Analysis of Driver Safety Performance Using Safety State Model

EDWARD J. LANZILOTTA

A significant component in the pursuit of safety is estimation of risk probability. In transportation systems virtually all safety-related events and outcomes involve an intermediate event known as an accident. The safety state model is a probabilistic model that is used to estimate the probability of an accident as a function of the human–machine system state. By using a discrete Markov network, the safety state model forms a framework for capturing the human–machine and human–human interactions in a transportation system. The observed data are used to calibrate the model, which is subsequently used to estimate the risk probability performance of other human operators. The theoretical development of this model is reviewed. In addition, motivation and background, as well as advantages and disadvantages with respect to existing quantitative methods of risk probability estimation, are discussed. Finally, the applicability to driver performance analysis is discussed.

A current trend in the automobile industry is an emphasis on safety. Automobile manufacturers are implementing features intended to improve highway safety. The motivation is well justified because the magnitude of damage, injury, and death on highways remains a significant problem.

A key component of safety in the human–machine system is the human driver. As controller of the vehicle the driver must monitor the state of the vehicle (position and speed) and the surrounding environment (including other vehicles, as well as roadside elements), make decisions with regard to control actions, and actuate the vehicle controls to carry out the control decisions. The driver is solely responsible for the state of the vehicle: the control decisions and actuation determine the resulting position and speed of the vehicle in the environment. The control actions of the driver play a significant role in the effectiveness of any designed-in safety device or system. In the worst case a driver can counteract or override the effects of a designed-in safety system because of either a lack of training or a higher level of risk-taking.

In this work the driver–vehicle system is modeled as a closed-loop control system (*1,2*), with the vehicle as the "plant" and the human driver as the "controller." The driver senses the state of the vehicle and environment and provides control input to the vehicle. The control strategy is a time stream of decisions made by the driver, which govern the actuation response to the driver's perception of the system state. These decisions are the result of some combination of rational thought and instinctive response, and the basis for the decisions is typically obtained through a combination of training and experience.

Measuring driver performance is a challenge, especially with respect to safety-related decision behavior. Current driver evaluation methods focus on perception and actuation, because these task skills are observable and testable (e.g., through eye examinations and simplistic road tests). However, the decision-making component of the human–machine control system plays a significant role in governing the interaction between a vehicle and its environment. These interactions ultimately determine the safety of the driver's actions.

This research is focused on modeling system behavior in ground transportation systems. In particular, researchers are interested in evaluating the decision behavior of vehicle drivers. This behavior is evaluated in terms of risk probability as a function of the state of both the vehicle and its environment during operation. A probabilistic model is used to represent the state of the human–machine system. This model, termed the *safety state model,* provides a framework for observing system behavior and driver decision behavior as a function of time. The observed data are used to calibrate the model, which is subsequently used to predict the risk probability as a function of the system state. By using the resultant relationship between system state and estimated risk probability, the safety state trajectory of a human–vehicle system can be transformed into a risk probability trajectory as a function of time. The risk probability trajectory is used as the basis for evaluation of driver performance with respect to safety. In the event of an accident the safety state trajectory provides a chain of events leading to the accident, which may be used for determination of causality.

DISCUSSION OF SAFETY AND RISK ASSESSMENT

Before developing the theory of the safety state model consider the terms used and the meanings associated with them. Lowrance (*3*) defines safety as the "judgment of acceptability of risk." This definition provides a working framework for the pursuit of safety, which includes subjective and objective components. The subjective component, which is the judgment of acceptability, evaluates whether a given level of risk is acceptable to the society that is affected. Policies are set on the basis of that judgment. These policies determine the trade-off between a level of risk and the resources expended to reduce that level of risk. Risk judgment is typically performed by policy makers.

The objective component is risk assessment. A variety of definitions of risk can be found in the literature. Rowe (*4*) defines risk as "the potential for unwanted negative consequences of an event or activity," alluding to the notion of chance. Lowrance (*3*) includes the probabilistic component explicitly, defining risk as the "measure of probability and severity of adverse effects." Rescher (*5*) echoes that idea: "Risk is the chancing of negative outcome. To measure risk we must accordingly measure both of its defining components, the chance and the negativity." Gratt (*6*) specifies the relationship between probability and severity in risk assessment by stating that the "estimation of risk is usually based on the expected result of the

Human-Machine Systems Laboratory, Department of Mechanical Engineering, Massachusetts Institute of Technology, Room 3-355, 77 Massachusetts Avenue, Cambridge, Mass. 02139.

conditional probability of the event times the consequences of the event given that it has occurred." Wharton (7) offers that "a risk is any unintended or unexpected outcome of a decision or course of action," including both positive and negative outcomes.

