the middle of the texture on the second beat of measure 2. In all there are six statements of the motive packed into these three measures. The density is extraordinary; the music of the introduction is motivically saturated. Play these measures and listen for each statement of the motive. The music that follows can be heard as an unpacking of material so intensely presented in the introduction.

Chapter 2
Pitch-Class Sets

Pitch-Class Sets

Pitch-class sets are the basic building blocks of much post-tonal music. A pitch-class set is an unordered collection of pitch-classes. It is a motive from which many of the identifying characteristics—register, rhythm, order—have been boiled away. What remains is simply the basic pitch-class and interval-class identity of a musical idea.

In Example 2–1, you see five short excerpts from a piece by Schoenberg, the Gavotte from his Suite for Piano, Op. 25. In each excerpt, a single pitch-class set (D♭, E, F, G) is circled. That pitch-class set is expressed musically in many different ways. It is the melody that begins the piece and that ends the first section (measure 7). It is heard as a pair of dyads at the beginning of the second half of the piece (measure 16) and as a chord (measure 24). Finally, it returns as the last musical idea of the piece (measure 27).

Something similar happens at the beginning of Stefan Wolpe's *Form for Piano* (see Example 2–2). The piece begins with a simple melodic statement of six pitch-classes: A♭–F–B♭–A–G–E. In the music that follows, the same pitch-class set is stated twice more, but with the notes in different registers, rhythms, and order.

The possibility of presenting a musical idea in such varied ways—melodically, harmonically, or a combination of the two—is part of what Schoenberg meant by his well-known statement, "The two-or-more-dimensional space in which musical ideas are presented is a unit." No matter how it is presented, a pitch-class set will retain its basic pitch-class and interval-class identity. A composer can unify a composition by using a pitch-class set (or a small number of different pitch-class sets) as a basic structural unit. At the same time, he or she can create a varied musical surface by transforming that basic unit in different ways. When we listen to or analyze music, we search for coherence. In a great deal of post-tonal music, that coherence is assured through the use of pitch-class sets.

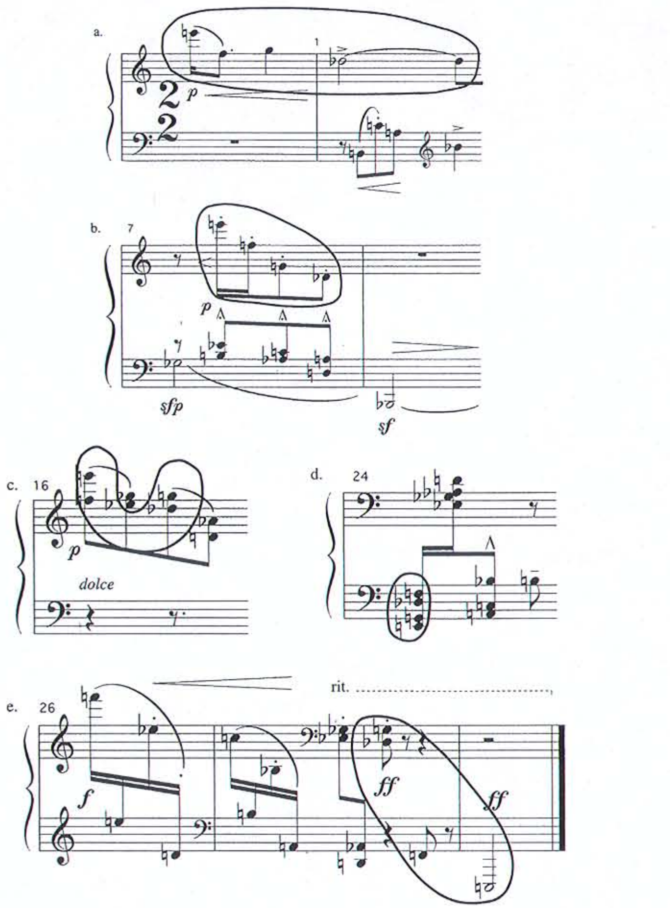

Example 2–1 A single pitch-class set expressed in five different ways (Schoenberg, Gavotte from Suite for Piano, Op. 25).

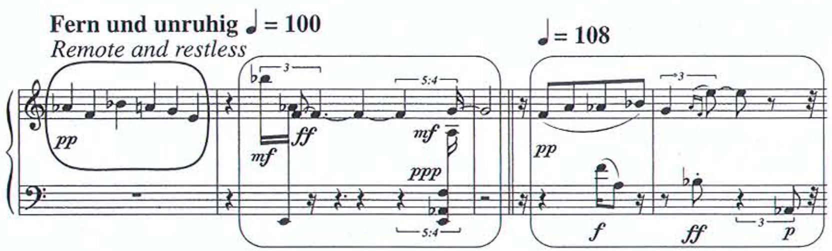

Example 2–2 Three statements of a single pitch-class set (Wolpe, *Form for Piano*).

Normal Form

A pitch-class set can be presented musically in a variety of ways. Conversely, many different musical figures can represent the same pitch-class set. If we want to be able to recognize a pitch-class set no matter how it is presented in the music, it will be helpful to put it into a simple, compact, easily grasped form called the *normal form*. The normal form—the most compressed way of writing a pitch-class set—makes it easy to see the essential attributes of a set and to compare it to other sets.

Consider the first three measures of the third of Schoenberg's Five Orchestral Pieces, Op. 16. Example 2–3 contains a two-piano reduction of a passage that is richly orchestrated and contains thirty-six distinct instrumental attacks.

Example 2–3 A complex surface, but only five pitch classes (Schoenberg, Orchestral Piece, Op. 16, No. 3).

Our task is to boil the sonority down into its normal form. First, we eliminate all duplicates and consider only the pitch-class content. There are only five different pitch classes in the passage: C, G♯, B, E, and A. Next, we write those pitch classes as though they were a scale, ascending within an octave. There are five ways of doing this, and our problem is to choose the smallest stack (the most compact and compressed representation of the set). (See Example 2–4.)

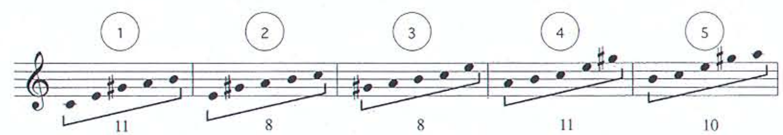

Example 2–4 Finding the normal form.

The first and fourth orderings span eleven semitones from lowest to highest, while the fifth ordering spans ten semitones. Clearly these are not the smallest ways of stacking these notes. Either the second or third ordering would be better, since both span only eight semitones. Now we have to choose between the second and third orderings. In situations like this, our preference is for the one with the larger intervals toward the top of the stack, and whose notes thus cluster toward the bottom. The normal form is the ordering that is most packed away from the right side. The third ordering has only four semitones from its first note to its second-to-last note (G♯–C), while the second ordering has seven semitones from the first to the second-to-last (E–B). We thus prefer the third ordering. The normal form of the sonority from Example 2–3 is [G♯,A,B,C,E]. We will use square brackets to indicate normal forms.

In some ways, the normal form of a pitch-class set is similar to the root position of a triad. Both are simple, compressed ways of representing sonorities that can occur in many positions and spacings. There are important differences, however. In traditional tonal theory, the root position of a triad is considered more stable than other positions, the inversions of the triad being generated from the root position. The normal form, in contrast, has no particular stability or priority. It is just a convenient way of writing sets so that they can be more easily studied and compared.

