TEMPO CURVES CONSIDERED HARMFUL*

A critical review of the representation of timing in computer music

Peter Desain & Henkjan Honing

Center for Knowledge Technology
Utrecht School of the Arts
Lange Viestraat 2B
NL-3511 BK Utrecht

In the literature of musicology, computer music research and psychology of music, timing and tempo measurements are mostly presented in the form of continuous curves. The notion of these tempo curves is dangerous, despite its widespread use, because it gives its users the false impression that a continuous concept of temporal flow has an independent existence - a musical or a psychological reality, and that time can be perceived independent of the events carrying it. But, if one bases a transformation or manipulation of timing on the implied characteristics of such a notion, this is doomed to fail.

Introduction

Music representation is an important domain in computer music because it directly influences the behavior of computer music systems. A careful choice has to be made, based on the right representational characteristics of the domain. There have been several discussions (see e.g. Dannenberg, Dyer, Garnett, Pope & Roads, 1989; Honing, 1991) and concrete proposals for the representation of music (e.g. ANSI, 1989). Timing turns out to be one of the most complex and pluriform aspects of music to capture in a representation. It often forces designers to make pragmatical choices at the cost of generality, flexibility and, as we are about to show, even musical validity. This paper focuses on the representational aspects of timing and will review some well-known representations of timing used in computer music.

In a musical performance music, some of the required timing, tempo, and changes thereof, are explicitly notated in the score (if available) in the form of a global tempo marking or an *accelerando* or *rallentando* sign, but many more, and significantly more

subtle variations are performed either based on structural features marked in the score (like bar lines) or the interpretation of the piece by the performer. These timing deviations, that are collectively called *expressive timing* or *rubato*, have been shown to be continuously variable and reproducible (Shaffer, Clarke & Todd, 1985), and clearly related to the musical structure (Clarke, 1988; Palmer, 1989). It is the latter observation that lacks in the representations of timing used in computer music; without a link to musical structure these representations are quite useless.

The representation of timing and tempo with patterns and curves

Before putting the 'Tempo Curve' to the test, it might be helpful to set out some terminology normally associated with the representation of timing and tempo in the different fields of music research.

To refer to expressive timing, in computer music the term *micro tempo* is often used, comparable to the term *local tempo* (or velocity) used in the psychology of music (i.e. the tempo changes from event to event, expressed as a ratio of a score time interval and a performance time interval). For clarity, the term *timing* would be more appropriate here. It specifies the timing deviation on a note-to-note basis and is often referred to as the *expressive timing profile*,

* A more elaborate version of this article will be published as (Desain & Honing, 1991a), telling the story of a psychologist, a mathematician, and their musical friend putting the Tempo Curve to the test. It also includes an extensive list of references.

Figure 1. Piano-roll notation of Figure 2 and a "tempo track" specifying the local tempo for each note of a performance at a) tempo 60, and one at b) tempo 90.

timing pattern or *rubato pattern*. In these patterns, the points are usually connected, either stepwise with straight line segments or with a smooth interpolation, yielding a *timing curve*. Only the first (stepwise) representation maintains a proper relation with the time map in which points are connected with line segments. These continuous *time maps* are used by Jaffe (1985) and most people of the computer music community. They relate performance time to score time. Time maps can be superimposed, using one for each voice. They can also be constructed for uniformly spaced units in the score, like bars or beats. The corresponding duration patterns then form a true *tempo pattern*.

It are the explicit and implied characteristics of these timing and tempo curves (the Tempo Curve, for short) that we want to put to a critical test.

Putting the Tempo Curve to the test

Figure 1a shows an example of a "tempo track" - a Tempo Curve that resides inside the modern sequencer programs. On the horizontal axis runs the so-called *score time*, indicating where one is in the score in bars or beats. The vertical axis indicates *local tempo*. A note that is played longer than prescribed in the score will yield a value below 60 here. Likewise, a note that is played shorter will result in a value above 60. Thus the peaks in this pattern represent speeding up.

These kinds of graphs, in different forms (see also Figure 3), gained such a

widespread scientific and practical use, that they almost gained an existence of their own, suggesting that Tempo Curves

- make musical sense
- are useful representations in computer music systems
- exist as a mental representation (in the head of the performer or listener)

If these graphs indeed represent something, then it must be possible to actually *use* them. For instance, when applying a transformation to this representation, it should yield a reasonable performance.

We will take as an example a change of global tempo, applied to the performance of a melody (the theme from the six variations composed by Ludwig van Beethoven on the duet *Nel cor più non mi sento*, see Figure 2). This is easy to accomplish with a sequencer program: they all have a tempo knob that can be adjusted freely.

