Presented at the AIAA Guidance, Navigation, and Control Conference, Portland, OR, 8/10/99
AIAA-99-4145

CALCULATING DEPENDENT SURVEILLANCE UPDATE RATES BY MODELING THE TIME-DEPENDENCE OF INFORMATION VALUE

Stephen Atkins*
NASA-Ames Research Center, Moffett Field, CA 94035-1000

R. John Hansman, Jr.†
Massachusetts Institute of Technology, Cambridge, MA 02139-4307

Abstract

The time-dependence of information has historically been managed though manual procedures, typically developed heuristically, specifying how frequently or under what conditions aging information should be refreshed. In the design of Automatic Dependent Surveillance (ADS) applications, when aircraft should automatically broadcast state information (e.g., position, velocity, or intent) must be addressed more rigorously. To answer questions about how frequently or under what conditions aging information should be updated, when the cost of seeking information must be accounted, a novel model of time-dependent information value is developed for a class of proceduralized decision problems, common in aviation. The model combines elements of classic, time-invariant information value theory, which is incapable of modeling the relationship between information value and information age, with estimation techniques. The time-dependent information value model is applied to investigate the rate at which aircraft should broadcast state information in an ADS environment. Results are presented for aircraft flying on parallel trajectories as well as on crossing trajectories. Results also provide general insight into the time-dependence of information value and its implications on information management processes.

Introduction

A variety of new automation systems, made possible by datalink communication and other new technologies, will require the measurement, communication, and display of time-varying information. Historically, the time-dependence of information has been managed though manual procedures, typically developed heuristically, specifying how frequently or under what conditions aging information should be refreshed. Automatic dependent surveillance (ADS) - aircraft reporting state measurements or intent information to controllers or other aircraft via datalink - has been proposed as an enabling technology for reducing separation requirements and increasing operational flexibility (e.g., free flight) in both domestic and oceanic airspace. The cost for using datalink communication for frequent updates necessitates a more rigorous methodology for designing when aircraft should automatically broadcast state information (e.g., position, velocity, or intent) in an ADS environment. This paper proposes a formal measure of information value as a quantitative foundation for making decisions concerning when new information should be sought.

However, existing time-invariant theories of information value are incapable of modeling the relationship between information value and information age. Atkins [1] reviews a selection of relevant literature. To answer questions about how frequently or under what conditions aging information should be updated, when the cost of seeking information must be accounted, a novel model of time-dependent information value is developed for proceduralized decision problems, common in aviation. The model combines elements of classic, time-invariant information value theory with estimation techniques to model the effect of time. Proceduralized decision problems are characterized by an established procedure or rule that specifies the correct decision or action as a function of a set of relevant state variables. The decision maker's task, therefore, is to estimate these state variables; poor state estimates may result in the decision maker choosing a decision that is incorrect according to the procedure. Information is valued by its effect on the decision maker's ability to estimate the relevant state variables, measured in the context of the proceduralized decision problem.

*Research Engineer.
† Professor, Aeronautics and Astronautics.

The time-dependent information value model is used to investigate when aircraft should broadcast state information in an Automatic Dependent Surveillance (ADS) environment, for two encounter geometries. The model is used to determine the optimal number and timing of measurements when two aircraft are on crossing trajectories, and the optimal periodic update rate when two aircraft are on parallel trajectories.

Time-Dependent Information Value

This section outlines a novel model for describing the time-dependence of information value. Atkins [1] presents a more thorough development.

Proceduralized Decision Problems

This paper considers *proceduralized decision problems*, a class prevalent in aviation, in which an established procedure or rule specifies the correct decision or action as a function of one or more relevant state variables. Decision problems for which a pilot is required to obey a Federal Aviation Administration (FAA) regulation or airline operating procedure belong to this class.

The procedure specifies, for each point in the space of relevant state variables, which of a set of possible actions is correct. The number of actions between which the decision maker must choose is typically small and, therefore, each of the possible actions is appropriate for a set of points in the state space. The points in the state space for which the procedure calls for a particular action comprise a *region*. *Threshold surfaces* are the boundaries around these regions. Note that the correct action does not depend on where the state vector lies within a region.

The decision maker's task in a proceduralized decision problem is to identify the region in which the state lies (or will lie at a future time of interest). This paper assumes that, if the state variables are known accurately, all decision makers will choose the action called for by the procedure. However, uncertainty with respect to the state variables may cause the decision maker to choose an incorrect action.

Separation of Estimation and Context

The solution of a proceduralized decision problem is found by solving an estimation problem - by identifying the region of the state space in which the state vector lies. To measure the value of new information to a proceduralized decision problem requires first understanding how the new information will affect the solution to the estimation problem and, second, how that change in the knowledge about the state will affect the decision outcome.

Figure 1 models the steps in solving a proceduralized decision problem. The process of selecting an action a, given the information I, is separated into two cascaded steps: constructing a model $f_{\mathbf{x}|I}$ of the relevant states from the information, and using that model to select an action. Note that Shannon's [2] information theory and estimation theory [3] are concerned with the sensitivity of the state model $f_{\mathbf{x}|I}$ to the information I.

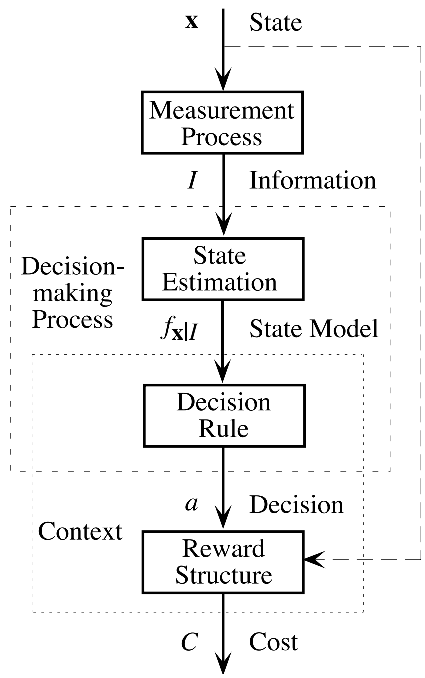

Figure 1. Separation of proceduralized decision problems into cascaded state estimation and decision selection.