Here is the step-by-step procedure for putting a set into normal form:

1. Excluding doublings, write the pitch classes as though they were a scale, ascending within an octave. There will be as many different ways of doing this as there are pitch classes in the set, since an ordering can begin on any of the pitch classes in the set.
2. Choose the ordering that has the smallest interval from first to last (from lowest to highest).
3. If there is a tie under Rule 2, choose the ordering that is most clustered away from the top. To do so, compare the intervals between the first and second-to-last notes. If there is still a tie, compare the intervals between the first and third-to-last notes, and so on.
4. If the application of Rule 3 still results in a tie, then choose the ordering beginning with the pitch class represented by the smallest integer. For example, (A, C♯, F), (C♯, F, A), and (F, A, C♯) are in a three-way tie according to Rule 3. So we select [C♯, F, A] as the normal form since its first pitch class is 1, which is lower than 5 or 9.

Now let us reconsider the sonority from Schoenberg's Orchestral Piece (Example 2–3), this time using pitch-class integers and following the procedure just outlined.

1. The five possible orderings are:

0 4 8 9 11
4 8 9 11 0
8 9 11 0 4
9 11 0 4 8
11 0 4 8 9

Notice that each of these orderings is ascending (or clockwise, if you prefer to think of it that way) within a single octave (the first and last elements are less than twelve semitones apart). Having arbitrarily started with the ordering beginning on 0, we just proceed systematically: The second element moves into the first place and the first element goes to the end as we move down the list.

2. We calculate the interval from the first element to the last by subtracting the first from the last:

First ordering: 11 – 0 = 11
Second ordering: 0 – 4 = 12 – 4 = 8
Third ordering: 4 – 8 = 16 – 8 = 8
Fourth ordering: 8 – 9 = 20 – 9 = 11
Fifth ordering: 9 – 11 = 21 – 11 = 10

3. We discover a tie between the second and third orderings.

4 8 9 11 0
8 9 11 0 4

We compare the intervals between their first and second-to-last elements:

Second ordering: 11 – 4 = 7
Third ordering: 0 – 8 = 4

Since 4 is smaller than 7, we conclude that the third ordering [8,9,11,0,4] is the normal form. There is no need to use Rule 4.

In many cases, rather than following this step-by-step procedure, it will be possible to determine the normal form simply by inspecting the set displayed around a pitch-class clockface. Figure 2–1 displays Schoenberg's set from Example 2–3, and it should be apparent that the normal order is the one that starts on 8.

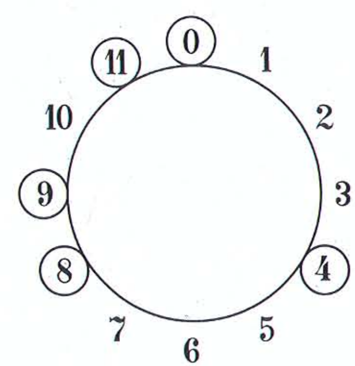

Figure 2–1

Example 2–5 shows the four chords that begin Carl Ruggles's *Lilacs.* Beneath the music, the pitch classes in each chord are identified on a pitch-class clockface. It should be possible, at a glance, to identify the normal forms (which are written beneath the clockfaces).

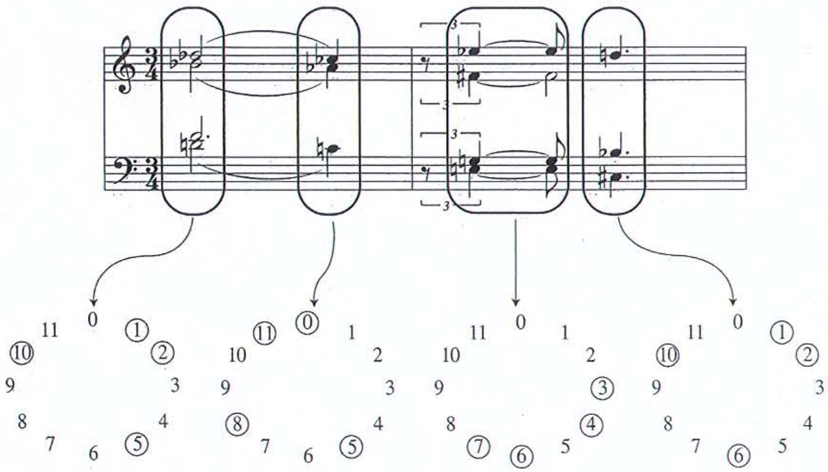

[10, 1, 2, 5]   [5, 8, 11, 0]   [3, 4, 6, 7]   [10, 1, 2, 6]

Example 2–5 Using a pitch-class clockface to determine normal form (Ruggles, *Lilacs,* mm. 1–2).

Or, you might find it easier to play, or imagine playing, the notes on a keyboard, using just one hand and playing them like a scale, ascending within an octave, starting in turn on each note. The arrangement that involves the smallest hand span, and has the biggest gaps toward the top, is the normal form. Whichever method you use, just keep in mind that we are always trying to represent sets in the simplest, most compressed way.

Transposition (T_n)

Traditionally, the term *transposition* refers to the transposition of a line of pitches. When, for example, we transpose "My Country, 'Tis of Thee" from C major to G major, we transpose each pitch, in order, by some pitch interval. This operation preserves the ordered pitch intervals in the line (i.e., the contour of the line). Because contour is such a basic musical feature, it is easy to recognize when two lines of pitches are related by transposition.

Things are different when we transpose a line of pitch classes rather than a line of pitches. We will now be adding pitch-class intervals to each pitch class in the line. Example 2–6 contains the main melody that opens the first movement of Schoenberg's String Quartet No. 4 and a transposed statement of the melody from the middle of the movement.

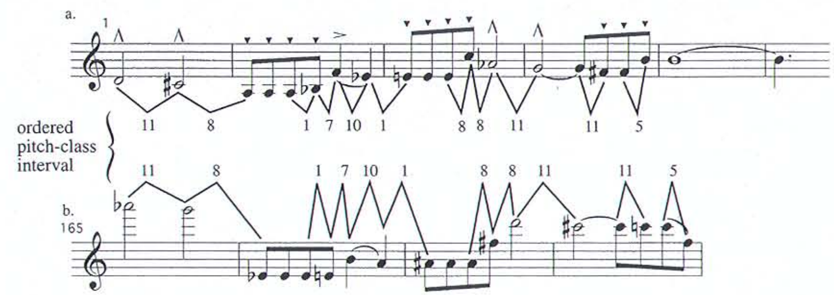

Example 2–6 Two transpositionally related lines of pitch classes (Schoenberg, String Quartet No. 4).

The contours of the two lines are different, so they sound superficially dissimilar. But notice two important features of pitch-class transposition. First, for each pitch class in the first melody, the corresponding member of the second melody lies the same pitch-class interval away—in this case, 6. Second, the ordered pitch-class interval between adjacent elements of the lines is the same in both cases. Both lines have the interval succession 11, 8, 1, 7, etc. That is why, despite their obvious differences, they still sound very similar to one another. (Their shared rhythm helps, too.) The two lines are pitch-class transpositions of one another.

We can describe the same relationship using integer notation. In integer notation, the first melody is: 2, 1, 9, 10, 5, 3, 4, 0, 8, 7, 6, 11. By adding 6 to each integer (mod 12), we produce the transposed version from the middle of the movement (see Figure 2–2).

	2	1	9	10	5	3	4	0	8	7	6	11	(melody beginning in measure 1)
+	6	6	6	6	6	6	6	6	6	6	6	6	
=	8	7	3	4	11	9	10	6	2	1	0	5	(melody beginning in measure 165)

Figure 2–2

The second line is a pitch-class transposition, at pitch-class interval 6, of the first line. We will represent the operation of pitch-class transposition as T_n, where T stands for transposition and n is the interval of transposition (also known as the "transposition number"). Thus, the second line is related to the first at T_6.

