Intelligent Vehicle-Highway System Safety: Problems of Requirement Specification and Hazard Analysis

ANTHONY HITCHCOCK

When accidents occur on a fully automated freeway, driver error will rarely be relevant. Design should not err in ensuring that equipment failure will not cause hazards. An attempt to demonstrate satisfactory methods of requirement specification and hazard analysis is described. The case treated is a single automated lane on a freeway that also has lanes for other vehicles and shares on- and offramps with them. When hazards are specified, it appears that the system configuration is determined. A number of concepts have emerged and seem likely to be basic to all hazard-avoiding designs. These concepts are identified and described. No method has been found for demonstrating that the set of hazards is complete. Peer criticism of those proposed is therefore earnestly sought because, as explained, these are the axioms on which the logical structure that demonstrates safety is to be based.

One form of the intelligent vehicle-highway system (IVHS) described by Mobility 2000 (*1*) is the advanced vehicle control system (AVCS). In its more advanced forms (AVCS-2 and AVCS-3), this term refers to systems in which, for a journey on a freeway, the driver of a vehicle surrenders control to a fully automated system, possibly retaining control of the route or merely specifying a destination. AVCS-2 and AVCS-3 are not seen as feasible options for off-freeway traffic; it would be neither possible nor desirable to automate the movements of pedestrians or other local traffic, and it does not seem feasible to create a hazard-free system containing such elements.

The proposed method begins by considering factors that can go wrong with a system and evaluating the consequences, together with an estimate of their probability. It will be worthwhile to attempt to design out some of these possibilities, whereas others are not sufficiently likely. By means of this preliminary hazard analysis, hazards of different degrees of seriousness are defined. Provided nothing has been overlooked, a safe system will result if the design is such that the hazards do not arise even if the system fails. This method, known as hazard and operational analysis (HAZOP or HAZAN) (*2*), is similar to one widely used in the chemical, nuclear, petroleum, and aerospace industries.

For fully automated freeways, there is no current experience from which probabilities can be assessed. However, there is obviously no way of guaranteeing that a car will not fail in operation, particularly if the owner is responsible for maintenance. If a vehicle fails in such a way that it decelerates more rapidly than a following vehicle can by braking, there is a clear danger of collision. Shladover (*3*) has shown that, in such a case, a collision can be avoided only if vehicle spacings are much greater than normal. Shladover also states that, if the vehicles are very close together (say 1 m or less), the collision only involves a small velocity change (delta-V) and is likely to result in property damage only. From this assumption comes the concept of widely separated platoons, in which vehicles are separated either by a very small or a very large spacing.

It is postulated that catastrophic hazards are those system states that are precursors of high–delta-V collisions, whereas critical hazards are the precursors of low–delta-V hazards. The design should exclude the former, even if a vehicle or some other part of the system fails or is circumvented deliberately. The incidence of critical hazards should be reduced as far as reasonably possible.

A set of hazards has been formulated. There does not seem to be any systematic method of demonstrating that the list is complete, although it is reasonably simple to demonstrate that a proposed condition is or is not a hazard. The list is therefore presented for peer criticism.

It is now permissible to say that a system is safe if it can be proven to avoid the defined hazards (except perhaps in some known cases of low probability or an excessive number of coincidental simultaneous failures). If the system is fully defined, the question of whether it is safe can in principle be answered definitively as a result of a process of mathematical logic. If a system can be shown to be safe in so unambiguous a fashion, it is clearly desirable to do so. There is thus a need to verify the definitions of hazards, which have the same relation to this analysis as do Euclid's axioms to his theorems.

One possible system that will be subjected to such analyses is described. To avoid the hazards that have been defined, the basic configuration seems to be constrained to a handful of possibilities. Further, some concepts have emerged that seem likely to be common to all or many such systems. Although the hazard analysis has barely begun (and is not reported), it is worth recording these results.

