Challis, B. P. and A. D. N. Edwards (2000). *Weasel: A system for the non-visual presentation of music notation*. Computers Helping People with Special Needs: Proceedings of ICCHP 2000, pp. 113-120, Karlsruhe, Germany, Osterreichische Computer Gesellschaft.

Weasel: A System for the Non-visual Presentation of Music Notation

Ben P Challis and Alistair D N Edwards
Department of Computer Science, University of York, YO10 5DD, UK
bpc@cs.york.ac.uk alistair@cs.york.ac.uk

Abstract
A computer-based approach to the delivery of music notation for blind people is described. Existing non-visual alternatives present musical information in a serial fashion which places a high cognitive demand upon the reader who, in effect, must read all the information regardless of its significance to any particular learning task. The Weasel notation system has been designed to address this problem using interactive high-level tactile representations, synthetic speech and audio playback.

1. Common music notation and non-visual alternatives

Common music notation (CMN) is the most widely accepted method for presenting music in both educational and professional settings. The main strength of CMN lies in its graphical nature where the reader can visually process great quantities of information simultaneously. Using this graphical approach, a music reader can 'sort' the musical data with considerable ease. The way in which a musician of any level chooses to approach an unfamiliar piece of music for the first time is highly individual. A particularly able music reader can be capable of sight-reading the majority of musical notes, symbols and textual instructions whilst simultaneously producing a performance based on that information. In contrast, a less confident reader will work only with information that is useful within a given task. For example, learning the order of pitches or durational values, concentrating on instrument specific technique such as fingering, bowing and picking or perhaps focusing on dynamic changes within an extract. Filtering out the information which is not significant to a learning task is a simple visual task when using CMN.

Figure 1. An example extract of common music notation (CMN).

Although CMN is without doubt a very powerful and flexible system of music notation, the fact that it so visually dominated means that blind musicians and music learners cannot use it at all; instead, they must rely on non-visual alternatives. With existing non-visual approaches, such as braille music and Talking Scores, there are compromises to be made. Either all the information within the original CMN will be translated producing large quantities of serial instructions or some of the information will need

to be removed with a potentially high cost in terms of loss of performance accuracy. Although the latter might be acceptable during the early stages of learning and might even facilitate basic structural outline to be included, as with Ockelford's suggested adaptation of braille music [6], in the long-term this could not provide the basis for serious musical study. In addition, there would also need to be many different editions of each piece depending on the nature of the learning task in hand.

1.1 Braille music

Besides being difficult to learn, braille music also requires the general ability of being able to read conventional braille when the reality is that only a small percentage of the blind population can read braille [2]. At even the most fundamental level there are complications for the braille music learner as the characters used for the musical alphabet of *ABCDEFG* correspond to the braille letters of *IJDEFGH*. Perhaps the most obvious drawback of braille music is the serial nature in which it is presented; the reader must read everything, regardless of its significance within a learning task. This, in addition to the way in which the braille music system produces large amounts of braille cells, makes it very difficult for the reader to have any idea of structural 'shape' within a piece. With CMN, the reader uses the number of bars-to-a-line, lines-to-a-page, repeated sections and dynamic flow to build up a basic outline of the progression of the piece; this is lost with braille music. The braille music shown in Fig. 2 is a translation of part of the extract of CMN shown in Fig. 1. Already, it can be seen that a considerable amount of cells are generated from even a small extract and that there is easy no method for filtering out unnecessary information.

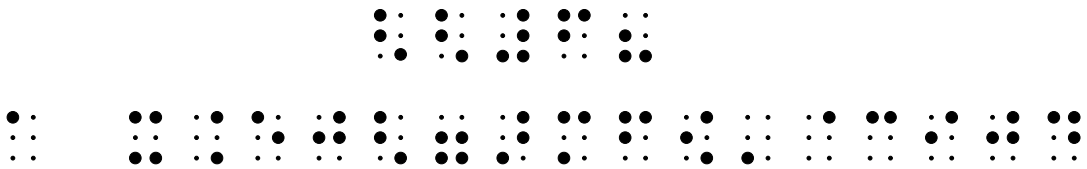

Figure 2. An example of braille music notation. This is a translation of the lead-in and first bar of the extract of CMN shown in Fig. 1. The five cells at the top show the key-signature and time-signature and would only appear at the start of the piece.

