

About the importance of auditory alarms during the operation of a plant simulator

Matthias Rauterberg

*Human–Machine Interaction Research Group, Institute for Hygiene and Applied Physiology,
Swiss Federal Institute of Technology, Clausiusstrasse 25, CH-8092 Zurich, Switzerland*

Received 30 April 1996; received in revised form 31 October 1996; accepted 31 October 1997

Abstract

An experiment was carried out to estimate the effect of auditory alarms on the work of a plant operator in the context of a computer simulation. The process simulator was implemented so that each of eight machines (computer numeric controlled [CNC] robots) produced sounds to indicate its status over time. Each sound was designed to reflect the ‘real-world’ semantic of the actual breakdown event. As many as 32 different auditory alarms plus six normal machine sounds could be played at once. We attempted to design the auditory alarms so that none would be masked (rendered inaudible) by other auditory alarms. Eight students of computer science operated our process simulation program of an assembly line with the eight CNC robots. Relevant information of disturbances and machine breakdowns was given in a visual (test condition 1), and in visual and auditory form (test condition 2). The results indicate that the additional feedback of auditory alarms significantly improves operator performance and increases some mood aspects positively. © 1998 Published by Elsevier Science B.V.

Keywords: Audible alarms; Auditory feedback; Computer simulation; Human–computer interaction

1. Introduction

In general, human perceptions are selective. Humans do not react equally to all the stimuli impinging upon them, instead they focus on a few. This perceptual focusing is called *attention* [1]. Through attentive processes we keep in focus selected stimuli and resist distracting stimuli. The two dominant senses are the *visual* sense and the *auditory* sense. Each sense plays a different role in controlling attention: the eye is a directed sense and focuses attention, the ear is an all-round sense (‘omnidirectional’, Ref. [2], p. 102) and




Fig. 1. The different roles of the visual and auditory senses.

guides the visual attention (see Fig. 1). The auditory sense can take input from any direction and, unlike the visual sense, does not have a scanning potential as a valid index of selective attention. Of all the stimuli around us only those that our higher mental processes tell us are relevant to the psychological processes going on at the moment are selected for attention. Some sort of attention mechanism selects for further processing those sensory inputs that seem most important or pertinent. The following physical properties of a stimulus are important in gaining attention: *intensity*, *size*, *contrast* and *movement*. Certain internal variables, such as motives and expectations, are equally important in determining which stimulus attracts our attention.

Saunders and McCormick [3] (p. 71) recommended the following guidelines to improve selective-attention task performance:

1. Where multiple channels must be scanned for signals, use as few channels as possible, even if it means increasing the signal rate per channel.
2. Provide information to the person as to the relative importance of the various channels so that attention can be directed more effectively.
- 3...6. If multiple visual channels are to be scanned, put them close together to reduce scanning requirements.
7. If multiple auditory channels are to be scanned, be sure they do not mask one another.
- 8...

The dominance of the visual sense opposes the instinctive tendency that humans have to switch attention to stimuli in the auditory modality. These stimuli are intrusive and the

peripheral receptors have no natural way to shut out auditory information. Humans cannot close their ears in the same way that they can close their eyes, nor can they shift the attention of their auditory sense in the manner that they shift their gaze. As a consequence, auditory devices are generally preferred to visual signals as warning indicators [4].

Sound is a familiar and natural medium for conveying information that we use in our daily lives, especially in the working environment [5]. The hearing of sounds (e.g. alarms) is based on the perception of events and not on the perception of sounds as such [6]. For this reason, sounds are often described by the events they are based on. The following examples help to illustrate the important kinds of information that sound can communicate [7]:

- *Information about abnormal structures*—e.g. a malfunctioning engine sounds different from a healthy one and/or alarms in supervisory control environments give feedback about anomalous states.
- *Information about events in space*—all audible signals out of our visual field can direct the visual attention (e.g. footsteps warn us of the approach of another person).
- *Information about invisible structures*—all *hidden structures* that can be transformed into audible signals can be perceived through their characteristic sound pattern (e.g. tapping on a wall is useful in finding where to hang a heavy picture).
- *Information about physical events*—all specific semantics of real-world *events* can be perceived through their characteristic sound pattern (e.g. we can hear whether a dropped glass has bounced or shattered).
- *Information about dynamic change*—all specific semantics of real-world *dynamics* can be perceived through their characteristic sound pattern (e.g. as we fill a glass we can hear when the liquid reaches the top).




