

Alarm Reliability

What If an Alarm Goes Off and No One Hears It? BY AVI HAREL

The first time I used sound for alerting was in May 1948, just a few seconds after I was born. Presumably, the sound I generated was intended to notify my mother about the existence of a new creature to care for, day and night, for years. Later, I used sound occasionally to warn my parents about particular troubles that I experienced: pains, hunger, thirst, fatigue, boredom, and so on.

Who were the sound users, and for what purpose did they use it? We were three users altogether: I was the system, designed by God, but also a user of the sound—I generated the sound to report about the troubles I experienced; my mother used the sound to learn about my troubles; and God, the system designer, used sound to let my mother know when I was in trouble.

I cannot remember the exact date of the first time I used a sound alarm generated by a human-made system. It was probably as a pedestrian trying to cross the street. Presumably, a car driver was using the horn to warn me about being in danger. Here again we were three users of the system sound: the car driver who operated the horn; I (the audience), who used the sound to avoid a mishap, and the car designer, who designed the car with a horn so that the operator could alert me about the threat.

In 1981, I was in charge of the requirement specifications of a control system that Rafael, the Armament Development Authority (ADA) of Israel, proposed for the national security services. Here again there were three types of users: the operator (the system administrator) who sets the sound parameters; the commanders who use the sound as “customers,” and we, the

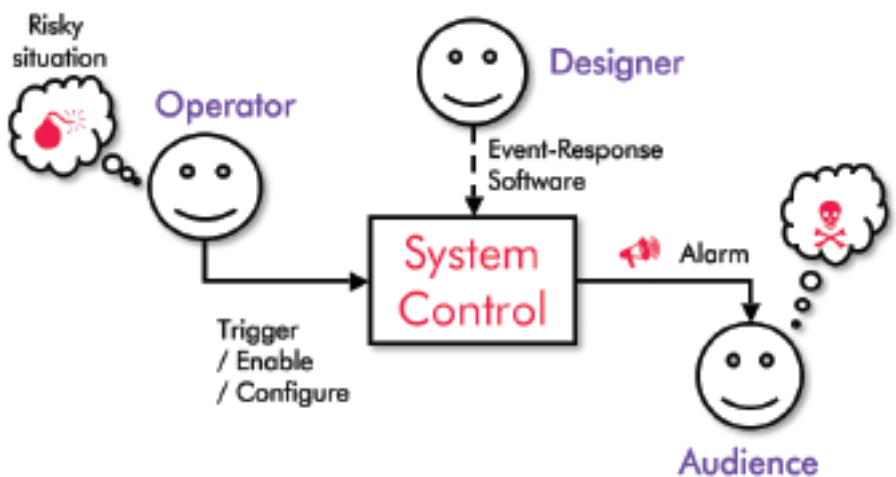
designers from Rafael. The documents that we wrote described the processes that the system should support. We knew that the operator should control the sound. However, the specifications did not mention when or how to use sound in the system operation.


There are three types of users involved in the previous examples the operator who controls the sound generation; the audience, the

target of the alert; and the designers who define when and how the operators will control the sound and what the audience will hear.

Which type of user was I? Apparently, I was each of them. In the first example, I was the operator. In the second, I was the audience. In the third, I was the designer. Apparently, most of the experience that sound designers have—that is relevant to sound

The Users of Sound





design—is from being the second type of user, namely, the audience of other systems that they observe and use.

Sound designers must learn when and how to use sound to alert about risky situations; how to ensure that the audience will attend the alarm and recognize the risk; what sound to provide in various situations; when to start and stop the sound; and how to adapt the sound attributes (pitch, level, rhythms, etc.) to the hazard; and more.

Risky Systems

Sound is used extensively in user interface design for entertainment and IT applications. For entertainment applications, such as games and videos, the sound quality is a key factor of the user experience. However, its role is often secondary: users can enjoy many games or videos even when the sound is turned off.

For IT applications (such as office applications), sound design is pragmatic; it is commonly used to provide additional feedback and to alert users about exceptional situations. The effect of using sound on the user experience is not dramatic; many users of such applications keep the speakers off.