1.2 Talking Scores

Talking Scores[1] are spoken descriptions of music notation recorded onto audio cassette. As with braille music, all the symbolic and written instructions are translated resulting, again, in large amounts of information presented in a serial fashion. A translation of the same extract as Fig. 2 would be:

> "The key signature has one flat and the time signature is six eight. Bar one, quaver rest, fifth octave d quaver, b flat quaver. Bar two, begin repeat, mezzo forte, a dotted crotchet tied to a quaver, b flat quaver, c quaver"

[1] Talking Scores are a development of the Royal National Institute for the Blind, UK.

Although the prerequisite of being able to read braille is removed the cognitive demand placed upon the reader still remains high and the learner's reading speed is directly limited by the speed of the spoken description.

2. The Weasel project

The Weasel music notation system is a computer-based system for the non-visual delivery of musical information. Using a combination of PVC overlays, an Intellikeys[2] touchpad, speech synthesis and audio output, blind musicians and music learners are able to explore pieces of music using strategies that are more flexible to their requirements. The central feature of the Weasel music notation system is the provision of a high level tactile representation of a page of music with which readers can interact, allowing them to focus on specific learning tasks. Fig. 3 shows an example outline of a typical Weasel overlay. The system produces audio playback or spoken descriptions of music information using synthetic speech. The main purpose of the system is to provide a flexible method by which a user can read and organise the musical information they wish to work with.

2.1 Design

The design of the tactile interface for the system was based on three preliminary design principles:

1. A consistency of mapping between the visual and non-visual representation should be maintained to enable blind and sighted users to use common terms of reference.
2. The tactile representation should focus on data that is static. Only static display technology can achieve the levels of granularity required to create fine textures and tactile patterns. Therefore, it is more appropriate that the more dynamic elements of the display be presented using audio which can accommodate such change.
3. Height should be used to provide a method of filtering information.

From the beginning of the Weasel project, it was apparent that user testing of an interface that involved tactile interaction would benefit from reference measurements of each user's existing ability at performing various tactile tasks. This way, if a user were to exhibit poor ability at interacting with the main tactile interface then comparisons could be made between those observations and the results of that individual's tactile aptitude pretest. A pretesting method was established based on research into tactile graphics including contrasting line types in line tracing [1], identifiable areal symbols [4] and point symbol design [3,5]. The overall design of the tactile aptitude pretest produced a number of useful guidelines for raised-line diagram design. These were transferred into the design process for the Weasel interface and are as follows:

1. *line width* – three easily identifiable line widths of 0.5mm, 2mm, 4mm
2. *line height* – three easily identifiable heights of 0.5mm, 2mm, 4mm
3. *line type* – the frequency of dots in broken lines is a significant factor when discriminating between different line types and confusion will arise between lines that have similar frequencies, even if this is from gaps rather than dots.

[2] Intellitools Inc, 55 Leveroni Court., Suite #9 Novato, CA. 94949

4. *parallel lines* – a minimum gap of 2mm is required between parallel lines to prevent them from being perceived as one.
5. *line tracing* – a line width of 1-2mm is appropriate where line tracing is to be involved e.g. following a guideline.
6. *point symbols* – symbols larger than a finger tip become complex to explore and are therefore slower for recognition. Small, simple symbols are more appropriate such as hollow and filled outline shapes (circle, triangle, square, diamond). Symbols with more than four sides quickly become ambiguous feeling generally circular in shape.
7. *areal symbols* – general texture is more immediate than specific pattern such that a user is likely to discriminate between the overall textural property of patterns (rough, smooth) before exploring the pattern in detail.

