

TOWARDS A CALCULUS FOR EXPRESSIVE TIMING IN MUSIC

Peter Desain & Henkjan Honing

[Published as: Desain, P., & Honing, H. (1991).
Towards a calculus for expressive timing in music.
Computers in Music Research, 3, 43-120.]

INTRODUCTION

In Desain and Honing (1992a) we argued that a simplistic notion of a tempo curve of a musical performance is a dangerous and harmful theoretical construct. Although the use of a tempo curve to describe time measurements is perfectly sound, the notion itself is often presented as a cognitive or musical concept. And tempo curves do not have any right to exist in those domains. In the above article, this was concluded from the fact that when it is used as a basis of transformations, inevitably the results make no musical sense. The cause of this failure can often be attributed to the lack of structural information in the tempo curve. For example, in changing the overall tempo of a performance, by manipulating the tempo curve alone, all time intervals of equal length between two notes are scaled in the same way. But some notes may constitute a particular kind of ornamentation, whose duration should be more or less unaffected by tempo. As a result the timing of the piece becomes unmusical. And there are many more examples of transformations that cannot be done on isolated tempo curves. Because the article had an essentially negative tone - identifying the problems and their causes - we felt compelled to follow it up with a study of possible solutions.

This paper is an attempt to identify ways in which structural knowledge can be used to enable expression transformations on musical performances that do make musical sense.

In past research we considered expression merely as deviations of attributes of performed notes from their value notated in a score. This definition, however useful in the initial study of expressive timing, soon lost its attractiveness. In general, listeners can appreciate expression in music performance without knowing the score. And a full reconstruction of the score in the form of a mental representation is often impossible. Take for instance the notion of loudness of notes. Should a listener be required to fully reconstruct the dynamic markings in the score before it is possible to appreciate the deviations from this norm as expressive information added by the performer? Such a nonsensical conjecture indeed follows from a rigid definition of expression as deviations from the score. But it is possible to find ways of defining expression on the basis of performance information only. The more so since it became possible to model the quantization of performed note durations into discrete categories (Desain & Honing, 1991), and therefore even the extraction of performed tempo is possible directly from the performance itself.

In this paper we will base expression on the notion of structural units in a working definition: *expression within a unit is defined as the deviations of its parts with respect to the*

norm set by the unit itself. An example might make this more clear. Lets take, for instance, a metrical hierarchy of bars and beats; the expressive tempo within a bar can be defined as the pattern of deviations from the global bar tempo generated by the tempo of each beat. Or, take the loudness of the individual notes of a chord; the dynamic expression within a chord can be defined as the set of deviations from the mean chord loudness by the individual notes. Using this intrinsic definition, expression can be extracted from the performance data itself, taking more global measurements as reference for local ones, assuming that the structural units themselves are known. Thus the structural description of the piece becomes central, both to establish the units which will act as a reference, and to determine its subunits that will act as atomic parts whose internal detail will be ignored. A generalization of this concept can also deal with expression arising from the interplay of two or more voices.

It will be clear by now that any other connotations of the concept of musical expression, its link to human affect and extra-musical indexicality, however interesting, will be ignored here completely.

Before the details of the calculus are presented it might be fitting to give some explanation for undertaking for this work. First of all, we think that the research of expression in music is in need of measurement instruments that can cope with the enormous complexity of performance data and that are much more sophisticated than tempo curves. Some of the proposed transformations can be used as an “auditory microscope” by exaggerating expression at certain structural levels, like amplifying the timing lead, the melody often has over the accompaniment. Some of the tools presented can be used as “expression scalpels” for trimming away certain kinds of expression that might obscure other phenomena, like removing the tempo deviations within each beat, but holding the timing patterns of the beats themselves invariant. Other tools can “transplant” musical expression from one piece of music to the other, say from a theme to its variation. The availability of this ‘machinery’ will deepen our understanding of the intricacies of music performance expression.

A further motivation is the practical applicability of this work in systems for computer music. Especially the music editors and sequencer programs that are commercially available nowadays which are in need of better ways to treat musical information in musical ways. Expressive timing should not be considered a nasty feature of performed music, as it is in nowadays multi-track recording techniques where tempo, timing and synchronization are treated as technical problems. Instead expressive timing has to be regarded as an integral quality of performed music whereby the performer communicates structural aspects of the music to the listener (Clarke, 1988). We hope that our work can inspire new music software based on this view.

OVERVIEW OF THE CALCULUS

Characteristics

The calculus has the following important characteristics:

The calculus is described here only for different brands of expressive timing. Dynamics could be formalised along the same lines, but for clarity we restrict ourselves to the domain of expressive timing. Other attributes that carry expression, like intonation, vibrato and timbre may require a different treatment.

The types of expression have to be computable to be within reach of this calculus. One must be able to calculate the expression at every level of the structural hierarchy, given the expression of their components (e.g. the timing of a chord must be computable when the timing of the embedded notes is given). One also must be able to state ways to effectively set the expression of the components once the expression of the whole is given (i.e. propagate a shift in timing down the hierarchy, to the basic objects carrying the expression). Types of expression that do not have this characteristic - or are not yet formalised as such- cannot be described.

Both performance and "score" timing of individual notes are clearly defined. Notes require attributes that can be measured more or less directly from the performance data like the note onset time and the offset time. At least the onset time must be clearly specified, which makes the calculus less appropriate for expressive performances by instruments for which onset times are not so clear cut. Secondly, the metrical note duration (the timing of the note as notated in the score) must also be available as a note attribute - either via quantization or by matching a performance to a known score. These processes are considered preprocessing here. Although the reference to score duration, score onset and score offset times is less appropriate in the context of our definition of expression - we will use this terminology, for lack of better terms.

The "score" timing of rests is clearly defined. Perhaps surprisingly, the rest plays a key role in some transformations. So we assume that it either can be inferred from the performance timing (Longuet-Higgins, 1976 shows a way of doing so), or it is recovered via the matching of a performance and a known score.

All proposed transformations are structure preserving. This means that the calculus is restricted to true expressive transformations: the score timing of the notes is known and fixed, and transformations will leave this and the structural description invariant.

The behaviour of musical material under expressive transformations is determined uniquely by its structural description and the type of expression.

The transformations are defined on a hierarchical structural description uniquely linking all material. Ambiguous structural descriptions (e.g. two or more possible structural descriptions) or incomplete descriptions cannot be dealt with. The obvious need for knowledge representations containing multiple structural descriptions (metrical, phrase, and rhythmical grouping structures, different analysis etc.) is not denied. We just require that such representations be preprocessed to select only one complete structural description. This is not a real restriction since transformations based on different kinds of structural knowledge of the same piece can always be done

in sequence. Re-inserting the trimmed structural descriptions into a transformed piece is trivial because the transformations preserve the structure.

Naturally, the higher-level structural description of the piece must be consistent with the performance timing. For example, a structural description of the piece in which two notes are given a certain sequential time order (one after the other) - can only fit a performance in which at least the onset of the corresponding notes obey the same order. The precise rules will be given when the structural descriptions are introduced.

The transformations are defined to apply to a certain level of the hierarchical structural description, ignoring details from lower levels and keeping higher levels invariant. Means to select such a level are assumed. In sophisticated realisations of the calculus this may entail a match language ("the first bar of the piano solo that begins with a C") or a graphical representation. In this paper we will simply assume that each musical object has a name as attribute and defines a structural level as the set of objects with a certain name.

Although the calculus separates different kinds of behaviour, it entails no musical knowledge of the transformations themselves. Accordingly, the proposed knowledge representation does support for example, arbitrary descriptions of the metrical structure of a piece, but has no knowledge of "the best structural analysis". To give a second example: the proposed knowledge representation does support ways to modify timing (a)synchrony between voices, but it has no knowledge about correct or effective ways of using this in musical performance.

The calculus also does not model cognition. It does not state how, for example, voice-leading helps auditory streaming, how phrase final lengthening beyond a limited range disables rhythm perception, or how structure is communicated by the expressive timing profiles. However, this work constitutes a solid basis for formalised theories about these issues, providing a powerful representation in which they can be expressed.

Representation

Several concepts are used in the calculus:

Musical objects are either of a basic nature or form a structural description of a collection of musical objects. Basic musical objects consist of notes and rests. Notes are the only musical objects that carry the expressive information. Structural descriptions form collections of musical objects. They may describe hierarchical time intervals like metrical-, phrase- or rhythmical grouping, they can group the notes of chords and ornaments together, or form large horizontal slices through the piece, describing the separate voices etc. Mere collections (sets) of objects are too meager a basis for most transformations, therefore, structural descriptions specify the intended relations in time between these objects as well (Honing, 1992). Most transformations can be defined if two orthogonal characteristics of the structural description are given: the temporal nature and the ornamenting quality. The first describes whether a sequence or a parallelism (a so called successive or simultaneous construct) is represented. The second describes whether the musical object is considered an ornament attached to

another object or not. Ornaments are shielded from certain modifications and refer to another object for certain attributes. These two binary characteristics result in four concrete types of structural description that will be described in detail later.