Sometimes, however, the effect of sound design on the user experience is very dramatic. One of the incidents that Steven Casey describes in his extraordinary book *Set Phasers on Stun* is such an example. The true tale “A Memento of Your Service” is about a supertanker operating in the wrong mode in which the rudder was disconnected from the steering wheel, resulting in an ecological disaster. Casey mentions that an initial indication for the exceptional situation was the lack of clicking from the gyro compass, which was typical during course changes. Unfortunately, the ship’s captain became aware of the lack of clicking only after it was too late.

Control Systems

The subsystem of the supertanker that controls the navigation is an example of a control system. A common characteristic of control

systems is that users interact with the system in bursts; in normal situations they can feel relaxed and idle, but in exceptional situations the interaction becomes effortful and intensive. Examples of control systems include:

- ⊙ Certain mission-critical industrial systems such as production line control systems
- ⊙ Safety-critical industrial systems including chemical process control and power plant control systems
- ⊙ Transportation control systems including air traffic control, urban traffic control, and railway control systems
- ⊙ Medical monitors used in hospital emergency rooms
- ⊙ Command and control systems used in the military, police, emergency services, and rescue forces
- ⊙ Security systems such as fence control or surveillance systems

Alerting

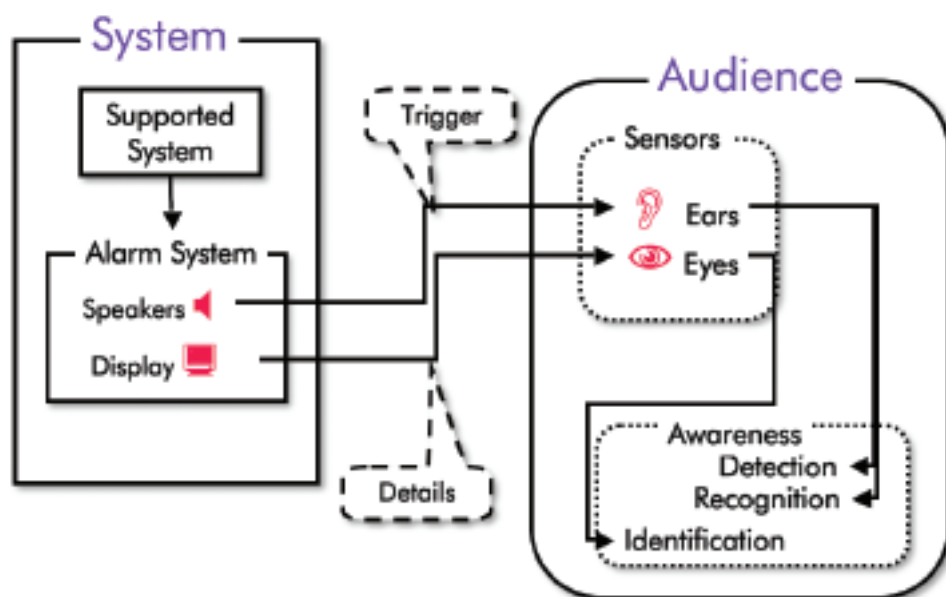
A main function of control systems is to alert the users about deviations from normal

conditions. Consider the example of the supertanker operating in the wrong mode, in which the rudder was not connected to the steering wheel. The initial indication for the wrong mode was not very obvious—it was the missing clicking typical of course change. Clearly, the mode indication was not by design. It was not in the operational specification. Apparently, none of the operation designers thought about how the ship’s captain would notice that the systems were in the wrong mode. Had they considered this at design time, they would probably have specified a clear indication of the exceptional situation, including both sound and visual cues.

Alert Design

The first step of alert design is risk analysis. This activity is application specific and performed by subject matter experts. The output of this activity is a list of potential hazards about which the audience needs to be warned.

