Music Perception
Winter 2002, Vol. 20, No. 2, 117–149

Parallelism as a Factor in Metrical Analysis

DAVID TEMPERLEY & CHRISTOPHER BARTLETTE
Eastman School of Music

A model is proposed of the effect of parallelism on meter. It is well-known that repeated patterns of pitch and rhythm can affect the perception of metrical structure. However, few attempts have been made either to define parallelism precisely or to characterize its effect on metrical analysis. The basic idea of the current model is that a repeated melodic pattern favors a metrical structure in which beats are placed at parallel points in each occurrence of the pattern. By this view, parallelism affects the *period* of the metrical structure (the distance between beats) rather than the *phase* (exactly where the beats occur). This model is implemented and incorporated into the metrical program of D. Temperley and D. Sleator (1999). Several examples of the model's output are presented; we examine problems with the model and discuss possible solutions.

Received September 24, 2001, accepted June 5, 2002

One of the most widely studied problems in music cognition is the problem of metrical analysis or "beat-finding": inferring the metrical structure from a piece of music. A large number of studies have focused on this issue, reflecting a variety of approaches, aims, and ways of defining the problem (Allen & Dannenberg, 1990; Chafe, Mont-Reynaud & Rush, 1982; Desain & Honing, 1992; Large & Kolen, 1994; Lee, 1991; Longuet-Higgins & Steedman, 1971; Parncutt, 1994; Povel & Essens, 1985; Rosenthal, 1992; Tanguiane, 1993; Temperley & Sleator, 1999; Todd, 1994). Some researchers have assumed an input that is completely regular or "quantized" (Longuet-Higgins & Steedman, 1971; Povel & Essens, 1985); others have sought to accommodate the fluctuations in pulse that are characteristic of performed music (Large & Kolen, 1994; Rosenthal, 1992).[1] Some researchers have sought to model the process of inducing a beat as a piece unfolds in left-to-

1. The great majority of metrical models operate on "note" information, although a few have attempted to handle actual acoustic input (Tanguiane, 1993; Todd, 1994).

Address correspondence to David Temperley, Eastman School of Music, 26 Gibbs St., Rochester, NY 14604. (e-mail: dtemp@theory.esm.rochester.edu)
ISSN: 0730-7829. Send requests for permission to reprint to Rights and Permissions, University of California Press, 2000 Center St., Ste. 303, Berkeley, CA 94704-1223.

right fashion (Lee, 1991; Longuet-Higgins & Steedman, 1971), while others have focused on producing the preferred metrical structure given the entire piece (Rosenthal, 1992; Temperley & Sleator, 1999). Some models have produced just a single level of beats, while others have generated a complete "metrical structure" of several levels. In terms of approach, some studies have worked within the rule-based paradigm of classical artificial intelligence (Lee, 1991; Longuet-Higgins & Steedman, 1971); others have applied connectionist techniques (Desain & Honing, 1992; Large & Kolen, 1994), and still others have adopted a "preference-rule" approach (the current study is in this category, as will be explained later). At a still more fundamental level, perhaps, we might distinguish between psychologically oriented studies, seeking to model how human listeners perform metrical analysis (and perhaps considering psychological evidence bearing on this question), and studies that approach the subject from an artificial-intelligence or engineering viewpoint, seeking to solve the problem by whatever means seems most elegant and effective—though of course these two aims are closely related and often convergent.[2]

Despite these differences, there is considerable agreement on one important issue, namely, the factors or criteria that are involved in the determination of metrical structure. Three criteria are universally accepted and are incorporated in one way or another in virtually all models of meter. First, there is a preference for beats to coincide with notes (or, more precisely, onsets of notes). Second, there is a preference for aligning beats with longer notes. (For psychological evidence on these points, see Povel & Essens, 1985.) Let us consider the problem of deriving just a single level of beats (Figure 1). We would favor hearing A over hearing B, because the former aligns beats with more notes; we favor hearing A over hearing C, although every beat coincides with a note in both hearings, because the notes hit by hearing A are (on balance) longer than those hit by hearing C. Third, there is a preference for regularity of beats—whether this regularity is assumed to be perfect (in the case of models assuming "quantized" input) or imperfect, as in the case of more naturalistic models. In Figure 1, for example, we would strongly resist an analysis such as hearing D, which began with beats on every third sixteenth-note beat and then switched to every fourth one. Although developing a model of meter incorporating these three simple criteria might seem like a straightforward task, in fact it is not; finding the correct operational definition of each criterion, and the optimal balance between them, has proven to be a formidable challenge (as indicated by the large amount of work that has been devoted to it). The difficulties only increase when the problem is expanded to include polyphonic as well as

2. Although some studies embrace either the psychological goal (Parncutt, 1994; Povel & Essens, 1985) or the engineering goal (Chafe et al., 1982), others embrace both goals (Rosenthal, 1992) or remain noncommittal between them (Longuet-Higgins & Steedman, 1971). Our own stance on this issue will be clarified later.

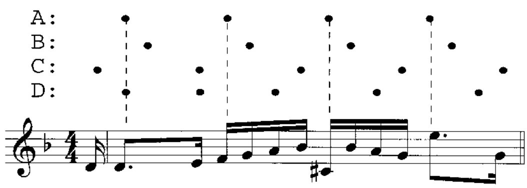

Fig. 1. Bach, Partita for Violin in D minor, Allemande, measure 1, showing four possible metrical analyses.

monophonic music—something that very few models of meter have done (the model of Temperley and Sleator, 1999, is one exception).

Besides the three criteria of note onsets, note length, and regularity, students of metrical analysis have identified several other factors that are involved in the perception of meter. One is grouping: when a series of notes form a group or phrase, there is a tendency to hear the strongest beat near the beginning of the group (Povel & Essens, 1985). In Figure 1, for example, if we take as given the notated quarter-note beat (Analysis A), we tend to hear the odd-numbered beats (the first and third) as stronger than the even-numbered ones (the second and fourth); this is because the first beat is closer than the second to the beginning of the phrase. Another factor is harmony: we prefer to hear beats at points of harmonic change (Temperley, 2001). In Figure 2, the fairly clear change of harmony at each notated downbeat is a strong factor favoring these points as strong beats. Both grouping and harmony are problematic because of their interactive relationship with meter; meter is a factor in both grouping and harmonic structure as well as being influenced by them (see Temperley, 2001, for discussion). Intensity (loudness) and melodic accent (i.e. factors of melodic shape or contour) also play some role in metrical analysis, although these factors appear to be fairly minor.[3]

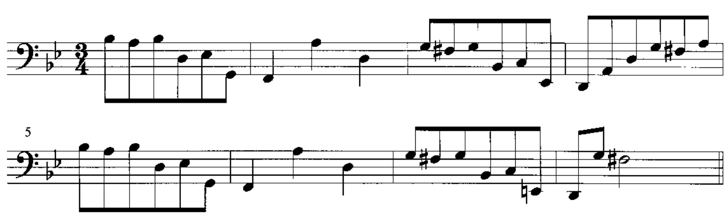

Fig. 2. Bach, Suite for Violoncello in G major, Minuet II, measures 1–8.

3. Regarding intensity, a study of performance expression by Drake and Palmer (1993) suggests that the difference in loudness between metrically strong and weak events tends to be quite small. Regarding melodic accent, several definitions of this concept have been proposed; however, in a study of these, Huron and Royal (1996) found that none had more than a small correlation with metrical strength as indicated in musical scores.

A further factor in metrical analysis, and the subject of the current study, is parallelism: the effect of repeated patterns on the perception of meter. A clear and compelling example of the role of parallelism in meter is the opening of Beethoven's *Moonlight* Sonata (Op. 27, No. 2), shown in Figure 3. The obvious hearing of this pattern features a strong beat on the first chord, followed by beats at every third subsequent note. Although the strong beat on the first chord could be explained by the note-onset criterion (favoring beats at points of several note onsets as opposed to just one), what accounts for the hearing of beats on subsequent G♯'s? One might attribute this to some inherent feature of the pattern, for example, the fact that the G♯'s are lowest in register. However, if one imagines the three-note pattern repeating indefinitely (without the left-hand part), it is quite easy to hear it with the accents on the E's, or even on the C♯'s. Rather, the metrical accentuation of the G♯'s appears to be due to the fact that a melodic pattern is present here; given that the first instance of the pattern is clearly accented in a certain way, we tend to impose the same accentuation on subsequent instances as well.

In examining the effect of parallelism on meter, we must be careful to isolate it from other factors. Consider Figure 4; we certainly hear a parallelism here between the first three quarter-note beats and the second three, and one might suppose that this favors a hearing of triple meter. (Note that in this example, the melodic pattern is not repeated exactly, but is transposed down a step—a frequent occurrence in cases of parallelism, as we will see.) In this case, however, the triple meter hearing can quite easily be accounted for by other factors: in particular, the large chords on the first and fourth beats encourage us to hear them as metrically strong. The effect of parallelism on meter is demonstrated most clearly in cases, like the *Moon-*

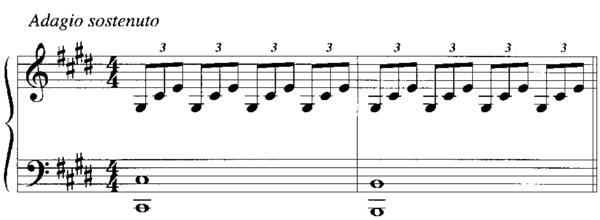

Fig. 3. Beethoven, Sonata Op. 27 No. 2, I, measures 1–2.