Now we must consider the possibility of transposing not a line but a *set* of pitch classes. A set is a collection with no specified order or contour. As a result, transposition of a set preserves neither order nor contour. The four pitch-class sets circled in Example 2–7 (from Webern's Concerto for Nine Instruments, Op. 24, second movement) despite their obvious dissimilarities, are all transpositionally equivalent.

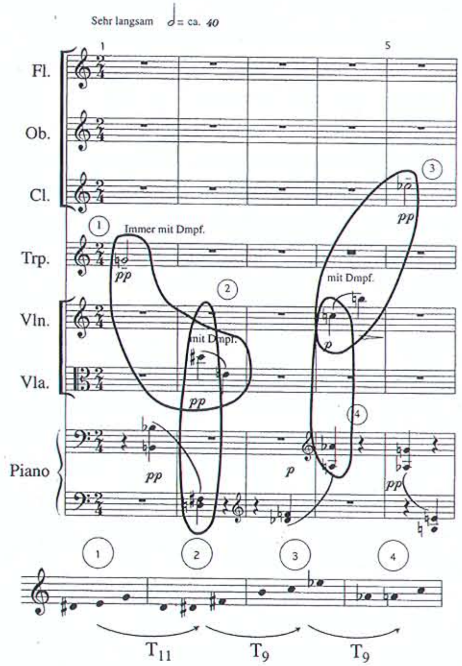

Example 2–7 Transpositionally equivalent pitch-class sets (Webern, Concerto for Nine Instruments, Op. 24).

Beneath the score in Example 2–7, each of the sets is given in normal form—that makes it easy to compare them and to see their transpositional equivalence. To get from Set #1 to Set #2, each pitch class in Set #1 moves down one pitch-class semitone onto a corresponding pitch class in Set #2. To illustrate this, an arrow is drawn from Set #1 to Set #2 and labeled T_{11}. Similarly, Set #2 moves to Set #3 by T_9, and Set #3 moves to Set #4 also by T_9. This will be our usual way of identifying the T_n that connects two sets: we draw an arrow from the first set to the second, and label the arrow with the appropriate T_n.

Because they are transpositionally equivalent, the four sets in Example 2–7 contain the same unordered pitch-class intervals; each of them contains a 1, a 3, a 4, and no others. That gives them a similar sound. Transposition of a set of pitch classes changes many things, but it preserves interval-class content. Along with inversion (to be discussed in the next section), transposition is the only operation that does so and, as a result, it is an important compositional means of creating a deeper unity beneath a varied musical surface.

Now let's look more closely at two melodic fragments from Example 2–7 (see Example 2–8).

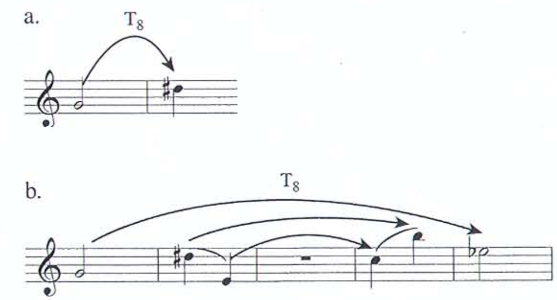

Example 2–8 Transposition within and between two melodic fragments (Webern, Concerto for Nine Instruments, Op. 24).

The first melodic interval is ip+8, and we can imagine the G and D♯ as related by T_8. The very same T_8 is the transposition that leads from the first three-note melodic fragment to the second. Each note in the first fragment moves up eight semitones to a corresponding note in the second fragment. In integer notation, 7 (G) plus 8 is 3 (E♭); 3 (D♯) plus 8 is 11 (B); and 4 (E) plus 8 is 0 (C). The same musical gesture that connects the first note with the second thus also connects the first fragment with the second. In Example 2–8, the arrows indicate the *mappings* brought about by transposition. In moving from the first fragment to the second, T_8 maps G onto E♭, D♯ onto B, and E onto C.

Now we need to discuss more specifically how to transpose a pitch-class set and how to recognize whether two pitch-class sets are related by transposition. To transpose a set, simply add a single pitch-class interval to each member of the set. For example, to transpose [5,7,8,11] by pitch-class interval 8, simply add 8 to each element in the set to create a new set: [1,3,4,7]. (See Figure 2–3.)

$$\begin{array}{rrrrr} & 5 & 7 & 8 & 11 \\ + & 8 & 8 & 8 & 8 \\ \hline = & 1 & 3 & 4 & 7 \end{array}$$

Figure 2–3

More simply, $[1,3,4,7] = T_8\,[5,7,8,11]$. We read this equation either "[1,3,4,7] is T_8 of [5,7,8,11]" or "T_8 maps [5,7,8,11] onto [1,3,4,7]." By *mapping*, we mean transforming one object into another by applying some operation. Here, applying T_8 to 5 transforms it into, or maps it onto, 1; T_8 maps 7 onto 3; and so on. If the first set was in normal form, the transposition of it will be also (with a small number of exceptions related to Rule 4 for determining normal form).

If two sets are related by transposition at interval n, there will be, for each element in the first set, a corresponding element in the second set that lies n semitones away. In our example above, for each element in the first set, [5,7,8,11], there is a corresponding element in the second set, [1,3,4,7], eight semitones away. Discovering this one-to-one correspondence is easiest when the two sets are both in normal form. The first element in one set corresponds to the first element in the other set, the second to the second, and so on. Furthermore, transpositionally related pitch-class sets in normal form have the same succession of intervals from left to right. Both [1,3,4,7] and [5,7,8,11] have the interval succession 2–1–3 (see Figure 2–4).

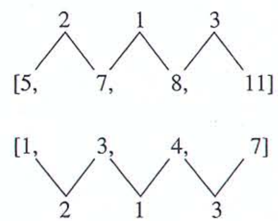

Figure 2–4

Say you are looking at the passage from Stravinsky's *Agon,* shown in Example 2–9, and you suspect there may be some relationship between the two circled sets (beyond the shared pitch classes B♭ and B). First put each in normal form. Both have the interval succession 1–2–1, so we know they are related by transposition. Now compare the corresponding elements. Each member of the second set lies three semitones higher than the corresponding member of the first set. To put it another way, each element in Set 2 minus the corresponding element of Set 1 equals 3 (see Figure 2–5). To put it most simply: Set 2 = T_3 (Set 1). T_3 also is the relationship between the two highest notes in the solo violin, from B♭ in m. 427 to the sustained D♭ in mm. 427–428.

We can represent these transpositional relationships using a combination of *nodes* (circles that contain some musical element, such as a note or a set) and *arrows* (to show the operation that connects the nodes). (See Figure 2–6.) The same operation that moves the music from note to note also moves it from set to set. As a result, while the contents of the nodes in Figures 2–6a and 2–6b are different, the two *networks* are the same. We will make frequent use of networks of this kind to represent musical motion in post-tonal music.

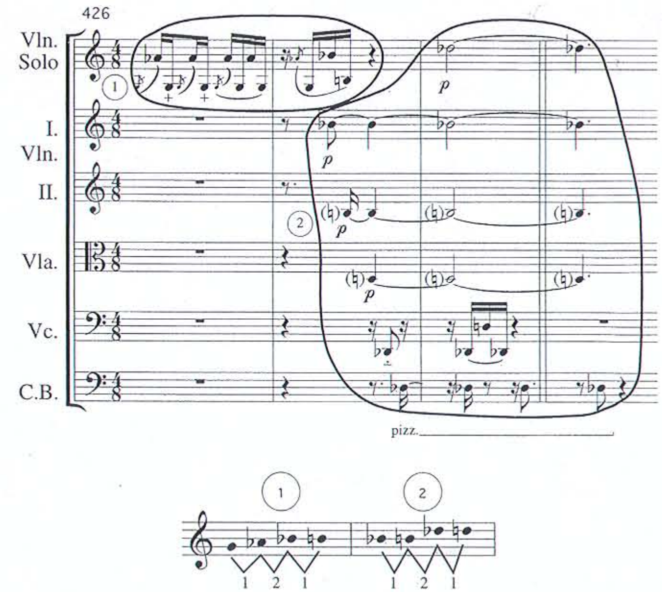

Example 2–9 Transpositionally equivalent pitch-class sets (Stravinsky, *Agon*).

