from the graph (this, interestingly enough, has been true of the previous RECREL graph as well). At this point, let us compare the transitive tuples formed by RECREL with those of ATMEMB and REL. Of the transitive tuples in the current RECREL graph, $\{\mathbf{B}, \mathbf{I}, \mathbf{J}\}$ and $\{\mathbf{G}, \mathbf{H}, \mathbf{I}\}$ are also present in the ATMEMB 0.790 graph in figure 3.33. $\{\mathbf{C}, \mathbf{H}, \mathbf{I}\}$ and $\{\mathbf{E}, \mathbf{G}, \mathbf{H}\}$ are added to this list if we use the ATMEMB 0.733 graph in figure 3.34 as our point of comparison. By contrast, each of the six three-element RECREL transitive tuples in figure 3.48 are also present as three-element tuples formed by REL 0.752 (these are described earlier between the descriptions of figures 3.40 and 3.41, and are shown as part of the more inclusive graph REL graph in figure 3.41).

Our next cutoff mark for RECREL similarity will be at the value 27. Set pairs related by this criterion are shown in figure 3.49. For the first time, the two “squares” of related sets, first shown in figure 3.46 above, are completely connected, forming two overlapping #4 transitive tuples, $\{\mathbf{A}, \mathbf{C}, \mathbf{E}, \mathbf{H}\}$ and $\{\mathbf{E}, \mathbf{G}, \mathbf{H}, \mathbf{I}\}$. Additionally, another #4 transitive tuple is formed at the right side of the graph, including sets $\mathbf{B}, \mathbf{C}, \mathbf{H}$, and \mathbf{I} . \mathbf{F} has now appeared on the graph for the first time, connected to set \mathbf{G} (at RECREL = 25), but \mathbf{D} is still found to be unrelated to the others, as it was in the REL and ATMEMB transitive tuples. Our final graph of RECREL values will be at the cutoff value 28. This is shown in figure 3.50. At this point, \mathbf{D} is still not connected to any other set in the piece. In fact, one has to look as far as the RECREL value 30 to find any connection to \mathbf{D} whatsoever. All other sets are members of at least one #3 or larger transitive tuple (and \mathbf{J} is the only set that is a member of exactly one).

Using RECREL, **D** and **J** are not only the least connected sets in the piece (this is more obvious from the summary of the relationships, found in figure 3.51), but they are also form one of the two least-related set pairs in the piece (RECREL = 41).¹⁴⁵

This RECREL graph also compares well with the embedding graphs at the beginning of this chapter. The only #5 transitive tuple in the RECREL graph (**{C, E, G, H, I}**) contains five of the six sets at the center of the inclusion graphs: only **A** is missing, but it is a member of another interlocking transitive tuple **{A, C, E, H}**. Interestingly, every transitive tuple that includes **A** also includes **H**, its complement. The reverse is not true, however. There are two transitive tuples (including the largest one) that include **H**, but not **A**. In fact, **H** is in a three-way tie for the second most connected set in the group while **A** is in a three-way tie for the least connected (not including **D**) set in the group. So, while the effects of cardinality are neutralized rather well by RECREL, the two complementary SCs are not, at least in this particular case, similarly related to other sets.

As we mentioned at the beginning of this section, the three subset-based measures, ATMEMB, REL, and RECREL suggest largely similar readings of the set relations in this piece. All three consider **D** to be dissimilar to every other set in this piece (**D** is also found to be dissimilar using ASIM and IcVSIM). The three indices also agree that sets **E** (7-35), **H** (7-23), **G** (8-22), and **I** (7-11) form a single transitive tuple, though they cannot agree whether **A** and/or **C**, their abundantly-saturated subsets, are likewise eligible for

¹⁴⁵The sets **A:J** are also related by RECREL = 41.

membership in the well-connected group.¹⁴⁶ All three indices also agree on the transitive tuple $\{\mathbf{B}, \mathbf{I}, \mathbf{J}\}$, a grouping suggested at the right side of the inclusion graphs.

As with SATSIM2 and ASIM, each of these indices suggests that \mathbf{I} is the most connected pcset in the group. In fact, using ATMEMB, REL, and RECREL (and also ASIM¹⁴⁷), \mathbf{I} is closely related to all of the other pcsets, except for \mathbf{A} and \mathbf{D} (using SATSIM2, \mathbf{I} is closely related to both \mathbf{A} and \mathbf{D} , but *not* closely related to \mathbf{E} and \mathbf{G}). It might therefore be seen as a sort of nexus of the group. This is certainly justifiable from the inclusion tables in figures 3.2 through 3.5 where \mathbf{I} is the sole link between the center group of pcsets ($\mathbf{A}, \mathbf{C}, \mathbf{H}$, and \mathbf{G}) and the rightmost group (\mathbf{B} and \mathbf{J}).¹⁴⁸

\mathbf{D} , as we already mentioned, is found to be the least connected (indeed the only unconnected) pcset in the group by each of these subset-based measures. We can once more find justification from figures 3.5 and 3.6 which show that \mathbf{D} is indeed only remotely connected to the inclusion network. On those inclusion graphs, it is connected to neither of the #6 pcsets and to only one of the three #7 pcsets. \mathbf{D} is embedded in both of the #8 pcsets (\mathbf{F} and \mathbf{G}), but that connection is fairly remote given that each of these measures prioritizes embeddings where the larger and the smaller set classes are closer in cardinality (this is particularly true of RECREL). We can therefore wonder why \mathbf{F} is not also detached from the similarity network. It is the only other

¹⁴⁶ \mathbf{C} , to be specific, is not embedded in \mathbf{E} , though it is in \mathbf{H}, \mathbf{I} , and \mathbf{G} . \mathbf{A} is embedded in all four sets and is the complement of \mathbf{H} .

¹⁴⁷In the case of ASIM, we suggested that cardinality was the primary reason \mathbf{I} was not related to \mathbf{A} and \mathbf{D} . Each of these other three indices (and particularly REL and RECREL), however, has a more cardinality neutral track record. This is shown effectively in the value group matrices of ASIM, ATMEMB, REL, and RECREL provided in Castrén 1994 on pp. 60, 85, 91, and 130, respectively.

¹⁴⁸One can question whether \mathbf{I} should be *so* closely related to pcsets \mathbf{E} and \mathbf{F} which are to the left of center and whether \mathbf{I} should be so distantly related to \mathbf{A} which, along with its complement \mathbf{H} , might instead be seen as the nexus (in Fortean terms) of the centermost group on the inclusion table.

pcset which is detached from the more selective inclusion table (based upon the proportional SATEMB values) in figure 3.6. **F** is an 8-note pcset which embeds none of the #6 or #7 pcsets in the group. Thus, it would seem to be even less connected than **D**. Despite this evidence, we can see that using these three measures **F** is connected (variously) to **A**, **G**, **H**, **I**, and **J**.

There are several possible reasons for this apparent discrepancy. First, it is possible that these measures are not as cardinality neutral as they might seem. There are, on average, fewer close relations to the smaller pcsets than to the larger pcsets in the group using each of these measures. However, the fact that **J**, a #7 pcset, is the second least connected pcset using REL and RECREL would seem to refute that claim. (With ATMEMB the average number of closely related pcsets increases almost linearly with the cardinality of each pcset.) Another possible reason for the discrepancy has to do with the fact that many more comparisons are performed in REL, RECREL, and ATMEMB when both pcsets are relatively large. When one compares a #5 pcset to a larger pcset, these indices can only examine the dyad-, trichord-, and tetrachord-class content of each pcset—even if the other set is, for example, an octachord. When one compares two octachords, however, these indices examine the total abstract subset content from dyad-classes up to septachord-classes. There are, thus, many more points of possible similarity (and, of course, difference).

Another important difference between an inclusion network (such as those in figures 3.2 through 3.5) and a subset-based similarity index is that the former is created by comparing only the sets listed with each other. Most subset-based similarity measures examine the total subset content. Therefore, two

large pcsets such as **F** and **G** could be dissimilar with respect to their embeddings of the pcsets used in this piece, but could be quite similar with respect to their total abstract subset content.¹⁴⁹ Finally, perhaps the most telling difference between inclusion networks and subset-based similarity indices, is the fact that the former can potentially be used to compare sets based upon their respective *subset and superset* content. The latter compare only mutual *subset* embedding and not the degree to which two SCs might be embedded in mutual supersets. We will follow up on this idea in the next chapter on future avenues for research.

Let us now return to the data proved by ATMEMB, REL and RECREL. An analysis based upon values from any of these three measures might well concentrate on the relative connectedness of each of the pcsets to their respective similarity graphs. We have already mentioned that **I**, the penultimate segment in the piece is the most connected pcset in each graph. **D**, on the other hand, might be viewed as the odd segment out. It is heard only once in the piece, in mm. 10-11, and the contour used in its pitch realization is markedly different from the material that precedes or follows it. **F** is also distinctive in the piece, and, as we have already shown, equally detached in a superset/subset-complex-based reading. Our AvgSATSIM n measure (a saturation-based index which compares “total” abstract subset content) is the only index of those employed that suggests **F** (not **D** or **J**) is the most detached pcset in the supernetwork of relations. In fact, as we shall see in the next section, AvgSATSIM n yields no close relations to **F** within the group.

¹⁴⁹Lewin’s REL is the only similarity measure that allows for a flexible TEST group of subset classes. Had we entered the ten set classes used in this piece as our TEST group, the results from REL would largely be isographic to those in our inclusion tables.

Figure 3.52 provides a diagram of the fifteen pcsets in the order that they appear in the piece. Rather than graphing the related pcsets using each the three subset-based measures separately, we will show only one graph that includes information from all three indices. This is easily done because the high degree of common relationships found using *ATMEMB*, *REL*, and *RECREL*. As before, lines between adjacent pcsets indicate a close resemblance. Thicker lines will indicate that all three measures agree on that particular relationship; thinner lines indicate that only one or two of the indices agree (in such cases the index name(s) will appear beneath the dotted line).

The only pcset that has no adjacent relations under these three measures is **D**, which also has no close relations whatsoever.¹⁵⁰ The greatest number of connections is found using either *RECREL* or *ATMEMB*; *REL* therefore yields the smallest number of connections. We cannot produce a symmetrical reading of the piece (similar to the *SATSIM2* reading) using any of these measures. With *ATMEMB*, we also cannot relate the beginning to the anacrusis of m. 23 (complementary pcsets **A** and **H**) as we did with *IcVSIM* and *SATSIM2*. The only very clear dividing place is either before or after **D** (mm. 10-11). If we place a sectional divide after **D**, a new section begins on **C2**. This seems perfectly sensible, particular given the relative lack of motivic connections between **C1**, which immediately precedes **D**, and **C2**. Because no other sectional divisions are strongly implied, our reading might well feature an introductory passage (mm. 1-11) that leads up to one very distinct set (**D**);

¹⁵⁰Also, one realization of **A** is also only trivially related to the other realization of **A**.

followed by a longer passage, occupying the latter two-thirds of the piece, where there is a (relatively) smooth flow from one set type to another. Still another reading of these data might claim that the three realizations of set class 6-8 (**C**) signal the boundaries from one section to another. **C1** comes at the end of the first section, just before the anomalous SC **D**. **C2** initiates the third section (in this case **D** serves as a single-event section), a less unified section than the first. That section ends with **G** (in m. 21) and the final section begins at **C3** (m. 21), which is connected to **H**. **H**, then, is connected by two of the three measures to **B2**, which, in turn, is connected to both of the last two sets. The last section, then, is one unbroken chain of similar sets from **C3** to **J**.¹⁵¹ As we first mentioned when discussing this sectional boundary in the ASIM-based analysis, **C3** is an apt set to initiate the final section because of its close pitch class and set class parallels to **A1**.¹⁵² Furthermore, the last note of the first section is D^\sharp (this is true regardless of whether the first section ends with **C1** or **D**), enharmonic to the last note of the piece, $E\flat$.

3.4.4 AvgSATSIM n

As we already mentioned, AvgSATSIM n is the single measure that finds no close relations to SC **F**. En route to discussing the isolation of **F**, we will follow the same protocol used in discussing the other measures of resemblance. An comparison matrix of AvgSATSIM n values is provided in figure 3.53, followed in figure 3.54 by a sorted list of the values in that matrix. Unlike the other subset-based indices, but like SATSIM2, **A** and **H**, the two

¹⁵¹The fact that ATMEMB and REL found **H** close to **B**, but did not find **A** close to **B** further illustrates the different relations enjoyed by complementary sets **A** and **H**.

¹⁵²Recall that both **A1** and **C3** begin with the same pitch classes, and **C3** is a literal superset of **A1** ($C3 = \{A1 + A^\sharp\}$).

complementary SCs are (by far) the *most* similar pair.¹⁵³ This close similarity arises because saturation of a set's elements (either intervals or subset classes) reflects the structural similarities (and, in the case of interval-class saturation, equivalency) of complementary sets.

The first cutoff point that we will use is $\text{AvgSATSIM}_n = 0.150$, shown in figure 3.55.¹⁵⁴ There is one three-member transitive tuple in this initial graph: $\{\mathbf{A}, \mathbf{E}, \mathbf{H}\}$. Like the other subset-based indices, the present relations compare well with the networks of embeddings shown earlier in this chapter. The six sets that participate in one of these very close AvgSATSIM_n relations are the same ones that comprise the center of our inclusion graphs in figures 3.2 through 3.6. No sets from the far right or left of the inclusion graph are yet related by AvgSATSIM_n .

The next dividing points are at $\text{AvgSATSIM}_n = 0.163$ and $\text{AvgSATSIM}_n = 0.183$. Transitive tuples formed under these standards for relatedness are shown in figure 3.56, with thin lines connecting sets related by the more conservative cutoff point and the thicker line between \mathbf{H} and \mathbf{I} signaling the more liberal cutoff. Including the relation of \mathbf{H} and \mathbf{I} , we now have four #3 transitive tuples, each of which has two sets overlapped with another (this is shown by the “quilt” of right triangles). \mathbf{D} and \mathbf{F} are not yet connected to any other sets in this graph and only \mathbf{J} and \mathbf{G} are members of #2 transitive tuples, but not members of any of the #3 tuples.

¹⁵³This set pair is the 11th most similar (of 45 comparisons) using RECREL, the 15th most similar using REL, and the 24th most similar (or 21st least similar) using ATMEMB.

¹⁵⁴This is admittedly not exactly a clear breaking point, but we wanted our first group of relations to be roughly the same size as those in the first graphs of the other indices for easy comparison. Also, the next most salient break after the initial set pair comes between the AvgSATSIM_n values 0.157 and 0.163, a span which is immediately surpassed given that the next most similar value is 0.183.

The final cutoff point, shown in figure 3.57, is at $\text{AvgSATSIM}_n = 0.223$. As we indicated earlier, only **F** (8-13) is still unrelated to any other SC in this graph. All the other SCs enjoy relations in transitive tuples of at least three elements. The graph in figure 3.57 outwardly expands the quilt shown in figure 3.56. There are now six right triangles; the top two triangles (#3 tuples) from the previous graph are now connected to form the single #4 transitive tuple. For the first time since the SATSIM2 transitive tuples, we now see **D** related (though not strongly) to the group. It also bears mention that while complementary SCs **A** and **H** were the first sets that AvgSATSIM_n deemed similar, they do not act in tandem in all the AvgSATSIM_n transitive tuples. In fact, **H** enjoys more connections than does **A** in this graph (as it did with the other subset-based indices). This is seen easily in figure 3.58, which contains a summary of the AvgSATSIM_n relations using our most recent cutoff point. While **F** is the most remote pcset in this group, there is no single clear AvgSATSIM_n nexus set for the piece. **C**, has the greatest number of connections (at 6), but **I**, and **H** have almost as many (at 5 each). **I**, of course, is the pcset that was best connected under each of the other subset-based measures. **H**, however, is the abstract complement of the first pcset in the piece (**A**) and is well connected to both the right and left sides of the inclusion networks while, quite obviously, maintaining the extremely strong connections to pcsets **A** and **C** in the center of the graph. **C** strongly embeds both **H** and **I** and, in addition to being centrally located on our inclusion graphs, is the most frequently heard pcset in the piece, realized three times in distinct segments.

A wider comparison between the inclusion networks and the AvgSATSIM n graphs reveals a high degree of similarity between the graphs. This suggests that AvgSATSIM n is a fairly accurate reflection of the inclusion network model. For example, **A** is closely AvgSATSIM n -related to **C**, **D**, **E**, and **H**. While there can be no direct subset relation between **A** and **D** (since they are both pentachords, the smallest cardinality pcsets on that particular inclusion graph), the two pcsets are abstract Rp relations (sharing a common 4-note SC, 4-23[0257]). **C**, **E**, and **H** are all supersets of **A**.¹⁵⁵ **G** is closely AvgSATSIM n -related to **E**, **H**, and **I**, the three septachords in the piece which it embeds. Many more comparisons between AvgSATSIM n and the inclusion networks may be asserted. We will, however, leave them to the reader to draw.