In the case of transportation safety risk assessment is most often considered in terms of fatalities, personal injury, and property damage: the results of an accident. Based on the definitions of risk, risk assessment in transportation systems can be divided into two subcomponents: the probability and the severity of an accident. The risk probability of an accident estimates the relative likelihood that such an event will occur. The severity estimates the ultimate outcome of the accident in the terms of interest (i.e., injuries and deaths) for a given set of conditions regarding the accident. In a sense the accident event represents a demarcation point in time: all of the events leading to an accident contribute to risk probability, whereas those that occur after the accident are in the domain of severity. These components of risk assessment parallel the concepts of "active safety" and "passive safety" devices, respectively. "Active safety" is a term typically applied to those devices or systems that assist in preventing accidents (such as anti-lock braking systems and traction control), whereas "passive safety" devices are those that reduce the severity of an accident when it does occur (such as airbags and door guard beams). Risk assessment is typically performed by systems analysts. This research is focused on estimation of risk probability.

Risk assessment is, in effect, a subset of reliability engineering, which is focused on estimating the probability and effects of system failures. When a system failure can result in injury or death to a human it becomes a safety issue. Assessment of system reliability with respect to a failure of this type is risk assessment.

Risk probability, especially in transportation systems, is not a static quantity. Instead, risk probability varies as a function of the state of the system, which includes the state of the vehicle as well as the state of the environment. The system state in transportation systems is quite dynamic with respect to time. The driver is responsible for a constant stream of control decisions, and the actions resulting from those decisions determine the state of the vehicle in relation to the state of the environment. Thus, through these control decisions the driver has a profound impact on the risk probability of the vehicle system. Many accident scenarios are the result of compounding several hazard conditions, each of which may be relatively innocuous when it occurs in isolation. Some of these hazards may be due to driver errors (8), whereas others may be due to machinery failures in vehicle or wayside equipment. The collected set of potential hazard conditions leading to a particular accident scenario can be considered a system state. Because this state varies with time and the risk probability is a function of this state, risk probability can also be considered a function of time.

Time is an integral component of risk probability. The risk probability can be modeled by probability theory as the relative likelihood of the occurrence of an accident. However, the risk probability of an accident only makes sense if its occurrence is compared with the alternative event, which is the nonoccurrence of an accident. Since nothing "happens" during the nonoccurrence, the event can only be considered with respect to some fixed metric. The safety state model considers the probability of an accident with respect to a fixed time frame, known as a *time slice.* Thus, the risk probability represents the relative likelihood of an accident in a single time slice. On average, it also represents the percentage of time slices that result in an accident. An alternate form of expression is in terms of the mean time (number of time slices) between the occurrence of accidents. This form is commonly known as the mean time between failures (MTBF) and is used extensively in the field of reliability engineering.

Even if the risk probability for a rare event is very small (alternately, the MTBF is very large), probability theory asserts that the event will eventually occur, given enough opportunity (i.e., time). From this, it can be seen that the only way to avoid a probabilistic event is to "get out of the game" before that event occurs. (In fact, this is what happens to most people with respect to rare catastrophic events—the human lifetime is much shorter than the time period in which one could expect to experience a single occurrence.) Thus, the concept of risk exposure is as follows: given a constant risk probability, the expected number of failures over a prescribed period rises with the size of the period. To reduce the overall risk of an undesirable event, one must reduce either the risk probability or the risk exposure.

Estimating the risk probability of transportation accidents is quite difficult for several reasons. First, accidents are relatively rare occurrences and are difficult to predict. In addition, the events and behavior of highest interest for risk probability estimation are those that immediately precede the accident event: attention needs to be focused on a time period that is identified by an event that occurs at the end of the period. Finally, a compound set of hazards and events typically leads to an accident.

A guiding motivation in this work is the notion that near collisions are far more common than actual accidents. If the capability of identifying near collisions and the conditions that lead to them exists, responses (in either design, operating procedure, or policy) can be formulated to reduce the occurrence of near collisions and in the process reduce the number of accidents. A dynamic estimation of risk probability provides a mechanism for identifying system states corresponding to near collisions.

SAFETY STATE MODEL

The safety state model is an extension of the more familiar event tree and fault tree models. An event tree is a representation of possible scenarios that can occur from a fault-precipitating event (9). A fault tree, by contrast, works backward from a system failure to identify the logical combination of all of the potential causes of that failure (*10–12*). The safety state model has been inspired by these methods of system safety analysis and represents a step forward in generality.