Classic information value theory [4, 5] measures the impact of receiving information directly on the cost that results from a decision. By measuring the sensitivity of the cost C to the information I, the classic definition of information value obscures the estimation process, which is central to proceduralized decision-making. Therefore a different metric is introduced to more directly measure the impact of information on the goal of improving the decision maker's knowledge about the relevant state variables.

Information value is interpreted as the effect of information on the ability to model the relevant state variables (i.e., the sensitivity of the model $f_{\mathbf{x}|I}$ to the information I), *measured in the context of the proceduralized decision problem.* The value which a piece of information has to the decision maker is defined as the change in the probability that the decision maker will choose an incorrect decision, weighted by the difference in the costs for the correct and incorrect decisions. Note that, since the decision maker is constrained to operate within the procedure, if the procedure is bad (i.e., non-optimal in a minimum cost sense), classic information value theory may assign negative value to information, even though the information reveals the true state of the world. The current approach will value this information in the context of the bad procedure (i.e., how the information helps the decision maker choose the decision prescribed by the procedure).

The next section introduces a general model of the decision maker's knowledge about a set of state variables, in relation to the threshold surfaces, including the effects of time and new information on this knowledge. The following sections model the *context* within which information is valued, and apply these ideas to define the information value metric.

Uncertainty and Uncertainty Significance

Knowledge about a set of relevant state variables $\mathbf{x}$, given the available information I, is modeled by a conditional joint probability density function (PDF) $f_{\mathbf{x}|I}$, describing the relative probabilities that the state vector will take on each of the feasible values. This approach allows the uncertainty as well as the state estimate to influence the decision. The probability that the state $\mathbf{x}$ is in region r_i, denoted $P[\mathbf{x} \in r_i]$, equals the integral under the PDF, over all $\mathbf{x}$ in the region r_i.

The state model having a large uncertainty will not necessarily affect the outcome of a proceduralized decision problem. Figure 2 shows three possible models for a single state variable x, in relation to a proceduralized decision threshold θ. The threshold defines two *regions*: $x < \theta$ and $x > \theta$. Although the standard deviations of the PDFs in parts *a* and *b* are equal, in *b* there is no uncertainty concerning whether the state is greater or less than the threshold θ. Although the standard deviation of the PDF in *c* is smaller than that in *a*, the two state models exhibit similar uncertainty concerning whether the state is greater or less than the threshold. These two comparisons illustrate how the significance of state model uncertainty to a proceduralized decision problem depends on both the magnitude of the uncertainty and the proximity of the expected state to the threshold.

Therefore, rather than considering the magnitude of the uncertainty, a measure of how significant that uncertainty is *in the context of the proceduralized decision problem* is required. State model uncertainty is significant to the outcome of a proceduralized decision problem when the region in which the state lies is not known with confidence.

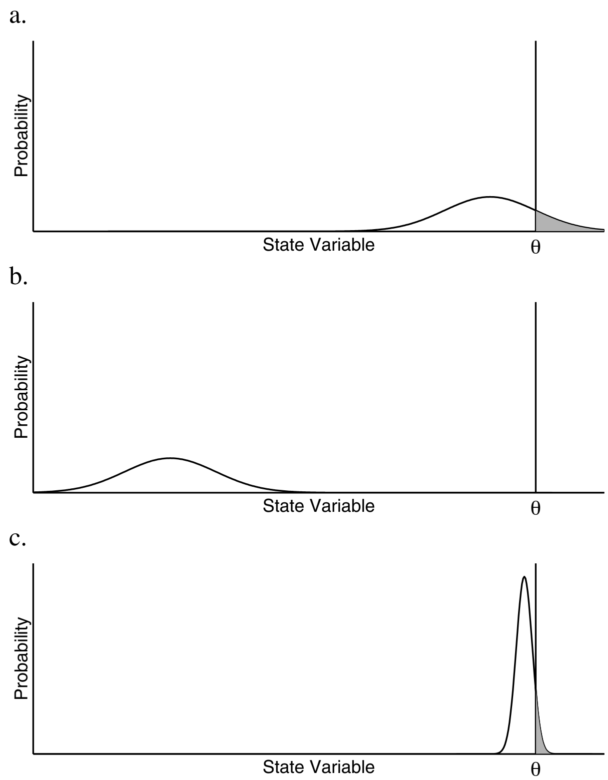

Figure 2. The difference between the magnitude and the significance of uncertainty

Modeling Time-Dependence

A measurement describes the condition of a state variable at the time the measurement was taken. As the measurement ages, it less accurately reveals the current or future condition of the state. Shortly after a measurement, the decision maker is confident that the state lies within a narrow range. However, when the state is predicted further into the future, the decision maker is equally confident only that the state lies within a much broader range. Therefore, to maintain the situation (i.e., state) awareness necessary to make effective procedural decisions, the decision maker must receive repeated observations of the state. Figure 3 illustrates monotonically increasing uncertainty in state models predicted at various times after a perfect measurement was taken, as well as the effect of a new measurement. New information conditionally modifies the a priori state model, both adjusting the expected state and reducing the uncertainty.

