Reprint: 4th International Workshop on Human Error, Safety, and System Development, Linköping, Sweden, June 11-12, 2001.

Describing, Predicting, and Mitigating Dissonance Between Alerting Systems

Lixia Song and James K. Kuchar
Department of Aeronautics and Astronautics
Massachusetts Institute of Technology
Cambridge, MA 02139 USA

Abstract

The potential for conflicting information to be transmitted by different automated alerting systems is growing as these systems become more pervasive in process operations. Newly-introduced alerting systems must be carefully designed to minimize the potential for and impact of alerting conflicts (or dissonance), but little is currently available to aid this process. The development of a model of alert dissonance would therefore be beneficial in terms of providing both a theoretical foundation for understanding conflicts and as a practical basis from which specific problems can be addressed. This paper identifies several different types of alerting dissonance that may occur. Using a state-space representation, a generalized model of alerting operation is developed that can be tailored across a variety of applications. Concepts of static and dynamic dissonance are introduced and directly related to the model. A preliminary human factors study is described to illustrate how differences in alert information can translate into reduced operator confidence. An example problem is also presented to demonstrate the application of the state-space model to identify where dissonance may occur and to develop countermeasures to prevent or mitigate conflicts.

Introduction

Automated alerting systems are becoming increasingly pervasive in time- and safety-critical operations, with applications spanning aerospace vehicles, automobiles, chemical and power control stations, air traffic control, and medical monitoring systems. As these applications are pushed toward higher safety and capability, new alerting systems have been introduced to provide additional protection from hazards. Accordingly, there has generally been an evolutionary, incremental addition of alerting systems to these applications over time. Because it is costly to completely redesign and recertify automation, new alerting systems are typically independent enhancements that do not directly affect the operation of existing subsystems.

The addition of alerting systems to an already complex operation carries several liabilities [1]. First, there is an increase in the amount of information processing required by the human operator, who now must be trained and able to respond rapidly to more information. There is also a potential for simultaneous alerts from the different systems, possibly overloading or confusing the human. This is a classic human factors challenge found in many work environments [2,3]. These alerts could also be conflicting in the sense that the information they provide suggests different actions be taken to resolve problems. Figure 1, for instance, shows an example conflict between alerting information: one system commands the operator to climb while the other commands a descent. The development of tools to formalize and better understand these types of conflict (or dissonance) issue is the focus of this paper.

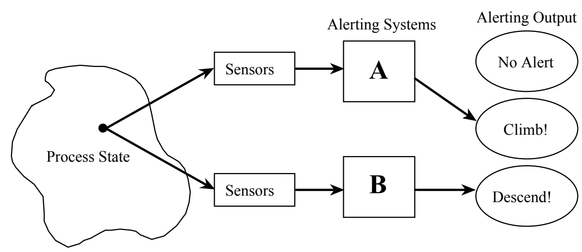

Figure 1: Schematic of an Alerting Conflict

Aviation alerting systems are one example area where these types of multiple alerting system conflicts are becoming more likely. In the era of "steam-gauge" aircraft that relied on electro-mechanical instruments (before the 1980s), nearly all alerting functions on aircraft were used to monitor autoflight controls and internal components such as engines, hydraulics, or electrical systems. One comprehensive study in 1977 found over 500 different alert displays and functions on the Boeing 747 flight deck [3]. The study also showed a history of exponential growth in the number of alerting functions on board aircraft. This trend was mitigated through the introduction of more advanced processing and electronic display technology in the 1980s. This technology allows for multifunction "glass cockpit" displays, reducing the number of separate lights and gauges, and enabling more comprehensive and integrative monitoring of systems, rather than requiring a separate display for each aircraft subsystem.

Since the 1970s, with improved sensor and communication capabilities, aircraft alerting systems have been increasingly focused on external threats such as terrain, other air traffic, and weather. Several of these external-hazard systems are now being augmented by the addition of newer, more capable alerting systems. The Ground Proximity Warning System (GPWS), for example, was mandated on U.S. transport aircraft in the mid-1970s. GPWS uses measurements of the height of the aircraft above terrain to predict whether there is a threat of an accident, and is susceptible to occasional false alarms or late alerts. In the late 1990s the Enhanced Ground Proximity Warning System (EGPWS) was introduced to provide earlier and more accurate warnings of terrain threats. EGPWS uses an on-board terrain database and includes a graphical display of the terrain field around the aircraft. Due to cost and certification issues, GPWS has been retained on aircraft and EGPWS has been added as a separate, independent system that does not change the operation of GPWS. EGPWS is a safety enhancement beyond the system that is required by law to be carried on an aircraft. The result, however, is that there are now two separate systems, each monitoring terrain threats and each with different alert threshold criteria and displays.

In the area of air traffic collision alerting, the Traffic Alert and Collision Avoidance System (TCAS) has been mandated on U.S. transport aircraft since the early 1990s. TCAS warns the pilots to an immediate collision threat and provides escape commands and guidance. Recently, other collision alerting systems have been under development to enhance safety and capability for closely-spaced approaches to parallel runways [4,5]. Specialized systems are required for parallel approach capability since TCAS was not developed with this type of operation in mind and would require major modifications to work in that environment. Thus, there could soon be two separate traffic collision warning systems on aircraft. Additionally, systems similar to TCAS but using enhanced sensor information and different, more strategic alerting criteria are being developed [6]. These systems, too, could result in multiple alerting systems monitoring traffic threats.

Multiple automobile alerting systems are also being proposed and developed, with functions including obstacle avoidance, roadway departure and lane-change warnings, and intersection collision warning systems [7]. These will also need to be carefully integrated to avoid alerting overload or conflicts.

Conflict Management

To date, conflicts between automation have been largely managed through prioritization. Each alerting system can be prioritized, and if more than one alerting system is triggered, the lower priority alerts may be inhibited or only displayed passively (i.e., without separate attention-getting signals). Several complex prioritization schemes have been investigated for the various alerting systems on board an aircraft [8,9]. Terrain, for instance, is placed at a higher priority than other air traffic, with the rationale that all else being equal, it is less likely that an aircraft would collide with another aircraft than it would hit terrain. Prioritization can run into trouble, however, if two alerts are both valid but the operator is only receiving or responding to one. Still, prioritization can help reduce sensory and cognitive overload of the human during a time of high stress.

An alternate way to mitigate the effect of alerting system conflicts is through operator training. Pilots are trained, for example, that EGPWS and GPWS use different decision-making logic, and that alerts from the two systems may not (and in fact probably will not) occur in concert. In more severe cases, however, training may fall short. For instance, two accidents of Boeing B757 aircraft in 1996 (the first near Puerto Plata, Dominican Republic, and the second near Lima, Peru) involved simultaneous, conflicting alerts in the cockpit. Both accidents were caused by clogged air data systems that resulted in alerts that the aircraft was flying too fast (from one system) and too slow (from a second, independent system). This led to significant confusion in the cockpit as to which alert to believe, and ultimately led to the accidents. Although trained procedures are in place to manage such conflicts, it is clear that training should not be relied upon too greatly.

