Audio Engineering Society

Convention Paper

Presented at the 119th Convention
2005 October 7–10   New York, New York USA

Development of Auditory Alerts for Air Traffic Control Consoles

Densil Cabrera[1], Sam Ferguson[2], and Gary Laing[3]

[1] Faculty of Architecture, University of Sydney, NSW 2006, Australia
densil@arch.usyd.edu.au

[2] Faculty of Architecture, University of Sydney, NSW 2006, Australia
samferguson@ihug.com.au

[3] Brisbane Operations Centre, Airservices Australia, Australia
gary.laing@airservicesaustralia.com

ABSTRACT

This paper documents a project that developed a hierarchical auditory alert scheme for air traffic control consoles, replacing a basic system of auditory alerts. Alerts are designed to convey the level of urgency, not provoke annoyance, be easily distinguished, minimize speech interference, and be easily localized. User evaluations indicate that the new alert scheme is highly advantageous, especially when combined with improved visual coding of alerts. The alert scheme was implemented in Australian air traffic control centers in July 2005.

1. INTRODUCTION

Air traffic control consoles convey to operators complex and important information, which needs to be understood, prioritized and responded to. Information includes the position, altitude, speed and route of aircraft, display of airways, sector boundaries, military and restricted airspaces, the presentation of automated system-detected non conformance alert and emergency events, and facilities for ground-air and ground-ground speech and text communication. Although all of the required information is conveyed through a combination of speech communication and computer graphical display, there can be considerable benefit in alerting controllers and their supervisors to important events.

The combined use of auditory and visual alerts provides redundancy – a controller not looking at the screen may miss a visual cue, but not an aural cue. Similarly, an auditory alert that is playing simultaneously with other alerts or in an environment with high ambient noise levels may be missed, whereas the visual cue will not. As a controller becomes busier, the presentation of information in one mode may become overwhelming, and the use of an alternative medium can reduce overloading and assist processing and workload handling. Sensitivity to aural and visual stimuli will vary between individuals, with some individuals being more visually perceptive, and some more aurally perceptive – and an aural and visual alerting scheme is advantageous for both sensory biases.

Nevertheless, an excess of alerts can lead to desensitization or habituation, where there is a tendency to tune out frequent repetitive sounds. Excessive alerts may induce annoyance and distraction, and there is the potential for the seriousness of an alert to be underestimated or 'lost' in the context of many alerts for routine events. Furthermore, speech communication is an essential part of the controller's work, and it is important that auditory alerts do not degrade this. In a large operations room, excessive auditory alerts from multiple consoles have the potential to create a build-up of ambient noise which could degrade the work environment.

Airservices Australia is a government agency responsible for air traffic control in 11% of the world's air space, including Australia. Airservices Australia also manages upper-level airspace (above 25,000 ft) under contract to the neighboring Pacific Island Flight Information Regions of the Solomon Islands and Nauru and lower-level airspace in the Pacific Ocean region at six airports for the United States' Federal Aviation Administration. Each year, Airservices Australia manages air traffic operations for more than three million flights carrying some 47 million passengers. The Australian air traffic management system is based on the use of an automated computer based air traffic control system, called The Australian Advanced Air Traffic System (TAAATS). An integrated traffic presentation, comprising radar symbols, flight plan symbols (for aircraft outside radar coverage) and automatic dependence surveillance symbols (aircraft automated satellite position reporting outside radar coverage) is viewed by air traffic controllers on a large computer screen, or Air Situation Display (ASD). Controllers receive auditory and visual notification of emergencies and alerts, via the ASD, in order to highlight an issue or situation requiring attention. Additional information (such as meteorological data) is conveyed by smaller screens.

In 2003, Airservices Australia resolved to improve the alert and emergency notification system through changes to visual and aural presentations to controllers, so as to reduce the likelihood of alarm desensitization and non-safety related distraction, and to increase the functionality of situational awareness alerts, including safety alerts. The key changes made were to use a color-coded scheme in the visual display, and an improved auditory alert scheme. In both cases, the schemes were designed to convey intuitively the urgency of the event represented. This paper reports on the development, evaluation and implementation of the improved auditory alert scheme.

1.1. Air traffic control operations environment

Because the work of an air traffic controller is a key component of air safety, and can be highly demanding and potentially stressful, the operations rooms of air traffic control centers are carefully designed environments [1]. The individual consoles, lay-out of multiple consoles, supervisor position, lighting conditions and levels, air quality and thermal conditions, and acoustical conditions are all considered in an effort to optimize the work environment.

The TAAATS consoles have four screens displaying a large amount of information. The main and largest screen shows aircraft symbols, route and airspace information. The smaller screens provide supporting data and communications control and information. A conventional computer keyboard and mouse are used. The controller wears a headset (one earpiece) with a microphone for verbal communication. Figure 1 shows one of these consoles.

Figure 1 A TAAATS air traffic control console, with a manikin in the controller's position.

The main visual display, which shows aircraft symbols, route and airspace information, was changed as part of this project. Previously, all alert status codes were displayed identically as yellow text, regardless of urgency. Now, the highest urgency status codes are displayed in yellow text with a red background, and the emergency codes (for aircraft-originated notification of an emergency or critical event) are displayed with red text on a yellow background. Less urgent status codes are displayed with yellow text (no background). These color codings were developed by Jenkins (Steven Jenkins and Associates Pty Ltd, Photometrics). Figure 2 illustrates the format of the information in the main visual display.

Figure 2 Example of the main visual display (with the new color scheme).

In operations rooms, TAAATS consoles are arranged to form aisles, with a supervisor's console for each aisle. As an example, Figure 3 shows the layout in a small Terminal Control Unit operations room with just one aisle – however larger operations centers have three aisles with a total of forty or more consoles.

The room acoustical conditions in air traffic control operations rooms were investigated as part of this project. The operations rooms are acoustically absorptive, with carpeted floors and a suspended ceiling of sound-absorbing tiles. The reverberation time in the Sydney operations room was less than 0.5 s in all octave bands measured (125 Hz – 8 kHz), decreasing in the high frequency range to about 0.3 s. Reverberation time in the larger control rooms of Brisbane and Melbourne was not measured, since the conditions were clearly similarly benign and the rooms are in constant use.

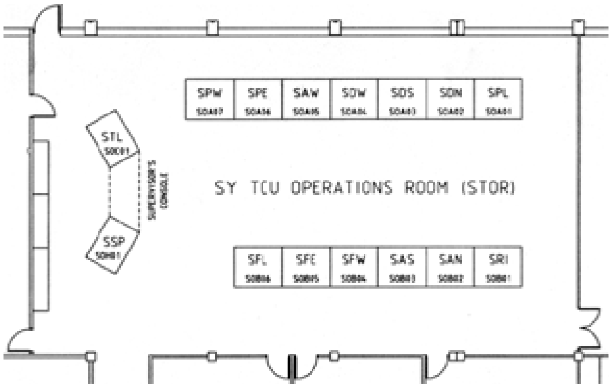

Figure 3 Configuration of the Terminal Control Unit operations room in Sydney. Each box represents a console position.

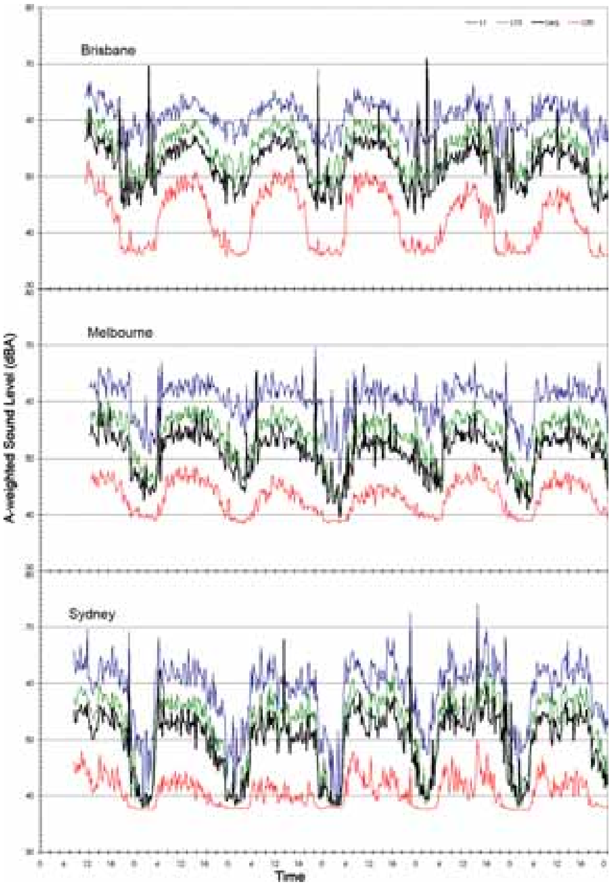

Figure 4 Sound logger measurements over six days at the operations rooms in Brisbane, Melbourne and Sydney. The logger integration time was 'fast' (125 ms), and each data-point represents a 15 minute period, showing L_1, L_{10}, L_{eq} and L_{90}.