Expressive magnitudes are values of expressive measurement on a certain scale. The scales themselves are of course crucial in modeling effective transformations, in cognitive and musical senses. For example, a tempo scale on which a transformation to make something twice as fast actually yields a double perceived tempo is quite useful. But for the sake of simplicity we abstract from the perceptual processes and the instruments that generate the sound, and will just assume simple physical measurements of time and other expressive attributes.

Expression maps describe the expressive patterns of structured musical objects at a certain structural level. They consist of a section for each musical object at that level. A section lists the expressive values for all components of that object. They come in different brands - consistent with the type of musical structure where they were extracted from. Expression maps can be extracted from and applied to musical objects, with possibly a modification of the map in between.

Expression types are sets of procedures to extract a particular type of expression map from a musical object, to impose it on a musical object, and to modify the map. They capture the difference between expressive tempo, asynchrony, and articulation. They may become fairly sophisticated, like a brand of expressive tempo that knows how to keep the articulation of an individual note invariant when the timing of the note onsets is changed.

Modifications are defined as operations on expression maps. They may scale, interpolate, or do any other operation on the map. They are often designed such that certain characteristics are kept invariant, e.g., the total duration of a section while changing the timing of the parts.

Transformations are defined as operations on musical objects. They are often direct generalizations of the expression map modifications - first extracting the map, applying the modification and imposing the modified map. They also handle the selection of the level of structural description on which to apply the transformation. Furthermore, they may have means to maintain consistency among the affected level and other musical material, e.g., making an accompaniment obediently follow the transformation in expressive tempo applied to the melody.

Implementation

Part of the work described in this paper was done in the design of the POCO system (Honing, 1990) for which a scaling operation of expressive timing linked to structural descriptions was implemented. But, in evaluating this rather complex piece of software, better abstractions arose. Especially the design of a set of data structures for

music that capture the differences in behaviour under transformation proved beneficial. Which again illustrated the adage:

“Get your data structures correct first, and the rest of the program will write itself.” (David Jones, quoted in Bentley, 1988)

Because the constructs interact heavily, and because it should be easy to add unforeseen new constructs (like a new type of expression), musical objects, expression maps and expression types are implemented as classes in an object-oriented language. In that way it is easy to express modifications and transformations as polymorphic operations that will behave according to the type (the class) of their arguments. The slicing-up of knowledge in classes means answering questions like: which part of the extraction procedure of an expressive tempo map of a sequential musical object is specific for expressive tempo only and should be stated within the expression type; which part only depends on the sequential nature of the musical structure, and should be part of the class for sequential musical objects; and which part describes the creation of an expression map and belongs to that class?

Although a good Object-Oriented Language (we used CLOS, a Lisp-based system) provides one with the programming-constructs needed to express these concepts, the actual process of factoring knowledge into these polymorphic procedures is still a difficult one, especially because during the design of the best structure of the classes - allowing for the most elegant factoring of the procedures - cannot be completely foreseen. This forced us to go through several re-design rounds before the concepts stabilized in their present form.

The following CLOS (Keene, 1989; Steele, 1990) constructs were used heavily in the implementation: multiple inheritance (forming class dependencies that are more complex than simple hierarchies), multi-methods (functions that are polymorphic in more arguments), mix-in type of inheritance (grouping of partial behaviour in an abstract class that must be mixed in with other classes to supply that behaviour to their instances), method combination (providing ways of combining partial descriptions of behaviour of one method for more classes). Together they make it possible to extend the system by adding program code only, instead of rewriting it.

The calculus will be incorporated in POCO. The other tools available in POCO, like score-performance matchers, multiple structural descriptions, storage and retrieval from standard MIDI-files, playback and editors for music text formats etc., will support a comfortable use of the calculus with real performance data. An implementation in the form of the microworld is given in the appendix and aims at conciseness and elegance. Luckily, this goal only occasionally conflicts with computational efficiency.

The following five paragraphs will describe the calculus in more detail. The reader interested in the more general aspects of the calculus is advised to continue reading below “Transformations”.




Figure 1. Classes of musical objects and their interrelations.

MUSICAL OBJECTS

Musical objects come in different brands. Some types are specific enough to describe an object completely (the instantiatable or concrete classes). Other types are used as a descriptive grouping of likewise behaviour (the abstract classes). The types of musical objects and their interrelations are shown in figure 1.

Basic musical objects

Basic musical objects are notes and rests (we use the word PAUSE to avoid the name clash with the Common Lisp primitive REST). In examples we will use notes with clearly observable onset and offset times (called PERF-ONSET and PERF-OFFSET) measured in ms. from the beginning of the performance. Both notes and rests have as a property a time position in the score (called SCORE-ONSET and SCORE-OFFSET) measured in any kind of (beat)-count (a rational number). These score times are calculated automatically from the supplied score durations of notes and rests via the structural descriptions. This facilitates easy creation of large scores.

Rests are essential and cannot just be ignored, as is done in some low-level representations (e.g. the Midi-file standard). They are central e.g. in dealing with articulation - a short note followed by a rest behaves differently under transformation than a longer note played in a staccato way.

Structured musical objects

Multilateral structures

In research on music perception and cognition a distinction is often made between successive temporal processes that deal with events occurring one after another, and simultaneous temporal processes that handle events occurring around the same time (e.g. Bregman, 1990; Serafine, 1988). For the first type of events of the expressive means can be rubato - the change of tempo over the sequence. In the second one the expressive means can be chord-spread and asynchrony between voices, both more timbral aspects. These processes work differently in perception. Since we want to imply differences in behaviour mainly by differences in structural description a way should be found in which both these constructs can be represented.

We propose to use for this purpose the simple time structures S and P that functioned well in (Desain & Honing, 1988; Desain, 1990; Desain & Honing, 1992b). If a collection of musical objects is formed such that they occur one after another they are described as a successive structured object named S (for Sequential). If a collection of musical objects occur at the same time they can be collected in a simultaneous structured object called P (for Parallel). These structures serve as a general way to represent a collection together with the temporal relation between the components, as stated in the score. We call the objects *multilateral* because their components are considered to be of equal importance, and are to be treated as such in expressive transformations.

The score times of a structured object and its parts are constrained by consistency rules. They are described separately in frames 1 and 2. These constraints are enforced by specifying only notes and rests with a score duration. The constraints propagate these automatically when a structural description is created and set all score onset and offset times.

In calculating expression, the previous and subsequent context of musical objects is sometimes needed. For instance, consider articulation: possibly defined as the overlap between the sounding parts of a note and the next one, i.e. the time difference between the offset of the note itself and the onset of the "next" note. Besides "next material" a link to "previous material" is foreseen to be needed as well, e.g. in the calculation of local accent patterns. To formalize and generalize this notion of "previous" and "next" material a definition of the left and right context of a musical object is given. This notion also reflects the fact that some expressive values cannot be calculated because some contexts are not available or carry no expression e.g. the tempo of the last note in a piece, or the performance onset of a voice that starts with a rest. Expressive transformations must thus expect to come across missing values in an expression map. The notion of context is explicitly represented in the program as attributes of the objects themselves. This is possible because the structural description is invariant and so are the contexts. Another possibility would be to represent them implicitly, recovering them by search via a bi-directional part-of link between musical objects. Alternatively, they could be represented tacitly, i.e. supplying them by a general control structure that walks through structured musical objects.

Collateral structures (ornamented objects)

Some musical objects contain components that should maintain a dependency relation to one another. If such a *collateral* pair is transformed, the transformation should be carried out on the main component only, the submissive one obediently following the main component's transformation, but not being transformed itself. An ornamented musical object like a graced note (a note preceded by a grace note), is a good example of a collateral object. For example, in the scaling of the expressive tempo of a melody which contains a graced note, the data on which the expressive transformation is carried out (in this case the performance onset) stems from the main object. The grace note is ignored. When in the actual transformation the graced note pair is stretched or compressed and moved to an other point in time, only the main note will undergo that operation. The ornament will just follow its shift in time.

A second use of this concept is made when the relation of an ornament to its main object, within such a collateral couple, is considered to be expressive, and a potential source of expressive transformations. In this case, the main object stays invariant, and only the ornament undergoes transformation. Take for example the asynchrony between the performance onset of a grace note and the note it is attached to. This time interval can be modified by appropriate means, resulting in local changes of the timing of the grace note - but keeping the timing of the main note invariant.

Collateral (ornamented) objects can again have two kinds of temporal nature: successive or simultaneous. The first one is called APPOG (for *appoggiatura*). It describes a "time-taking" ornament where the ornament is considered to finish when its main object starts (all in terms of score times). The second is called ACCIA (for *acciaccatura*). It can represent a so called "time-less" ornament that is supposed to start at the same time as the object it is attached to. Note that both parts of a collateral pair are musical objects themselves and can have internal structure. The concepts of APPOG and ACCIA ornamented objects are an elaboration of the PRE and POST objects that were introduced in (Desain & Honing, 1988). Consistency rules for score times and context are described in frames 3 and 4.

S, a multilateral successive structure




Figure of S object

Consider an S structure of n components C_i with $0 \leq i \leq n-1$.

