M100C
Notes on Pitch-Class Set Theory - 1

INTRODUCTION. Set Theory is a branch of mathematics. So why do we musicians bother with it? Quite simply, numerical abstractions allow us to model many aspects of posttonal music in efficient and informative ways. This doesn't make the music itself "mathematical," just the model. Similarly, while there's nothing intrinsically mathematical about the weather, we use a simple mathematical model to represent today's temperature — and meteorologists use more complex models to predict tomorrow's.

"Posttonal" refers to a variety of musical idioms that have emerged in this century and are not conventionally tonal. We prefer this term to another — "atonal" — because it reminds us that tonal centers may still operate in less conventional ways in some of the music we will encounter. However, we will normally focus on features other than tonality in this music, with the goal of learning how to use new tools to make sense of these less familiar features.

Just as we can use common tonal theory to account for certain aspects of posttonal music, we could approach conventionally tonal music from the perspective of pitch-class set theory. Indeed, tonal theory *does* make use of mathematical models. (And if you find some new concept in this course to be particularly challenging, try to think back to when you first came to grips with the idea in tonal theory that a third plus a third equals a fifth!)

PITCH. If we want to adopt from mathematics a model suitable for representing pitch, we have to know in advance what attributes pitch possesses in the music we intend to study. For instance, pitch is capable of varying continuously between a low and a high extreme (sing a low note and gradually slide the pitch upwards to a higher note — see?), so we might decide to model it with the *real numbers:* a space that includes numbers like 0, 4, $-3\frac{1}{2}$, 99.95, π, and so on.

But most of the music we'll encounter is conceived primarily in terms of the discrete steps of the chromatic scale. And the model we normally adopt for this structure is the *integers* (...–3,–2,–1,0,1,2,3,...). This collection of numbers shares several attributes with the chromatic gamut of pitches. Both structures are *ordered:* the numbers from most negative to most positive and the pitches from lowest to highest. And both structures are *equally spaced:* the "distance" between adjacent integers is constant, as is the "distance" between chromatically adjacent pitches.

We can take "distance" out of its quotation marks, as this concept too is coherent for both structures: to measure the distance going from one integer or pitch to another, we count the number of steps along the way. Of course, in the case of integers we're used to calculating this value by subtracting, and if we associate each pitch with an integer, then we can use the same calculation to compute the distance between two pitches. That's a simple — perhaps trivial — example of the utility of our model. It also prompts an important cautionary remark: when we do what has just been described, we are *not* subtracting one pitch from another (whatever that might mean). Rather, we're taking advantage of the fact that, in the abstract world of our model, subtraction computes distance.

Parenthetically, it's worth bearing in mind that the pitches "in between" the steps of the chromatic scale aren't irrevocably lost given our choice of an integer model. Mathematicians have lots of clever ways to get from the integers back to the real numbers, and we could build these into our model if we wanted to.

How, precisely, do we associate pitches with integers? One way is to assign 0 to C_4; then increasingly positive integers correspond to increasingly higher pitches above C_4, and increasingly negative integers correspond to increasingly lower pitches below the same reference. Of course, making C_4 our reference is an arbitrary choice, and if some other pitch is a tonal center or a center of symmetry in a particular context, we might want to choose it to be 0. Then the integer value associated with each pitch will be an instant measure of the pitch's distance from the center. If that's something we care about, then we'll build our model accordingly.

Another possibility is to assign 0 or 1 to the lowest pitch that interests us; then we can work with exclusively positive numbers. But it's still good to have the complete set of integers available, rather than limiting ourselves *a priori* to the counting numbers (1,2,3,...): we don't want our model to break down if we encounter a lower pitch than we expected.

Finally, note that *enharmonically equivalent* pitches are associated with the same integer. Thus if C_4 is 0, then so are $B\sharp_3$ and $D\flat\flat_4$.

PITCH SET. An unordered collection of pitches. We write a pitch set by listing its elements within curly brackets {}. Because pitch sets are unordered, the following notations all represent the *same pitch set:*

{–13, 0, 2} {–13, 2, 0} {0, –13, 2} {0, 2, –13} {2, –13, 0} {2, 0, –13}

Deciding what pitches belong together in the same set is an important step in any set-theoretic analysis. This process is called segmentation. We might segment a passage so that each melodic group, or each chord, formed a separate set.

PITCH INTERVAL. There are two ways we can think of pitch intervals: as a *melodic* displacement that carries us from one pitch to another, and as a *harmonic* combination of two pitches. A melodic interval is *directed:* going up from the lower pitch and going down from the upper pitch are two distinct melodic shapes. We can express this distinction by using a positive integer for an ascending interval and a negative integer for a descending one. As mentioned above, we can compute these values by subtraction.

The ***directed pitch interval*** from x to y is $(y - x)$.

A harmonic interval is *undirected* — that is, we do not hear it as ascending or descending. The only difference between the ascending interval from x to y and the descending interval from y to x is its positive or negative sign. So to compute an undirected interval, we compute the directed interval i in either direction and then discard the sign. This is called taking the *absolute value* of the directed interval; the result, $|i|$, is always a positive integer.

> The ***undirected pitch interval*** between x and y is $|y - x| = |x - y|$.

OCTAVE EQUIVALENCE. An important aspect of pitch *not* captured by our integer model is the sense in which the directed interval 12 (i.e. an octave) takes us from any given pitch to another that sounds like it. In tonal theory, we give these pitches the same letter names (so an octave above C_4 is C_5). We say that two such pitches are *octave equivalent,* and of course the same relationship holds between pitches separated by multiple octaves.

PITCH CLASS. We can sort the elements of the chromatic scale into twelve classes (sets, actually) by including in each class all the pitches that are octave equivalent to one another. One such class will consist of 0 and its octave-equivalent relatives: {...,–24,–12,0,12,24,...}. Another class will consist of 1 and its octave-equivalent relatives: {...,–23,–11,1,13,25,...}. And so on.

We call the twelve sets that result *pitch classes,* and we can construct a new model to represent them. What makes our pitch class model different from our pitch model is that pitch classes are *unordered:* since each class includes both high and low pitches, it doesn't make sense to say that one is higher than another. (Which is higher, C or C♯? It depends which C and which C♯!) We will keep the integers 0–11, but each of these values now represents a whole set of integers: the pitch class to which it belongs. Thus we have no need for integers larger than 11; 12, for instance, would represent the same class as 0. Likewise, we avoid integers smaller (more negative) than 0; for example, we use 11 in place of –1. The choice of C = 0 is common, but other assignments are sometimes useful.

It is often helpful to imagine the pitch classes arranged like a clockface, with 0 in place of 12 (note that I've written "t" for ten and "e" for eleven, so each pitch class consists of a single digit):

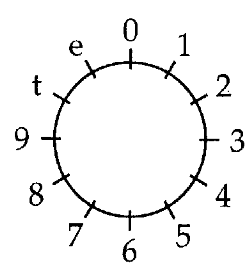

The important thing to remember is that computations done with pitch classes will normally need to "wrap around" on the clockface. For instance, to compute e + 4, we start at e and advance four positions clockwise: e + 4 = 3. Fortunately, we can safely perform arithmetic in the usual way and then reduce the result "mod 12" to produce a pitch-class result. Reduction mod 12 involves adding or subtracting 12 from a value as often as necessary to bring it within the range of 0 to 11:

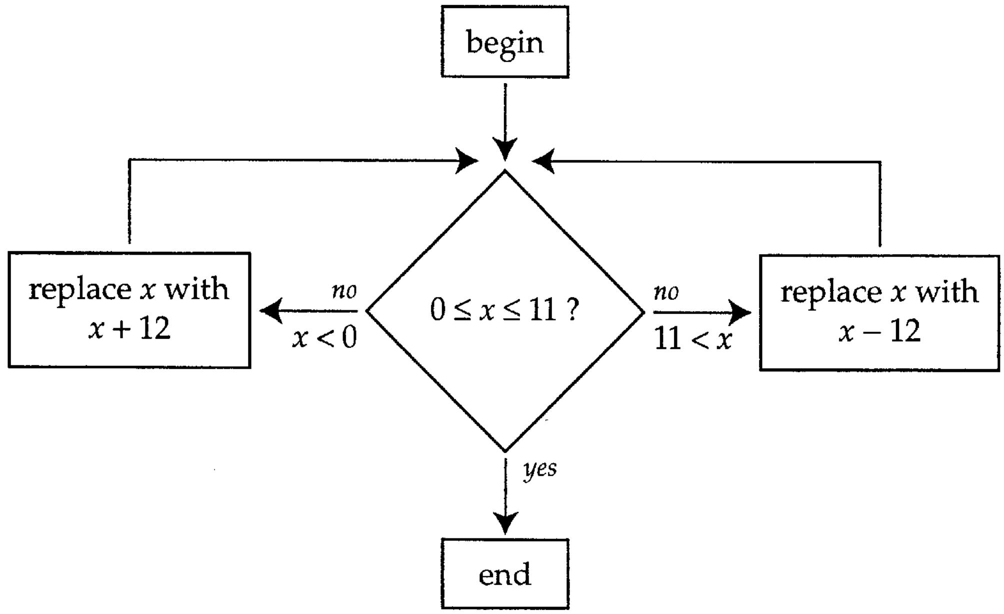

PITCH CLASS SET. A set of pitch classes, which means a set of sets of pitches. Consider a C major triad: we encounter this familiar structure in all kinds of different chordal inversions, spacings, and doublings; what they all have in common is that they consist of combinations of C's, E's, and G's. The pitch class set {C,E,G} (equivalent with C = 0 to {0,4,7}) represents a way to think of all of these chords as equivalent: it shows what they all have in common. However, we are not asserting that they are *identical!* A root position C major triad and a second inversion one have a family resemblance, but important differences in their sounds allow them to function differently — for instance, only the former sound is appropriate at the end of a conventional full cadence (in the key of C).

What is true of C major triads is true of pitch class sets in general. Each pitch class set stands for a large number of pitch sets that have a family resemblance to one another without sounding identical.

Parenthetically, how many distinct (and non-empty) pitch class sets are there? Every distinct subset of {0,1,2,3,4,5,6,7,8,9,t,e} is a distinct pitch class set, and it can be shown that the number of non-empty subsets of an n-element set is $2^n - 1$; here, $n = 12$ and there are therefore 4,095 distinct pitch class sets.[1]

PITCH CLASS INTERVAL. As was the case for pitch intervals, we can distinguish two kinds of pitch class intervals: a melodic and a harmonic case. We'll have to think carefully about how each case works. Melodically, it no longer makes sense to think of ascending versus descending shapes, because pitch classes are unordered. But we can think of the directed pitch class interval

[1] Need convincing? Consider that another way to write any pitch class set is to use a twelve-position array, and to store a "1" in the nth position if the pitch class $n - 1$ belongs to the set, or a "0" otherwise. Then each of the strings from 00000000000 to 111111111111 represents a distinct pitch class set. But these strings are also the numbers 0 to 4,095 in binary notation. That's 4,096 different strings, or 4,095 if we discard 000000000000 (which represents the empty pitch class set). There. Now you're either convinced or confused.

from pitch class x to pitch class y as a class of directed pitch intervals, where each pitch interval goes from some pitch in class x to some pitch in class y. And we can represent this class of intervals with the member that falls in the range of 0 to 11.

> The (directed) ***pitch class interval*** from pitch class x to pitch class y is $(y - x)$ mod 12.

Harmonically, these twelve pitch class intervals collapse into the seven values 0 to 6 as follows:

> The undirected pitch class interval or ***interval class*** between pitch class x and pitch class y is the smaller of the following two values:
>
> $$(y - x) \bmod 12$$
> $$(x - y) \bmod 12.$$

It is helpful to remember that a directed pitch class interval i and its intervallic "inversion" $(12 - i)$ are members of the same interval class.

M100C
Notes on Pitch-Class Set Theory - 2

TRANSPOSITION. Transposition is a familiar kind of *transformation:* it turns one object into another, in this case by "moving" it a specified distance. The "objects" can be pitches or pitch classes, or sets of either of these (or they can be other related structures). If the original object is X, and the distance of the transposition is *n*, and the result of transposition is Y, then we write $T_n(X) = Y$.

The notion of distance will depend on the nature of the object being transposed. If we're transposing a pitch, then we can move it either up or down, a distinction captured by the notion of *directed pitch interval.* If the direction of transposition is up, then the distance of transposition will be a positive value; if down then negative. Thus we define pitch transposition as follows:

If x is a pitch, then $T_n(x) = x + n$. (See illustration below.)

If we're transposing a pitch class, we must remember that these values wrap around. We therefore replace the notion of up versus down with the notion of clockwise versus counterclockwise, a distinction captured by the notion of *directed pitch class interval.* With this in mind, we define pitch class transposition as follows:

If x is a pitch class, then $T_n(x) = (x + n) \bmod 12$. (See illustration below.)

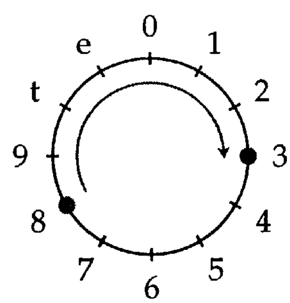

Pitch class transposition. $T_7(8)$ means transpose pitch class 8 by directed pitch class interval 7.

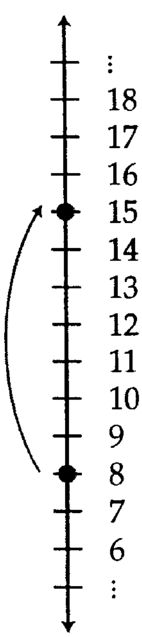

Pitch transposition. $T_7(8)$ means transpose pitch 8 by directed pitch interval 7.

Finally, to transpose a *set* of pitches or pitch classes, we transpose member of the set (each pitch or pitch class) and gather the results into a new set:

$$\mathrm{T_n}\{x_1,x_2,x_3,\ldots\}=\{\mathrm{T_n}(x_1),\mathrm{T_n}(x_2),\mathrm{T_n}(x_3),\ldots\}$$

☞ *Read this next part carefully!*

Tn SET CLASS. Just as a pitch class was a set of (octave-equivalent) pitches, a *Tn set class* is a set of (transpositionally equivalent) pitch class sets. "Transpositionally equivalent" means that each set in the class can be transposed to produce every other set in the class.

A concrete example will make this concept clearer. Let's begin with a pitch class set we considered earlier: the C major triad. By transposing this set, we can produce a D-flat major triad, a D major triad, and so on, so all of these results are members of the same Tn set class as our initial C major triad. Of course, transposing this initial set will never give us a minor triad, or an augmented triad, or a chromatic trichordal cluster, so none of these other chord formations belongs to the set class we're building. And that's the idea behind Tn set classes: each one is a distinct chord formation, like major triad or chromatic trichordal cluster. Of course, we don't always encounter these things as chords. As a chord, the chromatic trichordal cluster is not to everyone's liking (although it gets better when voiced as a stack of pitch interval 11s rather than pitch interval 1s). But this formation is perfectly likeable as a scale fragment for building melodies. (Think of the countermelody from the slow middle movement of Gershwin's *Rhapsody in Blue.*)

Parenthesis time again (meaning you can skip this paragraph if you're feeling anxious). Recall our earlier calculation that there were 4,095 distinct non-empty pitch class sets. It can be helpful to regard the Tn set classes as a classification of this large number of sets into a smaller number of categories representing distinct chord formations. The reduction in numbers is not what you might at first think — that is, it's not $4{,}095 \div 12 = 341\frac{1}{4}$ — because not every chord formation has twelve distinct transpositions.[1] The whole-tone scale, for instance, has only two distinct transpositions. There's no uncomplicated shortcut for calculating what the total number of distinct Tn set classes will be. But you can make a list and count, or get a computer to do this for you. The result is reported to be 351 distinct Tn set classes (and one more for the set class whose only member is the empty set).[2]

Tn PRIME FORM. It would be convenient to have a uniform way of naming Tn Set Classes. If you say "the Tn Set Class that includes the E-flat major triad," and I say "the Tn Set Class that includes the F-sharp major triad," it might take us a moment to be sure we're talking about the same animal. Of course, in this case there's a handy generic name available: the major triad. But there's a shortage of such names. So generally we take the same approach in naming set classes that

[1] Your first hint, of course, is that suspicious " $\frac{1}{4}$"!

[2] Robert Morris, *Composition with Pitch Classes: A Theory of Compositional Design* (New Haven: Yale University Press, 1987).

we took for naming pitch classes: we choose a representative member of the class. (Be sure you recall that this is so for pitch classes — for instance, the pitch class 2 represents the class {…,–22,-10,2,14,…}.)