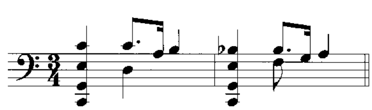

Fig. 4. Bach, Suite for Violoncello in C major, Sarabande, measures 1–2.

light Sonata, where the internal structure of the pattern itself does not favor any position as metrically strong.

In what follows, we will propose a model of parallelism and its role in metrical analysis. Although our methods in this study are computational, our larger interest is cognitive: by approaching parallelism from a computational viewpoint, we hope to gain insight into how it might be represented psychologically. (We will return to this larger issue in the final section of the article.) Our concern will be with "common-practice" music—Western art music of the 18th and 19th centuries; whether parallelism and its effect on meter differ significantly in other musical idioms is a question beyond our scope. We will begin by considering some other research relating to parallelism and meter.

Earlier Models of Parallelism and Meter

Among the computational studies discussed earlier, very few have attempted to model the effect of parallelism on meter; indeed, little attention has been given to the modeling of parallelism (repeated patterns within a piece) in general. However, two models do require discussion. The model of Steedman (1977) builds on the earlier metrical algorithm of Longuet-Higgins and Steedman (1971). This model analyzes a piece in a left-to-right fashion; it assumes a "quantized" input (in which all durations are represented as multiples of a low-level beat). It starts by generating a single metrical level (often the metrical level implied by the length of the first note, although not always), and then generates higher levels based on certain cues—for example, a long-short-short ("dactyl") pattern tends to imply a strong beat at its beginning.

In Steedman's extended version of this model, repetition is taken into account in the following way. If a pattern beginning on a beat B1 is then repeated beginning on a following beat B2, a new metrical level is established with beats at B1 and B2. Consider Figure 5 (one of Steedman's examples). The model first establishes the sixteenth-note level (as in Longuet-Higgins & Steedman's model). Proceeding to the right, when it gets to the sixth note it realizes that notes 4–6 are parallel to notes 1–3 (i.e., similar in intervallic pattern); this leads it to establish a new, dotted-eighth metrical level with beats at note 1 and note 4. The model then adopts this dotted-

Fig. 5. Bach, *Well-Tempered Clavier* Book II, Fugue No. 4, measures 1–2.

eighth metrical level as its basic unit; from then on, it considers and compares only "bites" beginning at each dotted-eighth beat. It finds that notes 13–15 are parallel to notes 1–3, leading to a higher (dotted-whole-note) metrical level with beats at notes 1 and 13.

An important feature of the Longuet-Higgins/Steedman model (and Steedman's extended version) is what Steedman calls the "principle of consistency": once a metrical level is established, it cannot be abandoned. Essentially, this assumes that syncopations—for example, long notes on weak beats—will never occur until the correct meter has been firmly established. Although this principle appears to hold up well for Bach fugues, it does not always apply; certainly there are cases where we seem to "revise" our initial metrical analysis of the first few events of a piece based on subsequent evidence (see Jackendoff, 1991, and Temperley, 2001, for discussion and examples). Another problem should also be mentioned, relating specifically to Steedman's parallelism model. Steedman correctly observes that, given a repetition of a pattern, we tend to assume that "the metrical accent will fall at the same point in both figure and repeat" (1977, p. 560)—this is the same point made earlier in connection with the *Moonlight* Sonata. This implies a metrical level whose beat interval is the same as the distance between corresponding points in the figure and the repeat. However, Steedman then goes further to assume that beats are located at the *beginning* of the figure and the repeat. In effect, Steedman suggests that the parallelism determines not only the period but also the phase of the metrical level. Figure 6 shows a case where this assumption is problematic: we clearly hear the repeated rhythmic and intervallic pattern in the melody (marked as motive X), and this is surely a factor in our perception of the meter. Yet each occurrence of the pattern begins just after a quarter-note beat, rather than on the beat. It seems to us that Steedman's model would incorrectly label the beginning of each X motive as a strong beat, thus correctly inferring a quarter-note level of meter but misidentifying the phase.[4] An alternative solution is to assume that parallelism determines only the period of a metrical level (the distance between beats), not the phase (where exactly the

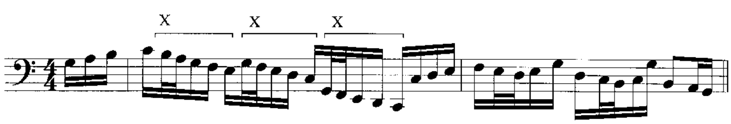

Fig. 6. Bach, Suite for Violoncello in C major, Allemande, measures 1–2.

4. It is also possible that Steedman's model would have already committed to the correct eighth-note level by this point (given the dactyl in notes 4, 5, and 6); it would therefore not even consider pattern X, as this begins on a weak sixteenth-note beat. Of course, this result is not satisfactory either.

beats occur); the phase, rather, depends on the location of "phenomenal accents"—long notes, loud notes, changes of harmony, and so on. In the case of the *Moonlight* Sonata, for example (Figure 3), the three-note melodic pattern suggests only that strong beats are three notes apart; the exact placement of these beats is determined in this case by the left-hand octave, adding a strong accent to the first note of the first pattern instance and thus (by parallelism) of all subsequent instances as well.[5]

A further study deserving discussion is Mont-Reynaud and Goldstein's (1985) model of rhythmic pattern recognition. The focus of this study is not on the role of parallelism in meter (though it does discuss this role in a general way), but rather, on parallelism itself: a model is proposed which searches for melodic patterns in a piece.[6] Given a piece represented as a sequence of symbols (notes), the model looks at each range of notes within the sequence, and compares it, essentially, with every other range of the same size. It "grows" patterns recursively, first looking for two-note patterns, then expanding to three-note patterns where possible (on the reasoning that a repeating three-note pattern ABC must be built on a two-note pattern AB), and so on. An additional step is then needed of eliminating redundant patterns. In the case above, for example, the model would first find ABC and BC independently and would then realize on a second pass that BC is redundant. The authors then consider the question of elaboration—how to identify cases where one pattern is an elaborated version of another.

Mont-Reynaud and Goldstein's method of pattern recognition is more exhaustive than Steedman's, and it seems to offer an effective approach to finding all significant patterns that occur within a piece. The problem of redundant patterns is a difficult one and greatly increases the amount of computation necessary. If a pattern exists of symbols N_1-N_j, the program will also independently find N_2-N_j, N_3-N_j,...N_{j-1}-N_j, then eliminating all but the first (and longest) pattern. Another problem with this approach arises with self-overlapping patterns such as ABCABCABCABC. Presumably Mont-Reynaud and Goldstein's model would identify ABC as a repeated

5. The distinction between period and phase is nicely illustrated by cases where the period of the meter is clear but the phase is not. One famous example of this is the opening of Beethoven's Sonata Op. 14 No. 2, I; a repeated pattern in both melody and accompaniment makes it clear that some kind of half-note beat level is present, but it is not clear where the beats occur, leaving the passage metrically ambiguous.

Another metrical model that appears to have incorporated parallelism is that of Rosenthal (1992). Rosenthal says that his model prefers analyses that "subdivide the performance in a motivically plausible manner," such that "a recurring motive will occupy the same position in the measures in which it occurs" (1992, p. 70). However, no explanation is given of how this is formalized or implemented.

6. The authors focus mainly on rhythmic patterns, but the possibility of applying their method to pitch patterns is also discussed.

Two other interesting models of melodic pattern recognition, Rolland (1999) and Cambouropoulos (2001), came to our attention too late to be included in this discussion.

pattern, but also BCA and CAB—again leading to considerable redundancy. This redundancy might be useful for some purposes; for example, when it is necessary to explicitly identify what the "motives" are in a piece, then all three possibilities (ABC, BCA, and CAB) would have to be considered and evaluated. For the purpose of modeling the effect of parallelism on meter, however, all of this may be unnecessary. What is important here is simply that there is a repetition at a distance of three notes (it doesn't matter which note the pattern is thought to start on), implying a metrical level with the same period. In the following model, we will propose a way of identifying such repetitions which does not depend on the explicit identification of motives.

Turning to the experimental psychological literature, we find little that bears directly on the role of parallelism in metrical analysis. An important study concerning parallelism itself is that of Deutsch (1980). Trained listeners were played melodies and asked to write them down; some melodies featured repeating three-note or four-note patterns (e.g., Figure 7a), while others did not (e.g., Figure 7b). Listeners were able to learn the patterned melodies much more easily than the unpatterned ones. Deutsch also found that rhythm plays a role in the perception of parallelism; a pitch pattern can be learned much more easily when combined with a rhythmic pattern that supports the pitch pattern rather than conflicting with it. (Similar results have also been obtained by Handel, 1973, and Boltz and Jones, 1986.) Such studies certainly show the psychological reality of parallelism and also show an important perceptual role for it: it aids in the encoding and memorization of music. However, they tell us little about the role of parallelism in metrical analysis. (We would guess that parallelism *did* affect metrical analysis in these cases; undoubtedly, listeners tended to infer a metrical structure that was parallel with the repeated pattern, for example, inferring strong beats on every third note in Figure 7a. But this issue is not directly addressed by these studies.) Thus, while studies by Deutsch and others provide strong evidence for the psychological reality of parallelism itself, the role of parallelism in metrical analysis remains to be explored.[7]

Fig. 7. Melodies that contain repeated patterns (A) are more easily learned than melodies that do not (B). From Deutsch (1980).

7. Deutsch and Feroe (1981) also propose a model for the encoding of repeated patterns in music; however, they offer no proposal for how such encodings are formed—or, for example, why one encoding is preferred over another.