Assume that component C_i has score onset time Son_i , score offset time $Soff_i$ and that the whole structure has score onset time Son and score offset time $Soff$. Then the following must hold:

$$Son = Son_0$$

$$Soff = Soff_{n-1}$$

$$Soff_i = Son_{i+1}, \text{ for } 0 \leq i \leq n-2$$

Assume that component C_i has performance onset time Pon_i and that the whole structure has performance onset time Pon . Then the following must hold:

$$Pon = Pon_0$$

$$Pon_i \leq Pon_{i+1}, \text{ for } 0 \leq i \leq n-2$$

Assume that component C_i has left context L_i and right context R_i and that the whole structure has left context L and right context R . Then the following holds:

$$L = L_0$$

$$R = R_{n-1}$$

$$R_i = C_{i+1}, \text{ for } 0 \leq i \leq n-2$$

$$L_i = C_{i-1}, \text{ for } 1 \leq i \leq n-1$$

Frame 1. Description of a S structure.

P, a multilateral simultaneous structure

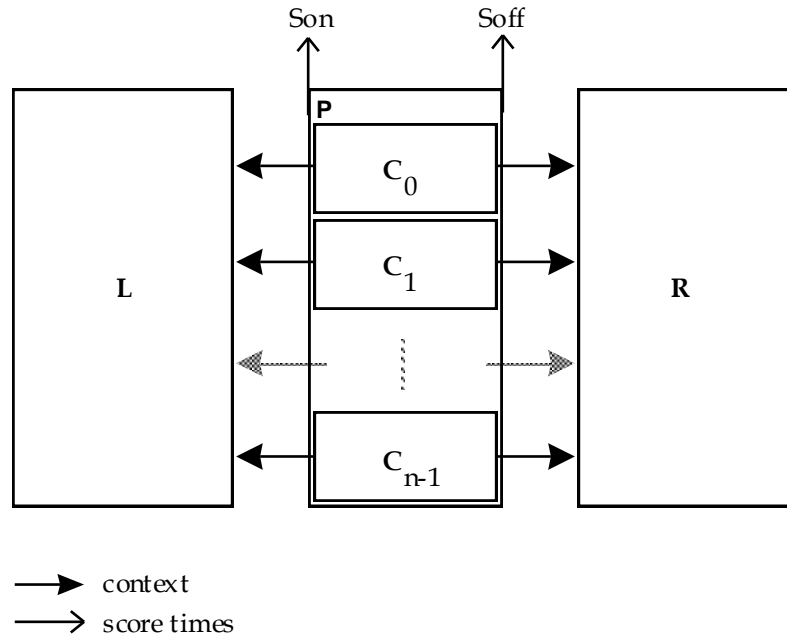


Figure of P object

Consider a P structure of n components C_i with $0 \leq i \leq n-1$.

Assume that component C_i has score onset time Son_i and score offset time $Soff_i$ and that the whole structure has score onset time Son and score offset time $Soff$. Then the following must hold:

$$Son_i = Son, \text{ for } 0 \leq i \leq n-1$$

$$Soff_i = Soff, \text{ for } 0 \leq i \leq n-1$$

Assume that component C_i has performance onset time Pon_i and that the whole structure has performance onset time Pon . Then the following holds:

$$Pon = \text{MIN}_{0 \leq i \leq n-1} Pon_i$$

Assume that component C_i has left context L_i and right context R_i and that the whole structure has left context L and right context R . Then the following holds:

$$L = L_i, \text{ for } 0 \leq i \leq n-1$$

$$R = R_i, \text{ for } 0 \leq i \leq n-1$$

Frame 2. Description of a P structure.

APPOG, a collateral successive structure

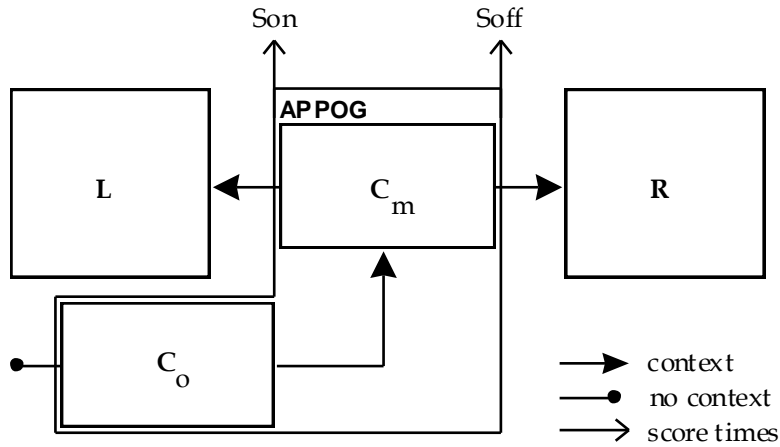


Figure of APPOG object

Consider a APPOG structure of a ornament component C_o and a main component C_m .

Assume that component C_o has score onset time Son_o , score offset time $Soff_o$, that component C_m has score onset time Son_m and score offset time $Soff_m$ and that the whole structure has score onset time Son and score offset time $Soff$. Then the following must hold:

$$\begin{aligned} Son_m &= Son \\ Soff_m &= Soff \\ Soff_o &= Son_m \end{aligned}$$

Assume that component C_o has performance onset time Pon_o , component C_m has performance onset time Pon_m and that the whole structure has performance onset time Pon . Then the following holds:

$$\begin{aligned} Pon &= Pon_m \\ Pon_o &\leq Pon_m \end{aligned}$$

Assume that component C_o has left context L_o and right context R_o , component C_m has left context L_m and right context R_m and that the whole structure has left context L and right context R . Then the following holds:

$$\begin{aligned} L &= L_m \\ R &= R_m \\ R_o &= C_m \\ L_o &= \text{undefined} \end{aligned}$$

Frame 3. Description of an APPOG structure.

ACCIA, a collateral simultaneous structure

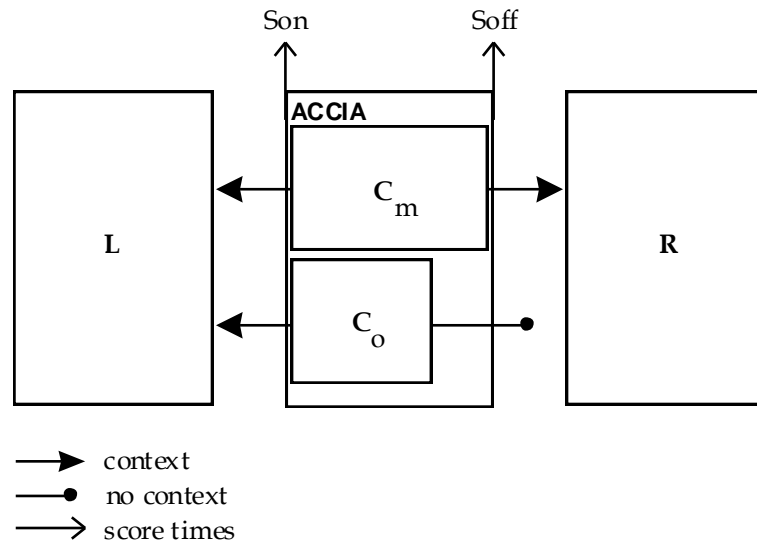


Figure of ACCIA object

Consider a ACCIA structure of a ornament component C_o and a main component C_m .

Assume that component C_o has score onset time Son_o , score offset time $Soff_o$, that component C_m has score onset time Son_m and score offset time $Soff_m$ and that the whole structure has score onset time Son and score offset time $Soff$. Then the following must hold:

$$Son_o = Son_m = Son$$

$$Soff_m = Soff$$

Assume that component C_o has performance onset time Pon_o , component C_m has performance onset time Pon_m and that the whole structure has performance onset time Pon . Then the following holds:

$$Pon = Pon_m$$

Assume that component C_o has left context L_o and right context R_o , component C_m has left context L_m and right context R_m and that the whole structure has left context L and right context R . Then the following holds:

$$L = L_m = L_o$$

$$R = R_m$$

$$R_o = \text{undefined}$$

Frame 4. Description of an ACCIA structure.

EXAMPLE OF THE REPRESENTATION OF A MUSICAL OBJECT

In figure 2 a fragment of a score is shown that will serve as a basis for the examples at the end of this article. It is the score of the last bars of the theme of six variations over the duet *Nel cor più non mi sento*, by Ludwig van Beethoven (with some adaptations), which is the same material used to study tempo curves in (Desain & Honing, 1992a). It contains examples of several kinds of musical structure: chords, voices, sequences, bars and beats, phrases and two types of ornaments. Figure 3 shows a graphical notation indicating two structural descriptions: a metrical hierarchy and a separation into voices. The way these structures are specified in Lisp is given in the appendix.




Figure 2. Score of the last bars of the theme of six variations over the duet *Nel cor più non mi sento*, by Ludwig van Beethoven.




Figure 3. a) Structural description of the metrical hierarchy of the score in figure 2, and b) Structural description of the voices in that piece.

REPRESENTING EXPRESSION

There are three kinds of expressive timing: expressive tempo, expressive asynchrony and expressive articulation. The first two are based on performance onset times only, the third is based on performance onset and offset times (see figure 4).