Set 2:	10,	11,	1,	2
Set 1:	– 7,	8,	10,	11
	= 3	3	3	3

Figure 2–5

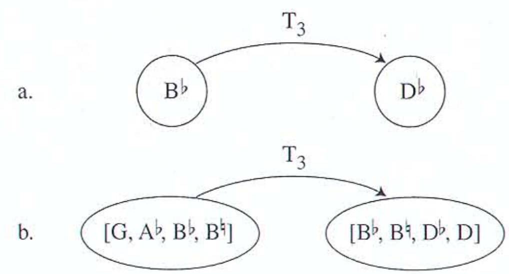

Figure 2–6

We have seen that if we transpose Set 1 at T_3, we map it onto Set 2. Conversely, we would need to transpose Set 2 at T_9 to map it onto Set 1. That is, each element in Set 1 minus the corresponding element of Set 2 equals 9 (see Figure 2–7). To put this relationship most simply: Set 1 = T_9 (Set 2).

Set 1:	7,	8,	10,	11
Set 2:	−10,	11,	1,	2
	= 9	9	9	9

Figure 2–7

If a and b are corresponding elements in two sets related by T_n, then n equals either a – b or b – a, depending upon which set you use as your frame of reference. Notice that these two intervals of transposition (a – b and b – a) add up to 12. (Try to figure out why this should be so.) The two sets in Example 2–9 are different in many ways, but they are transpositionally equivalent. Post-tonal music makes extensive use of this kind of underlying equivalence.

Inversion (T_nI)

Like transposition, inversion is an operation traditionally applied to lines of pitches. In inverting a line of pitches, order is preserved and contour is reversed—each ascending pitch interval is replaced by a descending one, and vice versa.

Inverting a line of pitch classes is similar in some ways. By convention, when we invert a pitch class we invert it around 0. Pitch class 3, for example, which lies 3 above 0, inverts into –3, 3 below 0. In other words, the inversion of 3 is 0 – 3 = –3 = 9. Figure 2–8 summarizes the possibilities.

pitch class (n)	*inversion (12 – n)*
0	0
1	11
2	10
3	9
4	8
5	7
6	6
7	5
8	4
9	3
10	2
11	1

Figure 2–8

In fact, inversion is a compound operation: it involves both inversion and transposition. We will express this compound operation as T_nI, where "I" means "invert" and "T_n" means "Transpose by some interval n." By convention, we will always invert first and then transpose. In Figure 2–8, we inverted and transposed at T_0. Thus, for example, $T_0I(3) = 9$. That is, we invert the 3—that gives us 9. Then we add the transposition number 0 to it, which again gives us 9. We also could transpose by intervals other than 0. For example, $T_5I(3) = 2$. To verify this, we first invert the 3, which gives us 9. Then, to transpose, we add the interval 5, which gives us 2. Remember, always invert first and then transpose. (Sometimes in the theoretical literature the expression I_n is used instead of T_nI, but the meaning remains the same: first invert around 0, as in Figure 2–8, then transpose by interval n.)

Example 2–10 shows two melodies from the beginning of Schoenberg's String Quartet No. 4. These lines of pitch classes are related by inversion.

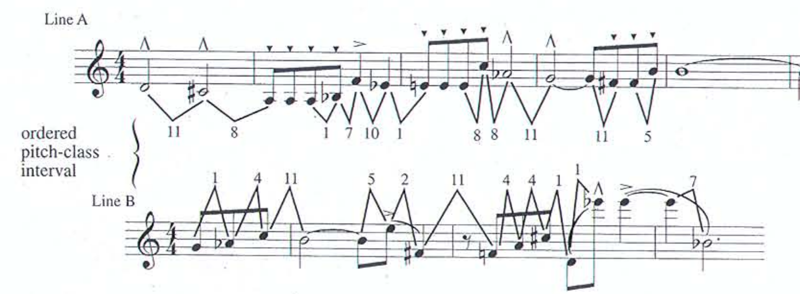

Example 2–10 Two inversionally related lines of pitch classes (Schoenberg, String Quartet No. 4).

Each pitch class in Line B is related by T_9I to the corresponding pitch class in Line A. The first pitch class in Line A corresponds to the first pitch class in Line B, the second to the second, and so on. Let's take one example to verify this. The second note in Line A is C♯, or 1. To perform the operation T_9I on 1, we first invert the 1—that gives us 11. Then, $T_9(11) = 8$. The corresponding note in Line B is, in fact, 8 (A♭). Now let's perform the same operation, T_9I, on the 8 (A♭). Invert the 8—that gives us 4. Then transpose by 9—that gives us 1. So, just as $T_9I(1) = 8$, so $T_9I(8) = 1$. That's because T_nI is its own *inverse,* the operation that undoes the effect of an operation. If you invert something (a note, a line, or a set) by some T_nI and want to get back to where you started, just perform the same T_nI again. By contrast, if you transpose something at T_n, you will need to perform the complementary transposition, T_{12-n}, to get back where you started. For example, to reverse the effect of T_3I, perform T_3I again, but to reverse the effect of T_3, perform T_9.

As with transposition, inversion of a line of pitch classes preserves the ordered pitch-class intervals, only now each interval is reversed in direction. In Line A, the succession of ordered pitch-class intervals is 11–8–1–7, etc. In Line B it is 1–4–11–5, etc. This can probably be seen more clearly using pitch-class integers (see Figure 2–9).

Now we come to the inversion of a *set* of pitch classes. Example 2–11 shows a familiar passage, the opening of Schoenberg's Piano Piece, Op. 11, No. 1. Three sets,

Line A:	2	1	9	10	5	3	4	0	8	7	6	11
ordered pitch-class interval:	11	8	1	7	10	1	8	8	11	11	5	
ordered pitch-class interval:	1	4	11	5	2	11	4	4	1	1	7	
Line B:	7	8	0	11	4	6	5	9	1	2	3	10

Figure 2–9

each involving a combination of soprano and alto notes, are circled and given in normal form beneath the music.

Example 2–11 Three equivalent pitch-class sets (Schoenberg, Piano Piece, Op. 11, No. 1).

Compare the first two sets. Set 2 has the same intervals reading from the top down as Set 1 does reading from the bottom up. Sets that can be written in this way, as mirror images of each other, are related by inversion. Now compare Sets 1 and 3. Again, they are written as mirror images of each other, and thus are related by inversion. Sets related by inversion have the same interval-class content; all three sets in Example 2–9 contain a 1, a 3, a 4, and no other intervals.

Figure 2–10 summarizes the relationships among these sets and uses arrows to indicate the relevant mappings. When sets related by inversion are written as mirror images of each other, the first note of one maps onto the last note of the other, the second onto the second-to-last, and so on. In comparing Set 1 and Set 3, for example, G maps onto B, G♯ onto B♭, and B onto G. The sets are related by T_nI, and to figure out the correct value of n, we take any note in one set and try to map it onto the corre-

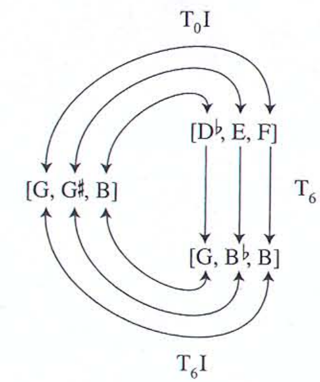

Figure 2–10

sponding note in the other. If we invert G, for example, we get F, and must transpose it at T_6 to map it onto B. Similarly, G♯ and B invert to E and C♯ and must be transposed at T_6 to map them onto B♭ and G. T_6I thus maps Set 1 onto Set 3. It also maps Set 3 onto Set 1—that's why the arrows point in both directions. By the same logic, Sets 1 and 2 are related at T_0I. Sets 2 and 3, both related by T_nI to Set 1, are related to each other by transposition, at T_6.