A Preference Rule Model of Meter

The current proposal builds on a computational model of meter that is described elsewhere (Temperley, 2001; Temperley & Sleator, 1999). This model is designed to handle unquantized (performance-generated) input as well as quantized; it is also designed for polyphonic input. The input assumed is a list of notes, giving the ontime and offtime (in milliseconds) and pitch of each note. The program begins by quantizing all event ontimes and offtimes to time-points called "pips," spaced 35 ms apart; only pips are considered as possible beat locations.[8]

The output of the system is a framework of levels of beats. The system begins by generating a main beat or "tactus" level; it then generates two lower and two upper levels. (We adopt the convention of calling the tactus level 2, upper levels 3 and 4, and lower levels 1 and 0.) The system's output consists of a note list (containing the notes as originally inputted, except quantized to pips) and a beat list, with each beat having a time point and a level number (the highest level at which that beat is present; see Figure 8).

The operation of the model is based on "preference rules"—an approach to musical analysis first proposed by Lerdahl and Jackendoff (1983).[9] In a preference rule system, a large set of possible ("well-formed") analyses are considered and evaluated by certain criteria; the preferred analysis is the one best satisfying the criteria. In the current case, the well-formedness constraints state that for each pair of adjacent beat levels, every beat at the higher level must also be a beat at the lower one, and exactly one or two lower-level beats must elapse between each pair of higher-level beats. For the tactus level (since it is generated first), the only well-formedness constraint is that intervals between adjacent beats are limited to a range of 400 ms to 1600 ms.

The system has three main preference rules, which are exactly the principles described earlier as the essential criteria of metrical analysis.

> Event Rule: Prefer to locate beats (especially strong beats) at onsets of events; the more events at a timepoint, the better a beat location it is.
>
> Length Rule: Prefer to locate beats at the onsets of long events.

8. Like other parameters of the program, the value of 35 ms was found to be optimal through trial and error. Quantization to pips is desirable from a computational viewpoint (because the speed of the program depends on the number of beat locations that must be considered), but there is a more musically substantive reason for it as well. In live performance, notes that are intended (and understood) to be exactly simultaneous are often not played that way; by quantizing to pips, the program attempts to make such notes exactly simultaneous.

9. Lerdahl and Jackendoff's model of meter incorporates parallelism as a factor; however, they make no attempt to define parallelism rigorously, or to precisely characterize its effect on meter.

```
Note      0     225   55
Note    225     450   57
Note    450     675   59
Note    675     900   60
Note    900    1012   59
Note   1012    1125   57
Note   1125    1350   55
...
```

Meter program (reads in note list, adds beats)

```
Note      0     210   55
Note    210     455   57
Note    455     665   59
Note    665     900   60
...
Beat      0   1
Beat    105   0
Beat    210   2
Beat    315   0
...
```

Fig. 8. The input and output of the meter program. "Note" statements consist of an ontime and an offtime (both in milliseconds) and a pitch; "Beat" statements consist of a time point (in milliseconds) and a level number.

In calculating the "length" of a note, the program uses what Temperley and Sleator (1999) call its "registral interonset interval": the time interval to the onset of the next note within the same register, which is defined as a range of 9 semitones above or below. The length of a note is then defined as the maximum of its actual duration and its registral interonset interval.

> Regularity Rule: Prefer for beats at each level to be roughly equally spaced.

At the tactus level, regularity is measured by comparing each beat interval with the previous beat interval and imposing a penalty proportional to the difference between them (the enforcement of regularity at other levels will be discussed later). Although gross irregularities are highly penalized by the regularity rule, small tempo fluctuations of the kind that often occur in performance are penalized only very slightly.

Consider just the tactus level. Any analysis at this level (that is to say, any row of beats) can be evaluated by using the three preference rules just given. Each rule assigns a numerical score indicating how well it is satisfied by the analysis; the preferred analysis overall is the one with the highest total score. For the event rule and length rule, the score reflects the number of events aligned with beats under the analysis as well as their lengths (adjusting for the fact that analyses with more beats will hit more events). Consider Figure 1 again: Analysis A is preferred over Analysis B, as it aligns more beats with event onsets. (The event rule also favors Analysis A over many other possible analyses—not shown—in which few or none of the beats coincide with event onsets.) Analysis A and Analysis C are equally preferred by the event rule, but the length rule favors Analysis A because the events coinciding with beats in Analysis A are longer that in Analysis C. Finally, Analysis D is penalized by the regularity rule, as the beats are not regular. Of the four analyses above, then, Analysis A would be preferred.

Because of the regularity rule, the optimal analysis for a short segment of a piece may depend on what happens elsewhere (either before or afterwards). This is why the model must consider complete analyses of a piece, rather than evaluating small segments in isolation. However, the number of possible analyses grows exponentially with the number of possible beat locations in the piece. A more intelligent search procedure is needed to find the best-scoring analysis. This search procedure is not our primary concern here, though we will return to it briefly later; for further discussion, see Temperley (2001).

At levels above and below the tactus, the event and length rule operate in exactly the same way. The regularity rule operates rather differently; here, penalties are imposed for changes in relationships between levels, rather than changes in absolute beat intervals. For example, if one pair of tactus beats is divided duply at the next level down, and the next triply, a penalty is imposed.

The model also incorporates two other weak criteria. In selecting the highest level, the model prefers to locate the first highest-level beat near the beginning of the piece (on the first beat of the next level down, rather than the second or third). Second, the model slightly favors duple over triple relationships between levels. The model also has the capability of using harmonic information to guide metrical analysis (preferring beats at points of harmonic change), but this must be provided in the input; this feature is not used in the "standard" version of the program assumed here.

The model as just described was tested on a systematically chosen corpus of polyphonic excerpts from the Kostka-Payne theory workbook (Kostka & Payne, 1995). Using quantized inputs generated precisely from a score, the model labeled tactus level beats correctly in 94% of measures and achieved slightly lower rates (86-94%) for all other levels. Using inputs generated from performances by a skilled pianist (a doctoral student in piano at Ohio State University) on a MIDI keyboard, the program labeled tactus beats correctly in 85% of measures and achieved rates of 71–86% for other levels. (See Temperley, 2001, for further details about these tests.)

The tests just described indicate that the model, while showing promise, has considerable room for improvement. The limitations of the model can be demonstrated informally as well; examples can easily be found in which the criteria available to the model are clearly insufficient to yield the correct analysis. The *Moonlight* Sonata is one case (Figure 3). The triple meter of this passage is immediately evident, as opposed to—for example—a duple meter assigning strong beats to every second note. However, it can be seen that none of the rules of the model (as stated earlier) confer an advantage to the triple analysis over the duple one. (The event rule favors a beat on the first note, but says nothing about where subsequent beats occur; the regularity rule favors some kind of regular pattern of beats, but is indifferent between a duple or triple grouping.) Figure 5 provides a similar example. In general, it can be seen that any sequence of isochronous (same-duration) notes is likely to be highly metrically ambiguous to the model; it will have little basis for choosing one analysis over another. In such cases, it tends to simply choose the fastest tactus level within its allowable range (400–1600 ms), which may lead to the correct choice but often will not.

Characterizing Parallelism

We now present a model of parallelism and its effect on meter. The model described here analyzes a piece for parallelism and then incorporates this information into the metrical analysis process described in the preceding section. (The model is capable of handling monophonic inputs only; we will return briefly to the problem of polyphonic music later on.) The problem can be decomposed into two subproblems: (1) How do we characterize parallelism—what *is* parallelism, exactly—and how do we search for it in a piece? (2) How do we use this information as input to metrical analysis? These two problems are the subject of this section and the next.

We have established that parallelism involves some kind of repetition. But repetition of what? Certainly parallelism can involve an exact repetition of a sequence of notes, as seen in the *Moonlight* Sonata. However, it can also involve repetition of an intervallic pattern at different pitch levels—what is sometimes known as a "sequence." This is seen, for example, in Figures 4, 5, and 6. In Figures 5 and 6, a pattern is shifted along a diatonic (major or minor) scale; in such cases, the pattern is repeated in terms of intervals on the scale, but not necessarily in terms of chromatic intervals (i.e. number of half-steps). (In Figure 5, the first three-note motive features chromatic intervals of -1 +1, and the second features -2 +2.) On the other hand, sequences can also occur where a chromatic intervallic pattern is repeated exactly; the melody of Figure 4 offers an example of this, where a pattern of 0 -3 +2 in measure 1 is repeated (transposed down two chromatic steps) in measure 2. Parallelism can also involve a repetition of *contour*: the pattern of ups and downs in a melody. In Figure 9, a three-note pattern is clearly evident in the first three measures, though it can be seen that the chromatic intervals vary from one pattern instance to the next and even the diatonic intervals vary. In terms of diatonic intervals, the first instance (G-B♭-G) features +2 -2, whereas the second (D-G-D) features +3 -3. Clearly, parallelisms of this kind must be allowed for as well. Intuitively, it seems that exact repetition and intervallic repetition (sequences) are the strongest form of parallelism, followed by contour repetition, followed by rhythmic repetition. (We are assuming that cases of intervallic repetition and contour repetition involve rhythmic repetition as well. This is generally the assumption in musical parlance; to speak of a sequence generally implies that the rhythmic pattern is repeated along with the intervallic pattern. Of course, one could have repetition of an intervallic pattern without rhythmic repetition, but cases of this are relatively rare and not particularly salient in perceptual terms.)