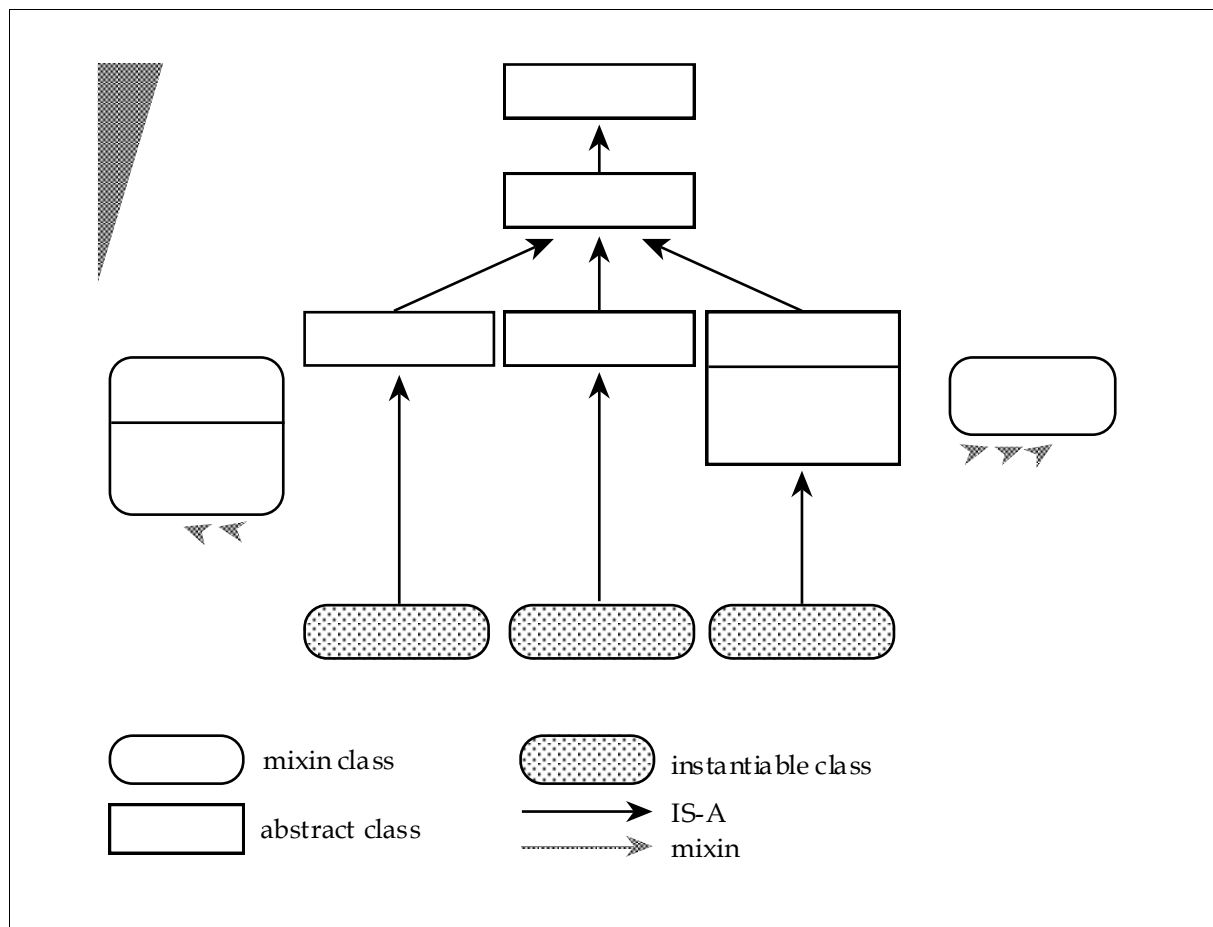


Figure 4. Expression type hierarchy.

One could imagine sophisticated algorithms that calculate the onset of a note and of parallel structures on the basis of their perceptual onset (P-center; see Vos & Rasch, 1981). But for clarity we use a very simple definition of onset times, which was already given in the frames 1 to 4. In that way, all musical objects have performance onset times and so can be used as units on which tempo and asynchrony measures are built.

Expressive tempo

The notion of tempo is relevant only for successive structures. It is defined as the ratio of score duration and performance duration. This velocity-like notion the inverse of the notion of a tempo factor, as is used in the psychology of music literature.

Expressive asynchrony

The notion of asynchrony is relevant only for simultaneous structures. It is defined as the difference of performance onsets. It is thus independent of score times.

Expressive articulation

Expressive articulation uses the performance offsets of individual notes. It simply assumes that they are given. A definition of performance offset of structured musical objects is not needed. Articulation is also independent of score times.

Articulation can be defined in several ways - but it is hard to find a way that will suffice in all circumstances. In the legato range the absolute overlap time of the sounding part of a note and the next one seems a good candidate for an articulation scale. In the staccato range the absolute sounding duration of the note seems the most prominent perceived aspect. In the intermediate range the relative sounding proportion is a good measure. For the moment we cannot do better than to supply these three concepts of articulation expression (overlap-, duration- and proportion-articulation) - leaving it for the user to choose the most appropriate one (see frame 5). For a multilateral structure the expressive articulation value is taken to be the average articulation of its parts. For a collateral structure the expressive articulation value is defined to be the articulation of its main part.

Definition of articulation

Consider a note with performance onset P_{on} , performance offset P_{off} and performance onset of its right context P_{on_r} . There are three alternative definitions of articulation A given:

overlap articulation $A = P_{off} - P_{on_r}$

duration articulation $A = P_{off} - P_{on}$

proportion articulation $A = \sqrt{F(P_{off} - P_{on}, P_{on_r} - P_{on})}$

If a multilateral structure with articulation A has components C_i for $0 \leq i \leq n-1$, and C_i has articulation A_i then:

$$A = \text{MEAN}_{0 \leq i \leq n-1} A_i$$

If a collateral structure has articulation A , and its main component C_m has articulation A_m then:

$$A = A_m$$

Frame 5. Definition of articulation expression.

Estimate onsets

Because sometimes the performance onset of missing objects (like the virtual note after the end of the piece) or the performance onset of a rest are needed, we devised a set of procedures that estimates these missing values on the basis of performance onsets that can be found in the context, using a linear interpolation or extrapolation method. The set of procedures forms a mix-in class that can be combined with any expressive timing type enabling that kind of expressive timing to deal - in all operations - with missing values. Estimation is derived from the same structural level as the transformation itself. For example, a transformation on a beat structure in need of a missing expressive value at the end of the piece (cf. the onset the final barline in a score) will be estimated on basis of the two previous beats -not on the basis of any internal detail. In the case of extreme tempo variations, as occur in a final retard, the estimation feature cannot work well. In this case it is better not to use it.

Articulation invariance

When moving the onsets of notes around (e.g. in modifying the performance onsets) it is quite annoying that the articulation of the individual notes also changes - an effect that is very easy to perceive and which may well overshadow subtle modifications of onset timing. Therefore a set of procedures can be mixed-in with expressive tempo and asynchrony. They are given a chance to calculate the articulation of individual notes before onsets are changed and to reinstall it afterwards. This will insure that articulation is kept invariant under transformations of onset timing (see figure 12).

EXPRESSION MAPS

An abstraction of the expression of an object is useful for many operations because it can hide the irrelevant details of the structure and provides a means to transfer expression from one object to another. Therefore expression maps were introduced. They describe expression of musical objects at one level of a structural description. All objects at the level described must have the same structural type. Maps contain a list of sections, one for each of those musical objects. A section lists the expressive values of the components of that musical object. Of course maps may be partial - consisting of several sections with gaps in between, or even have missing values within a section.

Onset timing

The application of a (modified) map of performance onsets on an object works as follows. First, all objects at the indicated level are found, paired with their corresponding sections. Then each section is applied to its object. This means that the components of that object are provided one by one with a new onset from that section.

This setting of onsets is handled differently according on the structural type of the component. If this component is a note, the onset is set directly. For S components the whole structure is stretched between that onset and and the next onset (the onset of the succeeding component). A P component is set to the provided onset, but keeps its

internal asynchrony invariant and truncates at the next onset. In the case of a ACCIA component, the main structure is set to the onset, with the ornament following the displacement of the main structure. Finally, for a APPOG component, the main structure is stretched between that onset and the next onset, with the ornament also simply following the displacement of the main object.

Now we have indicated how an expressive timing is applied to components of structured objects - it remains to be shown how such a change propagates when these components again are embedded structured objects themselves. This fairly complex process depends on the type of the embedded structured object and mirrors the decisions given above: S components are stretched, P components are shifted and truncated, and ornaments follow the shift of their main components.




Figure 5. Propagation of change of onset within an S structure for different component types. This figure shows the propagation process for an S structure containing different types of structural components. We assume the components are moved around by an arbitrary transformation, parametrized by a factor. In this figure it is shown how this change is propagated to the internal

structure of different kinds of components. The first component is an S structure and the onsets of its internal parts (lines marked with white circles) are stretched along proportionally. The second sub-structure is an APPOG structure and one can see that the onset of its ornament (line marked with upward pointing white triangles) shifts along with the main object. The third sub-structure is an ACCIA structure and the onset of its ornament behaves likewise. Note that the onset of the ornament is allowed to shift freely (line marked with downward pointing white triangles), even the order of notes is allowed to change here. The fourth sub-structure is a P structure and the onset of its components (lines marked with squares) are shifted and truncated at the end (the right context note; line marked with x's).