Background noise in the control rooms in Brisbane, Melbourne and Sydney was monitored over a 1 week period. Results show a daily fluctuation in noise level in all three control rooms. L_{90} rarely exceeds 50 dBA and L_{10} rarely exceeds 60 dBA in each of the control rooms. In all control rooms, the ambient sound decreases substantially at night (with L_{90} of 40 dBA or less). Speech (coordination between controllers and communication with aircraft) is the main contributor to the background noise, with speech levels increasing in busy periods. Much of this speech is from conversations between controllers. The measurements are illustrated in Figure 4.

1.2. Principles of auditory alert design

The term 'auditory display' is used for the field concerned with conveying information to people using non-speech sound. Among other things, this interdisciplinary field includes the design of navigation aids and computer interfaces for visually impaired people, the sonification of complex data in situations where auditory analysis may be superior to visual analysis, and the design of sound in human computer interaction applications. Hence the present work can be considered to be an application of auditory display. The auditory display movement, expressed primarily through the International Community for Auditory Display (www.icad.org), is partly a recognition of the great need for good quality design of sound in everyday and specialist applications, including emerging applications made possible through increased computing power. The present literature in the field indicates that many broad principles of display design are well understood, but that much more research is required to refine these principles.

In the following subsections of the paper, the design of auditory alerts to communicate urgency, to minimize stress, to optimize segregation and grouping, to maximize localization, to minimize speech interference, and to optimize loudness are outlined. Additionally, methods for representation of concepts through sound are briefly mentioned.

1.2.1. Urgency

Patterson [2, 3] presented a procedure for the construction of nonverbal aural signals having acoustic parameters that can be manipulated to allow for various types of signal design. Patterson's alerts consist of groups of pulses, and the pulse rate, fundamental frequency and interval between groups are used to encode urgency. This work was applied to aircraft flight deck alerts.

Edworthy *et al.* [4] and Hellier *et al.* [5] examined how urgency can be encoded through sound signal parameters: a faster modulation rate, higher pitch or less regular harmonics conveying a greater sense of urgency. They developed models of perceived urgency, based Stevens' power law, with exponents for pulse speed, number of repeating units, fundamental frequency and inharmonicity. Based on these data, Hellier *et al.* presented guidelines suggesting how to manipulate the acoustic parameters of an aural alerting signal in order to produce a 50% increase in perceived urgency level, a doubling of perceived urgency level, and a tripling of perceived urgency level.

For example, an alerting signal's tempo, or speed, may be manipulated in the following ways to achieve changes in the alert's level of perceived urgency:

- To increase the alert's level of perceived urgency by 50%, the alert's pulse rate (frequency of pulses) could be increased by a factor of 1.3;
- To double the alert's level of perceived urgency, the pulse rate could be increased by a factor of 1.6;
- To triple the alert's level of perceived urgency, the pulse rate could be increased by a factor of 2.2.

Loudness also conveys urgency [5] but is of more limited use in hierarchical alert scheme design because loudness discrimination is relatively imprecise, and alerts generally need to be at an optimum loudness (reliably audible above the background noise masking threshold, but not so loud as to produce distraction and annoyance effects) [2].

Operator response time tends to decrease with increasing perceived alert urgency [6]. However, Ahlstrom and Longo [7] advocate limiting the number of urgency levels to four. A large number of alert levels can create a confusing auditory display [8].

Acoustical parameters of the alert signal do not fully determine the perceived alert urgency. The context of the alert presentation also can have a substantial effect [9, 10].

1.2.2. Stress

As noted by Guillaume *et al.* [9, 10], auditory alerts need to be designed for their context. A home burglar alarm may convey a high sense of urgency, but also is designed to create an unpleasant and stressful acoustic environment for the burglar – and as such would be entirely inappropriate for an air traffic control auditory display. Indeed, minimizing the stress and annoyance associated with the auditory alert scheme was a major consideration of the present project. However, the notions of urgency and stress are linked, so it is a delicate problem to design alert signals that convey high urgency whilst minimizing stress.

The concept of sensory unpleasantness seems likely to be related to stress. Aures [11] and Zwicker and Fastl [12] identify general contributors to sensory unpleasantness, namely loudness, sharpness, roughness and pitch strength (tonalness). However, the main envisaged application of their work is in machinery and appliance noise, and so their model may only be broadly indicative for other applications. Sharpness has the strongest effect on sensory pleasantness, with high sharpness values corresponding to unpleasant sound. The loudness of sounds does not influence unpleasantness directly for values less than 20 sones – above this value increasing loudness degrades pleasantness (note that a 1 kHz tone presented at a free field sound pressure level of 80 dB has a loudness somewhat less than 20 sones). Nevertheless, sharpness and roughness models are influenced a little by loudness, so there is some loudness effect for lower sound pressure levels. Roughness reduces pleasantness, whereas tonalness adds to pleasantness. However, the role of tonalness in sounds that are inherently annoying is reversed, with strong pitches contributing significantly to annoyance [13]. This is likely to be due to the tones making the sound more noticeable, and therefore harder to ignore. Pitch height is not included directly in their model (although it can be related to loudness, sharpness and pitch strength).

Sharpness is generally modeled as a weighted psychoacoustic centroid [12], and could be thought of as an adaptation of spectral centroid in psychoacoustics. The most important implication of such models is that strong high frequency spectral content contributes to sharpness, and hence to sensory unpleasantness. Roughness is generally modeled as rapid fluctuations in the stimulation of auditory filters [12], which could occur through amplitude modulation, frequency modulation or beats. Sirens are often characterized by their sharpness, sometimes roughness and usually loudness and fluctuation strength.

High sound pressure levels are associated with stress – with a multitude of studies showing both acute and chronic physiological effects [15, 16],

The temporal envelope of an auditory alert can induce stress. Patterson [2] identifies a startle effect, which can occur when an auditory alert takes an operator by surprise. This has physical effects on the operator, and is likely to reduce rather than enhance performance. Hence Patterson advocates the use of gradual attacks in alert temporal envelopes.

Neuhoff and Heckel [17] examine a potential problem with gradual attacks, referred to as auditory looming. Many sounds in our environment have a sudden attack and gradual decay (this is the envelope of a damped resonance excited by an impulse, and is found not only in many sound sources, but also in the reverberation of the environments in which they are heard). When this typical envelope is reversed, the result can seem unusual and have a greater command over the listener's attention, and even alarm the listener. Neuhoff and Heckel relate this reaction to auditory distance perception – the reverse temporal envelope is associated with a rapidly approaching object, provoking an instinctive defensive reaction in the listener.

1.2.3. Segregation and grouping

The field of auditory scene analysis is concerned with how people analyze the sound environment into distinct auditory events, sound sources, groups of related sounds, and temporal streams of related sounds. Conversely, it is also concerned with how people discriminate between such groups, referred to as segregation. Bregman [18] provides the seminal work on auditory scene analysis, in which principles of auditory grouping and segregation are overviewed.

A primary characteristic of auditory alerts is that they be segregated from the rest of the auditory scene. In particular, they should be distinctively different from background noise, and from speech. Secondly, the auditory alerts should be perceived as a group – like a single perceptual channel of information. Thirdly, the various auditory alerts used should be segregated within the group. Finally, each auditory alert should be perceived as a single stream since it represents one concept.

The first and second criteria can be solved together, by selecting a type of sound for the alerts that is dissimilar to the background sound environment (including speech). Alert schemes often use proto-musical sounds, such as tones, tone combinations or modulated tones. Noise-like signals are less distinguishable from typical background sound environments. A consistent theme or style for the various alerts should help to satisfy the second criterion.

In order to satisfy the third criterion, there must be substantial differences between the alerts. When two alerts are heard simultaneously they should not fuse into one sound. This can be facilitated by avoiding simple integer relationships for frequencies – both fundamental frequencies of harmonic tones (pitch), and frequencies of repetition (rhythm). Ability to discriminate between sound parameter values (such as between pitches) varies between individuals [19], so there is a redundancy advantage in differentiating the various alerts using multiple signal parameters. Subtle differences between alerts should be avoided if they are to be reliably differentiated.

The final criterion is easy to satisfy through alerts that do not have excessive internal variety or disparate features. Nevertheless, perverse instances are easy to imagine.

Burt [20] investigated the design of auditory alerts for aircraft flight decks. She found that four distinctive alerting sets, each conveying three levels of urgency, could be designed for this application. Proto-musical features, such as rhythm and melody, could be used to define these sets.

1.2.4. Localization

There can be an advantage in being able to localize the source of the alert, especially in environments with multiple work-stations – in which an operator needs to know from which work-station an alert originates. In single operator environments, localization can be used to convey spatial information associated with an alert (e.g. an alert in a car cabin could come from the left or right, depending on the side of the car that was in the vicinity of danger).