Figure 3.59 provides a diagram of the fifteen pcsets in the order that they appear in the piece. As before, lines have been drawn between adjacent sets that are found to be similar (in this case by the cutoff AvgSATSIM n = 0.223). Unlike the other subset-based graphs, **D** is not only found to be similar to *some* other sets in the group, but to **C**, the set class between which it is sandwiched. **D**, it would seem, is therefore no longer a sensible dividing point. It is now **F** and **G** that are isolated from any connections to adjacent set classes (and **F**, we recall, is found to be dissimilar to all other SCs in the piece). In the music, **F** flows seamlessly into **G**, making that an unlikely formal dividing point. However, as we mentioned earlier in our discussion of the

¹⁵⁵There are two other supersets which do not make the cut-off point to be considered a close AvgSATSIM n relation with **A**. These are **F**, an eight-note superset which contains only one of a possible six (in any #8 set) **A**-type subset; and **G**, which contains five of six possible **A**-type subsets. AvgSATSIM n (**A**, **F**) = 0.355, the highest value in AvgSATSIM n comparison matrix; AvgSATSIM n (**A**, **G**) = 0.251, just a 0.033 away from the last cutoff value for close AvgSATSIM n relations (all cutoffs are, to a certain degree, arbitrary).

other subset-based measures, **C3**, which follows and is dissimilar to **G**, provides a rather elegant dividing point by virtue of its pitch and rhythmic similarities with **A1**. The first five segments (**A1** — **A2** — **B1** — **C1** — **D**) form a clear parallel to the last five segments (**C3** — **H** — **B2** — **I** — **J**), particularly given the similarities between **A**, **C**, and **H** and the fact that the third set of each is a form of **B**.

3.5 Resemblance determined by the cyclic saturation similarity index (CSATSIM)

Before attempting to draw some general conclusions from our examinations of the various interval-class- and subset-class-based indices, we will explore the results of one final saturation-based index. CSATSIM, which is described in detail in chapter 2, bases its comparisons on the similarity of *i*-cyclic subset embeddings in two pcsets. Thus, while it processes more than just the interval (or #2 subset) content, only those subset-classes which can be formed from repeated transposition by a single interval *i* (this is a simple definition of an *i*-cyclic pcset) are used in the comparison. Figure 3.60 gives a comparison matrix of these sets using CSATSIM. The values in this matrix are sorted and listed linearly in figure 3.61.

As with all the other saturation-based indices (and only IcVSIM among the non-saturation-based indices¹⁵⁶) the two complementary sets, **A** and **H**, yield the closest relation: in this case, the very small CSATSIM value 0.086. Figure 3.62 provides a list and graph of the transitive tuples formed by CSATSIM

¹⁵⁶We recall from our earlier discussion that any two complementary sets of #5 and #7 will be found to be as similar using IcVSIM because the difference between the ICVs of any two #5 and #7 complements will be exactly the same. The greater the difference in size among complementary sets, the less related they will appear using IcVSIM.

values 0.117. As one can see, there are only three relations at this cutoff and no real suggestion of larger group formations. We progress without further comment, therefore, to figure 3.63, which shows our second cutoff point: CSATSIM 0.144.

The three sets that formed the top segment in figure 3.62 have now grouped together into a single #3 transitive tuple, one of two in figure 3.63. The other #3 tuple connects sets **B**, **F**, and **I**. While **B** and **I** are in a subset/superset relationship and are related using all eight measures that we have employed, **B** and **F** are *not* in a subset/superset relationship and are related—indeed closely related—by only the CSATSIM and IcVSIM indices (and, logically, none of the subset-based indices).

Two distinct families of CSATSIM set types begin to emerge in figure 3.64: the five sets on the right of the graph form one of the families (which are admittedly weakly drawn at this point), with six similar set pairs among them; the five sets on the left of the graph form the other family, featuring seven similar set pairs among them. There is only a single connection between the two halves of the graph (**H** and **C**). All the sets have now been included (in the last graph **D** enjoyed no connections) and all but **J** are members of at least one #3 transitive tuple.

We could stop at this point, having connected each of the ten distinct set classes to our graph, but we are still operating under a comparatively conservative cutoff for similar versus dissimilar sets.¹⁵⁷ Moving the cutoff point up to CSATSIM 0.194 further clarifies these two distinct families of relations. The relations that fall within this cutoff are graphed in figure 3.65.

¹⁵⁷This is conservative compared to our final cutoff point for each of the other seven similarity indices. Such correspondences are revealed by comparing the final cutoff value number to the average value returned by each respective measure for the group of sets at hand.

On the right side of the graph, there are now eight connections (of ten possible) among the five sets; on the left side, there are nine connections among the five sets. Each half is almost isographic, very nearly forming a pentagram (five pointed star) within a pentagon (the graphic image for a #5 transitive tuple). While the connections among the families are robust, there are only three connections between the two families (of 25 possible). The relations among all the sets can be summarized by only six transitive tuples. A sorted list of these CSATSIM relations is provided in figure 3.66.

The two distinct families of set types found in our interpretation of the CSATSIM values begs the question: what are the distinguishing features of these two families? By what criteria, in other words, has CSATSIM found the five sets on the right of figure 3.65 to resemble each other (and likewise the five sets on the left)? Why, too, are there so few resemblances between the two families? Figure 3.67 contains cyclic saturation vectors (the data used in calculating CSATSIM) for the ten set classes under consideration. Looking only at the top line (boldfaced) of each saturation vector, we see that the two families of sets both have very little 4-cycle content, and a moderate amount of 2-cycle content. The sets in the top family have, on average, very few tritones, while the sets in the bottom family generally have a low-to-moderate number of tritones. The top family's sets have very little 1- and 3-cycle content, while the bottom family's sets have a higher 1- and 3-cyclic content. Finally, and most strikingly, each of the top family's sets has maximal or near-maximal embeddings of 5-cycle fragments; the bottom family's sets each have relatively low-to-moderate 5-cycle content.

One might think that a sharp difference in only one cyclic dimension (and a moderate difference in two other cyclic dimensions) would not have such a great impact in the creation of set-type families. However, as we showed in chapter 1, large amounts of any particular interval class (or maximal embeddings of a particular ic-cycle) effects the amount of all or many of the other interval classes (see figure 2.5—a matrix that shows the maximal and minimal ic content generated in all cardinalities by each cyclic set type).

Therefore, maximal to near-maximal difference in the relative amount of only a single ic (or ic cycle) is easily enough to determine a high degree of difference (dissimilarity).

The two set classes that enjoy the fewest close CSATSIM relations are, not surprisingly, two of the most distinct in the group: **E** (7-35) and **J** (7-8). **E**, the diatonic set, is comprised of a single unbroken 5-cycle fragment; **J**, on the other hand, is the only set in the group that embeds the *minimal* number of 5-cycle fragments for its cardinality.¹⁵⁸ Not surprisingly these two sets are not only the most distinct from the group, they are also the most distinct from each other, yielding the CSATSIM value 0.414. One could assert that these two sets—one a 5-cycle set, the other an “anti-5-cycle” set—are not only the yin and yang of the piece, but that they form a sort of nexus for their respective families of pcsets (as shown on the last CSATSIM graph in figure 3.65).

On the other hand, given that **E** is *not* related to each of the other four sets in its family and **J** is similarly *not* related to each of the other four sets in its family, we might look toward sets that *are* similar to all members of their family and dissimilar to all members of the other family. There are, in fact, sets that

¹⁵⁸7-8 [0234568] contains two instances of ic5 ([05] and [38]) that are not from adjacent places along a complete 5-cycle. This is the fewest number of ic5s possible in a 7-note set using the most cyclically fragmented possible realization.

meet those criteria: Both **A** and **G**, on the left side of the graph in figure 3.65, are related to each of the sets in the left-side family and to none of the set in the right-side family; **B**, on the other side, is related to each of the sets in its family, and to none in the left-side family. Since **A** and **B** are also the first two distinct sets in the piece, claiming that they form the primary agents for contrast in the piece seems rather elegant. Either pair of sets (**A/B** or **E/J**) could be singled out as the emblems for their respective families.¹⁵⁹ Which one chooses is matter of analytic interpretation.

In many ways, CSATSIM produces results which have affinities to both the saturation-based interval-class comparison (SATSIM2) and the saturation-based total subset-class comparison (AvgSATSIM n). This should not be too surprising given the analogous structures of the three indices. Figure 3.68 provides a side-by-side comparison of the most inclusive graphs that we used for each of these measures. The two families suggested by CSATSIM have been “boxed” in each graph for easy comparison.

Family I sets **A**, **H**, **E**, and **G** form a transitive tuple in both the SATSIM2 and CSATSIM. In the AvgSATSIM n graph, those pcsets all appear in close proximity, but the relation between **A** and **G** is (only barely) not strong enough to warrant a connecting line. In all three graphs, **{B, C, I}** and **{B, I, J}** are overlapping networks of closely related pcsets. Of all the graphs/lists of transitive tuples we have examined, CSATSIM presents the best case for two distinct (i.e., non-overlapping) families of pcsets. The two CSATSIM families are not only well-defined, but, as we mentioned, nearly isographic, featuring one four-pcset transitive tuple maximally overlapping with one three-pcset

¹⁵⁹**G/B** could also function in this way.

tuple on the right of the graph and two maximally-overlapping four-pcset transitive tuples on the left half.¹⁶⁰ As we mentioned earlier, there are only three connections between the two families in the CSATSIM graph. By contrast, there are six connections between the two families in both the SATSIM2 and AvgSATSIM n graph. This is almost as many (and in one case more) connections between families as among either family in these two graphs.

Using either SATSIM2 or CSATSIM, **H** (or **A/H** in the case of SATSIM2) and **I** enjoy the largest number of resemblances to other sets in the piece. (Using AvgSATSIM n , **H** and **I** are the second most strongly connected pcsets.¹⁶¹) Almost every transitive tuple in each graph contains either **H**, **I**, or both (AvgSATSIM n contains two exceptions: one of which involves pcset **F**, which is closely related to no other pcset). From this observation, we might refine one of the analytic narratives formulated earlier in our discussion of the CSATSIM values. While the piece might still be understood as beginning with two dissimilar pcsets, **A** and **B**, that set the stage for the contrasts heard in the rest of the piece,¹⁶² **H** and **I** (and to a lesser degree **C**) could be seen as the synthesis of their material into larger structures. **H**, in the words of Howard Hanson, is a “tonal expansion” (i.e., complement) of **A**. **I**, on the other hand, might be seen as the bridge between the contrasting **A/H** and **B**-type pcsets which directly precede it. **I**, as the penultimate segment in the

¹⁶⁰By maximal overlap, we refer to the maximal number of shared pcsets that two networks can contain without the smaller network being subsumed by the larger network. To calculate the maximal overlap, simply subtract 1 from the number of sets in the smaller network. For example, when a three-pcset network is the smaller group, the maximal overlap it can have with a larger (or same sized) network is 2.

¹⁶¹They are second to set **C** (6-8).

¹⁶²In fact, none of the eight similarity indices produce a close relation between these two pcsets.

piece, also serves as a bridge to the final pcset, **J**, which is by some accounts the most anomalous pcset in the piece.¹⁶³

While **H** and **I** are members of the two contrasting families within the CSATSIM-based graph of pcset relations, they resemble each other according not only to CSATSIM, but to all eight of the indices discussed herein. This connection is made vivid by their very similar musical realization. The contour of **H**, expressed as a cseg (contour segment), is <0,6,4,5,3,2,4,1>; the contour of **I** is <0,4,2,5,3,1,6,5>, where 0 represents the lowest note of each segment and 6 represents the highest note of each segment.¹⁶⁴ In contour space (c-space), there are only three possible intervals (cints): ‘0,’ where the two contour pitches (cps) are the same; ‘+,’ where the second is registrally higher than the first; and ‘-,’ where the second is registrally lower than the first. In csegs **H** and **I**, the pattern of adjacent cints is the same: <+, -, +, -, -, +, ->. Not surprisingly, these two csegs are judged to be quite similar using Marvin and Laprade’s contour similarity (CSIM) algorithm,¹⁶⁵ which compares all the cints (adjacent and non-adjacent) in each cseg. These cints are displayed in the top right triangle of the comparison matrices in figure 3.69. This figure also shows the derivation of the CSIM value of **H** and **I**. These are the only two segments with the same pattern of adjacent cints, and are the most closely related pair of contours in the piece (using the Marvin/Laprade CSIM algorithm).

¹⁶³If one allows dynamics into the mix, then **J** would almost certainly be considered the most anomalous segment, as it is the only one marked *forte* (actually, “poco più *f* e poco più mosso”), whereas the rest of the piece is marked “sempre *p* e molto tranquillo.” While this would factor into an overall reading of the piece, it does not play a role in our pcset-based analytical commentary.

¹⁶⁴The term “cseg” was coined by Marvin and Laprade 1987, 228. It is an “ordered set of contour pitches (cps) in contour space (c-space), a “type of musical space consisting of elements arranged from low to high disregarding the exact intervals between elements.” (Marvin and Laprade, 255) The notions of c-space and cps were developed in Morris 1987, 26-34. Morris’s “contour” is the same as Marvin/Laprade’s cseg.

¹⁶⁵Marvin and Laprade, 237.

Returning now to the pitch class content and the families formed by CSATSIM, we will conclude our analysis by discussing the temporal arrangement of the pcsets *vis-à-vis* the CSATSIM 0.194 values. Figure 3.70 shows a temporal arrangement of the fifteen pcsets. Beneath the letter-name designations is a Roman numeral: “I” indicates that it appeared in the left-side family in figure 3.68; “II” indicates that it appeared in the right-side family. Lines drawn between two adjacent sets indicate a close CSATSIM relationship.

There are only four connections between adjacent SCs in this piece, discounting the repeated **A**s at the beginning and **E**s in the middle. We will therefore progress to the task of relating non-adjacent sets in a manner which seems musically relevant. As we did with SATSIM2 (see figure 3.28), we can read a symmetrical design, the center point of which is SC **F**. The top graph in Figure 3.71 shows two readings that are nearly identical to those in figure 3.28. The slurs above the graph suggest the purely symmetrical reading from **F** outward to **B1** near the beginning and **I** near the end. The slurs beneath the top graph suggest connecting the first SCs (**A1** and **A2**) to their complement, **H**, toward the end; **B** at the beginning to **B** at the end, and **C** at the beginning to **I** at the end. In the first scenario, the first and last SCs are left out of the mix; in the second scenario, the last SC is left out. The bottom graph on figure 3.71 presents a slight variation on the second reading. In this case, **A1** is connected not to its complement, but to **C3** (which is similar to **A1** in terms of ic-cycle content, pc content, contour, and rhythm). **A2** is then connected to **H**, its complement, and the rest is the same as in the first graph.

A different reading is provided in figure 3.72. This one follows the idea that **A** and **B**—the first two SCs—form the essential agents of contrast throughout the piece and that every subsequent SC resembles one *and only one* of the initial two. The graph shows all slurs beginning on either **A** (**A1** and **A2** are treated together here—they are not only member of the same SC, but they also utilize exactly the same pitch classes) or **B**. Every set is connected on this graph and the segments in the piece are roughly shown to alternate between **A**-type sets (maximal or near-maximal 5-cycle sets) and **B**-type sets (minimal or near-minimal 5-cycle sets).

3.6 Conclusions

Let us now compare the close relations found using all eight similarity indices. Figure 3.73 provides a grand comparison matrix that shows which similarity indices closely relate which SC pairs in the piece. Of the 45 SC pairs shown on the matrix,¹⁶⁶ there are nine pairs that all eight indices agree are closely related (using the most liberal cutoff points considered herein). These are **{B, C}**, **{B, I}**, **{B, J}**, **{C, H}**, **{C, I}**, **{E, G}**, **{E, H}**, **{G, H}**, and **{H, I}**. Three #3 transitive tuples are formed by all eight indices: **{B, C, I}**, **{C, H, I}** and **{E, G, H}**. Among these transitive tuples formed by every index, each has at least some degree of overlap with another. This is depicted in figure 3.74—a graph of these universal transitive tuples. There are seven pairs of SCs that none of the indices found to be closely related. These are **{A, B}**, **{A, J}**, **{B, D}**, **{B, E}**, **{C, J}**, **{D, E}**, and **{D, J}**.

¹⁶⁶In a ten-element group, there are $9+8+7+6+5+4+3+2+1 = 45$ distinct pairs of elements.