Articulation expression

In comparison, to set the articulation expression to a structured object is much simpler. When a section of an articulation map is applied to a multilateral or collateral structure the articulation of its components are set to their respective values from the section.

The propagation of a (modified) articulation value to a component works as follows. If that component is a note, a new offset is calculated from the articulation value and set directly, taking care to maintain reasonable offset times (e.g. not shifting before its onset). If that component is a multilateral structure, its articulation is calculated (the mean articulation of its components) and the difference with the required articulation is propagated as an increment to all components. If it is a collateral structure, its articulation is calculated (the articulation of its main component) and the difference with the required articulation is propagated as an increment to both main and ornament components.

OPERATIONS ON EXPRESSION MAPS

Operations on expression maps work section by section. In each section the expression of a structured musical object is represented. The operations delivers a new section to be applied to that object. Care was taken to maintain structural consistency in all operations even in case of extreme parameter values. Of course expression transformations are intended as subtle changes and truncation or extreme normalization should in practice never occur.

Scale maps

Scaling expressive tempo

Scaling tempo is done in an exponential way. Inverse tempi are considered to be related by a scale factor -1; twice as slow is considered to be the mirror image of half as slow. This exponential scaling of expressive tempo mirrors the exponential nature of notated note durations.

Scaling the expressive tempo of an S section

The scaling of the expressive tempo of a multilateral successive structure works as follows. Assume the structure has n components named $C \setminus S \setminus DO3(i)$ with $0 \leq i \leq n-1$.

Assume component $C \setminus S \setminus DO3(i)$ has score onset time $Son \setminus S \setminus DO3(i)$ and performance onset time $Pon \setminus S \setminus DO3(i)$. Assume the right context of the structure (and thus the right context of component $C \setminus S \setminus DO3(n-1)$) is object $C \setminus S \setminus DO3(n)$. It has score onset $Son \setminus S \setminus DO3(n)$ and performance onset time $Pon \setminus S \setminus DO3(n)$. A section of the expressive tempo map of the structure contains all Son_i and $Pon \setminus S \setminus DO3(i)$ including $Son \setminus S \setminus DO3(n)$ and $Pon \setminus S \setminus DO3(n)$. The scale operation on such a section delivers a new section with performance onsets $Pon \setminus S \setminus DO3(i)'$ according to the following rules:

Define the score inter-onset interval $\Delta Son \setminus S \setminus DO3(i)$ and the performance inter-onset interval $\Delta Pon \setminus S \setminus DO3(i)$ and the local tempo $T \setminus S \setminus DO3(i)$ for $0 \leq i \leq n-1$ (a better term would be velocity) as:

$$\begin{aligned} \Delta Son \setminus S \setminus DO3(i) &= Son \setminus S \setminus DO3(i+1) - Son \setminus S \setminus DO3(i) \\ \Delta Pon \setminus S \setminus DO3(i) &= Pon \setminus S \setminus DO3(i+1) - Pon \setminus S \setminus DO3(i) \\ T \setminus S \setminus DO3(i) &= \setminus F(\Delta Son \setminus S \setminus DO3(i), \Delta Pon \setminus S \setminus DO3(i)) \end{aligned}$$

This ratio is scaled by an exponential factor f .

$$T \setminus S \setminus DO3(i)' = T \setminus S \setminus DO3(i) \setminus S \setminus UP3(f)$$

Then new raw performance durations $\Delta Pon \setminus S \setminus DO3(i)''$ are calculated:

$$\Delta Pon \setminus S \setminus DO3(i)'' = \setminus F(\Delta Son \setminus S \setminus DO3(i), T \setminus S \setminus DO3(i)')$$

These are re-normalised such that the total performance duration is kept invariant.

$$\Delta Pon \setminus S \setminus DO3(i)' = \frac{\Delta Pon \setminus S \setminus DO3(i)''}{\setminus F(Pon \setminus S \setminus DO3(n) - Pon \setminus S \setminus DO3(0), \setminus I \setminus SU(i=0, n-1, \Delta Pon \setminus S \setminus DO3(i)''))}$$

Starting at the same point, the new performance times are given as:

$$Pon \setminus S \setminus DO3(i)' = Pon \setminus S \setminus DO3(0) + \setminus I \setminus SU(j=0, i-1, \Delta Pon \setminus S \setminus DO3(j)')$$

Frame 6. Scaling the expressive tempo an S section.




Figure 6. Scaling the expressive tempo of an S section. This process is shown for a specific set of performance onsets Pon_{i} . In this figure the horizontal axis is the performance time P . On the vertical axis the scale factor f is given. Thus at the horizontal line at scale factor 1 the performance times Pon_{i} are shown as markers on the line; they are identical to the original performance times Pon_{i} . This operation (with scale factor 1) is the identity transformation with respect to the performance timing. At the horizontal line at scale factor 0 the performance times Pon_{i} are identical to the score times So_{i} (modulo normalization to the total performance duration). This operation (with scale factor 0) effectively removes the expressive timing of the performance. At factor .5 a diminished expressive timing profile will result, and at factor 2 an exaggerated rubato can be obtained. At negative values of the scale factor the expressive profile is inverted: a slower tempo becomes faster and vice versa. At extreme values of the scale factor the note that is played at the slowest tempo in the performance will gain almost the whole performance time interval spanned by the structure, pushing other notes to zero duration.

When the performance onset Pon_{n} is not available, the scale transformation uses Pon_{n-1} instead, and scales the tempo of the section with regard to the onset of the last component in the section - instead of the onset of the right context. This tempo scaling method works well for S constructs with many components and small tempo deviations.

Scaling the expressive tempo of an APPOG section

The scaling of the expressive tempo of a collateral successive structure works as follows. Assume this structure has a main component with score onset time $Son\ S\ DO3(m)$ and performance onset time $Pon\ S\ DO3(m)$ and a preceding ornament component with score onset time $Son\ S\ DO3(o)$, and performance onset time $Pon\ S\ DO3(o)$. Assume the right context of the structure (and thus the right context of component $C\ S\ DO3(m)$) is object $C\ S\ DO3(r)$. It has score onset $Son\ S\ DO3(r)$ and performance onset time $Pon\ S\ DO3(r)$. An APPOG time map section contains this score and performance data. The scale operation on such a map delivers a new map with performance onsets according to the following rules:

Define the main and ornament score inter-onset interval $\Delta Son\ S\ DO3(m)$, $\Delta Son\ S\ DO3(o)$ and the main and ornament performance inter-onset interval $\Delta Pon\ S\ DO3(m)$, $\Delta Pon\ S\ DO3(o)$ as:

$$\begin{aligned}\Delta Son\ S\ DO3(m) &= Son\ S\ DO3(r) - Son\ S\ DO3(m) \\ \Delta Son\ S\ DO3(o) &= Son\ S\ DO3(m) - Son\ S\ DO3(o) \\ \Delta Pon\ S\ DO3(m) &= Pon\ S\ DO3(r) - Pon\ S\ DO3(m) \\ \Delta Pon\ S\ DO3(o) &= Pon\ S\ DO3(m) - Pon\ S\ DO3(o)\end{aligned}$$

The ornament tempo $T\ S\ DO3(o)$ and the main tempo $T\ S\ DO3(m)$ are calculated as:

$$\begin{aligned}T\ S\ DO3(o) &= \sqrt{f(\Delta Son\ S\ DO3(o), \Delta Pon\ S\ DO3(o))} \\ T\ S\ DO3(m) &= \sqrt{f(\Delta Son\ S\ DO3(m), \Delta Pon\ S\ DO3(m))}\end{aligned}$$

$T\ S\ DO3(o/m)$ is tempo of the ornament relative to the main tempo. This factor is scaled by an exponential parameter f , and a new ornament tempo $T\ S\ DO3(o)'$ is calculated:

$$\begin{aligned}T\ S\ DO3(o/m) &= \sqrt{f(T\ S\ DO3(o), T\ S\ DO3(m))} \\ T\ S\ DO3(o)' &= T\ S\ DO3(m) * T\ S\ DO3(o/m)\sqrt{f}\end{aligned}$$

This gives a new performance duration $\Delta P\ S\ DO3(o)'$, which yields the new performance times $Pon\ S\ DO3(m)'$ and $Pon\ S\ DO3(o)'$:

$$\begin{aligned}\Delta P\ S\ DO3(o)' &= \sqrt{f(\Delta Son\ S\ DO3(o), T\ S\ DO3(o)')} \\ Pon\ S\ DO3(m)' &= Pon\ S\ DO3(m) \\ Pon\ S\ DO3(o)' &= Pon\ S\ DO3(m) - \Delta P\ S\ DO3(o)'\end{aligned}$$

Frame 7. Scaling the expressive tempo of an APPOG section.




Figure 7. Scaling expressive timing of an APPOG section. This process is shown for a specific set of performance onsets. Note that only the performance timing of the ornament is affected. At scale factor 1 the timing of the ornament is identical to the original timing. At scale factor 0 the ornament is performed at the same tempo as the main object (in this particular example the score duration of the ornament is half that of the main component). This operation (with scale factor 0) effectively removes the expressive way in which the ornament is performed, relative to the main component. At factor .5 a diminished expressive timing effect will result, and at factor 2 an exaggerated effect will be obtained. At negative values of the scale factor the expressive timing is inverted: a performance of the ornament at a lower tempo than the main component becomes one at a faster tempo and vice versa.

