SOUNDFIELDS AND SOUNDSCAPES: REIFYING AUDITORY COMMUNITIES

Iain McGregor, Alison Crerar, David Benyon

School of Computing
Napier University,
10 Colinton Road, Edinburgh, EH10 5DT,
Scotland, UK
i.mcgregor, a.crerar,
d.benyon@napier.ac.uk

Catriona Macaulay

Department of Applied Computing
Duncan of Jordanstone College of Art
University of Dundee,
Dundee, DD1 4HN, Scotland, UK
catriona@computing.dundee.ac.uk

ABSTRACT

This paper reports progress towards mapping workplace soundscapes. In order to design auditory interfaces that integrate effectively with workplace environments, we need a detailed understanding of the way in which end users inhabit these environments, and in particular, how they interact with the existing auditory environment. Our work concentrates first on mapping the physical soundfield, then overlaying this with a representation of the soundscape as experienced by its active participants.

The ultimate aim of this work is to develop an interactive soundscape-mapping tool, analogous to the modeling tools available to architects. Such a tool would be of use to designers of physical, augmented and virtual environments and usable without professional musical or acoustical expertise.

1. INTRODUCTION

In the study of Human-Computer Interaction (HCI), sound has been largely ignored as an output medium in favour of visual displays. The notable exception to this being work on ‘ear-cons’ (auditory icons) and sonification [1]. Indeed the design of the auditory environment in general has received relatively little attention, even in areas such as interior design.

We see a significant change on the horizon for HCI and sound and this has prompted an interest in the design of soundscapes. With the trend towards personal, mobile, ‘always on’ technologies coupled with the proliferation of ubiquitous technologies embedded in physical environments, it will not be sufficient for future designers to focus on the interaction of a single person with a device. HCI will involve the design of dynamic ‘information spaces’ [2] and the complex interactions of multiple users with multiple intelligent artifacts. In this scenario, an important part of realizing more hands-free and spontaneous interactions within such information spaces, will be designing sounds that convey appropriate information and that work together in some harmonious fashion. We see sounds as both inputs and outputs, not just helping people monitor the environment or being alerted to exceptional events. We see the ‘orchestration of the soundscape’ as being indispensable to the fabric of future information spaces.

At present the design skills of sound engineers are not easily transferred. Our motivation is to find methods and techniques for making the design of soundscapes available to interaction designers: we are looking for a mapping tool that can identify the attributes of sounds in any environment.

Besides creating a list of the sound events within an environment, there are a number of other potential uses for soundscape mapping. The most obvious is that of an intelligent noise map, where rather than removing sound-generating objects from the environment, we can establish what is necessary or desirable, and take action about those sounds which are considered background and redundant by a sufficient number within the environment.

Soundscape mapping could also be used to test how an additional sound-generating object would affect the pre-existing environment, for example, might it either physically or conceptually mask other sound events? Likewise, soundscape mapping could be used to test augmented environments, for example, in order to understand how wearing a single earpiece affects the interpretation of the previous environment, or how complex auditory interfaces affect traditional working practices.

Listing the sound sources and how they are used is also useful when developing auditory interfaces, in that it can inform the designer about what the interface has to sonically compete with and what it can replace. A virtual soundscape map could be made in order to ‘test’ a spatial environment, to establish whether the spatial aspects are being perceived as expected. A more limited method could be used to check whether all of the auditory elements of an auditory interface are appropriate and clearly heard with different hardware and operating conditions.

This paper explains the development of a method for classifying sounds and producing soundscape maps. Section 2 touches on seminal work by others in this area and introduces Macaulay and Crerar’s sound classification scheme. Section 3 describes how the scheme was first utilized with fourteen volunteers in private workspaces. Section 4 reports a more extensive use of the scheme in the context of a busy open-plan office, the methods employed to collect data, and the observations made. Interesting observations from this work and proposed future developments are described in Sections 5 and 6.

2. THE SOUNDSCAPE CONCEPT

The term ‘soundscape’ is derived from ‘landscape’ and can be defined as the auditory environment within which a listener is immersed. This differs from the concept of ‘soundfield’, which can be defined as the auditory environment surrounding the sound source, which is normally considered in terms of volume, duration, location and frequency range.

The concept of the soundscape is not new: Grano first differentiates between the study of sound and noise in 1929 [3]. The principle then lay comparatively dormant until 1969, when Southworth tried to establish how participants perceived the sounds of Boston and how this affected the way they saw the city [4]. Schafer [5] and Truax [6] (in the form of the World Soundscape Project) attempted to formalize the concept using visual representations such as soundmarks rather than landmarks. Rodaway [7], however, raised concerns about Schafer's reliance on visual metaphors, without doubting the relevance and importance of reifying the soundscape.