While one can certainly argue that the points of mass agreement indicate a greater degree of certainty that two pcsets are either similar or dissimilar,¹⁶⁷ the converse is less logically sound. In other words, one cannot really say that an index is faulty just because it was the only one (or only one of two) that found a particular pair of pcsets to be closely related. What seems like an anomalous relation, spotted by only one index could, in fact, be the result of a problematic feature of all the other indices. Rather than view these in terms of good and bad, however, it is more appropriate to imagine each of these indices as looking at a pair of pcsets through a particular lens. Some lenses magnify certain features while diminishing the effect of other features. Each index produces values that represent its particular frame of reference. Thus, while each of the indices appears to measure the same thing—the precise degree that two pcsets are related—each is also (and perhaps more importantly) a reflection of its own construction.

¹⁶⁷This is, in part, the approach taken by Ian Quinn in his recent paper, “Evolution A Forte-iori: On Similarity, Relations, Similarity Relations, and the Taxonomy of the Harmonic Menagerie,” presented at the Music Theory Society of New York State 1997 annual meeting, Rochester, NY.

Chapter 4

Similarity and dissimilarity of set types in Olivier Messiaen's *Petites esquisses d'oiseaux*, second movement ("Le Merle noir")

Messiaen's 1985 composition *Petites esquisses d'oiseaux* is a suite of six short pieces. Unlike his larger solo piano pieces devoted to bird music—the massive thirteen-movement collection *Catalogue d'oiseaux* (late 1950s) and the lengthy single-movement *La fauvette des jardins* (1970), Messiaen set the songs of only a single bird in each piece. The odd numbered pieces are all settings of the robin ("Le Rouge-gorge"); the second piece, which we shall examine, is a setting of the blackbird ("Le Merle noir"); the fourth piece is a setting of the song thrush ("La Grive musicienne"); and the final piece is a setting of the sky lark ("L'Alouette des champs"). The score to *Petites esquisses d'oiseaux* also contains a limited amount descriptive commentary detailing the birds' surroundings, particularly when compared to the very specific markings in both the *Catalogue d'oiseaux* and *La fauvette des jardins*. Musically, however, all three of Messiaen's bird-related works for solo piano are stylistically similar and, as we will briefly discuss later in this chapter, each employs similar harmonic materials.¹⁶⁸

4.1 Form

An annotated score to "Le Merle noir"¹⁶⁹ is provided as example 4.1. There are four large sectional divisions in the piece: each is divided into three

¹⁶⁸The harmonic elements in Messiaen's late piano concerto, *Un Vitrail et des Oiseaux*, written in 1986, also share a great many commonalities with this piece.

¹⁶⁹We will continue to refer to the second movement of *Petites esquisses d'oiseaux* simply as "Le Merle noir." It is worth noting, however, that Messiaen has also composed a separate piece for

distinct smaller sections. The three smaller sections recur in the same order, but slightly altered, four times to form the larger sections. We will use Roman numerals to label the larger sections and will append A, B, and C to label the three small sections. The three smaller sections are delineated not only by their thematic materials, but also by unique tempi and dynamics. The A sections each have the metronome marking $\text{♩} = 80$ and a constant *forte* dynamic marking, except for the final chord of each A section which is marked *piano*; the B sections are all marked $\text{♩} = 112$ and played *fortissimo*; and the C sections are played *pianissimo* at $\text{♩} = 88$. The form, then, is unambiguous and can be represented as follows (because this piece is unmetered measure numbers do not necessarily indicate relative section length¹⁷⁰):

		mm.			mm.			mm.			mm.
I	A	1-4	II	A	10-13	III	A	20-23	IV	A	31-36
	B	5-6		B	14-16		B	24-25		B	37
	C	7-9		C	17-19		C	26-30		C	38-40

Messiaen's depictions of the blackbird's song occur only in the B sections. The A and C sections are entirely homophonic. The chords that comprise these two sections are commonly called "color chords," used by Messiaen to paint a background picture of the birds' habitat, particularly the wind or waves.¹⁷¹ Each of the four A sections employ between 11 and 17 chords; three of the four C sections are only two chords long. The first chord in each C section is a four-note tone cluster (4-1 [0123]) in the lowest register of the piano, and the remainder of the chords are registrally-higher open-

flute and piano called *Le Merle noir* (1951). Messiaen also set the blackbird's song in other works, most notably in his *Catalogue d'Oiseaux* (1956-58).

¹⁷⁰That is not to say that the piece is *unmeasured*. The very specific metronome markings allow one to determine relative section length quite accurately.

¹⁷¹Sherlaw Johnson, 116.

spaced chords. Both the low clusters and the higher chords which follow are played under a single sustain pedal. The third C section (IIIC) is the longest of the four, consisting of five open-position chords and two low pedal clusters. In general, each of the three smaller sections (A, B, and C) grows longer with each return.

In sections 4.3 through 4.5 of this chapter, we will examine the four episodes of the three smaller sections. Rather than progressing in order, dealing with the A sections first, then the B and C sections, our analytic exposition will move from the A sections to the C sections, then finally (and rather briefly) to the B sections. There are multiple reasons for this atemporal ordering: primary among them are the observations that the A and C sections are most closely allied (both are homophonic), and both are quite distinct from the fast flourishes (representing the bird songs) that characterize the B sections. Further, the B sections are the most complex to segment. We will also raise some concerns about the segmentation of the C sections, but these issues are not nearly as difficult as those introduced in the B section. Therefore, the order of analysis is a progression from the easiest to analyze using a similarity index to the most difficult.

4.2 Choice of similarity measures for use in analysis

In this piece, and in fact throughout his *oeuvres*, Messiaen frequently employed cyclic or near-cyclic sets.¹⁷² Such sets sometimes arise as the products of his well-known use of symmetrical collections, particularly his modes of limited transposition. While the modes are used only infrequently in

¹⁷²Headlam 1996 discusses cyclic and near-cyclic sets (his “cycle+” collections) at length in the context of Alban Berg’s atonal music (c.f., pp. 67-77).

Messiaen’s bird music, chords derived from these symmetrical collections rather common elements of his harmonic vocabulary. Primarily for this reason, we elect to use the cyclic saturation similarity measure (CSATSIM) to compare the events in this short piece.¹⁷³ As in chapter 2, we will use the CSATV weighting value 1.2 in our comparisons.

4.3 The A (♩ = 80) sections.

Each of the four A sections of “Le Merle noir” is primarily comprised of chords that are members of SCs 7-20 [0125679] and 6-15 [012458], the first occurring roughly four times as frequently as the second.¹⁷⁴ As these are the first two distinct SCs in the piece, we will label occurrences of 7-20 and 6-15 chord types **A** and **B**, respectively. On the annotated score in example 4.1, each chord is labeled with both its set class and a letter name for ease of comparison. The letters are consecutive within like sections (e.g., the final chord in IA is labeled **C**; the final chord in IIA is labeled **D**), not consecutive between temporally-adjacent sections (thus, the first SC of section IB—which immediately follows IA—is *not* labeled **D**). Each of the chords in the A sections is an **A**- or **B**-type chord, except for the final chords (or, in the case of the last A section (IVA), the final three chords), which are members of SCs that are heard only once in the piece. We will call the final chord of each A section the “cadential chord” because of the harmonic and rhythmic closure that each conveys.

¹⁷³I do not mean to infer that CSATSIM is *only* useful where cyclic sets are used, but rather that it is particularly useful in such cases.

¹⁷⁴There are 38 chords that are members of SC 7-20 (**A**) and 10 chords that are members of SC 6-15 (**B**) in “Le Merle noir.”

The 54 chords in the four A sections are thus members of only eight distinct set classes. As we mentioned, the majority of these chords are labeled **A** or **B**; the other chord types occur only once throughout the piece. SCs **C**, **D**, **E**, and **H** are the cadential chords to sections IA, IIA, IIIA, and IVA respectively, and **F** and **G** occur as the antepenultimate and penultimate chords of IVA. Figure 4.1 shows a comparison matrix with all CSATSIM values among this group of SCs. The first four chords are the non-cadential chords (**A**, **B**, **F**, and **G**) and the last four are the cadential chords (**C**, **D**, **E**, and **H**). The box in the upper left of the matrix includes all comparisons among the non-cadential chords; the box in the lower right includes all comparisons among the cadential chords. Figure 4.2 presents the non-zero values from the comparison matrix, sorted from lowest (most similar) to highest (least similar).

As we discussed in chapter 3, there are various ways to interpret these numbers. For example, to find out whether $\text{CSATSIM}(7-20[\mathbf{A}], 6-15[\mathbf{B}]) = 0.163$ is a relatively close relation, we need to establish a context for CSATSIM values. We could begin by examining the arithmetic mean (average) of all CSATSIM relations of the value group #7:#6, as well as the average CSATSIM values for 7-20 compared to all hexachord classes and the average for 6-15 compared to all septachord classes (the value groups 7-20:#6 and 6-15:#7). The #7:#6 CSATSIM value group average is 0.232 and the values range from 0.011 to 0.617. Clearly, then, 0.163 is more similar than average. More specifically, the average CSATSIM value in the value group 7-20:#6 is 0.210 and 24.5% of the remaining hexachord classes (12 of 49) enjoy closer CSATSIM relations to 7-20 than does 6-15; the average CSATSIM value in the 6-15:#7 value group is 0.195 and 40.5% of the

remaining septachord classes (12 of 37) enjoy closer CSATSIM relations to 6-15 than does 7-20.

To get a broader view of what CSATSIM values are both possible and common for each SC, we can also examine the statistical summaries provided in appendix G. This appendix lists the prime form and Forte number for each SC as well as the average CSATSIM value found in comparing it to all SCs of cardinality 2 through 10 (the value group #2 .. #10:#2 .. #10), and the lowest and highest possible CSATSIM values for all possible CSATSIM pairs (in the value group #2 .. #10:#2 .. #10). The average CSATSIM value for 7-20 is 0.249; the average CSATSIM value for 6-15 is 0.248.

In the sphere of the relatively small number of SCs in these sections of music, perhaps our best context for CSATSIM values is once again obtained by comparing all the SCs in the group to each other. As we did in our analysis of the first movement of Stravinsky's *Three Pieces for Clarinet Solo*, we will begin by focusing on the very closest relations in the group of sets. We will then gradually broaden our focus, examining increasingly liberal cutoff points for "similarity." Our cutoffs will be determined by the relatively large gaps in CSATSIM values on the sorted list in figure 4.2. The first large gap that we notice is between 0.061 and 0.111. If we filter out set pairs which yield CSATSIM values greater than 0.061, we are left with the single relationship shown in figure 4.3. As in the figures shown in the last chapter, set pairs connected by a line segment on the diagram at the right of figure 4.3 are those which meet this particular standard for relatedness. The left of figure 4.3 shows the single transitive tuple formed by this particular standard.

The single set pair **B/F** is the only one to meet this most conservative cutoff point. While there are many appearances of **B** (6-15) in the A sections of this piece, the only appearance of **F** (6-14 [013458]) is in the penultimate measure of the final A section (IVA). Not only are the mutual cyclic embeddings of **B** and **F** similar, but, in m. 35, the two sets are also temporally adjacent in the music and are realized in a manner which makes the close connection obvious, particularly when the right-hand and the left-hand parts are examined separately. The first two trichords in the right hand of m. 35 are members of the same set class, as are the first two trichords in the left hand part. Transformationally, the right hand trichords move down a semitone from the m. 35 realization of **B** to **F** ($T_{-1}(\mathbf{B}) = \mathbf{F}$); the simultaneous motion in the left hand part can be described as an inversion about F#3 (the F# below middle C). Such split transformations (interpreting the motion in each hand separately to describe the voice-leading relationships of the whole) have been modeled by Shaun O'Donnell in his recent dissertation, "Transformational Voice Leading in Atonal Music."¹⁷⁵

Returning now to figure 4.2, the next largest gap in CSATSIM values is between 0.111 and 0.166. As with our first cutoff point, there is only one additional set pair that falls within this range. Figure 4.4 shows both CSATSIM relations where the value is less than or equal to 0.111. In addition to relating **B** and **F**, we now also have a relationship between **E** and **H**. **E** (5-27 [01358]) is the "cadential" chord of the section IIIA; **H** (4-26 [0358]) is the cadential chord of section IVA. While they are not connected temporally

¹⁷⁵O'Donnell 1997. Such split transformations have also been described by Joseph Straus. These two chords (**B** and **T**) could also be described as a Klumpenhouwer Network, another, more specific, type of split transformation, first detailed by David Lewin (Lewin 1990 and 1994).

(as are **B** and **F**), they occupy structurally-analogous positions in the form of the piece. As one can see from the list of CSATVs in figure 4.8, both sets have relatively high 5-cycle contents and relatively low amounts of all other cyclic adjacencies. **H** is also quite clearly an abstract subset of **E**.¹⁷⁶

The next gap in CSATSIM values yields several more connections among set classes in these A sections. Set pairs that yield CSATSIM values 0.166 are shown in figure 4.5. All A-section sets except **G** (6-21 [023468]) are now connected to at least one other set on the graph.¹⁷⁷ There is now one transitive tuple of cardinal 3, relating sets **A**, **B**, and **F**. These three sets are all non-cadential sets, and, with the exception of set **G**, which appears once in the last A section, they are the only non-cadential A-section sets. The four sets on the right of the diagram in figure 4.5 are the four cadential chords. While they do not form a complete transitive tuple, each is related to at least one other cadential chord and, with the exception of the relationship asserted between **F** and **E**, there are no connections between the cadential and non-cadential chords.

Moving the cutoff point up to CSATSIM relations 0.212, the next logical gap in figure 4.2, figure 4.6 shows that the connections among the cadential chords have become better established. The four cadential chords are members of two overlapping three-element transitive tuples (shown by the two right triangles at the far right of figure 4.6). Set class **G**, unconnected in figure 4.5, has now joined the group of non cadential chords (if weakly) by virtue of its relationship with **B**. More interestingly, our first four-element

¹⁷⁶Furthermore, the specific realizations of these two SCs draw only from the 5-cycle 7-35 “white note” collection.

¹⁷⁷Described another way, all sets except **U** are now part of at least one transitive tuple of cardinal 2.

transitive tuple connects **A**, **B**, **F**, and **E**. **A**, **B**, and **F** are all non-cadential chords, but **E** is the cadential chord of section IIIA. There is therefore a strong distinction between the cadential and non-cadential chords in all the A sections except the third. Given that the relationship between **A** and **B** is stronger than that between either **A** and **E** or **B** and **E**, we can still assert that a greater distinction will be made in comparing either **A** or **B** to **E** than in comparing **A** and **B** to each other.

If we raise the similarity cutoff point to $CSATSIM = 0.232$, we find that all eight A-section sets are found to be members of at least one transitive tuple of three or more elements. These relationships are shown in figure 4.7. All the non-cadential sets (**A**, **B**, **F**, and **G**) now form a single transitive tuple. The only relationship that is not asserted among the cadential chords is one between **D** and **H** (the last chords of sections IIA and IVA). The connection between them— $CSATSIM(D, H) = 0.253$ —would be included were we to move the threshold for similarity up one more notch. As is, the groups of non-cadential and cadential chords cluster together, with the only overlaps occurring between **E**, the cadential chord of the third A section and three of the non-cadential chords.

Shifting our focus away from the *most* related sets to examine the *least* related sets, we see that the four highest $CSATSIM$ values are between sets **C** and **F**, **C** and **G**, **C** and **A**, and **C** and **B**. The intersecting set is clearly **C** (5-35 [02479]). While **C** is closely related to all the other cadential chords (by $CSATSIM = 0.212$ —see figure 4.6) it forms the most dissimilar relations with all the non-cadential harmonies, and the set to which it is least related is **B**, the only one to which it is temporally adjacent. **C** is so distinct primarily because

it is the only cyclic set in this group. As one can see from figure 4.8, which contains the CSATVs of all the A-section chords, **C** is a 5-cycle set and each of the other cadential chords also has a relatively high amount of 5-cycle adjacencies (this, we will recall, was the strongest reason for the connection between **E** and **H** at CSATSIM = 0.111).¹⁷⁸

Returning now to the statistical summaries of sets **A** and **B** that we detailed at the beginning of the present section (4.2). The average CSATSIM values for sets **A** and **B** when compared with all other #2 through #10 set classes were 0.249 and 0.248 respectively. We will call the quality of being closely related to relatively few other SCs *singularity*. There are only three hexachord classes that are *less* singular than 6-15; 6-15 is therefore the fourth least singular hexachord class. There are eleven other septachord classes which are less singular than 7-20. Both SCs are relatively non-singular for groups of their respective cardinalities.