In the system considered, there is one automated lane in each direction, and that lane operates at a capacity of at least twice that of a nonautomatic lane. (If the system is to be economic, its capacity must be greater than that of a nonautomatic one.) If there is a network of such lanes, changes between links are not made automatically. This assumption

Program on Advanced Technology for the Highway, Institute of Transportation Studies, University of California (Berkeley), 1301 South 46th Street, Building 452, Richmond, Calif. 94804.

seems reasonable for the initial stages of exploitation of the technology. In Mobility 2000 terminology (*1*), this system is an AVCS-2 system.

APPROACH TO HAZARD ANALYSIS

Modular Structure

An AVCS system consists of a large number of separate, distributed control systems that interact with each other over time. Some of these systems contain hardware and produce mechanical output—for example, the control system that controls the steering axle. Others, like the program segment in a roadside controller that admits a vehicle, are all software in that their output is information sent to other parts of the system. It seems useful in the current context, however, to treat this difference as of little importance and to regard all modules simply as modules, with precursors, effects, and stimuli to other modules, possibly at later times, without regard to the electronic, electric, or mechanical nature of the means by which the effects are achieved.

This modular property of the requirement specification is vital to freedom from hazard in practice. Different manufacturers will achieve the desired effects in different ways, and, as time passes, many improvements will be made to each module. Modular construction of the specification eliminates the need to reassess the safety of the whole system every time a change is proposed.

Thus, it is both necessary and possible to regard the system as an assembly of modules, each of which must be fully specified insofar as its precursors, effects, and consequents are concerned and which must in total add up to a complete specification. A complete specification means that, for any conceivable configuration, either the specification will indicate unambiguously what happens next or it will be possible to demonstrate by deductive logic that no possible combination of circumstances could have preceded that configuration. This requirement is formidable, and it is not discussed further here.

The modular approach described implies an analogy between a distributed system and a software suite, the modules being paralleled by software procedures. This approach, in turn, opens to application the body of techniques for safety-critical software described in the literature (*4*).

Safety-Critical Subsystem

Some parts of an IVHS system are essentially concerned with actions that are critical to safety; others are not. A design principle recommended in relevant standards (*5,6*) and adopted here is that the safety-critical subsystem be kept as small as possible and rigorously separated from the rest, so that interaction between the two occurs only at clearly defined modules.

The focus here is not on the non-safety-critical part of the system, which does not mean that it is peripheral to the system as a whole. The entire subsystem concerned with controlling the movement of vehicles so that capacity is maximal (i.e., link control) can, and should, be kept separate from the safety-critical elements. So should the process (i.e., strategic control) by which drivers indicate where they want to get off and are given information about any parts of the system that are not functioning properly. These processes are vital to the proper functioning of the system but must be made peripheral to the hazard analysis by appropriate specification of the requirement.

One component of the system that is not modular and whose status as part of the safety-critical subsystem can be argued is the relevant law. Law is certainly part of the system, for the mechanisms may not prevent what the law admits as permissible. Thus, the designer must specify legal and administrative requirements along with software procedures or bounding conditions for stability of a control system. If an action violates the law, it may be hoped that its frequency will be reduced; however, it would be foolish to assume that the action will not occur. Indeed, the system should be designed, so far as possible, to make it difficult (as well as detectable and provable beyond reasonable doubt) to indulge in dangerous joy riding or otherwise to use the system in a way that endangers others.

There are possibilities of more organized wrongdoing. Someone could think it fun to sabotage the system—for example, by causing all police cars to run round in ever-diminishing circles or to track the movements of an erring spouse. Others may see the traffic control system of the wealthiest cities in the world as a natural target for terror. The system should be designed to protect against these possibilities, too.