To invert a set, simply invert each member of the set in turn. For example, to apply the operation T_5I to the set [1,3,4,7], just apply T_5I to each integer in turn. Remembering to invert before we transpose, we get ((12 – 1) + 5, (12 – 3) + 5, (12 – 4) + 5, (12 – 7) + 5) = (4,2,1,10). Notice that if we write this new set in reverse order, [10,1,2,4], it will be in normal form. Generally when you invert a set in normal form, the result will be the normal form of the new set written backwards. There are many exceptions to this rule, however, so beware! When in doubt, use the step-by-step procedure outlined earlier in this chapter.

Index Number (sum)

The concept of *index number* offers a simpler way of inverting sets and of telling if two sets are inversionally related. The first two sets from Example 2–11, written using integer notation, are [7,8,11] and [1,4,5]. Remember that when we compared transpositionally related sets, we subtracted corresponding elements in each set and called that difference the transposition number. When comparing inversionally related sets, we will *add* corresponding elements and call that *sum* an index number. When two sets are related by transposition and are written so that they have the same intervals reading from left to right (that will always be true when they are written in normal form), the first note in one set corresponds to the first note in the other, the second to the second, and so on. When two sets are related by inversion and are written so that they are intervallic mirror images of each other (that will usually, but not always, be true when they are written in normal form), the first note in one set will correspond to the last note in the other, the second to the second-to-last, and so on. In comparing Set 1 and Set 2 from Example 2–11, the sum of the corresponding notes is 0 in each case (see Figure 2–11).

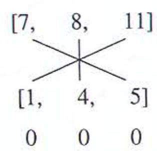

Figure 2–11

The sets are thus related at T_0I; 0 is the index number.

Figure 2–12 shows the first and third sets from Example 2–11: [7,8,11] and [7,10,11].

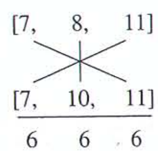

Figure 2–12

Again, the corresponding elements have a fixed sum, in this case 6. These two sets are related at T_6I. Each set is T_6I of the other. Any two sets in which the corresponding elements all have the same sum are related by inversion, and that sum is the index number.

Let us put this relationship in more general terms. If $T_nI(a) = b$, then $n = a + b$. In other words, inversionally related elements will sum to the index number. To find the index number for two elements, simply add them together. Conversely, to perform the operation T_nI on some pitch class, simply subtract it from n, since if $n = a + b$ then $a = n - b$. To perform the operation T_4I on [11,1,2,6], for example, subtract each element in turn from 4: $(4 - 11, 4 - 1, 4 - 2, 4 - 6) = (5,3,2,10)$, or [10,2,3,5] in normal form.

It may seem strange that addition plays such an important role in talking about T_nI. The idea of *subtracting* two notes, of figuring out the difference between them, makes clear musical sense. But what can it mean, say, to *add* an E to an F? Why is it that the sum of E and F is precisely the value of n that maps E onto F and F onto E under T_nI? To understand why, imagine the E and F on a clockface (Figure 2–13).

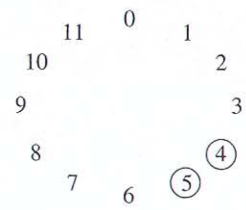

Figure 2–13

The E is at +4. If we invert it, we send it over to –4 (see Figure 2–14).

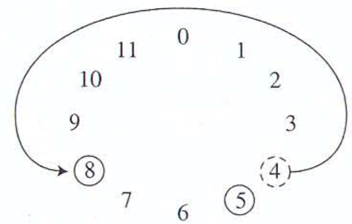

Figure 2–14

Now to get the inverted E to map onto the F we have to transpose it 4 (which gets us back to 0) plus 5 (which gets us to F). (See Figure 2–15.)

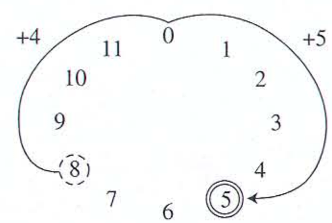

Figure 2–15

So T_9I maps E onto F. By the same logic, if we invert F, it goes from +5 to –5. Now to get it to map onto E, it has to be transposed at n = 9. So T_9I maps F onto E and E onto F.

Inversion (I^x_y)

There is another way of talking about inversion: I^x_y, where x and y are pitch classes that invert onto each other; they may be any pitch classes and they may be the same pitch class. Let's take I_B^G as an example—that's the inversion that maps G and B onto each other (see Figure 2–16).

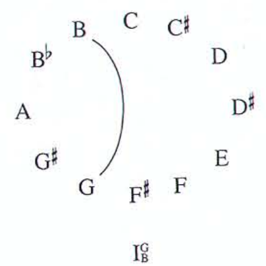

Figure 2–16

That same inversion also maps C onto F♯, C♯ onto F, D onto E, and D♯ and A onto themselves (see Figure 2–17).

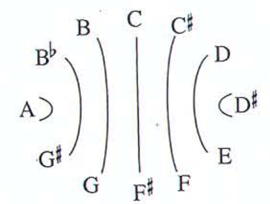

$I_A^A = I_{G\sharp}^{B\flat} = I_G^B = I_{F\sharp}^C = I_F^{C\sharp} = I_E^D = I_{D\sharp}^{D\sharp}$

Figure 2–17

By specifying any one mapped pair, we are simultaneously specifying all of the others. The inversion described in Figure 2–17 could thus be called I_A^A, $I_{G\sharp}^{B\flat}$, I_G^B, $I_{F\sharp}^C$, $I_F^{C\sharp}$, I_E^D, or $I_{D\sharp}^{D\sharp}$, and it would not matter which of the notes was written on top or on bottom. All of these labels are equally valid—which one we choose depends on the specific musical context.

Look back at Example 2–11 and at Sets 1 and 3. They are related by the inversion that maps the G and B onto each other, so a musically appropriate label in this case would be I_B^G. $I_{B\flat}^{G\sharp}$ would also be appropriate, because the same inversion that exchanges G and B also sends G♯ onto B♭. There are five other possible labels (as in Figure 2–17) but none seems musically relevant in this particular instance. One advantage of I_y^x over T_nI is that it does not emphasize inversion around C as 0, which may not have anything to do with the musical context.

In Example 2–12, the first line from Luciano Berio's *Sequenza for Solo Flute,* the highest six notes (excluding the grace notes) can be understood as two three-note

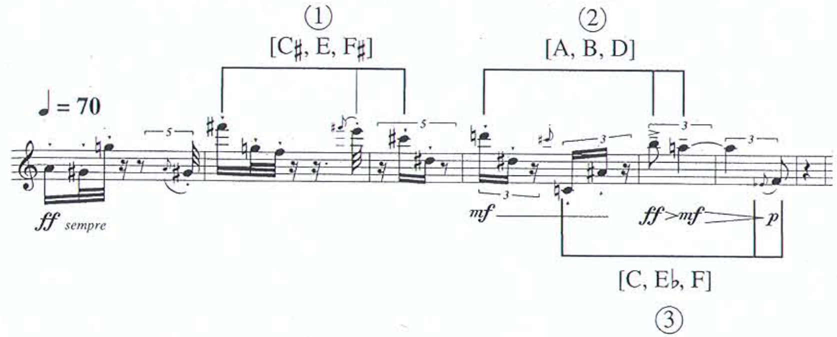

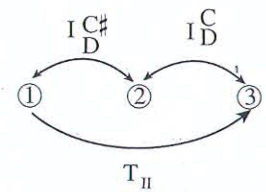

Example 2–12 Three equivalent pitch-class sets (Berio, *Sequenza for Solo Flute*).

sets related by inversion, and the lowest three notes (including the grace note) can be understood as related to one of these by inversion and the other by transposition. The first and second sets are related by inversion at $I_D^{C\sharp}$, a label that emphasizes the notes where they overlap registrally. The labels I_E^B and $I_A^{F\sharp}$, would also have been accurate, and all three labels are equivalent to T_3I (or I_3). The second and third sets are related by inversion at I_D^C, a label that emphasizes the notes at the registral extremes. The labels $I_{E\flat}^B$ and I_A^F would also have been accurate, and all three labels are equivalent to T_2I (or I_2). The first and third sets are related by transposition—the inversion of an inversion is always a transposition.