Studying the soundscape in addition to considering soundfields allows us access to what Augoyard [8] defines as the six fundamental dimensions of "sonic phenomena":

1. Physical signal
2. Environment
3. Perception
4. Cultural Representations
5. Sonic actions
6. Social interactions

This conforms to a traditional scientific observation from multiple perspectives rather than purely recording the physical phenomenon, allowing us to tap into what Rodaway refers to as the "perceived environment" [7].

While in our work we do not use Schafer's terminology of 'keynote', 'soundmark' and 'signal', we do consider three out of four of his sound contexts: acoustics, psychoacoustics and semantics. Aesthetics have not been considered at this stage, but it is intended to include them at a later date.

2.1. Soundscapes and Interface Design

To date, auditory elements within interfaces have tended to be short, discrete sounds, without any context or background, with more adventurous use of engineered soundscapes being the preserve of the computer games industry. Frustrated by the lack of appropriate auditory models for the interaction designer, Macaulay and Crerar studied the work of Brewster, Feld, Gaver and Truax as a basis for formulating a soundscape classification more appropriate to the field of Human Computer Interaction (HCI) [9]. The resulting model provides interactive systems designers with a framework for classifying sounds, which is a preliminary step in the move away from today's visually saturated interfaces.

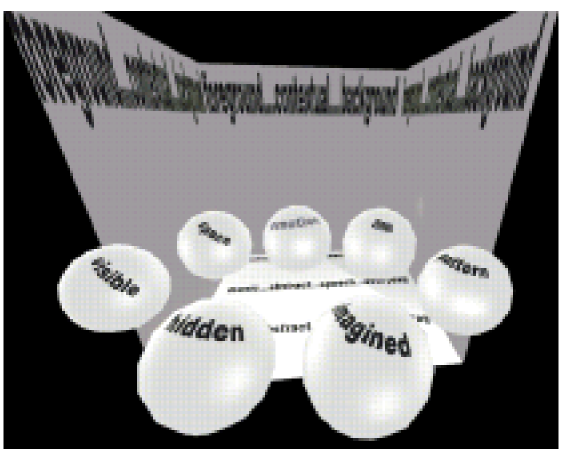

Figure 1: Visualization of sound classification reproduced from [9]

Macaulay and Crerar proposed a method of classifying constituents of soundscapes based upon (i) sound type, (ii) information category and (iii) acoustical information (see Figure 1). The sound type is broken down into *music*, *speech*, *abstract* and *everyday*. (Subsequently we have found that the 'abstract' and 'everyday' concepts are more readily described as 'other known' and 'other unknown'. Moore [10] points out the Boolean nature of the perception of sound, as either being perceived as speech or not).

The information categories are: *visible*, *hidden*, *imagined, patterns*, *passing of time*, *emotions* and *position in Euclidean space*. This allows us an insight into the information content provided to the soundscape inhabitant.

Finally the model includes acoustical information, or the level of listening: which can be either *foreground*, *contextual* or *background*. Foreground sounds are those with which the listener actively engages, contextual sounds provide an underpinning to the foreground, and background are all of the other 'ambient' sounds, often not attended to.

3. CLASSIFYING THE SOUNDSCAPE

In order to take this sound classification scheme, test its validity and develop it into a usable method, fourteen academic volunteers were recruited for a study of their workplace soundscapes. In the first part of the study participants were asked to spend fifteen minutes listening to the environment and making notes of the sounds heard, without reference to their physical properties (frequency, sound pressure level (SPL), location, etc.). After this they were asked to classify the sounds in terms of the Macaulay/Crerar model described above. This produced responses varying from 7 to 21 separate auditory events from the participants, despite a considerably larger number being heard by the first author (who is an auditory expert).

On the whole this method proved to be a very simple and effective way of eliciting responses from regular, intermittent and new inhabitants of the auditory environments. However, two main problems were found: a) participants routinely stopped what they were doing in order to listen and therefore no longer generated sounds, and b) participants frequently forgot to include the sounds that they themselves created.

Consequently in the second part of the study the first author noted all of the sounds heard over a longer period. After the event he questioned the participants about how they would classify the sounds using the Macaulay/Crerar scheme. A considerably larger list of sounds resulted; in one case fifty-nine distinct individual sounds were noticed within one hour. The participant was aware of all of the sounds and was happy to give in some cases quite elaborate 'natural' descriptions.