Alert Design – Channel Allocation



The second step of alert design is channel allocation, which means deciding which perceptual channel of the audience will be in charge of the mental activities involved in alarm processing. This step consists of the following parts:

- ⊙ Capturing the alarm—Ensuring that the audience will notice the exceptional situation
- ⊙ Risk recognition—Providing hints about the risk level associated with the alarm.
- ⊙ Hazard identification—Providing details required to identify the sources of the risky situation

Typical designs rely on both the visual and audio channels of the audience. Sound is normally used to attract the attention of the audience to the exceptional situation and to indicate the kind of hazard. The video channel is used to get details about the situation.

The final step of alert design encompasses the detailed design—including sound design—intended to ensure hazard detection and recognition, and visual design, intended to enable hazard identification.

Alarm Reliability

To make sure that the users are aware of the exceptional situation, we need to ensure that the system generates alarms that are audible and well distinguished from background noise and from other operational sounds, and that the sound breaks any mental barriers. Unlike IT applications, in which the users may work with the sound turned off, sound is essential for reliable alerting of control systems. The reason is that users of control systems are not dedicated to the interaction, and therefore they might not observe any visual alerts that the system provides. The users may be idle, as during night shifts (a design challenge commonly known as “the vigilance problem”), or very busy doing something else (actually their main duties) during other hours. In any case, it cannot be assumed that they always attend the control system. Therefore, the key to reliable alerting is good sound design.

It is a measure of trust. Can the audience be sure that they will be notified about the exceptional situations and that they will actually notice the alarm? If they cannot trust the system alarm they are in a continuous alert situation, which ensures that they will get tired and be liable to misperceiving alert situations. Therefore, alert reliability is essential to gaining attention during exceptional situations.

Barriers to Alarm Reliability

What can go wrong with the system alarm? What are the typical situations in

which the audience might fail to perceive the alarm? Typically, control systems have four potential sources for alarm failure: technical, operational, environmental, and mental.

Technical

In case of a technical problem, such as when the speakers are disconnected, and there is no sound at all.

Operational

When the sound is disabled, because somebody turned it off, for example, to enable noise-free team discussion.

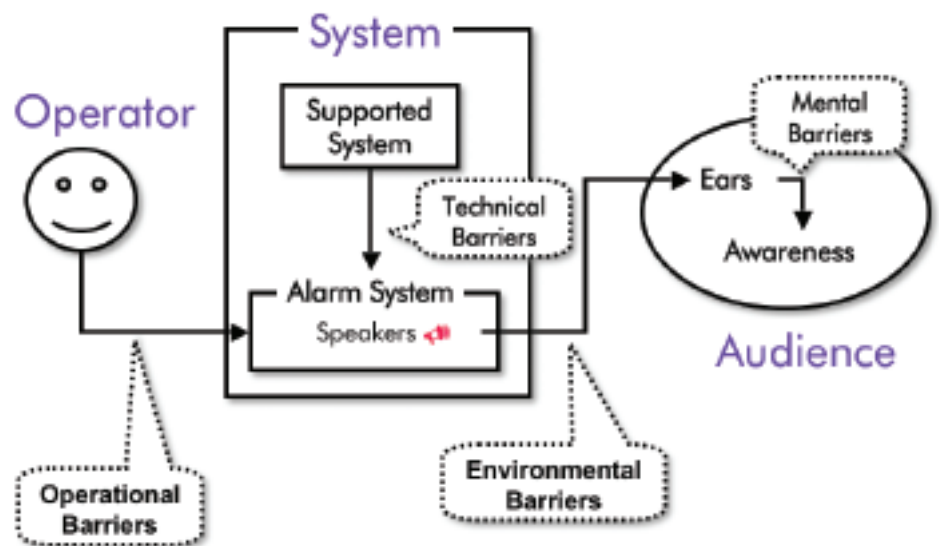
Basic Sound Design

Users might miss the alert even when it is audible because they are busy doing something else that requires their full attention. The design challenge is to shift the user focus from what they are doing to the alert, even when the users are operating under stress, such as in an emergency.

To attract the user’s attention to the alarm, the alarm sound should be well distinguished from the audio signals that the users receive regularly during normal operation.