Unlike the first two set classes, SC **C** (5-35) — the first cadential chord — has a very high degree of singularity. Its average CSATSIM value (when compared to all other #2 through #10 set classes) is 0.376. It is, in fact, in a tie with its M-relation, 5-1 [01234], as the most singular pentachordal set class.¹⁷⁹ That close ties exist between **C** and the other three cadential chords (using data from the CSATSIM 0.212 graph) reveals quite a lot about the strength of those comparisons. **E** and **H**, the second most similar set pair among the A-section chords are also both very singular set classes; While the

¹⁷⁸And, as a result of the high 5-cycle content, each of these cadential sets also has relatively low 4- and 6-cycle contents. See figure 2.5 in chapter 2.

¹⁷⁹M-related SCs have the same average CSATSIM value. This is understandable, since the M operator effectively exchanges the values in the ic1 and ic5 columns of an ICV or CSATV (1-cycles and 5-cycles share the same periodicity). Therefore, $CSATSIM(A, B) = CSATSIM(C, D)$ for all $A, B, C,$ and D where $A = T_n M(C)$ or $A = T_n MI(C)$ and $B = T_n M(D)$ or $B = T_n MI(D)$.

CSATSIM(**E**, **H**) value 0.111 appears very close, it is not as close any many of the CSATSIM values that we will see in our examination of the C-section chords. However, appendix G reveals that of all the 38 pentachordal SCs, **E** (5-27 [01358]) is the one which is most closely related to **H** (4-26 [0358]). Furthermore, **C** is the fourth closest #5 SC to **H**. While the CSATSIM(**E**, **H**) and CSATSIM(**C**, **H**) values of 0.111 and 0.162 suggest fairly close relations, they perhaps belie the fact that these are two of the closest relations possible among these very distinct set types. While we would not want to infer that the 0.111 relation between **H** and **E** is, in some sense, closer than a 0.080 relation between two other sets that have a lower degree of singularity. We do want to note that there are fewer SCs that are related to **H** and that the composer has chosen cadential SCs that are as similar as possible without reusing the same SC.¹⁸⁰

We will now take a step back from the CSATSIM values and examine, or in some cases re-examine, the shared features of the SCs that group together in our prior charts. The CSATVs of the eight **A**-section chords are provided in figure 4.8. The four cadential chords (**C**, **D**, **E**, and **H**) are all distinguished by their rather high ic5 content. The first cadential chord (**C**) is, as we mentioned, a 5-cycle chord, with all five members forming an unbroken 5-cycle fragment. The others are not pure 5-cycle chords, but they largely share the same ic profile: put broadly, each of the chords is high in ic5 content, low in ic1, ic4, and ic6, and each has a moderate amount of ic2 and ic3.

¹⁸⁰We do not believe that ordinal lists of CSATSIM values are as successful a determinant of SC similarity as the system we have been using because singular and less-singular SCs would necessarily have the same number of close relations.

The four non-cadential chords (**A**, **B**, **F**, and **G**) are more difficult to characterize, though they are, in general, more closely related to one another. Each of the chords is rather low in ic1, ic2, and ic6, and each has a moderate amount of ic5. None of the chords, however, has a very high degree of any particular interval class. This feature, or perhaps we should say this “flat distribution” of ics is closely associated with a low degree of singularity. Indeed, we can make the very broad generalization that chords (particularly tetrachords and larger) which are either cyclic sets or similar to cyclic sets are frequently more aurally distinctive (hence more singular) than those which are not highly saturated with one or more ics.¹⁸¹ Such cyclic or near-cyclic sets are quite common in certain repertoires—particularly in music from the first part of this century— and, perhaps consequentially, are generally more easily identifiable by ear than pcsets which do not resemble cyclic sets.¹⁸²

4.4 The C (♩ = 88) sections.

Each of the four C sections is initiated by a four-note *pianissimo* tone cluster in the very bottom register of the piano that serves as a pedal chord throughout the section (in three of the four C sections, the remainder of the

¹⁸¹Strong parallels can be drawn between the idea of singularity within a particular system for set categorization (in this case CSATV) and the idea of fuzzy entropy used by fuzzy logicians. Fuzzy entropy is the point at which an event is as similar as it is dissimilar to a particular constant (or point of reference). See Kosko 1993, 126-135.

¹⁸²Or, put another way, there are non-set-theoretic names in the musical vernacular to help us identify these rather common groups. 1-cycle sets are frequently called clusters; 2-cycle sets of #6 or smaller are called whole tone collections; 3/6-cycle (multiply cyclic) sets are frequently identified as diminished triads, diminished seventh chords, or members of an octatonic collection; 4-cycle sets are frequently related to the augmented triad, and a scalar arrangement of [014589] is called “hexatonic” by some, and others who use jazz vernacular sometimes refer to it as the “augmented scale;” and 5-cycle sets are frequently identified as quartal or quintal harmonies. Another potential reason that cyclic sets are readily identifiable is that the number of distinct sub-SCs in cyclic SCs is fewer than the number of sub-SCs in non-cyclic SCs.

section is only a single chord). This invites a question of segmentation: do we consider these clusters to be part of the chord or chords which follow them and are sounded over the same pedal? The fact that they resonate simultaneously suggests that we should, however given that the attacks are distinct, that they are separated by a large interval, and that the four lowest notes of the piano do not speak distinctly, particularly when played extremely quietly suggest distinguishing these chords in the groups to be analyzed.¹⁸³ Accordingly, we will distinguish the low 4-1 [0123] cluster from the higher chords which sound over it.

The chords in the C sections are members of seven distinct set classes (the low 4-1 cluster plus six others). Not surprisingly (and appropriately), the low clusters are rather distantly related to the other SCs. There is, however, a very close relationship between each pair of higher chords. Once again in example 4.1, each chord in the four C sections is labeled with both its set class and a letter name for ease of comparison. Figure 4.9 shows a CSATSIM comparison matrix; the box drawn on it includes all the higher chords and excludes the relationships between the higher chords and the four-note cluster. Figure 4.10 provides a list of all the CSATSIM values from figure 4.9, sorted from lowest to highest.

Unlike the sorted list of CSATSIM values in figure 4.2, here there is really only one very clear gap in values, and that is between 0.179 and 0.330, a much more dramatic dearth of values than we saw among the pairs of A-section chords. This, then, will be our single dividing point for similar versus

¹⁸³In fact, given the extremely low register and volume of these clusters, we might regard them simply as “sounds” and not as “harmonic chords.” As such, we might have just cause for omitting them from our analysis entirely.

dissimilar set pairs. As in the A section comparisons, the chords in these sections fall into two distinct families. Rather than distinguishing the final chord of each section from the others, this section distinguishes the first chord from the others. The average CSATSIM comparison of the cluster chord (**I**) to the others is a high 0.360; the average comparison among pairs of the other chords in the C sections (**J, K, L, M, N, and O**) is an extremely low 0.096. Figure 4.11 shows a list and graphic representation of the transitive tuples formed by SC pairs related by CSATSIM ≥ 0.179 .

There are, as one can see, only two transitive tuples at this cutoff point: **{I}**, the chromatic cluster, and **{J, K, L, M, N, O}**, the remainder of the chords. The six values that are (considerably) higher than our cutoff are those that relate **I** to the other SCs. Given this extremely distinct split in SC types, we do not gain much more information by looking at more conservative similarity cutoff values. The distinction between and among the two set-class families represented by these transitive tuples can be discerned by examining the sets' CSATVs. These are shown in figure 4.12.

While the cluster (**I**) is a 1-cycle set which maximizes ic1 and minimizes ics 4, 5, and 6, the six higher chords do not even come close to maximizing any of the ic cycle types. The fact that they are all rather neutral with respect to any particular interval class makes them rather non-singular. One could certainly say that this lack of singularity is their strongest common trait.

4.5 The B ($\text{♩} = 112$) sections.

Arguably, the B sections present the greatest analytical problems in this piece, both with regard to pcset relatedness and segmentation. Messiaen's

settings of bird song (found only in these B sections) are, according to the composer, based upon his “free” field transcriptions. It is not surprising, therefore, that much of the pcset material seems non-systematic—particularly when compared to Messiaen’s homophonic passages which both precede and follow the short outbursts of blackbird calls. While it is difficult to find groups of related sets in parallel places in the B sections, as we have done in the A and C sections, Messiaen does utilize his third mode of limited transposition for most of the opening gestures. This mode is equivalent to SC 9-12 [01245689a]—the nine-note 4-cycle set class. Section IB is entirely comprised of notes from a single realization of 9-12, and the pc material in sections IIIB and IVB draws heavily from the same realization of 9-12.

Aside from a few recurrences of 8-24 [0124568a]—one of only two #8 subsets of 9-12—there are very few closely-related pairs of pcsets in these sections.¹⁸⁴ Further, the segmentation process is not nearly as obvious as in the other two recurring short sections of this piece (A and C), and one must address this problem before investigating the relatedness of the musical groups. Do we process these polyphonic settings of the bird calls melodically, harmonically, or both? In other words, do we hear the melodies separately, or do we perceive the combinations of the melodies as a single entity? How we answer these questions naturally effects the way that we segment the notes into sets to be treated analytically.¹⁸⁵

¹⁸⁴This point might emphasize the fact that simply because two sets have a mutual superset does not mean that they themselves are similar.

¹⁸⁵In fact, one could reasonably say that they detect the melodic entities more than the harmonic entities during a given performance and hear the harmonies more clearly than the linear elements during some other performance. One’s attention might even shift back and forth between these (and other) entities. In our analysis, when we privilege simultaneous over successive events, we are admittedly making an arbitrary analytical decision. Such decisions almost invariably arise during the analytical process and, while they can be systematized using various segmentational algorithms, we prefer our

Let us examine some concrete examples of the segmentation issues that arise. In mm. 5-6, should we treat the right- and left-hand parts separately or in tandem? In either case, do we place grouping boundaries at the sixteenth notes that conclude some gestures? If so, are we comfortable with the difference in the sizes of the groups (seven notes followed by three notes followed by two notes)? If we treat the two hands in tandem, does m. 5 contain one group or two? While the former might be the more obvious answer, it is easily heard as two groups: the first four (pairs of) notes separate from the last three (pairs of) notes. This segmentation is justifiable because the fifth note pair exactly duplicates the first note pair and it is approached by the largest leap in the gesture. Similar questions arise in mm. 14-16: is measure 15 (or measure 16) one group or two? Rhythmically, it would seem like only one, but the fourth note-pair is a repetition of the third, suggesting a possible reading of two groups of three pairs (forming two distinct hexachords) rather than one group of six pairs (which form a single ten-note SC).

By extension, we can question whether mm. 24-25 group into one, two, or three segments (again assuming that we are grouping the right- and left-hand parts in tandem). The leap down to the B \flat /G dyad from the third to the fourth note pair is identical to that in m. 5. We should also consider whether the E/E \flat dyad at the end of the measure constitutes a separate group. If so, can we effectively compare a dyad class to a much larger set class? (And, for that matter can we compare the ten note SC in mm. 14-16 to smaller SCs very effectively?) Does it matter if we cannot (i.e., do all parts of an analysis need

present approach. Different, and even overlapping, segmentations of the same music might well evoke different analytical results. Such alternate readings would doubtless be interesting, but they not in the scope of this study.

to be absolutely cohesive)—particularly since this gesture does seem rather different from anything else in the music?

The answers that any particular analyst provides to these questions will have tremendous impact on his or her analytical reading. At the beginning of the first B section, for example, if we hear the entirety of m. 5 as a single segment, the set class is 9-12 [0124569ab], a very singular 4-cycle set. If, however, we hear the right-hand and left-hand parts separately, the set classes yielded are 5-30 [01468] and 5-27 [01358], respectively—two rather similar set classes (yielding a CSATSIM relation of 0.171), but neither is cyclically similar to 9-12.¹⁸⁶ If we hear the two hands acting in tandem, but perceive a division at the fifth note (as suggested above), then the two SCs heard are 8-24[0124568a] and 6-14 [013458]. 8-24 is a subset of and quite similar to 9-12 (CSATSIM = 0.156), but 6-14 is similar neither to 8-24 nor to 9-12 (CSATSIM(6-14, 8-24) = 0.455; CSATSIM(6-14, 9-12) = 0.420).

From this perspective, it would seem that either separating the hands or treating the entire measure as a whole would yield more coherent readings. However, we do not believe that similarity comparisons are necessarily (or even probably) a very effective and musically salient means of determining a preferred segmentation. It might well be possible to segment many pieces by like set class. Ideally, though, musical segmentation should arise as the result of careful listening or rehearsal, not from processing a score using some abstract criterion or criteria. When making decisions regarding musical segmentation, it is almost inevitable that we consider various criteria. When

¹⁸⁶Some inclusion-based models of similarity might find these 5-30 and 5-27 to be more similar to their parent mutual superset, 9-12. However, Rahn's ATMEMB, Lewin's REL, and our AvgSATSIM_n all agreed that, among these three set classes, 5-30:5-27 forms the closest relation, followed, in order, by 9-12:5-30 and 9-12:5-27.

such criteria suggest conflicting readings (as frequently happens), it is up to the analyst to decide which factor (or factors) is (are) most compelling in any particular musical scenario. Even if a complex algorithm could be developed to make such decisions, it would not, in our opinion, be desirable because of the creative flexibility that analysts would lose.¹⁸⁷

The segmentation that I have come to prefer yields some interesting relationships, but does not suggest divisions into families of related set types nearly as neatly in the A and C sections. In all three B sections, I group the two hands together at all times. In section IB (mm. 5-6), I perceive divisions after the fourth thirty-second note and after each of the three sixteenth notes (thus four segments of 4 + 3 + 3 + 2 attack points). In section IIB (mm. 14-16), I tend to process m. 14 as a single segment (4 attack points), m. 15 as two segments (3 + 3 attack points, divided at the sixteenth notes), and m. 16 as two more segments (3 + 3 attack points, divided at the repeated notes). In IIIB (mm. 24-25), I perceive two segments in m. 24 (4 + 5 attack points)¹⁸⁸ and three more segments in m. 25 (4 + 3 + 3 attack points). IVB (m. 37) is perhaps the least ambiguous of these sections. In it, I perceive only two segments (8 + 4 attack points), separated by the sixteenth rest.

Figure 4.13 provides a list of all the set classes yielded by my segmentation of these four B sections. Figures 4.14 through 4.17 provide CSATSIM comparison matrices for my segmentation of sections IB through

¹⁸⁷This point will be addressed more critically in the final chapter.

¹⁸⁸While it would be rhythmically consistent to group the repeated E/E^b dyad at the end of m. 24 apart from the notes that immediately precede it, I do not really hear the two-element group as a separate entity.

IVB respectively. There are, as one can readily see by the underscored numbers relatively few small numbers (e.g., below 0.2) in these four matrices.

With few close connections present among the four B sections, we widen our focus to relationships *between* the B sections. Figure 4.18 provides a CSATSIM comparison matrix for all the B-section sets, allowing us to assert relationships between as well as among the four B sections. Once again, however, we see that there are not a great number of close relationships to be found. Furthermore, the close relations present are frequently not between sets in parallel places in the structures of their short sections (e.g., opening, closing sets).

For example, three of the four B sections open with set class 8-24. The only B section that does not (IIB) opens with 8-3 [01234569]. $CSATSIM(8-24, 8-3) = 0.440$, a very distant relationship. In fact, none of the SCs in section IIB are particularly closely related to 8-24. Similarly, each of the four B sections ends with a different set class (4-21 [0246], 6-z49 [013479], 5-22 [01478], and 7-26 [013579], respectively). 4-21 is similar to none of the other section-ending SCs, yielding CSATSIM values 0.473, 0.463, and 0.245 respectively. **W**, **X**, and **Z** are all relatively similar to each other,¹⁸⁹ though, so we can say that three of the four section-ending and section-beginning SCs are similar to each other (in the case of the section beginning SCs, the similarity is maximal given that the three similar sets are all members of the same set class). Though interesting, this is by no means as compelling as the relations found between structurally parallel chords in the A and C sections.

¹⁸⁹ $CSATSIM(W, X) = 0.195$; $CSATSIM(W, Z) = 0.185$; $CSATSIM(X, Z) = 0.179$.

Figure 4.19 contains a sorted list of all the CSATSIM relations among B-section sets (those listed in figure 4.13 and labeled in figures 4.14 through 4.17). Figure 4.20 shows transitive tuples formed by CSATSIM relations where the value is less than or equal to 0.153 (the gap in values between 0.153 and 0.174 is one the largest in figure 4.19). While, as one can see, there are a number of closely-related set pairs, there are only two pairs that occur within the same B section. **P** and **S** both occur in IB and **U** and **W** both occur in IIB. Neither are adjacencies within their respective sections, however.

