A Human Factors Perspective
Auditory Alarm Signals

Stephen B. Wilcox

About the Author

Stephen B. Wilcox, PhD, FIDSA, is a principal of design science, a human factors consultancy, and an author of the alarm systems section of ANSI/ AAMI HE75:2009, Human Factors Engineering—Design of medical devices. *E-mail: sbw@dscience.com*

"There is nothing more practical than a good theory."
–Kurt Lewin

IEC 60601-1-8: 2006[1] (hereafter, "1-8") has made designing a set of auditory alarm signals a matter of following a set of relatively simple and specific rules that define the acoustic characteristics of the pulses, the temporal characteristics of the bursts, and so on. Now, just about anyone, using simple sound-generation software, can design an acceptable alarm by following those rules—*acceptable* being defined as "meeting the requirements."

Are We Satisfied With the Results?

"*We*" includes the manufacturers and consultants who develop alarm systems, the organizations who purchase and maintain them, and the medical professionals who use them. The answer is pretty obvious when AAMI's very call for papers on the topic of alarms mentioned that: "[Alarms] frequently malfunction or are turned off, ignored, or unheard, earning a top spot on lists of the most frequent and serious problems seen with devices."

Most would agree that the two biggest problems with auditory-alarm signals are, in order:

- The number of "false alarms", and
- The lack of integration or even harmonization among the various alarm systems of different devices.

While human factors professionals alone cannot really solve these two problems, the discipline of human factors can shed some light on the current problems.

'False' Alarms

What every medical professional working in the relevant environments experiences every day has, of course, been thoroughly documented by study after study that show a very small percentage of auditory-alarm signals to actually require action.[2-4]

From a human factors point of view, the prevalence of alarm signals that don't require action (not just "false alarms") undermines one thing that we know how to do well—create an alarm signal that is *detectable*.

We go to great lengths to make sure that our signals can be heard, and then we, in effect, assure that they won't actually be heard by sounding them over and over when no action is required. It is a truism to say that humans have evolved perceptual systems in order to guide their actions. It follows that our perceptual systems are attuned to what is meaningful, defined as what requires some sort of action. What doesn't require action, we "tune out"—so-called *habituation,* one of the most basic and unavoidable human characteristics and one that, presumably, makes it possible for us to operate in the world by "editing" what we use our limited attentional capacity for (see Styles[5] for an in-depth overview of these issues, i.e., the psychology of "attention").

So, in sum, the "alarm fatigue" phenomenon is not just a matter of fatigue or annoyance, or even just a matter of conscious ignoring of alarm signals; it's a matter of (because of basic characteristics of human perception) failing to hear the signals in the first place, despite the fact that they're detectable.

Integration of Alarm Systems

More than 25 years ago, Kerr[6] wrote, "It is clearly no longer sufficient to consider each warning device in isolation." Here we are in 2011, and the intensive care unit (ICU) and the operating room (OR) are still filled with dozens of devices that were developed completely independently, each with its own alarm system, unrelated to any of the other alarm systems.

While efforts are underway to address this problem, until the various institutional, economic, regulatory, and reimbursal barriers to such integration are overcome, our alarm systems will be far from perfect, no matter what else we fix. And, as the technology makes distributed alarm systems more prevalent, and as more devices move into the home, it becomes even more important to integrate the signals from multiple devices. Although this is not a problem that human factors professionals are going to solve, there is one other factor that human factors professionals *can* control: the design of the auditory-alarm signals themselves.

Detectability

Sound is simply vibration of air, i.e., pressure variation. This vibration can be fully described by analyzing it into its sine-wave components, a so-called *Fourier transform*. It turns out that, as discovered by the French mathematician and physicist Jean Baptiste Joseph Fourier, any wave form can be broken down into its sine-wave components. Thus, any sound, including an alarm signal, is, at some level, a collection of often hundreds of sine waves, each with a particular frequency and amplitude, as well as phase relationships to the other sine waves.

In the meantime, the ear contains structures which begin to vibrate as a function of the particular frequencies that impinge on them. Simplifying a bit to avoid getting overly technical, I can say that the cochlea of the "inner ear" is basically a curled-up, fluid-filled tube that varies in diameter from one end to the other, so it vibrates in different locations depending on the frequencies of sound that reach it. Thus, it transforms frequency of the vibration into a position along the length of the cochlear tube. Then receptors that are positioned along the length of the cochlea pick up this vibration and transport the "information" to the brain so that we can hear which particular combination of frequencies has contacted our ears. In the meantime, the strength of the signals from the receptors varies as a function of the amplitude of the waves, so we have a way to perceive the loudness as well as the frequency of the sound that we hear. (See Warren[7] for a good overview of the mechanisms associated with auditory sensation.)