The results illustrated the requirement for a form of dynamic classification. For example, within a single auditory event, a sound could move from foreground, through contextual to background and could have multiple information categories as time progressed.

The concepts of 'everyday' and 'abstract' sound allowed us to distinguish between novice and experienced soundscape inhabitants. In a complicated auditory environment, it would be expected that the novice would classify a greater percentage of the sound events as abstract, (unknown or unidentifiable), whereas it would be expected that very few sound events would be classified as abstract by the experienced listener. Experienced listeners would also consider a greater percentage of the sounds as being background, as they have learned to ignore them, or give them little relevance. Tracking changes in the sounds reported and in the classification of them could be an effective way of monitoring how long it takes an individual to interpret accurately all of the auditory elements of a system, in a real or virtual environment.

4. PROPOSED METHOD FOR MAPPING SOUNDSCAPES

Having tested Macaulay and Crerar's classification in isolation, we developed an ancillary method for supplying further acoustic and semantic information. This took the form of an observational study followed up by interviews. The resultant method was applied to a single location: in this case an open-plan office with regular, intermittent and new inhabitants. This environment was suitable for conducting a few informal unrecorded observations, in order to establish the type of problems, which might arise prior to the formal study.

The method is intended to produce sound maps of different 'scales' and emphases. We can overlay the soundfield with different listeners' perspectives, specifically regular, intermittent and new inhabitants of the auditory environment under study. Through further refinement it will be possible to concentrate on when the listener is trying to perform a specific task, or how their perception of the environment changes, for example when they are winding down at the end of the day.

There were three phases to this open-plan office study. For the first phase, a two-hour period was chosen. A recording was made in stereo, which aided subsequent logging of sound events. To achieve this, a pair of omni-directional tie-clip microphones were mounted on a stand 2m high separated by 15cms. Recording was straight onto DAT tape at 48kHz 16 bit, this gave us a dynamic range of 96dB, and a frequency range of 20Hz-24kHz, considerably exceeding the normal frequency range of listeners. Schafer [11] noted that the use of a microphone inhibits the 'cocktail party' effect [12], and thus we found that listening back to the recording allowed a more accurate impression of how the soundscape is heard by new inhabitants, and helped to reduce reliance on the real-time auditory interpretation of the first author.

At the completion of the recording, where possible, recordings and SPL readings of individual sound events were taken, in order to be able to calculate an approximate sound pressure level (both A and C scale) and frequency range. A still photograph was also taken of the sound source in order to aid interpretation, as well as recording the physical position and directivity pattern. This all proved essential in assisting considered responses about the multitude of sound events during the subsequent interviews with the inhabitants of that soundscape. In addition, the first author noted sound events during the two-hour recording period, keeping track of their time-code, and position (with reference to a grid previously prepared representing the room). These notes were collated and a candidate list of the sound events was produced.

In phase two, during the same two-hour time period a week later, the participants were asked to list all the sound events they heard and record the position of the source according to the grid representing the room. They were asked only to pause momentarily to make notes rather than stop and actively listen. This was to ensure that all of the sounds that the participants usually created during their work would still be present, thus ensuring that they could be heard by both themselves and their colleagues. This did mean that certain sound events were missed, but it did allow for object-oriented descriptions of the events, (describing the object rather than the sound event or cause), and also helped to prevent all the participants stopping at once, as not all of them heard every event, and naturally individuals were loath to notate every single occurrence of repetitive events.

While no two periods will ever have the same soundscape, using two consecutive periods in this way ensured that a realistic recording was obtained in the first phase of research, while during the second note-taking phase participants could ask questions freely and thereby not spoil a recording in progress.

A week later, the third phase of the work was carried out, again during the same time period. Participants were questioned about all of the sound events that they had notated previously, as well as those experienced by the observer. We encouraged interviewees to give phenomenological descriptions prior to classification (Moran [13] defines phenomenology as a method of describing things as they "appear to consciousness" (p.6)); this allowed participants to introduce factors that might be missing from the classifications. The questioning took three forms using: Macaulay and Crerar's classification; Delage's interactive functions [14], and finally Gaver's hierarchical descriptions [1]. Gaver's descriptions were only requested when the sound type was classified as either abstract or everyday.

4.1. Interactive Functions

Questioning the inhabitants of the soundscape about an individual sound event's interactive function(s) gives us an insight into its perceived semantics. Not only can we see where listeners share interpretations, but also where there is a mismatch between the intended design of a sound event and its common interpretation.