Additionally, it may be possible to modify operation so that dissonance is unlikely. One means of trying to ensure compatibility of parallel approach alerting systems with TCAS, for example, is to modify air traffic control procedures so that the likelihood of a simultaneous TCAS alert and parallel traffic alert is very small. Finally, it may be necessary to modify the design of the logic in the new (or existing) alerting system to reduce the potential for dissonance as much as possible.

Each of these mitigation methods incurs some costs as far as overall system performance is concerned. For example, prioritization and inhibition essentially hide part of the available information from the operator. This reduces the benefit of having the additional alerting system components, since their information is not transmitted to the operator. Reducing the likelihood of dissonance by modifying the process' operation (e.g., increasing separation between vehicles to reduce alert probability) may decrease overall system efficiency and capability. Finally, modifying the design of the logic is complex, costly, and may have other negative impacts on system performance. To minimize these negative effects, mitigation strategies should only be employed where necessary. Small degrees of alerting conflict may not impact operator performance, so it is important to understand how a given degree of dissonance translates into operator performance degradation. Then, mitigating actions can be taken to address those conflicts that are most critical.

To date, management of potential dissonance between systems has occurred without a structured understanding of the specific issues involved. The identification of the potential for dissonance and the development of mitigation methods would be greatly facilitated through the application of a coherent, formal model that articulates the design issues. Such a model would have three benefits. First, it would aid in understanding the different types of dissonance that may occur. Second, the model would help in identifying when or where the different types of dissonance could occur in a given operation. Third, the model may be used to design and evaluate mitigation contingencies to prevent or preclude dissonance from occurring.

This paper presents an initial model of multiple alerting system interactions that can be used to identify and describe dissonance. Several different types of dissonance are defined, each of which may require a different mitigation approach. How conflicts affect human decision-making is also discussed using a simple human factors study. A mathematical method for analyzing dissonance situations is then presented and applied in an illustrative example.

General Alerting System Background

A significant body of research has focused on the design and use of automation, with the goal of determining how automation should be implemented to work harmoniously with the human operator [10-13]. Endsley, for example, presents arguments that the human's preconceptions and mental models have a direct effect on how automation improves or detracts from Situation Awareness (SA) [10]. Automation, then, must be carefully designed and implemented to support the human. If not properly applied, automation can degrade SA by reducing the human's involvement in monitoring and control functions.

In the late 1990s, Pritchett and Hansman explored the concepts of *consonance* and *dissonance* between an alerting system's decisions and a human operator's internal model of a threat situation [14]. Their work and observed incidents in the field have shown that a mismatch or dissonance between the human and automation could lead to undesirable behavior from the human including increased delay in taking action, failure to take action at all, or even implementing an action contrary to the automation's command. One focus of the development of alerting systems should therefore be to ensure that the information that is conveyed, the timing of alerts, and the commands or guidance provided are as much in agreement or in consonance with the human as possible. But, there certainly may be cases in which dissonance is unavoidable: for example when the human is completely unaware of a threat and so does not feel there is a problem when in fact there is. In such cases, it is important to provide corroborating information with the alert so that the human rapidly understands the rationale behind the alerting decision and so comes into consonance with the automation as quickly as possible.

We move specifically into the issues related to dissonance between two or more alerting systems. The focus here, then, is on the automation, yet it is critical to remember that ultimately it is the human's understanding and interpretation of the automation's displays that affect whether dissonance has an impact.

Alerting System Model

All alerting systems generally perform four functions, shown in Figure 2: monitoring, situation assessment, attention-getting, and problem resolution. First, information about the process under control and relevant hazard states must be monitored through a set of sensors. Each alerting system may use a different set of sensors, and thus may form a different view of what is truly occurring in the process and environment. Based on this observable information, the alerting system assesses and categorizes the situation into one of several threat levels or alert stages. If the alert stage is sufficiently high, the human operator is alerted to the problem. This attention-getting function can range from a simple aural or visual cue (e.g., a tone or illuminated light), to displays that indicate the cause for the alert (e.g., a textual or verbal readout such as "Generator Failure"), to displays that also indicate how to correct the problem. The attention-getting signal also provides an indication of the urgency of the problem. This urgency may be conveyed implicitly through the general type of hazard that is being encountered, or it may be more explicitly conveyed by the structure of the alarm signal. For example, a chime sound is

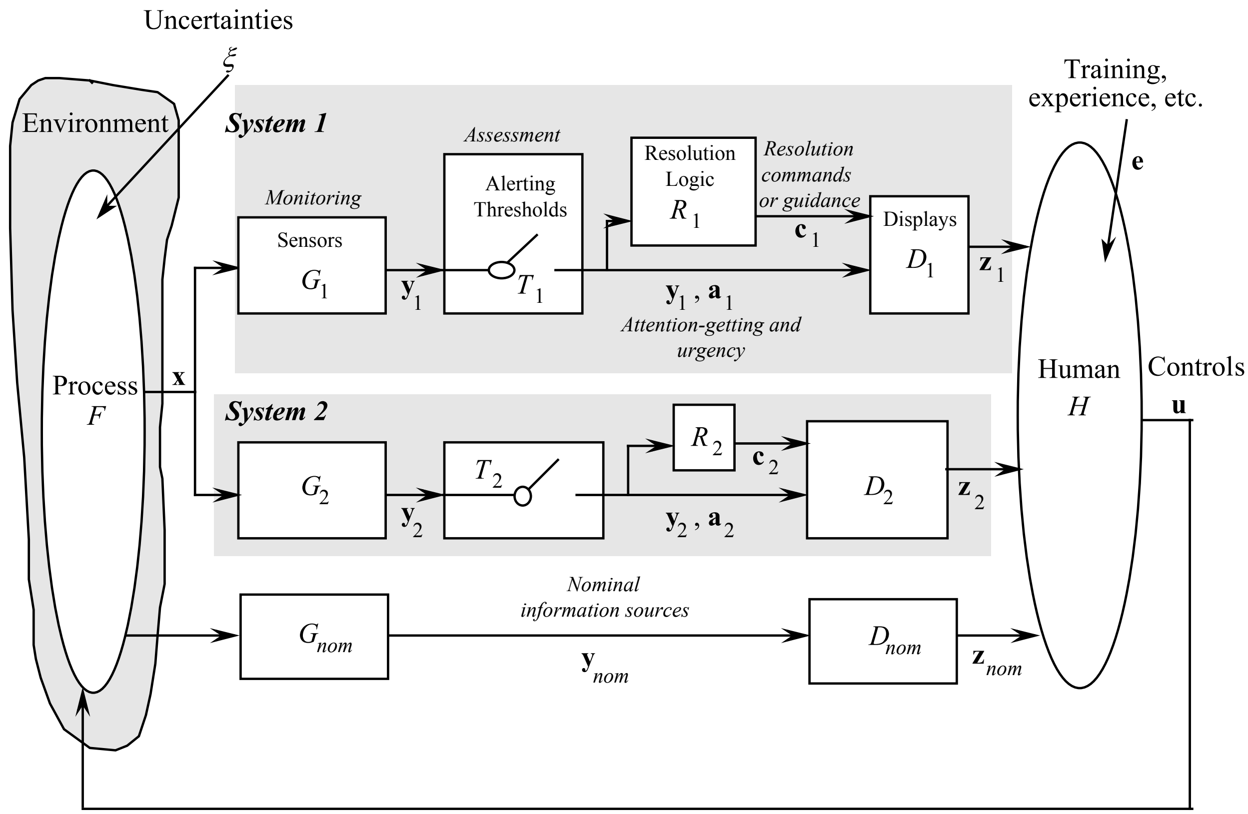

Figure 2: Block Diagram of Alerting System Elements

often used for low-urgency alerts, whereas a buzzer or wailing alarm may be used in more threatening situations [8,9].