Sound is defined by composition of tones, each consisting of sound attributes: pitch,

Barriers to Alarm Reliability



Environmental

- ⊙ When the sound is too low, below hearing threshold, because somebody reduced it when it was disturbing.
- ⊙ In case of temporarily noisy conditions, such as when operating a vacuum cleaner or when there is construction nearby.

Mental

- ⊙ When the sound is too weak to wake up the users during a night shift.
- ⊙ When the users disregard the sound, because of mental blackout due to emergency stress (a phenomenon called “tunneling effect”), or when they are too busy doing something else.

Here is the main challenge for sound designers: find ways to work around these problems to ensure that the audience will be aware of existing technical problems and of audibility limitations, and help them notice the alarm even when they are very tired or very busy doing other tasks.

level, rhythm, duration, etc. Sometimes the composition of tones forms a tune. For example, cellular phone companies enable users to set tunes as a convenient means to identify the callers. How can we decide which values to set for these attributes? How do we select tunes for alarms?

The traditional methodology for software development is incremental. You learn what existing systems can do and you build a new system that has more features and works better. This approach is inadequate for sound design. One problem is that there are only a few good designs and many poor designs. Most existing systems do not handle even the most frequent failure modes.

But, a more severe limitation is that it is difficult to decide which of the existing designs is good and which is bad.

Fortunately, we have a better, reliable source of knowledge about alarms: the safest way to handle alarms is by imitating nature. This approach is commonly used in “artificial

intelligence,” where we apply our knowledge about natural processes to designing artificial systems, making them look “intelligent.”

Learn from nature how to set the sound attributes. Note how babies call their parents when they are in trouble; examine how parents cry “watch out” to warn their children, and how a bird warns his spouse when a cat is getting too close. Typically, the level, pitch, and rhythms of the alarming sounds are higher than in normal communication. But more important than the physical attributes is the impact of sound on its audience. When designing sound for entertainment we think of tunes, melodies, and their entertaining effects. Alerting sound, on the contrary, should be annoying for the audience. It should make them stop what they are doing and pay attention to the warning signs.

Hazard Recognition

Basic sound design is about a singular hazard. It targets the first activity in alarm processing, namely, to capture the audience’s attention. In almost all practical systems this is insufficient because the audience needs to distinguish between various situations. For example, the alarm tunes about possible penetration to a secured base can distinguish between detection of suspicious objects and instances of hardware failure. Also, if a camera detects an object moving close to a surrounding fence, the alarm can be set to play a tune that sounds nice when the object’s direction is parallel to the fence; or dissonant when the direction is towards the fence.

The alarm sound can also provide hints about the alarming event. For example, the sound attributes can reflect attributes of the hazard. When a suspicious object is detected near the fence, the pitch can be inversely proportional to the object size, so that small objects will sound light and large objects will sound heavy. The rhythm can be directly proportional to the object’s speed, so that the rhythm of a fast object will sound fast, and so on.

Multi-Sensory Alarms

The amount of system-generated annoyance the audience can tolerate should be regarded as a resource of limited capacity. If this resource is wasted the audience becomes insensitive to the alarm. For example, consider a chemical plant in which the system alerts about exceptional parameters in the production line—such as too high or too low temperature, pressure, or percentage of specific composites, etc.

In a typical failure of a chemical process, more than one parameter might deviate from normal conditions. For example, in case of a leakage from a valve, the temperature in a tank may drop below normal and the pressure may rise above normal. In addition, the temperature and pressure of subsequent tanks may deviate from the normal. If the system alerts about each of the parameters individually, regardless of the other parameters, then the audience might become overwhelmed with alarms which might hamper the problem solving. Careful failure analysis is required to automatically identify the source of the exceptional situation and to alarm about the source rather than about the exceptional sensory data.

When designing sound for entertainment we think of tunes, melodies, and their entertaining effects. Alerting sound, on the contrary, should be annoying for the audience.

Continuous Alarms

Continuous alarms are alarms associated with exceptional situations that prevail for a long time. Consider an example of a surveillance system, in which the system alerts about exceptional situations, such as people staying in a forbidden zone. In a case in which people would stay there for a while—for example, doing maintenance—after the initial alarm, the system operator might turn it off because the continuous alarm is disturbing. Later, after the maintenance work is finished, the operator might forget to return the alarm back to the operational working mode. The alarm system is disabled, but the operator is unaware of it.

