

Towards some Organising Principles for Musical Program Auralisations

Paul Vickers

School of Computing & Mathematical Sciences, Liverpool John Moores University
Liverpool, UK

James L. Alty

LUTCHI Research Centre, Loughborough University
Loughborough, UK

Abstract

Early studies have shown that musical program auralisations can convey structural and run-time information about Turbo Pascal programs to listeners [3, 4, 10]. Auralisations were effected by mapping program events and structures to musical signature tunes, known as *motifs*. The design of the motifs was based around the taxonomical nature of the Turbo Pascal language constructs [3]. However, it became clear that as the musical complexity and grammatical rigour of the motifs increased, their discernability by the average user decreased. Therefore, from the lessons learnt from our work we propose a set of organising principles for the design and construction of musically-based program auralisations. These organising principles are aimed towards providing accessible auralisations to the average programmer who has no formal musical training.

1 Introduction

We have developed CAITLIN¹, a system for adding musical auralisations to Pascal programs with the aim of assisting the debugging process [4, 10]. CAITLIN's auralisations were designed to reflect the hierarchical nature of Turbo Pascal constructs [3]. Three small studies showed that as the musical consistency of the auralisation design increased, the comprehensibility of the motifs to listeners decreased (in some cases significantly). We did not consider that there was anything inherently unsound about the use of musical frameworks for program auralisation and so the results from the studies were analysed to see what lessons could be learnt.

We saw the need for developing a set of guidelines or organising principles which would steer a redesign of CAITLIN's motifs. The intention of the system has always been to assist the novice programmer with debugging tasks. Further, it is likely that most programmers will have very little formal musical training and this must be taken into account.

2 Experimentation

Three experiments were carried out at John Moores University and Loughborough University over a period of approximately eighteen months. The experiments yielded useful information about program auralisation. First, the earlier tests indicated that auralisation can be used in communicating information to programmers [10].

Comments received after the experiments showed that subjects were generally receptive to the idea of program auralisation and believed that it could be made to work. Useful information regarding selection of timbre, complexity of melodic structures and intuitiveness of mappings was gleaned from these tests. It should be noted at this point that all three experiments required identification to be made solely on the basis of aural cues. In a real debugging situation, programmers would obviously have detailed knowledge of the program being auralised.

Figures 1 and 2 show the overall results from the second and third experiments. Both experiments involved four separate tests. In the first test subjects were required merely to state whether the auralisations they were hearing belonged to iteration or selection constructs; hence, in this test any given response was either wholly correct or wholly incorrect. In the second test the two experiments, subjects were presented with a number of individual auralisations and were asked to identify the construct to which it belonged. A correct response was one where the exact construct was identified. However, if a subject identified a FOR...TO loop, say, as a FOR...DOWNTO loop then this was judged to be a correct identification of the construct's *sub-class* (see [3] for an explanation of class and sub-class taxonomy in Pascal constructs). If a FOR...TO loop were identified as a WHILE loop by a subject then this is considered to be a correct *class* identification as both constructs belong to the iteration class. The third test was similar to the second, except this time subjects were

¹ For further details and recent results, see <http://www.cms.livjm.ac.uk/www/homepage/cmospvick/caitlin/>

presented with pairs of constructs in sequence; in the fourth test, the auralisations included examples of sequential and nested constructs. Interestingly, the nesting appeared to have no effect on subjects' ability to identify constructs from their auralisations.

The charts in Figures 1 and 2 then show the number of such responses for each of the four tests in the two experiments.