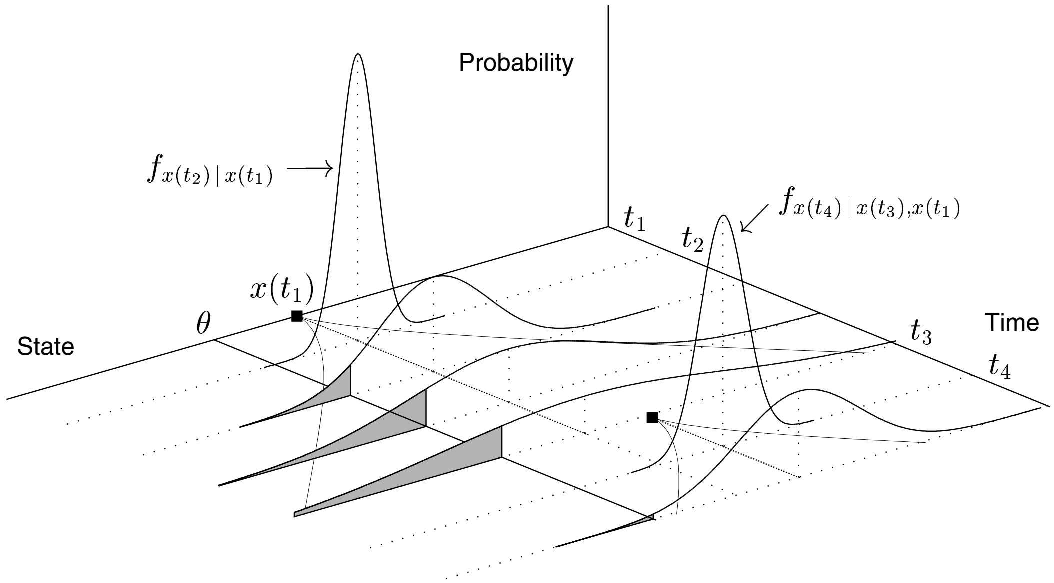

Figure 3. Time-dependence of the state model.

To forecast how the state may evolve in the future requires a model of the state dynamics; the rate at which uncertainty increases depends on this model. If the state dynamics and the external forces acting on the state are known deterministically, the future state trajectory could be predicted without any expected error. Since they are not, a random process is used to model the future state trajectory. The trajectory which the state actually follows will be one sample realization of the random process [6, 7].

Modeling Contextual Dependence

Figure 1 separated the process of solving a proceduralized decision problem into two cascaded steps - modeling state variables and applying a proceduralized decision rule. This separation revealed the value of information as its effect on the ability to model the state variables, measured in the *context* of the proceduralized decision problem (i.e., how improvements in the state model affect the likelihood that the decision maker will select the action that is called for by the procedure, and the relative costs of correct and incorrect decisions). The previous two sections have introduced a time-dependent model for the state variables. This section models the context within which information is valued. The following section formally defines the information value metric.

Payoff Matrix and Penalty Matrix

A *payoff matrix* is used to describe the context within which information is valued. Consider a proceduralized decision problem in which a decision maker must choose between M decisions $\{a_1, a_2, ..., a_M\}$. Threshold surfaces, defined by the procedure, divide the space of relevant state variables into N regions $\{r_1, r_2, ..., r_N\}$. The *outcome* of decision a_j, denoted $\{a_j, r_i\}$, is the consequence of choosing action a_j when the state is in region r_i. The payoff matrix assigns a cost K_{ij} to every possible outcome $\{a_j, r_i\}$. For example, K_{21} is the cost of choosing a_1 when the state is in region r_2. In general, the elements of the payoff matrix will be time-varying, the implication of which will be explored in later sections.

Assume that the procedure is optimal with respect to minimizing cost. The *correct* decision for a particular region (i.e., the decision called for by the procedure) is the decision that yields the minimum cost that is achievable when the state lies in that region. When the state is in region r_i, a_j is the *correct* decision if $K_{ij} = \min_k K_{ik}$, and is *incorrect* if there exists another possible decision a_k for which $K_{ik} < K_{ij}$. Define K_i^* to be the minimum cost that is achievable when the state is in region r_i (i.e., $\min_j K_{ij}$). The minimum costs are generally not equal (e.g., $K_1^* \neq K_2^*$), implying the state being in one region of the state space is inherently more costly than it being in the other.

The *penalty* for a decision outcome $\{a_j, r_i\}$ is the amount by which the cost for that decision outcome exceeds the cost for the outcome of the correct decision (i.e., K_{ij} - K_i^*). Since incorrect decisions are caused by uncertainty in the state model, the penalty is the cost of uncertainty. The *penalty matrix* gives the penalties for each possible decision outcome (i.e., every combination of decision and region). Table 1 shows the penalty matrix for the case $M = 3$, $N = 3$, and assuming a_i is the correct action when the state is in region r_i. For example, when $\mathbf{x} \in r_1$, the penalty for choosing action a_2 (an incorrect decision) rather than a_1 (the correct decision) is $K_{12} - K_{11}$.

Table 1. Penalty matrix.

N regions	M decisions: $a = a_1$	$a = a_2$	$a = a_3$
$x \in r_1$	0	$K_{12} - K_{11}$	$K_{13} - K_{11}$
$x \in r_2$	$K_{21} - K_{22}$	0	$K_{23} - K_{22}$
$x \in r_3$	$K_{31} - K_{33}$	$K_{32} - K_{33}$	0

Decision Model

In general, the decision maker will select his decision based on the procedure, the probabilities $P[\mathbf{x} \in r_i]$ that the state is in each of the regions, and costs for each of the possible decision outcomes (i.e., the payoff matrix). Since it is desirable to not limit the definition of information value to a particular decision policy, such as minimizing the expected cost, the decision is modeled by defining the probabilities $P[a=a_j \mid I]$ that the decision maker will choose each of the possible decisions. How the probabilities are determined is application specific, and will be illustrated in the case study.