Problem resolution may also be performed either explicitly or implicitly by the alerting system. In explicit systems, additional command or guidance information is presented to the operator. This may be a verbal message (e.g., "Climb!") and/or may include a visual display indicating the type of action to be taken and the aggressiveness with which that action should be taken. In more advanced systems, continuous guidance may be provided to aid in the resolution action. In implicit systems, the human operator may have a trained response to a particular alert stage, or may just decide at that time what action is most appropriate. Also shown in Figure 2 is the nominal information path by which the human operator obtains information about the controlled process and the environment. This information builds the human's internal model of the situation that may conflict with the conditions implied by one or both alerting systems.

From a more precise mathematical standpoint, we will denote **x** as the state vector representing the complete set of physical parameters that describe the dynamics of a hazard situation. In the case of TCAS, for example, **x** represents the three-dimensional position and velocity vectors of each aircraft involved.

In general, the complete state vector is not available to the alerting system logic, but is observed through a set of sensors. The resulting information that is observable to the alerting system is included in the vector **y**. The alerting systems use possibly different sets of observable states defined by different functions G_i operating on **x**. For the i^{th} alerting system,

$$\mathbf{y}_i = G_i(x) \tag{1}$$

For TCAS, **y** is a vector including the range, range rate, relative altitude, and relative altitude rate between two aircraft.

Alert Stages

Using the information in **y**, each alerting system applies a set of threshold functions or other logic, *T*, to map the situation into an alert stage. The alert stage is represented by the vector **a**, and specifies the level of threat according to that alerting system:

$$\mathbf{a}_i = T_i(\mathbf{y}_i) \tag{2}$$

The logic used by the alerting system to determine the appropriate alert stage and to provide guidance may vary from simple thresholds based on exceeding some fixed value to more complex algorithms involving a number of states. Many alerting systems work with two stages: non-hazardous and hazardous. More complex systems use a series of stages, each corresponding to a higher level of danger and urgency.

Alerting systems may categorize both the status of each individual hazard under observation, and also specify an overall threat level. TCAS does this, for example, by using different graphical icons depicting the threat posed by each nearby aircraft on a traffic display. Additional aural and visual displays are then used to indicate the overall threat level and whether any action is required. Thus, there may be two different types of alert stage, one for each individual hazard and one for the overall system. The *hazard alert stage* is defined as a discrete categorization of the level of threat posed by a given hazard under observation by a system. The *system alert stage* is the resultant overall level of threat posed by all the hazards under observation by that system. In TCAS, the system alert stage is equal to the maximum of all individual hazard alert stages. That is, the system as a whole takes the worst-case threat and uses its threat level. It could be desirable in other applications; however, to use a different method of translating hazard alert stages into system alert stages.

With TCAS, there are four hazard alert stages:

Stage 0 = No threat. The other aircraft is denoted by a hollow white diamond on the display.

Stage 1 = Proximate traffic. The other aircraft is shown as a filled white diamond on the display.

Stage 2 = Caution. The other aircraft is shown as a solid yellow circle.

Stage 3 = Warning. The other aircraft is shown as a solid red square.

There are three corresponding system alert stages for TCAS:

Stage 0 = No threat. No additional information is provided.

Stage 1 = Traffic Advisory (TA). A Master Caution light is illuminated in amber and an aural "Traffic, Traffic" alert is issued in the cockpit. Stage 1 is active if there is a caution hazard stage active but no active warning hazard stages.

Stage 2 = Resolution Advisory (RA). A Master Warning light is illuminated in red, an aural resolution command is issued (such as "Climb! Climb!") and the required climb angle or climb rate is shown on a cockpit display. Stage 2 is active if any hazard is in the warning stage.

Resolution Commands

Based on the alert stage and on the other information on the situation, the alerting system may produce resolution information, **c**:

$$\mathbf{c}_i = R_i(\mathbf{y}_i, \mathbf{a}_i) \tag{3}$$

The vector **c** includes the type of resolution action to be performed (e.g., turn or climb) and the magnitude of that maneuver. There are a variety of forms of resolution commands, depending on the complexity of the maneuver to be performed. Figure 3 shows three different possible styles for the same general command in which a turning-climb maneuver is desired. Figure 3a represents a case in which a specific target state is conveyed along with a single, specific trajectory to follow to achieve that end state. In Figure 3b, a target state is specified, but the means to achieve that state is not. Finally, Figure 3c shows a verbal command; the target state is only implicitly understood by the operator. Which command should be used in a given situation depends on the degree to which the automation can correctly model and predict the appropriate response. In complex, poorly-structured problems, the implicit command may be the most reasonable as it allows the human to bring to bear his or her intuition and other information to solve the problem. In well-structured problems, however, an explicit command may facilitate the human in implementing the most effective response.

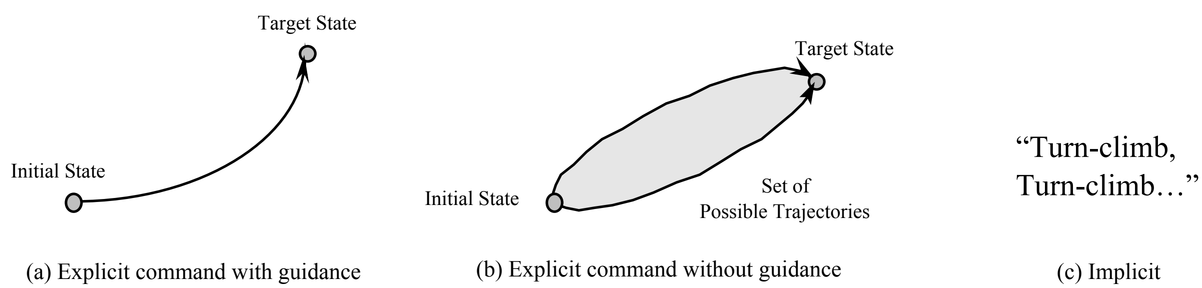

Figure 3: Example of Different Command Styles

(a) Turn-climb simultaneously (b) Turn and climb sequentially

Figure 4: Command Sequencing

Additionally, a complex command can be interpreted as either a simultaneous or sequential process. Figure 4 shows two different interpretations with “turn climb” as the example. Figure 4a describes a simultaneous turning-climb path, while Figure 4b shows a sequential case: first turn, then climb. Given all the possible combinations of alert stages and command types, it is clear that there is a rich design space for alerting systems. As a consequence, it is possible (and even probable) that two different alerting systems will apply different alert stage or command definitions to a similar problem. This may lead to dissonance as is discussed in a later section.