The textual representation is of most use when the operator is familiar with the domain area and can demonstrate much experience and knowledge in that domain area [8]. In comparison, more concrete and realistic (visual and audible) representations of information that the user of an interactive system can query are of most use when the domain area is new and unknown.

The parallel use of different media and the resulting parallel distribution of information, for example, by simultaneously showing a predecessor through a concrete representation and its explanation through audio distribution, leads to a denser sharing of information. In this case, the operator of a complex system can dedicate his or her attention solely to the visual information, which has parallel audio support. This reduces the need to change the textual or other visual delivery and prevents the overflow of visual information.

Sounds can be utilized to improve the operators' understanding of visual predecessors or can stand alone as independent sources of information. Gaver et al. [9] used sounds as diagnostic support applied to the direction of a process simulation, but they did not prove the hypothesis that an interface with audible feedback is superior to an interface without sound feedback. The authors describe only some global impressions of different operator reactions to sound feedback.

In the context of supervisory control an alarm is a signal that informs the operator of a dangerous or problematic process state. Wanner [10] classified alarms in two categories: programmed and not-programmed alarms. The first alarm class is divided by Riera et al.

Table 1
Dimensions and examples of visual and auditory alarm design

		Representational	Abstract
Visual	• Concrete	Picture, e.g. 	Speed → speedometer, e.g. 
	• Signified	Symbol, e.g. 	Danger → alarm flasher, e.g. blue light of a police car
Auditory	• Verbal	Speech, e.g. 'Stop the machine!'	Speech, e.g. 'Attention, please!'
	• Spatial	Onomatopoeia and mimic, e.g. event-generated sound pattern	Tone, e.g. beep-beep-beep ...

[11] into two groups: (1) *breakdown alarms*, which correspond to internal failures of components, and (2) *process alarms*, which shown an abnormal performance of a process. The not-programmed alarms are not defined at the time of system design but these audible cues are used by the operator (e.g. abnormal noise, smoke, steam, explosion, etc.).

Stanton et al. [12] (p. 85) classify alarms by their input modality (visual vs. auditory) and their information processing code (verbal vs. spatial). Spatial alarm processing requires a manual response to maximize performance while a verbal alarm requires a vocal response. Typical problems with alarms are "the avalanche of alarms during a major transient or shift in operating mode, standing alarms, alarm inflation, nuisance alarms, and alarms serving as status messages" ([12], p. 87). To extend Stanton's classification schema for alarms to include the approach used in this investigation we follow Edworthy and Adams [13] (p. 76) in their differentiation between representational and abstract alarms (see Table 1).

A pictorial icon (see the concrete telephone picture in Table 1) is in most cases a representational sign for the represented object. However, to find an adequate pictorial representation for dynamic phenomena (e.g. flow, time, speed, etc.) is a very difficult task. Arnheim [14] introduced an important difference between *generalization* and *abstraction*. The telephone example shows very clearly the generalization dimension (see Table 1): the concrete *picture* represents only a very special telephone type, but the telephone *symbol* can signify a large class of very different ones. Abstraction, in the sense of Arnheim, means something totally different: a projection of a concrete phenomenon to a perceivable structure (e.g. the speed of a vehicle given by the needle position in a speedometer, the time given by a watch, etc.). In this sense an abstract phenomenon can be very concrete! One of the most abstract alarms is the visual alarm flasher or a static alarm lamp. The