Given a set class X, we choose a representative set from it as follows:

1. Start with *any* set S that's a member of X.

2. Write the pitch classes of S as pitches ascending within an octave, leaving plenty of space for step 3.

3. Using each pitch from step 2 as a starting point, write the remaining pitch classes of S as pitches ascending within an octave. You will have a separate ascending ordering built up from each pitch class in S.

4. Choose the ordering that has the smallest interval from first to last. If there's a tie, choose the ordering that has the smallest interval from first to second-to-last. If there's still a tie, choose the ordering that has the smallest interval from first to third-to-last, and so on.

5. If step 4 yields a single ordering, transpose so that it ascends from 0. If step 4 yields more than one ordering (if the tie never breaks), choose any one of them and transpose it to zero.

The set selected by the above procedure is called the *Tn prime form* of X. The logic behind the prime form procedure is simple: we want to represent the set class with its "smallest" member. That's why we choose a set that includes 0, and whose other elements are relatively close to 0.

The example below shows two pitch sets A and B. What is the Tn prime form of the set class to which each belongs? The solutions are shown in stages. Tn prime forms are conventionally written in square brackets [] with a "Tn" subscript to distinguish them from the more common TnI prime forms (which we haven't gotten to yet). Often the commas between elements are omitted. As the example below illustrates, A belongs to $[02369]_{Tn}$ and B belongs to $[0268]_{Tn}$.


A
B
A
9
t
e
9
9
6
8
6
3
5
→
0 2 3 6 9

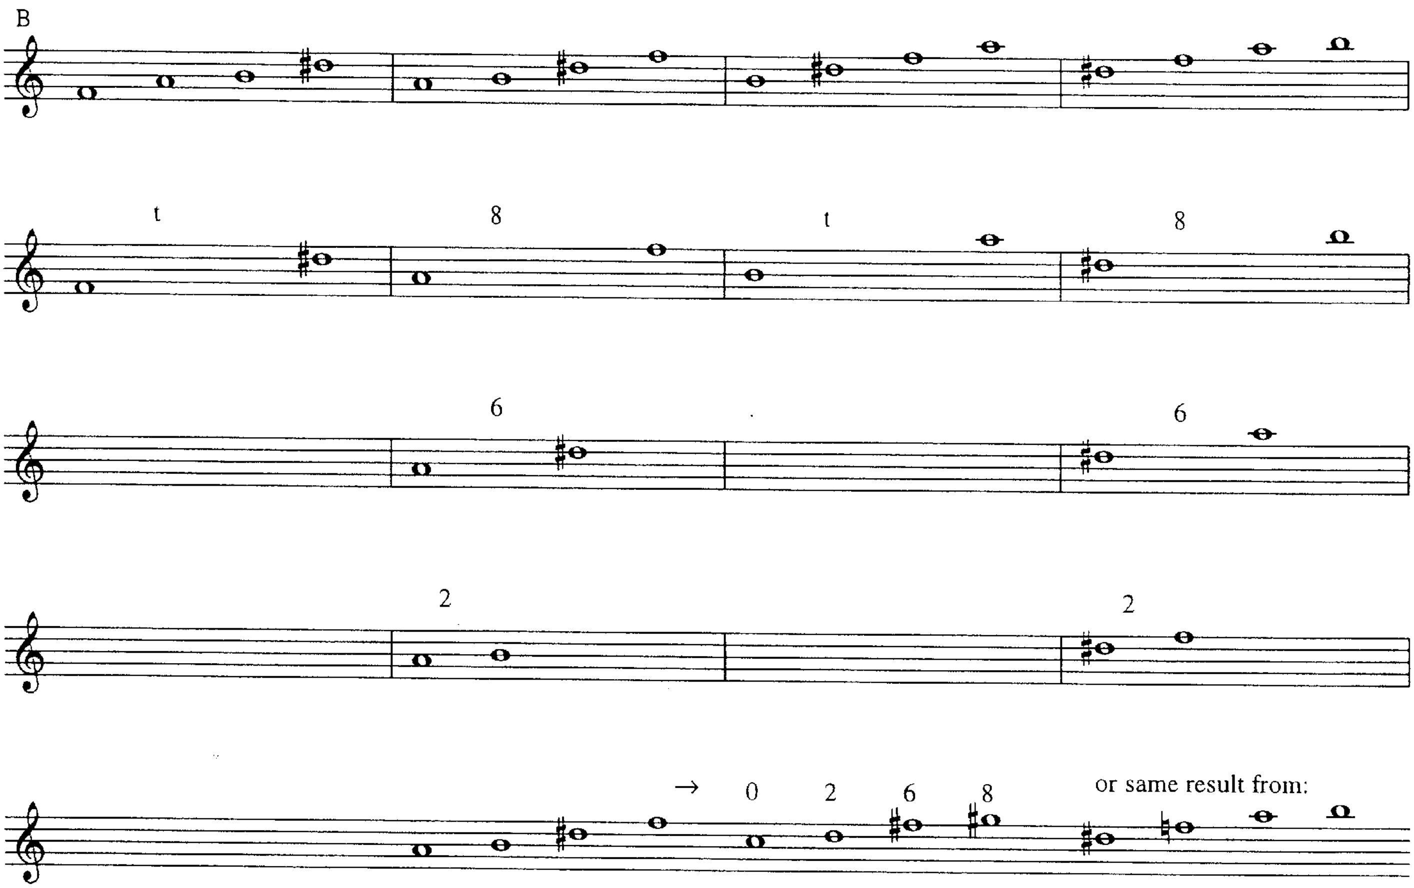
B
t
8
t
8
6
6
2
2
→ 0 2 6 8
or same result from:

with corrections, p. 3

M100C
Notes on Pitch-Class Set Theory - 3

INVERSION. To invert something, we measure its distance from a point of reference, and then we move it an equal distance in the opposite direction. The reference point is often called the *center of inversion.* Let's consider three examples involving pitches (rather than pitch classes).

First, to invert the pitch set {10, 13, 14} around the center 8, we measure the distances 10 – 8 = +2 and 13 – 8 = +5 and 14 – 8 = +6. Reversing the direction of these intervals gives us –2 and –5 and –6. Finally, moving away from the center 8 by these new amounts gives us 8 – 2 = 6 and 8 – 5 = 3 and 8 – 6 = 2. Thus {10, 13, 14} inverted around 8 is {2,3,6}.

For our second example, we will invert the same pitch set {10,13,14} around a new center: 0. Relative to this new center of inversion, we measure the new distances 10 – 0 = +10 and 13 – 0 = +13 and 14 – 0 = +14. Reversing the direction of these intervals gives us –10 and –13 and –14. And finally, moving away from the center 0 by these new amounts gives us 0 – 10 = –10 and 0 – 13 = –13 and 0 – 14 = –14. Thus {10, 13, 14} inverted around 0 is {–14,–13,–10}.

Before continuing with our third example, we can make some useful observations. First, we're at a clear advantage when the center of inversion is the pitch 0, because then we just reverse the signs (from + to – and vice versa) on the pitches we're inverting: x inverted around 0 is $-x$. Second, when we invert the same set around different centers, we get different transpositions of the same result. For instance, our first result above was {2,3,6}, our second was {–14,–13,–10}, and {2,3,6} = T_{16}{–14,–13,–10}. This suggests a shortcut way of computing the inversion around a nonzero center: invert around 0 and then transpose.

Before we can use this shortcut, we need to tackle one question: by what interval do we transpose the inversion-around-0, if we're seeking the inversion-around-c ? Let's compare the results:

x inverted around 0 is $\underline{\underline{-x}}$

x inverted around c is $c - d$, where d is the distance from c to x

$d = x - c$

x inverted around c is $c - (x - c) = c - x + c = \underline{\underline{2c - x}}$

This surprising result is correct: to invert x around c, we invert around 0 (giving $-x$) and then transpose by $2c$ (giving $-x + 2c$).

The pitch x inverted around the pitch c is $2c - x$.