Event Tree Analysis

Event tree analysis is used for human reliability analysis. The purpose of the method is to identify the probability of system failure from the occurrence of a precipitating event. From an event tree it is possible to detect points in the failure process where human reliability is problematic and to use that knowledge to suggest improvements to manual or automated procedures.

Event tree analysis starts at a precipitating event. From the occurrence of that event, branches are constructed to all of the possible next events. Each branch has a probability associated with the occurrence of the next events. Then, from each of the next events, tree limbs are constructed for subsequent events, with associated probabilities. Once the tree has been completed the overall probability of each possible event path can be calculated. Swain and Guttman (9) explicitly state that there should be no more than two branches from each node, representing a binary decision process; others (*11*) allow for event tree construction with more than two branches from an event node, corresponding to partial failure.

A generalized example of an event tree is shown in Figure 1. Figure 1 shows that at the precipitating event on the left the operator can choose to take Action *A* or not. If Action *A* is not chosen a failure will result. If Action *A* is chosen then a second decision point occurs, at which point the operator can choose to take Action *B* or not. If Action *B* is not taken there is a subsequent decision to take Action *C* or not. If Action *C* is not taken a failure will result. If Action *C* is taken the operator will be at the same decision point as if Action *B* was chosen earlier. Thus, Action *C* is known as a corrective action. The probability of system failure can be expressed as

$$P(\bar{A}) \cup [P(A) \cap P(\bar{B}) \cap P(\bar{C})] \cup [P(A) \cap P(B) \cap P(\bar{D})]$$
$$\cup [P(A) \cap P(\bar{B}) \cap P(C) \cap P(\bar{D})]$$

The failure process is thus transformed into a combination of individual failure probabilities, which are more easily determined. The overall probability of failure can then be evaluated through the mathematical combinations of these individual failures, as determined from the event tree.

Event tree analysis is especially well suited for analysis of systems that are procedural by nature, because it effectively measures deviation from an ordered sequence of events. In this regard it has proven to be very useful in the nuclear power industry. However, the binary decision form is not applicable to systems that offer several choices at each decision point. Even when the nonbinary form is used, the event tree method becomes unwieldy as the number of decision options rises. In addition, unless each decision point explicitly includes a time limit, event tree analysis cannot capture the time relationship between events. For these reasons event tree analysis is limited with regard to estimating the risk probability of drivers.

Fault Tree Analysis

Fault tree analysis is another method commonly used in human reliability analysis. In contrast to event tree analysis, fault tree analysis is considered backward looking. The analysis begins with the occurrence of the failure and works backward to identify the combinations of contributing factors.

An important feature of fault tree analysis is the logical combinations of preceding conditions. Through the use of Boolean operators ("and" and "or") the fault tree describes the combinations of precursor faults leading to a system failure. In addition, fault tree analysis can be used as a quantitative method by assigning probabilities to the various failure events. A generalized example of a fault tree is shown in Figure 2. In this example the system will fail either if Event *A* has occurred or if Event *B* occurs, which results from the combined occurrence of Events *C* and *D*. The probability of this system failure can be expressed as $P(A) \cup P(B)$, which is equivalent to $P(A) \cup [P(C) \cap P(D)]$. Boolean algebra provides a powerful tool for logically combining events and hazards that can lead to a failure scenario. Although fault tree analysis does not explicitly exclude "gate-to-gate" connections (*11*) (which correspond to complex Boolean logic expressions), convention dictates that the use of these constructs is avoided.

As with event tree analysis a weakness of the fault tree is coverage: in the case in which contributing events are independent the analysis becomes quite cumbersome, because all combinations of hazards must explicitly be included. Ansell (*13*) notes that fault tree analysis (as well as FMEA, another risk assessment technique) "suffers from a narrowing of our vision of the system by either limiting the number of failure modes for a component or the types of risks considered. They both implicitly rely on the correctness of the technology or science on which the model is built. This is reinforced by cognate dissonance; only perceived possible risks can be guarded against." Fault tree analysis is also weak with regard to capturing the time relationships between events that contribute to an accident.

Fault tree analysis has been applied to estimation of driver risk probability (*14,15*). These studies have been successful in using fault tree methods to identify a framework of causation. However, in neither case was the research directed toward estimating the risk probability as a function of system state. Based on the weaknesses that have been discussed, fault tree analysis is not well-suited to this purpose.

Structure of Safety State Model

The safety state model, an extension of event tree and fault tree analysis, is now described. By assuming that the conditions contributing to system failure are truly independent, the safety state model can be viewed as a generalization of event and fault tree analyses.