Scaling expressive asynchrony

Asynchrony occurs when two or more simultaneous musical objects - prescribed to happen at the same score time - have unequal performance onsets. The differences can be scaled linearly but care has to be taken not to disrupt the timing of higher levels.

Scaling the expressive asynchrony of a P section

The scaling of the expressive asynchrony of a multilateral simultaneous structure works as follows. Assume the structure has n components named $C(i)$ with $0 \leq i \leq n-1$. Component $C(i)$ has performance onset time $Pon(i)$. Assume the right context of the structure (and thus the right context of all components) has performance onset $Pon(n)$. A parallel time map of the structure contains all $Pon(i)$ including $Pon(n)$. The scale operation on such a map delivers a new $Pon(i)'$ according to the following rules:

Let the global performance onset Pon and the performance onset asynchronies $\Delta Pon(i)$ be defined as:

$$Pon = \min_{0 \leq i \leq n-1} Pon(i)$$

$$\Delta Pon(i) = Pon(i) - Pon \text{ for } 0 \leq i \leq n-1$$

The asynchronies are scaled by a multiplication factor f :

$$\Delta Pon(i)' = \Delta Pon(i) * f$$

New performance onsets Pon_i' are calculated, shifting such that the global performance onset is kept invariant ($\min(Pon_i') = \min(Pon_i) = Pon$). The result is truncated such that the onsets never move beyond $Pon(n)$. Of these two safeguards the first applying in case f is negative, the second applying in case f is large compared to the ratio of the asynchronies and performance duration of the whole structure. Together they ensure consistency with higher-level structural descriptions by keeping the components within the bounds of the structure.

$$Pon(i)' = \min(Pon(n), Pon + \Delta Pon(i)' + \min(\Delta Pon(i)'))$$

Frame 8. Scaling the expressive asynchrony of a P section




Figure 8. Scaling expressive timing of a P section. This figure shows this process for a specific set of performance times $P \setminus \text{DO3}(i)$ (say a chord performed with some spread). At scale factor 1 the performance onsets $Pon \setminus \text{DO3}(i)'$ are identical to their original $Pon \setminus \text{DO3}(i)$. At scale factor 0 all $Pon \setminus \text{DO3}(i)'$ occur synchronously at the minimum of their originals (i.e. removed chord spread). At factor .5 a diminished chord spread will result, and at factor 2 an exaggerated chord spread can be obtained. At negative values of this factor the spread is inverted: first notes becoming last and vice versa. At extreme values of the scale factor the notes are restrained from moving out of the chord structure into the next musical object by truncation. Note that the whole operation is independent of score times.

Scaling the expressive asynchrony of an ACCIA section

The scaling of the expressive asynchrony of a collateral simultaneous structure works as follows.

Assume the structure has a main component with performance onset time $Pon\backslash\backslash DO3(m)$ and an ornament component with performance onset time $Pon\backslash\backslash DO3(o)$. A time-map of the structure contains $Pon\backslash\backslash DO3(o)$ and $Pon\backslash\backslash DO3(m)$. The scale operation on such a map delivers new performance onsets according to the following rules:

Let the performance onset asynchrony ΔPon be defined as:

$$\Delta Pon = Pon\backslash\backslash DO3(o) - Pon\backslash\backslash DO3(m)$$

The asynchrony is scaled by a multiplication factor f , and a new performance onset $Pon\backslash\backslash DO3(o)'$ is calculated:

$$\Delta Pon\backslash\backslash DO3()' = \Delta Pon * f$$

$$Pon\backslash\backslash DO3(o)' = Pon\backslash\backslash DO3(m) + \Delta Pon\backslash\backslash DO3()'$$

$$Pon\backslash\backslash DO3(m)' = Pon\backslash\backslash DO3(m)$$

Frame 9. Scaling the expressive asynchrony of an ACCIA section.




Figure 9. Scaling expressive timing of a ACCIA section. It shows this process for a specific set of performance times (a note preceded by an acciaccatura). At scale factor 1 all performance onsets are identical to their original. At scale factor 0 the ornament occurs synchronously with the main note (removed asynchrony). At negative values of this factor the order of onset of ornament and main note is inverted. Note that the ornament is allowed to shift freely - even outside the bounds of the whole ACCIA structure.

Scaling expressive articulation

The articulation of a note is interpreted (scaled) relative to the articulation of the structure that it forms part of. For multilateral structures this is the average articulation. If thus the first note in a bar is played with more overlap than the other notes, a removal of the overlap articulation expression (a zero scale factor) will set the overlap of all notes to the mean overlap of the notes in the structure. And exaggerating the articulation expression (a scale factor larger than 1) will move the individual overlaps away from the mean - but maintaining the average overlap of all the notes in the bar. Of course all articulation types maintain reasonable performance offsets in the case of extreme values (i.e. note offsets will not shift before their onsets).

Scaling the expressive articulation of a multilateral section

Assume a multilateral structure has n components C_i with $0 \leq i \leq n-1$. Component C_i has articulation A_i (see frame 5 for the calculation of A_i). A section of the expression map of the structure contains all A_i . The articulation A of the structure itself is defined as:

$$A = \text{MEAN}_{0 \leq i \leq n-1} A_i$$

Let the expression deviations be

$$\Delta A_i = A_i - A \text{ for } 0 \leq i \leq n-1$$

The deviations are simply scaled by a multiplication factor f

$$\Delta A_i' = f * \Delta A_i$$

The scale transformation delivers new articulations A_i' by adding the new deviations to the reference articulation A such that the articulation of the whole structure is kept invariant (mean $A_i' = A$).

$$A_i' = A + \Delta A_i'$$

Keeping the expressive values in a reasonable range can only be done while applying them to the individual notes.

Frame 10. Scaling the expressive articulation of a multilateral section

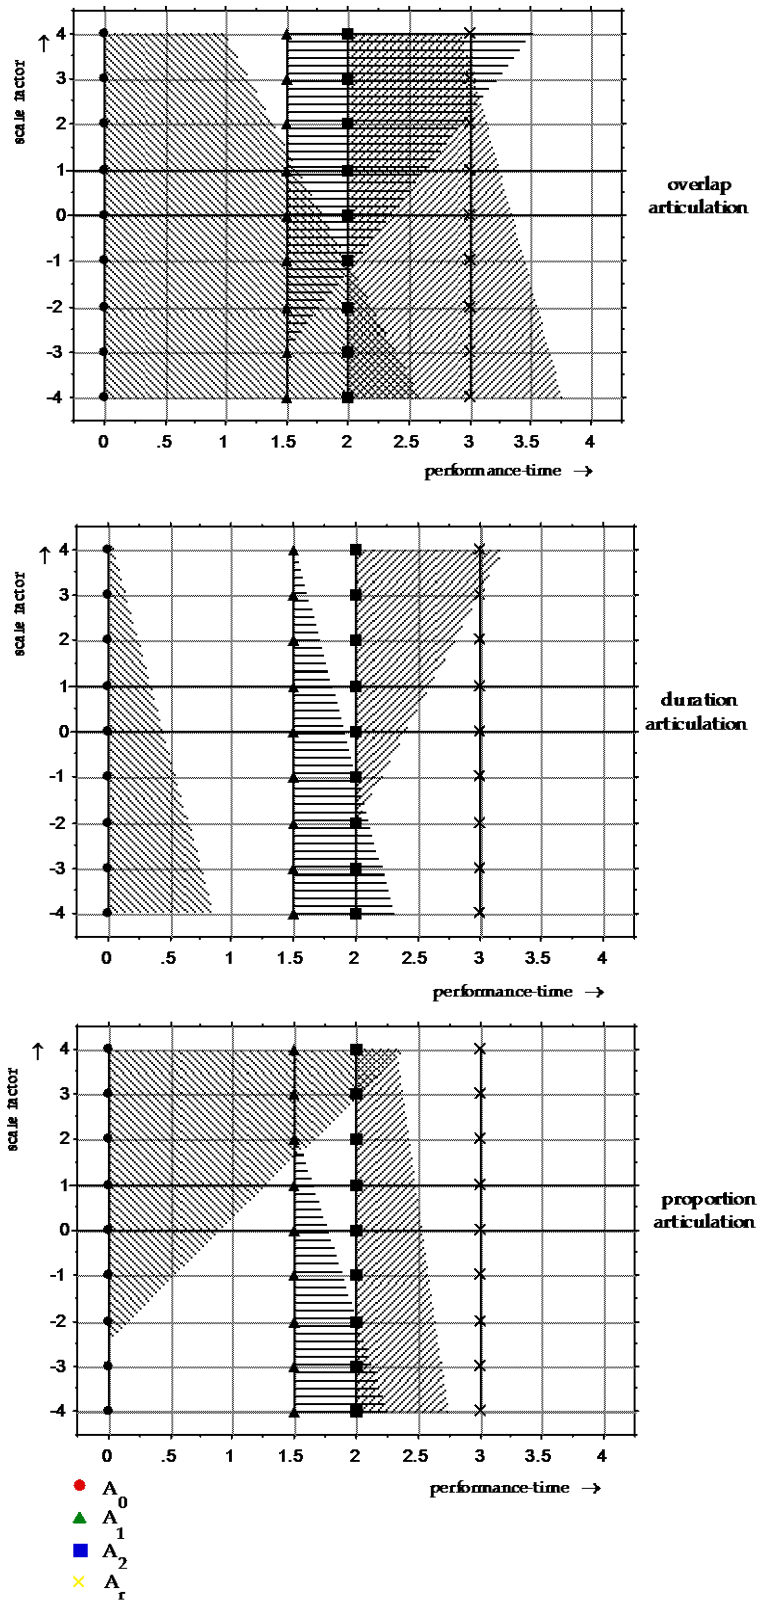


Figure 10. Scaling of an S section with three different kinds of articulation. It shows the scaling of three types of articulation for a multilateral structure, in this instance an S structure with a specific

set of performance onset and offset times. Here, at scale factor 1 articulations A_i' are identical to the original performance. At scale factor 0 all A_i' are scaled to the mean articulation A . At a scale factor above 1 the deviation of each A_i' with respect to A is exaggerated, with negative values constituting an inverse deviation: legato notes become more staccato and vice versa. Note that the mean articulation A is always kept invariant.

