

CAITLIN: A Musical Program Auralisation Tool to Assist Novice Programmers with Debugging

[Paul Vickers, Liverpool John Moores University](#)
[James L. Alty, Loughborough University](#)

Abstract: In the field of auditory display relatively little work has focused on the use of sound to aid program debugging. In this paper, we describe CAITLIN, a pre-processor for Turbo Pascal programs that musically auralises programs with a view to assisting novice programmers with locating errors in their code. A discussion follows of an experiment which showed that programmers could use the musical feedback to visualise and describe program structure. We then present conclusions and a discussion of future work. Keywords: Auralisation, audiolisation, auditory-display, musicode.

Introduction

The term software visualisation suggests the idea of investigation using the visual sense alone, however, the aim of software visualisation is simply to improve the understanding of software ([Domingue, Price, & Eisenstadt, 1992](#)). Therefore, it makes sense to use sound if it possesses properties that facilitate software comprehension. Previous research shows that sound is a useful tool in the presentation of information to users. Examples include: Edwards' sound-enhanced word-processor for the blind ([1989](#)), Gaver's auditory icons ([1986](#)), Blattner's earcons ([1989](#)), and Gaver's audio-enhanced graphical user interface ([1989](#)). Sonification, or the mapping of data to sound ([Scaletti, 1994](#)), demonstrates that large data sets can be represented using sound ([Bly, 1982](#); [Mezrich, Frysinger, & Slivjanovaski, 1984](#); [Scaletti & Craig,](#)

[1990](#)). Audification ([Kramer, 1994](#)), or the direct conversion of data to sound, has been used to analyse large sets of seismic data that would be more difficult to visualise using graphics ([Hayward, 1994](#)).

Program Auralisation

Within the field of auditory-display research, program auralisation is beginning to attract interest. Auralisation is the representation of program data (including execution state) using sound ([Jackson & Francioni, 1994](#); [Kramer, 1994](#)). The majority of efforts has been concerned with the auralisation of specific algorithms. Brown and Hershberger ([1992](#)) used music to enhance and complement an animation of a bubble sort algorithm. Other work has shown how sound can be used as the primary medium for visualising the state of parallel algorithms ([Francioni & Rover, 1992](#); [Jackson & Francioni, 1994](#)).

The case for using music to aid debugging is supported by Francioni et. al. ([1992](#)), although they felt that a visual presentation was also needed to provide a context or framework for the audio sound-track. Program state was captured as a set of trace data that were subsequently auralised and animated. Jackson et al ([1994](#)) have suggested ways in which certain aspects of parallel programs can be auralised, but again the auralisation is of a 'post-mortem' nature.

Auralisation projects of note include the InfoSound system ([1990](#)) by Sonnenwald et. al., DiGiano et al's LogoMedia ([1992](#)), Jameson's Sonnet system ([1994](#)), Bock's Auditory Domain Specification Language (ADSL) ([1995](#)) and Mathur's LISTEN system ([1993](#)). One stark omission from the existing literature is empirical evidence that program auralisation is actually useful. Preliminary experimentation showed that music can be used to visualise algorithm state ([Alty, 1995](#)). In this experiment we gauged reaction to

tasks that required the interpretation of musical output.

The first task showed that subjects were fairly accurate at estimating the difference (in semitones) between two musical pitches. The second task required subjects to sketch the perceived shape of short musical sequences. Subjects were generally able to pick up the basic shape being presented. As the experiment used musical sequences generated by a bubble sort the results indicate that music can aid visualisation of algorithm state.

The CAITLIN Project

We now need some good experimentation to determine what is possible and practicable with program auralisation. The CAITLIN project aims to determine whether musical feedback can help novices to debug programs. Music might also be used to assist visually impaired programmers. A pre-processor was constructed to auralise novices' Pascal programs. Experimentation will be undertaken to elicit evidence to support the claim that musical feedback is useful in enhancing the debugging process.

Our auralisations are deliberately based on musical techniques. Music is an extremely powerful medium for the delivery of large amounts of data in parallel, using techniques such as counterpoint and polyphony. Separate musical ideas can be delivered in parallel without confusing the listener if certain syntactic and semantic rules are followed. It makes sense to investigate whether music can be usefully employed in program auralisations. Other reasons for using music in program auralisations have been considered in earlier work (Alty, 1995), as have some of the arguments against using music (Alty, 1995). One of the more enduring criticisms is that quantitative information cannot be conveyed by sound or music.