It is assumed that the law will require the following, but it is not assumed that all drivers will conform:

1. Every vehicle entering the automatic system must have a license certifying that control systems conforming to a legally defined specification were present and working on some date past and that they have not been subsequently invalidated for some wrongdoing. (The period of the license is not critical to this discussion.) The control systems must be present and functional within defined tolerances at the time of attempted entry. It will be an offense to carry equipment on the vehicle that can issue signals similar to those specified for the system. (There will need to be a body of law and procedure for approving new and improved control devices.)
2. It will be proper for the license to be invalidated mechanically, with the possibility of appeal to a human court (and, presumably, financial compensation if the automatic checking device made an error).
3. It will be an offense to fail to be ready to resume manual control on reaching the designated point of departure or the end of the system. There will be a storage place, known as the dormitory (explained later) at the extreme downstream end or at convenient intermediate points. (It might be urged that those who end up in the dormitory create a rebuttable presumption that they have more than 80 mg or 100 ml of ethanol in the bloodstream, but this element is not a safety-critical one.)
4. It will be an offense to enter the automated system while carrying an external load or with a trailer that does not have appropriate control and communication equipment.

5. Devices that will be carried by emergency vehicles to permit abnormal movements may not be used on other vehicles, nor may their signals be forged.

Considerable controversy about the evidential nature of the system records in the absence of human corroboration will likely take place.

Definitions

It is assumed that the following vehicle controls will maintain the platoon formation:

1. A lateral control system will keep the vehicle in its lane by means of a passive reference line in the road. The reference can be coded to give data that will reduce lateral jerk by warning of changes in curvature and also to indicate the positions of off-turns.
2. A longitudinal control system will use an active sensor to determine the distance to the vehicle ahead and its time derivatives. The range at which the sensor can detect a vehicle ahead may be limited, either because it will not be able to detect a vehicle around a horizontal or vertical curve or by its nature. The maximum distance at which the sensor can certainly detect an object ahead is defined as the sensor-range spacing. This distance may depend on the weather. The system supervisors may choose to select a shorter spacing.
3. A vehicle-to-roadside communication system will be included. It will be of short range so that the noise level at any point ascribed to the presence of other communications (e.g., from parallel lanes or other vehicles at some distance) is small.

The first two controls are chosen in the light of current developments. The idea that communication ranges should be short to reduce noise is the author's. If this feature turns out to be critical, it should be checked.

It is also useful to define the following quantities, all of which are weather-dependent and most of which depend on the speed of the vehicle concerned:

1. Full platoon braking is the maximum deceleration that a platoon can reasonably be expected to undergo without collisions or a break in formation. (If each vehicle in a platoon brakes at its maximum rate, the rates will differ, and there will be low–delta-V collisions.)
2. Platoon spacing is the distance that must separate two platoons if the leading one suffers a deceleration considered to be the worst likely in case of catastrophe, the other simultaneously undergoes full platoon braking, and there is to be no collision between platoons.
3. Sensor-range spacing is the maximum distance at which the sensor can detect an object ahead (as defined previously).
4. Sensor-range speed is the greatest speed from which a vehicle can come to rest within sensor-range spacing.
5. Manual spacing is the minimum spacing at which drivers feel able to control their vehicles.

The actual values of these quantities for any time and place may be inserted by the system controllers or may be derived from automatic weather recorders, surface friction monitors, and so on.

Hazards

An objective of design is to prevent accidents. An accident, however, has many contributing factors, some of which are random and many of which are outside the designer's control. Therefore, hazards are defined as those factors that, if avoided, will mean that accidents cannot occur. Here, a distinction is made between catastrophic hazards, which are to be avoided in all cases, and critical hazards, which are to be avoided when possible but whose consequences are less serious.

Catastrophic hazards are those conditions that necessarily precede collisions at high delta-V as follows:

- *Hazard A.* A platoon (or single controlled vehicle) is separated from one ahead of it, or from a massive stationary object in its path, by less than platoon spacing.
- *Hazard B.* A vehicle, not under system control, is an unmeasured and unknown distance in front of a platoon or a single controlled vehicle.
- *Hazard C.* A vehicle is released to manual control before the driver has given a positive indication of acceptance.
- *Hazard D.* A vehicle is released to manual control at less than manual spacing from the vehicle ahead of it or at such a relative speed that a spacing less than manual spacing will be realized within, for example, 2 sec.