Consider a collection of *n* conditions that could possibly contribute to an accident scenario. These conditions include actions taken by the driver (such as acceleration or braking), the state of the

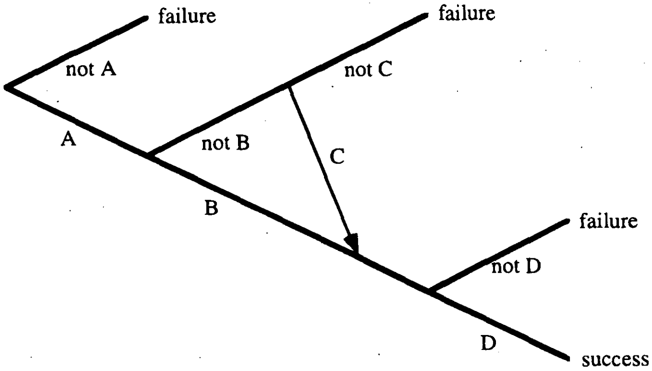

FIGURE 1 Event tree.

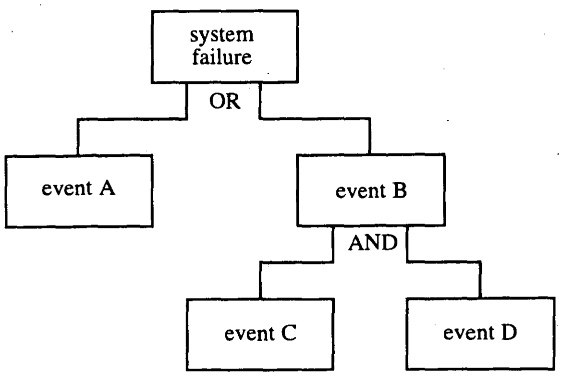

FIGURE 2 Fault tree.

driver (such as fatigue or impairment), the state of another vehicle in the system (such as vehicle ahead braking), and the state of the system environment (such as a red traffic light ahead). These conditions are constrained to be binary conditions. That is, the condition is defined such that it has only two possible values. The complete set of possible combinations of such a set of conditions can be represented by a binary word which is n bits long. The total number of possible combinations is 2^n.

Now consider each of the possible combinations (i.e., each number in the set of possibilities) to represent a state in a Markov network. These states are identified by the associated number, which is within a range of 0 to $(2^n - 1)$, inclusive. The accident scenario is identified as an additional state, labeled 2^n. The state number is termed the *safety state* of the system, and the resultant Markov network is known as the *safety state network.* An example of a three condition safety state network is shown in Figure 3.

The state transition of the Markov network is defined to occur at regular time intervals, with the time period of the interval fixed at h. The value of h is set such that only one condition may change its state (within reasonable probabilistic bounds). At each state transition instant (i.e., at the end of each state transition interval), the model will transition from the current state *S(i)* to the next state *S(j)* with probability $p_{i\to j}$. The probability $p_{i\to j}$ represents the holding probability for the state *S(i)*, which is the probability that the state will not change at the next transition. The collection of transition probabilities for a given state ($p_{i\to j}$, $j = 0, 1, 2, \ldots 2^n$) represents a probability distribution, and the sum of these probabilities must be 1 (Equation 1).

$$\sum_{j=0}^{2^n} p_{i\to j} = 1 \tag{1}$$

The safety state corresponding to the accident scenario [state *S(2^n)*] is a trapping state, which means that once it is entered the process can never exit that state. This notion correlates with the reality that the occurrence of an accident is permanent and cannot be undone.

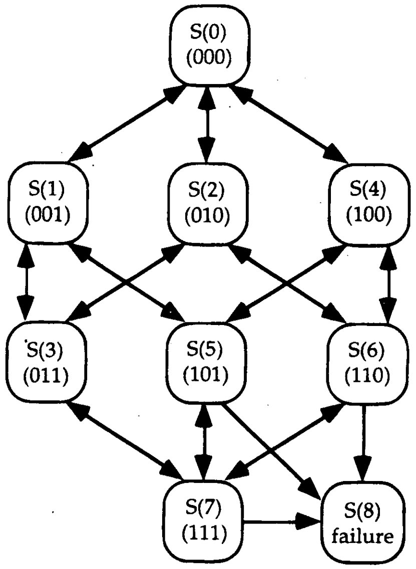

FIGURE 3 Safety state network.

As a trapping state the accident state has a holding probability of 1 and the probability of transition to any other state is zero.