Expected Cost and Expected Uncertainty Cost

The *expected decision cost*, $C|I$, is the cost that is expected to result from the decision problem, where the expectation is over the region in which the state lies and the decisions which may be chosen by the decision maker.

$$C \mid I = \sum_{j=1}^{M} P\left[a = a_j \middle| I\right] \sum_{i=1}^{N} P\left[\mathbf{x} \in r_i \middle| I\right] K_{ij} \qquad (1)$$

Consider the cost that is expected to result when $P[a=a_1] = 1$ and $N = 2$.

$$C \mid I = P\left[\mathbf{x} \in r_1 \middle| I\right] K_{11} + P\left[\mathbf{x} \in r_2 \middle| I\right] K_{21} \qquad (2)$$

Let a_1 be the correct decision when the state is in region r_1, which implies $K_{11} < K_{1j}$ for all $j \neq 1$. Also let $K_{11} > K_{21}$, which would occur if the state being in region r_1 is inherently more costly than it being in region r_2. Assume $P[\mathbf{x} \in r_1 \mid I] = P[\mathbf{x} \in r_2 \mid I] = 0.5$; the a priori expected decision cost equals $0.5\ K_{11} + 0.5\ K_{21}$. Assume new information I_2 reveals that the state lies in region r_1 (i.e., $P[\mathbf{x} \in r_1 \mid I_2, I] = 1$) and, therefore, the original decision a_1 is correct. After receiving this new information, the expected decision cost equals K_{11}. Although the new information reduces the uncertainty in the state model, it increases the expected decision cost. Therefore, expected decision cost is not a useful measure of the consequence of state model uncertainty on a proceduralized decision problem.

The role of new information is to increase the likelihood that the decision maker will choose the *correct* decision, by improving his ability to model the relevant state variables. Therefore, the fundamental quantity on which information value should be defined is the consequence of uncertainty in the state model, measured in the context of the proceduralized decision problem. The concept of *expected uncertainty cost* is introduced as a measure of the consequence of state model uncertainty on a proceduralized decision problem.

The *expected uncertainty cost*, $R|I$, is the penalty that is expected to result from the decision problem.

$$R \mid I = \sum_{j=1}^{M} P\left[a = a_j \middle| I\right] \sum_{i=1}^{N} P\left[\mathbf{x} \in r_i \middle| I\right] \left(K_{ij} - K_i^*\right) \qquad (3)$$

The penalty is the unnecessary cost which results because the decision is incorrect. The expected uncertainty cost is interpreted as the amount by which the cost is expected to exceed the minimum achievable cost, because the uncertainty in the model of the state may cause the decision maker to choose an incorrect decision. Notice that the expression for the *expected uncertainty cost* (3) is equivalent to that for the *expected decision cost* (1) with the *cost* of the decision outcome replaced by the *penalty*.

Definition of Information Value

Given initially available information I_1, the model of the relevant state variables is $f_{\mathbf{x}|I_1}$ and the a priori expected uncertainty cost is $R|I_1$. Given new information I_2, the a posteriori expected uncertainty cost becomes $R|(I_2, I_1)$. Note that receiving the new information may cause the decision maker to choose a different decision from that which he would have chosen with only the original information. The value of information I_2 is defined to be the magnitude of the change in the expected uncertainty cost, $\left| R \mid I_1 - R \mid \left(I_2, I_1\right) \right|$.

The absolute value function is included to positively value all changes in expected uncertainty cost. Information should have positive value when it allows the state to be estimated more accurately, and that improvement in the state model is significant to the

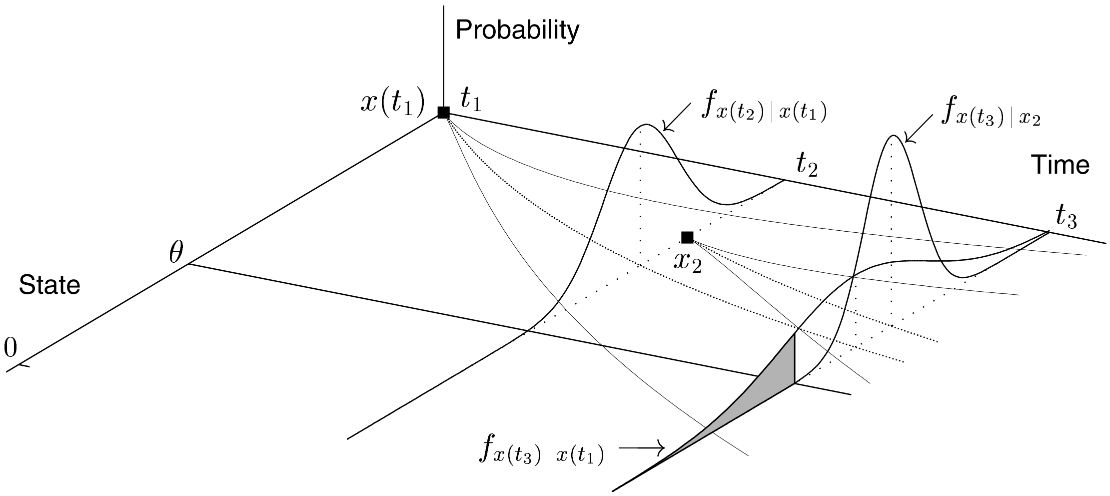

Figure 4. The effect of a measurement at time t_2 taking on a particular value x_2, on the state model at time t_3.

proceduralized decision problem (i.e., increases the likelihood that the correct decision will be chosen). By more accurately revealing the true state of the world, additional information is generally expected to reduce the expected uncertainty cost. However, if the true state is closer to the threshold between two regions than predicted by the a priori model, new information may increase the expected uncertainty cost. A positive value is attributed to information that increases the expected uncertainty cost because it reveals: (1) the a priori state model was misleadingly confident about how well the region in which the state lies is known, (2) the larger uncertainty should be taken into account when the decision is chosen, and (3) additional information should be sought to reduce the expected uncertainty cost.

Often, the decision whether or not to seek the new information must be made prior to knowing the content of the information. Figure 4 illustrates the effect of one possible measurement at time t_2 on the model of a state variable at time t_3. The *expected information value* of a measurement taken at time t_2, V, is defined as the expectation over the feasible measurements of the values conditioned on those measurements being received.

$$V(I_1, t_2) = \int_{I_2} f_{I_2|I_1} \left| \, R|I_1 - R|(I_2, I_1) \, \right| \, \mathrm{d}I_2 \tag{4}$$

The dependence of V on I_1 and t_2 is shown explicitly, as a reminder that information value depends on the decision maker's a priori knowledge and the time at which the information is measured. The value of information also depends on the model of the state dynamics.