Complete Set of Transmitted Information

Referring back to Figure 2, the vector $\mathbf{z}$ combines all the information that is displayed to the human operator by the alerting system. In general, $\mathbf{z}$ includes signals designed to attract the operator’s attention, the alert stage, and information to resolve the situation. The function D describes the display mapping from the state estimates available to the alerting system ($\mathbf{y}$) to the information provided to the human operator ($\mathbf{z}$) based on the alert stage ($\mathbf{a}$) and resolution information ($\mathbf{c}$):

$$\mathbf{z}_i = D_i(\mathbf{y}_i, \mathbf{a}_i, \mathbf{c}_i) \qquad (4)$$

For TCAS, the information in $\mathbf{z}$ includes a traffic display in the cockpit, aural messages, lights, and any resolution command and guidance information.

In addition to the alerting systems, there may be other, nominal information sources that provide information to the operator. This information is included in the vector $\mathbf{y}_{nom}$, which is then modified by the nominal displays D_{nom} as shown in Figure 2. Cockpit instruments, air traffic control communications, views through the windscreen, and aeronautical charts are examples of nominal information sources for a pilot. The operator is also affected by other factors such as the pilot’s internal model of the situation, knowledge of the alerting system’s role, prior training, fatigue, and previous experience. Past exposure to false alarms, for instance, has been observed to be a factor in delaying responses to alerts [15]. This modifying information is included in the vector $\mathbf{e}$. The function H then maps the observable states (via all the alerting systems and nominal information sources) to the control inputs $\mathbf{u}$. That is,

$$\mathbf{u} = H(\mathbf{z}_{nom}, \mathbf{e}, \mathbf{z}_1, \mathbf{z}_2) \qquad (5)$$

Ultimately, it is how the inputs to the pilot (as contained in $\mathbf{z}_{nom}$, $\mathbf{z}_1$, $\mathbf{z}_2$, and $\mathbf{e}$) are used to develop a control strategy that determines whether there is a conflict between the information elements being used. In this context, Pritchett and Hansman’s work examined dissonance between $\mathbf{z}_1$ for a single alerting system and the other information provided to the human in $\mathbf{z}_{nom}$. Here, we focus on dissonance across the information provided by two different alerting systems, as contained in $\mathbf{z}_1$ and $\mathbf{z}_2$.

To complete the control block diagram, the process’ state derivatives are determined from a generalized function, F, of the current state, operator’s inputs, and modeling or system dynamics uncertainties, ξ:

$$\dot{\mathbf{x}} = F(\mathbf{x},\mathbf{u},\xi) \qquad (6)$$

Model Of Multiple Alerting System Conflicts

Having introduced a general state-space model for multiple alerting systems, it is now possible to more formally state the types of dissonance that may occur. At a high level, all alerting systems can be thought of as mapping a set of measured or estimated states of a controlled process into discrete alert stages and discrete or continuous hazard resolution commands. Dissonance may occur whenever a given state maps into two different alert stages or two different resolution commands, or when the time-derivatives of these mappings differ.

Table 1: Alerting System Conflict Types

Conflict Type		Example Dissonant Situation	
		System 1	*System 2*
Alert Stage	system alert stage	no threat	warning
	hazard alert stage	aircraft A is a threat	aircraft B is a threat
Resolution	dimension	turn	climb
	polarity	climb	descend
	magnitude	turn 5°	turn 30°

Table 1 provides a listing of different forms of dissonance. Each row in Table 1 corresponds to a type of dissonance with certain properties. The right side of the table provides an example situation with two alerting systems in which that category of dissonance is present. For example, having one system command "climb" while a second system commanded "descend" would be a resolution polarity conflict. Each of these forms of dissonance is discussed in more detail below. It is also worthwhile dividing dissonance into *static* and *dynamic* forms, each of which are defined in the following sections.

Static Dissonance

When $\mathbf{z}_1 \neq \mathbf{z}_2$ at a given time for two alerting systems, static dissonance may exist. Breaking $\mathbf{z}$ into its components, first consider alert stage conflicts. Differences in system alert stage (first row of Table 1) can be present without causing dissonance if the two alerting systems have different roles. For example, EGPWS is designed to provide an earlier warning of terrain proximity than GPWS. Should this happen, there is probably no dissonance from the pilot's point of view, even though the alert stage from EGPWS is at a higher level than that from GPWS. If the opposite occurred, however, there may be dissonance because the pilot may not understand why EGPWS does not rate the terrain as a threat while GPWS does.

Another type of dissonance can occur when there is a difference in the hazard alert stage for a given threat, even if the system alert stages are consistent (second row of Table 1). This could happen, for example, in a case with two traffic alerting systems monitoring two different aircraft. If system 1 rates aircraft A as a low threat and aircraft B as a high threat while system 2 does the opposite, then both systems may agree with the same high-threat system alert stage, but the underlying hazard alert stages for each threat are different. The operator then may distrust one or both systems since they are disagreeing on the cause for the system alert stage.

Dissonance can also occur due to the resolution information contained in $\mathbf{z}$. Recalling Figures 3 and 4, the resolution information can be thought of as trajectories of varying levels of abstraction that are intended to direct the human operator to a safe target state. If two trajectories are in different dimensions, then there may be dissonance (e.g., a case where system 1 commands a change in altitude but system 2 commands a change in heading). If two commands are in the same dimension, then dissonance may still be present due to different polarities or magnitudes of the commands. If two systems are both commanding a change in altitude, but system 1 commands a climb and system 2 commands a descent, there is clearly a conflict. Or, if system 1 commands a much stronger climb than system 2, there may be dissonance.

Given the wide variety of commands that can be issued as illustrated in Figures 3 and 4, there may be subtleties in the commands that affect whether certain differences are perceived to be dissonant or not. The general concept, however, is that the resolution trajectories implied by the command (whether implicit or explicit) should not be disjoint; otherwise, dissonance is likely.

Dynamic Dissonance

The above types of dissonance are static: they exist at a given point in time. Since the situation is constantly changing, however, dynamic dissonance may also occur. In dynamic dissonance, it is the *change* in alert stage or *change* in resolution information over time that produces a conflict; that is, when $\dot{\mathbf{z}}_1 \neq \dot{\mathbf{z}}_2$. Consider two collision alerting systems, where one system initially indicates no threat while the second system indicates a high degree of danger and a warning is issued. This is static dissonance. However, if the first system upgrades the alert stage to a caution while the second system downgrades the alert stage, also to a caution, dynamic dissonance exists. Even though the two systems now agree about the proper alert stage, the human may be uncertain as to whether the situation is improving or getting worse due to the dynamic dissonance. Dynamic dissonance may also occur when the magnitude or direction of a resolution command changes. Two climb commands, but one which is increasing in magnitude while the other is decreasing in magnitude could lead to a similar type of confusion.