The argument is that while most individuals can tell if

a note increases or decreases in pitch, only trained musicians are able to determine exact intervals with any accuracy. However, quantitative information can be meaningfully described and presented in terms of its overall magnitude without needing to know its exact, discrete values. For instance, just as an unmarked thermometer allows visual judgement of the relative magnitudes of various temperatures, musical pitch enables aural gauging of the relative values of different data. Nature already provides an acoustic thermometer in the form of the striped ground cricket. By listening to the cricket's chirps, one can predict ambient temperature (Pierce, 1949); for example, twenty chirps per second equates to a temperature of 88.6 Fahrenheit.

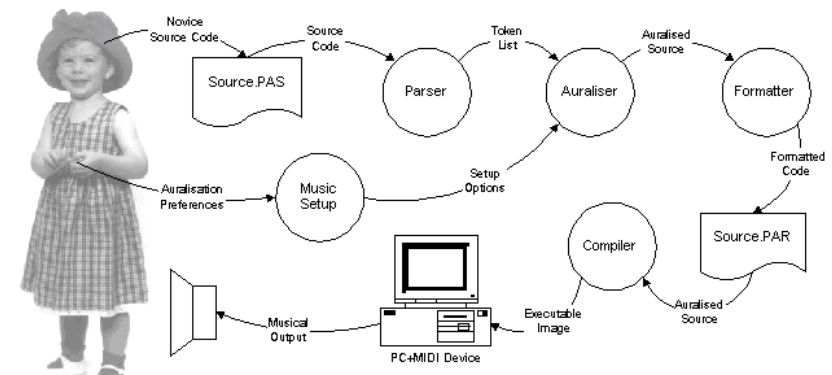


Figure 1. CAITLIN's architecture

The CAITLIN Pre-Processor

The first stage of the project has been the construction of the CAITLIN system. CAITLIN is a non-invasive pre-processor that allows a novice programmer to auralise a program written in Turbo Pascal. Figure 1 shows the basic architecture of the system in terms of its functional units and their linkages. Musical output is via MIDI to a multi-timbral synthesiser. The system can be implemented on a relatively modest platform

comprising a personal computer and sound card with a General-MIDI compatible instrument set.

CAITLIN is non-invasive in that it leaves the source program unchanged. The auralisations are effected by adding library routine calls to a copy of the program. The enhanced copy of the source program is compiled to produce an auralised executable image ([Figure 1](#)). Because CAITLIN is designed to assist in debugging executable programs and not to help compile code, it will only accept a source program free of syntax errors.

Upon running CAITLIN ([Figure 2](#)) the user is presented with a screen similar in concept and layout to that of the Turbo Pascal Integrated Development Environment (IDE). A menu option allows the user to load a source program into memory. The program is then parsed and stored as tokens in memory. After loading, the user can opt—via a menu ([Figure 2](#))—to auralise and then compile and run the auralised program, or musicode.

Auralisation is done at the construct level. A WHILE loop is auralised in one way and REPEAT, FOR, CASE, IF...THEN...ELSE and WITH constructs in others. The user may select, for each construct, the nature of the auralisation to be applied. This is fairly simple-minded, allowing selection of musical scale (e.g., major, minor, etc.), default note length (e.g., eighth note), MIDI channel and instrument. The speed at which the music is heard is controlled by a user-definable tempo variable. All options can be saved to a configuration file.

The auralisation comprises three basic parts for each construct: a musical signature tune (leitmotif) to denote the commencement of the construct, a musical structure representing the execution of the construct, and a signature to signal exit from the construct. The contents of the musical structure within the construct will depend upon the construct's characteristics.

Different constructs have different features that will be represented in various ways. To this end we have introduced the notion of the point of interest (POI). A point of interest is a feature of a construct, the details of which are of interest to the programmer during execution. For example, the IF construct has four POIs:

1. entry to the IF construct;
2. evaluation of the conditional expression;
3. execution of selected statement; and
4. exit from the IF construct.

For each construct type, the first and last POIs always denote entry to and exit from the construct respectively.

To enable the listener to distinguish between the POI-1 of FOR, WHILE and REPEAT loops we defined a short signature tune for each construct type. When a program is auralised, the tune associated with FOR statements is inserted prior to each FOR loop, and so on. A construct's last POI is auralised by playing a complement to its signature tune (such as playing it in reverse).

By defining a program in terms of its points of interest we build up an understanding of how each element and also how the whole program should sound. For example, we know that each FOR loop will be heard as a sequence of

- playing of signature tune, followed by
- repetition of music denoting iterated statement execution, followed by
- playing of modified signature tune.