Bernard Delage [14] and Helen Engelen in 1998 as part of a series of seven discussion meetings involving: architects; acousticians and electro acousticians; sound and visual designers; computer scientists; composers and scenographers, developed amongst other things a list of interactive functions associated with sound events, and specifically identity, namely: *warning, assisting, incitement, monitoring, reassurance, forgiving, guiding, protecting, relaxing*.

The study reported here threw up an additional interactive function of 'confirmatory' (the printer had finished printing) although this could perhaps be accommodated under 'reassurance'. Some participants classified a number of sound events as redundant, the use of 'redundant' rather than 'noise' proved useful as a number of participants referred to a considerable number of the everyday and abstract sounds as 'noise'.

4.2. Hierarchical description of simple sonic events

Gaver [1], with reference to 'everyday listening', described sound events in terms of 'interacting materials': from our perspective it does not matter what the actual materials were, more what they sounded like. This is intended to reflect the way in which we interpret the cause of the sound, rather than being concerned with volume, pitch, time etc. Gaver splits them up into three categories: vibrating, aerodynamic and liquid, each with their own subcategories: *impacts*, *scraping* and *others* for vibrations; *explosions* and *continuous* for aerodynamic; and *dripping* and *splashing* for liquid.

We are using this classification as it stands in order to extend Macaulay and Crerar's 'sound type'. We hope to expand on Gaver's use of 'other' with reference to vibrations as the work progresses.

5. MAPPING SOUNDSCAPES

We have described preliminary use of a soundscape mapping method applied to a single workplace environment, in this case the Departmental Office of the School of Computing at Napier University. The environment had five permanent inhabitants, and a large number of intermittent inhabitants (students and academic staff). The predominant sound was that of speech, with a single music source and a variety of everyday sounds such as telephones, traffic and typing.

We found that intermittent inhabitants have a considerable impact on the soundscape, starting conversations upon entering the room and continuing them while not even facing the person they are conversing with, and doing this irrespective of whether other inhabitants are disturbed by them. The practice of the 'out loud' [9] is of benefit to the regular inhabitants, who make queries from one end of the room to the other, as well as voicing non-specific comments, which are intermittently listened to, sometimes causing laughter. One regular inhabitant, through the use of a barely audible web radio station, customized the auditory environment around her desk, carving out a virtual space. This was the only instance of auditory personalization and was almost always masked by other sounds throughout the time periods studied.

There was a dramatic disparity between the number of sounds that occurred and those reported by the participants, in some cases a ratio of 3:1. This can be partially explained by the fact that the participants are adept at relegating sounds to the 'background', but is also due to the need to learn to ignore the high level of auditory interruptions experienced throughout the day in that particular shared environment.

Intermittent inhabitants have the greatest effect upon the soundscape through their desire to hold conversations across the room. But the regular inhabitants are adept at altering the volume of their speech according to who is present. During the period under study one participant became aware of how loud the street sounds were from an open window and closed it. During a preliminary observational session, the participants became aware of how loud a hinged countertop was (over 100 dB (A)), and subsequently stopped using it. But apart from these instances the regular inhabitants had little control over their auditory environment, and have shown little interest in wanting to explicitly influence their auditory environment.

5.1. Information Cartography

The next stage of the research is to develop an appropriate computer-based visualization to capture and dynamically interpret all of the data. Whilst visualizing the soundscape might at first seem counter-intuitive, it does allow psycho-acoustic aspects to be overlaid on the acoustical environment. After investigating a variety of forms of visualization such as Isobel maps, waveforms, spectrograms and musical notation we initially decided upon graphs of SPL vs. frequency, in order to establish potential masking and easily compare the physical aspects of a variety of sound events. However it is not suitable for capturing spatio-temporal data or any of the sound classifications we have discussed. The best candidate seems to be a Geographical Information System (GIS), which will allow the inclusion of all of the missing elements within a single map.

The process of animated cartography has already been used successfully by Servigne et al. [15] for mapping city noise. We believe that this concept can be extended by combining traditional cartography with that of information cartography. This combination should allow us to present seven dimensions at any time, whilst also allowing 'snapshots' and aggregations. Old [16] points out that information cartography provides five standard functions: pan, zoom, 3D rotation, layering and transparency, as well as nominal, ordinal and quantitative features.

Ultimately we envisage linking the resultant maps to the relevant recordings, enabling end users to listen to individual elements within the soundscape as well as to the overall Soundscape with all its auditory complexity. The eventual goal is to enable users to navigate their way through the auditory map, choosing from alternative points of listening, as well as direction, and auditioning the overall characteristics of the soundscape from the perspective of regular, intermittent or new inhabitants.