The following conditions are critical hazards—those that can precede low-delta-V collisions (it is uncertain whether Hazard E should be categorized as catastrophic or critical):

- *Hazard E.* A vehicle under automatic control in the transition lane is less than manual spacing behind a vehicle not under automatic control.
- *Hazard F.* A vehicle within a platoon suffers an electrical or mechanical failure that causes deceleration greater than full platoon braking, does not respond to controls (longitudinally or laterally), or fails to communicate.
- *Hazard G.* A vehicle goes too fast when joining a platoon.

Other hazards exist, such as illegal equipment that can pass false messages to the system controls, interference with the control computers, explosives, heavy weights dropped from bridges, and other deliberate acts. Design features may be desirable to circumvent such activity, but, because they are unlikely to interact with the design of the system as a whole, they are not relevant here.

With these exceptions, these hazards may be sufficient for analysis. That is, if these conditions are avoided, there will be no injury accidents resulting from the actions of vehicles in the system. The system can have no impact on the actions of vehicles it does not control, and these vehicles may be involved in accidents.

Design Consequences

Because the automated lane has a capacity more than twice that of a normal lane, it must be possible to enter it at the side—it cannot achieve capacity if fed from one end only. There is thus no need for a special entrance at the upstream end, which simplifies upstream extension of the system. At the downstream end, however, provision must be made for those drivers who fail to take control from the system (Hazard C). The system must bring them to rest in a specially designated length of the freeway, which is called the dormitory. Vehicles in the dormitory are packed together closely and are moved up as drivers take control of their vehicles and drive off.

In normal operation, at the ends as well as at any intermediate point, entrance and exit will be from a side lane, which is called the transition lane, or TL. Similarly, the automated lane is called the AL.

To avoid Hazard A, no stationary object must be allowed to enter the AL. For example, this situation would happen if an accident occurred on the manually operated lanes and wreckage was pushed onto the AL. Therefore, there must be a barrier between the AL and the TL. Access must be provided either at a limited number of gates or along a limited length of the TL that is shielded from the rest of the freeway by a fence. The latter, however, does not avoid Hazard A; if there were a low-speed collision, the debris might extend from the AL or TL to the other and act as a stationary obstacle. Therefore, only the former case should be considered, as shown in Figure 1. (The arrangement does not avoid Hazard A completely because the gates might admit wreckage from an accident at exactly the wrong place.)

A vehicle will have to make its way across several lanes of traffic before joining a platoon on the AL, and it must do so again when it leaves the freeway. The entrances and exits therefore do not need to be associated with actual on- and off-ramps and need not be confined in length to the length of any additional lane provided at the on-ramp. Indeed, it must not be possible to run out of room on the TL before joining the AL, as it is with physical on- and off-ramps. If none of the gates making up the logical on-ramp (LONR) can admit a vehicle, the driver can revert to manual control and approach the next gate. If a hazard would otherwise arise, a vehicle can also be refused exit at the gates of the logical off-ramp (LOFR). The terms LONR and LOFR encompass not only the physical gates but the total control subsystems associated with joining and quitting the AL.

As a vehicle (controlled or not) enters the gate in the barrier, there is a danger that it will strike the barrier and come to rest partly on the AL. This situation would induce Hazard A, unless the post is instrumented to warn the control system to reduce speeds and unless a vehicle only enters to join onto the rear of a platoon that has just passed. This safeguard would imply that vehicles only enter the AL after they are controlled by the system. It means that Hazard B is avoided, and, by making it possible to test vehicles on the TL before they join the AL, also avoids Hazard F on the AL, where its effects are the most serious.