For our third example, let's look at the same process from the other direction. Say we encounter the pitches {–1,0,3} and we recognize that they're an inversion of {10, 13, 14}. (The most likely way we would recognize this is by noticing that {–1,0,3} is a minor third with a half step attached below it, while

{10,13,14} flips this pattern: it's a minor third with a half step attached above it.) What is the center of inversion that turns {10,13,14} into {–1,0,3}? We can simplify the question by concentrating on just one of the pitches involved — let's pick 13, which inverts to become 0:[1]

13 inverted around c is $2c - 13$

13 inverted around c is 0

$2c - 13 = 0$

$2c = 13$

$c = 13 \div 2 = 6\frac{1}{2}$

Again, this surprising result is correct: the center of inversion in this case is halfway between the pitches 6 and 7. We can easily extend the integers to include halves, but as long as the music we're studying limits itself to the familiar chromatic scale, there's no compelling need for such an extension. Instead, we normally keep our world fraction-free and proceed as follows: we define one "official" case of inversion, which is inversion around 0, and then we team it up with transposition to take care of the other cases.

If x is a pitch, then $T_nI(x) = n - x$.

Still, thinking in terms of a center of inversion — even a fractional one — is sometimes valuable; you should understand both ways of modelling inversion.
Finally, we can translate this process into one involving pitch *classes.*

If x is a pitch class, then $T_nI(x) = (n - x) \bmod 12$.

Where a pitch and its inversion balance up and down from their center of inversion, a pitch class and its inversion balance clockwise and counterclockwise from theirs. And this leads to one difference between pitch inversion and pitch class inversion. As pitches, 10 and 6 balance around 8; as pitch classes, 10 and 6 balance around 8 but also around 2 (6 is four steps clockwise, 10 is four steps counterclockwise). Likewise, pitch classes 0 and 1 balance around both $\frac{1}{2}$ and $6\frac{1}{2}$. In general, pitch class inversion around c is equivalent to pitch class inversion around $(c + 6) \bmod 12$.

TnI SET CLASS. A set of pitch class sets related to one another by transposition and/or inversion. Earlier, we characterized each Tn set class as a distinct chord formation, and we observed that one such set class was the familiar major triad $[047]_{Tn}$, while another was the minor triad $[037]_{Tn}$. Now we're merging the contents of these two classes to produce one more inclusive one.

[1] Make sure you see that it's 13 which inverts into 0, while 10 inverts into 3 and 14 inverts into –1. You have to compare the minor-third-plus-half-step patterns described above.

There are several good reasons for putting a chord and its inversion into the same set class.[2] For one thing, pitch class sets related by inversion do represent "similar" sonorities. That is, while we sometimes regard major and minor triads as "opposites," it is nevertheless clear that they resemble one another far more than either resembles a chromatic trichordal cluster. Another advantage of working with larger, more inclusive set classes is that there are fewer of them. The 4,095 non-empty pitch class sets represent 352 Tn set classes but only 224 TnI set classes. For the composer who works habitually with these materials, the 224 TnI set classes are a small enough number to learn by heart, with some appreciation for the range of sonorities that each set class provides. As a 100C student, you will *not* be asked to memorize the 224 TnI set classes! But you are expected to become familiar with both the 12 distinct trichordal (3-element) TnI set classes, and the 29 distinct tetrachordal (4-element) TnI set classes.

TnI PRIME FORM. To name the TnI set classes, we elect a representative member from each class, using the same logic as we did for Tn set classes. We seek the "smallest" set in the class: one that includes 0 and whose other elements are relatively close to 0. The only difference in the case of TnI set classes is that we choose our representative from a larger field of candidates.

this is corrected

Given a TnI set class X, we choose a representative set from it as follows:

1. Start with any set S that's a member of X.
2. Find the Tn prime form of the Tn set class that includes S, and call it A.
3. Find the Tn prime form of the Tn set class that includes $T_0I(S)$, and call it B.
3. Compare A and B. Choose the one that has the smallest interval from first to last. If there's a tie, choose the one that has the smallest interval from first to second-to-last. If there's still a tie, choose the one that has the smallest interval from first to third-to-last, and so on.
4. If the tie in step 3 never breaks, then A and B are identical. Choose either one. Of course, you'll notice this before you invest any time in step 3, right?

TnI prime forms are conventionally written in square brackets [], often with no commas between elements: [037], [0134679t]. If we wish to emphasize that we mean a TnI set class (as opposed to a Tn one), we may add a "TnI" subscript: $[02479]_{TnI}$.

[2] Now is a good time to remind the reader that "inversion" means inversion around a center — the transformation defined earlier in this handout. It is *not* the process of placing different chord members in the bass that we call "inversion" in tonal theory.

INTERVAL CONTENT OF A PITCH CLASS SET. The *interval vector* of a pitch class set is simply a tally of how many times each interval class 1–6 occurs in the set. Consider the tetrachord {0126}. To determine its interval content, we inspect it as follows (note that *we exclude the interval class 0s* between each element and itself):

the interval class between	*this*	*and*	*this*	*is*
	0		1	**1**
	0		2	**2**
	0		6	**6**
	1		2	**1**
	1		6	**5**
	2		6	**4**

If you perform the same test for any transposition of the same tetrachord T_n{0126} or any inversion of it T_nI{0126}, you will get exactly the same result. Thus the interval vector reports a property of the whole T_nI set class; it is not limited to a particular pitch class set within that class.

The usual way to report a T_nI set class's interval content takes the form of a string of numbers interpreted as shown:

(# of interval class 1)(# of interval class 2)...(# of interval class 6)

Returning to our example with $[0126]_{TnI}$, we find (inspecting the column printed above in **bold**) that **1** occurs 2 times, **2** occurs 1 time, **3** occurs 0 times, **4** occurs 1 time, **5** occurs 1 time, and **6** occurs 1 time. Formatting this information as suggested above, finally, gives us the interval vector 210111.

COMMON-TONE THEOREMS - 1. Turns out that the interval vector can answer a very important question for us. Suppose we start with a pitch-class set X and then we want to move to a new transposition of the same set, T_nX. The Very Important Question is, How many pitch classes will X and T_nX have in common?

Let's consider this question first in terms of the pitch-class set A = {0126}, since we know so much about this set. If we take T_1A, then

0	*moves to*	1
1	*moves to*	2
2	*moves to*	3
6	*moves to*	7

summary: two common pitch classes (1, 2) and two new ones (3, 7)

In contrast, if we take T_3A, then

0	*moves to*	3
1	*moves to*	4
2	*moves to*	5
6	*moves to*	9

summary: no common pitch classes and four new ones (3, 4, 5, 9)

Comparing the results shown above for T_1A and T_3A, we can distinguish two cases:

1. A pitch class moves into the position formerly occupied by another element.

2. A pitch class moves into a previously unoccupied position.

The more times case 1 happens, the more common tones we get. The important thing is this: case 1 is going to happen whenever the interval of transposition is an interval contained in the original set. For instance, pitch classes 1 and 2 are both members of the original set A, and the distance between them is interval class 1. Now when we transpose pitch class 1 by this distance, it's going to land on pitch class 2. *Et voilà*, a common tone. In contrast, no element of A stands at a distance of 3 from any other member. So when we transpose A by the distance 3, each element will land in a previously unoccupied position. And the transposed result will thus have no pitch classes in common with the original.

Thus we arrive at the following very important result (but be sure to read on for one equally important exception!):

> The number of common tones between any set X and any transposition of the same set T_nX is equal to the number of times that the interval n occurs in X.

When n is a value from 1 to 5, we can read the number of common tones directly from the interval vector of X. When n is a value from 7 to e, we look at the interval vector in the position for $12 - n$. (Why? Because you get the same number of common tones whether you transpose clockwise or counterclockwise, which is obvious if you think about it this way: the number of common tones between X and T_1X is the same as the number of common tones between T_1X and X. And just as T_1X is X transposed one position clockwise, so X is T_1X transposed one position counterclockwise.) Thus the number of common tones between {0126} and T_t{0126} is 1 (looking at the entry for 12 – 10 in the interval vector 210111). T_t{0126} = {t,e,0,4} — it works!

Of course transposition by a distance of zero gives us the original set: every tone is a common tone. We could include a position in the interval vector for interval class 0, but it wouldn't tell us anything we don't already know. So interval 0 doesn't really complicate anything. The one and only complication we encounter is with the tritone, because of its symmetry. The transposition that moves C into F-sharp, for instance, will also move F-sharp into C: that's two common tones for the price of one tritone. As a result:

To determine the number of common tones at T_6, we must *double* the number of interval class 6s shown in the final position of the interval vector.

COMMON-TONE THEOREMS - 2. Just as we are sometimes interested in knowing the number of common tones between a set X and some transposition of it T_nX, we may likewise ask the analogous question about X and some inversion T_nIX. To determine the common tones for transposition, we needed the interval vector. Now, to determine the common tones for inversion, we shall need another specialized tool. (Of course, we don't "need" these tools; there's always trial and error.)