We propose the following scheme for representing these intuitions. The rhythm of a melody can generally be represented in terms of a low (fast) level of beats, such that every note spans an integer number of beats at that level; we call these Pbeats. (Complications may arise in cases where low-level beats are not all metrically equivalent—for example, where a quarter-note beat is divided sometimes in three and sometimes in four; we will avoid such cases here.) Pbeats are similar to "beats" as output by the metrical program, except that beats have a time point and a level, whereas Pbeats have only a time point. Let us assume that Pbeats are specified in the input to the parallelism model. (This assumption may seem problematic, considering that the function of the model is to help *determine* where the beats are, but we will return to this later.) Given a list of Pbeats and a list of notes (whose onsets, we will assume, always coincide with Pbeats), one can

Fig. 9. Bach, Sonata for Violin in G Minor, Presto, measures 1–11.

assign each Pbeat a pair of numbers: an "onset-status" number (1 if it contains the onset of a note, 0 if it does not) and a pitch number indicating the pitch of the note beginning or continuing at that point (assuming the usual convention of middle C = 60). For example, the first two measures of the Bach Minuet in Figure 2 could be represented as follows (assuming eighth notes as the Pbeat level):

(58,1) (57,1) (58,1) (50,1) (51,1) (43,1) (41,1) (41,0) (57,1) (57,0) (50,1) (50,0)

We do not distinguish here between rests and continuations of a note. Every note is assumed to continue until the next note begins.

We then generate a series of "phase statements." Each phase statement is identified with a number, *D*, representing a certain distance between Pbeats (in terms of the number of Pbeats elapsed—so the distance between adjacent Pbeats is 1). The phase statement then consists of a series of numbers (parallelism values, or "PVs"). Each PV represents the similarity between two Pbeats (in terms of the events they contain) that are *D* Pbeats apart; the *n*th PV in the list represents the relationship between the *n*th Pbeat of the piece and the Pbeat *D* Pbeats later. Consider first the case where both Pbeats contain note onsets. If the two Pbeats are the same in terms of their pitch interval to the previous note (assuming diatonic intervals, not chromatic), the corresponding PV in the phase statement is 3. (We compare only pitch intervals here; the rhythmic interval to the previous note is ignored.) If the two intervals are different, but alike in contour (i.e., either both ascending or descending), the PV is 2; if they are different in contour, the PV is 1. If the two Pbeats are both non-onsets, the PV is 2; if they are unlike in their onset status (one is an onset and the other is not), the PV is 0. These values were simply found to work well through trial-and-error testing. Figure 10 provides an illustration.

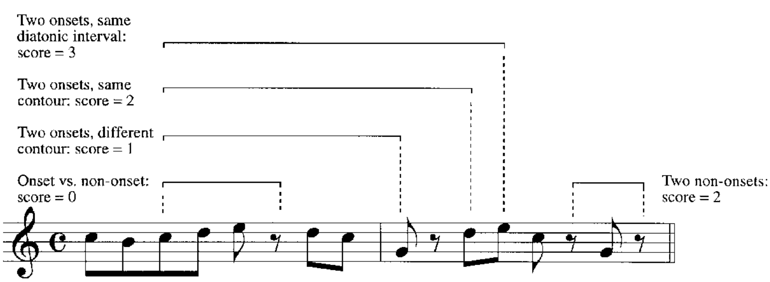

Fig. 10. Any two Pbeats (eighth-note beats in this case) can be compared for similarity, according to what happens there: whether or not an event-onset occurs, and if so, the preceding interval leading to the event. The example shows the scores ("PVs") that would be assigned for different cases. (The melody shown here was composed especially for this example.)

Figure 11 shows phase statements for the first four measures of the Bach minuet excerpt in Figure 2. (The first PV of each statement represents the first Pbeat of the piece compared with a later Pbeat; in this case the two Pbeats cannot be compared intervallically because there is no prior note. In this case, then, the score is simply 1 if the two Pbeats are alike in onset status or 0 otherwise.) Consider first the Phase 2 statement. The fact that the third and fourth PVs are 3 represents a repeated (diatonic) interval pattern of +1 -5 (this is shown by brackets above the staff). From "eyeballing" the numbers it can be seen that the Phase 2 numbers are generally higher than the Phase 3 numbers; this means that there is more repetition at a distance of 2 Pbeats than 3 Pbeats, and hence, that a metrical period of 2 Pbeats (a quarter-note beat) is favored over one of 3 Pbeats (a dotted-quarter-note beat) (exactly how this is enforced will be explained in the following section). Moving to higher-numbered phase statements, it can be seen immediately that the numbers for phase distance 12 are particularly high; the intervallic similarity between measure 3 and measure 1 is reflected in a sequence of 3s. Two other examples are shown in Figures 12 and 13, along with phase statements of particular interest. In Figure 12 (the Pbeat unit is the eighth-note here), we find a great deal of rhythmic repetition (at a distance of four beats)—the eighth-eighth-quarter motive occurs repeatedly, although with little similarity of interval and only moderate contour similarity. This is reflected in the Phase 4 statement, which

Phase	Parallelism values
1	1 1 1 1 1 2 0 0 0 0 0 0 1 1 1 1 1 2 1 2 3 1 1
2	1 2 3 3 1 0 1 2 1 2 1 0 2 2 3 3 1 1 1 2 1 2
3	1 1 1 2 0 1 0 0 0 0 2 0 1 1 1 2 2 1 1 1 2
4	1 2 1 0 2 0 2 2 2 0 1 0 2 2 1 1 2 1 3 2
5	1 3 0 1 0 2 0 0 1 0 2 0 1 3 2 1 2 2 1
6	1 0 2 0 1 0 1 0 2 0 1 0 1 1 2 1 1 1
7	0 1 0 2 0 1 3 0 1 0 2 0 2 1 2 2 2
8	1 0 1 0 2 2 1 0 2 0 2 0 3 1 1 1
9	0 2 0 1 1 1 2 0 1 0 1 0 3 3 2
10	1 0 2 2 3 3 1 0 1 0 1 0 1 1
11	0 1 1 1 1 1 2 0 2 0 1 0 2
12	1 3 3 3 3 3 3 0 2 0 2 0
13	1 1 1 1 1 2 1 0 2 0 1
14	1 2 3 3 1 1 1 0 1 0
15	1 1 1 2 2 1 1 0 2
16	1 2 1 1 2 1 3 0

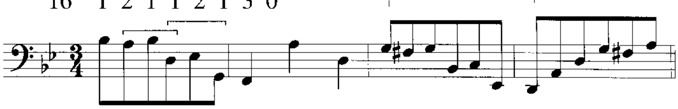

Fig. 11. Bach, Suite for Violoncello in G Major, Minuet II, measures 1–4, showing phase statements. The numbers ("PVs") in the phase statements indicate a comparison between the corresponding beat and another beat a certain distance later, depending on the "phase" (e.g., Phase 8 corresponds to a distance of 8 beats).

Fig. 12. Bach, Suite for Violoncello in C Major, Gavotte, measures 1–3 (melody only), showing phase statement for $D = 4$.

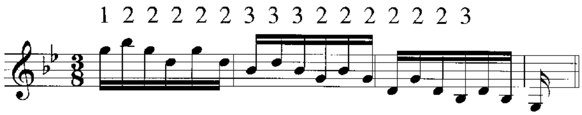

Fig. 13. Bach, Sonata for Violin in G Minor, Presto, measures 1–3, showing phase statement for $D = 3$.

consists mostly of 1s (indicating two onsets of different contour) and 2s (indicating two non-onsets, or two same-contour onsets). Finally, the Phase 3 statement of Figure 13 (the opening of the excerpt shown earlier in Figure 9) consists mostly of 2s, indicating a repeated pattern of contour with little repetition of interval at that phase.

A "parallelism finder" was designed that takes, as input, a list of Pbeats and notes, and outputs this input along with the phase statements just described—each phase statement gives its D number, followed by its PVs. (If the number of Pbeats in the input is N_{pb}, then the number of PVs in a given phase statement is $N_{pb} - D$. For example, consider the phase statement for $D = 4$ in a piece with 20 Pbeats. The first PV compares Pbeat 1 and Pbeat 5; the 16th PV compares Pbeat 16 and Pbeat 20; the 17th PV, if there was one, would be superfluous because there is no 21st Pbeat.)

A little explanation is needed of some other aspects of the parallelism finder. One is the calculation of diatonic intervals. As mentioned earlier, two intervals are generally considered the same if they correspond to the same diatonic interval within the currently established scale framework. To determine this decisively would require knowledge of the key—a significant problem that we will not explore here.[10] A simpler and still highly effective solution is available, however. We can observe that, as long as the intervals involved are the usual diatonic ones, their diatonic category can be calculated without knowledge of the key: a chromatic interval of +1 or +2 is always a second (major or minor), an interval of +3 or +4 is always a third, and so on. Thus we can say that a pair of intervals is the same if each interval is either +1 or +2 (though one might be +1 while the other is +2), or if each interval is either +3 or +4, and so on. Table 1 shows these corre-

10. Steedman's (1977) model uses knowledge of the key to determine diatonic intervals.

TABLE 1
Correspondences Between Chromatic and Diatonic Intervals

Chromatic Interval	Diatonic Interval
0	
1, 2	1
3, 4	2
5, 6	3
6, 7	4
8, 9	5
10, 11	6
12	7
13, 14	8
etc.	etc.

spondences. One problematic case is the tritone, chromatic interval 6, which could be either a diatonic interval 4 (a diminished fifth) or diatonic interval 3 (an augmented fourth) depending on the context; we allow interval 6 to be either diatonic interval 3 or 4. Chromatic intervals such as the augmented second (chromatic interval 3 but diatonic interval 1) are mishandled by this scheme, but this rarely seems to cause problems in practice.

A second issue concerns the number of phase statements generated. In principle, the program can generate phase statements of any distance—up to the number of Pbeats in the piece. In practice, however, phase statements are needed only up to the largest distances that could possibly occur between adjacent beats at any metrical level (this will be explained later). This means that only about 20–30 phase statements are usually required.