Implications on the Human Operator

The discussion to this point has focused on the specific, mathematical conditions that can lead to dissonance. But, even when there is a mathematical difference between $\mathbf{z}_1$ and $\mathbf{z}_2$ or their derivatives, there may not be a perceived dissonance from the human operator's standpoint. If the difference in alerting system information is still consistent with the operator's other knowledge about the process and environment, a conflict will probably not be perceived. If the operator clearly understands the rationale behind the alerting system behavior, then it is possible that a conflict will not result in confusion over which action to take. Such an understanding can be further developed through training. Still, in high-pressure situations, even small conflicts may grow into significant confusion, delay, or implementing an incorrect response.

One critical consideration of dynamic dissonance is that its impact may depend on how rapidly the changes in alert information occurs. Recall the earlier example where one system initially indicated a high degree of danger and a second system indicated no threat. If both systems change to a moderate-caution level simultaneously, it is likely there would be a stronger perceived dissonance than if one system changed to caution followed by a significant delay before the second system also indicated caution.

It is important to gain a better understanding of how differences between the information conveyed to the human ultimately translate into perceived dissonance, and then how that dissonance translates into human performance. There has been little work in the past in the specific area of alerting systems, so more research on these human factors aspects is warranted.

Preliminary Human Factors Evaluation

As an initial step, a basic experiment was carried out to compare the confidence of a human operator when faced with different alert stage conflict situations. The hypothesis was that the larger the difference in alerting signals, the poorer the decision performance and the less confidence there would be in that decision. Lower confidence may cause a delayed or incorrect response and increased workload as the human tries to determine which system is correct.

In this preliminary study, human subjects observed the static states of two alerting systems as indicated by two lights on a computer screen. The subjects were told that each alerting system monitored the same threat, but no other specific context was provided. Each system was in one of three alert stages as indicated by the color of the indicator light on the computer screen. State 0 corresponded to a no alert case (green light), state 1 indicated a caution (yellow light), and state 2 indicated a warning (red light). The subjects were told that each system operated independently and produced a correct decision 90% of the time. Additionally, a non-symmetric payoff matrix was provided to the subjects that weighed the cost of a missed detection more than a false alarm, and the gain from a correct rejection as more than that from a correct detection.

The subject's task was to observe the alert condition for five seconds and then indicate his or her estimate of the true condition (0, 1, or 2) with the intent to minimize the cost incurred from the payoff matrix. Additionally, each subject then indicated his or her confidence in that threat level estimate using a 100-point percentage scale (from "not at all confident" to "completely confident").

Seven graduate student subjects participated in the study. Nine different alert stage conflict situations were shown to each subject, in a random order. These situations are described here using a notation such that (0, 1), for example, indicates that system A was in alert stage 0 while system B was in alert stage 1. The set of nine situations was then: (0, 0); (1, 1); (2, 2); (0, 1); (1, 0); (1, 2); (2, 1); (0, 2); and (2, 0). In the results, symmetric situations were found to be statistically similar. That is, there was no distinguishable difference between conditions (0, 1) and (1, 0). The same can be said for condition pairs (1, 2) - (2, 1), and (0, 2) - (2, 0). For simplicity, only one of each pair will be used in the descriptions below; (0, 1) therefore represents also the (1, 0) case.

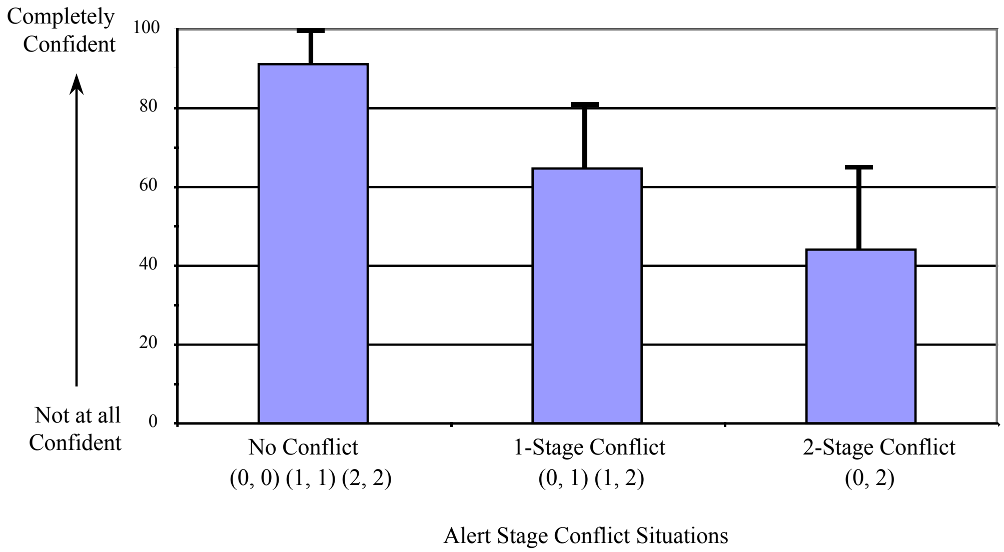

Figure 5: Confidence in Alerting

Using Bayes' rule, the payoff matrix, and the assumption that conditions 0, 1, and 2 were equally-likely *a priori*, the optimal decision (0, 1, or 2) can be determined to maximize the expected payoff for each situation. For example, for case (1, 2), the best strategy was to believe that the true condition was warning (2). From the subjects' estimates of the true threat levels, when the two systems matched [cases (0, 0), (1, 1) and (2, 2)] they chose the optimal strategy 95% of the time. In the conflict cases [(0, 1), (1, 2) and (0, 2)], the optimal strategy was chosen only 71% of the time. There was no statistical difference in this decision performance between (0, 1), (1, 2) and (0, 2) cases.

Figure 5 shows the confidence rating results, grouped into three categories: no conflict [cases (0, 0), (1, 1), and (2, 2)], 1-stage conflicts [cases (0, 1) and (1, 2)] and 2-stage conflicts [case (0, 2)]. In the no-conflict cases, the operators had high confidence (mean of approximately 90%). Mean confidence in the 1-stage conflict cases dropped to approximately 60%, and to approximately 40% in the 2-stage conflicts. These differences are all statistically significant ($p < 0.05$) using a paired t-test. Within each grouping in Figure 5, there was no statistically-significant difference. That is, confidence in the (0, 1) case was similar to the (1, 2) case.

These observations support the hypothesis that a conflict between alerting systems can reduce confidence in the decisions made by an operator. Granted, this study was quite simple, but it still serves as a demonstration of the general concept linking dissonance to decision quality.

A number of more focused and complete studies are required to develop a better understanding of the potentially complex interactions between alert dissonance and the resulting effect on human performance. In particular, it will be necessary to examine human behavior in the presence of dissonance in more realistic alerting settings. One key area for study is the connection between dynamic characteristics of alert conflicts and the perception of dissonance.