The collection of probability distributions for the entire set of possible safety states can be represented in matrix form (Equation 2). The row of this state transition matrix (P) represents the current state, whereas the column number represents the next state.

$$P = \begin{bmatrix} p_{0\to 0} & p_{0\to 1} & \cdots & p_{0\to 2^n} \\ p_{1\to 0} & p_{1\to 1} & \cdots & p_{1\to 2^n} \\ \vdots & \vdots & \ddots & \vdots \\ p_{2^n\to 0} & p_{2^n\to 1} & \cdots & p_{2^n\to 2^n} \end{bmatrix} = \begin{bmatrix} P_0 & P_F \\ 0 & I \end{bmatrix} \tag{2}$$

Note that because of the structure imposed earlier in the development, the state transition matrix can be partitioned in a convenient manner. The rightmost column ($P_{i\to j}$, $i = 0, 1, 2, \ldots, 2^{n-1}$) represents the failure probabilities from any of the nonfailure states and can be represented by the vector P_F. The lowest row is the probability distribution for the failure state, which is all zeros with the exception of the holding probability. The remaining submatrix represents the transitions among the nonfailure states (termed *operational states* because they represent the set of states that are possible during nonfailure operation). The submatrix of operational state transition probabilities is labeled P_O.

When considering this network topology it is important to keep in mind the issues of scale. The number of states in the safety state network grown as a power of 2 with an increasing number of conditions, and the number of elements in the state transition matrix grows as the square of the number of safety states. So, for example, a 10-condition network has roughly 1,000 states and 1,000,000 elements in the state transition matrix. When applying this method to actual systems the analyst must keep in mind the effects of scale and choose the conditions carefully to avoid having an unnecessarily large and unwieldy safety state network.

Estimating Risk

Although the state transition matrix itself is interesting, the ultimate power of this model lies in the ability to estimate the probabilities of future states. Consider the current state *S(i)* to be represented as a vector $\bar{\theta}(k)$ of dimension 2^{n+1} by 1, in which the ith element is 1 and the remaining elements are zero. (In this notation k represents the transition number as the process progresses in time.) One can calculate the probability distribution of the next state, shown as $\bar{\theta}(k + 1)$, using Equation 3.

$$\bar{\theta}^T(k + 1) = \bar{\theta}^T(k)P \tag{3}$$

Using this strategy one can look beyond the next transition to determine the probability of reaching a given state in any number of transitions in the future (Equations 4). This is a powerful concept, and by using this concept one can evaluate the probabilistic behavior as far in to the future as one would like. Future probabilistic behavior can be summarized in the $\Phi(\tau)$ matrix, which expresses the ability to transition from one state to another in τ transitions (Equation 5). Note that the $\Phi(\tau)$ matrix can be partitioned in exactly the same manner as the state transition matrix (Equation 2).

$$\begin{aligned} \bar{\theta}^T(k + 2) &= \bar{\theta}^T(k + 1)P = \bar{\theta}^T(k)P^2 = \bar{\theta}^T(k)\Phi(2) \\ \bar{\theta}^T(k + \tau) &= \bar{\theta}^T(k)P^\tau = \bar{\theta}^T(k)\Phi(\tau) \\ \bar{\theta}^T(k + \infty) &= \bar{\theta}^T(k)P^\infty = \bar{\theta}^T(k)\Phi(\infty) \end{aligned} \tag{4}$$

$$\Phi(\tau) = \begin{bmatrix} \varnothing_{0\rightarrow0}(\tau) & \varnothing_{0\rightarrow1}(\tau) & \dots & \varnothing_{0\rightarrow2^n}(\tau) \\ \varnothing_{1\rightarrow0}(\tau) & \varnothing_{1\rightarrow1}(\tau) & \ddots & \vdots \\ \vdots & & & \vdots \\ \varnothing_{2^n\rightarrow0}(\tau) & \dots & \dots & \varnothing_{2^n\rightarrow2^n}(\tau) \end{bmatrix} = \begin{bmatrix} \Phi_O(\tau) & \Phi_F(\tau) \\ 0 & I \end{bmatrix} = P^{\tau} \quad (5)$$

The real goal is to determine the probability of reaching the failure state from any given current state. Because the Markov network is a finite single-chain network the ultimate state will be the failure state. (This corresponds to the fatalistic notion that, given enough time, a probabilistic failure will eventually happen.) However, the theory of Markov processes provides a mechanism for calculating the mean time to another state from a known state. In the case of the safety state network, identification of the mean time to the failure state is the point of interest.