Let's use the example of A = {0126}. Recall that we construct the inversion T_nIA by subtracting each element of A from *n*. If we want common tones, then we want the result of subtraction to match a value in the original set. That is, we want some element *a* of A such that $n - a = b$ is also an element of A (it could be identical to *a*, or it could be another element of A). And this will happen whenever $n = a + b$.

Therefore, the tool that will tell us about common tones under inversion is an *addition table* (as usual, we perform the addition mod 12):

+	0	1	2	6
0	0	1	2	6
1	1	2	3	7
2	2	3	4	8
6	6	7	8	0

From such a table, we can then compile a *sum vector*, which simply tallies the number of occurences of each sum 0 through e in the addition table:

(# of 0 sums)(# of 1 sums)...(# of e sums)

In the case of {0126}, the sum vector is 223210222000; this tells us that T_2I{0126} will have three common tones with {0126}, T_5I{0126} will have zero common tones with {0126}, and so on. Generalizing, we discover the following principle:

The number of common tones between any set X and any inversion of the same set T_nIX is equal to the number of times the sum *n* is occurs in the addition table for X.

M100C
Notes on Pitch-Class Set Theory - 6

LARGE SETS. Relatively large sets, like octachords (eight elements) and nonachords (nine elements), play an important role in much posttonal music. For example, the octatonic scale [0134679t] is a set that figures prominently in the music of Igor Stravinsky, Olivier Messiaen, and Béla Bartók among others. The important properties of the octatonic scale include the following:

1. It contains multiple instances of several diatonic elements. For instance, the subsets {1346}, {4679}, {79t0} and {t013}correspond to steps $\hat{1}$–$\hat{4}$ of different minor scales; the subsets {047}, {37t}, {6t1}, and {914} are different major triads; and it also includes minor triads and dominant seventh chords.
2. It is inversionally symmetrical; for instance, T_1I{0134679t} = {0134679t}.
3. It is one of what Messiaen calls the "modes of limited transposition,"[1] which means that at certain intervals of transposition, every tone is a common tone: T_3{0134679t} = T_6{0134679t} = T_9{0134679t} = {0134679t}.

There are other large sets with similar properties. One is the nonachord [01245689t], another favorite of Messiaen's. The properties of this set include the following:

1. It contains multiple instances of several diatonic elements. For instance, the subsets {0245}, {4689}, and {8t01} correspond to steps $\hat{1}$–$\hat{4}$ of different major scales; the subsets {158}, {269}, {590}, {6t1}, {914}, and {t25} are different major triads; and it also includes minor triads and dominant seventh chords.
2. It is inversionally summetrical; for instance, T_2I{01245689t} = {01245689t}.
3. It is another of Messiaen's modes of limited transposition: T_4{01245689t} = T_8{01245689t} = {01245689t}.

We can be systematic in exploring these large sets. For instance, we can always recognize modes of limited transposition according to the following principle:

> A set X is a *mode of limited transposition* if the number of common tones produced at some non-zero transposition of X equals the number of elements in X. If the number of elements in X is *j*, then we simply look at the interval vector for X; if one of the first five entries in the vector is *j*, or if the sixth entry is *j*/2, then X is a mode of limited transposition.

---

[1] Messiaen, *The Technique of My Musical Language*, transl. John Satterfield (Paris: Leduc, 1956).

For example, the interval vector for X = {01245689t} is 666963; the fourth entry of this vector is 9, which equals the number of elements in X. This confirms that X is a mode of limited transposition.

Unfortunately, the sheer size of the larger sets complicates our study of them. For example, figuring out the interval vector of a nine-element set, or figuring out its prime form so we can look up its interval vector in a table, is a laborious process. (Indeed, Messiaen appears to have constructed his own list of modes of limited transposition by trial and error, and it is not complete.) Fortunately, we can determine a number of properties of any large set, including its interval vector, by inspecting a related, smaller set instead.

COMPLEMENTATION. The *complement* of a pc set X is the set of all pcs that do *not* belong to X. Thus the complement of any nonachord will be a trichord, the complement of any octachord will be a tetrachord, and so on.

It turns out that a set and its complement have an important relation, which we can understand by thinking as follows. Imagine a small set containing the single pc 0; this of course gives us a large complementary set containing the eleven pcs {123456789te}. While the small set doesn't include any interval classes (other than class 0, the unison), the large set includes many instances of *every* interval class. For example, the large set includes lots of class 4's; the only members of this class that the large set does *not* include are the two that involve pc 0, since this one pc is missing from the large set.

Now imagine that we want to transfer one additional pc from the large set to the small one, and suppose we want to do this in such a way that the large set ends up including as many interval class 4's as possible. There are two cases:

1. If we give pc 4 or 8 to the small set, then the large set only loses one of the interval class 4's that it used to include. (Specifically, if we transfer pc 4, then the large set loses {48} but it never had {04}. If we transfer pc 8, then the large set loses {48} but it never had {80}.)

2. If we give a pc other than 4 or 8 to the small set, then the large set loses two of the interval class 4's that it used to include. (For instance, if we transfer pc 5, then the large set loses {15} and 59}.)

Thus it turns out the best way to keep the maximum number of interval class 4's in our large set is to give an interval class 4 to our small set.

If you think about scenarios like this one, you can convince yourself of the following principle: the more instances of an interval classes that a particular set contains, the more its complement will contain as well. And in fact, we can be much more quantitative about the relationship:

If the interval vector of X is *abcdef,* and Y is the complement of X, and Y contains *n* more pcs than X, then the interval vector of Y is

$$(a+n)(b+n)(c+n)(d+n)(e+n)(f+n/2)$$

The above principle is sometimes known as the *complement theorem.* Let's put it to work. Suppose we want to determine if the set Y = {0135679e} is a mode of limited transposition. To make use of the result developed above (in a dashed box), we need to know the interval vector of Y. Rather than figuring out the interval vector (or the prime form) for this large set directly, we can use the complement theorem. The complement of Y is X = {248t}. Here are the interval classes that X includes:

the interval class between	*this*	*and*	*this*	*is*
	2		4	2
	2		8	6
	2		t	4
	4		8	4
	4		t	6
	8		t	2

Tallying these results, we find that the interval vector for X is 020202. Y has four more pcs than X, so the complement theorem tells us that the interval vector for Y should be $(0+4)(2+4)(0+4)(2+4)(0+4)(2+4/2) = 464644$. Finally, the number of common tones between Y and T_6Y equals twice the number of interval class 6's in Y, which is $2 \times 4 = 8$, and this equals the number of elements in Y. So yes, Y is another mode of limited transposition: $T_6Y = Y$.

INTERESTING FACTS. There is 1 mode of limited transposition with ten notes, 1 with nine notes, 3 with eight notes, and 4 with six notes. There are no seven-note modes of limited transposition. Of these 9 modes of limited transposition, Messiaen identified 7.

THE Z RELATION. Sounds like a spy novel, but it's not quite that exciting. Earlier, we observed that a set and its T_n and T_nI transformations all have the same interval vector. (That's why we think of the interval vector as reporting a property of a whole TnI set class, rather than just a particular set.) But it turns out the converse is not true: *two sets with the same interval vector are not necessarily members of the same TnI set class.* When people first began to study the resources of the chromatic scale from a set-theoretic perspective, they assumed incorrectly that sets with the same interval vector *would* always belong to the same TnI set class. As a result, some classes were overlooked in early stages of research. In the first widely circulated complete list of TnI set classes,[1] these once-overlooked sets are tacked onto the end of each portion of the list and the letter Z is appended to their names. (Can you guess what Z stands for? I can't.)

Two distinct set classes with identical interval vectors are called Z *relatives* of one another. There's one Z-related pair of tetrachords, three Z-related pairs of pentachords, and 15 Z-related pairs of hexachords. To find the tetrachords and pentachords, you'll have to hunt in the TnI Set Class List (distributed earlier) for matching interval vectors; Z-related hexachords are listed across from another in that list. Because a hexachord and its complement are the same size (6 notes each), the complement theorem tells us that they'll have the same interval vector. It's always true: a hexachord and its complement will without exception have identical interval vectors. So there's two possibilities: either a hexachord and its complement are members of the same TnI set class (which happens in 20 cases), or they're Z relatives of one another (which happens in 15 cases).