This is illustrated by [Figure 2](#). The code window in [Figure 2](#) shows a listing of a simple program

employing two FOR loops. The auralisation employed in this example is straightforward. Code is inserted by CAITLIN so that each iteration of the outer loop generates a pitch of an ascending scale—the scale type being selected by the user. The inner loop plays a descending scale as the loop counter in this case is decremental. The user only has access to the original source code and does not see the expanded auralised source.

Experimentation

We carried out preliminary experimentation on CAITLIN, by first familiarizing subjects (eight faculty members) with the types of auralisation used by having them listen to ten examples. Each example was accompanied by a narrative description of the program in question with source code available on request. Each sample auralisation could be repeated as many times as required.




Figure 2. CAITLIN main screen

Following the familiarisation session subjects were presented with nine auralisations. For each auralisation, they were asked to describe the structure of the program it represented. Only audio cues were available, and the output of the programs was not shown. Also, no facility was provided for changing any of the system parameters (such as instrument used

etc.). The entire process took approximately 25 to 30 minutes.




Figure 3. Preliminary experimental results

The [results](#) suggest that, on the whole, the subjects were able to visualise program structure using only the auralisation. Most subjects specified exactly the program structure represented by the auralisations. Where subjects did not give an exact description, they could describe the essence of the structure, they were scored as 'nearly' correct (e.g., specifying a FOR...TO loop rather than a FOR...DOWNT0 loop). Where the answer bore no correspondence to the actual structure, then a score of 'no idea' was given. It is worth noting that the one subject who scored five 'no ideas' and four 'nearly' corrects claimed to lack familiarity with western music. More thorough experimentation should determine whether it was really this unfamiliarity or simply a lack of familiarity with CAITLIN itself that was to blame.

Programs 8 and 9, which scored the fewest correct responses, contained combinations of IF and IF...ELSE...IF constructs. The poor scoring on these two examples is interesting because we have neglected (contrary to our previously stated guidelines) to provide an auralisation for the IF construct's final POI. Thus, it is impossible to determine in all but a very few cases when an IF statement terminates. Program 8 involved nested IF statements. Deducing this required the listener to spot that the inner selection was played at an octave higher than the outer one, and no example of nested selections was given in the familiarisation session. CAITLIN also fails to signal when an ELSE path exists for a

selection whose conditional expression yields true.

Conclusions

In general, programmers understood what CAITLIN was doing and could follow the execution of simple programs. However, the ambiguity surrounding the IF statement shows that it is important to auralise exit from a construct as well as entry to it. The auralisation of the first and last POIs helped subjects to differentiate nested and sequential program structures. The background drone provided during execution of WHILE and REPEAT loops also assisted with this.

Instrument selection is important, as subjects commented that it was easy to deconstruct auralisations in the mind when the timbres used for the various constructs were markedly different. Careful attention must be paid to signature tune construction. One subject was unable to distinguish between the entry and exit signatures of the FOR loop. Although this subject gave nine correct responses, more complex examples might cause confusion, especially when such loops are nested. The signature tune used for the REPEAT loop was more intricate than other signatures, and did appear to confuse several subjects.

A proportion of the non-correct responses appear to be caused by subjects incorrectly remembering what each auralisation represented. One subject identified one auralisation as both a REPEAT loop and as a WHILE loop in consecutive test programs. He described his uncertainty as being caused by not remembering which tune was which. A longer familiarisation session might have improved his score.

The feedback from these preliminary experiments is being used to develop the next version of CAITLIN which will be used by novice programmers. Further experimentation will determine whether the novice programmer can use auralisations to help debug programs.

[The CAITLIN tutorial](#) provides access to these examples and the nine test auralisations used in the experiment.

Visit the [CAITLIN home page](#).