Each overlay consists of two areas; an interactive structural overview (Fig. 3) and a control section below it (Fig. 5). In many respects, this looks like a page of CMN but without any musical events within the bar areas. References such as "the second bar of the bottom line" or "the last bar of each line" are equally valid within both the visual and non-visual representation. A horizontal guideline is positioned at the bottom of each line of music with vertical barlines meeting this to define a number of bar areas for that line. The barlines are slightly higher than the guideline which is approximately 1mm high. This is to help the user locate a specific bar quickly. Double barlines (often used to represent the start or end of a section) are spaced 2mm apart and repeat marks are represented using a high dotted line (4mm). First-time and second-time bars use contrasting textures within the whole bar; first-time are smooth whilst second-time are rough.

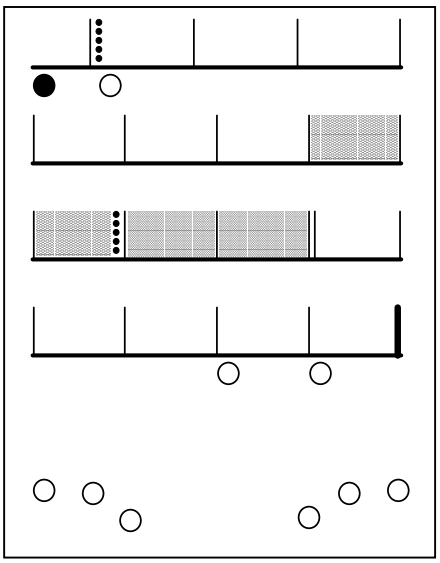

Figure 3. A specimen outline of a Weasel overlay. 32cm high by 21cm wide.

The bar areas are 'active' so that when a user presses in one the musical content is sounded. There are four modes for exploring the music depending on the nature of the task. *Bar* mode allows the user to select a single bar and listen to its contents. *Browse* mode allows the user to add new bars to a queue to be played back so that continued playback can be achieved. *Skip* mode allows the user to instantly move playback to the beginning of a new bar ceasing the playback of any current bar. The last mode, *block* mode, allows the user to define a section of bars to be rehearsed repeatedly.

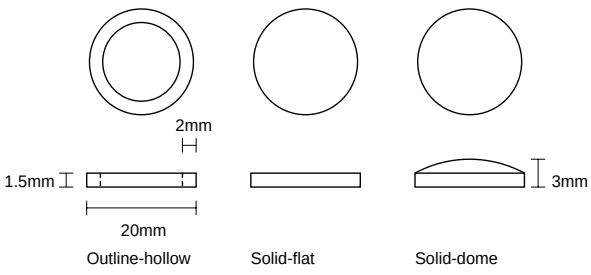

Figure 4. Detail of the three circular symbols.

Circular symbols are located beneath the guideline. There are three types: outline-hollow, solid-flat and solid-domed (Fig. 4). The outline-hollow symbols represent the presence of a dynamic marking. When pressed, synthetic speech is used to present the contents e.g. "mezzo forte". A solid-flat symbol shows the existence of a list of text items belonging to a bar, one of which is a dynamic marking. The user presses the

symbol once to hear the currently active item within the list (dynamic marking by default). By double-clicking the symbol the next member of the list becomes active in place of the current. This eventually wraps-around to return to the dynamic marking within the list. The solid-domed symbol is similar to the solid-flat in that it shows the presence of a list of text items. However, there is a significant difference in meaning in that there is no dynamic marking within this type of list.

At the bottom of the overlay there is a control section (Fig. 5) which allows the user to navigate through a menuing system for selecting different modes of interaction and information retrieval. The controls are designed to accommodate the index, middle and ring fingers of each hand. The left hand changes *options* and the right hand changes *items*. Each option has a number of items associated with it. The options are playback mode, tempo, metronome and keyboard click. For example, the playback mode *option* has the *items* of *bar, browse, skip* and *block*. The user can cycle through available *options* or *items* by pressing the leftmost or rightmost of the finger locators which are displayed as hollow-outline symbols. When the centre symbol of either hand is pressed the current setting for that control (*option* or *item*) is spoken using synthetic speech. The list of *options* and *items* do not wrap-around in the way the previously described lists do, instead the user will simply reach the end of the list and will have to travel in the opposite direction to reach the other end.

Figure 5. Control section of tactile overlay (outline hollow symbols).