Robust sound localization requires rich spectral content extending to high frequencies [21]. While this may seem at odds with the need for low sharpness alerts (in order to minimize stress), it is possible to provide the high frequency energy in a brief impulse, like a ‘click’. This approach was taken in the present project.

A more detailed discussion of localization is in section 1.4.

1.2.5. Speech interference

Speech intelligibility relies primarily on consonant discrimination, making the octave band centered on 2 kHz the most important for intelligibility. This is reflected in the weighting functions for speech intelligibility predictors such as Speech Intelligibility Index [22] or Speech Transmission Index [23], and indeed the older Articulation Loss of Consonants model was concerned only with the 2 kHz octave band [24]. The implication is that there should be a substantial advantage in having auditory alerts that do not unduly raise the masked threshold in this frequency range.

1.2.6. Audibility and loudness

For an auditory display to be effective, it must be audible. However, the limits of human hearing, as determined in ideal conditions and represented as the well-known equal loudness contours, do not give a practical indication of the range of sound signals useful for auditory display, especially when considering alert design. In practical situations, audibility is limited by masking, which is established by background noise and speech.

Human factors design criteria for United States air traffic control specify a signal to noise ratio of at least 10 dBA [7]. Nevertheless, a controller should be able to adjust the audio signal level depending on the circumstances.

The upper limit for alert loudness is certainly less than the threshold of pain, and certainly less than a level that could cause a temporary threshold shift or hearing loss after long term exposure. The upper limit is constrained further by the need to minimize stress and speech interference.

In the present project, the upper limit for alert loudness was also constrained by the audio hardware configuration used.

A design question is whether it is absolutely essential that an alert be heard in all instances, or whether visual information can substitute for an auditory alert in an exceptionally adverse background noise environment. The cost of highly robust audibility may be speech interference and annoyance. For a prioritized alert scheme, it may be less important for a low priority alert to be always heard than a high priority alert. Ahlstrom and Panjwani [25] note that is possible that less critical alarms in the airway facilities environment may be turned down with little or no operational impact.

1.2.7. Associations

Due to a lifetime of experience with sound around us and in our culture, there are strong associations between certain sound parameters and concepts. The 'looming' sound effect already mentioned seems to be an example of this – where distance perception and loudness perception are associated. Another example is that size can be represented with loudness [26]. One intriguing example studied by the authors is the pitch scale (ranging from 'low' to 'high'), which can be used effectively to represent elevation (i.e. the physical scale from low to high) [27]. In these examples the association is so deeply rooted that it may not be separable from perception.

In speech, meaning is conveyed not just through the words, but also through prosody – that is, the modulation of speech rate and pitch. Opposite meanings can be created through different intonations – for example a rising 'yes' can indicate questioning, whereas a falling 'yes' can indicate satisfaction. Rapid speech indicates emotional arousal (eg excitement), and the loudness and fundamental frequencies of speech also tend to rise with excitement [28]. Similar tendencies exist in music [29]. The aforementioned urgency encoding of alerts appear to be closely related to these prosodic and musical features.

Guillaume *et al.* [9, 10] show that there can be a substantial learning effect for auditory alerts when the meaning of the alert is understood by operators. Hence the urgency encoding of alerts, in terms of their acoustical parameters, does not fully determine the urgency conveyed to an operator. Accord between the alert acoustical parameters and their associated meaning is desirable, so that cognitive and perceptual processes work together.

1.2.8. Earcons and auditory icons

This somewhat quirky pair of terms has emerged in the auditory display field to denote two key forms of representation using sound.

The term 'auditory icon' is used for recognizable sounds of things or events, used to represent a closely related concept. For example, a recording of car tires on a rumble strip could be used as an alert sound for car driver navigation (to alert the driver of deviation from a lane) [30]. Keller and Stevens [31] identify various representation systems that auditory icons can be based on.

Earcons are sounds designed to represent concepts, without using sounds otherwise associated with the concept. For example, a tone or series of tones might be used as an alert for some event. The sounds used in the present project mainly fall into the earcon category, in part, because a small number of sounds is used to represent a large number of concepts.

Of course, speech signals are also possible to use as auditory alerts. However, in the air traffic control environment, speech is already used for communication, so speech alerts would interfere with this activity. Speech alerts can be useful in contexts where there are so many types of alerts, or alerts are so rare, that the meaning of abstract sounds could be forgotten.

1.3. Comparable auditory alert schemes

An auditory alert scheme is used for air traffic control in the United States. The Standard Terminal Automation Replacement System (STARS) uses six auditory signals, based on tones with fundamental frequencies in the range from 800 to 1600 Hz [32]. In one case, the tone warbles (alternates between the two frequencies of 1600 and 2000 Hz). Tone frequency and duration provide redundant encoding of urgency. Terminal air traffic control auditory alerts used in the United States were reviewed by Ahlstrom [33].

Various auditory alerting systems for aircraft flight deck environments have been proposed or are in use. Patterson [2, 3], Burt [20] and Guillaume *et al.* [9, 10] are among those who have taken systematic approaches to design and evaluation of these systems.

Brock *et al.* [34] developed a hierarchical auditory alert system for a 'legacy' navy decision support system console. The sound source for this console was simply a buzzer, which could be switched on or off. The urgency of the alert was encoded by the rhythm and speed of the buzzer pulses, using the principles of Edworthy *et al.* [4] and Hellier *et al.* [5]. Unfortunately, the unpleasant sound of the buzzer itself limited the user appeal of this auditory alert system.

While there exist no standards for auditory alert signals for air traffic control systems, standards exist for auditory danger signals, for example, for work environments with industrial machinery [35, 36]. The signal parameters specified in these standards appear to be incompatible with the low stress work environment required for air traffic control systems.

1.4. Pre-existing auditory alert scheme

Prior to this project, the consoles were supported by a two level auditory alerting system, consisting of either a long or short 1.7 kHz pure tone. This auditory alert system was far from optimal, with controllers reporting that it conveyed little or no information about alert priority, and was annoying and difficult to localize. Some controllers were surprised to hear that there were two alert signals (long and short) when the auditory alerts were discussed in the briefing sessions that were part of the present project. The problem with localization, reported by controllers in Sydney, was that a controller could hear an alert, but mistakenly believe that it came from another controller's console.

Pure tones can be very difficult to localize because they contain very little information for the hearing system to interpret [21]. Binaural difference cues can be helpful for left-right discrimination of pure tones, and the localization process differs between low and high frequency ranges. In the low frequency range (below about 700 Hz), inter-aural time differences (expressed as phase differences) provide effective left-right discrimination. At higher frequencies, ambiguities occur, limiting the usefulness of inter-aural time differences, although head movements can help to resolve such ambiguities. At high frequencies, especially above 2 kHz, inter-aural level differences become large enough to be used for left-right discrimination. The pure tone frequency of the pre-existing auditory alert is in the frequency range between these two localization cues, where left-right discrimination is at its poorest. With the absence of head movements, pure tones, at any frequency, provide no information for localization around cones of confusion (eg distinguishing front from back), because high frequency spectral cues are used for this – requiring rich spectral information. Instead, narrow band signals are associated with localization illusions [21, 27]. Furthermore, in the presence of reflections and standing waves, such as occur in normal room environments, pure tone localization for any frequency can be no better than chance [37].

The pre-existing system also was limited by having the loudspeaker beneath the console desk top. A small hole in the desk top, used for the keyboard cable, provides the only possible direct sound route between the loudspeaker and operator's ears. As can be seen in Figure 1, the keyboard can completely obscure this hole – depending on the controller's preference for keyboard position. Relocating the loudspeaker was not within the scope of the present project. However, experiments in a simulator showed a substantial benefit of a line of sight loudspeaker position in terms of direct to indirect sound energy ratio.

The fact that controllers wear an earpiece on one ear is likely to degrade their localization ability. Nevertheless, some localization should be possible with appropriate sound design. The occlusion of one ear increases the need for rich spectral information in the auditory alerts.

2. NEW AUDITORY ALERT SCHEME

2.1. Design Goals

The TAAATS system currently has 16 non-conformance events or emergency notifications that require aural alert presentation to controllers – but a much smaller number of auditory alerts was desired in order to avoid confusion. Therefore a categorization of the events and emergencies was determined, with four levels of priority. One auditory alert was to be used for each level of priority. From highest to lowest priority, the four prioritized categories (P1-P4) are:

- P1 – System-detected Safety Net critical event (short term collision avoidance alert, danger area infringement warning, and minimum safe altitude warning);

- P2 – aircraft-originated notification of emergency or critical event (for example, hijack notification or aircraft radio failure);

- P3 – system-detected non conformance event, high priority (for example, divergence between predicted and actual aircraft route); and

- P4 – system-detected non conformance event, low priority (for example, overdue position report).