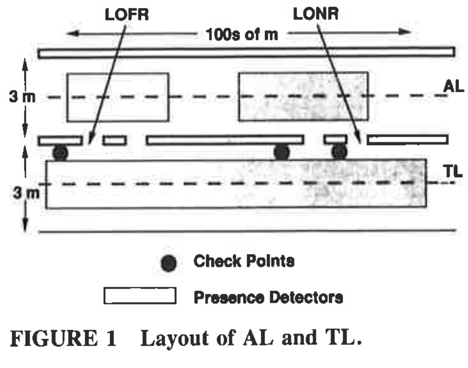

FIGURE 1 Layout of AL and TL.

Thus, the TL must also be equipped to control vehicles whose drivers have indicated a wish to join, at least near the gates. Avoidance of Hazards C and D implies that the TL is also equipped to control a vehicle after it has left the AL. (This equipment includes the LONRs and LOFRs.) Further, with admission only at the rear of a platoon on the AL, LONR capacity will clearly be limited. It seems likely, therefore, that it will be desirable to preform partial platoons on the TL. Hazard avoidance permits this action, and it will be assumed in the following paragraphs.

The previous discussion may have implied that gates are used for entrance or exit, but not both. This restriction may not be necessary, except that the last gate must be an on-only gate. Then, a vehicle that has left the system but whose driver has not resumed control can be readmitted to avoid Hazard C.

The TL will clearly be used by manually controlled vehicles that wish to join or have just left the system. There is therefore no way of excluding vehicles that do not wish to join or are not equipped to do so, and it does not seem right to try. It follows that, to avoid Hazard B, there must be a length of TL, controlled by LONR and LOFR, in which the position of vehicles can be determined by presence detectors. The length will stretch some distance, perhaps 1000 m, upstream of the first gate.

This feature is necessary not only to avoid Hazards B, D, E, and G on the TL but to provide data by which the strategic control system can function. (Strategic control optimizes LONR and LOFR capacity, in part by advising which vehicles should preform platoons, but it is not safety critical.) The presence detectors also make it possible to identify which vehicle is seeking entry and to carry out tests on it (e.g., to ensure that its dimensions correspond to those on its license so that it has no external load).

Requirement Specification of Roadside Configuration

A general specification has been developed for a system that may meet the required criteria of hazard avoidance. The specification consists of the following (see Figure 1):

- An AL will be separated from the rest of the freeway by a barrier. The AL will have a lateral-control reference along its length and will be equipped with a means of communicating with the vehicles on it. There will be an appropriate communication protocol. At its downstream end, there will be a dormitory for storage of vehicles that have failed to resume manual control.
- A series of local control systems, made up of a LONR and a LOFR, will be associated with gates in the barrier. There will also be a length of instrumented TL adjacent to the AL, controlled by the LONR and LOFR. It will stretch upstream from a point a little downstream of the last gate (an on-only gate), to a point perhaps 1000 m upstream of the first gate. The TL will also bear a lateral-control reference and communication equipment. Vehicles entering the AL, pos-

sibly preformed into partial platoons, will join on-AL platoons only at the rear.

- The TL, and that part of the AL immediately adjacent to the gates, will be equipped with presence detectors and a control system that can track identified vehicles through the LONR and LOFR.
- There will be a control system for the AL. Although not necessary, it would seem convenient to associate the length of AL between two access points with the relevant LOFR and LONR to form a block. If so, there will need to be communication between adjacent blocks.
- A centralized strategic control system will propose how speeds will be regulated, through which gates vehicles will enter or leave the AL, and possibly more. It will not be safety-critical because its advice will be mediated by the local block controllers, who will only pass on its suggestions if it will not create a hazard to do so. Any human intervention that may become desirable in normal operation will be through this system. For example, one way in which the local controller may operate is for the vehicle control system to recognize three speeds as relevant to its own speed. One would be the speed of the vehicle ahead, which in a platoon will not be exceeded for very long by very much. The second would be a target speed, set by strategic control, to ensure that, for example, a closing member of a new platoon will come within sensor-range spacing or that the front vehicle of a platoon will reach an on-gate just far enough ahead of a joining vehicle that it can join onto the back. The third would be a maximum speed, set by the local controls to ensure hazard avoidance. It will override the target if demanded by hazard avoidance.
- A body of law regulating permissible behavior on the AL and TL and specifying the equipment that must, may, or must not be present on vehicles will also be part of the system.