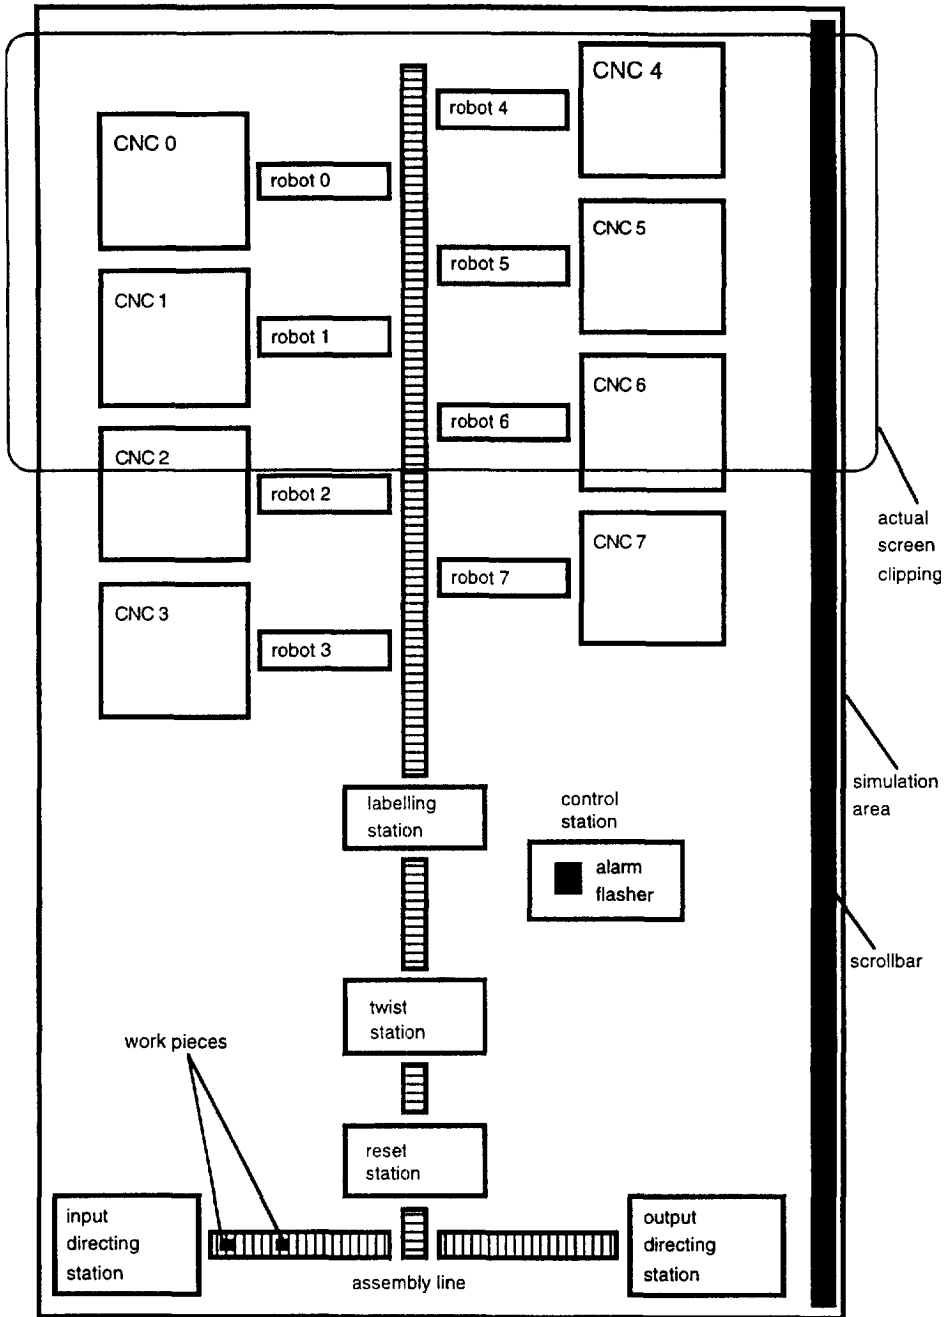


Fig. 2. The schematic view of the whole plant simulator. The rectangle shows the actual screen output that each operator can see at any given time.

semantic of these alarms must be given additionally in a written and/or iconized form, and must be learned by the operator ([13]). On the other hand, a concrete phenomenon can be an alarm sign by itself ([15], p. 171): a jammed or torn-off pipe, a spilled glass, etc. (see Wanner's not-programmed alarms).

Until now the extent to which auditory feedback in human-machine interfaces influences operator's performance has been unclear. Our main interest therefore was to test the hypothesis of Gaver et al. [9] that human operators in a 'real' process control situation can monitor multiple background activities simultaneously through continuous auditory sound feedback. We designed a system that produces visual (i.e. alarm flasher) and audible cues (i.e. representational and spatial auditory alert sounds) to support operators to monitor the status of ongoing processes.