The design of continuous alarms is delicate and requires careful analysis of the alerting situation. For example, consider the chemical plant in the previous section, after a first alarm. Typically, it takes some time for the operational team to find out the source of the problem and to fix it. If during that time the system keeps alerting, the alarm might disturb the team in the problem solving. On the other hand, if the alarm stops after a while, and the team is busy solving another problem, they may forget to take care of the first problem. Also, if the system provides continuous alarms, then the team might turn it off intentionally, in an attempt to focus on problem solving.

To enable the team to focus on problem solving, we need to stop the annoying sound. However, to remind the team about the continuous hazard, we need to provide annoying sounds every now and then.

Repeating Hazards

The first true tale in Casey’s *Set Phasers on Stun* is about the well-documented accidents of the radiotherapy equipment Therac-25. The machine provided an error message “Malfunction 54,” but the operator disregarded this message because too many similar messages were involved in normal operation of the machine. The result was serious radiation burns, since that message meant that the radi-

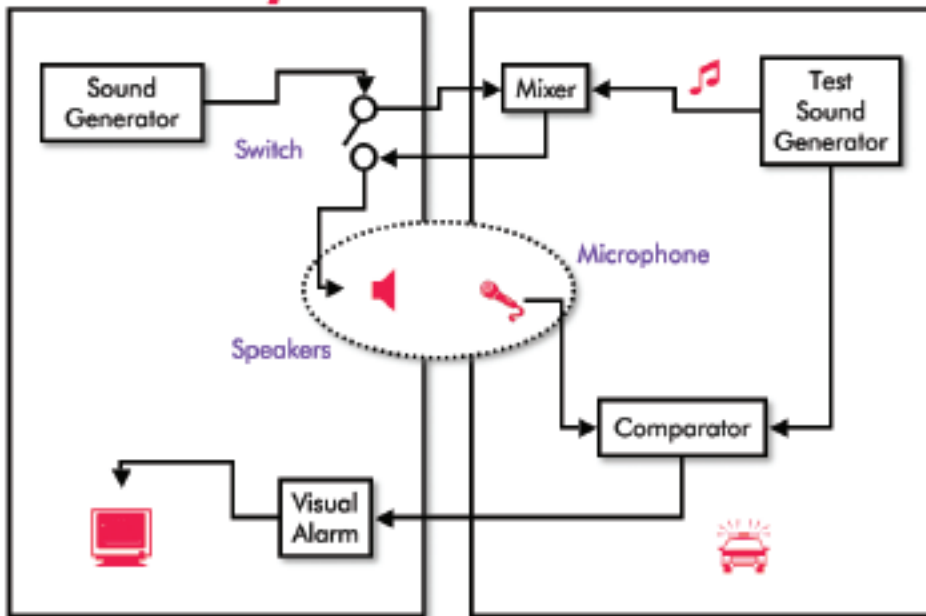
ation was not turned off when it should have been.

False alarms are a main barrier to operational vigilance. Casey sites another example of this effect in the true tale “Never Cry Wolf,” about a prisoner who escaped from jail just by crossing its fences, knowing that the guards would disregard the alert because they were used to false alarms generated by the wind and by wild animals. Terrorists often intentionally generate false alarms to reduce the sensitivity of security forces to the real alarms.

To ensure that the alarm sound alerts the users, the rate of false alarms should be minimized. How can we reduce the rate of false alarms? A well-known method according to “human detection theory” is to adjust the alert threshold. For example, suppose that the working temperature of a chemical process is in the range between 80 to 100 degrees, and that in a certain tank the temperature rises occasionally to temperatures higher than 100 degrees even during normal operation, resulting in false alarms. By changing the alert threshold to 105 degrees, the rate of false alarm should decrease. The problem with this approach is that by reducing the rate of false alarms, we increase the chance of missing real alerts. In the chemical process example, when the temperature rises above 105, the alarm may

Control System

Sound Tester



provide too short notice to enable recovery in cases of real hazards.