Before continuing, a general point arises here, namely the relationship between parallelism and motivic structure. Motivic structure is generally assumed to be some kind of network of segments in a piece that are similar or related. Although parallelism clearly relates to motivic structure, the representation of parallelism we have proposed differs from a motivic analysis in at least two fundamental ways. First of all, the representation proposed contains no explicit representation of motives or related segments. It could possibly be used as the basis for deriving such a representation, however; for example, any sequence of 3s in a phase statement will indicate a motive that is exactly repeated or transposed (diatonically or chromatically). Second, as noted earlier, the preceding representation captures only repetition across small time intervals; yet motivic connections can occur at indefinitely long time scales. It does not appear that such long-range motivic connections are often important for metrical analysis, but they may occasionally play a role; for example, it may be that once a motive is heard in a certain metrical context, there is a tendency to impose the same metrical interpretation on it when it occurs later, even much later, in the piece.

Using Parallelism as Input to Metrical Analysis

We now turn to the second part of the problem. The assumption is that the output of the parallelism finder—a list of notes, Pbeats, and phase statements—will be available to the metrical program as it analyzes a piece. How can this information be used to guide metrical analysis?

We begin with Steedman's formulation of the parallelism rule, quoted earlier: when a melodic pattern is repeated, "the metrical accent [should] fall at the same point in both figure and repeat." We can state this more generally, without committing to any segmentation of the input into "patterns," by saying that when repetition occurs at a certain distance, we want beats (at some level) to be separated by this distance (or "beat interval," as we call it in metrical terms) as well. Unlike Steedman—as discussed earlier—we do not assume that beats should occur only at the beginning of repeated patterns (this would not be possible under the current framework anyway, because no explicit "patterns" are identified). Rather, it seems that repetition at a certain distance should favor beat intervals of that distance anywhere in the neighborhood where the repetition occurs. This leads us to the following formulation of the parallelism rule as a preference rule:

> Parallelism Rule: Prefer beat intervals of a certain distance to the extent that repetition occurs at that distance in the vicinity.

Let us consider how to incorporate this rule into the metrical model proposed earlier. Suppose we are considering some tactus level—some row of beats—and we want to know how well it satisfies the parallelism rule. We examine each pair of neighboring beats in turn (first beats 1 and 2, then beats 2 and 3, and so on). Consider a hypothetical beat pair (B1, B2), where BI is the interval between them. We must first choose the phase statement whose time interval is closest to BI. Because the metrical program measures beat intervals in absolute time (milliseconds), rather than Pbeats, we must convert each phase statement's *D* value (measured in Pbeats) to an absolute time value. If we assume that all Pbeats are equidistant, this can be done simply by multiplying *D* by the time interval between any pair of adjacent Pbeats. (The case in which Pbeats are not equidistant will be discussed later.)

In this way, we associate each beat pair with a single phase statement. We then look at the PVs in that phase statement, within a certain time interval (the value we use is 1 s) on either side of B1. Note that these PVs tell us, essentially, the extent to which events in the vicinity of B1 are "paralleled" at an interval of BI later—that is, in the vicinity of B2. Essentially, we simply add all the PVs within this 2-s window; this then provides a "score" telling us the goodness, in terms of parallelism, of a pair of beats at B1 and B2. However, we also weight the PVs under a linear function reflecting closeness to B1. This linear function assigns a weight of 1 to PVs exactly at B1, descending to a weight of 0 for PVs 1 second away from B1. Figure 14 shows an example (using hypothetical PVs). The Pbeats are 250 ms apart, and the beat pair shown (between beats at time points 1500 and 3000) is being considered; presumably this phase statement would correspond to Phase 6, as that is the distance closest to BI (6 x 250 = 1500). The 2-s window around B1 is shown with a bracket; PV numbers within this window contribute their scores to the parallelism score, weighted according to the linear function described earlier.[11] (All PVs outside the 2-s window contribute a score of 0.) Using this method, any beat pair can be evaluated according to how well it is supported by the parallelism of the piece. An entire beat level can then be scored by summing the scores for each beat pair, and this score can be added in with the scores from the other preference rules. Because parallelism in a certain part of the piece affects only the scores for beat pairs in the vicinity, a change in the preva-

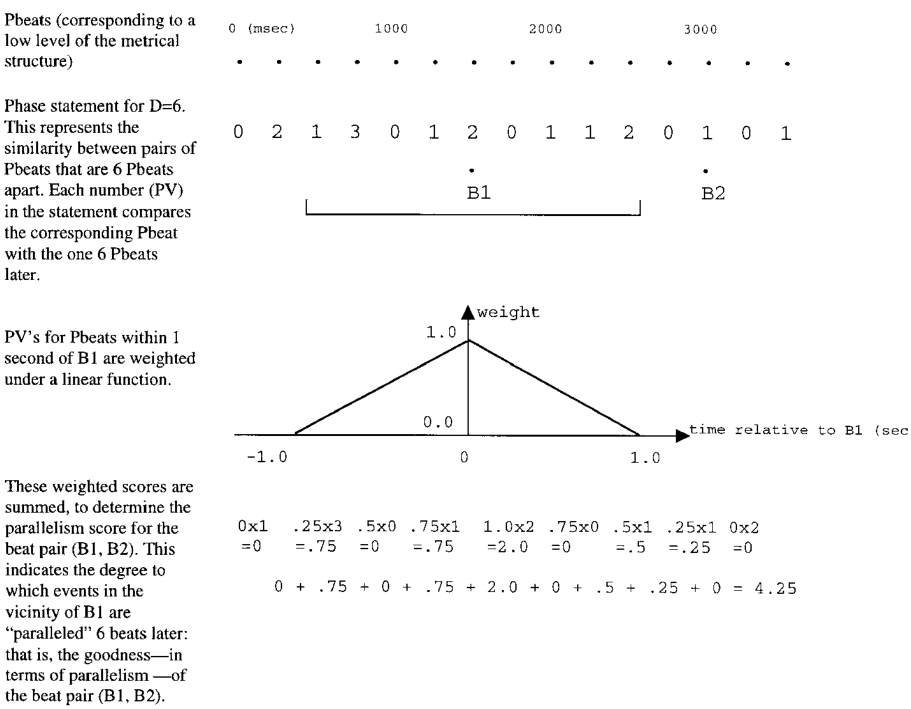

Fig. 14. An example showing the calculation of parallelism scores for the beat pair (B1, B2). (This is a hypothetical example, not reflecting any actual piece.)

11. The 2-s window around B1 may or may not overlap B2; in this case it does not.

lent parallelism distance from one section to another may result in a change in metrical structure, as is surely correct (an example of this will be seen later).[12]

As mentioned earlier, the program does not actually consider complete analyses, but uses an efficient search technique to find the best one; this is a variant of a search technique from computer science known as dynamic programming. As it happens, the method just described of examining individual beat pairs fits in nicely with this procedure. Consider just the tactus level. As explained earlier, all onsets and offsets are first quantized to pips, and beats may occur only at pip locations. The program then proceeds in a left-to-right fashion; at each pip, it finds the highest-scoring analysis of the piece so far ending with each possible beat interval, that is, ending with a beat at the current pip P_i and a beat at some earlier pip P_j. To do this it must consider adding the current beat pair, (P_j, P_i), on to each previous beat pair (P_k, P_j), in order to factor in the regularity rule; it takes the score for the best analysis of the piece ending in each (P_k, P_j) pair (which has already been calculated), and adds on the new score for (P_j, P_i). It records the score of the highest-scoring analysis ending in (P_j, P_i), along with the (P_k, P_j) pair that this analysis entails. When doing this, it factors in the parallelism score for (P_j, P_i). When the whole piece has been analyzed, it chooses the highest-scoring (P_j, P_i) near the end of the piece and traces this back through the piece to yield the preferred analysis.

Currently, the program applies the parallelism rule only for the tactus level and higher levels, not levels below the tactus; this was simply an intuitive decision based on the observation that parallelism does not usually seem to be an important factor in lower-level metrical analysis. The application of the parallelism rule at higher levels is the same as that described for the tactus. The search procedure for higher levels is similar as well; it is more constrained, however, given the well-formedness rules that every higher-level beat must be a beat at the immediately lower level, and exactly one or two beats at the level below must elapse between each beat pair at the current level.

One aspect of the model that might be questioned is the use of Pbeats as input to the parallelism finder. Because the whole point of the parallelism finder is to help determine the metrical structure, assuming Pbeats as input might seem circular. However, we would argue that this is (a) unavoidable and (b) not really problematic. Regarding the first point, some kind of integer representation of rhythmic values seems essential to our intuitions

12. It might be thought that the program should favor phases such as 2, 3, 4, and 6 as opposed to (for example) 5 and 7. We have not found much need for such a rule. Repeated patterns of 5 beats are extremely rare in common-practice music; when they do occur, it is arguable that they lead us to hear a quintuple metrical structure. However, the need for such a rule does arise occasionally; an example will be given later.

about parallelism. The idea of a repeating "rhythmic pattern" generally implies a pattern of durations measured as multiples of a low-level beat—*not* as a pattern of absolute time values; we can certainly perceive a repeating rhythmic pattern even when the tempo is fluctuating. We cannot think of any way of characterizing parallelism without some notion of low-level rhythmic units.[13] Regarding the second point, the parallelism finder only requires information about the lowest level of beats, and as mentioned earlier, parallelism rarely seems needed to determine this in any case. Our practical solution to this problem has been to first run a piece through the metrical program (ignoring parallelism), producing the usual note list and beat list; this output is read in by the parallelism finder, which treats the lowest level of beats as the Pbeats, ignoring all higher levels. The parallelism finder adds the phase statements; all this (notes, Pbeats, and phase statements) is then fed back into the meter program, although the Pbeats are now used only indirectly to determine the metrical structure (see Figure 15). (It is not assumed that Pbeats will necessarily be beats in the final metrical structure; they generally will coincide with the lowest metrical level of the final metrical structure, but may not always do so.) Indeed, something along these lines seems plausible for human perception as well; some kind of low-level metrical analysis is presumably done to identify patterns of repetition, which then guides higher-level metrical analysis—although a tighter integration of the two processes would surely be more cognitively plausible than what we have proposed.