There are twelve possible inversions, each of which brings about a unique set of mappings (see Figure 2–18). For each inversion, the mappings are indicated with curved lines on the pitch-class clockface and the possible labels, in the form I_y^x, are listed beneath. It is easy to translate from the I_y^x model to the T_nI model of inversion, because $x + y = n$. To find the relevant index number, just add together any pair of mapped notes. The index numbers for each of the twelve inversions are given above the clockfaces in Figure 2–18.

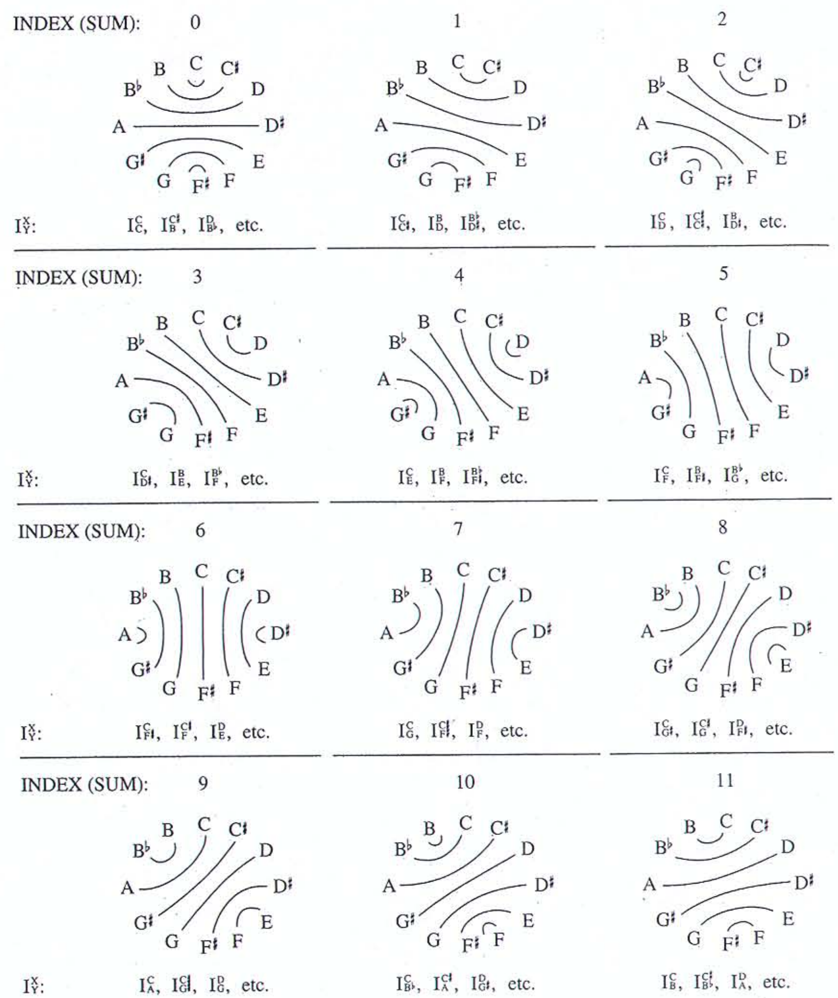

Figure 2–18

Set Class

Consider the collection of pitch-class sets in normal form shown in Figure 2–19. The first column begins with an arbitrarily chosen set, which is then transposed to each of the other eleven transposition levels. Thus, each of the twelve sets is related to the remaining eleven by transposition. The second column begins with an inversion of the set, then again transposes it systematically. In the second column as in the first,

[2,5,6]	[6,7,10]
[3,6,7]	[7,8,11]
[4,7,8]	[8,9,0]
[5,8,9]	[9,10,1]
[6,9,10]	[10,11,2]
[7,10,11]	[11,0,3]
[8,11,0]	[0,1,4]
[9,0,1]	[1,2,5]
[10,1,2]	[2,3,6]
[11,2,3]	[3,4,7]
[0,3,4]	[4,5,8]
[1,4,5]	[5,6,9]

Figure 2–19

each pitch-class set is related by transposition to the other eleven. The sets within each column are sometimes referred to as defining a *Tn-type*. A Tn-type is a class of sets that are all related to each other by transposition. Normally, a Tn-type contains twelve different sets. The twelve minor triads, for example, comprise a single Tn-type as do the twelve dominant-seventh chords.

Now consider all twenty-four of these sets together. Each of the twenty-four is related to all of the others by either transposition or inversion. They form a single, closely related family of sets. A family like this is called a *set class* (also referred to as a *Tn/TnI-type* or a *Tn/I set class*). [1,2,5], [5,6,9], [6,9,10], and twenty-one other pitch-class sets are all members of a single set class.

Normally, a set class will contain twenty-four members, like the one we just discussed. Some, however, have fewer than twenty-four distinct members. Consider the familiar diminished-seventh chord. If we write it out beginning in turn on each of the twelve pitch classes and then invert it and do the same, we quickly notice a good deal of duplication. If we eliminate all the duplicates, we find that this particular set class contains only three distinct members. Few sets are as redundant as this one (although one set, the whole-tone scale, is even more so). Most set classes contain twenty-four members; the rest have between two and twenty-four.

Set-class membership is an important part of post-tonal musical structure. There are literally thousands of pitch-class sets, but a much smaller number of set classes. Every pitch-class set belongs to a single set class. The sets in a set class are all related to each other by either T_n or T_nI. As a result, they all have the same interval-class content. By moving from set to set within a single set class, a composer can create a sense of coherent, directed musical movement.

The passage in Example 2–13 consists of a series of trichords in both hands of the piano part. Within the right-hand part, the trichords are all related by transposition; the same is true within the left-hand part. Each of the parts thus presents six trichords that belong to a distinct Tn-type. At the same time, each trichord in the right

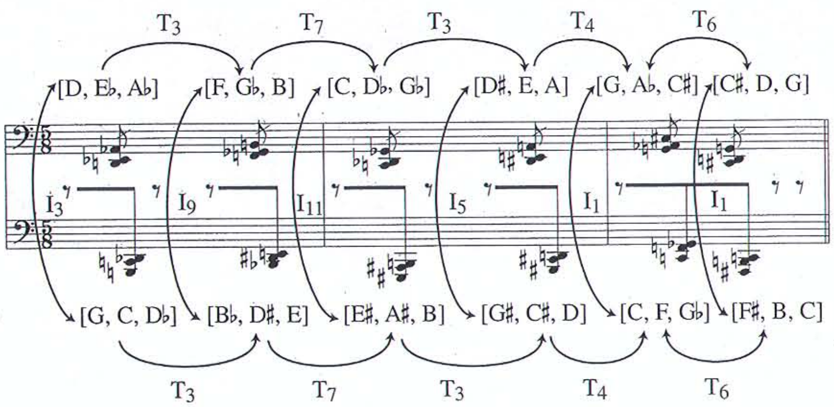

Example 2–13 Twelve trichords belonging to a single set class (Crumb, "Gargoyles," from *Makrokosmos,* Vol. 2, No. 8).

hand is related to each trichord in the left hand by inversion. As a result, all twelve trichords in the passage belong to the same set class (Tn/TnI-type).