Figure 4.21 contains a list and graph of the transitive tuples formed by CSATSIM relations less than or equal to 0.195. While many more set relations are asserted, there is only one more set pair that occurs within a particular B section: **X** and **Z** both occur in IIIB, but, as before, they are not temporally adjacent. Thus, even though the graph of transitive tuples does not look less interesting than the graphs created using A-section sets, we find that sets that occur at like points in the B sections are not, by-in-large, related to each other and neither are sets that occur within a particular section (and particularly in adjacent groupings).

We therefore cannot assert much, if any, unity among the set types at parallel or adjacent locations in the B sections of this work. This is not to say that these sections are unsystematic, only that Messiaen's settings of the blackbird calls are difficult to interpret using CSATSIM as vehicle for comparing pcsets. CSATSIM was, however, quite useful in interpreting the progression of set types in the A and C sections of this work (i.e., the sections where bird calls are not represented). It is entirely possible—likely, even—that other similarity indices or other means of analysis (e.g., those that examine

contour, texture, pitches in their register (p-space), etc.) might prove fruitful for these sections of music. We do not, after all, posit CSATSIM as *the* way to analyze a piece of music, but merely as a tool that is useful in certain situations. The use of CSATSIM to analyze the chords in the A and C sections was, by contrast, quite fruitful in establishing musical connections.

Chapter 5

Extensions and Conclusions

In this chapter, we will summarize and critique two theorists' work in the realm of pitch space (p-space)¹⁹⁰ and we will lay the groundwork for a saturation-based system of relating pitch sets (psets) in p-space, thus extending the saturation-based pc-space tools that form the heart of this dissertation. We will then address some criticisms that have been leveled against both similarity indices and pc-space set theory in music analysis. We will conclude by recapitulating the primary ideas put forth herein.

5.1 Similarity and equivalence in pitch space

To this point, we have only discussed tools for analysis that operate within the domain of pitch-class space. We saw this as a necessary constraint in creating a coherent and finite study—one which was bound together by the common theme of measuring the degree of *saturation* that one interval, cyclic fragment, or subset enjoys in a pcset. Despite limiting ourselves to the domain of pc-space, we do not subscribe to the idea that this domain alone is sufficient in producing convincing musical analyses. We do believe, however, that an enumeration of the pc-space characteristics of a set is very helpful in understanding much atonal music. This point will be expanded later in this chapter.

While the majority of recent atonal theory has operated in the realm of pitch-class space, a few theorists have begun to create a vocabulary that deals

¹⁹⁰As we mentioned in the first chapter, pitch-class space assumes octave equivalence; pitch space does not.

with other dimensions as well: most notably contour and pitch spaces. We touched on some of the work of Marvin and Laprade and Morris at the end of our third chapter, when we demonstrated that, in our Stravinsky analysis, two of the most related segments using pc-based similarity indices were realized using very similar contours.¹⁹¹ Observations such as this one do not necessarily conflict with pcset theory, but rather can complement it, making the resemblance of the two segments seem even stronger. Other theorists who have been involved in developing new methods for examining musical contour include Michael Friedmann, Larry Polansky, and Richard Bassein.¹⁹²

Developing a consistent and rigorous vocabulary for pitch space structures has proven problematic. This is the realm of music analysis where we consider pitches coupled with their register. Any two Cs are equivalent in pc-space; in pitch space (p-space), however, octave designation is an integral part of pitch identity. By expanding our analytical horizons outward, we account for the abstract intervallic (or subset) content and, in the process, create a portrait of a pitch set (pset) in its original spacing. An open and a close spacing of a particular triad, for example, would be differentiated in p-space.

John Rahn, in his 1980 textbook *Basic Atonal Theory*, distinguishes between pitches and pitch classes and provides a simple heuristic for

¹⁹¹Marvin and Laprade 1987 and Morris 1987.

¹⁹²See Michael Friedmann, "A Methodology for the Discussion of Contour: Its Application to Schoenberg's Music," *Journal of Music Theory* 29 (1985): 223-248; Elizabeth West Marvin, "The Perception of Rhythm in Non-Tonal Music: Rhythmic Contours in the Music of Edgard Varèse," *Music Theory Spectrum* 13 (1991): 61-78; Larry Polansky and Richard Bassein, "Possible and Impossible Melodies: Some Formal Aspects of Contour," *Journal of Music Theory* 36 (1992): 259-284; Robert Morris, "New Directions in the Theory and Analysis of Musical Contour," *Music Theory Spectrum* 15 (1993): 205-228; and Elizabeth West Marvin, "A Generalization of Contour Theory to Diverse Musical Spaces: Analytical Applications to the Music of Dallapiccola and Stockhausen," in *Concert Music, Rock, and Jazz since 1945*, ed. Elizabeth West Marvin and Richard Hermann (Rochester, NY: University of Rochester Press, 1995): 135-171.

numerically labeling pitches and intervals in both p- and pc-space.¹⁹³ In 1987, Robert Morris more completely detailed the construction of psets as well as p-space intervals, cycles, operations, and pset classes.¹⁹⁴ With the creation of pset classes, he provided a method of judging pset equivalence.

In both Rahn 1980 and Morris 1987, p-space elements were introduced en route to a much broader exposition of pc-space possibilities. In 1994, however, Morris presented and later published (in 1995) a methodology for determining not only equivalence among psets, but also similarity.¹⁹⁵ We will review the major points of this recent article and examine how some of the techniques for determining pset-class (PSC) similarity might be extended to form a saturation-based PSC index. Before doing so, we will summarize and critique a recent, and more informal, method for representing different pset spacings: Brian Robison's fractional interval-class vector.

5.2 Robison's fractional interval-class vector

In an effort to maintain some degree of pitch-class and interval-class identification while accounting for different spacings of musical structures, Brian Robison at Cornell University proposed a variation on the standard six-argument ICV.¹⁹⁶ As the name indicates, Robison's fractional ICV (fractICV) uses fractional values rather than simply integers. In Robison's vector, pitch intervals of one through six semitones receive the "full" value of 1. Intervals between seven and eleven semitones, are counted as some fraction of 1 (the

¹⁹³Rahn 1980, 19-31.

¹⁹⁴Morris 1987, 36-58

¹⁹⁵Robert Morris, "Equivalence and Similarity in Pitch and their Interaction with Pset Theory." Paper presented at the annual meeting of the Society of Music Theory: Tallahassee, FL, 1994.

Published in 1995 under the same title in *Journal of Music Theory* (see citation for Morris 1995).

¹⁹⁶Robison 1994.

fractional values of interval 7s are recorded in the ic5 column of the vector, the fractional values of interval 8s are recorded in the ic4 column, etc.); intervals between thirteen and eighteen semitones (notice that we skipped the octave) are counted as some still-smaller fraction of 1, etc.

Robison suggested three different algorithms for deciding just how much lower an interval 7 should be weighted compared to an interval 5 and how much lower an interval 17 should be weighted compared to an interval 7, etc. The three algorithms for decrementing the fractional values of increasingly large intervals are 1) “linear,” using a steady decrement of 0.05, 2) “moderately exponential,” and 3) “strongly exponential.” The two exponential systems are “based on reciprocals of e raised to $\frac{1}{5}$ and $\frac{1}{2}$ the index of transposition” (where e is “the base of the system of natural logarithms, approximately 2.718 in value”).¹⁹⁷ Figure 5.1 reprints Robison’s table of fractional interval counts.

Suppose, for example, we were weighting the interval of seventeen semitones ($i = 17$) using Robison’s three algorithms. The interval 17 is a member of ic5; it is weighted as follows using Robison’s linear scale. The interval of 5 semitones ($i = 5$) is the most compact member of ic5, it is recorded as a value of 1 in the ic5 column, using Robison’s linear scale. The next largest member of ic5, $i = 7$, is recorded as a value of 0.95 in the ic5 column. Finally, $i = 17$ (an octave plus a perfect fourth) is the next largest member of ic5 and it is recorded as 0.9 in the ic5 column. Using Robison’s moderately exponential decremental scheme, we take the reciprocal of e raised to $\frac{1}{5}$ (.2) of the index of transposition. The index of transposition for intervals between

¹⁹⁷Robison 1994, paragraph 9.

13 and 18 semitones is 2 (see figure 5.1). The product of 2 and 0.2 is 0.4; $e^{0.4} = 1.491$. The reciprocal of 1.491 ($\frac{1}{1.491}$) is 0.670: this is the “moderately exponential” value recorded in the ic5 column for each interval 17 present. Finally, to calculate the strongly exponential weighed value for interval 17, multiply the index of transposition (2) by $\frac{1}{2}$ (.5). The product of 2 and 0.5 is 1. $e^1 = 2.718$. The reciprocal of 2.718 ($\frac{1}{2.718}$) is 0.368: this is the “strongly exponential” value recorded in the ic5 column for each interval 17 present.

To illustrate Robison’s vector, let us examine the first piano chord from the first movement of Messiaen’s *Quatuor pour la fin du temps*.¹⁹⁸ The chord is spaced as follows {F3, G3, B♭3, C4, E♭4, B4, E5}, where C4 = middle C. This is a literal pitch transposition of the first chord in the Messiaen piece analyzed in chapter 4 (“Le Merle noir”).¹⁹⁹ It is therefore also a member of SC 7-20 [0125679]. The (conventional pc-space) interval-class vector of 7-20 is $\langle 4,3,3,4,5,2 \rangle$.

In this case, the four ic1s come from the following note pairs: {C4, B4}, {E♭4, E5}, {B♭3, B4}, {F3, E5}. None of these note pairs is literally separated by (pitch) interval 1. Rather, the first is separated by interval 11, the second and third by interval 13, and the fourth by interval 23. Using Robison’s strongly exponential decrement scheme (as he did in his subsequent letters to the MTO electronic mail list), the intervals 11, 13, 13, and 23 receive fractional values of 0.607, 0.368, 0.368, and 0.223 respectively. The fractional ICV value for ic1 is calculated by adding these together (1.565).²⁰⁰

¹⁹⁸Robison also demonstrated his vector using this chord in a letter to the MTO-talk electronic mail list on March 29, 1995.

¹⁹⁹Or, to be more historically correct, we should say that the latter is a transposition of the former.

²⁰⁰This can be expressed more formally as follows: let d = the decremental formula (either the exponential or linear schemes) and i be a particular p-space interval class. $d(i)$ produces the decremented

By contrast, there are five members of ic5s separating elements of this pset. They are found between the following note pairs: {F3, B♭3}, {G3, C4}, {B♭3, E♭4}, {B4, E5}, and {F3, C4}. The first four of these dyads produce intervals of 5 semitones. They are therefore maximally compact representations of ic5 and counted as one full ic5 each. The last dyad contains notes separated by 7 semitones; it is counted as 0.607 of an ic5. The ic5 argument in the fractional vector of this set is therefore the product of 4 and 1, added to the product of 1 and 0.607, totaling 4.607: a value that is rather close to 5, the actual number of ic5s.²⁰¹

If we compare the fractional interval-class vector values to those of the standard ICV values (as we do in figure 5.2), we can get some idea of the p-space distribution of a particular musical segment. Such a comparison is, unfortunately, not investigated by Robison (his article is admittedly a “work-in-progress”). Rather, he uses the vectors in their own right and, in doing so presents data that, in our opinion, is rather difficult to interpret.²⁰²

In a footnote, Robison mentions that he does “not mean it [the fractional vector] to imply a reduced degree of ic membership for compound intervals!” (emphasis Robison).²⁰³ Robison’s intentions notwithstanding, it seems to us that this is exactly what his numbers do suggest. The ic 2 argument of the

(weighted) value of interval i . Each of the six fractional ICV arguments (ic) is derived by taking the sum of each $d(i)$ where i is a member of a particular pc-space ic . Therefore, $\text{fractICV}_{ic} = \sum_i d(i)$.

²⁰¹Some set classes can never yield a “maximally compact” fractional ICV: one that is equal in every argument to the standard ICV. For example, [027] (3-9) contains two ic5s and a single ic2. If the SC is realized as a quartal harmony, so that the two ic5s are presented as perfect fourths, then the ic2 appears as the outermost interval (as a minor seventh). If the pitches are arranged so that the ic2 appears as a major second, then at least one of the two ic5s will be larger than a perfect fourth.

²⁰²In his *MTO* letter March 19, 1995, he does show pc-space ICVs in tandem with his fractional ICVs, though no formal comparison between the two is initiated.

²⁰³Robison 1994, footnote 9.

vector above is considerably higher than the ic1 argument, despite the fact that there are more ic1s than ic2s in the sets.²⁰⁴ Without comparing these fractional values to the original ICV, it is impossible to separate arguments that indicate a large number of wide intervals versus those which indicate a small number of narrow intervals. Thus, while the vector is derived from p-space intervals, it alone actually yields very little information regarding p-space intervals. There is, for example, no way to tell the width of the interval between the outermost voices.²⁰⁵ This criticism is valid to various degrees, depending on which decrement scheme is chosen.

Robison has, in essence, created a version of the pc-space ICV, weighted based upon p-space information. Perhaps the strongest evidence that this is vector is not functionally representative of p-space relationships is the exclusion of interval-class 0. Admittedly, most theorists do not include zero as an argument in the ICV because the number of interval 0 in a particular pcset is equal to the cardinality of the pcset.²⁰⁶ However, when distinguishing between a perfect fourth, a perfect fifth, and a perfect eleventh, for example, it is necessary to incorporate all intervallic relations including octaves.

Consider two psets {C4, G4, C5} and {G4, C5}, for example. They both have the same pitch-class content, but the first has an octave duplication and the second does not (the second is also a literal pitch subset of the first).

²⁰⁴Furthermore, the intervals being weighted are not spaced evenly. Interval 5 and the next member of ic5, interval 7, for example, are only separated by two semitones; interval 7 and the next member of ic5, interval 17, are separated by ten semitones. Robison's strongly exponential system of weighting accounts for this discrepancy better than do his linear and moderately exponential systems.

²⁰⁵In some cases, it would be possible to analyze the fractional ICV values and determine the smallest fraction, thus, largest interval, included. Even where this is possible, deconstructing these weighted values would not be a simple task.

²⁰⁶A notable exception is Robert Morris who includes the number of ic0 as the first place of his seven-argument vector in the SC tables in Morris 1987 and 1991.

Robison's solution for dealing with octaves is to remove one of the notes from consideration. He routinely deletes the duplicated pcs that form the larger intervals (that are members of non-zero pc interval-classes) relative to the other notes of the set. In the first set above, the C4 is deleted from consideration because it forms an interval 7 with the G4, whereas the C5 forms an interval 5 with the G4. The fractional interval-class vector of both sets above is therefore $\langle 0,0,0,0,1,0 \rangle$ (the same as their pc ICV). These two sets, then, are functionally identical using Robison's tool.

We believe that the fractional vector should, at very least, be expanded to seven arguments, including ic_0 (as Morris does with the ICV in Morris 1987 and 1991). The fractional vector of the first set using Robison's highly exponential decrements would be $\langle 2.607, 0, 0, 0, 1.607, 0 \rangle$, while the fractional vector for the second would be $\langle 2, 0, 0, 0, 0, 1, 0 \rangle$, a significant difference that more accurately represents the different cardinality and spacing of the two psets.

While the goal of combining the notions of p- and pc-space interval-classes into a single composite value for comparison is certainly admirable, we believe that this vector does not adequately represent either type of interval. Furthermore, the cardinality of the set is not clear by the fractional vector alone, making set comparisons difficult. It is, however, worth reiterating that this vector was first posited as a work-in-progress, and Robison is certainly to be congratulated for first attempting to conjoin p- and pc-space data.

5.3 Morris's Pitch-Space Similarity Measures

In Robert Morris's work (1987 and 1995), he introduces some new tools that deal exclusively with p-space pitches and intervals and that uniquely define pitch sets and pitch set classes. He also posits some tools that help integrate the realms of p- and pc-space. Before describing his apparatus, we should define a few terms as Morris uses them. A *pset* is an unordered set of pitches. All pitches are defined numerically: middle C (C4) is represented by the number 0; C#4 = 1; B3 = -1, etc. In p-space, *interval classes* are only contrasted with *directed intervals*. In other words, directed intervals are preceded with a sign (+ or -) while interval classes are unsigned.