6. CONCLUSIONS

The soundscape mapping method described here produced both a soundfield and a soundscape, which have highlighted important issues for the design of auditory displays. This allows us an insight into the way in which the soundscape is experienced by regular, intermittent and new inhabitants of the auditory environment, and incorporates semantics as well as psychoacoustics and acoustics in a single method.

Moore [10] refers to the difficulty of even attaining accurate measurements of loudness, never mind pitch, reminding us of its qualitative nature. The mere fact that we stop to listen takes the event out of context, and as we know from HCI, it is crucially important that we understand context, especially if this method is to be used by designers. As Porteous and Mastin found, due to the variability of individual perceptions, any form of classification is difficult to achieve [17]. Part of this research is to bring as many classification systems together as possible, and apply a range of these as appropriate to each type of environment.

While the method is time-consuming, it is commensurate with the amount of time taken to create a visual map, especially if we were to include all of the movements of the participants. We are still working on a form of visualization, and are keen to incorporate a method of describing the room acoustics, in order that the map could be used to predict how certain types of sound events might affect the soundscape. A further aim is to develop a method of easily assessing the auditory acuity of inhabitants, again so that we can overlay that onto the map. This technique could be manipulated if the participants had to use headsets, or were constantly using the telephone.

In future work we plan to apply soundscape mapping to a range of real projects to test out the generalisability and scalability of the classification method and associated tools. It will be interesting to see how the method can be used in environments where the sound events number thousands, or where a single sound masks all or most of the others, or even where all of the inhabitants are new to the environment. In all of this the overriding challenge will be to make the techniques intuitive and accessible to the widest possible range of end users.

In the more distant future, we might imagine intelligent environments with auditory components that adapt dynamically to changing circumstances (for example to co-occurring sounds, or to the mix of people who happen to be present). In the meantime there is much to be done in taking pressure off the over-loaded visual channel and exploring how sound can contribute to more useful, enjoyable and usable information artifacts.

7. REFERENCES

[1] W. W. Gaver, "What in the World do we Hear?" Ecological Psychology, vol. 5, pp. 1-29, 1993.

[2] Benyon, D. R. (2001) The New HCI? Navigation of Information Space. *Knowledge-based Systems* 14(8) 425 - 430

[3] J. Grano, "Reine Geographie," Acta Geographica, vol. 2, pp. 1-202, 1929.

[4] M. Southworth, "The Sonic Environment of Cities," in Environment and Behaviour, vol. 1, 1969, pp. 49-70.

[5] R. M. Schafer, The Tuning of the World. Toronto: McClelland and Stewart Limited, 1977.

[6] B. Truax, Acoustic Communication. Norwood: Ablex Publishing Corporation, 1984.

[7] P. Rodaway, Sensuous Geographies: Body, Sense and Place. London: Routledge, 1994.

[8] J. - F. Augoyard, "The Cricket Effect: Which tools for research on sonic urban ambiences?" presented at Stockholm, Hey Listen! Stockholm, 1998.

[9] C. Macaulay and A. Crerar, 'Observing' the Workplace Soundscape: Ethnography and Auditory Interface Design," presented at ICAD 98, International Conference on Auditory Display, University of Glasgow, Glasgow, 1998.

[10] B. C. J. Moore, An Introduction to the Psychology of Hearing, 4th ed. London: Academic Press, 1997.

[11] R. M. Schafer, "A l'ecoute de l'environnement," in Yearbook of Soundscape Studies 'Northern Soundscapes', Vol. 1, 1998, R. M. Schafer and H. Jarviluoma, Eds. Tampere: University of Tampere, Department of Folk Tradition, 1998, pp. 158-163.

[12] E. C. Cherry, "Some experiments in the recognition of speech, with one and two ears," Journal of the Acoustical Society of America, vol. 25, pp. 975-979, 1953.

[13] D. Moran, Introduction to Phenomenology. London: Routledge, 2000.

[14] B. Delage, "On sound design," presented at Stockholm, Hey Listen! Stockholm, 1998.

[15] S. Servigne, M-A Kang and R. Laurini, "GIS for Urban Soundscape: From Static Maps to Animated Cartography" presented at DM in UCE 2000, International Conference on Decision Making in Urban and Civil Engineering, Lyon, France, 2000

[16] L. J. Old, "Utilizing Spatial Information Systems for Non-Spatial-Data Analysis." Scientometrics, vol. 51, pp.563-571, 2001.

[17] .J. D. Porteous and J. F. Mastin, "Soundscape," Journal of Architectural Planning Research, vol. 2, pp. 169-186, 1985.