A second issue concerns the equidistance of Pbeats. If we think of Pbeats as corresponding to the lowest level of the metrical structure, they will generally not be exactly evenly spaced in a situation of real musical performance. However, the general scheme just sketched—in which a piece is fed first to the metrical program to generate Pbeats, then to the parallelism finder, and then back to the metrical program—does not assume perfect regularity, at any level or at any stage. The metrical program is designed to handle situations of fluctuating tempos, and can identify lowest-level beats in such cases with considerable (although not total) success. (If the metrical program is unsuccessful at this task—if the lowest-level beats are not correctly identified—then of course the performance of the parallelism finder may suffer.) If the Pbeats identified by the metrical program are somewhat irregular, this itself should not cause any problem; indeed, the parallelism finder pays no attention to absolute time values, except to determine which note onsets coincide with which Pbeats. Nor should it be problematic to feed these irregular Pbeats back into the metrical

13. If the lowest-level beats of a piece were perfectly regular, and all note onsets coincided perfectly with these beats, the parallelism finder could probably determine the Pbeats fairly easily from the note list. However, as we will explain, it seems better *not* to assume that Pbeats are perfectly regular.

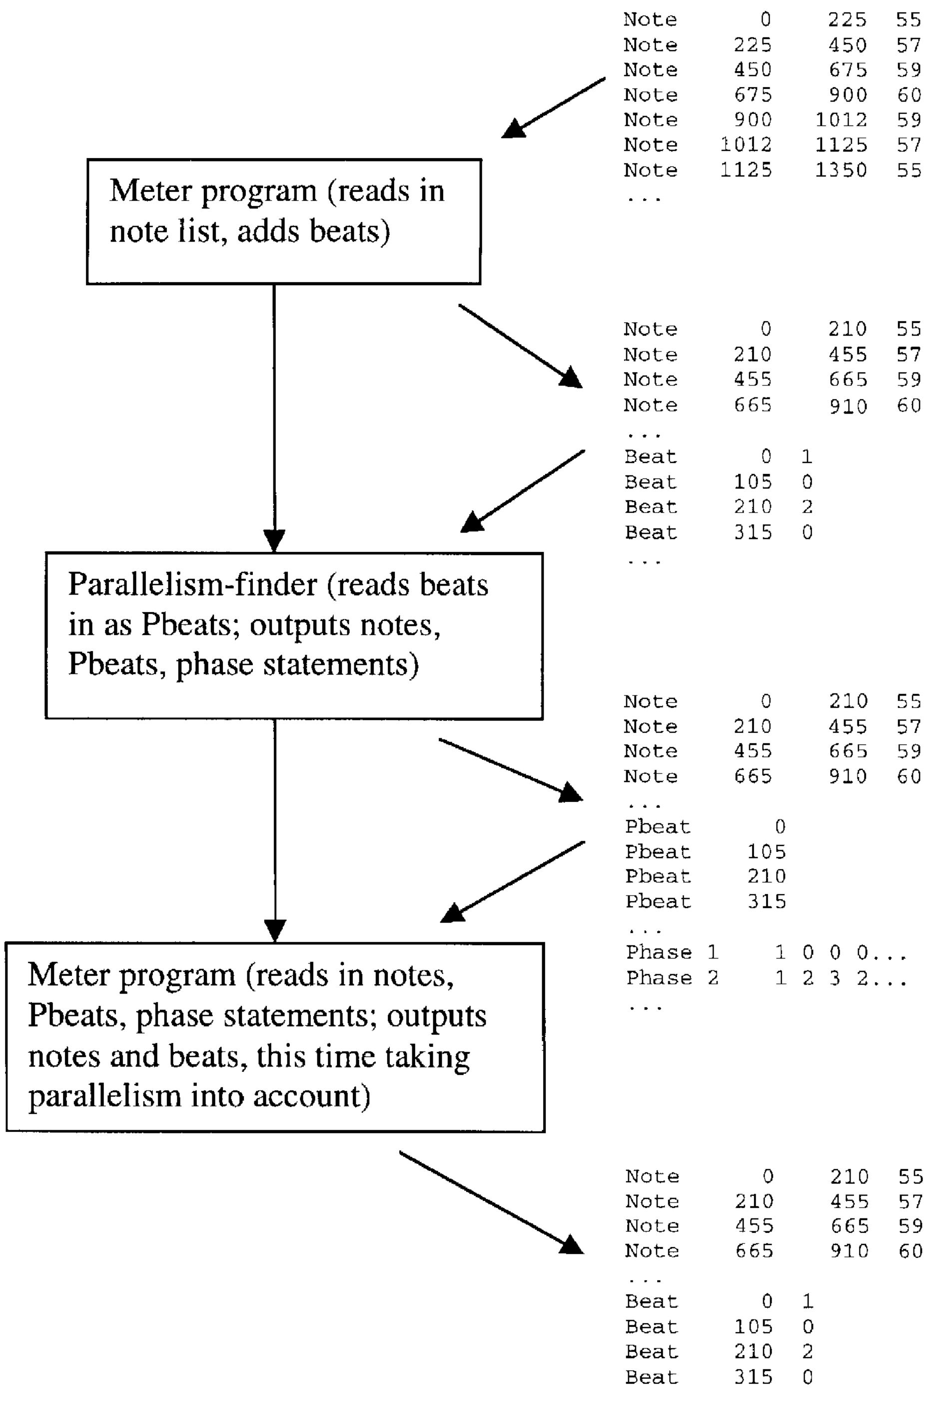

Fig. 15. The operation of the combined parallelism and meter programs.

program.[14] In short, unquantized data do not appear to pose any problem for the parallelism finder itself; such data may sometimes cause problems for the meter program as a whole, but this is not our present concern.

A Few Examples

In developing a computational model, it is desirable to perform some kind of test of the model's success. The question is whether the parallelism model proposed here improves the performance of the Temperley/Sleator meter program, and to what degree. Ideally, testing would be done by running the program with and without the parallelism component on a systematically chosen and varied corpus of common-practice pieces and then comparing the results. This approach was problematic for several reasons. Because the model works only on monophonic music, the test corpus would have to consist entirely of monophonic excerpts; this eliminates corpora that have been used for testing elsewhere, such as the Kostka-Payne corpus (Temperley, 2001). Moreover, because the Temperley/Sleator program achieves the correct analysis on a large majority of excerpts even without parallelism (see the test results cited in the section earlier), quite a large corpus of pieces might be required in order to obtain a significant level of improvement. There is another complication here as well (one that arises, in fact, with all metrical models): we are seeking to model human perception of meter, but the perceived meter of a piece may not always correspond exactly to the notated meter, although we assume that it generally does. For example, in the fugue subjects of Bach's *Well-Tempered Clavier* (a corpus that might otherwise be quite appropriate for testing), it is sometimes quite doubtful that the perceived meter would correspond to the notated one, at least on the basis of the fugue subject alone.[15] For these reasons, systematic testing was determined to be impractical at present (although we hope to undertake it in the future). Rather, the program was simply tested on an unsystematically chosen set of monophonic excerpts from Bach fugues and violin and cello suites; several of these are discussed

14. One issue does arise here: in deciding which phase statement is associated with a given beat pair (B1, B2), the metrical program looks for the *D* that is the closest to that pair's beat interval. If the Pbeat intervals are not all the same, it makes a difference how this value is calculated. The program does this by averaging the 10 Pbeat intervals in the vicinity of B1. In this way, fluctuations in Pbeat length—even quite large fluctuations—from one part of the piece to another should not cause severe problems, because Pbeat intervals are calculated in a local fashion.

15. Consider, for example, Book I, Fugues 1, 8, 14, and 19. This corpus has, in fact, been used for testing of other metrical models (Longuet-Higgins & Steedman, 1971; Steedman, 1977); in these cases, the notated meter was assumed to be correct.

later. The examples presented are all cases where the prior version of the program (as proposed in Temperley, 2001) produces an incorrect analysis; the issue, then, is whether incorporating the factor of parallelism can improve the program's performance. (It is of course possible that the parallelism component would *worsen* performance in cases where the earlier model was correct; we have not investigated this yet.)

The only parameter to be set (aside from those already discussed) was the weight of the parallelism rule relative to other rules; an optimal value of this was determined by trial-and-error testing, which was used on all the examples reported here. Regarding other parameters of the program (not relating to parallelism), the parameter set used is the same as that used in the tests reported in Temperley (2001), with one exception. On the monophonic passages used here, the tactus level produced by the program was excessively irregular; thus a greater weight was given to the regularity rule to ensure a regular tactus.