Aircraft on Crossing Trajectories

This section generalizes the time-dependent information value definition to measure the total value of multiple new measurements, and studies when new measurements should be taken to support the problem of predicting whether or not a conflict will occur between two aircraft on crossing trajectories. Several assumptions about the encounter geometry are made to simplify the necessary computations. Assume the nominal trajectories of two aircraft, A and B, intersect at a right angle, and that the aircraft will closely follow their nominal trajectories (i.e., will not deviate laterally or in altitude). Aircraft A's along-track position is measured at time t_1, with a known expected error, and its speed is modeled as an integrated, first-order Markov process [3]. The standard deviation of aircraft A's along-track position after 60 minutes, without an intermediate measurement, is 25 miles. The position and speed of aircraft B are known accurately at time t_1, and its speed is constant throughout the encounter. In this example, a conflict occurs if the separation between the aircraft is less than 5 miles at any time during the encounter; the a priori expected miss distance is 0 miles at time t_3. $t_3 - t_1 = 60$ minutes, in this example. Given an encounter geometry and a description of the uncertainties about the aircraft trajectories, several methods for calculating the probability that a conflict will occur are available in the literature.

Collective Expected Information Value of Multiple Measurements

When multiple new measurements are available, the total value of different measurement schedules (i.e., the number of measurements and the times at which they are taken) must be compared. The information value definition introduced in the previous section may be generalized to calculate the *collective expected information value* of multiple new measurements. This will be shown for two measurements, and extension to more than two measurements should be apparent. Rather than a single new measurement being taken at a time t_2, assume that two new measurements will be taken at times t_a and t_b.

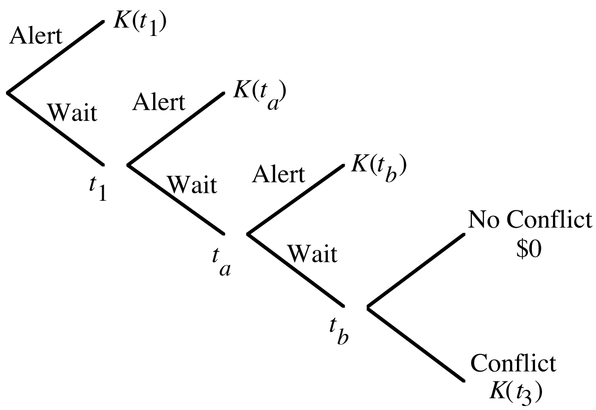

Figure 5. Decision tree for the decision problem at time t_1.

The cost $J(t)$ for alerting at time t represents the cost for an avoidance maneuver initiated at that time. Assume $J(t) = \$50 + \$100\left(\frac{t}{60}\right)^2$. Figure 5 shows the decision tree for the decision problem with which the controller is confronted at time t_1; the decision to either continue monitoring or intervene is made at times t_1, t_a, and t_b, when new information is received. Table 2 gives the expected uncertainty costs for the two possible decisions, at each of the times when a decision is made, assuming a minimum expected cost decision rule. $\mathrm{P}[C]$ is the probability a conflict will occur.

Table 2. Expected uncertainty costs.

	Alert	No Alert
t_1	$K(t_1)\left(1 - \mathrm{P}\left[\mathrm{C}\middle\|x_1\right]\right)$	$\left(K(t_a) - K(t_1)\right)\ \mathrm{P}\left[C\middle\|x_1\right]$
t_a	$K(t_a)\left(1 - \mathrm{P}\left[\mathrm{C}\middle\|x_a\right]\right)$	$\left(K(t_b) - K(t_a)\right)\ \mathrm{P}\left[C\middle\|x_a\right]$
t_b	$K(t_b)\left(1 - \mathrm{P}\left[\mathrm{C}\middle\|x_b\right]\right)$	$\left(K(t_3) - K(t_b)\right)\ \mathrm{P}\left[C\middle\|x_b\right]$

Assume the controller decides to continue monitoring at time t_1. The collective expected value of new measurements taken at times t_a and t_b, calculated at time t_1, equals the expectation (over the possible measurements) of the value of the measurement at time t_a (to the decision at t_a) plus the expectation (over the possible measurements at time t_b, given the measurement received at time t_a), of the value of the two new measurements (to the decision at t_b), where value is defined as the change in the expected uncertainty cost.

$$V\left(x(t_1), t_a, t_b\right) = \int_{x_a} f_{x(t_a)|x_1} \left[\begin{array}{c} \left|R|x(t_1) - R|x_a\right| \ + \\ \int_{x_b} f_{x(t_b)|x_a} \left|\mathrm{R}|x_a - R|x_b\right| \ \mathrm{d}x_b \end{array} \right] \mathrm{d}x_a \tag{5}$$

The plan for when the first measurement should be taken must consider the opportunity for taking a second measurement. Therefore, when the first measurement should be taken can be found by determining the times t_a and t_b whose combination maximizes the collective expected information value. However, Equation (5) calculates the total value that is expected prior to receiving either of the new measurements. After the first new measurement is received, the expected value of the second measurement, conditioned on the first, may change. Therefore, whether or when to take the second measurement must be reevaluated after the first measurement is received.

Results

Figure 6 plots the collective expected information value as a function of the times at which two new measurements are taken. If a single new measurement is taken, its expected value is largest when it is taken at 48 minutes. Taking the measurement at this time is the best compromise between the ability to make the correct decision, and the benefit for making the correct decision at an early time.

