Emergency alarms: the 2006 missile attack case study

Avi Harel, Ergolight

ergolight@gmail.com

Abstract

People often fail to react properly to traditional emergency alarms. Specifically, they often do not know when the alarm is applicable to them, and how they should react. This study explains the essentials of safety alarm design, considering the relevance to the audience, and the need to enforce proper reaction. The study is based on subjective and anecdotal understanding of typical people reaction to missile attacks in the 2006 Israeli war with Hezbollah. The method proposed for enforcing proper reaction is based on the cybernetic principle of learning from the animals. The conclusions were implemented in the alarms produced by the Iron Dome defense system integrated with the Israeli Home Front Command (IHFC).

Introduction

Until the 2006 war, the emergency alarms generated by the IHFC were based on public warning equipment similar to those used for air raid sirens in WWII. This equipment was used for both memorial days and emergencies: on memorial days it generated a monotonous tone, and in emergencies it generated sounds with sinusoidal pitch.

During the ongoing wars with the neighboring assaults, officers of the IHFC noticed that many civilians did not obey the emergency alarms, and that those who did, often took the wrong actions (Harel, 2006 A). Sometimes, the consequences were fatal. Typically, the authorities blamed the victims for infringing the safety instructions.

The missile attack case study

During the Israel–Hezbollah war in the summer of 2006 the northern part of Israel was bombarded massively. The Hezbollah militants fired about 4000 rockets to the population of northern Israel, about 150 a day, for 34 days. Sirens were heard day and night, warning about rocket attacks. People reacted to the alarms in various ways. Mostly, they went into shelters, however, some people disregarded the alarms, and others overreacted, hysterically.

The means to detect the attacking missiles did not enable prediction of the explosion location. Only a few of the alarms ended up in nearby explosions. After a few days of war, many people got used to the alarms, and disregarded them. 43 civilians were killed, 418 civilians were treated in hospitals, and another 875 treated for shock. No formal analysis of the root cause of these casualties was conducted.

Traditional root cause analysis (RCA)

There are several ways for the developers to escape RCA. According to Murphy's Law, failure is often attributed to bad luck. According to the theory of Black Swans, coincidence is not predictable (Taleb, 2008). According to theories of organizational accidents, organizations prefer to blame the user instead of admitting their own contribution to the accident (Reason, 1997).

In contrast, the engineering approach is that incidents should be attributed to design mistakes, namely, lack of protection. The organizational challenge is to investigate the circumstances, and to identify these mistakes thereof. In the case of the 2006 war, nobody did this kind of investigation.

A main conclusion from the war was that the organization should maintain a safety culture. Rather than blaming the victims, the organization should focus on developing means to prevent mishaps.

Reactive RCA

According to the traditional approach to emergency alarm, the instructions for the civilians were to rescue themselves into shelters. Apparently, most casualties are due to problems following these instructions. The objective of RCA is to identify such problems. In reactive RCA we look for the trigger for the failure, hoping to learn how to prevent such triggers. In the context of emergency alarms, we want to find out what was wrong with the instructions.

Apparently, the IHFC officers were not aware of the sources of the people's misbehavior and did not investigate the typical conditions associated with the ways people reacted to the alarms. They did not ask why so many people disregarded the alarms, and how they could decide on the best way to react. The answers to these questions should be based on a model of human behavior when facing hazards, or when under threat.

Subjective RCA

How do people perceive the risks of hazards and threats?

I live in Haifa, a city that was bombarded by Hezbollah. A few weeks prior to the war I submitted an article to the Usability Professional Association (UPA) magazine on the subject of alarm reliability (Harel, 2006). The article has a section with guidelines for designing sound such that people who hear it may sense the threat.

A few weeks after the war started, I received an acceptance notice from the UPA magazine, asking me to make some changes. While editing the section on false alarms, I heard an alarm. I was busy editing, and since I was used to false alarms, I did not react. After finishing the editing, it came to my mind that I was not reacting properly to the alarm. Naturally, I blamed the design, rather than myself: why did the alarm not stop my editing? Why did I disregard the alarm?

This insight made me reexamine the effect of alarms on human behavior. I noticed that alarms about real threats were the same as false alarms. Most of the alarms that I heard were irrelevant to the region of Haifa, and those that were relevant to Haifa sounded the same.

The implication of this observation to alarm design is that alarms about real threat should be distinguished from false alarms.

Anecdotal RCA

How do people recognize the risks associated with the various reaction options?

A few days after this event I went to the grocery store to buy some food. When I was on my way back home, I heard an alarm, and as I was used to false alarms, I disregarded it and kept going carelessly home, carrying the food with me, as usual.