The aims of the auditory alert scheme were to convey intuitively the urgency of the event, be easily distinguished, not induce stress or annoyance, and be easily localizable. A degree of design integrity was also taken as a goal – meaning that the alerts should be in a consistent style. This is also desirable in terms of facilitating appropriate auditory scene analysis. Based on initial discussions the alerts were to be abstract sounds, rather than identifiable as the sounds of particular sources. We initially aimed to have the alerts at a constant loudness, but this aim was abandoned following simulator trials and measurements.

The possibility of sounding the alerts via the earpiece (in addition to the loudspeaker) was considered, but it was considered preferable to restrict the audio in the earpiece to speech communication only.

2.2. Design process

2.2.1. Initial candidates

Initially, we developed candidate alarms for P1-P4 that could be classified thematically as beeps, tone sweeps, percussion, bell-like sounds, and sounds similar to water impact (eg drip sound). A variety of temporal and spectral envelopes for these candidates was prepared. Urgency was encoded in these prototypes using techniques including:

- Number of pulses in a group;

- Pulse repetition rate;

- Group repetition rate;

- Tone fundamental frequency; and

- Spectral balance (sharpness) of the sound.

An initial evaluation was conducted with an operational specialist from Airservices Australia. 'Beep'-like sounds were deemed inappropriate in that meeting, because they seemed likely to be annoying. Bell-like sounds appeared to have potential. Following that meeting, more prototypes were developed.

Refined candidates were evaluated and further refined in a meeting with several operational specialists from Airservices Australia, and the draft final candidates for P1-P4 were selected from these. As part of this meeting (held in a 5.1 format surround sound production studio), a spatial simulation of an air traffic control operations room was demonstrated, using various prototype auditory alerts together with ambient noise. In this meeting, sharpness was excluded as a parameter for urgency, because it appeared to induce annoyance, and seemed likely to lead to stress.

Rather than describing the prototype alerts in any detail, only the final alerts are described in this paper (Section 2.3).

2.2.2. First simulator trial

Approach

The first simulator trial was held over two days using five operational air traffic controllers. Three combinations of alerts were evaluated:

- New visual alerts with old auditory alerts;

- Old visual alerts with new auditory alerts; and

- New visual alerts with new auditory alerts.

The trial had nine 60 minute sessions – i.e. three each of the above combinations. Exercises were run as normal training exercises, with direction to the simulator pilots responsible for the generation of safety net critical events and aircraft emergencies. A spreadsheet detailing the time for each event was given to the simulator pilots to ensure that these were spread evenly initially. For later runs, different events were generated simultaneously, and from different consoles, in order to evaluate the effect of multiple auditory alerts in the room. P1 and P2 alerts were limited to a maximum of approximately 3 each per controller per run, as large numbers would be unrealistic, and could prejudice the findings with regard to the effect of the auditory alerts on the room environment.

A briefing for the participants was held on the day prior to the trial. This provided the controllers with background and information regarding the alerts changes and proposed end state operational model. The auditory alerts were demonstrated to participants in this briefing to give the controllers awareness of, and familiarity with, the new alert sounds prior to simulation.

Evaluation of baseline condition

Controller evaluations of the baseline condition (the old visual and auditory scheme, which the controllers used in their everyday work) were collected using questionnaires, in which each question was rated using an integer scale from -3 to 3, but also with scope for written comments. The scale values were defined as 'unacceptable', 'poor', 'marginal', 'neutral', 'fair', 'good' and 'excellent'. The first six questions were concerned with the visual display, and the final three questions with the auditory display. Since this paper is on the auditory alerting system, it concentrates on these final questions. Questions 7 to 9 are paraphrased and annotated (to minimize or clarify jargon) as:

7. How effective are the aural alarms in providing situational awareness and a hierarchy of importance between high and low priority alerts?

8. How effective are aural alarms in providing the ability to differentiate between emergencies (i.e. P2) and alerts (i.e. P1, P3 and P4)?

9. Comment on the ability to localize the aural alarm. That is, the ability to determine which console the sound is originating from.

Note that the terms 'aural alarm' or 'audible alarm' are used in internal documentation, whereas this paper adopts the term 'auditory alert' for the same concept.

The evaluation of the baseline condition showed that the visual display was more positively evaluated than the auditory display. In fact, ratings of the auditory display were universally negative (mean score of -2.4), whereas the visual display ratings tended to be neutral (mean score of -0.2).

Ratings for question 7–9 were identical in distribution, with two ratings of -3 (unacceptable) and three ratings of -2 (poor). With respect to Question 7, two controllers commented: "Same tone, get complacent about it" and "Very little difference in alarms, does little to enhance situational awareness". Three controllers commented on Question 8: "Little difference, alarms all sound the same", "All same – alert to look for visual cue" and "There is no differentiation". Three controllers commented on Question 9 (on sound localization): "Everyone in the aisle looks around to see whose console is beeping", "Very bad. Often can't tell unless you can see the visual indication as well" and "Absolutely impossible to tell".

Evaluation of new auditory alerts (old visual alerts)

The questions used following the three trial runs with new auditory alerts, but the old visual alerts are paraphrased as:

1. Rate the overall effectiveness of aural alarm in alerting controllers to the urgency or criticality of a triggering event – i.e. can you tell if that event is high or low priority?

2. Rate the effectiveness of aural alarms in providing awareness or hierarchy of importance between high (i.e. P1) and low (i.e. P4) priority alerts.

3. How effective are aural alarms in providing the ability to differentiate between emergencies (i.e. P2) and alerts (i.e. P1, P3 and P4)?

4. Rate the ability to localize, or determine from which console an aural alarm is originating.

5. Rate the ability to distinguish between the four different aural alarms.

6. Give an overall comparison of new aural alarms with the pre-existing scheme.

7. What is the level of annoyance or distraction caused by the new aural alarms?

For Question 5, the same rating scheme (from 'unacceptable' to 'excellent') was used. For the other questions the seven-value bipolar rating scheme was relative to the baseline condition: 'substantially worse', 'worse', 'slightly worse', 'no change', 'slightly improved', 'improved' and 'substantially improved'.

Results are shown in Figure 5, indicating a generally positive response to the new auditory alerts. However, results for Question 7 (annoyance and distraction) are equivocal.

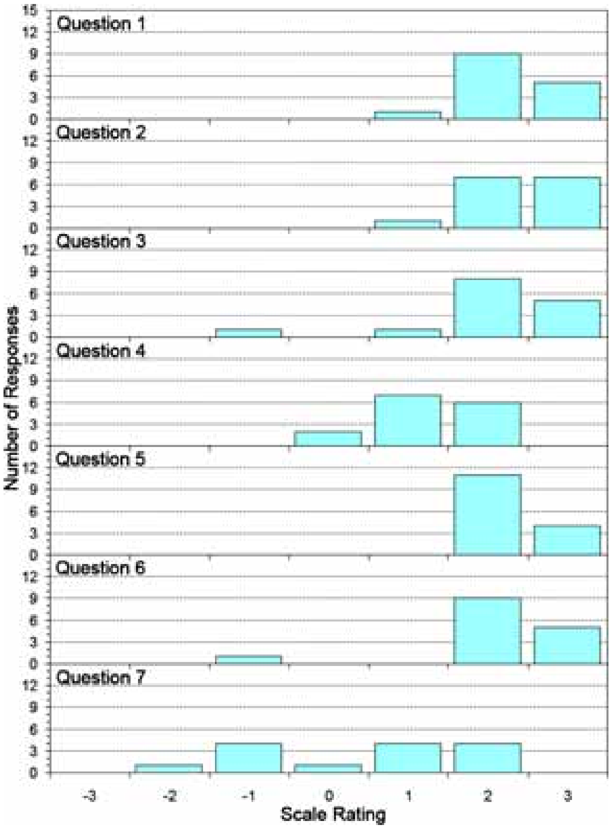

Figure 5 Rating distributions for the seven questions used to evaluate the new prototype auditory alerts, with the pre-existing visual alert scheme. With five participants and three trials, the maximum possible result is 15.

In the development of the alerts, an attempt had been made to factor in the frequency response of the console loudspeaker, including the fact that it was located beneath the desktop. However, even causal listening revealed that this attempt had failed. The result was that the P1 prototype was much louder than P2 and P3, and that P4 was quieter still. Comments by the participants were that P1 was too intrusive when the console audio gain was appropriate for P2 and P3, and that P4 was too quiet. The response to Question 7 was primarily related to this problem.

Mean rating scale values increased for every question from the first to the third run. This is partly due to the fact that the software volume control was changed between runs, to attempt to arrive at the best volume setting for the auditory alerts.

Comments from the participants are too numerous to report directly, but were generally positive. Much comment was made on the need to refine the playback levels of the auditory alerts. Comments on Question 4 (localization) identified the playback levels of the low priority alerts as an impediment to localization. Subsequently, further efforts were made to improve localization cues in the process described in Section 2.2.4.

While the four priority auditory alerts were included in this trial, additional information and feedback alerts still used the old 1.7 kHz pure tone. This was clearly perceived as a problem.