The diagnosing and treating of problems with the plant were aided by alert sounds in the form of an event-generated sound pattern (auditory-spatial-representational, see Table 1), in addition to the visual alarm flasher at the control station (visual-signified-abstract, see Table 1 and Fig. 2). We carried out an experiment allowing us to test our hypothesis in a laboratory environment with a high alarm rate during a supervisory control task.

2. Method

2.1. Subjects

Eight male students of computer science at the Swiss Federal Institute of Technology (ETH) took part in the experiment as untrained operators (mean age of 24 ± 1 years).

2.2. Simulator

The simulation is based on a flexible manufacturing system that produces cases made of aluminium ('work pieces' in Fig. 2). The whole system consists of eight simulated computer numeric controlled (CNC) manufacturing centres and eight loading robots for these centres. In the input directing station all work pieces are automatically directed on the assembly line. The assembly line transports each work piece through different stations to the CNC manufacturing centres and back to the output directing station. The whole plant was deliberately designed to be too large to fit on the computer screen so operators could only see about half the robots and CNC machines at any time ('actual screen clipping' in Fig. 2).

The status of a work piece in the flexible manufacturing system could be one of the following:

1. loading on the assembly line at the input directing station;
2. transportation on the assembly line;
3. fixation on the carrier at the reset station;
4. final fixation and twist on the carrier;
5. fixation on a pallet with three other work pieces at the robot;
6. processing one of two sides in the CNC station;
7. changing from one side to the other at the reset station;

Table 2

Sound types: auditory alarm, duration and size (kB = kilobyte)

Machine	Sound	Alarm	Duration (seconds)	Size (kB)
CNC 0-7	CNC sound	No	1.20	51
Robot 0-7	Robot sound	No	0.39	16
I (input) station	I station sound	No	0.41	17
O (output) station	O station sound	No	0.78	33
R (reset) station	R station sound	No	1.40	60
T (twist) station	T station sound	No	0.40	17
L (labelling) station	L station sound	No	0.49	21
CNC 0-7	No cooling	Yes	1.08	46
CNC 0-7	Jammed pipe	Yes	1.38	59
Robot 0-7	Lost piece	Yes	1.04	44
Robot 0-7	Tear off pipe	Yes	1.04	44
Control station	Warning	Yes	0.24	10

8. being provided with a serial number at the labelling station;
9. loading off the assembly line at the output directing station.

Steps (3) to (7) are carried out twice, once for each side of the work piece.

We implemented our simulator so that each of the machines made a sound to indicate its status over time. Each sound was selected to reflect as much as possible the real-world semantic of the actual event (auditory–spatial–representational, see Table 1). For instance, a splashing sound indicated that cooling liquid was being spilled. Because of the complexity of our system, as many as 38 different sounds can be placed at once (see Table 2).

We attempted to use sounds such that none would be masked (rendered inaudible) by other sounds. We followed the two strategies recommended by Gaver et al. [9] as useful in avoiding masking. First, sounds were spread fairly evenly in frequency so that some were high-pitched and others lower. Second, we avoided playing sounds continuously and instead played repetitive streams of mixed sound slices (500 ms each), thus maximizing the chance for other sounds to be heard in the gaps between repetitions. CNC 0 and CNC 4 were characterized by a high-pitched sound, and CNC 3 and CNC 7 were low-pitched (cf. Fig. 2).

The normal running of a machine was coupled with the characteristic sound pattern of the corresponding machine type. Instead of the normal sound, each machine breakdown generated a specific alert sound: the auditory alarm (see Table 2). This alert sound was not a traditional alarm signal. If a robot or a CNC centre breaks down, then this centre cannot process the pallet of four work pieces further on.

A machine breakdown that will not be repaired immediately leads to a jam on the assembly line. The most important, but not dangerous, consequence of an overlooked alarm is the significant decrease in the performance and productivity of the whole plant.