A pitch set class (PSC) is defined by what Morris calls the spacing of a pset. The spacing is represented by a series of adjacent intervals, calculated from the lowest to the highest notes. Let us return to our Messiaen chord {F3, G3, Bb3, C4, Eb4, B4, E5}, which we will call *X*. Using Morris's nomenclature, this set is represented as {-7, -5, -2, 0, 3, 11, 16}. The numerical presentation facilitates intervallic calculation (through simple subtraction). The spacing of this chord, or SP(*X*), is [2, 3, 2, 3, 8, 5]. The p-space inversion (or "dual") of this PSC is obtained by reversing the intervals from top to bottom [5, 8, 3, 2, 3, 2]. A pset and its inversion are members of the same PSC.²⁰⁷

Morris also created a tool for combining the idea of pitch-class (octave) equivalence with the intervallic spacing. The pitch-class interval series (PCINT) only uses numbers 0 through 11 as its arguments. If two psets share the same spacing (SP)—as do the first chord in both Messiaen's *Quatuor*

²⁰⁷The inversion of a specific pset is derived by inverting all pitches about the axis of middle C (0). This amounts to nothing more than a simple reversal of all the signed numbers. 1 inverts to -1, -1 inverts to 1, -27 inverts to 27, 0 inverts to 0, etc. The use of 0 as an axis of inversion is arbitrary, no doubt chosen for convenience. The inversion operation in pc-space (as used by Morris, Forte, and others) is also arbitrary, set at pc0.

pour la fin du temps and his “Le Merle Noir” from *Petites Esquisses d’Oiseaux*—they are, by definition, transpositions of each other. Two psets that yield a common PCINT, however, are not necessarily transpositionally (or even transformationally) related. For example, the PCINT [2, 3, 2] could be realized as psets {0, 2, 5, 7} or {-5, 9, 26, 57}, to cite just two possibilities. The SP of the first is [2, 3, 2]; that of the second is [14, 17, 31].

SP and PCINT are both ordered sets of intervals. In addition to these, Morris also introduced a partially-ordered p-space interval series. His figured bass classification (FB), like its tonal counterpart, only measures intervals calculated above the bass. FB also allows for octave (but not inversional) equivalence, but it does not specify the exact order that the intervals must appear within a set.

In tonal music, the figured bass symbol $\overset{6}{3}$ refers to a structure where there is some third and some sixth above the bass (these intervals can be either simple or compound). Either the sixth or the third could appear closest to the bass and there could be any number of octave duplications of the three essential notes (including the bass itself). In Morris’s FB, the same holds true. The FB intervals are ordered from smallest to largest, using numbers between zero and eleven.²⁰⁸ The FB 345 could be represented as pset {0, 3, 4, 5} or {4, 9, 20, 31}, to name just two of many possibilities.

For two examples of places where PSCs can be analytically informative, let us return to the Messiaen piece that we examined at length in chapter 4 (refer back to example 4.1). In the four A sections (mm. 1-4, 10-13, 20-23, and 31-36), the pc SC 7-20 [0125679] (denoted **A**) appears 38 times. Of these

²⁰⁸The letters “a” and “b” are used to substitute for the numbers 10 and 11 respectively in both PCINT and FB, effectively eliminating the need to separate arguments using commas.

pitch realizations of 7-20, there are only three distinct p-space SCs used: [2, 3, 2, 3, 8, 5], [3, 2, 3, 2, 6, 5], and [4, 3, 2, 5, 6, 5]. The first two (respectively heard in the first two chords in the piece) are very similar and could be described as split transformations of each other where the lowest five notes of the two sets are inversions of each other ($I[2, 3, 2, 3] = [3, 2, 3, 2]$) and the uppermost dyads are transpositions of each other. These two realizations of pc SC **A** are the only ones used in sections IA, IIA, and IIIA. The third distinct spacing of 7-20 used in this piece is found in section IVA at chords 9, 11, and 13 (the second chords in each of mm. 32-34).

This, then, is an example of several instances of a single pc SC, partitioned into a few distinct spacings. In the same piece, there is also an example of three different pc SCs that are realized using very similar spacings. Measure 35 contains three six-note chords. The first is a member of pc SC 6-15, the second is a member of 6-14, and the third is a member of 6-21 (these are denoted **B**, **F**, and **G** respectively on example 4.1). The spacings of the respective chords are [3, 2, 4, 4, 3], [2, 3, 4, 4, 3], and [3, 3, 4, 4, 2]. Each of these spacings contain the same numbers and types of intervals among adjacent notes (i.e., the same PCINT), though their order is different. In each, there are two adjacent interval 4s (forming an embedded augmented triad) in the third and fourth arguments.

Our CSATSIM index found these three pc SCs to be closely related to each other. The first two of the chords are the closest-related set pair among the four A sections (CSATSIM = 0.061). A more complete analysis than we undertook in the last chapter could certainly take matters of spacing into

account. Clearly, such p-space structures do not negate pc-space relations, but rather they complement them.

In order to calculate a degree of relatedness among the three chords in m. 35, it is helpful to have a similarity measure.²⁰⁹ According to Morris, “defining similarity relations between psets ... is of some merit since it does not have the problems associated with pitch-class similarity.”²¹⁰ The primary such problem that Morris articulates is that a single pitch-class SC “has a multitude of realizations in pitch and time.” Of course p-space SCs also have a variety of different realizations, but each of these is a literal pitch transposition or inversion of each other. The issue of different realizations in time applies equally to both types of SC.

Morris’s p-space similarity relation, called “pitch measure” (PM), compares two psets and returns the numbers of common pitches and p-space interval-classes. Remarkably, that thumbnail sketch of PM comprises almost a complete description of the measure. A simple illustration will elucidate both the data that PM provides and the objects that it compares.

Consider the following two psets: $X = \{0, 4, 7, 13\}$ and $Y = \{-7, -4, 0, 13\}$. The interval-class roster (complete list of interval-classes—adjacent and non-adjacent) of X is $\langle 3, 4, 6, 7, 9, 13 \rangle$; the ic roster of Y is $\langle 3, 4, 7, 13, 17, 20 \rangle$. $PM(X, Y)$ returns two numbers: the first is the number of common tones between X and Y ; the second is the number of common ics between X and Y . X and Y share two common pitches—0 and 13; the two psets share four common ics—3, 4, 7, and 13. Therefore, $PM(X, Y) = 2, 4$.

²⁰⁹These p-space tools provide different means of determining pset equivalence (i.e., different points of view from which to examine a pset), but they do not imply a mode of comparison.

²¹⁰Morris 1995, 223.

Let us examine the results of PM performed on the three m. 35 sets discussed above. For the sake of consistency with the previous chapters we will call these three sets **B**, **F** and **G**. Figure 5.3 shows them notated both musically and using pitch interval notation. It also provides their interval-class rosters and a calculation of the three different PM comparisons among them.

According to Morris, two psets are maximally PM similar if “they are as similar as possible but not identical. ... In the case of trichords, maximal similarity is therefore all (three) ics in common and two pitches in common.”²¹¹ Morris also describes a chain of transformations called FOLDSIM that produces networks in which adjacent psets are maximally similar under PM.²¹² The three sets shown above in figure 5.3 are *not* maximally similar. The pair **B**, **G** are the most similar, exhibiting both pitch and intervallic similarity, but should any of the other pairs should be considered dissimilar? In the “worst” case, there are no pitch intersections, but eleven of fifteen intervals the same.

PM is both simple and clear, requiring no numeric encoding or decoding to calculate or decipher the values produced. In both Morris 1987 and 1995, a vocabulary is introduced that enables us to begin discussing issues of p-space set relations. Rather like the important work of Allen Forte in the early sixties (especially his R_0 , R_1 , R_2 , and R_p similarity measures), this recent work

²¹¹Morris 1995, 231.

²¹²We will not discuss FOLDSIM in detail. It is a formalization of Jonathan Bernard’s *foldings* and *unfoldings*, that maps one pset to another by inverting one of the pitches in a set about another pitch in that same set. As Morris points out, an analogous process is described in Lewin 1987, 189 (there called FLIPEND and FLIPSTART). Additionally, this process can be likened to Lewin’s RICH (retrograde inversion chain), also described in Lewin 1987, 180-188 (just prior to his description of FLIPEND and FLIPSTART).

of Morris provides a starting point for set theoretic analysis using multiple musical dimensions.

As a prelude to our critical commentary on PM, we should mention that Morris's PM and his tools for labeling psets are not and do not pretend to be connected facets of a complete theory (or anything close to that) of pset relatedness. While he provides examples where these tools are analytically informative, the article (Morris 1995) is by no means prescriptive. Perhaps as a result of its open architecture, we feel free to suggest alternative modes of representation that might enable one to associate the worlds of p- and pc-space further. Also, we find Morris's structures highly suggestive of different approaches that one could take to the idea of pset similarity.

Before embarking on our own additions, we shall address a few problems. While the simplicity of PM—Morris's two-dimensional similarity relation—makes it relatively easy to calculate and employ, one faces a potentially difficult task in contextualizing its values. In the figure above, it seems clear that $PM(\mathbf{B}, \mathbf{G}) = 3, 12$ represents greater similarity than $PM(\mathbf{B}, \mathbf{F}) = 0, 12$ and that the latter is more similar than $PM(\mathbf{F}, \mathbf{G}) = 0, 11$. $PM(\mathbf{B}, \mathbf{G})$ is larger in both values than $PM(\mathbf{B}, \mathbf{F})$, and the values of $PM(\mathbf{B}, \mathbf{F})$ are larger in one case and the same in the other compared to $PM(\mathbf{B}, \mathbf{G})$. How would we compare PM values 3, 4 and 4, 3, however? Are two psets that have more intervals than pitches in common more or less similar than two psets that share more common pitches than intervals? Furthermore, how do we compare the PM values generated through the comparison of trichords with those generated through the comparison of tetrachords—or any other cardinality of psets?

The problem of non-contextual values and the associated trouble with comparing any two sets of different cardinality is, I believe, the greatest issue that must be addressed when using this pset relation. Indeed, the problem of comparing psets of different cardinality is a rather treacherous bridge to cross. With an indefinite number of intervals and notes and no accounting for octave doubling, how can one create anything close to a cardinality-neutral system of relations? We will suggest one possible answer to this question in the next section of this chapter. First, though, we will address the logistical issue of comparing larger sets using PM.

In Morris 1995, the author demonstrates his p-space tools using trichords exclusively. As might be obvious from figure 5.3 above, working out ic rosters of larger cardinality sets and/or sets that use larger intervals is much more time consuming than doing the same with trichords. Of course, this is a problem only if one is counting intervals by hand. If one uses a computer (as would be necessary for most of the pc-space similarity indices discussed herein²¹³), the problem disappears in part. Large sets do, however, yield very long ic rosters. A seven-note set, for example, produces twenty-one intervals. Comparing two #7 sets where any interval types are possible is significantly more time consuming than the simple comparison of two #3 sets. While this again could be done by computer, part of the elegance of Morris's PM is that it can be performed "on the fly" without technological assistance. Comparing (and deriving) two 21-place (or larger) lists of intervals may not be

²¹³Interestingly, Morris's ASIM is perhaps the sole example of a pc-space similarity measure that is easily calculable with pencil and paper alone.

conceptually challenging, but it is still a rather daunting task—particularly if performed repeatedly.²¹⁴

We suggest that one way to preserve some semblance of octave equivalence while still dealing with p-space interval counts is to change the mode of representation from a linear string of intervals (of any possible length) to more of a spiral-type model that measures distance using only twelve interval columns, but allows for any number of rows.²¹⁵ The top row would indicate the simple intervals from 0 to 11; the second row from the top would indicate intervals 12 through 23, etc. An example of this *p-space interval-class vector* (PICV) is shown along with a parallel example of Morris's ic roster in figure 5.4. The set used is once again the first piano chord from Messiaen's *Quatuor pour la fin du temps*.²¹⁶

While both Morris's interval-class roster and our PICV convey essentially the same information, we believe that the consistent design of ours makes it more amenable for comparison than Morris's. Our vector is designed similarly to the well-known p-space ICV. For the first row, it utilizes pre-defined columns for each interval from 0 through 11; subsequent rows represent a parallel tally of (increasingly larger) compound intervals. This avoids the necessity of repeating intervals: <2,2,3,3,5,5,5,5,...> becomes a 2 in the ic2 and ic3 columns and a 4 in the ic5 column, etc. Our vector, following

²¹⁴In describing his FOLDSIM procedure, Morris admits problems in talking about folding as a means toward maximal similarity in the case of #4 and larger sets (Morris 1995, 237).

²¹⁵While we consider this to be a spiral model conceptually, for ease of representation we display it as a multidimensional array.

²¹⁶In our PICV, we could have divided the rows of intervals at the tritone as did Brian Robison. We decided on the present system for two reasons: 1) it is our belief that octave-displaced (compound) intervals are far more similar in p-space than inversionally-related intervals (like 5 & 7); and, more importantly, 2) Robison's division at the tritone doesn't allow one to have a member of ics 0 and 6 at each level. This is because there are two forms of ics 1 through 5 in each octave, but only one form of ic0 and ic6.

the protocol of Morris 1987 and 1991, also counts the number of ic_0 in all octaves. In octave 0, the number of ic_0 will always equal the cardinality of the set. For the sake of completeness, counting all members of ic_0 it is an integral part of this design.

Our PICV is useful when working with Morris's SP or particularly his PCINT since one can easily tell how many p-space intervals of a given octave equivalence class are embedded in a set by adding together the values in that particular column. (This is a quicker than taking the interval number mod 12—as one would have to do using Morris's nomenclature.) The PICV would also be a useful first step in deriving a composite p-/pc-space vector such as that used by Robison because it keeps track of exact intervallic distance while facilitating translation into the familiar realm of pc ics. While we believe that our PICV offers a useful alternative representation of Morris's ic roster, its introduction here is really tangential to the larger issue of pset similarity.

These critical points aside, Morris has, to this point, been *virtually* the only theorist to devise p-space analytical tools.²¹⁷ His SP, FB, and PCINT structures and PM relation are certainly useful as they stand. More importantly, however, they provide a firm foundation for further development in this important and growing field.

5.4 Saturation and p-space similarity measures

Because the values produced by Morris's PM have not been placed into one or more specific contexts, it is—as we mentioned earlier—very difficult to compare them with each other in a clear and consistent manner. Similarly,

²¹⁷Two exceptions are Rahn 1980 and Hermann 1993.

Robison's fractional ICV attains a high degree of precision, but the values themselves are highly ambiguous because we do not actually know how many intervals of type *A* there are in any set *X*—or even, for that matter, how many *intervals* there are in any set *X*. This, too, is a problem of contextualization, and this is the point at which the p-space work of these two theorists intersects with the central concern of this dissertation—saturation. Saturation, after all, is all about context. In each of our saturation vectors, every argument *a* expresses not only whether *a* is present in set *X* and the number of *as* that are present in set *X*, but also how many *as* are both *possible* and *trivial* in any set of the same cardinality. Thus, the issue of context has been satisfied—at least in large part—even before a comparison between two sets is initiated.

In the spirit of the saturation vectors and indices introduced in chapter 2 and employed in chapters 3 and 4, we will suggest some ways in which saturation can be applied in p-space. Doing so will not only help to place the interval-class roster or PICV of a pset into a context, but will also enable us to construct relatively cardinality-neutral p-space similarity indices.

The most striking problem in applying interval saturation to p-space is the issue of maximal and minimal embeddings of elements. The problem is not determining what they are, but whether they are of any use. Since p-space is theoretically infinite, the maximum number of any interval class in any size set is cardinality-1. In a three note set, one can have only as many as two instances of *any* interval-class. In a cyclic space such as pc-space, a cycle of periodicity 3 can yield three of the same interval-class in a three-note set—{0,4,8}, for example. {0,4,8} in p-space, however, is constructed of two ic4s

and one ic8. The minimum number of any particular interval in any set of any cardinality is always zero. Because the range of psets can be infinitely wide, it is always possible to create a set that excludes any given interval.

If we limit ourselves to the range of human audition or to the 88 notes found on a piano (we might call this piano-space), then the space is no longer infinite and some generalizations regarding minimal and maximal interval content can be drawn—particularly if, as with pc-space, we know the cardinality of the set.

Given a particular pset size and a maximally-large interval that can be spanned, we can employ the idea of saturation. For example, if we are working in piano space, the largest interval is 88 and the pitches range from -39 (A0) up to 48 (C8). We can easily count the number of interval 86s that are possible in this space: three—one from -39 to 46, one from -38 to 47, and one from -37 to 48. This is somewhat helpful, but not terribly practical given that we rarely see chords—or even pieces—so large that they could approach maximal saturation of any particular interval. Given the still-enormous size of piano-space (and even the relative enormity of clarinet space or trombone space), our context is still quite vast and not terribly useful in describing relationships among most psets.

If, however, instead of thinking of the boundaries of the particular space, we individually consider the registral boundaries of a particular segment (or the registral boundaries of a particular series of segments) *vis-à-vis* the cardinality of that chord, then we have the potential to calculate interval-class (or even set-class) saturation meaningfully in p-space. For example, if our boundary interval is 15 and the cardinality of the pset is 4, there can be no

more than two instances of intervals 6 through 14 and no more than three instances of intervals 1 through 5. These maximal embeddings, like those in the pc-space interval-class saturation vector, are easily calculable using interval cycles—but we are getting a bit ahead of ourselves. We will explain these statements in greater detail as we formally define our system.