Evaluation of new auditory alerts (new visual alerts)

The questions used for the three simulator sessions with new auditory and visual alerts combined were:

1. Rate the overall effectiveness of the presentation in alerting controllers to the urgency or criticality of a triggering event – i.e. can you tell if an alarm is high or low priority?

2. Rate the effectiveness of the visual display to identify the aircraft on screen that is subject to the alert or emergency event.

3. Give an overall comparison of new presentation with pre-existing display.

4. How well do the aural and visual alerts complement one another?

5. How effective is the new hierarchy (P1-P4) at providing information about the urgency of an event?

6. Are the combined changes more effective than the new aural or visual changes alone?

7. To what extent did the presentation of aural and visual information overload or overwhelm?

Questions 1, 4 and 6 used a seven value bipolar scale from ‘substantially worse’ to ‘substantially improved’. Questions 2, 3 and 5 used a scale from ‘unacceptable’ to ‘excellent’. Question 7 used the scale (from -3 to 3): ‘substantially more’, ‘more’, ‘slightly’, ‘no change’, ‘slightly less’, ‘less’ and ‘substantially less’. The distributions of scale ratings are given in Figure 6.

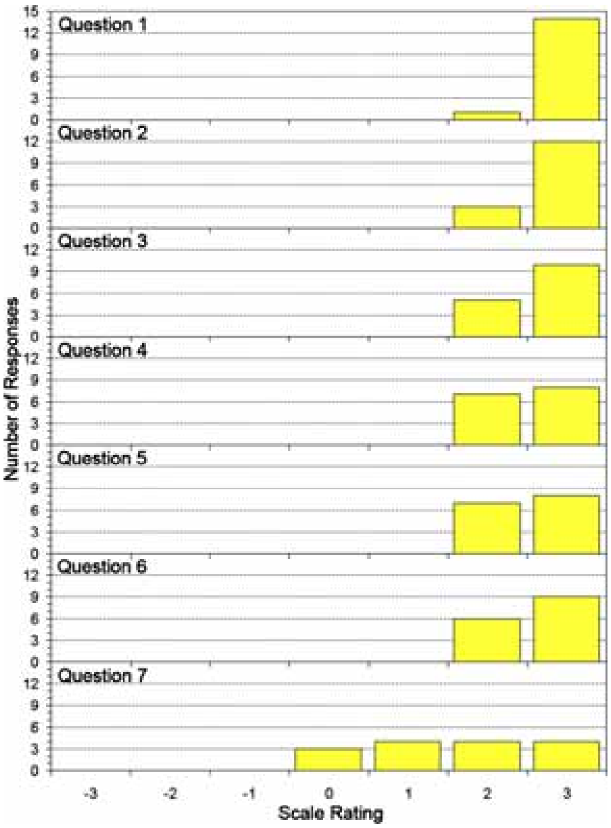

Figure 6 Rating distributions for the seven questions used to evaluate the new prototype auditory alerts with the new visual alert scheme. With five participants and three trials, the maximum possible result is 15.

As can be seen from the rating distributions, the results were overwhelmingly positive, but with some reservations with regard to Question 7.

Comments by participants were more positive about the new auditory alerts in these three runs, compared to the comments in the runs using the pre-existing visual alert scheme. The improvement offered by the new auditory and visual presentation was clear at the end of the two-day trial, when controllers commented that they did not wish to return to the pre-existing alert scheme.

Post-trial evaluation

A post-trial evaluation was held the following day, in which trial participants used their recollections of the two-day trial to rate the pre-existing display, the new visual alerts with pre-existing auditory alerts, the new auditory alerts with the pre-existing visual alerts, and the new visual and auditory alerts in combination. Only four of the five trial participants were available for this evaluation. Only the questions relating to the auditory alerts are presented here. Ratings for these were made on a seven value bipolar scale from ‘unacceptable’ to ‘excellent’.

The questions relating to auditory alerts, together with the results, are given below and in Figures 7 and 8:

1. Rate the overall effectiveness of presentation of alerts in notifying controllers of the urgency or criticality of a triggering event – i.e. can you tell if an alert is high or low priority?

Responses to this question gave a maximally positive evaluation of the new auditory and visual alerts in combination, as shown in Figure 7. Both the new auditory and new visual alerts, on their own, were also rated more positively than the pre-existing alert scheme.

2. Rate the overall effectiveness of the presentation of emergencies in providing situational awareness to controllers.

As illustrated by Figure 7, responses to this question had a similar tendency to those for Question 1.

3. Rate the ability to differentiate between different types of Alerts and Emergencies.

Responses showed that both the new visual alerts alone and the new auditory alerts alone gave a substantial improvement in distinguishing the types of events represented. The rating for the combination of new visual and new auditory alerts was maximally positive.

4. How well do the proposed aural and visual alerts complement one another?

This question received three ratings of ‘excellent’ (scale value of 3) and one rating of ‘good’ (2).

5. Rate the effectiveness of the presentations in providing situational awareness to controllers.

Again, the new visual or auditory alert schemes showed improvements on their own, and a maximally positive result when used in combination (Figure 7).

6. Rate the overall effectiveness of the new aural alerts for each priority level.

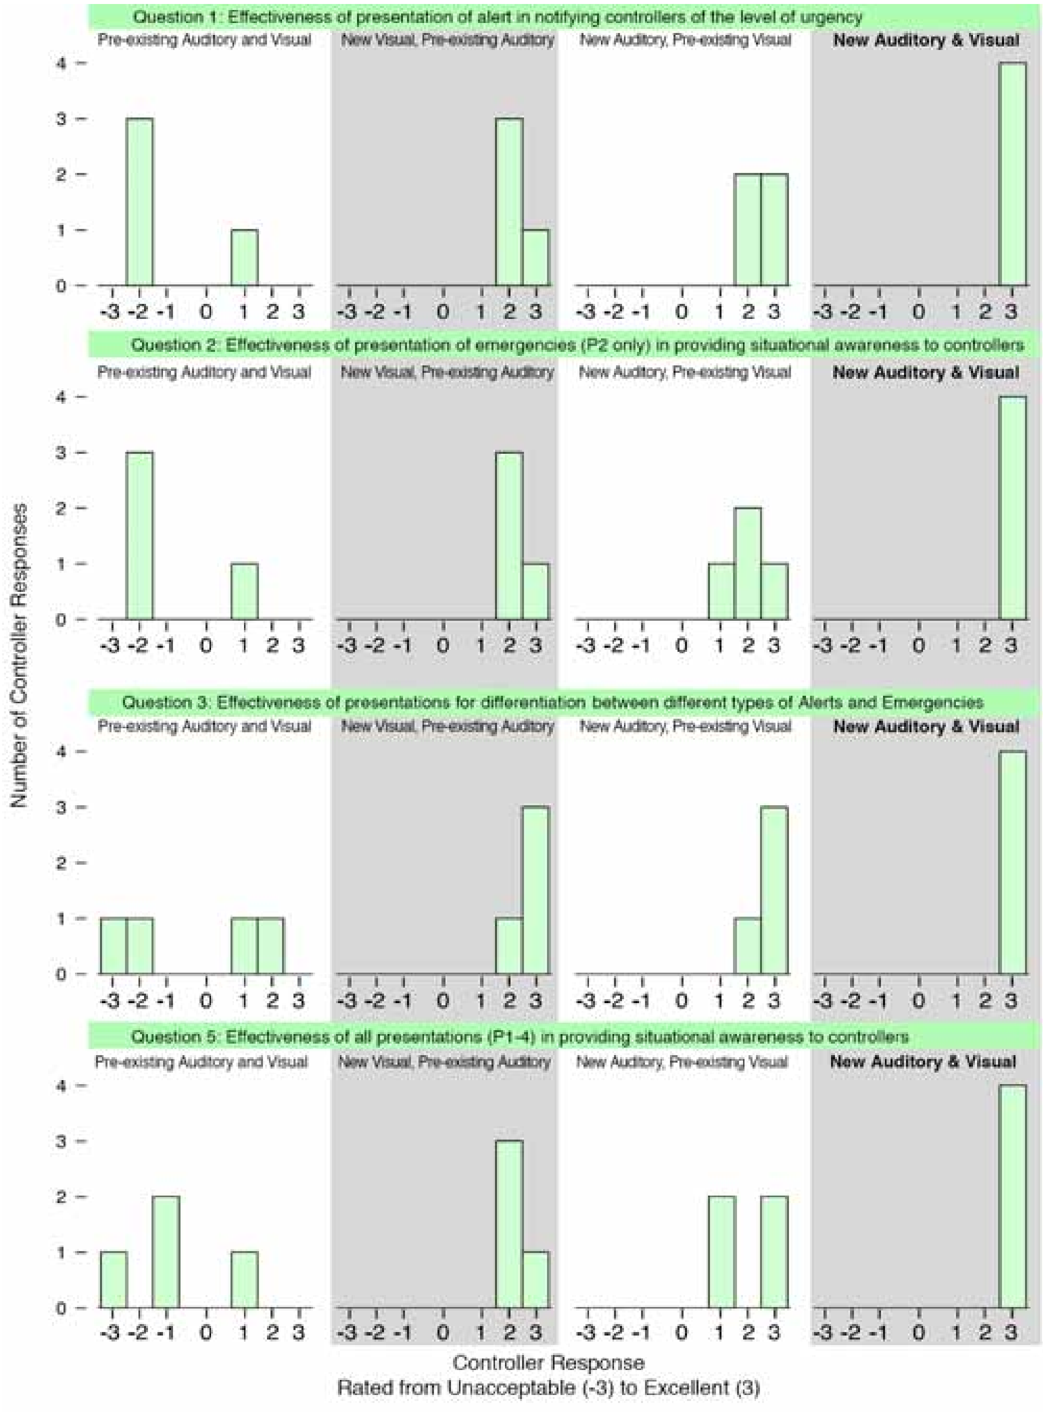

Figure 7 Rating distributions for Questions 1, 2, 3 and 5 of the post-trial evaluation. With four participants, the maximum possible result is 4.