2.3. Task

Subjects were instructed to operate a plant simulator and to ensure a high productivity rate. The task was to troubleshoot the whole manufacturing system. First, each subject had

Table 3

Breakdown types that lead to an audible alarm and their repair codes

Machine name	Breakdown type	Repair code
CNC 0-7	No cooling	3713
CNC 0-7	Jammed pipe	8319
Robot 0-7	Lost piece	1731
Robot 0-7	Tear off pipe	1733
Control station	Status request	8700

to detect that a breakdown had happened. Then he had to find the interrupted machine (robot or CNC machine). The actual breakdown event showed the operator how to repair the machine. The operator could get this information visually in a modal dialogue box with the status report at the control station or in an audible form through auditory alarm feedback.

A CNC machine could have two breakdown events ('jammed outlet pipe of cooling agent' or 'empty cooling agent'; see Table 3). A robot could breakdown with two different events ('lost work piece' or 'tear off a pressure pipe'; see Table 3).

The operator had to move the actual screen up and down by clicking with the mouse in the scrollbar area to go to the interrupted machine. He saw only the part of the whole plant given by the 'actual screen clipping' (see Fig. 2). Each interrupted machine could be repaired by entering an appropriate repair code (a four-digit number, see Table 3) in a repair dialogue box located at the machine. A mouseclick on the machine symbol would then pop up the repair dialog box. Entering the correct repair code transfers the interrupted machine in the normal state. If an incorrect repair code is entered, then no internal state change happens and the operator hears only a short beep. Each subject was trained to operate the simulator during a 15 minutes training phase beforehand.

The operators' view of the plant behaviour was that robots and CNC centres break down accidentally. The plant simulator was programmed so that all breakdowns appeared in the same sequence. This approach guarantees that the trials among operators are maximally comparable.

2.4. Procedure

The experiment had a two-factorial test design. Factor A was with or without audible feedback. Test condition 1 was only visual alarm feedback with a warning flasher and a modal dialogue box with the status information of each manufacturing system located at the operator control station (see Fig. 2). Test condition 2 was visual and audible feedback including the auditory alarms of each machine breakdown.

Factor B was a repeated measurement design for counterbalancing. Four subjects started the experiment with audible feedback (test condition 1) and repeated the same task without audible feedback (test condition 2). The other four subjects started without audible feedback (test condition 2) and repeated the task with audible feedback (test condition 1).

Each subject filled out a questionnaire to estimate the individual experiences with computers (answering time about 10 minutes). The subjects were introduced to operate the simulation tool through ‘learning by using’ (about 15 minutes) and instructed to maximize the productivity rate. The simulation for the troubleshooting task ran for exactly 20 minutes. Before and after each troubleshooting task the operator had to answer a workload questionnaire (eight scales with 36 items overall as monopolar rating scales; see Table 6; [16]). This workload questionnaire measured the mental workload at a rough estimate. After each troubleshooting task we measured the subjective satisfaction with a semantic differential (11 items as bipolar rating scales, e.g. ‘non-transparent’ versus ‘transparent’; see Table 5). Each individual session took about 90 minutes.

2.5. Material

We ran the experiment on an IBM compatible PC (Olivetti® i386, 25 MHz, 6 MByte main storage, 17" VGA colour screen) with an extra sound card (Logitech® 16 Bit, 44 kHz, stereo). A special simulation program was developed in Turbo Pascal® 1.0 to present the signals on the screen. Operators heard the auditory alarms via two small active speakers (maximal 3 watt). All machines on the left side (CNC 0,...,3 and robot 0,...,3; see Fig. 2) could be heard out of the left speaker. The right speaker produced the sound of all machines on the right side (CNC4,...,7 and robot 4,...,7). All other sounds of the machines in the center of the plant came out concurrently from both speakers.

2.6. Measures

Our first dependent variable was a point scale that measures the productivity of the plant. Each work piece that entered the assembly line at the input direction station counted as one point. One point was counted for each work piece side that was processed at a CNC machine. Each work piece that left the assembly line at the output direction station counted an extra point. Each work piece on the assembly line could therefore score from one to four points. The productivity score after 20 minutes simulation time was the sum over all work pieces that entered the assembly line.

The second dependent variable was the number of requested status reports at the control station.

The third and fourth dependent variables were the number of correct and the number of incorrect repairs.

In addition, the eight scales of the workload questionnaire and the 11 items of the semantic differential were recorded as dependent variables measuring operator satisfaction.