If two new measurements are taken, the collective expected information value is largest when one measurement is taken at 22 minutes and the other at 52 minutes. The expected value of the first measurement results because there is some chance that the measurement will allow the correct decision to be made at an early time. If the correct decision is to alert, the cost of the avoidance maneuver will be less than if the maneuver were initiated at a later time. Alternatively, if the first measurement reveals that the correct decision is to not alert, then the information cost of the second measurement may be avoided. Notice from Figure 6 that the expected value of a measurement taken at 52 minutes is fairly independent of whether or not a measurement is taken at 22 minutes. This occurs because, if the content of the first measurement is such that a second measurement is still necessary, the first measurement does not significantly reduce the uncertainty that exists before the second measurement is received.

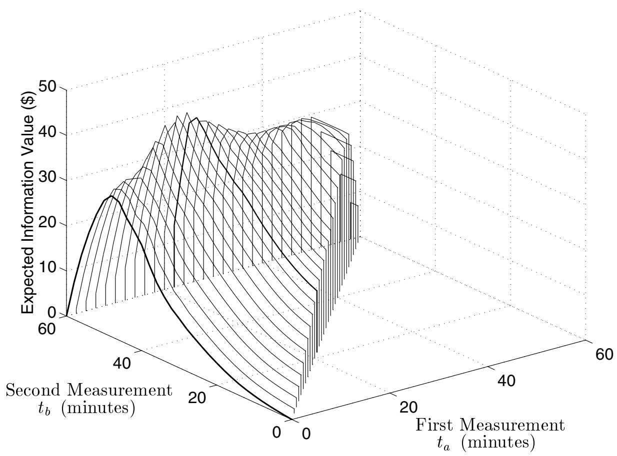

Figure 6. Collective Expected Information Value for two new measurements.

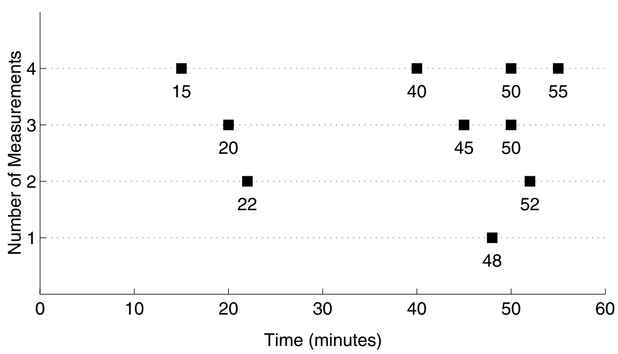

Figure 7. Optimal times at which to take 1, 2, 3, or 4 new measurements.

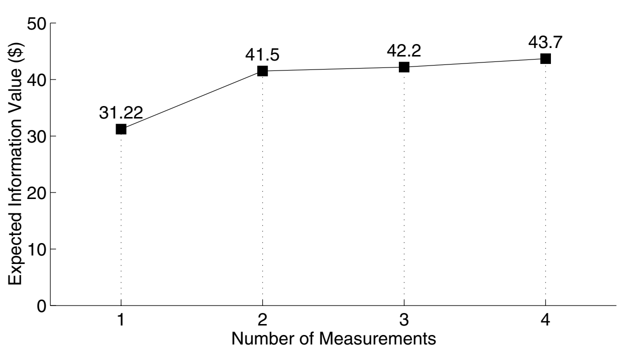

Figure 8. Maximum Collective Expected Information Values achievable from 1, 2, 3, and 4 new measurements.

Equation (5) may be further generalized to calculate the collective expected value of more than two measurements. Figure 7 shows the optimal times at which to take 1, 2, 3, or 4 new measurements; Figure 8 shows the collective expected information value for measurements taken at those times. For 3 and 4 new measurements, the intervals between the optimal measurement times decrease as the time remaining until the possible conflict decreases. However, the collective expected information value exhibits diminishing returns, as the number of measurements is increased. An additional measurement should be sought only if the resulting increase in the collective information value is greater than the cost of the measurement. Therefore, this approach may be used to determine how many new measurements should be taken, in addition to when those measurements should be taken, by calculating the net value for one, two, three, and so on, new measurements, where *net value* is defined as the collective expected information value minus the total information cost.

Note that the expectations underlying the results in Figures 7 and 8 are calculated prior to receiving any of the new measurements (i.e., at time t_1). After the first new measurement is received, the expected values of the future measurements, and the times at which future measurements have maximum expected value, may change. Given a new piece of information, the plan for what information should be sought in the future must be re-calculated, using the same method. The resulting computational demand makes this approach more applicable as an off-line analysis tool, than as an on-line algorithm (i.e., programmed to run in real-time on-board aircraft or in ATC computers).

Aircraft on Parallel Tracks

This section applies the time-dependent information value model to study the optimal tradeoff between the cost of information and the cost of uncertainty, when measurements are taken periodically. Consider two aircraft, A and B, flying at equal speeds on parallel oceanic tracks (or landing on parallel runways). A controller monitors the aircraft separation and intervenes if a situation which would otherwise result in a conflict arises. Since a conflict may occur at any time, the decision problem is continuous. However the controller's predictions for the future trajectories of the aircraft only change when new information is received. Therefore, if upon receiving a measurement the controller decides to continue monitoring, the controller will not intervene before receiving the next measurement.

To support monitoring or continuous control tasks, information must be repeatedly updated. In general, the optimal intervals between multiple new measurements vary as the context changes, as shown in the previous section. In the present problem, measurements are taken at equal intervals because, as long as the aircraft remain on their nominal parallel trajectories, the context in which information is valued remains constant; let Δt be the time between subsequent measurements. Note that if one of the aircraft deviates from its nominal path, the new encounter geometry resembles the problem studied in the previous section. Therefore, the objective of taking or communicating periodic measurements is to detect an aircraft "blunder." The update rate is chosen to balance the cost of latency in the detection against the cost of the measurements.