To calculate the mean time to failure one first needs to express the probability that the failure will occur on a specific state transition in the future. As shown in Equation 6, $\overline{\Psi}(\tau)$ is a vector quantity that provides the probability that the failure will occur on the τ^{th} state transition in the future as a function of the safety state. The mean time between failure as a function of the safety state is the expected value of the number of transitions until the failure (Equation 7).

$$\overline{\Psi}(\tau) = \Phi_F(\tau) - \Phi_F(\tau - 1) \quad (6)$$

$$MTBF = \overline{M} = \sum_{\tau=1}^{\infty} \tau\overline{\Psi}(\tau) \quad (7)$$

Knowing this, one can estimate the risk probability as the inverse of MTBF (Equation 8).

$$\overline{F} = 1/\overline{M} = \begin{bmatrix} 1/M_1 \\ \vdots \\ 1/M_n \end{bmatrix} \quad (8)$$

In summary, this section details the methodology used to derive the risk probability and mean time between failures given a static state transition matrix. Both the risk probability and mean time between failure are expressed as a function of the system's safety state. Provided there is an obtainable state transition matrix that characterizes average driver behavior, these results can be used to compare the performances of drivers in several ways. Some of these are discussed later.

Model Calibration

As shown, a state transition matrix that is characteristic of average driver behavior is required. To obtain that matrix (i.e., calibrate the model), observations from an existing system are used measuring the binary state of the individual conditions that are combined to form the safety state network. The measurement record marks the points in time at which each individual state change occurs. From these data a safety state trajectory can be constructed as a function of time. This trajectory includes all pertinent data for model calibration, including state occupancy times and state transitions.

By statistical analysis of these data the state transition probabilities are computed. First, state occupancy statistics are used to determine the holding probability, $P_{i\rightarrow i}$. The state transition statistics are then used to calculate the individual transition probabilities (Equation 9).

$$p_{i\rightarrow i} = (1 - p_{i\rightarrow i}) \left(\frac{\text{number of transitions to } S(j)}{\text{number of transitions from } S(i)} \right) \quad (9)$$

Driver Performance Measurement

Assuming that a state transition matrix of sufficient quality has been obtained, it is now possible to compare the performances of subject drivers. The safety state model is used to calculate the transformation function between the safety state and the estimated risk probability. This transformation function may be stored in lookup table form, with the safety state as the index in the table.

Data are collected on a subject driver in the same manner used for calibration data: by observing the safety state conditions and recording the changes of state. These data are combined to form a safety state trajectory as a function of time. By using the transformation function the safety state trajectory is then transformed into a risk probability trajectory.

The risk probability trajectory as a function of safety state is used directly for comparing driver performance. The average risk probability for the system average can be computed by taking the weighted average of the risk probability vector elements weighted by the relative time spent in that state (from the calibration data).

Several statistical approaches are useful for comparing the safety-related performances of subject drivers. These include instantaneous risk level, peak risk level, overall average risk level, windowed risk level, and cumulative risk level. In any of these approaches the performance of the subject driver can be compared with either the average risk probability (described in the previous paragraph) or the performance of other subjects.

The instantaneous risk level identifies the current risk probability estimation as a function of time. This measure is independent of history in that the specific safety state trajectory leading to the current state is not identified. (This is a general property of Markov processes.) It is a useful measure for evaluating performance in real time and could be used as feedback information to the driver as well.

The peak risk level identifies the highest level of risk that has occurred. Typically, peak measurements are made within some predefined time period, such as the duration of the test. This measure is useful in identifying the bound on the level of risk that a driver will take.

The overall average risk level provides an overall mean of the risk probability trajectory. To compute this the risk probability trajectory is integrated over time, and the resultant integration is divided by the integration time period. In the overall average the time period continues to grow. As a result the overall average represents a summary of the complete history of safety-related performance.

In contrast, the windowed risk level computes the average over a fixed time interval immediately preceding the current moment. For example, if the time window was defined as 10 min, the windowed average would provide the average safety-related performance for the last 10 min only. This measure "forgets" the past performance that is outside the defined time window and can be used as a means of measuring learning curves or fatigue characteristics.

Finally, the cumulative risk level represents the total amount of risk that has been taken. It is computed by taking the time integral of the risk probability trajectory and represents the expected number of accidents (in a Bernoulli sense). This measure is not intended to predict the occurrence of accidents—the expected value represents an average. However, it can provide useful insight into risk exposure. This measure would not be appropriate for subject feedback, because subject knowledge of this information could significantly affect the subject's performance.

Discussion of Results

The safety state model provides a method for estimating the dynamic risk probability as a function of system state. It has several strengths when compared with event and fault tree methods; however, it has weaknesses as well.