State-Space Model for the Analysis of Dissonance

The preceding sections developed concepts of static and dynamic dissonance by examining the similarities and differences between the information passed to the human operator in **z** and how that conflict may translate into an effect on human performance. An additional step is to formulate a means of identifying how these conflicts originate. By exposing those situations that lead to dissonance, the system design can be modified, operations can be changed, or the operators can be trained to work around the dissonance.

To expose those conditions where dissonance may occur, we begin by examining the state space of the alerting system and observing when alerts are issued. Consider a simplified one-dimensional problem in which the in-trail separation of two vehicles is monitored by two independent alerting systems placed in the trailing vehicle. The leading vehicle follows some path open-loop, while the trailing vehicle may receive alerts to speed up or slow down to maintain spacing. As a baseline, assume that alerting system 1 is set up to issue an alert if the two vehicles get too far apart. An alert from system 1 would command the trailing operator to increase speed to reduce the separation between vehicles. System 2 is set up to alert if the vehicles are projected to be too close within some amount of time, or if the vehicles are very close together and not diverging fast enough. An alert from system 2 would command the trailing operator to reduce speed and increase separation. We are then interested in identifying when dissonance between the two systems could occur.

In this case, the positions and velocities of the two vehicles make up the state space:

$$\mathbf{x} = [x_0, x_1, \mathrm{v}_0, \mathrm{v}_1]^{\mathrm{T}} \tag{7}$$

System 1 measures only the range between the vehicles, while system 2 uses both the range and range rate. So,

$$\mathbf{y}_1 = [r] = [x_1 - x_0] = G_1\mathbf{x}$$

$$\mathbf{y}_2 = [r,\ \dot{r}\,]^{\mathrm{T}} = [x_1 - x_0,\ \mathrm{v}_1 - \mathrm{v}_0]^{\mathrm{T}} = G_2\mathbf{x} \tag{8}$$

This example has a simple, binary alert stage for each system: 0 or 1. System 1 alerts ($\mathbf{a}_1 = 1$) when the range between vehicles is greater than a threshold distance R_1. In the notation we have developed, predicates (or inequalities) denoted f_{ij} are defined to divide the state space into subsets. When the state is inside the subset, the predicate is true; when outside, the predicate is false. Combinations of these subsets then form the alert stage space within the universe of the state space, U. Each resulting subset is denoted A_{ik} for the k^{th} alert stage of system i. So, for system 1, an alert occurs when the state is in region A_{11}. This is formally described by:

$$T_1 = \begin{cases} f_{11} : r > R_1 \\ A_{11} = f_{11} \\ A_{10} = U - A_{11} \end{cases} \tag{9}$$

According to Eq. (9), an alert is issued (state is in A_{11}) when condition f_{11} is true; this is equivalent to $r > R_1$. Otherwise, the state is in region A_{10}, which indicates that system 1 is in alert stage 0.

System 2 alerts ($\mathbf{a}_2 = 1$) when the vehicles are converging and projected to be less than a range R_2 apart within τ seconds, or if they are close together and diverging but at a slow rate ($r\dot{r} < K$, where K is some constant). This is similar to the logic used by TCAS [16], and can be described by the following threshold function:

$$T_2 = \begin{cases} f_{21} : \dot{r} < 0 \\ f_{22} : \dfrac{r - R_2}{-\dot{r}} < \tau \\ f_{23} : r\dot{r} < K \\ f_{24} : r < R_2 \\ A_{21} = (f_{21} \cap f_{22}) \cup (f_{23} \cap f_{24}) \\ A_{20} = U - A_{21} \end{cases} \tag{10}$$

Figure 6 shows the two alerting systems' alert spaces in the two-dimensional $\mathbf{y}$ space of r and $\dot{r}$. A "+" has been added to the active alert stage in the diagram for system 1 to emphasize that an alert from that system commands the trailing operator to accelerate and reduce spacing between vehicles. A "0" implies that no command or guidance information is displayed by the alerting system. A "–" is used to show where a command to the trailing vehicle to decelerate (increase spacing) would be given by system 2.

Static Dissonance

Having set up the basic alert stage regions in state space, we can analyze the two systems together as shown in Figure 7. When the two systems are combined, the intersections of their alert stages are denoted by the sets S_{mn} where m is the alert stage from system 1 and n is the alert stage from system 2:

$$S_{mn} = A_{1m} \cap A_{2n} \tag{11}$$

There are four possible combinations of alert spaces between the two systems: $S_{00} = A_{10} \cap A_{20}$, $S_{01} = A_{10} \cap A_{21}$, $S_{10} = A_{11} \cap A_{20}$, and $S_{11} = A_{11} \cap A_{21}$. To help identify potential dissonance, the "+", "–", or "0" notations from Figure 6 have been carried through. The notation, "+0", for example, indicates that system 1 commands an acceleration while system 2 does not display any command information.