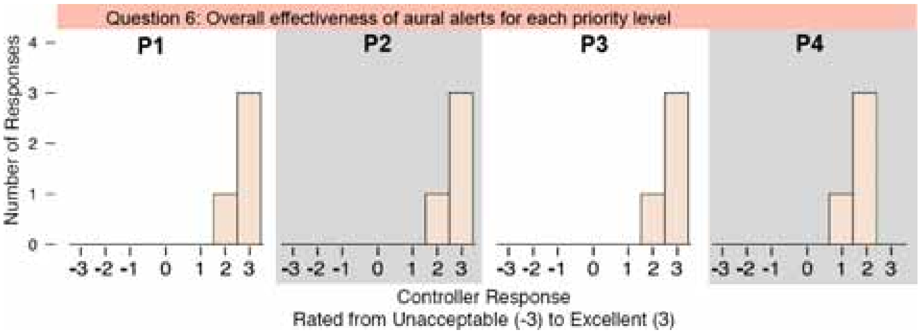

Figure 8 Rating distributions for Question 6 of the post-trial evaluation. With four participants, the maximum possible result is 4.

Responses to this question show good to excellent evaluations of P1-P3, but a fair to good evaluation of P4. The general opinion expressed was that P4 was too quiet.

The results in the trial as a whole show positive evaluations, but some problems were made clear: namely that the P4 auditory alert needed to be louder, the P1 alert loudness needed to be reduced for the software volume control setting of 3, and additional alerts were required to substitute for the default 1.7 kHz tone which was still playing for information alerts and feedback signals.

2.2.3. Information and feedback alerts

As mentioned above, the new aural alarms highlighted the incompatibility of the new alarms with the existing information and feedback alerts, which occur to notify the controller for timer activation and receipt of information such as meteorological data. The existing aural alerts for these (1.7 kHz pure tone) inappropriately conveyed a higher sense of urgency than the new Priority alerts. Additional information and feedback auditory alerts were therefore developed to complement the new Priority alerts.

The information and feedback events include:

- Timer alert – the timer is set as a memory aid or prompt when the controller has a time-related task, for example to remind the controller to do coordination or that restricted military airspace is becoming active.

- Aeronautical Information Function (AIF) – this is an alert of information receipt such as new meteorological data.

- Controller-Pilot Data Link Communications (CPDLC) alert – CPDLC is a text-based messaging system consisting of standard messages (e.g. ‘Request climb to FL350’) sent by satellite to enable direct communications between pilots and controllers. It is used where aircraft are outside normal radio coverage (e.g. in the middle of the Pacific Ocean). Alerts are used to notify the controller of newly received text communication from an aircraft.

- Keyboard error – this alerts the controller to an invalid keyboard entry.

The alerts designed for these were intended to be quite distinct in style from the four Priority alerts (P1-P4). Rather than conveying urgency, the alerts were intended to convey something of the meaning of the event represented.

The new timer alert is simply a short ‘click’, reminiscent of the ticking of a clock. It was to be repeated once per second until acknowledged.

The new AIF alert consists of a long and short tone with a rising pitch interval of a major third. The first tone is 850 Hz and the second is 1070 Hz. The tone timbre is close to pure, in contrast to the priority alert tones – except that broad spectral content is briefly present at the start of each tone. The sound might be characterized as a damped chime.

The CPDLC message notification alert consists of a short and long tone, with a falling pitch interval of a major third. The first tone is 1430 Hz and the second is 1135 Hz. The descending tone pattern was chosen in analogy to a message descending from an aircraft to ground. The tone timbre is the same as that of the AIF alert.

The keyboard error alert is a short duration descending sweep of a harmonic tone. This sound seems to represent the notion of an error well.

2.2.4. Refinements using a simulator

Prior to the first simulator trial, laboratory measurements of the loudspeaker had been made in an anechoic room, as well as for the loudspeaker under a desk (time-windowed to exclude late reflections and reverberant energy). The loudspeaker has a 55 mm diameter driver, enclosed in a box with external dimensions of 80 x 62 x 54 mm. As shown by Figure 9, positioning the loudspeaker under a desk reduces its transfer function to the operator's position in the high frequency range. Perhaps a positive effect of this is that the loudspeaker's frequency range could be thought of as extending to a lower frequency when positioned under the desk.

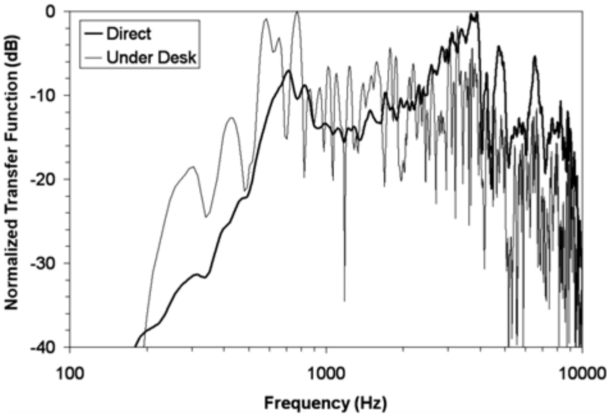

Figure 9 Indicative normalized frequency responses of the loudspeaker, based on impulses measured in the free field (direct), and with the loudspeaker positioned under a desk.

Equalization compensation of the P1-P4 alerts was done prior to the first simulator trial based on this transfer function measurement. Testing using a sample loudspeaker (but different digital to analog converter and amplification system) appeared to give reasonable results. However, as mentioned before, the relative playback levels in the first simulator trial clearly required adjustment. This failure to compensate for loudspeaker frequency response, which was evident in the simulator trial, may have been due partly to the methods used. Nonetheless, it became increasingly apparent that the entire audio system (i.e. the console, including the computer, sound card and loudspeaker in its normal position) needed to be taken into account, and that this system was noticeably non-linear (and therefore not amenable to compensation by fixed equalization). Therefore several hours were spent using a simulator console, with an iterative process of playing, measuring and refining the sound files.

The main purpose of this simulator measurement and refinement was to set the sound pressure level of the auditory alerts at an operator's head position for a default software volume control setting of 3 (on a stepped scale from 1 to 5). An attempt was made to quantify the gain steps of the software volume control, but it varied between sound files, reflecting the system's non-linearity. In very approximate terms, the five gain settings from 1 to 5 yielded -10 dB, -5 dB, 0 dB, +3 dB and +5 dB relative to a setting of 3.

Although there are many possibilities for measuring loudness, a simple approach was chosen for this project. The maximum A-weighted sound level, using 'fast' integration time (125 ms) was taken as a loudness indicator. This was measured with a sound level meter (Brüel & Kjær 2250) at the operator head position. Boxes blocked the under-desk path as a controller's legs might be expected to. If anything, this arrangement will have underestimated the sound pressure level, since the blockage offered by the boxes was more complete than that offered by a real person.

Changes to the sound files involved much more than simple gain changes, because in many cases the signal was already at or near maximum amplitude in the 16-bit digital file format. Changes included detailed modifications to the attack, more general changes to the envelope, filtering and dynamics processing. The information and feedback alerts were included in this measurement and adjustment process.

The final sound pressure levels determined from this process are given in Section 2.3.

2.2.5. Second simulator trial

The second simulator trial was conducted with the modified auditory alerts, two months after the initial trial. Four of the original five controllers participated in this trial. Due to software limitations, the CPDLC and Keyboard Error alerts could not be implemented – however the AIF and Timer alerts were implemented. Two one-hour simulation runs were conducted, and assessments were made at the end of the trial.

The questions asked, and results obtained, are given below.

1. Rate the overall effectiveness of the presentation in alerting controllers to the urgency or criticality of a triggering event – i.e. can you tell if an alarm is high or low priority, compared to current?

All participants rated this as ‘substantially improved’ using a 7-value bipolar scale. The comments by participants were: “Good aural differentiation between alarms, easy to tell which alarm is which”, “Easily discern urgency of alert”, “Very good”, and “Excellent, particularly if you’re not looking at the screen. The aural alert provides a degree of urgency”.