Near the grocery store there was a transporter parking, operated by volunteers of the civil guard. They used a megaphone to warn the people around there about upcoming rocket attacks. In this case, the warning was about the time remaining until the explosion. They said that it was 30 seconds. When I heard it, I started running home. I did not realize why I started running. I arrived home in the last seconds, yet I did know understand what made me change my behavior. Why did I disregard the alarm in the first place, and what made me start running after hearing the call from the megaphone?

Only after arriving safely home, I started to think of the two options that I had to choose from. In hindsight, my conclusion was that the information in the message received from the megaphone made the difference. It was critical for deciding what to do. At the time of the alarm, I was not sure about how I should react. It would be safer to run home if I could get there in time, and it would be safer to look for a nearby shelter or to lay by a nearby wall if I could not get home in time. I knew the way home well enough to estimate the time it would take to get there, but a priori I did not have any estimate of the time remaining until the expected explosion. Because I did not have the information required to decide on the safe option, I disregarded the alarm, and continued with my careless walk home. Only after hearing the megaphone, I realized that the option of running home was safer than the option of looking for a nearby shelter.

Human centered RCA

In traditional interaction design we empower the human operators, enabling them to access rarely used critical features, to cope with unexpected situations. These features are error prone in terms of feature accessibility or availability. A human factors version of Murphy's Law is:

The model of human reaction to alarms is based on a general paradigm, namely, the Human Factors version of Murphy's law (Zonnenshain & Harel, 2009):

> *If the system enables the users to fail, eventually they will!*

This paradigm emphasizes the system vulnerability to use errors, implying that it should be the developer's responsibility to eliminate the options of use errors. Specifically, in order to enforce special human behavior, the machine should provide information about the risks, and the information should be presented in forms considering the theory of human perception.

Use errors such as in this case study are due to inconsistent assumptions about the system situation, because the rules are implicit. They are not specified in the requirements documents, and the system does not distinguish between proper and improper situations. A preliminary version of this analysis was precented by Harel (2012).

Proactive RCA

The objective of proactive RCA is to highlight possible ways, and to propose principles and methods for enforcing the people awareness of the emergency, so that they can decide how to react in order to protect themselves. As Weinberg (1971) pointed out, technology-oriented engineers do not have the skills required to cope with this challenge. The human-centered design (HCD) approach proposed by Norman & Draper (1986) is more adequate for proactive RCA. The principles and methods may be based on studies, prior knowledge, applying practices from various domains, using subjective intuition, insight, stories and anecdotes.

Model-based RCA

According to the model of human behavior, people might act carelessly when they are not aware of the risks. In this case study people did not behave as expected for two reasons:

- Because they were not aware of the risks of disregarding the alarms
- Because they could not decide how they should react.

The operational model may be described as scenario-based, in which the scenarios correspond to desired user intentions, such as:

- The default scenario: the alarm is not relevant to me
- I can reach a proper shelter in the time slot that I have
- I can reach a second-best shelter for sure
- I may get enough protection by laying down.

To decide on the safest way, the user should know the places of the applicable shelter and have estimates of the time for getting at each of them.

Common design strategies

In traditional interaction design we focus on enabling the users to access rarely used error prone features, in terms of feature accessibility or availability. Common design strategies include:

- Feature-driven design: this is the basic approach, employed by engineers who focus on integrating features in the system design. In this case study, the features are those activated by the selection menu, the on-off control, and the pad used for parameter setting.
- Human-centered design: this is a methodology for developing systems that are usable, useful, and delightful to people, by focusing on the user's needs (Norman & Draper, 1986). The goal is to optimize the user's behavior, to match the intentions with the needs and to maximize the user's performance.
- HSI engineering: this is a new approach, focusing on eliminating operational confusion hampering seamless operation, by preventing error-prone activity.

A main barrier to reaching the user's needs is events attributed to user's errors, when the user's activity does not match the user's needs. The user's intention is a virtual variable, mediating the user's activity and the user's needs. In HCD we typically focus on ensuring that the user's intention complies with the user's needs. In HSI engineering we focus on ensuring that the user's activity complies with the intentions.

Cross-domain RCA

Often, we may come across incidents in other domains, with characteristics similar to those of our incidents. Eventually, the incidents of the other domains may have already been analyzed, which means that we can adapt conclusions from the investigations in the other domains. In the context of the 2006 war case study, we can think of alarms typically provided in emergencies of various sources, such as earthquakes, storms, floods, avalanches, etc. Solutions for emergency alarms from other domains may apply to the challenge of war alarms, and vice versa.