Other features may well be present; indeed, those just described are necessary but not sufficient for hazard-free economic operation. From here, it seems that the choices broaden. The specifications presented assume that there will be centralized strategic control (and perhaps self-monitoring), but that safety-critical controls will be local. The actual location of controllers has not been shown to be safety-critical, and may not be, but this arrangement is common on many systems with safety-critical elements.

Uniqueness of Configuration

Beginning with the concept of a single automated lane bearing platoons, and accepting the constraint of hazard avoidance, the configuration of the system is almost completely determined. Therefore it must be generally agreed that safety does indeed demand that the conditions defined as catastrophic hazards be avoided. Simple as they seem, they greatly constrain design choices.

The only other alternative is one in which Hazard B is avoided on the AL, not by insisting that vehicles be under system control before admission but by including presence detectors throughout the whole length of the AL. Such an arrangement would presumably not allow unequipped vehicles to enter the AL legally.

It may be possible to design a system as free of hazards as that discussed here. Avoidance of Hazard E will be difficult and seems likely to make it necessary for control to be passed to the system just before entry; in some reduced form, TL, LOFR, and LONR would continue to be relevant concepts. It seems probable, however, that much of the potential capacity advantage would be lost with such a system. Certainly no reduction in lane width inside the AL barrier would be possible.

However, this possibility is not conclusive and needs further examination. It again seems likely that the constraints imposed by hazard avoidance will be strict.

SYSTEM SPECIFICATION

Operational Modes

It will be a principle of design of any system that, should accidents happen or faults develop, the system will undergo graceful degradation. In other words, the system will pass through a series of modes in which it still operates, though less efficiently, but maintains its hazard-avoidance function. If an accident is severe or the failures are serious enough, vehicles must be brought safely to rest.

A series of operational modes can be envisaged in which any block may operate. The block is the relevant unit. Nothing smaller is possible because all vehicles on the AL in one block must remain there until they reach a gate, and it is clearly not always necessary to take exceptional action in one block because of difficulties in an adjacent one.

In designing a system for safety analysis, seven possible modes of operation have been chosen. This amount seemed initially to be unnecessarily complicated, but no argument has been found that any one of the modes is not required, if any intermediate mode between normal operation and full arrest is to be admitted.

The modes in which traffic continues to move, but at reduced speed, are ones in which all platoons and controlled vehicles on the AL move at sensor-range speed. This speed makes it possible for each platoon to stop safely if there is an obstruction.

The capacity of these modes will depend strongly on the sensor range. As Shladover (*3*) has shown, the capacity of an AL operating with platoon spacings defined as they are here is a function of speed, which has its maximum at 30 to 50 mph. The system controllers will usually set a maximum speed above this limit; however, if the detector range is sufficient, platoons moving at sensor-range speed will give a capacity not less than that at normal speed. Certainly, the loss of capacity may not be unacceptable.

The seven modes are as follows:

1. *Normal.* There are no problems requiring reduced speed or special observation.
2. *Natural.* The system is completely shut down. Vehicles travel on all lanes without any form of system control. This mode needs no hazard analysis.
3. *Sensor-Range–Continue.* Platoons on the AL are constrained to sensor-range speed but, at the end of the block, progress to the next one, which may be in normal mode.

4. *Sensor-Range–Exit.* Platoons on the AL are constrained to sensor-range speed and must leave the AL at the next LOFR.