Another basic factor is that the different locations of the two ears provide various means of localizing sounds, based, for example, upon slight timing and amplitude differences between the sounds reaching the two ears.

The "alarm fatigue" phenomenon is not just a matter of fatigue or annoyance, or even just a matter of conscious ignoring of alarm signals; it's a matter of (because of basic characteristics of human perception) failing to hear the signals in the first place, despite the fact that they're detectable.

So there's a nice convergence between what we know about acoustics, or the physics of sound, and what we know about the anatomy of the ear that allows us to explain how hearing works.

Here are some principles that emerge from this knowledge:

- The ear is more sensitive to some frequencies than others. The "sensitivity curve" to frequency takes an inverted-U-shaped form, so that sounds in the midrange are easier to hear than sounds at either extreme—high or low. This means that, if we present multiple frequencies all at the same amplitude, the highest and lowest frequencies are harder to hear and are perceived as softer than the ones in the midrange. Most people are most sensitive to sounds in the 1,000-4,000 Hz range. This sensitivity curve informs the frequency ranges specified by 1-8.
- Loss of hearing is frequency-specific and generally at the extremes. The inverse of the above is that sounds that are too extreme cannot be heard at all, and, in general, hearing loss progresses with age from the extremes, particularly at the highest frequencies. The requirement of multiple frequencies in 1-8 means that the inability to hear any particular frequency will not make an alarm signal impossible to hear.

The psychology of perception has made a good deal of progress in understanding (and therefore predicting) what can be detected. However, it has made much less progress understanding (and therefore predicting) the *content* of perception—*what* we perceive. Which leads back to the *purpose* of perception.

- Masking of sound is frequency-specific. The presence of one sound makes another sound with a similar frequency harder to hear, i.e., it masks it. This is another advantage of the 1-8 multiple-frequency requirement. It means that any given frequency may be masked, but hopefully not all. One other implication is that the required amplitude of an alarm signal can be reduced if it is attuned to the ambient sound, i.e., contains different frequencies from the ambient sound. In general, the masking phenomenon provides an argument for the value of testing all candidate alarm signals in the actual or recorded ambient sound environment to assure that they can still be clearly heard.
- Lower frequency sounds travel farther (everything else being equal) but are more difficult to localize. This is why subwoofers can be located anywhere in the room, and why you can only hear the deep bass of that two-block-away stereo system. A key implication is that alarm signals that have to be accurately localized in space should not be too low in frequency.
- *Timbre,* the characteristic of sound that allows us to differentiate between musical instruments (say a piano vs. a saxophone) can be identified in terms of differences in the patterns of the associated sine waves.

This is not necessarily an exhaustive list, but it illustrates that there is some light from the study of basic "psychoacoustics" to be shed on the design and evaluation of auditory-alarm signals.

Note, however, that all of these examples have to do with *detectability,* not with whether a sound is easy to learn or remember, or if it's annoying or not, or if it can be understood, or if it can be easily differentiated from other sounds. In other words, the study of basic psychoacoustics has made much less progress when it comes to handling the *meaning* of sound as opposed to its detectability.

The Meaning of Sound

The psychology of perception has made a good deal of progress in understanding (and therefore predicting) what can be detected. However, it has made much less progress understanding (and therefore predicting) the *content* of perception—*what* we perceive. Which leads back to the *purpose* of perception. The purpose of perception is *to guide action.* So the reason we have perceptual systems is to optimize action.

Although this seems obvious, you can read quite a few articles and monographs about perception without any evidence of a recognition that perception has the purpose of guiding action. I believe that this bias on the part of psychology has led to an overemphasis on *detection,* which has, in turn, filtered down to that branch of applied psychology that we call human factors, that, in turn, has affected standards such as 1-8.

The result is that following 1-8 will assure that you have alarm signals that are detectable, but following it will not assure that you have alarm signals that are intuitive, easy to learn (as Judy Edworthy[8] discusses), or that allow anything but a crude understanding of the underlying condition that is being signaled.

And what happens when multiple 1-8 alarm signals begin to annunciate? With more than two or three, our ability to comprehend what's going on breaks down.