After the 2006 war was over my wife and I visited our son, who worked in his post-doc for Fermilab, near Chicago. During our visit there we encountered an event of thunderstorm, typical of the Chicago region. The alarm was vocal, and the information provided with the alarm was about the relevance and the timing of the expected hazard. Subsequently, the thunderstorm approach has been incorporated in IHFC alarm systems in Israel, integrated with the Iron Dom system protecting the Israeli population from hostile rockets.

Human-centered alarm design

The impact of alarms

One conclusion from the 2006 war was that people sometimes behaved carelessly, because most of the alarms were irrelevant to the audience, and therefore were perceived as false alarms. Another conclusion was that people took the wrong action because they did not have any idea about the time remaining until the explosion, which they needed to decide how they should react (Harel, 2006 A). The emergency system did not provide the information about the risks, and the information required for deciding how to react to the alarm.

The impact of emergency alarms depends on technical and cognitive factors:

- Technical factors of the alarm system, about alarm generation: activation, audibility
- Cognitive factors of the audience: risk perception, recognition, and reaction.

These factors are discussed by Harel (2006).

Safety-oriented engineering

Safety-oriented engineering is a combination of two strategies (Doc 9859, 2009):

- The Reactive Strategy is a gradual development of the safety requirements, in response to risky events. It is most useful when dealing with technological failures, or unusual events.
- The Proactive Strategy includes identifying hazards before they materialize into incidents or accidents and taking the necessary actions to reduce the safety risks (i.e., Risk Mitigation Plan). A key action is the validation of risk reduction. Any change in the system involves introducing new risks.

In the design of warning systems, this implies that we consider all expected circumstances, such as the user awareness of a sensor fault, the user's attention to sound alarms, the user's awareness of a failure of the audio channel, visual backups for the audio channel, testing procedures, etc.

The challenge of sound design

Sound designers should learn when and how to use sound to alert in emergency; 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; how to apply the sound attributes: pitch, amplitude, rate, rhythms, and more.

People's reaction to alarms depends on their perception of the risks associated with the hazard and on their estimation of the benefits of the reaction options that they recognize and identify. People might not react to alarms that sound like daily signals. For example, people might miss signals generated by the police or by medical transportation, because they are busy doing something 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, and to divert it to the emergency.

An alarm has three functions:

- Risk awareness: ensuring that the audience will notice the exceptional situation.
- Hazard recognition: providing hints about the risk level associated with the alarm.
- Reaction support: providing details required to identify the sources of the risky situation.

The primary goal of alarm generation is to avoid situations of missed alarms, while reducing the rate of false alarms. Another goal is to provide the audience with the information required to maximize their protection, namely, to enforce safe reaction to the alarm.

Risk awareness

Careless people behavior is often observed when people do not have the information required for decision making (Harel, 2020). Alarms may indicate two awareness modes:

- Hazard: when the user is aware of a concrete risk, but does not have the details required for eliminating the risk
- Threat: when the user understands what should be done to eliminate the risk.

Emergency control

Because users are error prone, we need to minimize its part in the situation control, to reduce it to the bare necessity. The principle of minimal human involvement leads to the principle of direct mapping from intention to action, such that the impact of the action is done by automation (Harel, 2021).

In normal operation, we may assume that the human intentions comply with the situation perception. Therefore, we can trivially extend the principle of direct mapping to apply to the human needs:

> *In normal operation, the mapping from human perception to action should be automated.*

This may not work when the system is in an exceptional situation, in which the human perception is often biased by prior knowledge, gained in normal operation. Therefore, specifically, for emergency control, such as war alarms, we need to extend the principle. The emergency variant of the principle of minimal human involvement is:

> *In an emergency, the mapping from human needs to action should be automated.*

Emergency systems may employ special means to ensure that human perception complies with their needs.

Behavior shaping

The requirements from the alarm system should focus on those intended to ensure that the audience is aware of the risk and knows how to react to eliminate it.

- In order that people can trust the system, it should generate an alarm each time an explosion might be heard, and the alarm generation should be reliable.
- In order that people can understand the meaning of the alarms for them, and to avoid hysteric responses, the alarm should indicate the risk level of the hazard, namely, the expected explosion.
- In order to enforce proper reaction to the threat, the alarm should provide predictions of the explosion location, and of the time remaining until the explosion.

A model of alarm perception

Typical alarm designs rely on both the visual and audio channels of the audience. Sound is normally used to attract the attention of the audience to exceptional situations. The visual channel is normally used to provide details about the risk, as well as guidance for the proper reaction. Sound may be used to accelerate hazard recognition, and to facilitate the guidance for the proper support.

The perception may be biased by subjective properties, such as over confidence, or hysteric behavior.