The parameters of our system are expressed using the following variables:

b = the boundary interval: the interval between the outermost notes in a pset.

c = cardinality: the number of distinct pitches in a pset.

i = interval class: an undirected interval of any size.

l = length of a particular interval cycle. We will call cycles where $l = 1$ “trivial.”

The basic groundwork for ic saturation in p-space is established through the following theorems and illustrations:

Theorem 5.4.1: The maximum cardinality (c) of any pset of boundary b is $b + 1$.

Illustration 5.4.1: Let $b = 3$. If the lowest pitch is 0 and the highest pitch is 3, there can be as many as two pitches between the boundaries. Two inner pitches plus two boundary pitches equals four possible pitches ($b + 1$) in $b = 3$ (a space of boundary interval 3).

Theorem 5.4.2: The maximum number of *any* particular interval-class i possible in a set of cardinality c , irrespective of the boundary interval, is $c - 1$.

Illustration 5.4.2: Let $c = 4$. $c - 1 = 3$. Any pset of #4 can include no more than three of any ic. If one repeatedly stacks instances of any interval i ,

the number of notes in the set will always be one greater than the number of stacked interval i s.

Theorem 5.4.3: The maximum number of any particular interval-class i possible in a set of boundary interval b , irrespective of cardinality, is $b - i + 1$.

Illustration 5.4.3: If $i = 4$ and $b = 15$, $(b - i + 1) = (15 - 4 + 1) = (11 + 1) = 12$. The maximum number of ic4 in a pset of $b = 15$ is 12. To attain 12 ic4s, one must use all sixteen possible notes within the boundary (e.g., $\{0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15\}$). Let us examine two more examples of maximum interval saturation within the same boundary interval. If $i = 8$ and $b = 15$, $(b - i + 1) = (15 - 8 + 1) = (7 + 1) = 8$. The maximum number of ic8 in a pset of $b = 15$ is 8. Attaining maximum ic8 saturation within $b = 15$ also necessitates the use of all sixteen possible notes, though in this case each pitch participates in the formation of only one ic8 (the longest 8-cycle is #2, whereas the longest 4-cycle is #4). If $i = 14$ and $b = 15$, $(b - i + 1) = (15 - 14 + 1) = (1 + 1) = 2$. The maximum number of ic14 in a pset of $b = 15$ is 2. Maximum saturation is attainable using a pset as small as four notes (e.g., $\{0,1,14,15\}$), but the maximum remains at 2 even if the cardinality of the set is larger than 4.

Theorem 5.4.4: The number of distinct i cycles in a pset will always be equal to i , assuming the boundary interval (b) is at least as large as i .

This includes the number of “trivial” one-element i -cycles, if such exist.

Illustration 5.4.4: Let us assume that we have a fully chromatic pset that has 0 as its lowest element. Given an ic i , there will be a distinct i -cycle beginning at each pitch from 0 up to $i - 1$. The cycle beginning on i will be a subset of the cycle beginning on 0. For example, let us examine the distinct 3-

cycles in a set of $b = 10$: The first cycle is (0-3-6-9), the second is (1-4-7-10), and the third is (2-5-8). The cycle that would begin on 3 (3-6-9) duplicates some of the pitches in the 3-cycle beginning on 0.

Theorem 5.4.5: The longest unbroken i cycle possible in a pset of boundary interval b is $\frac{b+1}{i}$, irrespective of cardinality.

Because it is impossible to have fractional numbers of pitches, any remainder will always be rounded *up* to the next highest integer. The maximum number of adjacent i cycles is therefore equal to the maximum cycle length minus one ($\frac{b+1}{i} - 1$).

Illustration 5.4.5: Let $i = 4$ and $b = 15$. $\frac{15+1}{4} = 4$. There is no fractional remainder in this case, so we needn't worry about rounding up. The largest 4-cycle in a space of $b = 15$ is therefore 4 (comprising 3 i 4s).

Theorem 5.4.6: The cycles of i within a set of boundary interval b will either all be the same length, or there will be two possible lengths (when 5.4.5 produces a remainder): l_1 (the smaller) and l_2 (the larger).

The first case will occur when the equation in theorem 5.4.5 produces no remainder; the second case will occur when there is a remainder. Where there are two distinct lengths, the length of the larger is one greater than the length of the smaller ($l_2 = l_1 + 1$).

Illustration 5.4.6: When $b + 1$ (the maximum cardinality of an set of boundary b —see theorem 5.4.1) is equally divisible by the cyclic interval, all cycles will be of equal length ($\frac{b+1}{i}$). When it is not equally divisible, some cycles will be longer than others, though there will only ever be two sizes for a given b and i . For example, if $b = 14$ and $i = 5$, the maximum cardinality ($b + 1$) is 15. $\frac{15}{5} = 3$. There is no remainder, therefore all the 5-cycle lengths

will be the same size (#3). The cycles of $i5$ within $b = 14$ are (0-5-10), (1-6-11), (2-7-12), (3-8-13), and (4-9-14). If the boundary is reduced to 13, however $\frac{b+1}{i} = \frac{13+1}{5} = 2\frac{4}{5}$. There is a remainder and, as one can readily see, the last cycle would be shortened to (4-9), leaving four cycles of #3 and one cycle of #2.

Theorem 5.4.7: When there are two distinct i cycle lengths in boundary b , the number of larger cycles is equal to the remainder of $\frac{b+1}{i}$.

$\frac{b+1}{i}$ is the formula first presented in theorem 5.4.5. Given that there are always i number of i cycles (theorem 5.4.4), the number of smaller cycles is equal to the difference between i and the number of larger cycles (number of l_1 i cycles + number of l_2 i cycles = i).

Illustration 5.4.7: Let us return to the example used after theorem 5.4.4 where $i = 3$ and $b = 10$. $\frac{b+1}{i} = \frac{10+1}{3} = 3\frac{2}{3}$. This means that the maximum 3-cycle length is 4 ($3\frac{2}{3}$ rounded up) and the minimum cycle length is 3 ($4 - 1$). There are two (the remainder of $3\frac{2}{3}$) four-note 3-cycles and one (the difference between i and the remainder) three-note 3-cycles which, together, produce the maximum number of ic3s in any $b = 10$ pset. Each four-note cycle produces three ic i s and the three-note cycle produces two ic i s. $3 + 3 + 2 = 8$, the maximum number of ic3s in $b = 10$ (this number is also derivable through theorem 5.4.3: $b - i + 1 = 10 - 3 + 1 = 8$).

We now have formulas to derive the number and length of the ic cycles given both a particular interval type and a boundary interval. With this information, we can easily calculate the maximum and minimum numbers of any ic i within any boundary b given cardinality c . As an example of how this information would be weaned from the above formulas, let us re-examine the

3-cycles within $b = 10$. They are, as we noted, (0-3-6-9), (1-4-7-10), and (2-5-8). If we wanted to derive the maximum number of ic3s within $b = 10$, given a set of a particular cardinality, we could begin by taking pitches from the 3-cycle beginning on 0 (we'll call this "cycle 0"), then, when that cycle is exhausted, progress to cycle 1, and then on to cycle 2. A set of cardinal 3 within $b = 10$, for example, could have as many as 2 ic3s (0-3-6). A set of cardinal 4, would have one more ic3 (0-3-6-9). A set of cardinal 5, would, by contrast, not have one more ic3 because a new cycle is begun and, as with ic cycles in pc-space, intervals of type i occur only within i cycles. A set of cardinal 6, would allow an ic3 relation to the new element added in our #5 set: the maximum number of ic3 in a #6 pset where $b = 15$ would therefore be 4, etc.

To derive the minimum numbers of ic3 possible, we also use the cycles. But, instead of progressing linearly through each cycle before moving onto the next, we progress from the first pitch in cycle 0 to the first pitch in cycle 1, etc. When the first pitch from each i cycle has been used, we move back up to cycle 0 and take the next element that is cyclically *non-adjacent* to the first. We continue likewise with each cycle. Figure 5.5 provides the maximal and minimal numbers of ic3 in $b = 10$ for each cardinality 1 through 11 and shows their 3-cyclic elements.

As we saw in pc-space, the progression of minimal and maximal possible interval i embeddings is not co-linear with the size of the set. When a one i -cycle has been completed, no new ic i can be attained through the addition of only a single new pitch. When calculating minimum i content, we select notes that are from different cycles or which are non-adjacent within a cycle. When

every other note in each cycle has been used, at least one i must be formed with the next added pitch. If a cycle has an even number of elements and we have begun our process of generating sets minimally saturated with i by selecting the first note of each cycle, we progress by adding the last pitch of each cycle to our set. This insures that only one interval i is added with each additional note. If we choose to add an interior pitch in the cycle, two interval i s would be formed (one with the note which precedes the new pitch and one with the note that follows it). This process is demonstrated on the right side of figure 5.5.

The pset $\{0,2,4,5,8,9,10\}$ includes two ic3s (pitch dyads $\{2, 5\}$ and $\{5, 8\}$). From the information in figure 5.5, we can assert that the ic3 content is three short of the maximum and one more than the minimum possible for any #7 pset of $b = 15$. We can easily generate all minimal and maximal values for all i, b , and c using the formulas for cyclic length and number provided above. These minimal and maximal values could, as before, be used in a vector (perhaps a revision to the PICV suggested earlier in this chapter) and that, in turn, could be used as fodder for a similarity index.

Such a p-space similarity index would, we believe, be both flexible and relatively cardinality neutral,²¹⁸ without requiring any equation of p- and pc-space. While we have suggested how such a saturation vector and associated similarity index *might* be built, the actual formation of such a measure along with a survey of its values falls outside the scope of this study and will be left for future research.²¹⁹

²¹⁸Though it would not be pitch width (boundary interval) neutral.

²¹⁹Alternately, one could use other systems of determining some intervallic profile (involving elemental saturation or not) to form a p-space similarity index. We do not mean to infer that our preliminary work is uniquely capable of forming a cardinality-neutral measure of p-space set similarity.

5.5 Saturation in other dimensions

Just as the idea of saturation can be utilized as a means of contextualizing interval- or subset-class content in both pc- and p-spaces, it could theoretically be extended to other dimensions as well. In contour space (c-space), we might examine the saturation of ascending compared to descending intervals in a contour segment (cseg) or perhaps in a series of contour sub-segments (csubseg).²²⁰ Rhythmic sets in a piece could be compared with what is possible given a particular meter or group of rhythmic values. Alternatively, a context-sensitive version of rhythmic saturation could be created by examining what rhythms are employed in a piece and using those data as a basis for comparing the relative saturation of each rhythmic element in the piece. In a timbral space, we might examine the orchestration of passages relative to the orchestral forces available. Instruments or voices might be considered as separate variables or as members of orchestral families.

Potentially, saturation-based similarity indices could be defined for any of these spaces and, if one desires, indices that measure multidimensional relatedness could be created by combining such measures together into a composite. Such a composite index could either average the values (assuming that they are commensurable) from the sub-indices, or those sub-indices could be weighted according to one's preference. There is obviously much work yet to be done in constructing similarity indices (and not only saturation-based indices) which measure musical features other than pitch.²²¹ We

²²⁰Hermann 1994 and 1995 describes contour space as just *one* possible “preintervallic” space among many. (See Hermann 1995, paragraph 18.) Saturation is potentially useful in any such space.

²²¹Hermann 1993 offers perhaps the best effort at examining different spaces and creating a flexible resemblance measure to deal with them.

believe that examining the saturation of various types of elements can lead to relatively cardinality-neutral comparisons of many musical parameters.

5.6 Criticisms of similarity indices and pitch-class-based analysis

Until recently, most sophisticated publications and papers involving the analysis of atonal music have focused implicitly or explicitly on the realm of pitch-class space. That is not to say that register, orchestration, dynamics, rhythm and a host of other concerns were ignored (occasionally they were), rather that those parameters were not considered formally and methodically. In the past, such issues have always informed our decisions regarding the materials that we choose to analyze (i.e., our segmentation), they have not, by in large, provided a means of examination.

Many of the criticisms raised about the use of similarity measures in analysis are, explicitly or implicitly, also criticisms of pitch-class as the sole avenue of musical study. Perhaps the most recent and eloquent such criticism was written by Thomas Demske in the electronic journal *Music Theory Online* (*MTO*). Many of his criticisms were answered thoughtfully by Richard Hermann, also in *MTO*.²²² We will summarize some of the points of Demske 1995 and Hermann 1995 here and add a few more words to the ongoing debate.

Two issues raised by Demske warrant particular mention here. Demske is bothered by the lack of clarity regarding the meaning of values returned by similarity indices. He also believes that similarity measures do not adequately

²²²Demske 1995 and Hermann 1995.

take perception into account. Both points have considerable merit, and we will address them in turn.

Regarding the first point: Demske points out that in the realm of similarity, all sets are similar but to different degrees.²²³ What seems like a close relationship in one context might seem rather distant in another. How, then, can one use the values produced in a consistent and logical manner when it is not clear what they represent? We do not deny any of these claims; we also do not believe that they represent flaws. In addition to saying that all sets are similar, but to varying degrees, we might *also* say that all sets are dissimilar, but to varying degrees. This is one of the stronger points of intersection between the relatively new area of fuzzy logic and similarity indices. In both cases, we recognize that every evaluation is neither true nor false but true *and* false to varying degrees. We can have points of maximal and minimal similarity—represented by 1 and 0 (or 0 and 1) in most measures—but even these represent a point of view, a notion of equivalence, not an absolute binary opposition.

Regarding the problem of contextualizing values, we agree that this is something that must be solved in order for similarity indices to be used coherently in analysis. Our solution—not necessarily a perfect one—was to use the other values derived from comparisons within the piece and examine not one cutoff point for a similar versus dissimilar judgment, but many different ones, comparing them as we progressed. Naturally, we might have chosen pieces where all the sets are found to return very small similarity values

²²³He states: “Barring *ad hoc* provisions in the interpretation, all sets are similar with respect to each other. Only the degree of similarity varies. . . . An *a priori* cutoff point *X* would help, where by only [two sets with at least a particular similarity value] are considered “similar” in an absolute sense. (T_nI “equivalence” supports analogous appeals to the absolute.) But could such a point be meaningfully determined here?” (Demske 1995, paragraph 8.)

(indicating high degrees of similarity), or we might have used a piece where all the sets seemed quite dissimilar compared with what is possible. In such cases, our *ad hoc* system of drawing boundaries might have seemed unreasonably acute, rather like grading on an absolute curve in a class where everyone scored nearly perfectly.

It is for this reason that we always began our evaluation of the data by citing the value averages and high and low points not only for the group of sets at hand, but also for all sets in the same cardinality range, and for the sets at hand compared to all other sets (in or out of the piece). In our analyses, we didn't actually use these numbers for evaluation. We didn't need to. Had the statistical summary of the sets at hand been radically different from the more general value group statistics, we would have noticed and adjusted our methodology.

Demske's second point regarding the perceptual salience of measures of resemblance is both complex and broad. He begins by stating, "evaluating similarity clusters with perceptual criteria is complicated because we can attend to the same music in many different ways."²²⁴ Indeed, reliance on perception necessarily introduces all or most other musical parameters into the analytical mix. The way that a listener perceives pitch or pitch class relationships is surely colored by the musical texture: the way the pcs are spaced, the instrumentation, the durations, spatial and acoustical concerns, etc.

Hermann answered Demske's charge that similarity indices do not necessarily register the varieties of human perception by making two very important points. He said, "whether or not these relations actually do model

²²⁴Demske 1995, paragraph 14.

perception in a specific passage depends upon whether the passage exhibits its materials in such a way that encourages its perception with that tool by a reasonably experienced listener.”²²⁵ Furthermore, Hermann reasoned that “all analysts must make decisions about what strikes them as the most salient or important features of the piece and then select the appropriate ‘formal evaluation procedure’ designed to address those features. Certainly to run all possible theoretical tools at the piece in an analysis would quickly overwhelm the analyst with a plethora of analytical observations upon the data: most observations are likely to be true, but many may be of little significance, aural or otherwise. Thus, the selection of theoretical tools for analysis acts like a set of filters upon the analyst and the piece at hand.”²²⁶

We wholeheartedly agree with Hermann’s defense. If an analyst finds that other parameters in a musical texture suggest readings that conflict with his or her pcset analysis, then those other parameters should probably take precedence. Perhaps our admission that some unnamed “other parameters” can be used as justification for putting aside or overriding pc-based analytical data might seem to undermine much of what has been proposed herein. We do not think that it does, however. It was never our goal to advocate any particular methodology for atonal analysis—only to provide some new and, we hope, useful tools and to suggest ways in which they could be applied. We do not believe that an analysis must stem from any particular theoretical system, nor that an analysis must be limited to any single system.