2.2 User testing

To assess the system, a group of six musically able volunteers were first trained how to use the interface and then asked to perform a number of tasks within it. Capable musicians were used in preference to learners to help clarify the results in terms of ability in using the system for information retrieval rather than ability at learning to read music. Ideally, the entire subject group would have consisted of blind musicians but practical problems prevented this. This is an area that Stevens [7] has described as a common problem in this kind of evaluation. A solution is to use sighted people interacting non-visually, at least in the early phases of design where the intention is to involve more blind testers at a later stage of the design process. With this in mind, five of the group used were sighted users who wore blindfolds during both training and assessment and the sixth member was a congenitally blind student from a local college. It was anticipated that the inclusion of a blind member within the group would help validate the results obtained in terms of applicability to blind users in general.

Each subject was first trained in recognising and interacting with five tactile overlays each of which represented a short extract of music. Once familiar with the various symbols and controls the user entered into an assessment stage. At this point, two new overlays were presented to each subject who would then be asked to perform a number of tasks. The first of these tasks required the user to progress through the extract from start to finish describing the tactile components that were contained within each overlay and what these symbols represented. The purpose of this initial task was to establish whether the tactile components used within the overlays were easy to find and

recognise. The rest of the tasks required the users to carry out a series of step-by-step instructions in the way that a musician might be expected to rehearse a new piece of music. Two examples of these tasks are shown below:

1) Set to block mode, select bars 2 to 8, set tempo to 150 beats per minute, playback block.

2) Set tempo to 115 beats per minute, set to skip mode, play a little of the second bar of each line.

In total there were eleven of these navigational style tasks for each overlay. However, as with any assessment that involves working closely with a small group of users, many observations were made during these tasks and these contributed greatly to the overall results.

2.3 Results

In general, the users were able to perform the majority of tasks with relative ease and from that perspective the approach adopted within the Weasel project seems to be very promising. In the initial task of locating and identifying all the tactile components on the overlays, the users proved to be very efficient. Even in the second set of tasks, there were only minimal mistakes made, the most noticeable of which involved two subjects failing to include the top line of music when counting down to find a specific line. As a result of this, the task would be performed using the correct bars but located on the incorrect line. Using the touchpad in portrait orientation could be regarded as awkward as users are expected to stretch their arms to reach the furthest areas of the overlay.

A number of other areas of concern were also observed during training and assessment. One was that the basic visual-to-tactile mapping had created an excess of what is probably best described as *empty space*. The term *empty space* is used with reference to areas within an overlay that do not communicate anything. It was relatively easy for a user to place a finger into an area of the display and not locate an informative area. Once in this situation, it was easy to remain confused with little idea of how to progress.

As described earlier, two types of symbol found beneath the guideline could be pressed using either a single-clicking or double-clicking action. Observation showed the double-clicking procedure to be awkward to achieve for some of the users. A general lack of haptic feedback within a static display appears to make this style of interaction inappropriate. Users were quickly frustrated when they had to repeat the action a number of times before being successful.

In the list of items belonging to the play-mode option, the user progressed through the list in one direction if the leftmost *item* button was pressed and in the opposite direction if the rightmost was pressed. However, many of the available options, such as *repeats*, were either *on* or *off* and were not presented as lists. Instead, both the leftmost and rightmost buttons had the same function; the current state of the option was reversed. The users expressed a preference for the two states to be presented as two items and for the outer control buttons to always work in the same way.

Some users commented that the audible 'click', provided to reinforce pressing actions on the overlay, was annoying as it sometimes interfered with the rhythm of the musical playback. In *browse* mode this was particularly true where it was possible to press a bar during playback to add it to the queue therefore receiving an out of rhythm 'click' as playback continued. It seems likely that where the data being manipulated is already in the auditory domain, as is the case with most music applications, and the interface employs auditory reinforcement, there might be a fairly low user threshold for auditory overload. One user even found the use of a metronome facility quite distracting during the training sessions.