But, if one listens to such a speeded-up version of a performance of this theme from, say, tempo 60 to tempo 90, it will sound weird. Why? Because a sequencer will speed up everything by the same amount, whereas, in a real performance, the expressive timing profile is not just played faster (speeded-up walking is not the same as running, as can be seen on a video in fast-forward mode).

Comparing the timing of the performance played at tempo 60 with one played at tempo 90 brings out a lot of differences (see

Figure 1a and 1b; of course, a listening test will be even more revealing).

A prominent difference is that the rubato is adapted according to the tempo. When a piece is played at another tempo, other structural levels become more important, e.g. at a faster tempo the *tactus* (the level of the metrical structure where the beats pass at a moderate rate) will shift to a higher level, the fine subdivisions of the beat will get more “out of focus”, and phrasing of longer time spans will gain in detail.

Figure 2. Score used for the performances depicted in Figures 1 and 3.

Even more noticeable are the local effects. For some notes the duration is not changed at all in the faster performance. Some of these notes are grace notes. But not all grace notes behave like this. For example, in the theme, the two grace notes that cover an interval of a sixth, in bar 7 and 19 (see Figure 2), are timed like any other note: they are actually played in a metrical way. Thus there are grace notes that are notated the same way in the score, but which are performed differently; there is a difference between ornaments that either “crush in” notes (ornaments that take up very little or no time) or “lean on” notes (ornaments that have an relatively important melodic or harmonic function).

The fact that a simple sequencer program cannot play the onsets of such ornaments correctly might be forgivable, but there are still more problems. This kind of representation can also not deal with the articulation of notes, i.e. the proportion of the notes’ duration that is actually sounding. Especially when these notes are played staccato at a moderate tempo, they will sound too short at the faster tempo; the

sounding time interval cannot just be scaled to another tempo. In real performances articulation is adapted according to the tempo and with respect to the rhythmical and metrical structure that it is part of.

Furthermore, the chord-spread (the small timing asynchronies within a chord) should not be changed by the global tempo change transformation.

So in fact, sequencers cannot be used to *change* something as simple as the global tempo. This is because detailed knowledge

about structural levels, articulation, timing of ornamentations, and other structural information is indispensable. It turns out that a tempo knob on a commercial sequencer package cannot be used to adjust the tempo.

Summarizing this paragraph, we can now put together a small list of questions that you should

ask a designer of a computer music system claiming to support time and tempo transformations:

- what about articulation?
- how does it handle ornaments?
- how does it deal with chord-spread and the timing relations between multiple voices?

Timing is linked to structure

We have shown that structure is indispensable for a musical relevant representation of timing, and, in principle, timing can be linked to any musical structural concept. Before we continue our evaluation of the Tempo Curve, we will give an overview of well-known types of musical structure that timing can be linked to and that should be supported by a proper representation of timing.

A *chord* is such a musical structure. Small timing asynchronies within a chord (called “chord spread”) are perceived as an overall timbral effect - the actual timing pattern or order is hard to perceive.

A second group are *ornaments*, like grace notes and trills. They can be roughly divided in *acciaccatura*, so-called 'timeless' ornaments, and *appoggiatura*, ornaments that take time and can have a relatively important harmonic or melodic function. The former normally fall outside the metrical framework, the latter tend to get performed in a metrical way.

Expressive timing also takes place between the different *voices* in a piece. The independent timing of individual voices is sometimes hard to perceive because their components are immediately organised by the perceptual system into different streams. This is not the case with (almost) simultaneous onsets which result in clear timbral differences. These can be heard in ensemble playing where often the leading voice takes a small lead of around 10 ms.

The most obvious structural units in music might be the *metrical* ones, like bar and beat. This strictly hierarchical metrical structure may extend above and below these levels. Special expressive marking of the first beat in the bar, either by timing or articulation, is a common phenomenon.

Another important structural unit is the *phrase*. Phrases may not be ordered in a strict hierarchy (they might overlap), and may cut across metrical structure. Phrase final lengthening is the most well-known way in which they are treated in relation to timing.

Rhythmic structure is important because a large proportion of the timing variance can be attributed to rhythmical groups. Some standard rhythmical patterns, like triplets, even seem to have a preferred and often used timing profile.

And finally, any *associative relation*, e.g. between a musical fragment and its repetition, can be given intentional expression by using the same or different timing patterns.

Generative models

There are three well-known models generate expressive timing from a structural description of the music. The first is Clynes' model based on the so-called "composer's pulse" (Clynes, 1984). He proposes composer specific and meter specific, discrete tempo patterns linked to one or more levels of the

metrical structure. This composer's pulse is assumed to communicate the individual composers' personality. In this model all expressive timing stems from metrical units like bar and beat. Systematic research by Repp (1990) showed that the performance of the pulse is somewhat dependent on the rhythmical material of the piece it's applied to.