2. Rate the effectiveness of the alert or emergency presentation to identify the aircraft on screen that is subject to the alert or emergency event.

This question emphasizes the visual alerts, rather than the auditory alerts. Three participants rated this as ‘excellent’, and one rated it as ‘good’. The comments by participants were: “Color draws attention quickly”, “Improved presentation over old alerts”, “Very clear, easy to identify”, and “Particularly the STCA (Short Term Collision Avoidance Alert)”.

3. Give an overall comparison of new presentation with current display.

All four participants rated this as ‘substantially improved’. The participant comments were: “Much better than before”, “Improved visual and aural alerts lead to substantial improvement overall”, and “Easier to read. Draws attention to priority alerts”.

4. How well do the aural and visual alerts complement one another?

Three participants rated this as ‘excellent’, and one as ‘good’. The participant comments were: “Don’t tend to use visual as much as aural, as you can tell which alarm it is from sound”, “Complementary to each other”, “Good mix”, and “Perfect”.

5. To what extent did the presentation of aural and visual information overload or distract, compared to the pre-existing system?

Two participants rated this as ‘substantially less’, one as ‘less’, and one as ‘slightly less’. The participant comments were: “Less time spent on determining what type of alarm”, “Less distracting” and “Because there is a minute to process an MPR or ETO (P4) before an aural alert there will be a huge reduction in aural alerts on a procedural sector”.

6. Rate the overall effectiveness of new fixed alerts (Timer and AIF).

Three participants rated this as ‘excellent’ and one as ‘good’. The participant comments were: “Excellent, not obtrusive, but easily recognizable”, “Timer and AIF substantially improved – easy to discern what the alert is for, and less obtrusive”, “Like them a lot” and “Excellent”.

7. How do the new fixed alerts (Timer and AIF) complement the priority aural alarms (P1-P4)?

Three participants rated this as ‘excellent’ and one as ‘good’. The participant comments were: “Maybe hard to hear if other alarms are going”, “Really, really well”, and “Great”.

8. Rate the overall effectiveness of each of the new aural alarms.

Ratings for this question are shown in Figure 10. The participant comments were: “All seem to work well together”, “All really excellent”, and “Wouldn’t change a thing”.

Further to the views of controllers, observations were made on the substantially enhanced situational awareness for a supervisor in a control room using these alerts.

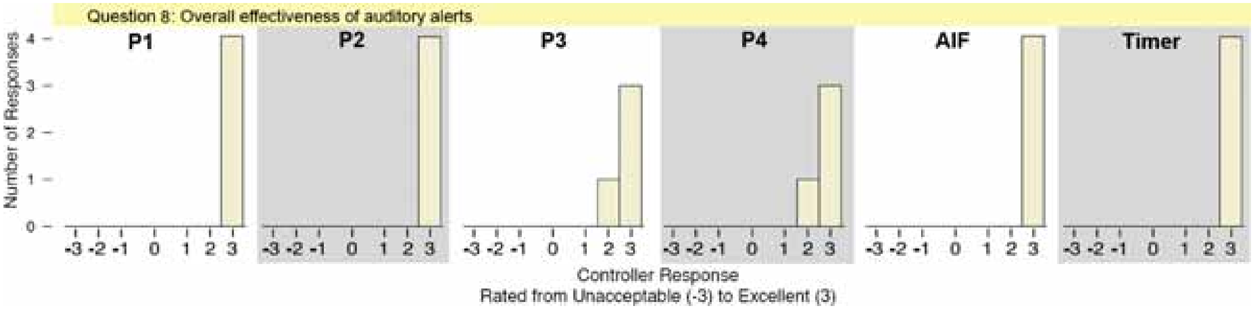

Figure 10 Rating distributions for Question 8 of the second trial. With four participants, the maximum possible result is 4.

2.3. Selected alert characteristics

It is not possible, in text, to convey fully the alert characteristics, but some of their salient characteristics are given in this section.

The four priority alerts (P1-P4) convey urgency redundantly through the parameters given in Table 1.

The fundamental frequencies of the alert signals were designed to maximize pitch segregation by having a large frequency interval and by minimizing harmonics in common. Hence an approximate tritone interval was used between the fundamental frequencies, with a quarter tone shift for what would otherwise have been octave related fundamentals. The intention was to make the sounds maximally distinguishable, even when played simultaneously (e.g. by two adjacent consoles).

Low frequency fundamentals are conveyed as virtual pitches [38], since the audio system is not sufficiently efficient to produce the actual fundamental frequencies at a useful sound level. Attenuating the fundamentals from the wave files enabled these alerts to be produced at a higher A-weighted sound level.

The number of pulses in a group, pulse rate and intervals between group repetitions were developed in the meetings with operational specialists. The urgency schemes of Edworthy *et al.* [4] and Hellier *et al.* [5] are not strictly applicable to the selected sounds, but suggest that a large and appropriate range of urgency is represented by the four Priority alerts.

	P1	P2	P3	P4
Number of pulses in group	14	5	2	1
Pulse rate	6.3 Hz	1.9 Hz	1.3 Hz	-
Pulse duration	160 ms	517 ms	750 ms	1500 ms
File duration	2.26 s	2.6 s	1.49 s	1.57 s
Interval between file repetitions	1 s	2 s	4 s	20 s*
Fundamental frequency	740 Hz	520 Hz	385 Hz	268 Hz
Sound level (vol. = 3)	70 dBA	63 dBA	57 dBA	58 dBA

Table 1 Key characteristics of the priority auditory alert signals. * For some P4 events, the auditory alert only sounds after a 60 s delay if unacknowledged. Sound levels are those measured as described in Section 2.2.4.

For the most part, the priority alerts have a fast attack and much more gradual decay. Such temporal envelopes can cause a startle reaction [2]. However, candidates with more gradual attacks were generally disliked, and the authors speculate that this may have been due to the auditory looming effect discussed earlier. The alert signals used have nominal sound pressure levels that are relatively low, and it seems probable that this reduces the likelihood of startle.

The attack portion of the impulses was designed to enhance auditory localization, since this had been identified as a problem with the pre-existing alert scheme. The attacks contain spectral content up to 10 kHz, although this broadband spectrum is brief in all cases. The repeated pulses of P1-P3 appear to provide particularly robust localization. Modifications made to the attacks and playback levels in the refinements described in Section 2.2.4 are likely to have significantly enhanced the localization cues. However, localization was not assessed in the final trial.

If L_{90} is taken as the background noise level, then the volume control setting of '3' almost always yields a signal level 10 dB above the background noise of the operational environment. However, a higher volume control setting may be desirable in busy periods. In refining the alerts, the sound pressure levels were limited by the hardware used, with lower frequency tones conveyed less efficiently by the hardware than higher frequency tones (even after high-pass filtering to remove the fundamental frequencies). The absolute sound levels are somewhat less than recommended for some auditory alert contexts [e.g. 2, 35], yet appear to be without problem in the operational environment, including in alerting supervisors to high priority events. The relatively low sound pressure levels should have the advantages of minimizing operator stress and background noise build-up.

A delay of 60 s for two of the three event categories that cause P4 auditory alerts means that there is a substantial decrease in the number of auditory alerts compared to the pre-existing system. If the operator acknowledges the visual P4 alert prior to the auditory alert, then no auditory alert will sound. Otherwise (after 60 s) the alert sounds every 20 s. This quieting of the control room environment was considered to be a major contribution of the new auditory alert scheme.

The information alerts, described in Section 2.2.3, are distinct in style from the Priority alerts, and therefore easily distinguished. Measured playback sound levels were between 56 dBA and 60 dBA for these alerts.

The auditory alert scheme designed for this project used sounds selected not only for their technical merit, but also for their design appeal. Although the sounds are by no means musical, they might be described as having greater musical qualities than other auditory alert schemes proposed in the literature and implemented. This bias towards timbral sound quality came in part from the musical training of the alert designers (both have their first degrees in music), but was reinforced in the meetings with operational specialists from Airservices Australia in which candidate alerts were selected and refined. Synthetic sound possessing a simple signal structure can have a perceived problem of 'artificiality', which in some circumstances limits its appeal [39]. The auditory display field is still grappling with the problem of relating empirical knowledge to aesthetics, and it seems likely that future design guidelines will be more strongly influenced by aesthetics.

2.4. Implementation

The new auditory alerts (P1-P4, AIF and Timer) were implemented in the Australian air traffic control system on the 28th July 2005, following five weeks of training.

The Keyboard Entry Error and CPDLC auditory alerts have not been implemented due to software limitations, but may be implemented in future.

Trials indicate that the auditory alert scheme of this project is highly advantageous. A post-implementation review will be undertaken to ascertain whether operational issues become apparent over the longer term, such as distraction and clutter.