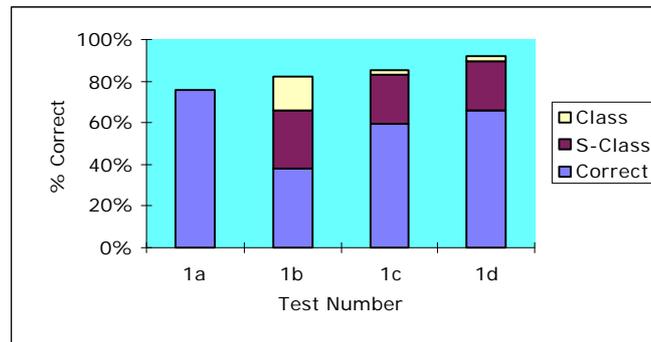


Figure 1: General results from experiment 2

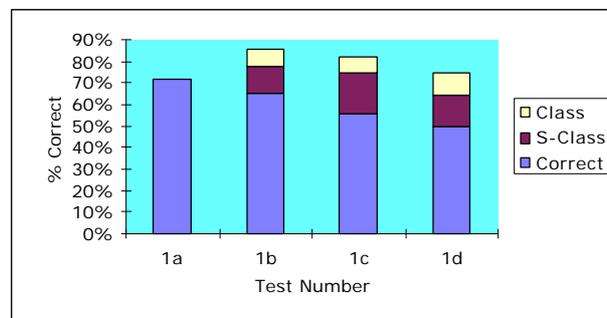


Figure 2: General results from experiment 3

Between the three experiments we observed an overall decline in the ability of subjects to correctly identify Pascal constructs from the auralisations alone ($t=-2.59823$, $p=0.0233$). The first two experiments showed some significant increases in correct responses across the duration of the tests (e.g. $t=-3.45794$, $p=0.04070$); the third experiment showed a significant decrease in accuracy of response over its duration ($t=3.19689$, $p=0.01867$). (We must add the caveat at this point that these results are from small samples (<10)). We also observed that whilst the second experiment showed an improvement in construct identification across its four tests, the overall accuracy for the study was lower than for the first study. In summary, we observed that:

- Accuracy of construct identification decreased significantly.
- Motif design changed from purely arbitrary mappings in the preliminary experiment to quite formalised musical structures in the third.

That is, as the motifs became more musical, more structurally consistent and more organised, accuracy of identification became worse. From this we draw the conclusion that whilst the organisational principles of motif design provide a useful framework for the representation of constructs, the average user does not discriminate well between subtle musical changes. Above all, the average user needs clear and unambiguous auralisations. Motifs should be constructed in such a way that their organising principles and musical grammar and structure become transparent and so allow the individual features of the different constructs to be made readily apparent. If designed well, then the motifs will be identifiable at three levels of perception (uniqueness, metaphorical and semantic). Page limits do not permit discussion of these concepts here and a more complete coverage can be found elsewhere [2]. Suffice it to say in this context that *uniqueness* (or *detectable musical mapping* [2]) corresponds to the construct auralisation's ability to be uniquely identified and not confused with another. The *metaphorical* level (or *perceptual context* [2]) is one where given a detectable mapping, the motif creates expectation on the part of the listener. The motif is interpreted in domain terms and meaning can be assigned to it, though the listener may not necessarily be able to reason about the global interactions. The *semantic* level (or *reasoning level* [2]) is where the listener develops high-level structures in the mind to understand the domain from a higher or more abstract viewpoint.

3 Review of Cognitive Aspects

To assist in the formulation of a set of organising principles it is useful to review some of the reported factors involved in musical perception. Watkins and Dyson assert that the “*perceptual organisation of tone sequences is facilitated if they are structured to approximate the schemata of melodies belonging to the musical idiom with which our listeners are familiar*” [11]. Western musical forms based around the seven-note diatonic scale are readily recognised around the world (especially in the computing community [1] and so this gives a good starting point for auralisation design. Watkins and Dyson go on to say that “*tone sequences that obey the constraints imposed by the system of Western scales and keys are more easily organised (cognitively), learned and discriminated than control sequences of comparable complexity*” [11].

3.1 Melody Recall

Experimentation supports the idea that familiarity with tunes makes them readily discriminable from unknown music even by musically-unsophisticated listeners [11]. Therefore, there is no intrinsic requirement that users of an auralisation system possess formal musical training.