In general, meter explains the timing deviations found in musical performance to a limited extent, especially in the Romantic period, when the large rubato patterns seem to communicate the phrase structure. That is the idea that Todd (1989) elaborated. He states that the tempo in each phrase follows the following pattern: speeding up through the middle of the phrase, and slowing down at the end. The last phenomenon has its own name: phrase final lengthening. The form of these tempo changes can be formalized by describing the beat length by sampling a parabolic curve. He then proposes that this process takes place at each level in a nested hierarchy of phrases and sub-phrases. The results at each level combine to yield the tempo for the whole piece.

This model is using tempo curves linked to the phrase structure. Most notable in phrase final lengthening is the interaction with the actual rhythmical material. If a rubato is very strong it may come close to altering the rhythm itself. Todd does not specify a method to yield appropriate parameter values for the curves, they have to be fitted to a real performance, and predictions about for instance the adaptation of the curves to global tempo is not explained. And again, it is only a partial model; all structure other than phrase structure is ignored.

It is quite remarkable that a formalized generative model of the link between rhythmical grouping structure and expressive timing does not exist. There is evidence for a systematic way in which certain rhythmical patterns are played: like a triplet in the context of a duple meter. But a general theory is still lacking, even though there is considerable evidence that rhythmical grouping is responsible for a large proportion of the timing variance.

Sundberg and his colleagues (Friberg et al, 1991) propose a generative model that mostly ignores hierarchical structural descriptions and concentrates on the surface structure of the music: local features and patterns found in the note-by-note description of the score. They describe a rule-based system to generate expression from a score, based on this surface structure. The research was done in an analysis-by-synthesis paradigm and captures expert intuition in the form of a large set of rules. An example of a rule is "Faster uphill": a duration of a note is shortened in performance if it is preceded by a lower pitched note and followed by a higher

yield a combined expressive timing profile. Moreover, all of these models make, either implicitly or explicitly, use of the Tempo Curve in describing the timing aspects independent of the musical material.

The use of a Tempo Curve as an independent reality

Now let's have a close look at the way most authors present timing patterns. A typical example is shown in Figure 3.

Again, on the horizontal axis, score time is represented, and on the vertical axis local tempo. Most timing or tempo measurements are presented in the form of a continuous curve instead of just a scattergram of

Figure 3. A Tempo Curve (local tempo versus metrical time).

pitched one. There is a large set of these rules available (ca. 25) and the set is still growing. It turns out that the use of one or two rules is quite effective, but larger rule-cocktails are hard to judge. It is unclear how rules interact and which classes of rules are dependent. Furthermore, the effectiveness of a rule depends heavily on the material, the piece it is applied too (the authors use another example for every rule).

To conclude with respect to these three generative models, one can say that they indeed model part of the link from structure to expression, but it is not at all clear how these incompatible structural descriptions

measurements, the measurement points are, like here, connected by straight lines. These curves more or less imply an independent existence, apart from the rhythmic material from which they were measured. But we have to realize that we cannot perceive timing or tempo without events carrying it (the psychologist James J. Gibson even wrote an article called "Events are perceivable, but time is not"). And vice versa: "filling up" a time interval by adding an event between two measured points is problematic, because it will change the perceived duration of the original interval.

A critical test should reveal the consequences of representing tempo curves as

an abstract entity. If it has indeed an existence of its own, we should be able to detach it from the rhythmic material and map it to another piece. For instance, it should be possible to map the timing of a theme to the variation of that theme (especially, if it has the same metrical and harmonic structure, and "only" the rhythmic and melodic material differ).

What we then need is a way to invent the tempo for the notes in a denser variation, between the local tempo measurements

applied to a group of notes in the variation, constantly making little jumps in tempo (see Figure 4a).

If we follow the authors that draw tempo curves of straight line segments, we would get a linear interpolation (see Figure 4b).

Wessel and his colleagues (1987) worked at even smoother interpolations, using splines (see Figure 4c). This, indeed, is a smooth way of using rubato and it will often sound better than the previous methods, but it is still far from a real performance, because it lacks the essential link to the structural details of the music.

Figure 4. a) Stepwise interpolation, b) linear interpolation and c) an interpolation using splines. (Black dots are measured, white dots are interpolated).

made for each note in the theme. Figure 4 shows some ways in which such interpolation could be realized.