3. CONCLUSION

With regard to the aim of this project, the new auditory alerts give controllers immediate awareness as to the criticality of the triggering event, irrespective of where the visual focus of the controller lies. This advantage is supplemented by the situational awareness available to the supervisor, who is immediately notified of any critical event by the recognition of a high priority audible alarm. The information and feedback alerts were found to be easily differentiated and to complement the priority alerts well, without creating confusion.

As an empirically-based field, auditory alert design is quite young, and no doubt will develop extensively in the coming years. Nevertheless, the design of auditory alerts has a long history, and auditory alerts make an important contribution to everyday living. In many cases, established auditory alert schemes are highly effective, and perhaps optimal for the application. Unfortunately there are many instances of poor auditory alert design – for example, where the alert provokes annoyance or confusion more than meeting its goal.

Such alerts can be highly detrimental to the sound environment. Now that there are widely available resources on auditory alert design principles, such poor design cannot be justified. Psychologists, audio engineers and musicians have between them a wealth of skill and knowledge which can be applied to effective auditory alert design.

This paper gives a practical example of how good audio design, albeit within a limited audio system, can have a profoundly positive effect on complex computer-based work environments such as an air traffic control operations room.

4. ACKNOWLEDGEMENTS

Many people contributed to this project, and the authors particularly acknowledge the contributions of Duncan Kennedy (Airservices Australia), Trevor O'Donnell (Directorate of Safety and Environment Assurance, Airservices Australia) and the trial participants. The changes to the visual display referred to in this paper were developed by Dr Steven E. Jenkins (Steven Jenkins and Associates Pty Ltd).

5. REFERENCES

[1] K. M.Cardosi and E. D. Murphy (Eds), *Human Factors in the Evaluation of Air Traffic Control Systems*, Final report Volpe Centre ATC Human factors Program, U.S. Department of Transportation Research and Special Programs Administration (April 1995).

[2] R. D. Patterson, "Guidelines for auditory warning systems on civil aircraft," CAA Paper 82017, London Civil Aviation Authority (1982).

[3] R. D. Patterson, "Guidelines for the design of auditory warning sounds," *Proceedings of the Institute of Acoustics*, vol. 11, no. 5, pp. 17-24 (1989).

[4] J. Edworthy, S. Loxley, and I. Dennis, "Improving auditory warning design: Relationship between warning sound parameters and perceived urgency," *Human Factors*, vol. 33, no. 2, pp. 205-231 (1991).

[5] E. J. Hellier, J. Edworthy, and I. Dennis, "Improving auditory warning design: Quantifying and predicting the effects of different warning parameters on perceived urgency," *Human Factors*, vol. 35, no. 4, pp. 693-706 (1993).

[6] J. L. Burt, D. S. Bartolome, D. W. Burdette and J. R. Comstock, "A psychophysiological evaluation of the perceived urgency of auditory warning signals," *Ergonomics*, vol. 38, no. 11, pp. 2327-2340 (1995).

[7] V. Ahlstrom and K. Longo, *Human factors design standard* (DOT/FAA/CT03/05), Atlantic City International Airport, NJ: DOT/FAA, William J. Hughes Technical Center (2003).

[8] J. A. Ballas, "Delivery of information through sound," in G. Kramer (Ed.) *Auditory Display: Sonification, Audification, and Auditory Interfaces*, Addison Wesley, Reading, MA, pp. 79-94 (1994).

[9] A. Guillaume, L. Pellieux, V. Chastres and C. Drake, "Judging the urgency of nonvocal auditory warning signals: Perceptual and cognitive processes," *Journal of Experimental Psychology, Applied*, vol. 9, no. 3, pp. 196-212 (2003).

[10] A. Guillaume, C. Drake, M. Rivenez, L. Pellieuz and V. Chastres, "Perception of urgency and alarm design", presented at the 8th International Conference on Auditory Display, Kyoto, Japan, 2002.

[11] W. von Aures, "Ein Berechnungsverfahren der Rauhigkeit," *Acustica*, vol. 58, pp. 268-281 (1985).

[12] E. Zwicker and H. Fastl, Psychoacoustics: Facts and Models, Springer-Verlag, Heidelberg, Germany, 1990.

[13] H. Fastl, "The Psychoacoustics of Sound Quality Evaluation." *Acustica/Acta Acustica* vol. 83, no. 5, pp. 754-764 (1997).

[14] W. Spieth, "Annoyance threshold judgements of bands of noise," *Journal of the Acoustical Society of America*, vol. 28, pp. 872-877 (1956).

[15] W. Babisch, "Health effects of extra-aural noise research," *Noise Health*, vol. 6, no. 22, pp69-81 (2004).

[16] G. W. Evans, P. Lercher, M. Meis, and W. W. Kofler, "Community noise exposure and stress in

children," *Journal of the Acoustical Society of America*, vol. 109, pp. 1023-1027 (2001).

[17] J. G. Neuhoff and T. Heckel, "Sex differences in perceiving auditory 'looming' produced by acoustic intensity change," presented at the 10th International Conference on Auditory Display, Sydney, Australia, 2004.

[18] A. Bregman, *Auditory Scene Analysis*, Cambridge (MA), MIT Press (1990).

[19] J. G. Neuhoff, R. Knight and J. Wayand, "Pitch change, sonification, and musical expertise: Which way is up?" presented at the 8th International Conference on Auditory Display, Kyoto Japan, pp. 351-355, 2002.

[20] J. L. Burt, *Empirical Studies Concerning Aural Alerts for Cockpit Use Leading to an Aural Alerting Signal Categorization Scheme*, PhD thesis, Virginia Polytechnic (1999).

[21] J. Blauert, *Spatial Hearing: The Psychophysics of Human Sound Localization*, rev. ed., MIT Press, Cambridge MA (1997).

[22] American National Standards Institute, *Methods for the Calculation of the Speech Intelligibility Index,* ANSI S3.5-1997 (1997).

[23] International Electrotechnical Commission, *Sound system equipment – Part 16: Objective rating of speech intelligibility by speech transmission index*, IEC 60268-16 (2003).

[24] V. Peutz, "Articulation Loss of Consonants as a criterion for speech transmission in a room," *Journal of the Audio Engineering Society*, vol. 19, no.11, pp. 915-919 (1971).

[25] V. Ahlstrom and G. Panjwani, *Auditory Alarms in the Airway Facilities Environment*, US Department of Transportation, Federal Aviation Administration, (2003).

[26] D. Cabrera and S. Tilley, "Parameters for the auditory display of height and size," presented at the 9th International Conference on Auditory Display, Boston, USA, 2003.

[27] D. Cabrera, S. Ferguson, S. Tilley and M. Morimoto, "Recent studies on the effect of signal frequency on auditory vertical localization," presented at the 11th International Conference on Auditory Display, Limerick, Ireland, 2005.

[28] K. R. Scherer, "Speech and emotional states," in *Speech Evaluation in Psychiatry*, J. K. Darby (ed.), Grune and Stratton, New York (1981).

[29] E. Schubert, "Modeling perceived emotion with continuous musical features," *Music Perception*, vol. 21, no. 4, pp. 561-585 (2004).

[30] D. McKeown, "Candidates for within-vehicle auditory displays," presented at the 11th International Conference on Auditory Display, Limerick, Ireland, pp. 182-189, 2005.

[31] P. Keller and C. Stevens, "Meaning from environmental sounds: Types of signal-referent relations and their effect on recognizing auditory icons," *Journal of Experimental Psychology: Applied*, vol. 10, pp. 3-12, (2004).

[32] R. A. Newman and K. Allendoerfer, *Assessment of Current and Proposed Audio Alarms in Terminal Air Traffic Control*, US Department of Transportation, Federal Aviation Administration, Report TN00/21 DOT/FAA/CT (2000).

[33] V. Ahlstrom, *An Initial Survey of National Airspace Issues in Terminal Air Traffic Control*, US Department of Transportation, Federal Aviation Administration (2003).

[34] D. Brock, J. A. Ballas and D. C. McFarlane, "Encoding urgency in legacy audio alerting systems," presented at the 11th International Conference on Auditory Display, Limerick, Ireland, pp. 190-197 (2005).

[35] International Organization for Standardization, *Ergonomics — Danger Signals for Public and Work Areas — Auditory Danger Signals*, ISO 7731 (2003).

[36] International Organization for Standardization, *Ergonomics - System of Auditory and Visual Danger and Information Signals*, ISO 11429 (1996).

[37] W. M. Hartmann, “Localization of sound in rooms,” *Journal of the Acoustical Society of America*, vol. 74, no. 5, pp. 1380-1391 (1983).

[38] E. Terhardt, “Pitch, Consonance, and Harmony,” *Journal of the Acoustical Society of America*, vol. 55, no. 5, pp. 1061-1069 (1974).

[39] E. Carterette, “Perception and physiology in the hearing of computed sound,” in *Structure and Perception of Electroacoustic Sound and Music*, S. Nielzén and O. Olsson, Eds., Amsterdam, Excerpta Medica (1989).