In experiments on subjects’ ability to recall melodies, Sloboda and Parker [9] observed that the most fundamental feature preserved in the recalled melody was its metrical structure. Further, musicians and non-musicians differed significantly only on one measure, that of the ability to retain the harmonic structure of the original. Therefore, it would be wise not to rely on ability to discriminate between harmonic structures in the auralisation motifs.

3.2 Music Contour

Contour was judged by subjects in our studies as being a useful aide-memoire for recalling the motifs. However, results reported by Edworthy [6] confirm Dowling’s suggestion [5] that contour becomes important when tonal context is weak or confusing. Contour is less important in familiar melodies and melodies retained over a period of time. The results from our preliminary experiments tend to bear this out. Whilst the motifs were designed according to specific principles, and thus motifs representing related constructs were based on the same melodic contour, this feature was of little concern to the listener. What is more important is that prior to using the CAITLIN system in earnest, subjects must first be exposed to the different motifs so that the tunes become familiar.

3.3 Transition Probabilities

It has been observed that certain pitch combinations occur more often than others in western music [12]. Analysis of such pitch combinations gives us the notion of the *transition probability* [11] which describes the likelihood of a particular attribute given an attribute in preceding elements. Listeners appear to be sensitive to transition probabilities that occur in western tonal melodies with commonly occurring transition probabilities being perceived as more musical [11, 12].

3.4 The Fifth-Span

A feature related to the transition probability is the *fifth-span* [7, 8]. The fifth span is determined as follows. Each note in the equal-tempered scale is assigned a sharpness value (q) which represents the number of fifths (an interval of seven semitones) between it and its keynote (the root note of the scale, e.g. C in C major). For the key of C major, the q values of all twelve notes of the scale are shown in Table 1. The keynote acts as a tonal centre, or anchor point, in western tonal melodies. The difference between the q -values of two adjacent notes in a melody is called the “fifth-span” (or δq) which is given in (1) as the absolute value of the difference in q -values of the two notes of the interval [7, 8].

$$dq = Abs(q_1 - q_2) \quad (1)$$

Note	D \flat	A \flat	E \flat	B \flat	F	C	G	D	A	E	B	F \sharp
q value	-5	-4	-3	-2	-1	0	1	2	3	4	5	6

Table 1: q values of equal-tempered notes (from Watkins & Dyson [11])

Watkins and Dyson [11] found that fifth-spans of more than six fifths are rare in western tonal music. They argue that this constraint may provide listeners with information about the key of a melody in which each interval suggests a range of likely keynotes for which the fifth-span of the interval is less than six. Looking at Table 1 it can be observed that the diatonic scale consists of intervals with a fifth span of five or less. Thus, for

the typical western listener, intervals will be expected to lie within the diatonic scale.

4 Towards Some Organising Principles

We now propose some organising principles for auralisation motif design which take into considerations the lessons learnt from the above.

The motifs will comprise intervals that lie within this expected range. Events that signal a significant change in state within a program (e.g., the controlling expression of a loop changing from the continuation value to the terminating value) could be represented musically by unexpected intervals (i.e. those with fifth spans of greater than 5).

1. Motifs should be structured according to a hierarchy which reflects the taxonomy of Pascal constructs (see [3] for details).
2. Listeners expect to hear diatonic intervals and so these should be used as the default interval in the motifs [11].
3. Non-diatonic intervals (or those with fifth spans greater than five) are “unexpected” [11] and thus could be used when state changing events occur.
4. Major mode to represent Boolean *True*; minor to represent Boolean *False*.
5. Percussive sounds to reinforce the occurrence of decision points.
6. Drones (sustained pitches) to signal presence in construct body (to provide a continuity of context).
7. Motifs must not rely on an ability to discriminate between harmonic structures; metrical structures have been shown to be more important [9].