Scaling the expressive articulation of a collateral section

Assume that a collateral structure has ornament and main components C_o and C_m . Component C_o has articulation A_o and component C_m has articulation A_m (see frame 5 for the calculation of A_o and A_m). A section of the expression map of the structure contains these values. The articulation A of the structure itself is defined as:

$$A = A_m$$

Let the expression deviation be

$$\Delta A = A_o - A$$

The deviation is scaled by a multiplication factor f

$$\Delta A' = f * \Delta A$$

The scale transformation delivers a new articulation for the ornament by adding the new deviation to the reference articulation A .

$$A_o' = A + \Delta A'$$

$$A_m' = A_m$$

Keeping the expressive values in a reasonable range can only be done while applying them to the individual notes.

Frame 11. Scaling the expressive articulation of a collateral section.




Figure 11. Scaling of an APPOG section with three different kinds of articulation. It shows three types of articulation scaling for an ornament (here an APPOG structure). At scale factor 1 the articulation A_o is identical to the original articulation of the ornament. At scale factor 0 A_o is identical to the articulation of the main component A_m . At a scale factor above 1 the deviation of A_o with respect to the main component A_m is exaggerated, negative values constituting an inverse articulation: legato ornament articulation become more staccato and vice versa.

Keeping articulation consistent in the scaling of expressive timing

In the scaling of timing of onsets we ignored the influence it should have on its offsets. To obtain some sort of articulation consistency we can use the three types of articulation (as described above) when scaling expressive tempo and expressive asynchrony. In figure 12, we use expressive tempo scaling for an S section as an example in illustrating the different types of articulation consistency.

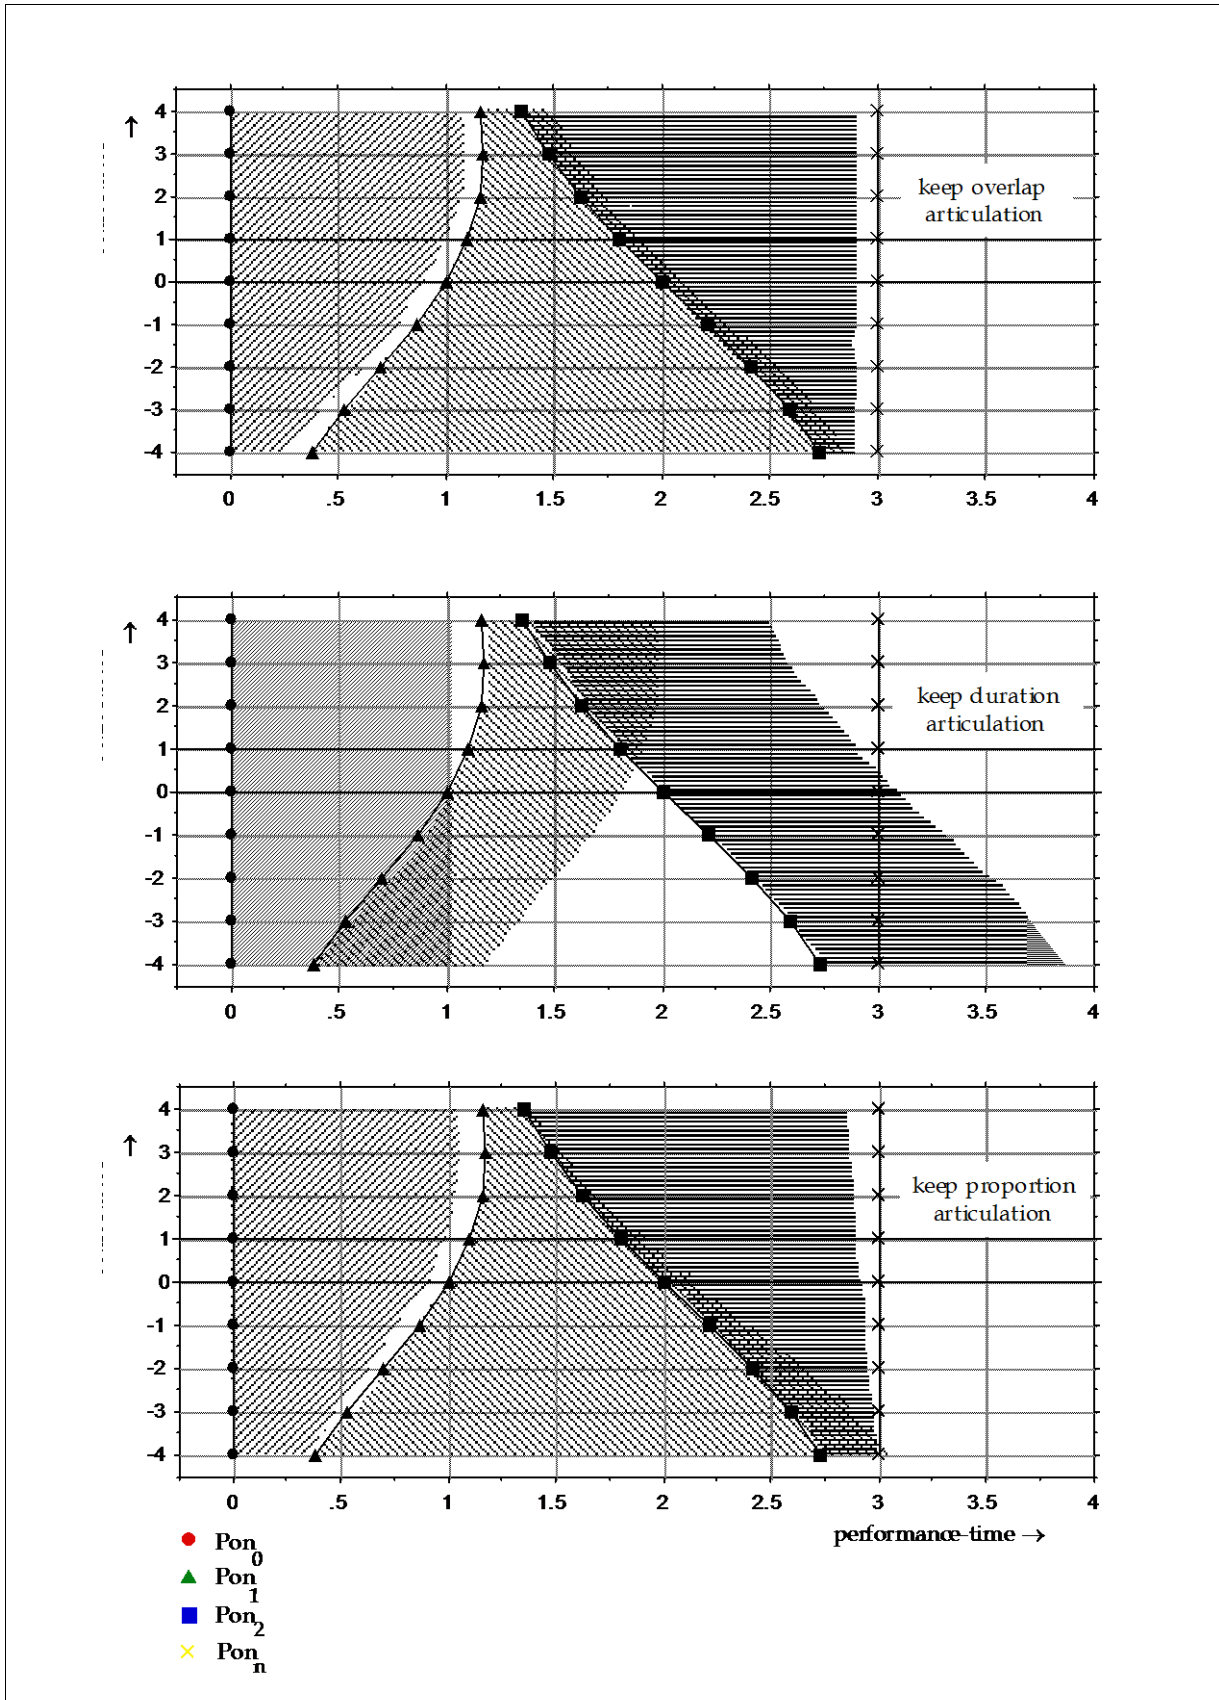


Figure 12. Scaling of an S section that keeps a particular type of articulation consistent. Shown for the same set of performance onsets as used in figure 6.