ALL-INTERVAL TETRACHORDS. The Z relation will turn out to matter the most in the case of hexachords, since we'll have to consider it when combining two hexachords to produce twelve-tone rows with certain special features. But for now we'll direct our attention to the single Z-related pair of tetrachords: $[0146]_{TnI}$ and $[0137]_{TnI}$. The interval vector for each of these set classes is 111111, which means they have another special property (in addition to being Z-related to one another): they contain one instance of each and every interval class. As a result, these set classes are known as the "all-interval tetrachords." Here's details:

interval class	1	2	3	4	5	6
in [0146]	0,1	4,6	1,4	0,4	1,6	0,6
in [0137]	0,1	1,3	0,3	3,7	7,0	1,7

It's impossible for a trichord (or smaller set) to be all-interval, because the total number of intervals it contains will be less than six. And it's easy for a set larger than a tetrachord to include at least one of each interval, but no such set will include *only* one of each interval, since the total number of intervals it contains will be more than six. So being "all-interval," in the sense of containing one-and-only-one of each interval class, is a special property possessed only by the two tetrachordal set classes identified here.

[1] In Alan Forte, *The Structure of Atonal Music* (New Haven: Yale, 1973).

SEGMENTS. So far our most common interpretation of pitch class sets has been as harmonic sonorities in which their individual elements simply coexist — although we've certainly seen such sonorities "arpeggiated" in various ways to produce melodic lines. But if we're particular interested in the melodic presentation of a pitch class set, then we'll usually want to pay more attention to the *order* in which pitch classes appear. We call an ordered set a *segment* and notate its elements in pointy brackets ‹› in the order in which they appear. For example, the segments ‹01257›, ‹75210›, ‹50172›, and ‹2072521› are among the infinitely many distinct orderings of the unordered set {01257}. Incidentally, we will almost always assume that the order of a segment will be expressed in *time* (so that the ordering proceeds from earliest to latest), but it can also be expressed in other dimensions, like loudness (from softest to loudest, say) or pitch height (from lowest to highest, say).

TRANSFORMATIONS ON SEGMENTS. The transformations familiar from our work with (unordered) sets can be applied to (ordered) segments as well, and it should be obvious how. To **transpose** or **invert** an unordered set, we transposed or inverted its individual elements; as long as we keep the results of these individual transformations in the right order, we can use the same technique to transform a segment. Check for yourself to be sure you can correctly calculate T_8‹2072521› and T_1I‹50172›.[1]

Along with transposition and inversion, another type of transformation is commonly applied to segments: **retrogression,** which reverses the order of a segment, running it from last to first. We symbolize retrogression with a capital "R" and call the result that it produces the *retrograde.* For example, the retrograde of ‹29e854› is R‹29e854› = ‹458e92›.

TWO CLOSELY RELATED NOTATIONS. Frequently, a piece will begin by presenting some prominent segment and continue with a series of transformations of the same segment. (Often, we call this material a "motive" and pay attention to its rhythm and other features in addition to its pitch structure, but let's keep things simple for now.) The are two common methods for labeling the various transformations we commonly encounter. Learn them both!

<u>METHOD 1</u>

X (or any neutral label) — the original segment
T_nX — X transposed *n* semitones clockwise
T_nIX — X inverted around zero and then transposed *n* semitones clockwise
(recall that an efficient way to compute this result is to subtract each element of X from *n*)
RX — the retrograde of X (but RT_0X is more common; see the very next item)
RT_nX — the retrograde of T_nX
RT_nIX — the retrograde of T_nIX

[1] And the answers are:
T_8‹2072521› = ‹t83t1t9›
T_1I‹50172› = ‹8106e›.

METHOD 2
P*n* — the original segment, transposed to begin with the pitch class *n* (P stands for "prime")
I*n* — the inversion of the original segment, transposed to begin with the pitch class *n*
R*n* — the retrograde of P*n*
RI*n* — the retrograde of I*n*

Suppose we begin with the segment ‹35215›. Here are some transformations labeled according to both methods.

	METHOD 1	METHOD 2
‹35215›	X	P3
‹02et2›	T_9X	P0
‹0890t›	RT_7X	Rt
‹42562›	T_7IX	I4
‹48746›	RT_9IX	RI6

TWELVE-TONE ROWS. One particular kind of segment has acquired special significance in a large repertory of twentieth-century music: segments that contain every pitch class exactly once. These are generally known as twelve-tone rows. There is an extensive specialized theory treating twelve-tone rows, but we can say plenty about them if we simply approach them as a special case of the segments we've already begun to study.

One of the first things to say about twelve-tone rows is that order is their *only* interesting feature. That is, the unordered content of every twelve-tone row is the same, {0123456789te}; transforming a particular twelve-tone row doesn't change its unordered contents, it just changes their order.

A useful way to view all of the transformations of a particular twelve-tone row is to construct its *matrix* (or "magic square"). There will be twelve distinct transpositions of the row (P0, P1, ..., Pe) and twelve distinct inversions of it (I0, I1, ..., Ie), and each of these can also occur in retrograde, which makes a total of 48 row forms. Let's use the example of ‹01392e4t7856›[2] (from Schoenberg's Suite for solo piano, Op. 25, his first thoroughly twelve-tone composition, dating from 1923) to see how these forms are arranged in the matrix.

To begin with, we set up a 12 × 12 grid, placing P0 in the top row and I0 in the leftmost column:

[2] Here I should acknowledge that my use of pointy brackets, while consistent with the notation for shorter segments, is uncommon. Instead, a twelve-tone row is most often written as string of pc numbers with no enclosing brackets whatsoever.

	P0 →											
I0 ↓	0	1	3	9	2	e	4	t	7	8	5	6
	e											
	9											
	3											
	t											
	1											
	8											
	2											
	5											
	4											
	7											
	6											

Now, since P0 contains all twelve pcs, we can build all twelve inversions I0–Ie in the columns of this grid, and since I0 contains all twelve pcs, we can build all twelve transpositions P0–Pe in the rows of the grid. Doing just this gives us the complete matrix:

0	1	3	9	2	e	4	t	7	8	5	6
e	0	2	8	1	t	3	9	6	7	4	5
9	t	0	6	e	8	1	7	4	5	2	3
3	4	6	0	5	2	7	1	t	e	8	9
t	e	1	7	0	9	2	8	5	6	3	4
1	2	4	t	3	0	5	e	8	9	6	7
8	9	e	5	t	7	0	6	3	4	1	2
2	3	5	e	4	1	6	0	9	t	7	8
5	6	8	2	7	4	9	3	0	1	t	e
4	5	7	1	6	3	8	2	e	0	9	t
7	8	t	4	9	6	e	5	2	3	0	1
6	7	9	3	8	5	t	4	1	2	e	0

row forms

P — left to right

R — right to left

I — top to bottom

RI — bottom to top

A good way to check your work: **the diagonal from top left to bottom right should be all zeroes.** This will be true for *every* twelve-tone matrix.

The row forms employed by Schoenberg at the beginning of Op. 25 are boxed in the above matrix: P4, Pt, and I6. EXAMPLE 1 shows how he has used each of these. The presentation of P4 is quite straightforward: it is unfolded one note at a time as a melody.

While we often encounter twelve-note melodies like this in pieces based on twelve-tone rows, they are not the only kind of material that can be fashioned from a row. (And it's a good thing they're not — a piece would get pretty tedious if it were nothing but a succession of melodies, each precisely twelve notes long.)

EXAMPLE 1 • Schoenberg's use of P4, Pt, and It at the beginning of Op. 25

One common device used to provide a greater variety of motivic material in twelve-tone music is the division of the row into smaller segments. Divisions into three groups of four notes and four groups of three notes are especially common. A 4+4+4 division may be implicit in Schoenberg's use of P4, but it is explicit in his use of Pt, where the 9th through 12th elements of the row are presented as a counterpoint to the 5th though 8th elements. And a similar segmentation of It gives rise to three layers of counterpoint: the lowest voice takes the 1st through 4th elements, the highest voice takes the 5th through 8th, and the middle voice takes the 9th through 12th.

Listening to the complete texture Schoenberg has built from P4, Pt, and It (EXAMPLE 2), we would have difficult time following these twelve-tone rows in their entirety. But it is easy to hear connections between many of the smaller segments that these rows comprise. For instance, it is easy to hear the imitative relationship between the first through fourth elements of P4 and Pt. Of course, this hearing is supported by details in Schoenberg's composition; most significantly, he has arranged the pitch classes of these segments so that they rise and fall through the same pattern of pitch intervals. (If instead he had presented Pt as B-flat, falling to C-flat, rising to D-flat, falling to G, then its connection to P4 would have been much less clear.)