Hermann notes that many of Demske’s criticisms seem to apply not to the idea of similarity measures in the abstract but to the fact that similarity had, in

²²⁵Hermann 1995, footnote 2.

²²⁶Hermann 1995, paragraph 12.

the examples Demske examined, always been measured in the realm of pc-space. Demske's claims of perceptual invalidity might then be taken as claims that pcset analysis itself has little perceptual validity—or at least that the salience of pcsets is largely determined by their realizations. Such a claim was also leveled in Morris 1995. In Morris's example 15a, he demonstrates that different surface realizations of 8-28 (the octatonic collection) can sound considerably different. His example 15a has been reproduced as our figure 5.6.

We do not question Morris's assertion that pcset-based measures of similarity and equivalence are not representative of the vivid distinctions among these sets. They are distinguished not only by the width of their spacing, but also by the types of intervals that fall between their adjacent (and also non-adjacent) pitches. The first set, for example, features one close spacing of SC 4-28 [0369] (diminished seventh chord and complement of the octatonic collection) in the treble clef part and another in the bass clef part. Most of the registrally-adjacent intervals are interval 3s (minor 3rds)—one of the two interval classes with which 8-28 is maximally saturated (ic6 being the other). By contrast, the third set in figure 5.6 is built of alternating interval 5s and 4s (perfect fourths and major thirds). While these intervals are obviously included in 8-28, that set is minimally saturated with both of them. That particular realization of 8-28 might be seen or heard as bringing out relatively “unlikely” intervallic formations.²²⁷

In Morris's figure 15b, he compares each of his four distinct pitch realizations of 8-28 to a realization of a very different pcset, spaced in roughly

²²⁷Probability aside, we acknowledge that composers frequently realize structures in ways that bring out their least common properties.

the same manner and utilizing as many common tones as possible. These realizations clearly demonstrate that distinct pc SCs can be realized in similar p-space settings. The chord that he compares to his third realization of 8-28 is a seven-note realization of 5-35 [02479], a five-pc quartal (or quintal) harmony.²²⁸ The lowest three pitches in both chords are precisely the same, suggesting parallels in both their spacing and number of common tones.

Morris's demonstration provides a very compelling reason for examining not only pitch class, but also pitch. We believe that a parallel argument could be made that looking only at p-space—or even at a composite of p- and pc-spaces—ignores other salient features. In figure 5.7 below, we provide four orchestrations of Morris's third realization of 8-28. We believe that these four orchestrations would each be distinct to a listener and that furthermore the different timbres and sound envelopes²²⁹ might be as distinct as Morris's four unique psets in figure 5.6. In particular, the final realization, which has the organ and marimba playing every other note in the chord, might lead one to “rediscover” the two embedded members of 4-28 in 8-28 since the marimba will die away, leaving the organ still sounding. This orchestration might even bring out the two diminished seventh chords even more than a single timbre playing the first of Morris's realizations that separates the two 4-28 solely by register.²³⁰

Our point here is not that we should examine all parameters before formulating any analytic opinion. Such a restriction would be unwieldy to

²²⁸Ic5 is therefore maximally saturated.

²²⁹Roughly speaking, the rate of attack, decay, sustain, and release of each instrument.

²³⁰Our claim that timbral data can further inform analytical statements does not refute Morris's claim that p-space sets can be analytically useful. P-space sets, after all, can distinguish pc-space sets, but not vice-versa. Similarly, orchestrated p-space sets can distinguish unorchestrated p-space sets, but not vice versa.

say the least. We are also not suggesting that Morris, Hermann, Demske, and others are misguided in seeking analytical tools that account for more than just pitch class. Rather, we are simply demonstrating that almost all current analytical tools (and analyses) are open to the sort of criticism that Morris levels against pc-space readings. Given the unlikelihood that some comprehensive system of analysis that examines all relevant parameters can or will (or even should) be introduced, we believe that pcset analysis holds up quite well as an efficient place to *begin*—if not *end*—many atonal analyses.²³¹

5.7 Conclusions

The central purpose of this study has been to introduce a fundamentally new way of considering the data that we use in formulating atonal analyses. In particular, we have concentrated on methods of examining pcsets based upon their intervallic and cyclic content and also by the number and types of subsets that they embed. Rather than simply tallying the numbers of each subset- or interval-class found within a particular pcset, we have examined those numbers relative to what is both minimally and maximally possible in the universe of pcsets of the same cardinality. These contextually-enriched values represent the degree to which object x is saturated in pcset X .

These saturation values were then used as fodder for a number of new similarity indices—algorithms that weigh the respective saturation values associated with two different pcsets and return a number between zero and one that indicates the degree to which the two pcsets resemble each other. We acknowledge that our own similarity indices are not the only algorithms

²³¹We do, however, concede that there are certainly many pieces where pc-set connections are not salient and therefore inappropriate.

by which saturation values can be measured. Systems based upon standard deviation, such as those designed by Isaacson and Hermann,²³² present but two of many other possibilities. Saturation-based data, by the nature of their design, make *a priori* adjustments for cardinality and we believe that they are helpful in relating sets (of equal or unequal size) no matter what technique is used in drawing the comparisons.

To demonstrate some of the similarities and differences among our own measures as well as among some of the most widely discussed measures posited by other theorists, in chapter 3 we examined a short piece by Stravinsky through the lenses of eight different indices. While the manner in which we interpreted the various data to form analytical narratives certainly admits a degree of subjectivity, we believe that some unique features of these indices were revealed through these readings. In chapter 4, one single measure—our cyclic saturation similarity measure (CSATSIM)—was used as the primary analytical tool in examining a slightly longer piece by Messiaen. Here, the techniques of set comparison developed in chapter 3 were put to use in demonstrating several distinctive families of set types used in that piece. Our interpretation of the CSATSIM data revealed that Messiaen employed similar types of sets at similar junctures in the musical fabric.

This final chapter addressed some criticisms that have been leveled against both pcset analysis in general and similarity relations in particular. We also suggested ways in which the idea of elemental saturation could be extended into other musical domains and provided foundations for a pitch-space-based similarity index that would weigh the degree to which psets are

²³²Isaacson 1990 and Hermann 1993.

similarly spaced. This saturation-based system, we believe, would allow for the design of the first similarity index that facilitates relatively cardinality-neutral comparisons of psets.

Much atonal analysis has relied—directly or indirectly—on data from functions and measures such as Lewin’s EMB and COV and Forte’s K and Kh. These tools all answer the question “is element x embedded in or does it embed some other element y ?” with either a yes or no answer (K and Kh) or a number that tells one how many x are embedded in or embed y (EMB and COV). While this is undoubtedly useful information, we believe that it is enriched by the frames of reference that our saturation-based tools provide.

Finally, we believe it important to reiterate our disclaimer that while we believe these saturation-based tools to be analytically useful in many situations, they are by no means useful in every situation. There is, after all, no single tool makes all other tools obsolete. It is up to each theorist and analyst to decide which are appropriate in any given circumstance.

Math and logic symbols

=	equal to
	does not equal
	equivalent to
<	less than
>	greater than
	less than or equal to
	greater than or equal to
	bidirectional mapping
	directional mapping
	directional mapping
	and
	or
	not (negation)
	if ... then
	if and only if ... then
	therefore
	is/are an element of
	is/are not an element of
	is a subset of
	is a proper subset of
	set intersection
	set union
U	the universal set
∅	null set
X	set X; uppercase letters (except U and V) are variables for sets of elements
\bar{X}	the complement of set X
#X	the cardinality of (number of elements in) set X
{ ... }	an unordered set of elements
[...]	an unordered set of elements in some prime form
< ... >	either a vector or an ordered set of elements
\sum	the summation sign
$\sum_{i=1}^n$	The direct sum of the elements from (the counter/subscript) $i = 1$ to n .
$\sum_i I$	The direct sum of the elements of a group I. (The counter/subscript i is a member of the group I.) This is called an indexed family

Bibliography

- Babbitt, Milton.
 1960. "Twelve-Tone Invariants as Compositional Determinants." *Musical Quarterly* 46: 245-59.
- Baumslag, Benjamin and Bruce Chandler.
 1968. *Theory and Problems of Group Theory*. Schaum's Outline Series. New York: McGraw-Hill.
- Bennighof, James.
 1987. "Set-Class Aggregate Structuring, Graph Theory, and some Compositional Strategies." *Journal of Music Theory* 31/1: 51-98.
- Block, Steven and Jack Douthett.
 1994. "Vector Products and Intervallic Weighting." *Journal of Music Theory* 38/1: 21-41.
- Carter, Elliott.
 1960. "Shop Talk by an American Composer." *Musical Quarterly* 46: 189-201
- Castrén, Marcus.
 1994. *RECREL: A Similarity Measure for Set-Classes*. Studia Musica 4. Helsinki: Sibelius Academy.
- Demske, Thomas.
 1995. "Relating Sets: On Considering a Computational Model of Similarity Analysis." *Music Theory Online* 1/2.
- Clough, John.
 1965. "Pitch-Set Equivalence and Inclusion (A Comment on Forte's Theory of Set-Complexes)." *Journal of Music Theory* 9: 163-171.
 1983. "Use of the Exclusion Relation to Profile Pitch-Class Sets." *Journal of Music Theory* 27: 181-202
- Ericksson, Tore.
 1986. "The IC Max Point Structure, MM Vectors and Regions." *Journal of Music Theory* 30/1: 95-111.

Forte, Allen.

- 1964. "A Theory of Set Complexes for Music." *Journal of Music Theory* 8: 136-83.
- 1965. "The Domain and Relations of Set-Complex Theory." *Journal of Music Theory* 9: 173-180.
- 1973. *The Structure of Atonal Music*. New Haven, Yale University Press.
- 1988. "Pitch-Class Set Genera and the Origin of Modern Harmonic Species." *Journal of Music Theory* 32: 187-334.

Hanson, Howard.

- 1960. *Harmonic Materials of Modern Music: Resources of the Tempered Scale*. New York: Appleton-Century-Crofts.

Headlam, Dave.

- 1996. *The Music of Alban Berg*. New Haven: Yale University Press.

Hermann, Richard.

- 1993. "A General Measurement for Similarity Relations: a Heuristic for Constructing or Evaluating Aspects of Possible Musical Grammars." Ph.D. diss., University of Rochester.
- 1994. "On 'Preintervallic' Spaces and on Their Interactions with Some Intervallic Spaces." Delivered at the Society for Music Theory annual conference.
- 1995. "Towards a New Analytic Method for Post-Tonal Music: A Response to Thomas Demske." *Music Theory Online* 1/3.

Hindemith, Paul.

- 1937. *The Craft of Musical Composition, Book I (Theoretical Part)*. Trans. by Arthur Mendel. New York: Associated Music Publishers.

Hoover, Mark.

- 1984. "Set Constellations." *Perspectives of New Music* 23/1: 164-179.

Howe, Hubert S.

- 1965. "Some Combinational Properties of Pitch Structures." *Perspectives of New Music* 4: 45-61.

Isaacson, Eric J.

1990. "Similarity of Interval-Class Content between Pitch-Class Sets: The IcVSIM Relation." *Journal of Music Theory* 34/1: 1-28.
1992. "Similarity of Interval-Class Content Between Pitch-Class Sets: The IcVSIM Relations and its Application." Ph.D. diss., Indiana University.
1996. "Issues in the Study of Similarity in Atonal Music." *Music Theory Online* 2/7.

Kaplan, Richard A.

1990. "Transpositionally Invariant Subsets: A New Set-subcomplex." *Integral* 4: 37-65.

Kosko, Bart.

1993. *Fuzzy Thinking: The New Science of Fuzzy Logic*. New York: Hyperion.

Lewin, David.

1959. "Re: Intervallic Relations between Two Collections of Notes." *Journal of Music Theory* 3/2, 298-301.
1960. "Re: The Intervallic Content of a Collection of Notes, Intervallic Relations between a Collection of Notes and its Complement: An Application to Schoenberg's Hexachordal Pieces." *Journal of Music Theory* 4/1, 98-101.
1977. "Forte's Interval Vector, my Interval Function, and Regener's Common-Note Function." *Journal of Music Theory* 21/2, 194-237.
- 1979-80a. "Some New Constructs involving Abstract PCsets, and Probabilistic Applications." *Perspectives of New Music* 18, 433-444.
- 1979-80b. "A Response to a Response: On PCset Relatedness." *Perspectives of New Music* 18: 498-502.
1987. *Generalized Musical Intervals and Transformations*. New Haven: Yale University Press.
1990. "Klumpenhouwer Networks and Some Isographies that Involve Them." *Music Theory Spectrum* 12/1: 83-120.
1993. *Musical Form and Transformation: 4 Analytic Essays*. New Haven: Yale University Press.
1994. "A Tutorial on Klumpenhouwer Networks, using the Chorale in Schoenberg's Op. 11, #2." *Journal of Music Theory* 38/1: 79-101.

Lipschutz, Seymour.

- 1976, *Theory and Problems of Discrete Mathematics*. Schaum's Outline Series. New York: McGraw-Hill.

- Lord, Charles.
1981. "Intervallic Similarity Relations in Atonal Set Analysis." *Journal of Music Theory* 25: 91-111.
- Martino, Donald.
1961. "The Source Set and Its Aggregate Formations." *Journal of Music Theory* 5: 224-73.
- Marvin, Elizabeth W. and Paul A. Laprade.
1987. "Relating Musical Contours: Extensions of a Theory for Contour." *Journal of Music Theory* 31/2: 225-266.
- Mead, Andrew W.
1984. "Pedagogically Speaking: A Practical Method for Dealing with Unordered Pitch-Class Collections." *In Theory Only* 7/5-6: 54-66.
1988. "Some Implication of the Pitch-Class/Order-Number Isomorphism Inherent in the Twelve-Tone System: Part 1." *Perspectives of New Music* 26/2: 96-163.
- Messiaen, Olivier.
1944. *Technique de mon Langage Musical*. Paris: Alphonse Leduc.
- Morris, Robert D.
1979-80. "A Similarity Index for Pitch-Class Sets." *Perspectives of New Music* 18: 445-460.
1982. "Set Groups, Complementation, and Mappings among Pitch-Class Sets." *Journal of Music Theory* 26/1: 101-144.
1987. *Composition with Pitch Classes: A Theory of Compositional Design*. New Haven: Yale University Press.
1990. "Pitch-Class Complementation and its Generalizations." *Journal of Music Theory* 34/2: 175-245.
1991. *Class Notes for Atonal Music Theory*. Hannover, NH: Frog Peak Music.
1995. "Equivalence and Similarity in Pitch and their Interaction with Pcset Theory." *Journal of Music Theory* 39/2: 207-243.
- O'Donnell, Shaugn
1997. "Transformational Voice Leading in Atonal Music." Ph.D. diss., City University of New York.
- Perle, George.
1954. "The Possible Chords in Twelve-tone Music." *Score* 9: 54-58.
- Perone, James.
1988. "The Projection Quality Index." *Interface* 17: 79-102.

- Quinn, Ian.
1997. "Evolution A Forte-iori: On Similarity, Relations, Similarity Relations, and the Taxonomy of the Harmonic Menagerie," presented at the Music Theory Society of New York State 1997 annual meeting, Rochester, NY.
- Rahn, John.
1979-80. "Relating Sets." *Perspectives of New Music* 18: 483-498.
1980. *Basic Atonal Theory*. New York: Schirmer Books.
1989. "Toward a Theory for Chord Progression." *In Theory Only* 11/1: 1-10.
- Regener, Eric.
1974. "On Allen Forte's Theory of Chords." *Perspectives of New Music* 13: 191-212.
- Robison, Brian.
1994. "Modifying Interval-Class Vectors of Large Collections to Reflect Registral Proximity Among Pitches." *Music Theory Online* 0/10.
- Sherlaw Johnson, Robert.
1989. *Messiaen*. Berkeley and Los Angeles: University of California Press.
- Starr, Daniel.
1978. "Sets, Invariance and Partitions." *Journal of Music Theory* 22/1: 1-42.
- Teitelbaum, Richard.
1965. "Intervalic Relations in Atonal Music." *Journal of Music Theory* 9: 72-127.
- Tenney, James.
1988. *Meta ≠ Hodos and META Meta ≠ Hodos*. 2d ed. Oakland, CA: Frog Peak Music.
- Zimmermann, Hans J.
1985. *Fuzzy Set Theory — and Its Applications*. Boston: Kluwer-Nijhoff Publishing.