Enforcing risk awareness

The results of the subjective RCA suggest that the sound attributes should indicate an estimate of the risk of a missile attack. How can we decide which values of the sound attributes may be adequate for indicating the risk level?

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

Wiener (1948) suggested that we can learn from the animals. The best method to design effective 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".

To enforce risk awareness the audience should sense the emergency. The implication to alarm design is that the emergency sound should incorporate a perceptual code indicating an estimate of the risk to the audience.

Attributes of emergency alarms

How should we define the tunes for alarms?

The attributes of an emergency sound can reflect attributes of the risk. For example, when a suspicious object is approaching the audience, 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 proportionate to the risk level, so that the audience can sense the risk by the rhythm. In the context of emergency alarms, the residual time is inversely correlated with the risk: the shorter is the residual time, the higher is the risk.

Risk perception

Following Wiener (1948), Harel (2006) proposed that we should learn from the animals how to set 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.

Risk recognition

Basic emergency alarm design is monolithic. The goal is risk awareness. The objective is 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 for medical alarms enable fast recognition of the emergency. In the context of homeland security, if a camera detects an object moving close to a border fence, the alarm can be set to play a nice melody when the object's direction is parallel to the fence; or dissonant when the direction is towards the fence, hinting that this might be an intruder.

In the context of emergency alarms, the residual time is inversely correlated with the risk: the shorter is the residual time, the higher is the risk.

Reaction support

Under stress, it is difficult to follow or estimate the time elapsed since the beginning of the alarm situation, and to predict the exact time of the explosion. Apparently, at least some of the casualties during the recent battles with Hamas could have been prevented, should the alarm include an indication of the time left until the explosion, in real time.

The information of emergency alarms should enforce safe reaction of the alarm audience. The information should indicate the threat proximity. In the context of war alarms, the information for the audience includes:

- Location proximity: Risk recognition (is it a general hazard or a threat to the audience?)
- Time proximity: Threat perception (when should we expect the explosion?)

To enable safe reaction, the system should provide an estimate of the time remaining until the hazard becomes a threat. Indeed, in subsequent versions of the IHFC alarm system based on the lesson obtained from the case study, a new alarm system has been developed, which provides the audience with estimates of the explosion location and time proximity.

Alarm reliability

In the context of alarms, the term reliability refers to the feature of reducing the risks of failure in the alarm generation or in the user perception and recognition of hazards.

Missed alarms

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 six potential sources for alarm failure: technical, operational, functional, environmental, cognitive, and blackout.

- 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 discussion
- Functional: Avoiding interference with ongoing activities
- Environmental: In case of a temporary noisy conditions, such as when operating a vacuum cleaner or when there are construction works nearby, or when the sound is too low, below hearing threshold, because somebody reduced it when it was disturbing
- Cognitive: Due to alarm fatigue, such as during a night shift
- Blackout: When the users disregard the sound, due to emergency stress (a phenomenon called "tunneling effect"), or when they are too busy doing something else.

The design challenges include ensuring early detection of missed alarms due to problems with any of these sources.

The alarm dilemma

Following the theory of statistical inference, we can define two types of alarming failure:

- Type I (Alpha) error, missed alarm, when the system does not alarm about a hazard
- Type II (Beta) error, false alarm, when an alarm is generated, but there is no actual hazard.

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.

The alarm dilemma may be phrased in terms of the conditions for enforcing safe reaction. The basic alarm dilemma is concerned about when and how to enable or disable the alarms.

- The risks of missed alarms: Alarms are essential for the users to prepare for the hazard. In many reported accidents, such as Bhopal, Therac 25, Torrey canyon, AF296, and PL603, an alarm could have saved the extreme costs of the mishaps.
- The risks of false alarms: Repeating instances of false alarms might result in desensitization of the audience, reducing their vigilance thereof, implying that the audience might not trust real alarms when they are required in the future.

False alarms are often initiated and and their rates are accelerated on purpose by hostile agencies, to reduce our vigilance.

False alarms

The amount of system-generated annoyance the audience can tolerate is limited by the attention capacity. If this resource is wasted, the audience becomes insensitive to the alarm. To attract the user's attention to the alarm, the emergency sound should be well distinguished from the audio signals that users receive regularly during normal operation.

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

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

The vigilance challenge

A particular case of missed alarms is on the transition from tranquility to emergency. In normal days, the operators of emergency alarms are likely to disable the alarm system. It has been noticed that sometimes the first alarm is missing because the alarm system was disabled in purpose, to avoid the inconvenience of panics due to false alarms.