Information Cost

If there is zero cost for seeking or using information, then information should be updated continuously. When there is a cost for seeking or using information, the benefit of the additional information must be balanced against that cost. In the previous sections, the value of information was defined as the difference between the expected uncertainty costs, with and without the information. To compare different pieces of information, the cost without the information was used as a consistent baseline from which to define the measure of value. In the continuous decision problem, this baseline, which would have the physical interpretation of never receiving new information, does not produce a meaningful measure of value. Therefore, the cost itself, where cost now includes both the uncertainty cost R and the cost of information C_I, is used to compare possible information update rates. Note that the cost of information must be measured in the same units as the cost of uncertainty (i.e., the units of the payoff matrix).

$$C_\mathrm{T} = C_\mathrm{I} + R \tag{6}$$

Since the decision problem is continuous, the average cost per unit of time will be used to compare various update rates. Superscripts are used to indicate the period of time over which a cost has accumulated: Δt denotes a cost which has accumulated over a single measurement interval, and dt denotes the average rate at which cost accumulates.

$C_\mathrm{I}^{\Delta t}$ is the cost for a single measurement. $C_\mathrm{I}^{dt}(\Delta t)$ is the average cost of information per unit of time.

$$C_\mathrm{I}^{dt}(\Delta t) = \frac{C_\mathrm{I}^{\Delta t}}{\Delta t} \tag{7}$$

Expected Uncertainty Cost for Continuous Decision Problems

To study the periodic update rate, assume the aircraft remain on their nominal trajectories. If the aircraft deviate onto crossing paths, the future information requirements can be determined as in the previous section. Only a single row of the payoff matrix, corresponding to the decision to not intervene, is required. Two outcomes are feasible - either a conflict occurs or a conflict does not occur, defined by the parameter θ. Let the rates at which cost accrues for these decision outcomes be: K_{11} = \$50/minute and K_{12} = \$0/minute, respectively. Note that, since the decision is continuous, the payoff matrix must be expressed in terms of the rate at which cost accumulates.

$R(\tau)$, the instantaneous rate at which expected uncertainty cost accumulates, for the decision to continue monitoring, equals the probability that the aircraft separation is less than the required separation θ times the penalty $K_{12} - K_{11}$ = \$50/minute. τ is the time since the last measurement, and $\theta = 10$ nautical miles, in this example. The \$50/minute penalty for the controller failing to intervene when a conflict will occur, due to uncertainty in the state model, represents the added cost to the aircraft for delaying the necessary avoidance maneuver (e.g., increased fuel-burn of a more aggressive maneuver).

A probabilistic model for the aircraft trajectories is required. Assume the aircraft are nominally separated by 50 nautical miles, and their future cross-track positions are described by a twice-integrated, first-order Markov process [3]. This model exhibits exponentially growing uncertainty in aircraft separation as information age increases; a parameter σ, expressed in nautical miles per minute^2, determines the rate at which the uncertainty in the state grows. Although a variety of other models could be suggested and studied at length, this model serves the objective of this example - to study the general properties of the relationship between information cost and update rate.

The *cumulative expected uncertainty cost* $R^{\Delta t}$ represents the cost that is expected to accumulate over a single measurement interval due to exposure to a conflict situation, and equals the integral of $R(\tau)$ from $\tau = 0$ to $\tau = \Delta t$. Define R^{dt} to be the average cumulative expected uncertainty cost over the measurement interval.

$$R^{dt}(\Delta t) = \frac{R^{\Delta t}(\Delta t)}{\Delta t} = \frac{1}{\Delta t}\int_0^{\Delta t} R(\tau)\, d\tau \tag{8}$$

Optimal Measurement Interval

Figure 9 plots the per-unit-time costs from Equation (6), as functions of the measurement interval, for the case $C_\mathrm{I}^{\Delta t} = \$5$, K_{12} = \$50/minute, and σ = 10 nautical miles/minute^2. When Δt is small, the fixed cost for a measurement is assessed over a short time, resulting in a high average (per unit of time) information cost. However, the frequent measurements achieve a small average cumulative expected uncertainty cost. When Δt is large, the average cost of information is small, while the average cumulative expected uncertainty cost is high. The optimal (i.e., minimum total cost) measurement interval, denoted by Δt^*, is a tradeoff between the cost of information and the expected cost of uncertainty. For the assumptions made in Figure 9, the optimal measurement interval is approximately 4 minutes. The corresponding minimum average total cost for monitoring and maintaining separation between aircraft which are following parallel tracks separated by 50 nautical miles is approximately \$1.50/minute. Parametric studies may be used to identify the assumptions to which this result is most sensitive, allowing effort to be concentrated on accurately identifying those parameters.

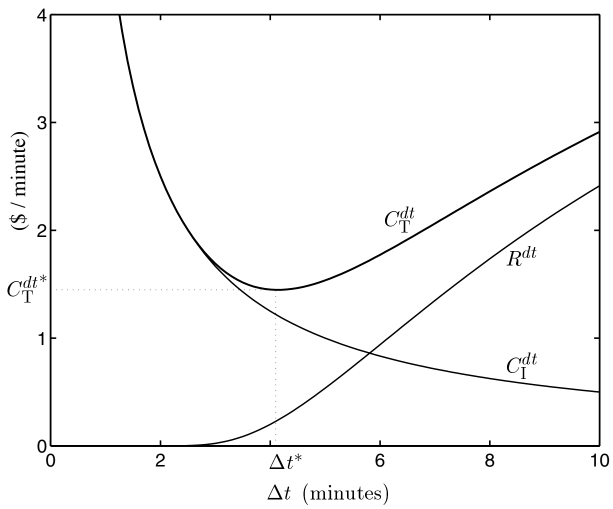

Figure 9. The optimal measurement interval trades off between information and expected uncertainty costs.

By plotting the average total cost C_T^{dt} versus the measurement interval Δt for four values of the per-measurement cost of information $C_I^{\Delta t}$, Figure 10 illustrates the sensitivity of the optimal measurement interval to the cost of information. Increasing $C_I^{\Delta t}$ makes information more expensive relative to the average cumulative expected uncertainty cost. Therefore, both the optimal measurement interval Δt^* and the minimum average total cost C_T^{dt*} increase (i.e., measurements should be taken less frequently) as $C_I^{\Delta t}$ increases.