The safety state model is fundamentally different from both event and fault tree analyses in the conceptual definition of the nodes. In both event and fault tree methods the nodes of the tree represent events. In the safety state model the nodes represent system states, and the transitions between the nodes represent events. This is significant, because any state can be reached by several different paths, which represent the occurrence of different events. Once a state has been reached (that is, once a given set of conditions is true), the path to that state (the order of events) does not matter. This is a fundamental concept in Markov process modeling.

Significant among the strengths of the safety state model is its generality. No assumed dependencies or interactions exist between any of the contributing conditions. All combinations of conditions are considered equivalently. Thus, the analyst needs to specify only the conditions that define the safety state, without defining the relationship between them. This serves to improve the robustness of the analysis by reducing the chance of human error (which might be due to either prejudice or oversight). This becomes more significant as the number of conditions increases.

As a result of its generality the method of safety state analysis is easily automated. This frees the analyst from any direct contact with the state transition matrix, which will be of formidable size in any significant problem. The ease of automation allows safety state analysis to be applied to problems that are too large and unwieldy for other methods. However, this strength is tempered by memory and computation requirements, which are significant.

One of the most significant weaknesses is in the area of data collection. In an operational system, the data required for calibration would be quite difficult to gather. For driver performance evaluation a broad range of in-car data would be required. Systems in current operation do not collect appropriate data, and there is debate whether measurements of this nature are even feasible within commonly accepted notions of personal privacy. An alternative approach is to use data collected from simulation systems. Simulation systems can be configured to provide a plethora of appropriate data, and networked simulators provide a mechanism for collecting data on interactions between two or more vehicles. Although there are debates in the driver performance community regarding the applicability of simulator-based results to operational systems, the author believes that simulation-based experiments provide the only currently available means for exploring the viability of this method.

AN EXAMPLE

To illustrate the method of safety state analysis, consider the following example: a frontal impact scenario in a rail system. In this example the goal is to identify the risk probability of striking another vehicle (or obstruction) with the front of the vehicle. The first step is to select the set of conditions that are considered contributory to this failure event. The following set of conditions are used:

- Condition 0: throttle actuated. This condition is true if the throttle is applied and false if the throttle is not actuated.
- Condition 1: brake applied. This condition is true if the brake is applied and false if the brake is not actuated.
- Condition 2: brake failure. This condition is true if the braking system of the train has failed in any way (even if the driver is not aware of this condition) and false if the braking system is functioning properly.
- Condition 3: overspeed. This condition is true if the train is traveling at a speed greater than that allowed by either the static speed limit or the signaling system and false if the train is being operated within the allowable speed bounds; the overspeed condition includes the case of passing a stop signal (entering an occupied block).
- Condition 4: obstacle. This condition is true if there is an obstruction on the track (possibly another vehicle) within the stopping distance of the train and false if the track is clear; for the purposes of this example the stopping distance is defined as the distance required to bring the train to a full stop from the current speed by using full-service braking and is a function of speed.

This set of conditions includes measurement of human control actions (throttle and brake actuation), vehicle state as a result of human control actions (speed condition), on-board equipment failures (brake failure), and external conditions (obstacle). Since each condition is binary, this set of conditions can be combined into a single number. The resulting set of safety states is given in Table 1, with the operational states numbered from 0 to 31. The check marks indicate the conditions that are true for each state. For example, the system is in state 14 when there is no obstacle present but the vehicle is over the speed limit, the brakes are applied, and the brakes have failed. An additional state is included in Table 1, state 32. This state represents the occurrence of the failure event, which in this example is a collision.

To calibrate the model data are collected from an operational system. Once they are collected these data are processed to determine the state transition matrix, which is used to calculate the risk probability as a function of the safety state. Subjects are evaluated by recording the safety state information, converting the resultant safety state trajectory into a risk probability trajectory, and comparing the risk probability trajectories, as discussed previously.

SUMMARY

The research described here is a response to a need for methods in which the safety-related decision performances of vehicle operators can be evaluated. Some of the difficulties in this area include event rarity, compound and interacting errors, and related difficulty in determining causality. This work identifies the human operator as a key component in the safety of transportation systems. The driver-vehicle system is modeled as a closed-loop control system, and a probabilistic model of system behavior is presented.