EXAMPLE 2 • P4, Pt, and It combined at the beginning of Schoenberg's Op. 25

With a little effort, you should be able to hear several more connections between tetrachordal segments of the various row forms in these measures. Can you hear the 1st through 4th elements of I4 as an "upside down" (inverted) version of the shape that begins the piece? (Actually, this particular connection is strongest for the 1st through 3rd elements.) Listen, too, for connections between the various transformations of the 5th through 8th elements, and more connections between the various transformations of the 9th through 12th elements.

What this exercise demonstrates is that, even in a piece where the segmentation is made less explicitly than what we've just seen, **the twelve-tone row may be best regarded as a source of relatively compact segments.** Tracing connections across the most prominent of these gives us a reasonable way to hear twelve-tone music as coherent. (At least, it's more reasonable than trying to memorize a twelve-tone row and keep track of transformations of the complete row!) And recognizing these smaller connections helps us understand the twelve-tone music of Schoenberg and his students as an outgrowth of their earlier atonal works (in which motivic connections between pc segments play a fundamental role where, of course, no twelve-tone rows are present).

By the way, there's one more segmental feature of Schoenberg's Op. 25 row that's too good to miss. What do you notice about the 1st through 4th elements of R4? (Remember, first of all, that R4 is the retrograde form that *ends* with pc 4; and if you've noticed only that elements 1–4 of it are a chromatic tetrachord, then you're not paying enough attention to their *order*.)

EXAMPLE 1 • Schoenberg's use of P4, Pt, and It at the beginning of Op. 25

P4

Pt

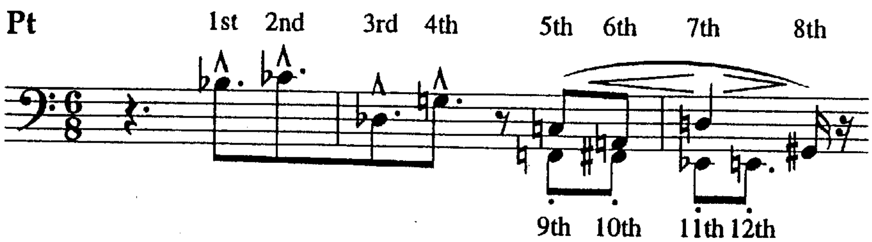

It

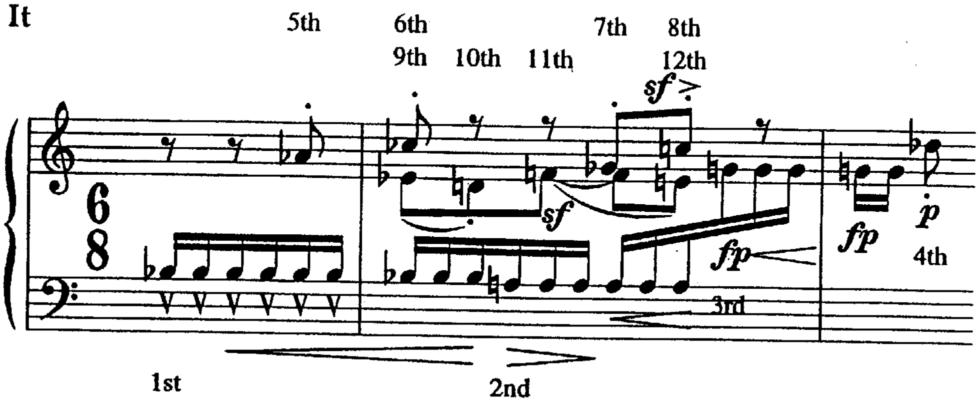

EXAMPLE 2 • P4, Pt, and It combined at the beginning of Schoenberg's Op. 25

SEGMENTAL DERIVATION. Earlier, we saw that a twelve-tone row may be treated as a bundle of smaller segments, some of which may be prominent in a particular musical setting. For instance, at the beginning of Schoenberg's *Suite* Op. 25 (and frequently throughout that composition), the row is clearly divided into three tetrachordal segments, each with independent motivic significance: 0139•2e4t•7856.

A logical extension of this segmental approach is to build twelve-tone rows whose segments are transformations of one another. Schoenberg's Op. 25 row is *not* such a row (for instance, no transposition, inversion, and/or retrogression will transform ‹0139› into ‹2e4t›). But Schoenberg's student Anton Webern was particular fond of rows whose segments are transformations of one another. Here's the row of Webern's *Concerto*, Op. 24: 0e3•487•956•12t. If we call the first trichordal segment X, then note the following:

‹0e3› = X
‹487› = RT_7IX
‹956› = RT_6X
‹12t› = T_1IX

As the above relationships indicate, the complete row in Webern's Op. 24 is *derived* from the trichord X. A summary of this structure is given in EXAMPLE 1, together with an excerpt from the music. Of the possible types of derived rows, ones built from transformations of an initial trichord or tetrachord are the most common.

EXAMPLE 1 • Webern, Op. 24: (a) derivation of P0 from initial trichordal segment; (b) opening measures based on Pe.

VERTICAL AND HORIZONTAL ROW USAGE. A twelve-tone row can serve as a source of both harmonic and melodic material. For instance, the first three elements of a row can be combined to form a chord, or presented in succession to form a melody. But music typically involves both melody *and* harmony. Two common methods of twelve-tone row usage place emphasis on one or the other of these dimensions.

Vertical row usage distributes a single row form across all of the layers of a musical texture. For example, the row given in EXAMPLE 2 (a) is treated vertically in EXAMPLE 2 (b). This type of row usage makes chords conform to the row structure. For instance, the three-note chords at the beginnings of measures 1 and 2 occur in the row as the 1st through 3rd and 8th through 10th elements, respectively. Likewise, the dyads formed between the upper voices occur as intervals of the row. (Note that these voices have been arranged to emphasize the heavy presence of interval class 2 in the row.)

EXAMPLE 2 • (a) a twelve-tone row; (b) the row used vertically; (c) the row used horizontally

The other side of the coin, when a row is distributed vertically, is that the resulting melodic fragments typically will *not* correspond to the pattern of the row itself. Take a look at the bass voice in EXAMPLE 2 (b). It consists of pitch classes that are non-adjacent in the row (specifically the 1st, the 8th, and finally the 3rd elements of the row), and the melodic intervals that it traverses (intervals 3 and 6) never occur between adjacent elements of the row. Furthermore, the uppermost voice ends by outlining a major triad (F-C-A), another element foreign to the row in use. Thus vertical row usage provides serial control over harmony but sacrifices serial control over melody.

In *horizontal* row usage, the situation is reversed. In EXAMPLE 2 (c), for instance, the individual melodic lines are each made to conform to a row form. (Here, the uppermost line is a segment of Ie, and the middle and low lines are segments of P6.) But serial control over melody comes at a price, for now the harmonic structure of the music bears little resemblance to the structure of the row. Note for instance the quartal chord at the beginning of the passage, the minor triad at the end, and the accidental "octave" that falls on the downbeat of measure 2; all of these elements are foreign to the original row. Conversely, there is in this passage hardly a trace of the interval class 2's that populate that row.

Of course, the flaws of vertical and horizontal row usage are less than fatal, and many composers of twelve-tone music choose to live with some combination of them. When a row is used vertically, melodic figures can conform to at least the approximate rise and fall (or *contour*) of established motivic shapes, and their motivic identities can be reinforced by the use of consistent rhythmic profiles. And when a row is used horizontally, careful rhythmic choices give the composer control over which pitches are reached simultaneously. (For instance, the "octave"in EXAMPLE 2 (c) could have been avoided if the bass's A had been delayed until later in the measure.) On top of all this, composers vary in how strictly they adhere to their twelve-tone designs. "Cheating" — by omitting row elements or stating them out of order — is not unheard of in the compositional process of Schoenberg or Berg. (Webern, on the other hand, was notoriously strict in his adherence to the rules of the twelve-tone game.)

COMBINATORIALITY. This fancy word describes a principle that became a consistent part of Schoenberg's twelve-tone practice and has been further developed by more recent composers. Combinatoriality can be seen as an attempt to achieve simultaneous serial control of the horizontal (melodic) and the vertical (harmonic) dimensions. The basic idea can be gleaned from EXAMPLE 3, an excerpt from the first movement of Schoenberg's Violin Concerto, Op. 36. In this passage, the solo violin plays a line constructed from P9, while the orchestra provides an accompaniment based on I2. This row usage is largely horizontal, since different textural layers (melody and accompaniment) present different row forms (although various layers within the accompaniment are handled vertically).