This contrasts with our ordinary auditory world. As I write this, I hear cars driving by, birds singing, an occasional telephone ringing, music in the background, a couple talking as they walk by, a cat purring at my feet, a dog barking in the distance, and my wife asking me a question. I don't have any trouble at all parsing these various sounds, understanding what and where they are, ignoring the ones that aren't relevant, and responding appropriately to the ones that I need to deal with.

Why is this so much easier than dealing with even a few alarm signals that sound simultaneously? I think, ultimately, the problem has to do with the way that psychologists have approached the study of auditory perception. Historically, the starting point for understanding auditory perception has been to examine the relationship between what the person hears on the one hand and acoustic "primitives" (i.e., sine waves) on the other. This has led to a number of productive insights, but it has also resulted in the creation of signals that have no inherent meaning, so are hard to learn, hard to differentiate, etc.

But there is another approach, as discussed in the next section.

Ecological Psychoacoustics

One result of starting with simple acoustic parameters is that psychologists, in studying perception, have traditionally put the locus of explanation for meaning "in the head." In other words, from the psychologist's perspective, the ear takes in simple, meaningless sine waves, then these sine waves are transformed from meaninglessness into meaningful experience via "mental processing," "information processing," etc. So the focus of perceptual psychology is to build models of this "cognition."

The practical problem with these models, from the point of view of a signal designer, is that they don't provide insight into what acoustic properties correlate with what aspects of meaningful perception. Indeed, many psychologists have even argued that such an endeavor is not even possible.

There is an emerging alternative, however, that is very intriguing and potentially much more useful. The approach originally championed by J.J. Gibson[9] involves turning the traditional approach on its head. Rather than starting with simple "primitives" for the "input" to the senses and trying to build a model of perception from there, Gibson argued for starting with natural experience and trying to find patterns in the "input" that correspond to experience—what Gibson called "building an ecological physics." Gibson and others have made quite a bit of progress in the area of visual perception with this approach, work that has led, for example, to the development of flight simulation.

The bottom line is that what is simple acoustically is complicated perceptually and vice versa, so that by creating signals via simple acoustic variables, we make signals impossibly complex for the humans that have to use them.

It is also being applied in the auditory area, yielding some very interesting results. (See Neuhoff[10] and Casey[11] for good introductions to this discipline of "ecological psychoacoustics.") A good deal of progress has been made in finding the acoustic patterns that correlate with natural sounds. A key finding is that the relevant parameters are not to be found in the pattern of frequencies or in simple changes in the pattern of frequencies (that is, in the very parameters that are specified in 1-8), but in higher-order patterns—changes in the changes of frequency patterns, what Casey calls "high-level change."

One simple study in this area was conducted by Vanderveer.[12] He presented 30 common

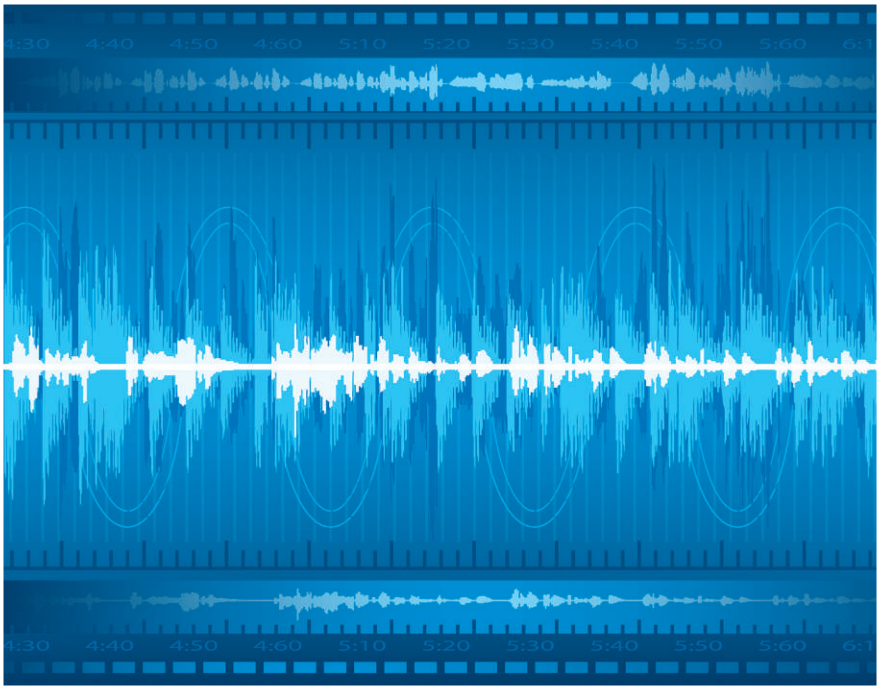

A good deal of progress has been made in finding the acoustic patterns that correlate with natural sounds.