Figure 16, the opening of the Allemande to Bach's Suite for Violoncello in C Major, provides a simple successful example of the model. (Just this portion of the piece was used as input to the program.) The original notation of the passage is shown in Figure 6. Without the input of parallelism, the program gets the analysis shown in Figure 16A. The eighth-note level is correct (the program regards this as Level 2, i.e., the tactus); however, the next level up (Level 3) is analyzed as triple, rather than duple. (The main factors here are the low C in the middle of the phrase and the B eighth note at the end, both of which the program wants to consider as metrically strong. Recall that the length of a note is determined by its "registral interonset interval"; the low C is a long note by this criterion.) Once the factor of parallelism is included, the program produces the analysis shown in Figure 16B. Level 3 is analyzed as duple, Level 4 is duple as well, and the phase of each level is correct; thus the metrical analysis produced corresponds exactly to the original notation, with the exception that the "whole-note" level is missing. As noted earlier, the intervallic pattern repeating at a distance of a quarter note (Phase 8, given the 32nd-note level as the Pbeat

Fig. 16. The Bach excerpt shown in Figure 6, showing (A) the program's analysis without parallelism, (B) the program's analysis with parallelism.

level) strongly encourages the program to have a metrical level of this period. Figure 9, the opening of the Presto from Bach's G Minor Violin Sonata, shows another successful example. Without parallelism, the program's analysis is completely wrong; it finds a quarter-note level rather than a dotted-quarter level as the tactus, with beats on every fourth sixteenth note. With parallelism, the metrical structure identified is perfectly correct.[16] (The program does not actually need as much of the piece as this to identify the correct structure; if given only the first three measures, it finds the correct analysis, whereas without parallelism it is once again incorrect.)

Figure 17 shows a case of partial success, the subject of Fugue 21 from Book I of the *Well-Tempered Clavier*. Figure 17A shows the correct metrical structure, as indicated in the notation. Note the parallelisms at the one-measure (dotted-half-note) level between the end of measure 1 and the end of measure 2, and also (even more strongly) between measure 3 and measure 4. Without parallelism, the program produces the analysis in Figure 17B; the tactus (quarter-note) level is correct, but Level 3 is incorrectly identified as duple. With parallelism, the analysis in Figure 17C is produced. Now the period of Level 3 is correct, but the phase is wrong; the program places a dotted-half-note beat on the second and fifth quarter-note beats rather than the third and sixth. This error cannot be blamed on the parallelism model itself. As discussed earlier, the parallelism model is only supposed to decide the correct *period* for each level; it is indifferent as to *phase*. The hope is that other factors will determine the phase correctly, but this does not always happen. Perceptually, the important cue here would

Fig. 17. Bach, *Well-Tempered Clavier* Book I, Fugue 21, measures 1–4. (A) The original notation. (B) The program's analysis without parallelism. (C) The program's analysis with parallelism.

16. The program identifies the level below the tactus as the eighth-note level rather than the dotted-eighth. Although this agrees with the notation (which specifies 3/8 rather than 6/16 as the time signature), it might seem surprising—perhaps even contrary to perception—given the strong three-sixteenth-note parallelisms in the first three measures of the piece. The reason for this is that, under the current implementation, parallelism affects metrical analysis only at the tactus level and higher, not at lower levels.

seem to be harmony (a factor not incorporated into the standard version of the program); there are clear harmonic changes on the second and third notated downbeats, leading us to infer strong beats there. Cases such as this—where the period of the meter is correct but the phase is not—suggest that although the parallelism component of the model is working fairly well, other aspects of the model require improvement (through the addition of other factors, or the improved implementation of existing ones).

Figure 18, the subject of Fugue 24 from Book II of the *Well-Tempered Clavier*, provides an interesting case of failure in the parallelism model itself. Figure 18A again shows the correct analysis. Without parallelism, the analysis in Figure 18B is produced; the tactus is incorrectly identified as duple rather than triple. Perceptually, the triple meter of this subject seems obvious enough; and it seems clear that at least one factor is the repeated motive (down an octave, then up an octave) in measures 3 and 4. Yet, in this case, adding parallelism does not produce the right analysis. Even increasing the weight of the parallelism factor beyond the normal level does not lead to the correct result. Why does the program fail here? One clue lies in the phase statements. Figure 18 shows the phase statements for Phase 4 (the incorrect quarter-note level) and Phase 6 (the correct dotted-quarter level). From eyeballing the numbers, it does not appear that the Phase 6 numbers are any better than the Phase 4 numbers. Phase 6 has seven 3s (indicating a same-interval match); Phase 4 has six 3s. Phase 6 is hardly preferable to Phase 4 from the point of view of the PVs, so the metrical program has little reason to favor the dotted-quarter pulse. What the Phase 6 statement does have is a string of 3s and 2s close together, indicating the repeated motive in measures 3–4. This suggests that, perhaps, the program should give more weight to an intense degree of repetition over a short period—that is to say, a *motive*—rather than to a moderate degree of parallelism spread out over a longer period (as is found for

Fig. 18. Bach, *Well-Tempered Clavier* Book II, Fugue No. 24, measures 1–6. (A) The original notation; phase statements for Phases 4 and 6 are shown above the staff. (B) The program's analysis, with and without parallelism.

phase distance 4).[17] However, we will not attempt to formalize this intuition any further here.

A final example illustrates an important feature of the metrical program (and the current model of parallelism): the ability to detect shifts in metrical structure from one section of a piece to another. As mentioned earlier, we found that the program's tactus level was generally too irregular on the excerpts used for testing, so the weight of the regularity rule was boosted. However, reducing the weight of the regularity rule sometimes produces interesting results. Figure 19 shows the entire first half of the Presto to Bach's G Minor Violin Sonata (the opening of which was discussed earlier). The excerpt was analyzed by using a relatively low weighting of the regularity rule. Dots above the score indicate the tactus level found by the

Fig. 19. Bach, Sonata for Violin in G minor, Presto, measures 1–54. The program's analysis is indicated above the staff.

17. One possibility would be to square the parallelism scores contributed to each beat pair, thus favoring more concentrated (as opposed to more diffuse) parallelism.

program. It can be seen that this analysis is correct in some aspects, incorrect in others. The beginning (mm. 1–17) is correct in both period and phase; from measures 17–31, the analysis is correct in period but one sixteenth note off in phase. (The fact that the final note of m. 18—and many of the subsequent measures—is not closely followed within register makes it seem "long" to the program.) The switch to a quarter-note pulse (with beats on every fourth sixteenth note) at measure 33, although contrary to the musical notation, is not implausible; a clear four-note sequence occurs in measures 33–35 (a kind of cross-rhythm), which, it could be argued, temporarily challenges the dotted-quarter beat and causes us to entertain a quarter-note beat instead. Less plausible is the quintuple pulse (with beats every five notes) in measures 36–42; this is not supported by any significant five-beat parallelism but seems to be the program's way of switching gradually from the 4-sixteenth-note period of the previous section to the 6-note period of measure 43 onward. Although it was noted earlier that there was rarely a need for a rule against five- and seven-beat groupings (see footnote 12), this a case where such a rule would help.[18] In this final section, the period is once again correct (and is strongly supported by parallelism), although the phase is once again wrong. For the most part, the parallelism rule seems to be working well in this case, detecting the prevalent six-beat parallelism in many parts of the piece but also the four-beat pattern in measures 33–35. The fact that the *phase* of the meter is so often incorrect shows (once again) that there is room for improvement in other aspects of the model.

Directions for Further Work

Aside from points made in the preceding section, there are several ways that the current model could be extended and improved. One obvious improvement would be to extend the program to handle polyphonic music—a natural further step, considering that the Temperley/Sleator program is designed for polyphonic input. This would appear to be a major undertaking, however; in particular, it would require a different approach to measuring parallelism from the one just described. One approach would be to look only at the top voice of the music (which could be crudely identified by looking at the highest note present at each moment and merging these together) and performing monophonic parallelism analysis on this. This approach seems unlikely to work well, however; very often, parallelism is

18. Measures 35–42 do feature a strong parallelism, but it is a parallelism at a distance of 12 beats—two measures—so it is of little help to the program in determining the tactus level.

most prominent in an inner (accompanying) voice, rather than in the melody. In the *Moonlight* Sonata, for example (Figure 3), it is the middle voice (which starts out as the top voice but is joined by a higher melody in m. 5) that establishes and maintains the repeated pattern that is primarily responsible for conveying the quarter-note level of the meter. A more robust approach would be to first identify the melodic voices of the texture (a highly complex process in itself—see Marsden, 1992; Temperley, 2001), analyze each voice independently for parallelism, and then somehow merge these analyses together. For example, the PVs in each phase statement could reflect the sum of PVs for the individual voices.

One might also wonder if the method proposed here for identifying parallelism within a single voice is really sufficient. In particular, parallelism can often involve a melodic pattern followed by an elaborated variant. The current program's success at handling this depends on the nature of the elaboration. Consider the Bach fugue subject in Figure 17A; one might regard this as a one-measure pattern (ending on the downbeat of m. 2) followed by an elaborated (and shifted) repetition of the pattern. In this case, part of the pattern (the last three notes) is repeated exactly (although transposed) in the repetition; the program identifies this and thus captures the partial similarity between the two segments. However, consider a case such as Figure 20, the Gavotte from the E Major Violin Partita (consider just the top line for the moment). The bracketed segments shown in the score clearly sound related; the second is an elaborated version of the first. But the elaboration here largely involves the insertion of new notes between existing notes, and this disrupts the beat-to-beat similarity between the two segments. (Only the F♯-E-F♯ segment of the original is exactly preserved in the repeat.) As a result, the parallelism between the two is mostly unnoticed by the program. To handle such a case, the program would need to identify the structural notes of each segment and perform the parallelism comparison on these; it would then realize that the structural notes of the two segments are essentially the same. Of course, identifying the structural notes of a melody is a highly complex and often very subjective matter; and

Fig. 20. Bach, Partita for Violin in E major, Gavotte, measures 1–7.

it is not clear that such subtle structural parallelisms are often a factor in the perception of meter, important and interesting though they are.