5 Redesign of Motifs

We have tried to use the above principles to redesign the set of motifs used by CAITLIN. For instance, Figure 3 shows a score for a WHILE loop with two iterations. To make it easier to identify whether a motif represents an iteration or a selection, we have added a rhythmic marker. All iteration motifs begin and end with a repeated triangle sound and selections begin with two ascending pitches played on a cow-bell and end with two descending pitches on a cow-bell (metaphorical for question and answer). For example, Figures 3 and 4 respectively show the scores for auralisations of a WHILE loop with two iterations and an IF...ELSE statement whose condition yields a false result.

Figure 3: A WHILE loop with two iterations

Figure 4: An IF...ELSE statement yielding False

6 Evaluation and Conclusions

The new auralisations which have been designed around the above organising principles are presently undergoing full study at Loughborough University. The experimental design is along the lines of previous tests, though with more formality and attention paid to training. We hope the results will show that the new motifs are readily understood and identified by users and that musical auralisations based upon a set of structured design principles can be used to help novice programmers with debugging their programs.

7 References

1. Alty, J. L., Can We Use Music in Computer-Human Communication?, in *People and Computers X*, D. Diaper and R. Winder, Eds. Cambridge University Press: Cambridge, 1995.
2. Alty, J. L. and Rigas, D. I., Communicating Graphical Information to Blind Users Using Music: The Role of Context, in *Proc. CHI98 Conference on Human Factors in Computing Systems (Los Angeles, CA, April 18-23, 1998)*, ACM Press, pp. 574-581.
3. Alty, J. L. and Vickers, P., The CAITLIN Auralization System: Hierarchical Leitmotif Design as a Clue to Program Comprehension, in *Proc. The Fourth International Conference on Auditory Display (Palo Alto, 1997)*, Xerox PARC, Palo Alto, CA 94304.
4. Alty, J. L., Vickers, P. and Rigas, D., Using Music as a Communication Medium, in *Proc. Refereed Demonstrations, CHI97 Conference on Human Factors in Computing Systems (Atlanta, GA, March 22-27, 1997)*, ACM Press, pp. 30-31.
5. Dowling, W. J., Melodic Information Processing and Its Development, in *The Psychology of Music*, D. Deutsch, Ed. Academic Press: New York, 1982, pp. 413-429.
6. Edworthy, J., Melodic Contour and Musical Structure, in *Musical Structure and Cognition*, P. Howell, I. Cross and R. West, Eds. Academic Press: New York, 1985, pp. 169-188.
7. Longuet-Higgins, H. C., Perception of Melodies, *Nature* 263, (1976), pp. 646-653.
8. Longuet-Higgins, H. C., The Perception of Music, *Interdisciplinary Science Reviews* 3, (1978), pp. 148-156.
9. Sloboda, J. A. and Parker, D. H. H., Immediate Recall of Melodies, in *Musical Structure and Cognition*, P. Howell, I. Cross and R. West, Eds. Academic Press: New York, 1985, pp. 143-167.
10. Vickers, P. and Alty, J. L., CAITLIN: A Musical Program Auralisation Tool to Assist Novice Programmers with Debugging, in *Proc. Third International Conference on Auditory Display (Palo Alto, Nov 4-6, 1996)*, Xerox PARC, Palo Alto, CA 94304, pp. 17-24.
11. Watkins, A. J. and Dyson, M. C., On the Perceptual Organisation of Tone Sequences and Melodies, in *Musical Structure and Cognition*, P. Howell, I. Cross and R. West, Eds. Academic Press: New York, 1985, pp. 71-119.
12. West, R., Howell, P. and Cross, I., Musical Structure and Knowledge Representation, in *Representing Musical Structure*, Cognitive Science Series, R. West, P. Howell and I. Cross, Eds., 1 ed. Academic Press Ltd: London, 1991, pp. 1-30.