Based on the work of Lowrance (*3*) an organization for the efforts involved in the pursuit of safety is identified. Safety-related work is divided into subjective and objective components. Risk assessment, the objective component, is further divided into two components. One component, risk probability, measures the probability of occurrence of a necessary intermediate event, whereas risk outcome measures the outcomes of these events in terms of the ultimate risks. In the case of transportation systems, the ultimate undesirable outcome is damage, injury, and death, and the intermediate event is

TABLE 1 Safety State Description

Safety State	Binary Form	Obstacle	Overspeed	Brake Failure	Brake Applied	Throttle Applied	Failure Possible?
0	00000						
1	00001					√	
2	00010				√		
3	00011				√	√	
4	00100			√			
5	00101			√		√	
6	00110			√	√		
7	00111			√	√	√	
8	01000		√				
9	01001		√			√	
10	01010		√		√		
11	01011		√		√	√	
12	01100		√	√			
13	01101		√	√		√	
14	01110		√	√	√		
15	01111		√	√	√	√	
16	10000	√					√
17	10001	√				√	√
18	10010	√			√		√
19	10011	√			√	√	√
20	10100	√		√			√
21	10101	√		√		√	√
22	10110	√		√	√		√
23	10111	√		√	√	√	√
24	11000	√	√				√
25	11001	√	√			√	√
26	11010	√	√		√		√
27	11011	√	√		√	√	√
28	11100	√	√	√			√
29	11101	√	√	√		√	√
30	11110	√	√	√	√		√
31	11111	√	√	√	√	√	√
32	100000						

an accident. Based on this organizational description, the focus of this research is in the area of risk probability assessment.

A probabilistic model for system behavior (the safety state model) is developed. This model is based on finite Markov processes, with event tree and fault tree techniques used as inspirations. From the safety state model a method for determining MTBF and risk probability is developed. Both of these quantities are expressed as a function of system state. A method for calibrating the safety state model is presented and is based on experimental data. Finally, a method for measuring individual driver performance and comparing it with average driver behavior and the behaviors of other individual drivers is presented.

In conclusion, the safety state model represents a unique method for assessing the safety-related decision performances of vehicle drivers. Future developments include experimental verification of the usefulness of the model with data gathered from a human-in-the-loop high-speed rail simulation system.

ACKNOWLEDGMENTS

The author extends his deepest appreciation to the many colleagues who took time to review this paper and to contribute their constructive commentary. This work was prepared with support from the Volpe National Transportation Systems Center, Cambridge, Massachusetts.

REFERENCES

1. Sheridan, T. *Telerobotics, Automation, and Supervisory Control.* MIT Press, Cambridge, Mass., 1992.
2. Forbes, T. W. (ed.). *Human Factors in Highway Traffic Safety Research.* Wiley-Interscience, New York, 1972.
3. Lowrance, W. W. *Of Acceptable Risk.* William Kaufmann, Inc., Los Altos, Calif., 1976.
4. Rowe, W. D. *An Anatomy of Risk.* John Wiley and Sons, Inc., New York, 1977.
5. Rescher, N. *Risk: A Philosophical Introduction to the Theory of Risk Evaluation and Management.* University Press of America, 1983.
6. Gratt, L. B. Risk Analysis or Risk Assessment: A Proposal for Consistent Definitions. In *Uncertainty in Risk Assessment, Risk Management, and Decision Making*, Plenum Press, New York, 1987.
7. Wharton, F. Risk Management: Basic Concepts and General Principles. In *Risk Analysis, Assessment, and Management* (J. Ansell and F. Wharton, eds.), John Wiley and Sons, Inc., New York, 1992.
8. Reason, J. *Human Error.* Cambridge University Press, Cambridge, 1990.
9. Swain, A. D., and H. E. Guttman. *Handbook of Human Reliability Analysis with Emphasis on Nuclear Power Plant Applications.* Sandia

National Laboratory Report NUREG CR-1278. U.S. Nuclear Regulatory Commission, Washington, D.C., 1983.

10. Lewis, E. E. *Introduction to Reliability Engineering.* John Wiley and Sons, Inc., New York, 1987.
11. McCormick, N. J. *Reliability and Risk Analysis.* Academic Press, 1981.
12. Gertman, D. I., and H. S. Blackman. *Human Reliability & Safety Analysis Data Handbook.* John Wiley and Sons, Inc., New York, 1994.
13. Ansell, J. Reliability: Industrial Risk Assessment. In *Risk Analysis, Assessment, and Management* (J. Ansell and F. Wharton, eds.), John Wiley and Sons, Inc., New York, 1992.
14. Joshua, S., and N. Garber. A Causal Analysis of Large Truck Accident Through Fault Trees, *Risk Analysis,* Vol. 12, No. 2, 1992.
15. Kuzminski, P., J. S. Eisele, N. Garber, R. Schwing, Y. Y. Haimes, D. Li, and M. Chowdhury. Improvement of Highway Safety I: Identification of Causal Factors Through Fault-Tree Modeling. *Risk Analysis,* 1995.

Publication of this paper sponsored by Committee on Simulation and Measurement of Vehicle and Operator Performance.