3. Results

First, we present the results of the four dependent variables that measure operators troubleshooting activities. There is a significant difference between the two test conditions for two of the four dependent measures (productivity score: MANOVA, $df = 1$, $F = 4.6$,

Table 4

Results for the four dependent variables that measure operators' troubleshooting activities for the two test conditions: with or without auditory alarm

Variable name	With auditory alarm	Without auditory alarm	<i>p</i>
Productivity score	70 ± 5.6	65 ± 5.3	0.05
Number of status reports	17 ± 5.8	23 ± 4.0	0.03
Number of correct repairs	36 ± 2.5	36 ± 2.3	0.99
Number of incorrect repairs	16 ± 11.0	9 ± 7.1	0.18

$p \leq 0.05$; number of status reports: MANOVA, $df = 1$, $F = 5.9$, $p \leq 0.03$; see Table 4).

Without auditory alarm feedback operators moved to the control station and requested the status report significantly more than in the test condition with sound feedback (see Table 4). We observed that the operators in the test condition with auditory alarm tended to go first to the control station to look for all breakdowns, and after that went through the whole plant to make repairs machine by machine. During this walk through they could remember all unrepaired machines by listening to the different sound pattern of each alarm type. On average the operators heard approximately five to ten different sounds per minute.

We observed a significant improvement through continuous auditory alarm feedback but also found that operators seemed to perceive the simulation with auditory alarms more non-transparent than without auditory alarms (see Table 5).

Operators felt significantly more self-assured (MANOVA, $df = 1$, $F = 6.9$, $p \leq 0.02$; see Table 6) and more socially accepted (MANOVA, $df = 1$, $F = 6.0$, $p \leq 0.03$; see Table 6) after working with auditory alarm feedback than without auditory feedback. Their readiness for endeavour, restfulness and motivation seemed to be increased in the test condition with auditory alarm feedback.

Table 5

Results for the eleven items of the semantic differential for the two test conditions: with or without auditory alarm (bipolar rating scale: (-) [-2, -1, 0, +1, +2] (+))

Variable name		With auditory alarm	Without auditory alarm	<i>p</i>
(-)	(+)			
Time-consuming	Time-saving	-1.1 ± 0.7	-1.0 ± 0.9	0.79
Rigid	Flexible	-0.9 ± 1.3	-0.8 ± 0.8	0.74
Circumstantial	Simple	0.5 ± 2.3	0.4 ± 3.1	0.89
Non-transparent	Transparent	0.4 ± 1.1	1.4 ± 0.6	0.06
Confuse	Unequivocal	0.1 ± 2.7	1.1 ± 1.0	0.18
Unclear	Clear	0.0 ± 2.6	-0.4 ± 1.4	0.60
Complicated	Uncomplicated	0.0 ± 1.1	-0.3 ± 1.9	0.71
Prescribed	Free	-0.5 ± 0.9	-0.4 ± 1.1	0.82
Unforeseeable	Foreseeable	0.0 ± 2.3	0.1 ± 1.8	0.87
Unsusceptible	Susceptible	-0.8 ± 1.1	-0.9 ± 1.0	0.78
Angry	Pleasing	-0.4 ± 1.7	-0.1 ± 1.3	0.71

Table 6

Results for the differences (after and before) of the eight scales of the workload questionnaire for the two independent test conditions: with or without auditory alarms (monopolar rating scale)

Variable name	With auditory alarm	Without auditory alarm	<i>p</i>
Readiness of endeavour	2.4 ± 4.1	-0.5 ± 4.1	0.20
Restfulness	1.3 ± 2.7	0.4 ± 3.3	0.59
Readiness for contacts	0.9 ± 2.5	-0.8 ± 2.2	0.22
Drowsiness	-1.1 ± 2.4	-1.5 ± 3.2	0.80
Self-assurance	1.8 ± 2.0	-0.6 ± 1.7	0.02
Social acceptance	0.1 ± 1.0	-1.1 ± 1.0	0.03
Feel excited	0.0 ± 6.1	-1.0 ± 5.9	0.74
Motivation	0.3 ± 2.2	-0.3 ± 1.0	0.13

4. Discussion

The sense of hearing is an all-round sense. This aspect is an important difference to visual perception, which is a directional sense. An auditory interface can be much physically larger than a visual interface (screen). Visually hidden aspects of parallel processes in the background can be made perceptible with auditory feedback [17]. The results of our experiment support this design approach of supporting operators with continuous sound feedback. Auditory feedback of concurrent processes that are important for task solving improves the usability of human-machine interfaces.