The vigilance challenge is to ensure that the operators will be aware of existing technical problems and of audibility limitations, and to help them notice the alarm even when they are very tired or very busy doing other tasks.

The intrusion dilemma

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

- The intrusive way: The operator is the watch dog. It is the operator'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 generated and notify the operators about it by messages displayed on a system screen. This can be accomplished by special sound detectors.

Engineering

The design of alarm system should focus on the effect on the user's behavior. Special facilities are required to ensure that users can perceive the alarms properly, so that they know how to respond when they face a real problem (link to article). Special modules may implement the requirements and methods for enforcing proper user's responses to the alarms.

Intra-domain learning

The Standards Institution of Israel (SII) committee for usability standards identified the need for standards for emergency alarms. They proposed to IHFC authorities to apply these conclusions but received no response. However, the IHFC officers decided to adopt the principles for designing emergency alarms after they were published by the association of Israel engineers.

It is quite natural that people hesitate before adopting new ideas. In this case, the IHFC authorities would not cooperate with the SII committee assigned to deal with this challenge.

As I failed to identify the responsible organization within the defense forces, I published an article in a conference on safety engineering, based on my personal experience, suggesting that alarms should include coded information about the predicted hazard location and timing. Fortunately, an IHFC officer was searching the internet for ways to enforce people obedience to defense instructions, and they found my article. These principles are being employed these days in the municipalities that suffer from rockets launched from the Gaza strip and southern Lebanon, aiming at civilians and soldiers.

Alarm development

The development of emergency may be done in cycles, of three steps each:

- The first step of alarm design is risk analysis. This activity is application specific, performed by subject matter experts. The output of this activity is a list of potential hazards, about which we need to warn the audience.
- 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.
- 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.

Inter-domain customization

There is much in common between various emergency needs, enabling applying this case study to alarming needs in other domains, including:

- Medical alarms, such as those used in monitors

- Control room alarms, such as those used in the Process Industry
- Vehicle driving ITS safety alarms
- Intruder alarms
- Disaster alarms, about environmental hazards.

Each of these domains needs to employ special modules, which sets it apart from the others. For example, the modules used for medical alarms should implement a feature of safe muting, to facilitate occasional consultation of the medical team, while maintaining a reminder about the state of alert. The modules used for driving alarms should employ a combination of physical sensing modalities and special audio alarms, to allow very fast hazard perception. The modules used to warn the public about upcoming disasters should provide the people with information about what they should do in order to save their lives.

The vision

Emergency alarms may be implemented by customized models based on generic rules such as those discussed in this article.

Sound designers should pay attention to the technical, operational, functional, environmental, cognitive, and blackout factors, looking for ways to improve the alarm reliability and to place less burden on the users.

References

Doc 9859 (2009). Safety Management Manual (SMM). International Civil Aviation Organization (ICAO) (http://www.icao.int/anb/safetymanagement/DOC_9859_FULL_EN.pdf).

Harel, A 2006. Alarm Reliability, *User Experience Magazine, Vol 5*., Issue 3.

Harel, A 2006 A. Using Sound for Alerting: Lessons from the War with Hezbollah. *User Experience Magazine*, Vol 5., Issue 3. Postscript

Harel, A 2012- Designing war alarms: a multi-disciplinary approach. *The Israeli Ergonomics Association on Human Factors Engineering for Military Systems, Netzer Sereny, Israel*.

Harel, A 2020. System Thinking Begins with Human Factors: Challenges for the 4th Industrial Revolution. in R.S. Kenett, R.S. Swarz and A. Zonnenshain (Eds), Systems Engineering in the Fourth Industrial Revolution: Big Data, Novel Technologies, and Modern Systems Engineering, Wiley, DOI: 10.1002/9781119513957.ch15

Harel, A 2021. Model-based Human Interaction Design. DOI: 10.13140/RG.2.2.22631.16807

Norman, DA & Draper, SW 1986. User Centered System Design: New Perspectives on Human-Computer Interaction", in *D. A. Norman and S. W. Draper (ed.)*, Hillsdale, New Jersey: Lawrence Erlbaum Associates, 1985

Reason, J 1997. Managing the Risks of Organizational Accidents, Ashgate.

Taleb, NN 2008. *The black swan*. Penguin Books.

Weinberg, GM 1971. *The Psychology of Computer Programming*, Van Nostrand Reinhold Company, New York.

Wiener, N 1948. Cybernetics or Control and Communication in the Animal and the Machine. MIT University Press, Cambridge, MA, 212.

Zonnenshain, A & Harel, A 2009 - Task-oriented System Engineering, *INCOSE Annual International Symposium, Singapore*, DOI: 10.1002/j.2334-5837.2009.tb00982.x.