The row forms that Schoenberg has combined fit together in a special way: the unordered pc content of the 1st through 6th elements of P9, plus the unordered pc content of the 1st through 6th elements of I2, equals all twelve pitch classes. This type of configuration, known as *hexachordal combinatoriality,* can be diagrammed as follows (where — this is important — the upper and lower row forms are different transformations of the same row):

upper row form	1st – 6th elements	7th – 12th elements
	↕ *no common tones*	↕ *no common tones*
lower row form	1st – 6th elements	7th – 12th elements

The result of hexachordal combinatoriality is that each layer (upper and lower) contains a true row form (so that there's serial control of the horizontal dimension), and half-of-a-row, when combined with the corresponding half of the other layer's row, adds up to all twelve pcs with no duplications (so that there's a degree of serial control of the vertical dimension — at least, "accidental octaves" are eliminated).

To demonstrate that a row combination is combinatorial, show that the first hexachord of one row and the first hexachord of the other row have no common tones.

Because combinatoriality depends on common-tone relationships, it is possible build twelve-tone rows with any combinatorial properties you like by using the common-tone theorems to help you choose appropriate hexachords. But as we've seen before, the common-tone theorem for inversion can be complicated to use. So composers may rely on trial and error (as Schoenberg did), or they may consult a table that lists all of the hexachords that are useful for building combinatorial rows. (The earliest such table was compiled by — surprise — Milton Babbitt. A particularly straightforward one has been published by Joel Lester.[1])

Schoenberg was especially fond of combining a P form and an I form of the same row (as he has done with P9 and I2 in the Violin Concerto).[2] This is known as *I combinatoriality,* since each of the rows is an inversion of the other. (Likewise, a combination of an R form and an RI form constitutes I combinatoriality, since once again each of the rows is an inversion of the other.) The other types of hexachordal combinatoriality are P combinatoriality (when each of the rows is a transposition of the other), R combinatoriality (when each of the rows is a retrograde of the other), and RI combinatoriality (when each of the rows is a retrograde inversion of the other). The four types of hexachordal combinatoriality are summarized in the following table:

[1] Joel Lester, *Analytic Approaches to Twentieth-Century Music* (New York: Norton, 1989).
[2] Actually, Schoenberg was more particular still: he preferred to combine a P form with an I form whose transposition level was a perfect fifth away. He seems to have thought of this combination as being somehow analogous to tonic and dominant. But remember, the total pitch content of every twelve-tone row is the same, so it's hard to extend the analogy very far.

type	*combines*
P combinatoriality	P with P, I with I, R with R, RI with RI
I combinatoriality	P with I, R with RI
R combinatoriality	P with R, I with RI
RI combinatoriality	P with RI, I with R

Finally, it's worth mentioning that one of these types of hexachordal combinatoriality is particularly simple: every row X is R-combinatorial with its retrograde RT_0X. To confirm this, pick *any* row, pair it up with its own retrograde, and proceed according to "To demonstrate that a row combination is combinatorial," in the box above.

Example 1

(a)

(b)

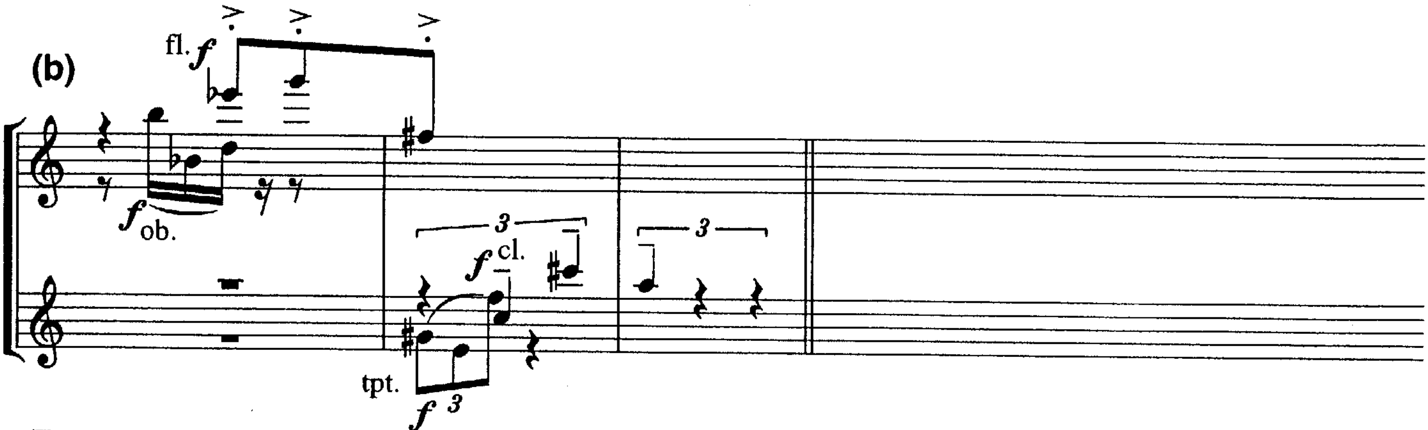

Example 2

(a)

(b)

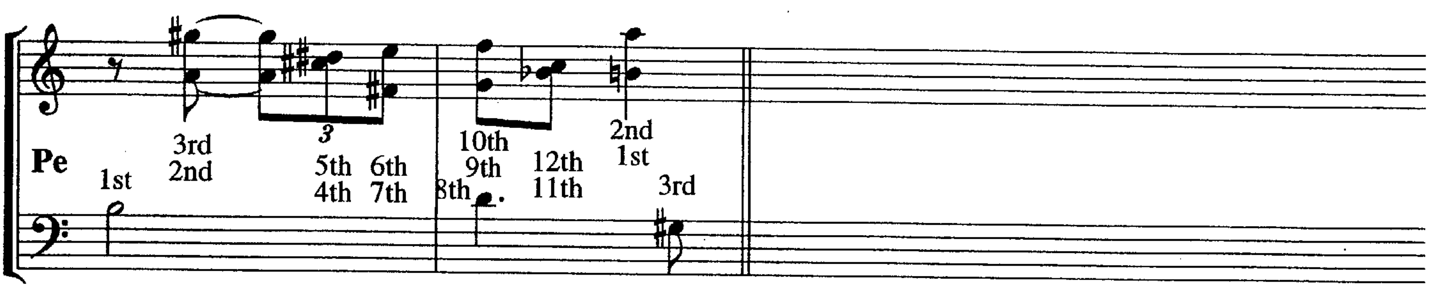

(c)

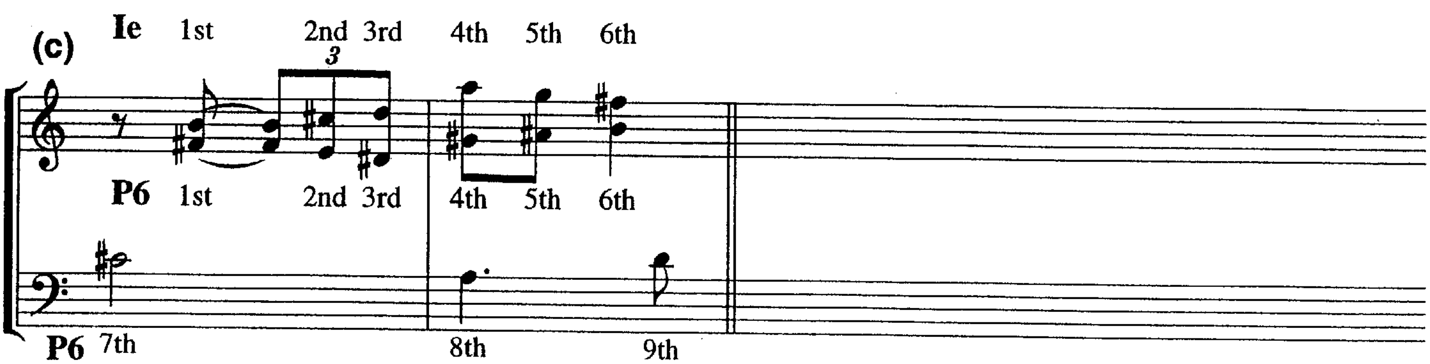

Example 3