Audition is a spatial sense; we can be aware simultaneously of many sounds coming from different locations. Spatial patterns in audition, however, are much more limited than those of vision. Audition is primarily a time sense because its main patterns are those of succession, change and rhythm. Auditory feedback typically arrives sequentially in time whereas visual pattern might be presented either sequentially or simultaneously. Of course many perceptual experiences depend on the operation of several senses at once; then the prominence of one sense over another becomes a matter for study [18,19].

Auditory feedback has poor referability, meaning that sound patterns usually cannot be retained continuously although they can be repeated periodically. Visual patterns offer good referability because the information usually can be stored in the display. One of the possible advantages of auditory feedback is its attention-demanding capacity; it breaks in on the attention of the operator. Visual stimuli, however, do not necessarily have this captive audience. The operator has to be looking towards the display in order to perceive the stimulus. Hearing is somewhat more resistant to fatigue than vision ([20], p. 427).

How many different concurrent sounds can be discriminated? Operators reacted to up to 38 different sounds in our simulation study. Momtahan et al. [5] showed that staff in operating rooms were able to identify only a mean of between 10 and 15 of 26 alarms. Nurses were able to identify only a mean or between 9 and 14 of the 23 alarms found in their intensive care unit. Momtahan et al. explain their results as being due to the poor design of auditory warning signals and poor training. Standardization of auditory feedback can minimize this perceptual problem.

Cohen [17] found that it is a difficult task to design tones “which tell the right story and are also pleasant and emotionally neutral.” Good auditory feedback needs sound patterns that are interpretable without visual redundancy (e.g. door creaks open, door slams). We have to look for sound patterns that stand for themselves. Given these sounds we have to map them in a metaphorical sense to new events introduced by technology (e.g. door creaks open \Rightarrow ‘start a process’, door slams \Rightarrow ‘stop a process’). For simulation tools that deal with real-world events we can easily use the corresponding real-world sounds.

The results of our study support the ‘real sound’ approach. To avoid boredom and fatigue, caused by always giving the same sound pattern, the design of tones for auditory feedback should be highly context sensitive, e.g. listening to everyday sounds is based on the perception of events and not on the perception of sounds in and of themselves (see [6]). Everyday sound created in a particular environment is neither a characteristic of any of the participating objects nor a characteristic of the event itself. Different everyday sounds are solely determined by their respective interaction and environmental conditions. Most everyday sounds are the result of one or more interactions between two or more objects in a definite place and in definite surroundings.

Every interaction of everyday objects possesses attributes that have an influence on the produced sound. A framework concept for the description of auditory feedback is needed in which auditory alarms can be represented as auditory signal patterns along several descriptive dimensions of various objects interacting together in a certain environment. This approach is especially appropriate for the design of auditory feedback signals of the process alarms. To make auditory alarms context-sensitive leads directly to a design strategy that reduces the number of context-free alarms (cf. the discussion of reduction techniques in Stanton et al. [12]).

5. Conclusion

The results of this experiment show that the performance of operating a plant simulator on a computer could be significantly improved when the feedback from machine breakdowns and other disturbances was given in an audible form as well as a visual form. We also observed a significant increase in different aspects of operator mood. Overall, we can say that operators feel better and less stressed with sound feedback than without sound.

We found that auditory alarm feedback was effective in the following way. Auditory alarm feedback helped operators keep track of the ongoing processes and auditory alarms allowed operators to track the activity, rate and functioning of normally running machines. Without auditory feedback, operators overlooked machines that had broken down. With auditory feedback these problems were indicated either by the machine’s sound ceasing or by various alert sounds. Continuous auditory feedback allowed operators to hear the plant as an integrated complex process. The sounds merged to produce an auditory pattern, in much the same way as the many sounds of everyday machines do.

Using non-speech sounds to provide system information is appealing for several reasons. First, by adding sound to the interface the bandwidth of communication can be significantly increased. Second, the information conveyed by sounds is complementary to that available visually and thus sound can provide a means for displaying information

that is difficult to visualize, especially with a physically limited screen size. Auditory alarm feedback can improve the usability of human–machine interfaces because most interfaces stress visual perception but auditory feedback can help to reduce eye strain and fatigue.