5. *Crash-Stop.* This mode is used when the presence detectors at the gates indicate a stationary object on the AL or when a signal from the barrier indicates that it has been breached. All platoons brake to rest at full platoon braking. However, if one platoon is less than platoon spacing from, but not at, the breach, each vehicle goes to maximum braking.

6. *Stop.* This mode can only be activated by the system controllers. It is used, for example, when the AL contains paramedics on foot or when debris must be cleared. A number of special functions are possible: a vehicle identifying itself as an emergency vehicle can enter the downstream off-gate and be guided backward down the AL, or a vehicle equipped with special markers (provided by the highway patrol) can be guided backward to the last on-gate of the previous block.

7. *Resume:* Again, this mode is activated only by the controllers. Its effect is to successively accelerate platoons moving in a block on one of the sensor-range modes or the stop mode back to normal-mode conditions, and then to allow the platoon to reenter normal mode.

There are a number of logical links between modes. For example, if a section goes to crash-stop mode, the upstream one must go to sensor-range–exit mode.

Further Development

The seven modes of operation introduce considerable complexity, and a great many choices, into a specification. As operational detail accumulates, the complexities multiply. A system has currently been specified, and fault tree analysis to prove the absence of unforeseen hazards has begun. In the operational area no parallel to the vast simplification that is possible by applying hazards analysis to the physical configuration has been found. When the analyses are finished, they will be reported.

CONCLUSIONS

This preliminary work has revealed a number of results of general importance about AVCS-2 and AVCS-3 systems, with the conclusion that the hazards listed must indeed be avoided.

The possible physical configurations of the system are extremely limited. The hazards do not constrain the form of the system completely, but the number of possible configurations is small and probably includes some that cannot achieve the required capacities. The use of the permitted configurations is necessary but not sufficient for hazard avoidance. There are also operational limitations.

The modular construction of the system and the concept of a strategic controller that is not safety critical, but whose commands are mediated through a local safety-critical controller, seem basic. (These features are also true for many other systems, such as area traffic control.)

The concepts of AL, TL, LONR, and LOFR certainly have wider application and may well be components of all hazard-free systems.

Operations will be of considerable complexity. Command sets and instruction sets are likely to differ among most cells of a matrix of vehicle modes versus modes containing between 50 and 100 cells. It does not seem possible to propose a simple, hazard-free system.

ACKNOWLEDGMENTS

The research reported here is part of the Program on Advanced Technology for the Highway, prepared under the sponsorship of the State of California Business Transportation and Housing Agency, Department of Transportation, and the U.S. Department of Transportation, FHWA and NHTSA. The author would like to thank his colleagues at the University of California (Berkeley), notably Anna Bozzini, Susan Graham, Robert Parsons, and Steven Shladover, for helpful discussion and other assistance.

REFERENCES

1. *Final Report of the Working Report on Operational Benefits.* Mobility 2000, Dallas, Tex., 1990.
2. T. A. Kletz. *HAZOP and HAZAN Notes on the Assessment and Identification of Hazards*, 2nd ed. Institution of Chemical Engineering, Rugby, England, 1986.
3. S. E. Shladover. *Operation of Automated Guideway Transit Vehicles in Dynamically Reconfigured Platoons.* Report UMTA-MA-06-0085-79-1, 2, and 3. UMTA, U.S. Department of Transportation, 1979.
4. N. Leveson. Software Safety: Why, What and How. *Computing Surveys*, Vol. 10, 1988, pp. 125–163.
5. *Procedures for Safety-Critical Software.* Report D-0178A. Radio-Technical Commission for Aeronautics, Washington, D.C., 1986.
6. *Software for Computers in the Application of Industrial Safety-Related Systems.* Report IEC 65A SEC94. International Electrotechnical Commission, Geneva, Switzerland, 1990.

Publication of this paper sponsored by Committee on Vehicle User Characteristics.