One way is to assume that the tempo stays constant during the period between two measured points. This is the method Mathews uses in his conductor programs (Boulanger, 1990). He links a change of tempo to one level of the metrical structure (in fact to any arbitrary event). But one can image that it will sound a bit "jumpy", since one tempo measurement in the theme will be

To summarize these proposals: they all have the important restriction of representing 'horizontal timing' only; 'vertical timing', i.e. the timing of events that happen at the same score time (like certain ornaments, chords, or the timing between different voices) is not explicitly represented and therefore cannot be transformed. Some proposals make a link between a Tempo Curve and one type of musical structure (Mathews to one level of a metrical structure, Jaffe to voices, Clynes to a metrical hierarchy, Todd to phrase structure, and Sundberg to local structure). But a proper representation of timing and tempo should at least support these and other musical structures.

Conclusion

We hope to have shown that we must be aware of the Tempo Curve. Of course, one should be encouraged to measure tempo curves and use them for the study of expressive timing. But it is a dangerous notion, despite its widespread use, because it lulls its users into the false impression that it has a musical and psychological reality, and that time can be perceived independent of the events carrying it.

There is no abstract tempo curve in the music nor is there a mental tempo curve in the head of a performer or listener. Any transformation or manipulation of timing on the implied characteristics of such a notion is doomed to fail. For example, the application of a tempo curve to material other than from which it was derived, yields awkward performances, and even a

simple change of a parameter like global tempo fails a simple listening test.

Research in expressive timing has shown that timing and structure are closely linked. Phrase structure and metrical structure make important contributions to the timing profile of performances as are local structural aspects like ornaments and chords. A representation of timing should be expressed in terms of these musical structures, it cannot just ignore them. That does not mean that generic models that represent timing in terms of some sort of structure, even when they describe just a fraction of the many aspects of expressive timing, do not constitute a valuable contribution to the field. They only have to be seen in a proper perspective in which their limitations are understood as well. It also does not mean that certain features in computer music software and commercial sequencers, like tempo knobs and tempo tracks, should be forbidden. Their mere existence at least makes the realization of their limited worth evident.

We are currently, compelled by our own argument, working on alternative representations that form a basis for transformations that do make more musical sense. A first result of this research is described as a calculus for expressive timing on the basis of structural descriptions (Desain & Honing, 1991b).

References

- ANSI (American National Standards Institute) (1989) X3V1.8M/SD-6 Journal of Development Standard Music Description Language (SMDL). San Francisco: Computer Music Association.
- Boulanger, R. (1990) Conducting the MIDI Orchestra, Part 1: Interviews with Max Mathews, Barry Vercoe and Roger Dannenberg. Computer Music Journal 14(2).
- Clarke, E.F. (1988) Generative principles in music performance. In Generative processes in music. The psychology of performance, improvisation and composition, edited by J. A. Sloboda. Oxford: Science Publications.
- Clynes, M. (1984) The secret life of music. In Proceedings of the 1984 International Computer Music Conference. San Francisco: Computer Music Association.
- Dannenberg, R., L. M. Dyer, G. E. Garnett, S. T. Pope, & C. Roads. (1989) Position papers. In Proceedings of the 1989 International Computer Music Conference. San Francisco: Computer Music Association.
- Desain, P. & H. Honing (1991a) Tempo curves considered harmful. In "Music and time", edited by J. D. Kramer. Contemporary Music Review. London: Harwood Press. (forthcoming).
- Desain, P. & H. Honing (1991b). Towards a calculus for expressive timing in music. Research Report. Utrecht: Center for Knowledge Technology. Submitted to the Psychology of Music.
- Friberg, A., L. Frydén, L. Bodin & J. Sundberg (1991) Performance Rules for Computer-Controlled Contemporary Keyboard Music. Computer Music Journal, 15 (2).
- Gibson, J. J. (1975) Events are perceivable but time is not. In The Study of Time, 2, edited by J.T. Fraser & N. Lawrence. Berlin: Springer Verlag.
- Honing, H. (1991) Issues in the Representation of Time and Structure in Music. In: Proceedings of the 1990 Music and the Cognitive Sciences Conference, edited by I. Cross and I. Deliège. Contemporary Music Review. London: Harwood Press. (forthcoming).
- Jaffe, D. (1985) Ensemble timing in Computer Music. Computer Music Journal, 9(4).
- Palmer, C. (1989) Mapping musical thought to musical performance. Journal of Experimental Psychology, 15(12).
- Repp, B. (1990) Further Perceptual Evaluations of Pulse Microstructure in Computer Performances of Classical Piano Music. Music Perception, 8(1).
- Shaffer, L.H., E.F. Clarke, & N.P. Todd (1985) Metre and rhythm in piano playing. Cognition, 20.
- Todd, N. (1989) A Computational Model of Rubato. In "Music, Mind and Structure", edited by E. Clarke and S. Emmerson. Contemporary Music Review 3(1).
- Wessel, D., D. Bristow, & Z. Settel (1987) Control of Phrasing and Articulation in Synthesis. In Proceedings of the 1987 International Computer Music Conference. San Francisco: Computer Music Association.