References

- Alty, J. L. (1995). Can We Use Music in Computer-Human Communication? In D. Diaper & R. Winder (Eds.), *People and Computers X*, Cambridge: Cambridge University Press.
- Blattner, M. M., Sumikawa, D. A. & Greenberg, R. M. (1989). Earcons and Icons: Their Structure and Common Design Principles. In *Human Computer Interaction*, 4, pp. 11-44.
- Bly, S. A. (1982). Communicating with Sound, In *Proc. CHI '82* (pp. 371-375). New York: ACM Press/Addison-Wesley.
- Boardman, D. B. & Mathur, A. P. (1993, Sept. 21). *Preliminary Report on Design Rationale, Syntax, and Semantics of LSL: A Specification Language for Program Auralization*. W. Lafayette, IN: Dept. of Computer Sciences, Purdue University.
- Bock, D. S. (1995). *Auditory Software Fault Diagnosis Using a Sound Domain Specification Language*, Ph.D. thesis, Syracuse University, Syracuse.
- Brown, M. H. & Hershberger, J. (1992). *Color and Sound in Algorithm Animation*, *Computer*, 25(12), 52-63.
- DiGiano, C. J. & Baecker, R. M. (1992). Program Auralization: Sound Enhancements to the Programming Environment, In *Proc. Graphics Interface '92* (pp. 44-52).
- Domingue, J., Price, B. A. & Eisenstadt, M. (1992). A Framework for Describing and Implementing Software Visualization Systems, In *Proc. Graphics Interface* (pp. 53).
- Edwards, A. D. N. (1989). Soundtrack: An Auditory Interface for Blind Users. *Human Computer Interaction*, 4(1), 45-66.
- Francioni, J. M. & Rover, D. T. (1992). Visual-Aural Representations of Performance for a Scalable Application Program. In *Proc. High Performance Computing Conference* pp. 433-440.
- Gaver, W. W. (1986). Auditory Icons: Using Sound in Computer

Interfaces. Human Computer Interaction, 2, 167-177.

Gaver, W. W. (1989). The SonicFinder: An Interface that Uses Auditory Icons. Human Computer Interaction, 4 (1).

Hayward, C. (1994). Listening to the Earth Sing. In G. Kramer, (Ed.) Auditory Display, I, Santa Fe Institute, Studies in the Sciences of Complexity Proceedings. (Vol. XVII, pp. 369-404) Reading, MA: Addison-Wesley.

Jackson, J. A. & Francioni, J. M. (1992). Aural Signatures of Parallel Programs. In Proc. Twenty-Fifth Hawaii International Conference on System Sciences, pp. 218-229.

Jackson, J. A. & Francioni, J. M. (1994). Synchronization of Visual and Aural Parallel Program Performance Data. In G. Kramer (Ed.), Auditory Display, Santa Fe Institute, Studies in the Sciences of Complexity Proceedings (vol. XVIII, pp. 291-306). Reading, MA: Addison-Wesley.

Jameson, D. H. (1994). Sonnet: Audio-Enhanced Monitoring and Debugging, in Auditory Display. In G. Kramer (Ed.), Auditory Display, Santa Fe Institute, Studies in the Sciences of Complexity Proceedings (vol. XVIII, pp. 253-265). Reading, MA: Addison-Wesley.

Kramer, G. (1994). Preface, In G. Kramer (Ed.), Auditory Display, Santa Fe Institute, Studies in the Sciences of Complexity Proceedings (vol. XVIII, pp. xxiii-xxxviii). Reading, MA: Addison-Wesley.

Mezrich, J. J., Frysinger, S. & Slivjanovski, R. (1984). Dynamic Representation of Multivariate Time Series Data. Journal of the American Statistical Association, 79(385), 34-40.

Pierce, G. W. (1949). The Songs of Insects. Harvard University Press.

Scaletti, C. (1994). Sound Synthesis Algorithms for Auditory Data Representation, In G. Kramer (Ed.), Auditory Display, Santa Fe Institute, Studies in the Sciences of Complexity Proceedings (vol.

XVIII, pp. 223-252i). Reading, MA: Addison-Wesley.

Scaletti, C. & Craig, A. B. (1990). Using Sound to Extract Meaning from Complex Data. In E. J. Farrel (Ed.), Extracting Meaning from Complex Data: Processing, Display, Interaction. (Vol. 1259, pp. 207-219). San Jose, California: SPIE.

Sonnenwald, D. H., Gopinath, B., Haberman, G. O., Keese, W. M., III & Myers, J. S. (1990). InfoSound: An Audio Aid to Program Comprehension. In Proc. Twenty-Third Hawaii International Conference on System Sciences, 11, pp. 541-546. IEEE Computer Society Press.

Authors

Paul Vickers

Course Leader for BSc (Hons) Computer Studies

fax: (+44 151) / 0151 207-4594

Tel & voicemail: (+44 151) / 0151 231-2283

<http://www.cms.livjm.ac.uk/www/homepage/cmspvick/index.htm>

School of Computing and Mathematical Sciences

Liverpool John Moores University

Liverpool, UK

p.vickers@livjm.ac.uk

James L. Alty

LUTCHI Research Centre

Department of Computer Studies

Loughborough University

Loughborough, UK

j.l.alty@lboro.ac.uk