Examples 2–14a and 2–14b contain passages that also involve motion from set to set, via transposition or inversion, within a single set class. In both cases, we can understand the music in terms of four equivalent sets, related in pairs by I_1, I_7, or T_6—the network in Example 2–14c applies to each, although the set class is different. Both passages thus draw on the unifying potential of set-class membership, and in strikingly similar ways.

A)

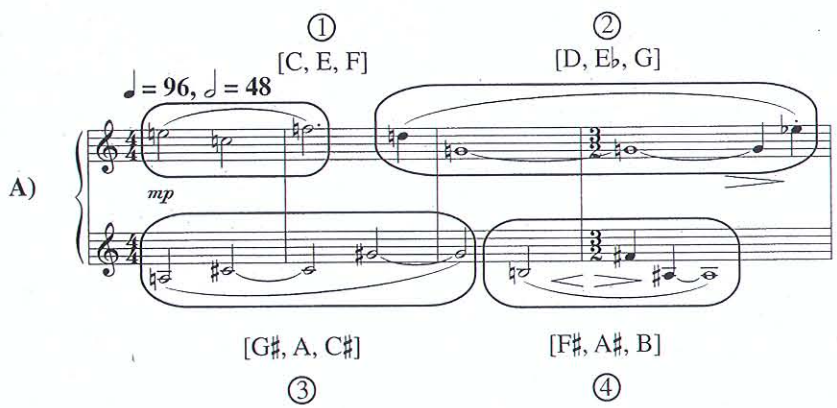

(continued)

B)

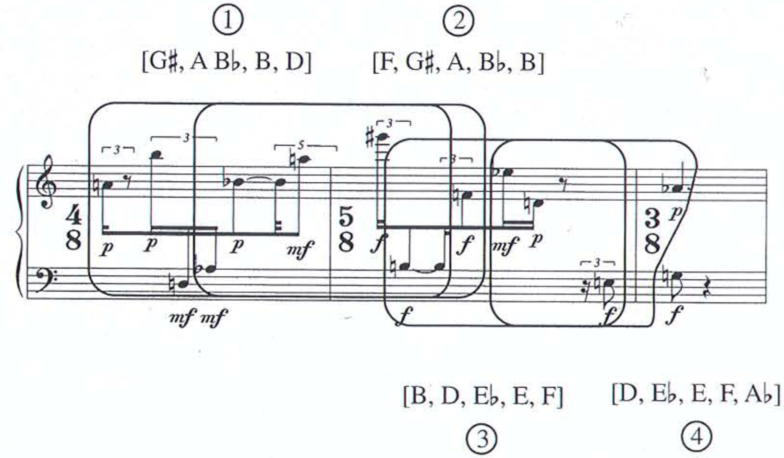

C)

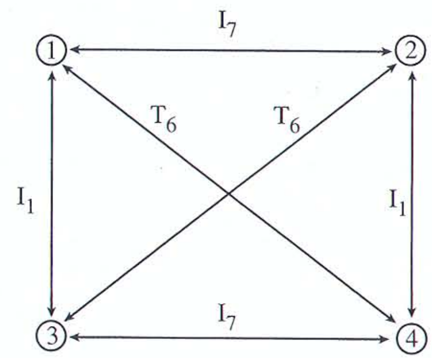

Example 2–14 Four sets belonging to the same set class: A) Wuorinen, *Twelve Short Pieces,* No. 3; B) Stockhausen, *Klavierstück* III; C) network analysis of both passages.

Let's look yet again at the opening section of Schoenberg's Piano Piece, Op. 11, No. 1, to see how a progression involving members of the same set class can create one taut strand in the larger compositional fabric (see Example 2–15). In the first three measures, a single continuous melody descends from its high point on B. In measures 4–8, the melody is reduced to a two-note fragment that reaches up to G three times. In measures 9–11, the opening melody returns in a varied form with a high point of G♯. These three notes, B–G–G♯, are separated in time but associated as contour high points. These are the same pitch classes as the first three notes in the piece and the sustained notes in measures 4–5.

Each note in this large-scale statement is also part of at least one small-scale statement of a member of the same set class. The B in measure 1 is part of the collection B–G♯–G. In measures 4–5, the G is not only part of the sustained chord (G–G♯–B) but also part of the registral grouping G–B♭–B. In measure 10, the G♯ is

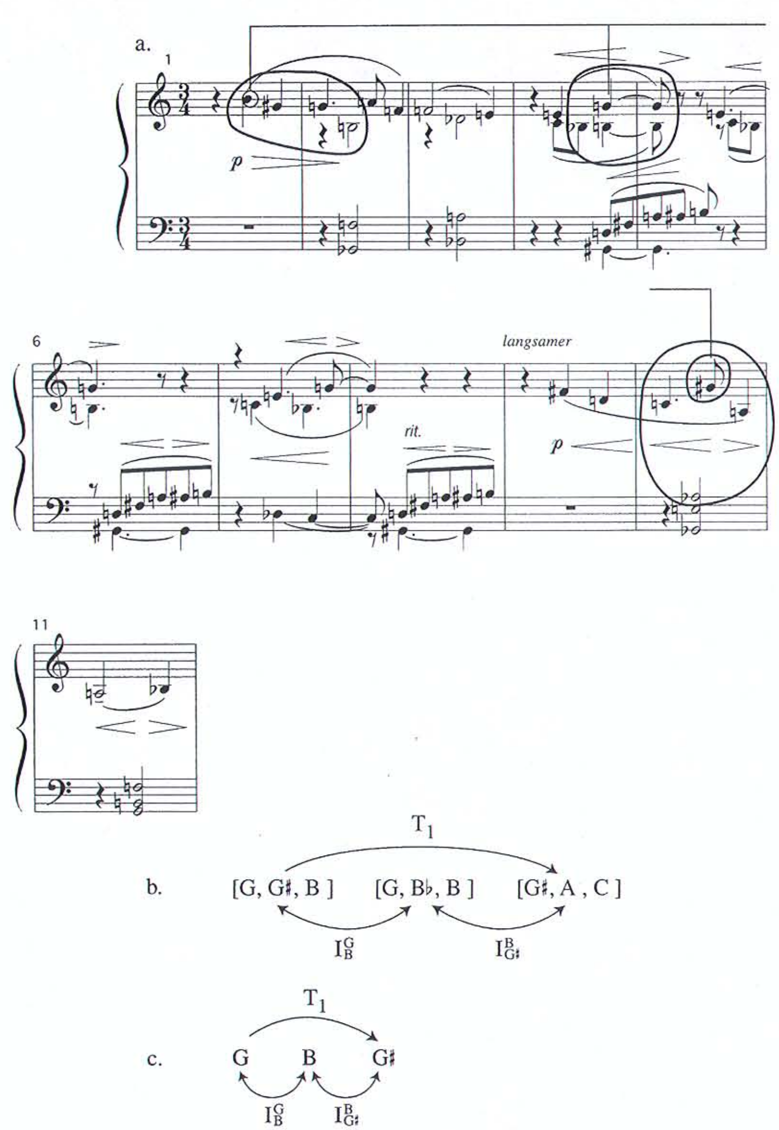

Example 2–15 Progression among members of the same set class.

part of the collection C–G♯–A. These three sets are circled in Example 2–15a and written in normal form in Example 2–15b. The operations that lead from set to set are identified. Example 2–15c shows that the same operations that lead from set to set also can be understood to lead from note to note within the first set. In that sense,

relationships embodied in the opening three-note motive are composed-out as that motive is transposed and inverted over the course of the passage.