Acknowledgements

We wish to thank the following persons for their generous support: Erich Styger for developing the simulation program, all students participating as test subjects and the anonymous referees.

References

- [1] D. Kahneman. Attention and Effort. Prentice-Hall, Englewood Cliffs, 1973.
- [2] C.D. Wickens, Engineering Psychology and Human Performance, Harper Collins, New York, 1992.
- [3] M. Sanders and E. McCormick, Human Factors in Engineering and Design, McGraw-Hill, New York, 1992.
- [4] R.D. Sorkin, Design of auditory and tactile displays, in: G. Salvendy (Ed.), Handbook of Human Factors, Wiley, New York, 1987, pp. 549–576.
- [5] K. Momtahan, R. Hetu, B. Tansley, Audibility and identification of auditory alarms in the operating room and intensive care unit. *Ergonomics* 36 (10) (1993) 1159–1176.
- [6] M. Rauterberg, M. Motavalli, A. Darvishi and H. Schauer, Automatic sound generation for spherical objects hitting straight beams based on physical models, in: T. Ottmann and I. Tomek (Eds.), Educational Multimedia and Hypermedia, Proceedings of ED-MEDIA'94, Association for the Advancement of Computing in Education, Charlottesville, 1994, pp. 468–473.
- [7] S. Mountford and W. Gaver, Talking and listening to computers, in: B. Laurel and S. Mountford (Eds.), The Art of Human–Computer Interface Design, Addison-Wesley, Reading, MA, 1990, pp. 319–334.
- [8] H. Marmolin, Multimedia from the perspective of psychology, in: L. Kjeldahl (Ed.), Multimedia: Systems, Interaction and Applications, Springer, Berlin and New York, 1992, pp. 39–52.
- [9] W. Gaver, R. Smith and T. O'Shea, Effective sounds in complex systems: the ARKola simulation, in: S. Robertson, G. Olson and J. Olson (Eds.), Reaching through Technology CHI'91 Conference Proceedings, Addison-Wesley, Reading, MA, 1991, pp. 85–90.
- [10] J.C. Wanner, Facteurs humains et sécurité, Séminaire: Erreurs humaines et automatisation, Ecole Nationale de l'Aviation Civile, 19–20 May 1987.
- [11] B. Riera, B. Vilain, L. Demeulenaere and P. Millot, A proposal to define and to treat alarms in a supervision room, in: Preprints of the 6th IFAC/IFIP/IFORS/IEA Symposium on Analysis, Design and Evaluation of Man–Machine Systems at Massachusetts Institute of Technology, Cambridge, MA, 27–29 June 1995, pp. 463–468.
- [12] N.A. Stanton, R.T. Booth, R.B. Stammers, Alarms in human supervisory control: a human factors perspective, *Int. J. Computer Integr. Manufac.* 5 (2) (1992) 81–93.
- [13] J. Edworthy and A. Adams, Warning Design. Taylor & Francis, London, 1996.
- [14] R. Arnheim, Visual Thinking, University of California Press, Berkeley, 1969.
- [15] W. Gaver, Auditory icons: using sound in computer interfaces, *Human–Computer Interaction* 1 (2) (1986) 167–177.
- [16] E. Apenburg, Befindlichkeitsmessung als Methode der Beanspruchungsmessung, *Zeitschrift für Arbeits— und Organisationspsychologie* 30 (1) (1986) 3–14.
- [17] J. Cohen, Kirk here: using genre sounds to monitor background activity, in: S. Ashlund, K. Mullet,

- A. Henderson, E. Hollnagel and T. White (Eds.), INTERCHI'93 Adjunct Proceedings, ACM, New York, 1993, pp. 63–64.
- [18] F. Hartman, Single and multiple channel communication: a review of research and a proposed model, *Audio-Visual Communication Rev.* 9 (6) (1961) 235–262.
- [19] S. McAdams and E. Bigand (Eds.), *Thinking in Sound*, Clarendon, Oxford, 1993.
- [20] E. McCormick, *Human Engineering*, McGraw-Hill, New York, 1957.