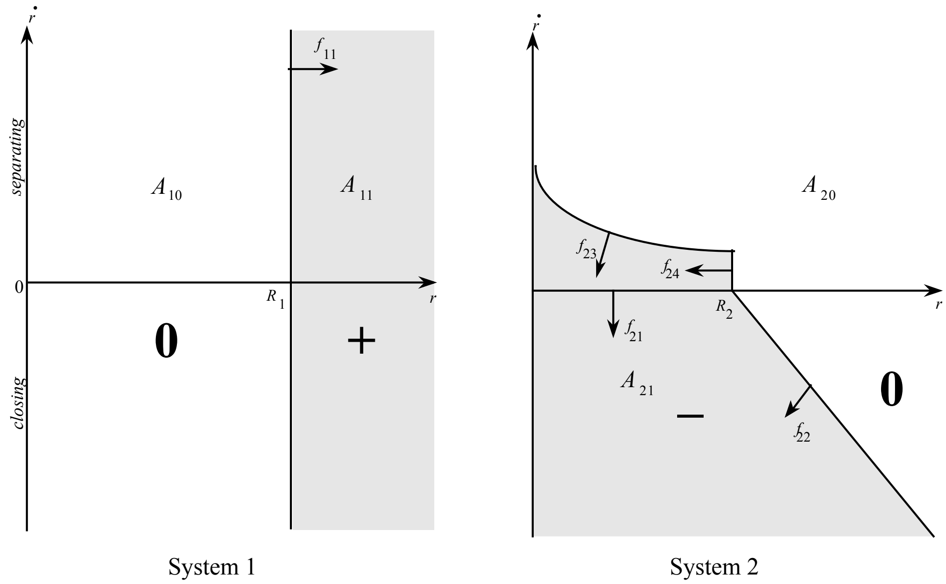

Figure 6: Example In-Trail Separation Alert Stage Mapping

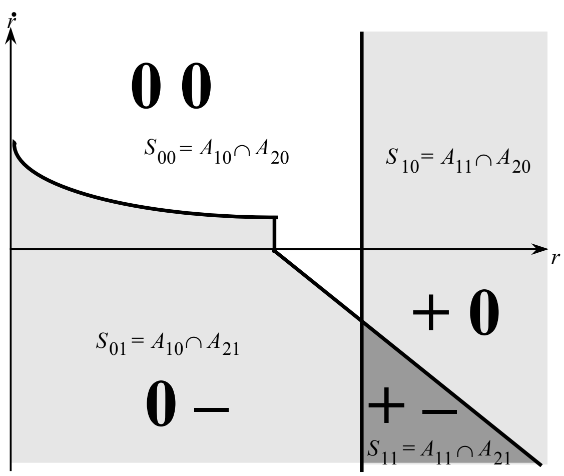

Figure 7: Combined In-Trail Alert Stages

To better visualize the potential conflicts, consider Figure 8, which shows the one-dimensional space of potential acceleration actions by the trailing vehicle for each alerting condition. Assume there is some limit on the potential acceleration of the vehicle, *a*. If System 1 is not alerting, then the operator is conceivably allowed to apply any acceleration he or she may desire within that acceleration limit. Thus, stage A_{10} can be thought of mapping to the action space [-*a* *a*]. If System 1 does alert, then the operator should accelerate the trailing vehicle. This corresponds to the action space (0 *a*]. Note that 0 is not included in the action space since some positive response is required following the alert. Similar mappings can be made for System 2. System 2 has the same action space as System 1 if there is no alert. However, an alert from System 2 commands the trailing vehicle to decelerate, corresponding to action space [-*a* 0).

With this notation, then, it is possible to observe static conflict situations. For example, S_{11} is a dissonant region because the intersection of the two systems' action spaces { (0 *a*] and [-*a* 0) } is empty.

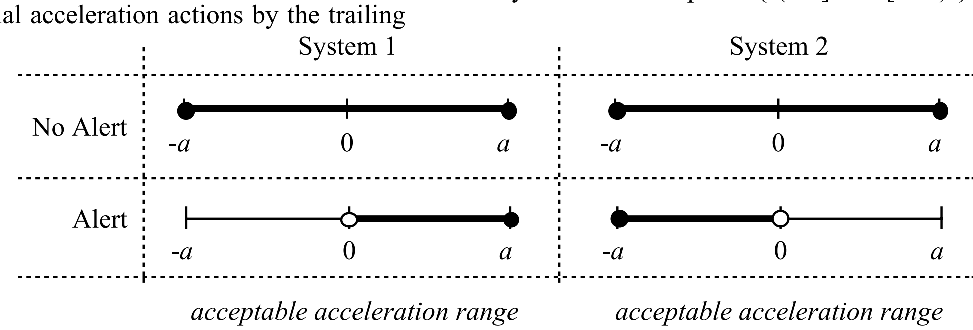

Figure 8: Action Spaces for Alerting Situations

That is, the two systems are issuing contradictory resolution commands (one to accelerate, the other to decelerate). Regions S_{01} and S_{10} would probably not result in dissonance, because the intersection of their action spaces is not empty. Although there is a disagreement in alert stages in S_{01} and S_{10}, the two systems have different roles and so would not be expected to operate simultaneously. So, "+0" or "0–" conditions would likely be acceptable.

The "+ –" dissonance in region S_{11} could be quite problematic. This corresponds to a case in which the vehicles are rather far apart but closing rapidly. The operator receives one alert to reduce spacing (from system 1) while system 2 is simultaneously commanding the operator to increase spacing. Depending on the relative strengths with which these commands are issued, the operator may be uncertain as to the correct action to take. The potential for dissonance could be reduced by either operationally preventing the two vehicles from entering S_{11} or by modifying one or both systems' decision thresholds to reduce the size of S_{11}. With reference to Figure 7, the second mitigation option could be implemented by modifying system 1, for example, such that its threshold does not apply if the closure rate is large. This is effectively prioritizing system 2 with a higher priority than system 1 whenever both systems would otherwise be triggered.

Dynamic Analysis

The analysis above for static dissonance does not completely describe the interactions between the two systems. It is also necessary to examine the process dynamics to see how dissonance may evolve over time. In this case, there is a specific physical coupling between the range and range-rate states, meaning that only certain trajectories are possible. Specifically, it is impossible to enter region S_{11} from the left; by definition, the negative range rate indicates that the range must be decreasing. So, the only way in which dissonance can occur is for the range to be decreasing at a large rate while in region S_{10}. In a specific problem, the possible trajectories in the S_{mn} diagram can be examined to determine whether it is possible to have the large range and closure-rates needed to enter region S_{11}.

As an example dynamic analysis, Figure 9 overlays several potential state trajectories on the state space diagram. Assume that the process dynamics are such that the relative speed between vehicles can be increased or decreased by an acceleration of no more than a certain amount. Starting at the state denoted **A** in Figure 9, for example, the future state trajectory must lie somewhere between the parabolic curves shown.

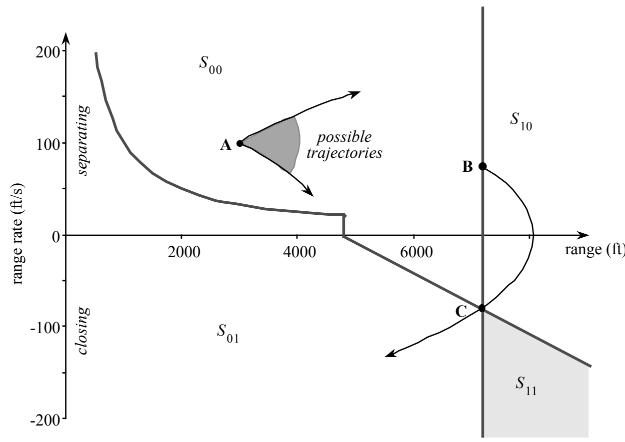

Figure 9: Dynamic Trajectory Analysis

Consider now starting at state **B**. Here, the vehicles are diverging and the state has just entered region S_{10}. The trailing operator receives an alert from system 1 to speed up and decrease spacing. If the vehicle is accelerated at its limit, the state will follow the trajectory shown, just crossing past region S_{11} through point **C**. If a lower magnitude of acceleration were used, the trajectory would lie to the right of that shown and could therefore enter region S_{11}.

The transition from region S_{10} to S_{01} that occurs at point **C** may initially appear to be a case of dynamic dissonance. As Figure 10 shows, however, when transitioning from S_{10} to S_{01}, there is a similar directionality of movement in the action spaces from each system. This implies that such a transition may not be dissonant since the operator will have a consistent change in the acceleration level to apply.

Similar dynamic analyses could be performed under different conditions and assumptions. The general approach, however, is one in which potential paths through the different alerting regions can be explored. This identifies what conditions may lead to static or dynamic dissonance. Additional effort can then be focused on those conditions to determine how likely they are, the impact of the dissonance, and to develop countermeasures to reduce the effect of the dissonance on operator performance.