sounds to people and asked them to identify the sounds (e.g., paper tearing, a person walking). He showed that, first of all, the task was quite easy. People were extremely accurate in identifying the sounds. Also, when asked to classify them, their categories were based upon the sources of the sounds, not their acoustic similarities. Another interesting finding was that, when asked what they were hearing, people described the sources, not the sounds. These findings contrast dramatically with Momtahan's[13] research described in this volume by Edworthy—that medical professionals do a poor job of recognizing the very alarm signals that they hear every day.

The bottom line is that what is simple acoustically is complicated perceptually and vice versa, so that by creating signals via simple acoustic variables, we make signals impossibly complex for the humans that have to use them. As Casey puts it:

"…there appears to be an inverse relationship between simplicity of stimulus and simplicity of perception."

Casey gives the examples of the sounds of coins dropping and glass breaking. We have no trouble immediately identifying these sounds. They are simple from the point of view of ordinary perception. However, their acoustic descriptions are quite complex, although, as he shows, not impossible.

In sum, it turns out that, as one would expect, if thinking about it from an evolutionary point of view, humans are much better at learning, differentiating, identifying, classifying, remembering, and understanding natural sounds than they are at learning, differentiating, identifying, classifying, remembering, and understanding sounds that have no inherent meaning. However, creating natural sounds (or natural-like sounds) requires manipulation of higher-order variables, not the simple variables that are typically manipulated by the signal designer and that are mandated by 1-8.

To repeat, human perception is inherently "tuned in" to various complex higher-order parameters of sound that correlate with natural sounds, but not particularly tuned in to simple parameters of sound such as frequency. It follows that creating signals based upon manipulation of simple acoustic parameters yields sounds that are hard for humans to deal with.

Conclusion

It is hoped that introducing a different way of thinking about auditory signals can start some readers down a different path that will eventually lead to radical (and better) types of alarm signals.

A key problem with alarm systems is that we've made it easier for ourselves and, therefore, harder for alarm users. It's easier to develop alarm systems device-by-device than to create integrated systems. It's easier to create alarms that are linked to simple parameter-based limits than to attune them to the complicated, multiple sources of information that medical professionals rely upon to make decisions.

Humans are much better at learning, differentiating, identifying, classifying, remembering, and understanding natural sounds than they are at learning, differentiating, identifying, classifying, remembering, and understanding sounds that have no inherent meaning.

And it's easier to manipulate simple acoustic parameters that we understand than complex higher-order variables that we don't fully understand.

As a consequence, people rather than machines have to integrate multiple alarm signals and draw conclusions from complex data sets; they have to consider alarm signals along with various other forms of information to decide what they have to do or not do, and they have to associate meaning with inherently meaningless sounds.

Regarding the latter, I applaud the work of Edworthy and her colleagues[14] on how to acoustically manipulate perceived "urgency" and the call for "earcons" that have shown empirical superiority to more traditional alarm signals.

However, that's only a start. We can benefit from a rethinking of our whole approach to auditory alarm signals, aiming instead to create *meaningful sonic environments*. What this would involve is taking advantage of the provisions in 1-8 that allow for auditory alarm signals that do not conform to the pulse and burst requirements that the standard provides. This is allowed if the alarm signal is based on a fundamentally different technology (synthesized speech is the example given) or if alarm signals are provided in addition to others that do conform to the pulse and burst requirements. (See subclause 6.3.3.1 and the associated commentary on subclause 6.3.3.1 in Annex A.)