Stretch maps

Sometimes it is useful to be able to keep a consistency in performance timing between voices when modifying one of them. Naming the modified material as the foreground and describing the rest as the background, the consistency requires that a series of performance onsets, at a selected background level, that happen between two performance onsets in the foreground are “stretched along” with the changes in the foreground. This feature is implemented by first extracting a timing map from the background, and “stretching” this map between the old and modified foreground map before it is reapplied to the background. The fore- and background must be parallel (must happen during the same score time interval) and have to be S structures. Maintaining the consistency between other kinds of structure remains a problem.

Interpolate maps

A more sophisticated notion of expression entails the difference in expression between two structured objects. The best known example is voice leading in ensemble playing (Rasch, 1979) whereby the leading instrument often takes a small but consistent timing lead (around 10 ms). Inter- or extrapolation between two extracted timing maps yields the possibility to scale this kind of expression.

Transfer maps

Sometimes it is useful to apply an expression map extracted from one object, to another object, possibly with a different structure, e.g. boldly applying the expressive timing of the melody to the accompaniment. This is supported via an operation on timing maps that uses the structure of one map but imposes expressive values of the other.

TRANSFORMATIONS

Transformations of musical structures are generalizations of the operations on expression maps. They handle the selection of a level of structural description, extract a map, do the operation and re-impose the map. However, they often become quite sophisticated because they also take care of maintaining consistency with a background (material that is not affected directly). The application of the modified map has its own complexity, whereby changes are propagated to lower levels depending on the types of musical structure encountered. Finally, in the setting of new performance onsets of the notes, also the offsets may change in order to keep the articulation invariant. Out of the wealth of possibilities we choose some examples to be illustrated further by means of figures. In the figures the performance onsets and/or offsets of the individual notes at different parameter settings are given. The structure of the musical objects transformed are shown underneath.

In the following examples the same performance of a Beethoven theme is used (the fragment as shown in figure 2), allowing for comparison of the different

transformations and to see the effect of applying the same transformation to different levels or types of structure. Note that for all the transformations the identity transformation is shown at scale factor 1.

Scale timing




Figure 13. Scaling the expressive tempo of bars in the Beethoven fragment. Underneath the figure a structural description of the fragment is shown in bars. Imagine what would happen if we asked a performer to emphasize his/her timing of the bars? One possibility would be to play the onsets of the bars, that were played slightly early, even earlier, and ones that were played late, later still. This particular transformation can be read from figure 12 as the lines with the black markers, indicating the component in the bar that carries the expressive timing. Both the performance onset of the first and the last bar of the enclosing bars' structure are not changed; the transformation is done at the level named "bars", with its timing kept invariant. The lines with white markers show the embedded material that follows the change of the performance onset of each bar. Note that the

timing of the ornamented notes does not change (they keep the same distance with respect to the note they cling to), as does the spread of the chords.




Figure 14. Scaling the expressive asynchrony of each bar in the Beethoven fragment. It shows the expressive transformation we might expect to happen when a performer is asked to exaggerate the asynchrony between the top-voice and bottom-voice at the onset of each bar. The figure shows the scaling of the asynchrony of the bottom voice onsets (the black squares), without changing the timing of the bars (lines marked with black circles and triangles). The embedded notes of the bottom voice (lines with white squares) just shift along with the expressive timing of their embedding structure. Here again, the ornament timing and the chord spread stay invariant.




Figure 15. Scaling the expressive asynchrony of each chord in the Beethoven fragment. It shows another expressive transformation that exaggerates the chord spread, turning them almost into arpeggio's at high scale factors. At scale factor 0 the chord spread is completely removed. The timing of the rest of the fragment stays unaltered.




Figure 16. Scaling the expressive tempo of the melody in the Beethoven fragment, a) without and b) with “stretching” the accompaniment. It shows that the timing of each note of the melody

becomes exaggerated with a higher scale factor. Here the accompaniment (lines marked with white squares) is not affected at all. Figure 16b, on the other hand, shows a musically more reasonable transformation: the accompaniment follows the movements of the transformed melody, e.g. slowing down when the tempo of the melody slows down. Here the accompaniment is kept consistent with respect to the original performance (compare with the onsets at scale factor 1). Note that note order can change between melody and accompaniment, because of the structural description in two parallel voices.

Keeping articulation consistent

In the above examples we showed the scaling of onset times and neglected what happened to the offset times. But, as we showed before, this cannot just be ignored in musically relevant transformations. We can select one of the described types of articulation to keep consistent, but we do not show this here (see figure 12 for a simple example).

Scale intervoice expression

When the expression between voices is scaled, two parameters are used. The first one selects a reference level of expression (0 designates the expression of the first, 1 designates the expression of the second, 0.5 is the mean of the two etc.). The second parameter determines in how far the voices are removed from that reference level (0 means completely on reference level, 1 means as in original performance, 2 means exaggerated with respect to the reference etc.).




Figure 17. Scaling the intervoice timing between the melody and the accompaniment in the Beethoven fragment, a) with the melody as reference, and b) with the mean of the melody and the accompaniment as reference. In figure 17a) intervoice timing (one type of intervoice scaling) is scaled with the melody voice as the reference. It shows the scaling of the asynchrony between the accompaniment and the melody, as found in the performance (see the horizontal line where

the scale factor is 1). Notes that are not synchronous (i.e. don't have the same score time) interpolate their change with respect to their surrounding performance onsets that are considered parallel (have the same score time). Note that the timing of the melody does not change because it is used as reference.

In figure 17b the mean of the melody and accompaniment timing is used as reference, resulting in displacements (with respect to this invisible reference) of both voices.

In both figures, the first event in the melody voice is unaffected since there is no measurable timing in the accompaniment (only a rest).

CONCLUSION

In this paper we have presented a proposal for a calculus that enables expressive timing to be transformed on the basis of structural aspects. The program implementing the calculus, will hopefully prove to be useful for formalised theories of music cognition. The proposed representation constructs allow for easy maintenance and extension. An object-oriented programming style proved a good choice for this kind of modelling. The algorithmic parts became reasonably simple, but the program can still be considered as quite complex, especially its elaborate knowledge representation. This algorithmic simplicity combined with structural complexity mirrors, in this respect, the widespread hypothesis that the complex expressive timing profiles found in musical performances are more readily explained as the product of a small collection of simple rules linked to a relatively complex structure, than as the result of a large collection of interacting rules, with hardly any structure.

This research again confirmed that music is a very rewarding field for experimentation with knowledge representation concepts.

ACKNOWLEDGEMENTS

We are very grateful to Eric Clarke who made it possible for us to work for two years on research in expressive timing at City University in London, and the British ESRC for their financial support (grant A413254004) during this period.

REFERENCES

- Bentley J. (1988) More Programming Pearls, Confessions of a Coder. Reading, MA: Addison-Wesley.
- Bregman, A. S. (1990) Auditory Scene Analysis: The Perceptual Organization of Sound. Cambridge, Mass.: Bradford books, MIT Press.
- Clarke, E. F. (1988) Generative principles in music performance. In J. A. Sloboda (Ed.) Generative processes in music. The psychology of performance, improvisation and composition. Oxford: Science Publications.
- Desain, P. & H. Honing (1988) LOCO: A Composition Microworld in Logo. Computer Music Journal 12(3): 30-42.
- Desain, P. & H. Honing (1991) Quantization of Musical Time: A Connectionist Approach. In P. M. Todd & G. J. Loy (Eds.) Music and Connectionism. Cambridge, Mass.: MIT Press.
- Desain, P. & H. Honing (1992a) Tempo curves considered harmful. To appear in Contemporary Music Review.
- Desain, P. & H. Honing (1992b) Time functions function best as functions of multiple times. To appear in Computer Music Journal 16(2).

- Desain, P. (1990) Lisp as a second language. Perspectives of New Music 28(1).
- Honing, H. (1990) POCO: An Environment for Analysing, Modifying, and Generating Expression in Music. In Proceedings of the 1990 International Computer Music Conference. San Francisco: Computer Music Association.
- Honing, H. (1992) 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).
- Keene, S. E. (1989) Object-Oriented Programming in Common Lisp: A Programmer's Guide to CLOS. Reading, MA: Addison-Wesley.
- Longuet-Higgins, H. C. (1976) The Perception of Melodies. Nature 263: 646-653.
- Rasch, R. A. (1979) Synchronisation in performed ensemble music. Acoustica 43, 121-131.
- Serafine, M.L. (1988) Music as Cognition: The Development of Thought in Sound. New York: Columbia University Press.
- Steele, G. L. (1990) Common Lisp, the Language. Second edition. Bedford, MA: Digital Press.
- Vos, J. & R. A. Rasch (1981) The perceptual onset of musical tones. Perception & Psychophysics 29(4): 323-335.