There are many occurrences of other members of the same set class in this passage, including the chord in measure 3 [A,B♭,D♭], the highest three notes in measure 3 [D♭,E,F], the mid-range notes in measure 4 [B♭,B,D] and the inner notes in the five-note figure in the tenor in measures 4–5 [F♯,A,A♯]. Indeed, the passage is virtually saturated with occurrences of members of this set class. It occurs as a melodic fragment, as a chord, and as a combination of melody and chord. It is articulated by register and, over a large span, by contour. An entire network of musical associations radiates out from the opening three-note melodic figure. Some of the later statements have the same pitch content, some the same pitch-class content. Some are related by transposition, some by inversion. All are members of the same set class. As in tonal music, but with even greater intensity, an initial musical idea grows and develops as the music proceeds. The mere presence of many members of a single set class guarantees a certain kind of sonic unity. But we often will be more interested in the ways in which the music moves from set to set within a set class than in mere set-class membership.

Prime Form

There are two standard ways of naming set classes. The first was devised by the theorist Allen Forte, who pioneered the theory of pitch-class sets. On his well-known list of set classes, he identifies each with a pair of numbers separated by a dash (e.g., 3–4). The first number tells the number of pitch classes in the set. The second number gives the position of the set on Forte's list. Set class 3–4, for example, is the fourth set on Forte's list of three-note sets. Forte's set names are widely used and appear in Appendix 1.

The second common way of identifying set classes is to look at all of the members of the set class, select the one with the "most normal" of normal forms, and use that to name the set class as a whole. This optimal form, called the *prime form,* begins with 0 and is most packed to the left. Of the members of the set class shown in Figure 2–19, two begin with 0: 034 and 014. Of these, (014) is the most packed to the left and is thus the prime form. Those twenty-four sets are all members of the set class with prime form (014). More familiarly, we say that each of those sets "is a (014)." In this book, prime forms will be written in parentheses with no commas separating the elements. T and E will stand for 10 and 11 in this compact format. A set class will generally be identified by both its Forte name and its prime form, and set class will usually be abbreviated as sc. Thus, the sets in Example 2–13 are all members of sc3–5 (016); the sets in Example 2–14A are members of sc3–4 (015); the sets in Example 2–14B are members of sc5–4 (01236); and the sets in Example 2–15 are members of sc3–3 (014). Often, we will omit the Forte-name, in which case these set classes would be identified as sc(016), sc(015), sc(01236), and sc(014).

To identify the set class to which some pitch-class set belongs, you will have to find the prime form of the set class. That process is usually referred to as "putting a set in prime form." Here is how to do it:

1. Put the set into normal form. (Let's take [1,5,6,7] as an example.)
2. Transpose the set so that the first element is 0. (If we transpose [1,5,6,7] by T_{11}, we get [0,4,5,6].)
3. Invert the set and repeat steps 1 and 2. ([1,5,6,7] inverts to [11,7,6,5]. The normal form of that set is [5,6,7,11]. If that set is transposed at T_7, we get [0,1,2,6].)
4. Compare the results of step 2 and step 3; whichever is more packed to the left is the prime form. ([0,1,2,6] is more packed to the left than [0,4,5,6], so (0126) is the prime form of the set class of which [1,5,6,7], our example, is a member.)

As with normal form, it will often be possible to determine the prime form just by inspecting a set displayed around a pitch-class clockface. Find the widest gap between the pitch classes. Assign zero to the note at the end of the gap and read off a possible prime form clockwise. Then assign zero to the note at the beginning of the gap and read off another possible prime form counterclockwise. (If there are two gaps of the same size, choose the one that has another relatively big gap right next to it.) Whichever of these potential prime forms has fewer big integers is the true prime form. Figure 2–20 illustrates with the four sets we used in Example 2–5 to determine normal form. It is basically a matter of visualization, and it will get easier with practice.

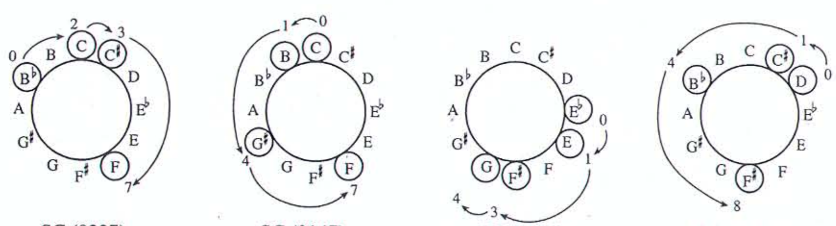

Figure 2–20

In Appendix 1, you will find a list of set classes showing the prime form of each. If you think you have put a set in prime form but you can't find it on the list, you have done something wrong. Notice, in Appendix 1, how few prime forms (set classes) there are. With our twelve pitch classes, it is possible to construct 220 different trichords (three-member sets). However, these different trichords can be grouped into just twelve different trichordal set classes. Similarly, there are only twenty-nine tetrachord classes (four-member sets), thirty-eight pentachord classes (five-member sets), and fifty hexachord classes (six-member sets). We will defer discussion of sets with more than six elements until later.

The list of set classes in Appendix 1 is constructed so as to make a great deal of useful information readily available. Any sonority of between three and nine elements is a member of one of the set classes listed here. In the first column, you will see a list of prime forms, arranged in ascending order. The second column gives Forte's name for each set class. The third column contains the interval-class vector for the set class. (This is the interval-class vector for every member of the set class, since interval content is not changed by transposition or inversion.) In the fourth column are two numbers separated by a comma; these numbers measure the transpositional and inversional symmetry of the set class—we will discuss these concepts later. Across from each trichord, tetrachord, and pentachord, and some of the hexachords, is another set with all of its relevant information in the reverse order. We will discuss these larger sets later.

Segmentation and Analysis

In the post-tonal music discussed in this book, coherence is often created by relationships among sets within a set class. It is possible to hear pathways through the music as one or more sets are transposed and inverted in purposeful, directed ways. Often, we find that there is not one single best way to hear our way through a piece; rather, our hearings often need to be multiple, as the different paths intersect, diverge, or run parallel to each other. To use a different metaphor, post-tonal music is often like a rich and varied fabric, comprised of many different strands. As we try to comprehend the music, it is our task to tease out the strands for inspection, and then to see how they combine to create the larger fabric.

One of our main analytical tasks, then, is to find the principal sets and show how they are transposed and inverted. But how do you know which sets are the important ones? The answer is that you cannot know in advance. You have to enter the world of the piece—listening, playing, and singing—until you get a sense of which musical ideas are fundamental and recurring. In the process, you will find yourself moving around a familiar kind of conceptual circle. You can't know what the main ideas are until you see them recur; but you can't find recurrences until you know what the main ideas are. The only practical solution is to poke around in the piece, proposing and testing hypotheses as you go. In the process, you will be considering many different *segmentations* of the music, that is, ways of carving it up into meaningful musical groupings.

When you have identified what you think may be a significant musical idea, then look carefully, thoroughly, and imaginatively for its transposed or inverted recurrences. Here are some places to look (this list is not exhaustive!):

1. In a melodic line, consider all of the melodic segments. For example, if the melody is six notes long, then notes 1–2–3, 2–3–4, 3–4–5, and 4–5–6 are all viable three-note groupings. Some of these groupings may span across rests or phrasing boundaries, and that is okay. A rich interaction between phrase structure and set-class structure is a familiar feature of post-tonal music.
2. Harmonically, don't restrict yourself just to chords where all of the notes are attacked at the same time. Rather, consider all of the *simultaneities,* that is, the notes sounding simultaneously at any particular point. Move through the music like a cursor across a page, considering all of the notes sounding at each moment.