There are at least three alternative approaches to developing alarm signals that are more meaningful:

- Using "earcons", that is, the auditory equivalent of icons, as described by Brewster, et al.[15] and McGookin and Brewster.[16] McGookin and Brewster have, in fact, demonstrated the perceivability of earcons when more than one are presented simultaneously.
- Using techniques developed by "Foley artists" who use all sorts of tricks to create sound effects for video and audio productions.[17]
- Manipulating higher-order acoustic parameters in line with the ecological approach to acoustics described by Casey.[11]

Even if new alarm signals are not meaningful, in the sense of recognizable as known sounds, the evidence is clear that the ear is not particularly attuned to simple acoustic parameters, but highly attuned to various higher-order parameters. This means that even to use acoustic parameters in an abstract way to convey meaning—e.g., to signal which device it is or which of several alarm conditions is relevant—it will be more effective to manipulate ecologically relevant higher-order acoustic parameters rather than simple acoustic parameters like basic frequency.

So the idea is to begin with a vision of what would make alarm signals meaningful, then to build meaningful alarm signals that conform to that vision, using one of the these three approaches. Of course, it will still be necessary to pay attention to frequency, amplitude, etc., in order to assure that the signals are detectable in the first place. And any alarm must be included in validation research to assure that it is effective.

Admittedly, trying to make alarm signals more meaningful will force the typical alarm signal designer out of his or her “comfort zone” and probably require collaboration with experts in these other fields. However, particularly in the absence of system integration that will reduce the number of alarm signals, alarm signals would be easier to learn, to understand, to remember, and to differentiate if the dimensions of variation were those higher-order parameters that human perception is attuned to rather than acoustic “primitives” that are easy to manipulate but inherently meaningless to the perceiver.

If true alarm-system integration is ever achieved, it will presumably be possible for auditory alarm signals to fulfill just one function—indicating to the user that he or she should go to the integrated visual alarm display, which would communicate the details. In such a world, *detectability* might be the only requirement for an alarm signal, making my discussion of meaningful signals moot. However, until the number of auditory alarm signals is dramatically reduced, making them more meaningful should make a real difference. ■

References

1. IEC 60601-1-8: 2006 *Medical electrical equipment—Part 1-8: General requirements for basic safety and essential performance—Collateral standard: General requirements, tests and guidance for alarm systems in medical electrical equipment and medical electrical systems.*

2. **O’Carroll TM.** Survey of alarms in an intensive therapy unit. *Anaesthesia,* 41, 742-744, 1986.

3. **Kestin IG, Miller BR, Moore CH.** Auditory alarms during anesthesia monitoring. *Anesthesiology,* 69, 106-109, 1986.

4. **Welch J.** Auditory alarms in intensive care. In N. A. Stanton and J. Edworthy (Eds.), *Human Factors in Auditory Warnings* (pp. 319-342). Brookfield, VT: Ashgate Publishing Co, 1999.

5. **Styles EA.** The *Psychology of Attention,* 2nd. Ed. London: Psychology Press, 2006.

6. **Kerr JH.** Warning devices. *British J. of Anaesthesia,* 57, 696-708, 1985.

7. **Warren RM.** *Auditory Perception: an Analysis & Synthesis (3rd Ed.).* Cambridge, UK: Cambridge University Press, 2008.

8. **Edworthy J.** Designing Effective Alarm Sounds. *BI&T,* 2011;45(4):290-294.

9. **Gibson JJ.** *The Ecological Approach to Visual Perception.* Boston: Houghton Mifflin, 1979.

10. **Neuhoff J, ed.** Ecological Psychoacoustics. Elsevier Academic Press, Amsterdam, 2004.

11. **Casey MA.** *Auditory Group Theory with Applications to Statistical Basis Methods for Structured Audio.* Unpublished Dissertation, MIT, Program in Media Arts and Sciences, 1989.

12. **Vandeveer NJ.** Ecological acoustics: Human perception of environmental sounds. *Dissertation Abstracts International.* 40/09B, 4543 (University Microfilms No. 8004002), 1979.

13. **Momtahan KL.** Mapping of psychoacoustic parameters to the perceived urgency of auditory warning signals [dissertation]

14. **Hellier E, Edworthy J.** On using psychophysical techniques to achieve urgency mapping in auditory warnings. *Applied Ergonomics,* 30, 167-171, 1999.

15. **Brewster SA, Wright PC, Edwards A.** Experimentally Derived Guidelines for the Creation of Earcons, in *Proceedings of HCI 95,* Huddersfield, UK, pp. 155-159, 1995.

16. **McGookin DK, Brewster SA.** An investigation into the identification of concurrently presented earcons. *Proceedings of the 2003 International Conference on Auditory Display,* Boston, MA, USA, July 6-9, 2003.

17. **Mott RL.** *Sound Effects: Radio, TV, and Film.* London: Focal Press, 1990.