3. Weasel II and future work

The results and observations made during the assessment of Weasel have led to a number of design changes that are currently being implemented in Weasel II. The touchpad is now used in a landscape orientation to provide a more accessible surface area to explore. To reduce the quantity of *empty space* within the interface a new strategy for exploration has been incorporated into the new series style of overlays. The guideline and bar areas are now presented as a single strip. The guiding strips are raised from the background by approximately 200 microns and each strip is no further away vertically from the next than 2cm. Once the finger leaves one guide-strip then its neighbour, if there is one, will be encountered immediately. Each line of bars is 2cm high (the approximate length of a fingertip) by 30cm wide with all the information required by the user included within each of these strips. The need for exploration to locate symbols is minimised; if there is a symbol present then it will be found by simply following the line.

The barlines are now wider at 1.2cm (the approximate width of a fingertip) and are at one of three different levels of height, each representing the number of performance instructions within the coming bar area. Although the symbols for repeat marks were identifiable within the first version of Weasel these have been replaced by 'ramps'. These travel from the top edge of a barline down into the adjacent bar area providing a more physical indication of the direction in which the repeat mark is pointing. Textured bar areas are still used to represent the first-time and second-time bars but third and fourth time bars are also being included as these are frequently used in song-line style books of popular music to help condense pieces onto one or two pages.

In the preliminary version of Weasel, the design emphasis was on the provision of an interactive structural overview of a page of music along with efficient controls for moving between different modes of information retrieval. For assessment purposes, only audio playback was implemented so that the basic framework could be assessed against the original criteria. More sophisticated means of filtering and delivering the musical information will be included in the new version. Spoken descriptions similar to the Talking Scores approach will be available along with different levels of audio retrieval e.g. left hand or right hand for keyboards, rhythmic elements only, pitch information only and combined pitch and rhythm. A refreshable braille display could easily be included if considered desirable. One final feature to be explored within Weasel II is the addition of a force-feedback mouse to allow the user to explore the contents of a bar, receiving haptic feedback about note and symbol positions within the bar. This would be horizontally and vertically allowing the user to easily explore harmonic relationships when polyphonic sections are encountered.
All the equipment used within the project is relatively inexpensive and is also easily available. However, consideration is also being given to how the overlays can be manufactured more easily as

these are currently quite time consuming to produce. As a result of this overall research, it is also envisioned that a more complete set of design principles for the general integration of tactile interaction into the human-computer interface will be achieved. These principles can then be applied to the design of other auditory-tactile interfaces to assess their validity within different applications. It is also anticipated that this, more complete, set of design principles could also have significant impact on the design of tactile diagrams in general.

4. Conclusion

The Weasel notation system provides the basic framework by which a blind musician can explore a piece of music in a structured way. The system provides flexibility to meet the requirements of individuals who will have different strategies for learning unfamiliar pieces. The main strength of the Weasel system is that no information is lost as part of a translation process. All of the original musical information is available at all times ready to be displayed in a form that suits the preference of the reader.

5. Acknowledgements

This research has been funded by the Engineering and Physical Sciences Research Council (award ref. 96309035). An Intellikeys touchpad was kindly made available by Intellitools Inc.

6. References

[1] BENTZEN, B. L. and PECK, A. F. (1979). Factors affecting the tracability of lines for tactile graphics. Journal of Visual Impairment and Blindness. 73(7), 164-269.

[2] BRUCE, I., McKENNELL A. and WALKER, E. (1991). Blind and partially sighted adults in Britain: the RNIB survey. Her Majesty's Stationary Office, London.

[3] LAMBERT, L. M. and LEDERMAN, S. J. (1989). An evaluation of the legibility and meaningfulness of potential map symbols. Journal of Visual Impairment and Blindness. October, 397-403.

[4] LEDERMAN, S. J. and KINCH, D. H. (1979). Texture in tactual maps and graphics for the visually handicapped. Journal of Visual Impairment and Blindness. 73(6), 217-227.

[5] NOLAN, C. Y. and MORRIS, J. E. (1971). Improvement of tactual symbols for blind children: Final report. American Printing House for the Blind, Louisville, Ky., USA.

[6] OCKELFORD, A. (1996) Music Matters. RNIB publications. ISBN 1-85878-071-3

[7] STEVENS, R. D. (1996). Principles for the design of auditory interfaces to present complex information to blind people. D.Phil. Thesis, Department of Computer Science, University of York.