Another method for reducing the rate of false alarms is by risk analysis, namely, by careful examination of possible scenarios and adjusting the alert threshold to the situation. For example, if the alarm system of a jail measures the size and speed of objects crossing the fence, then the system may avoid most of the false alarms triggered by wind or by small animals.

Sound Reliability

Did you ever wonder why monitors in emergency rooms beep continuously, as often seen in movies? Obviously, the annoying sound indicates that the particular patient requires special attention. Also, the continuous beeping ensures that the personnel can rely on the sound—that the monitor will actually provide an alarm should the patient's situation get worse.

But then why is it so annoying? Beep, beep, beep... Is it intended for the personnel, to ensure that they are vigilant? Or is it intended to encourage the patient to recover, to go back home, away from the beep, beep, beep...? Couldn't the monitor designers provide relaxing, elevator-style background music instead? Or do they care more about development costs, and less about the users? Beep implementation is straightforward. Playing tunes requires some extra development efforts.

The risks are that the users occasionally turn the sound off because it interferes with their ongoing activities. And consequently, they might not be alerted when needed.

Sound Assurance

How can we be sure that a system generates alert sounds? Can the system know that the speakers are disconnected or that the sound switch is in the "Off" position? Or, can the system tell when the team discussion is over, and therefore the sound should be turned back on?

The system cannot decide automatically when the sound should be enabled or disabled. This is the user's task. The only thing that we can do in the design phase is to help the users become aware of situations in which sound is disabled. We achieve this by providing continuous test sounds, such as the beeps of medical monitors. Practically, "sound assurance" means ensuring that the users can hear the test sounds. This can be accomplished in various ways:

- ⊙ The intrusive way—The user is the watch dog. It is the user's duty to always listen to the sounds and to notice the absence of test sounds.
- ⊙ Non-intrusive ways—The system can detect situations when sound is not being generated and notify the users using messages displayed on a monitor. This can be done with special hardware; for example, a sound tester made up of a test sound generator, a mixer, a sensitive microphone adjacent to the system speakers, and a comparator. The tester generates hardly-audible test sounds, mixes the test sounds with the system sounds, captures the mixed sound through the microphone, and provides visual alarms when no traces of the test sound are found in the mixed sound. This solution is still theoretical, waiting for its first implementation.

Audibility Assurance

Suppose that a system generates an alarm sound; the speakers are connected and sound is enabled. How can we make sure that it is audible, namely, above the hearing threshold and the background noise, so that the users can actually hear it? Sound audibility can be assured using the same test proposed for sound assurance:

- ⊙ The intrusive way—The same test sound used for sound assurance is used also to detect situations of non-audible alert sound.
- ⊙ Non-intrusive ways—Automatic sound level adjustment through special hardware; for example, a modified version of the yet-theoretical sound tester mentioned above. The modified version can compare the level of test sound with the level of background noise and user-adjusted threshold levels and provide visual alerts when it is too low.

Conclusion

The challenge for designers of control programs, especially of those used in safety-critical and mission-critical systems, is to enable carefree interaction so that the users do not need to worry about missing sound alarms. This enables the users to focus on their main jobs, and to handle emergency situations successfully.

Traditional sound design does not support this requirement sufficiently; users are required to continuously stay tuned to the test sound and to identify situations when the alarm sounds are missing or below hearing threshold.

As sound designers pay better attention to the technical, operational, environmental, and mental details, systems that require sound alerts will become more reliable, and will place less burden on their users. The relationship between user, operator, and designer will approach a better balance, one that will ensure safety and allow the alarms to fulfill their natural role.

ABOUT AVI HAREL



Avi Harel is a mathematician with thirty years of expertise in user interface design and development. Avi is the inventor of ErgoLight patents and award-winning software tools for incorporating human factors in system design (<http://ergolightsw.com/CHI/Company/Articles.html>). Avi is the founder and active manager of ErgoLight Ltd. (<http://ergolight-sw.com>).