So far, we have assumed that the effect of parallelism on meter is to favor beat intervals of the same period as the parallelism. However, parallelism also appears to affect metrical analysis in at least two other ways. For one thing, when a pattern is immediately repeated, there is a preference for the stronger beat to occur on the first occurrence of the pattern rather than the second. Temperley (2001) refers to this as the "first-occurrence-strong" rule. Figure 21 offers an example. Given just the first measure, one tends to hear a strong beat at the very beginning of the piece (following the tendency to hear strong beats near phrase beginnings). However, there is a clear parallelism between the second half of measure 1 and the first half of measure 2; by the rule just mentioned, this favors hearing beat 3 of measure 1 as a downbeat, stronger than beat 1 of measure 2 (this, of course, is contrary to the notated meter). A further, subtle, effect of parallelism on meter involves patterns that repeat with one or more notes held constant, while others change. In such cases, there is a tendency to hear the strong beats on the changing notes rather than the constant ones. In Figure 22, the fact that three notes of the four-note pattern remain the same (C-B-C) while the fourth one changes causes us to hear the changing note as metrically strong. This might be called the "new information" rule: when a pattern contains new information and old information, we tend to locate the metrical accents on the new information. This rule would also help the program in its analysis of the Bach presto (Figure 19). A six-note pattern in measure 25 is repeated in varied form in the following three measures; in this case, the fact that the last three notes of each measure (B♭-A-B♭) are repeated exactly each time should tell the model that the strong beat is

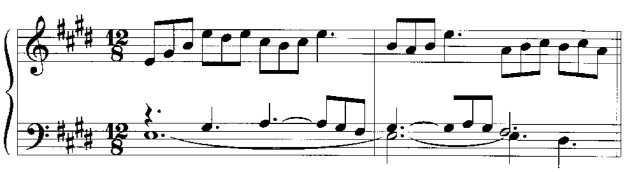

Fig. 21. Bach, *Well-Tempered Clavier* Book I, Prelude No. 9, measures 1–2.

Fig. 22. Bach, Prelude BWV 953, measures 1–3.

unlikely to be here, but rather, somewhere in the first three notes of the pattern (which are different on each occurrence).[19]

A final complication in the modeling of parallelism and meter concerns the fact that the influence is not just one-way. Parallelism can often influence metrical analysis; but once a metrical framework is established, this can affect the parallelisms that are heard.[20] In some cases parallelisms that go against an already established metrical framework go by almost unnoticed. For example, Figure 23 features an exactly repeated three-sixteenth-note motive, G-F-E♭ (marked with brackets), but because the two occurrences are not parallel with respect to the prevailing meter the parallelism is hardly heard. (Rather, we hear the first G-F-E♭ as parallel to the F-E♭-D beginning on the second quarter-note beat of the measure; and these two segments *are* metrically parallel.) On the other hand, cases such as the cross-rhythm in measures 33–35 of the Bach presto (Figure 19) suggest that parallelisms of sufficient strength can stand out against a previously established meter and can even cause us to modify our metrical perception accordingly. The complex relationship between meter and parallelism brings to mind the interaction of meter with harmony and grouping, mentioned earlier. Like parallelism and meter, harmony and meter have a complex interactive relationship; harmonic structure can influence meter, but metrical structure in turn exercises a profound influence on harmony. Of course, these kinds of interaction between levels present formidable challenges from a modeling point of view.

In short, there are a number of ways in which the current model could be improved. It could be expanded to handle polyphonic music and varied

Fig. 23. Bach, Suite for Violoncello in E♭ major, Allemande, measures 5–6.

19. The "new information" factor is discussed by Steedman (1977). This phenomenon is analogous to the phenomenon of "contrastive stress" in language: elements of a sentence are generally accented when they are new or unexpected. Consider the sentence "John chased the dog." Normally, one would say this sentence with the accent on the direct object *dog*. However, if your audience had been assuming that *Fred* had chased the dog, you would probably accent the subject: "*John* chased the dog." If your audience had assumed that John had petted the dog, you might say "John *chased* the dog." This is a bit different from the musical case—the linguistic case does not involve actual repetition of anything, and the linguistic case involves explicit accentuation of certain elements while the musical case does not—but the basic idea is the same.

20. It could also be argued—as noted earlier—that metrical analysis affects parallelism in a more basic way, in that some kind of encoding of lower-level beats is necessary for parallelism to be represented.

(elaborated) repetition; it could also be extended to incorporate other effects of parallelism—the "first-occurrence strong" rule and the "new information" rule. The "two-way" relationship between meter and parallelism also remains to be addressed. Clearly, a number of challenges remain in the modeling of parallelism and its effect on meter.

In closing, we should return to an issue that was raised briefly in the first section of this article: the relationship between computational modeling and psychology. Even if the current model were developed to the point of complete success in predicting listeners' judgments of meter, this would not prove that human perception of meter involved the same procedure. However, in the course of this computational investigation into parallelism and its role in metrical analysis, we have identified a number of aspects of the process that we believe any psychological model will have to incorporate. We have suggested that such a model must be sensitive to different kinds of repetition: rhythmic repetition, contour pattern, and exact (diatonic) intervallic pattern. We have also suggested that parallelism is best modeled by comparing beats at different distances, rather than by an explicit segmentation of the music into patterns or segments, a step that is not always independently motivated and that can also lead to problems of redundancy.[21] Regarding the effect of parallelism on meter, we have suggested that parallelism primarily affects the distance between beats (period), not the exact placement of those beats relative to the music (phase). Any model that meets these various requirements deserves consideration as a hypothesis about musical cognition—all the more so if it achieves reasonable success in producing the desired outputs. However, whether our model truly captures the psychological processes involved in the perception of parallelism and meter remains an open question; further work—in particular, experimental work—is needed to determine its validity in this regard.

References

Allen, P., & Dannenberg, R. (1990). Tracking musical beats in time. In *Proceedings of the 1990 International Computer Music Conference* (pp. 140–143). San Francisco: International Computer Music Association.

Boltz, M., & Jones, M. R. (1986). Does rule recursion make melodies easier to reproduce? If not, what does? *Cognitive Psychology, 18,* 389–431.

Cambouropoulos, E. (2001). Melodic cue abstraction, similarity, and category formation: A formal model. *Music Perception, 18,* 347–370.

21. We can sometimes have a strong sense of repetition at a certain distance, without having a strong sense of any particular segmentation. In measures 12–16 of Figure 19, for example, we clearly perceive repetition at a distance of one measure, but the segmentation is quite ambiguous; it is by no means obvious that segments begin at the beginning of each measure.

Chafe, C., Mont-Reynaud, B., & Rush, L. (1982). Toward an intelligent editor of digital audio: Recognition of musical constructs. *Computer Music Journal, 6*(1), 30–41.
Desain, P., & Honing, H. (1992). *Music, Mind & Machine.* Amsterdam: Thesis Publications.
Deutsch, D. (1980). The processing of structured and unstructured tonal sequences. *Perception & Psychophysics, 28,* 381–389.
Deutsch, D., & Feroe, J. (1981). The internal representation of pitch sequences in tonal music. *Psychological Review, 88,* 503–522.
Drake, C., & Palmer, C. (1993). Accent structures in musical performance. *Music Perception, 10,* 343–378.
Handel, S. (1973). Temporal segmentation of repeating auditory patterns. *Journal of Experimental Psychology, 101,* 46–54.
Huron, D., & Royal, M. (1996). What is melodic accent? Converging evidence from musical practice. *Music Perception, 13,* 489-516.
Jackendoff, R. (1991). Musical parsing and musical affect. *Music Perception, 9,* 199–230.
Kostka, S., & Payne, D. (1995). *Workbook for Tonal Harmony.* New York: McGraw–Hill.
Large, E. W., & Kolen, J. F. (1994). Resonance and the perception of musical meter. *Connection Science, 6,* 177–208.
Lee, C. (1991). The perception of metrical structure: Experimental evidence and a model. In P. Howell, R. West, & I. Cross (Eds.), *Representing Musical Structure* (pp. 59–127). London: Academic Press.
Lerdahl, F., & Jackendoff, R. (1983). *A Generative Theory of Tonal Music.* Cambridge, MA: MIT Press.
Longuet–Higgins, H. C., & Steedman, M. J. (1971). On interpreting Bach. *Machine Intelligence, 6,* 221–241.
Marsden, A. (1992). Modelling the perception of musical voices: A case study in rule–based systems. In A. Marsden & A. Pople (Eds.), *Computer Representations and Models in Music* (pp. 239–263). London: Academic Press.
Mont-Reynaud, B., & Goldstein, M. (1985). On finding rhythmic patterns in musical lines. In *Proceedings of the 1985 International Computer Music Conference* (pp. 391-397). San Francisco: International Computer Music Association.
Parncutt, R. (1994). A perceptual model of pulse salience and metrical accent in musical rhythms. *Music Perception, 11,* 409–464.
Povel, D.–J., & Essens, P. (1985). Perception of temporal patterns. *Music Perception, 2,* 411–440.
Rolland, P.-Y. (1999). Music content feature abstraction: Discovering patterns in musical sequences. *Journal of New Music Research, 28,* 334–350.
Rosenthal, D. (1992). Emulation of human rhythm perception. *Computer Music Journal, 16*(1), 64–76.
Steedman, M. (1977). The perception of musical rhythm and meter. *Perception, 6,* 555–570.
Tanguiane, A. S. (1993). *Artificial Perception and Music Recognition.* Berlin: Springer.
Temperley, D. (2001). *The Cognition of Basic Musical Structures.* Cambridge, MA: MIT Press.
Temperley, D., & Sleator, D. (1999). Modeling meter and harmony: A preference-rule approach. *Computer Music Journal, 23*(1), 20–27.
Todd, N. P. M. (1994). The auditory "primal sketch": A multiscale model of rhythmic grouping. *Journal of New Music Research, 23,* 25–70.