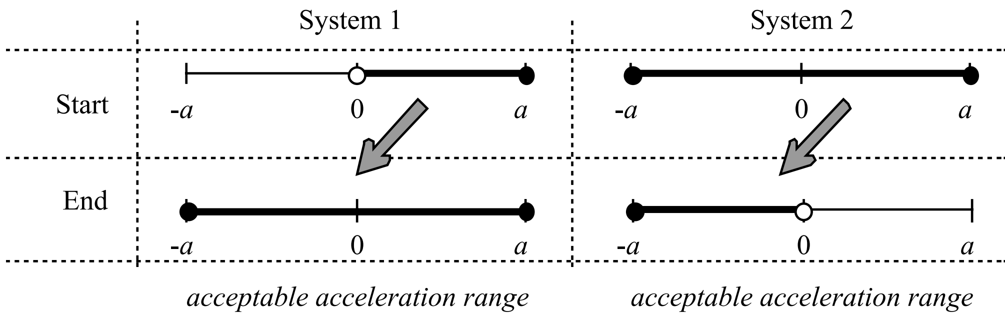

Figure 10: Dynamic Changes in Alerting Actions From S_{10} to S_{01}

Concluding Remarks

Alert system dissonance has not been a major concern in the past beyond the desire to minimize simultaneous alerts and prevent information overload. Conflicting alert information is likely to become more prevalent, however, as alerting systems continue to be injected into complex system operations. Several areas in aerospace have already been identified where dissonance is likely to occur if this issue is ignored, and certainly there are other regimes where similar problems are of concern.

To date, management of dissonance between systems has mainly involved inhibition of alerts, and has typically occurred without a structured understanding of the specific issues involved. This paper presents a more formal model that has several objectives. First, it aids in understanding the different types of dissonance that may occur. This will be useful in building a common terminology with which to compare and discuss alerting system conflicts. Second, by using state-space techniques, the model can be used to identify when or where each different type of dissonance could occur in a given operation. Ultimately, the model may be used to design and evaluate more advanced mitigation contingencies to prevent or impede dissonance from occurring. This may include modifying training and procedures to prevent operation in a dissonant region or changes to the logic used in new or existing alerting systems.

The use of a state-space representation provides a generic framework that facilitates articulating the specific information elements that are sensed, processed, and displayed by an alerting system. Based on this framework, formal definitions of static and dynamic dissonance are presented, including how this dissonance is connected to differences in alert stage or resolution command information. By drawing the mapping between process states and the resulting alert stages and resolution commands, it is then possible to identify conditions that lead to dissonance. Dynamic analyses can also be performed to expose time-varying situations that may exacerbate a problem.

A critical aspect of alerting dissonance is the impact of conflicting information on the human's situation awareness and decision-making processes. As discussed here, this impact depends on the specific application, situation, and human operator characteristics, and so it is difficult to develop a general model of human behavior at this time. As a start, we have focused on describing how conflicting information between alerting systems arises and changes with time, and very coarsely how dissonance may affect decision confidence. It will be important to examine how a conflict in information ultimately translates into human performance problems. The model presented here is therefore one preliminary piece of this larger set of work that is required, but does provide a foundation on which to build future efforts.

Acknowledgment

This research was funded by the NASA Ames Research Center under grant NAG2-1425.

References

[1] A. Pritchett, B. Vandor, and K. Edwards, "Testing and Implementing Cockpit Alerting Systems", *Reliability Engineering and System Safety*, in press.

[2] K. Momtahan, R. Tetu, and B. Tansley, "Audibility and Identification of Auditory Alarms in the Operating Room and Intensive Care Unit", *Ergonomics*, Vol. 36, No. 10, pp. 1159-1176, 1993.

[3] I. Carrick, "Incorporating Human Factors into Safety Systems in Scottish Nuclear Reactors", Proceedings of the Workshop on Human Error and Systems Development, pp. 27-33, Glasgow, Scotland, March, 1997.

[3] J. Veitengruber, G. Boucek, and W. Smith, "Aircraft Alerting Systems Criteria Study, Vol. 1: Collation and Analysis of Aircraft Alerting Systems Data", FAA-RD-76-222, National Technical Information Service, Springfield, VA, May, 1977.

[4] M. Waller and C. Scanlon, eds., "Proceedings of the NASA Workshop on Flight Deck Centered Parallel Runway Approaches in Instrument Meteorological Conditions", NASA Conference Publication 10191, Hampton, VA, December, 1996.
[5] J. K. Kuchar and B. D. Carpenter, "Airborne Collision Alerting Logic for Closely-Spaced Parallel Approach", *Air Traffic Control Quarterly*, Vol. 5, No. 2, pp. 111-127, 1997.
[6] W. E. Kelly, "Conflict Detection and Alerting for Separation Assurance Systems", 18th Digital Avionics Systems Conference, pp. 6.D.1-1 – 6.D.1-8, St. Louis, MO, October, 1999.
[7] W. Najm, J. Koopmann, and D. Smith, "Analysis of Crossing Path Crash Countermeasure Systems", Paper No. 378, 17th International Technical Conference on the Enhanced Safety of Vehicles, June, 2001, Amsterdam, Netherlands.
[8] G. Boucek, D. Po-Chedley, B. Berson, D. Hanson, M. Leffler, and R. White, "Aircraft Alerting Systems Standardization Study, Vol. 1: Candidate System Validation and Time-Critical Display Evaluation", DOT/FAA/RD-81/38/I, National Technical Information Service, Springfield, VA, January, 1981
[9] B. Berson, D. Po-Chedley, G. Boucek, D. Hanson, M. Leffler, and H. Wasson, "Aircraft Alerting Systems Standardization Study, Vol. 2: Aircraft Alerting System Design Guidelines", DOT/FAA/RD-81/38/II, National Technical Information Service, Springfield, VA, January, 1981.
[10] M. R. Endsley, "Toward a Theory of Situation Awareness in Dynamic Systems", *Human Factors*, Vol. 37, No. 1, pp. 32-64, 1995.
[11] T. B. Sheridan, *Telerobotics, Automation, and Human Supervisory Control*, MIT Press, Cambridge, MA, 1992.
[12] C. D. Wickens, *Engineering Psychology and Human Performance*, Harper Collins, 1992.
[13] N. Sarter and D. Woods, "How in the World Did We Ever Get Into That Mode? Mode Error and Awareness in Supervisory Control", *Human Factors,* Vol 37, No. 1, pp. 5-19, 1995.
[14] A. R. Pritchett & R. J. Hansman, "Pilot Non-Conformance to Alerting System Commands During Closely Spaced Parallel Approaches", MIT Aeronautical Systems Laboratory Report, ASL-97-2, Cambridge, MA, January, 1997.
[15] J. L. DeCelles, "Delayed Response to GPWS Warnings", *Air Line Pilot*, Vol. 61, No. 1, January, 1992.
[16] Radio Technical Committee on Aeronautics (RTCA), "Minimum Performance Specifications for TCAS Airborne Equipment", Document No. RTCA/DO-185, Washington, D.C., September, 1983.