Figure 11 shows the sensitivity of the optimal measurement interval to the penalty for delaying

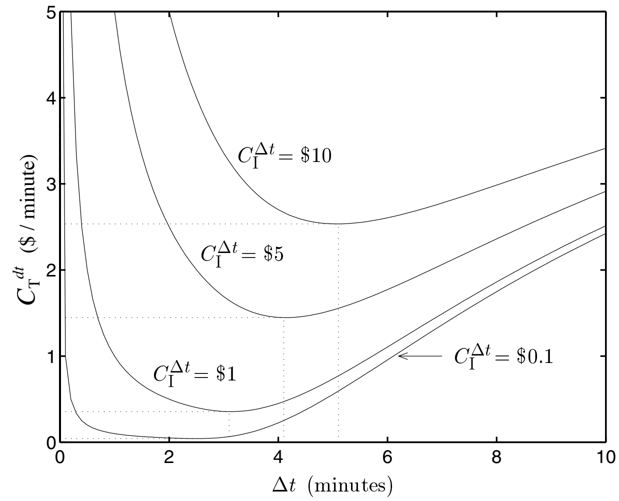

Figure 10. Sensitivity of the optimal measurement interval to the per-measurement cost of information.

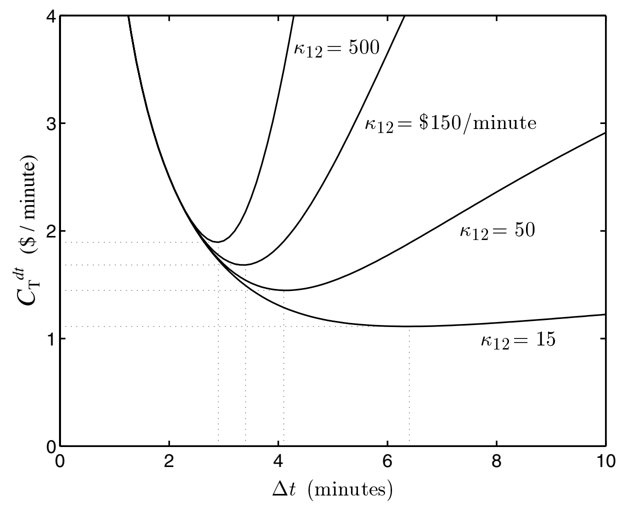

Figure 11. Sensitivity of the optimal measurement interval to the penalty for delaying intervention.

necessary intervention, by plotting the average total cost C_T^{dt} versus the measurement interval Δt, for four values of K_{12}. For larger values of K_{12}, the cumulative expected uncertainty cost increases faster. Therefore, increasing K_{12} decreases the optimal measurement interval, because information is less expensive relative to the cumulative expected uncertainty cost, making more frequent information updates cost effective. The optimal measurement interval exhibits similar sensitivity to the state dynamics model parameter σ, which is used to vary the dependence of $R^{\Delta t}$ on Δt. As σ increases, the optimal measurement interval decreases because the relative cost of information decreases.

Conclusions

To answer questions about when information should be measured, communicated, or displayed, a novel model of time-dependent information value was developed for a class of proceduralized decision problems. Proceduralized decision problems require the decision maker to estimate a set of relevant state variables, in the context of proceduralized decision thresholds. The new information value metric directly measures the impact of information on the decision maker's ability to estimate these state variables. The metric is based on the expected cost of the uncertainty in the state model - the amount by which the cost resulting from the decision will exceed the minimum achievable cost, due to the uncertainty causing the decision maker to choose a decision that is incorrect according to the procedure.

The model of time-dependent information value was used to determine how many measurements should be taken, and at what times, to support a controller monitoring and maintaining separation between two aircraft on intersecting trajectories. The time intervals between optimally spaced measurements were shown to decrease as the time remaining until the possible conflict decreases, and the cumulative information value, expected prior to receiving any of the measurements, exhibited diminishing returns with an increasing number of measurements. Since new information may change the expected value of future information, how many and when subsequent measurements should be taken must be re-planned. The computational demand of re-calculating when to take future measurements each time a measurement is received makes the approach best suited as an off-line analysis/design tool.

The time-dependent information value model was also used to identify the optimal periodic rate at which two aircraft flying along parallel tracks should broadcast position measurements. The solution trades off between the expected cost of uncertainty and the cost of information. Parametric studies were used to identify the sensitivity of the results to the assumptions.

More than providing specific results, these examples demonstrate the utility of the proposed time-dependent information value model as a tool for designing processes, both manual procedures and automatic systems, for managing time-dependent information. Proceduralized decision problems were successfully modeled as estimation problems by defining regions in the state space whose boundaries are significant to the procedure. The expected uncertainty cost was shown to be a useful basis for measuring information value for this class of decision problems. Finally, the time-dependence of the information and the decision problem were incorporated using estimation techniques.

References

[1] Atkins, Stephen C., "A Unified Methodology for the Management of Time-Dependent Information in Aeronautical Systems," Ph.D. Dissertation, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA, 1998.

[2] Shannon, Claude E., "Communication in the Presence of Noise," *Proceedings of the Institute of Radio Engineers*, vol.37, no.1, pp.10-21, 1949.

[3] Brown, Robert G., "Introduction to Random Signal Analysis and Kalman Filtering," New York: John Wiley & Sons, 1983.

[4] Howard, Ronald A., "Information Value Theory," *IEEE Transactions on Systems Science and Cybernetics*, vol.2, no.1, pp.22-26, 1966.

[5] Sheridan, Thomas B., "Reflections on Information and Information Value," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 25, no. 1, pp. 194-196, 1995.

[6] Drake, Alvin W., "Fundamentals of Applied Probability Theory," New York: McGraw Hill, 1967.

[7] Gelb, Arthur, "Applied Optimal Estimation," Cambridge, MA: MIT Press, 1974.