

# Intelligent Vehicle-Highway System Safety: A Demonstration Specification and Hazard Analysis

## A. HITCHCOCK

A complete specification and fault-tree analysis on one concept for an automated freeway system are described. The example chosen includes a single automated lane on a freeway that also admits manually driven vehicles. Proper execution of all maneuvers is independently verified by infrastructure instruments. There is maximal intelligence within the infrastructure. It is planned to do this again with different starting assumptions. This is only the first example in a set of two or three. Hazards have been specified. A safety criterion has been chosen by way of example. A design has been specified. A fault-tree analysis is also described. This analysis attempts to verify that the design does satisfy all criteria. The example demonstrates that full specification is possible and that design errors (there were four) can be detected by fault-tree analysis. After further development, the technique of full specification and fault-tree analysis can become a basis for safety standards that will apply to both design methods and verification of conformity to safety criteria. The initial assumptions were few and broad. Along with the need to avoid violation of the hazard criteria, they determine a very small set of possible system design structures. There is more than one safe way to design an automated road, but there is not an abundance of options.

Hitchcock has indicated (1) that the safety of an automated freeway can be examined, at the conceptual stage, thus

1. The first stage is to identify the safety-critical subsystem (S-CS). This should be of modular design so that each module can communicate with others only by defined protocols. In particular, each module in the S-CS has a fully specified interface with modules not in the S-CS. Thus, malfunction of a non-safety-critical module cannot cause danger. The work of Varaiya and Shladover on system architecture is relevant here (2).
2. Next, the S-CS must be specified completely. This means that what happens next in any condition of the system whatsoever may be determined.
3. Hazards must be specified. Here, a catastrophe is a collision that may cause death or injury, that is, a high-delta-V collision, and hazard is a condition in which a failure can result in a catastrophe.
4. A safety criterion is now selected. Fault-tree and other forms of analysis should be used to verify that it is satisfied (3).

A demonstration specification and a fault-tree analysis have been completed and reported (4,5). The safety criterion was that, for a hazard to occur, two independent faults must occur. Because a hazard is the precursor of a high-delta-V crash, three independent faults are necessary before anyone can be hurt as a result of automation. The specification and analysis are too complex to be reproduced here in detail. They demonstrate that the process just described is sufficient to ensure safe design at the conceptual stage.

Ultimately, the logical patterns described must be quantified. Safety criteria can then be stated in terms of catastrophe rates. The reliability of the critical components must be known in order to estimate casualty rates. Alternatively, the critical components can be identified and the required reliability specified.

If there is agreement about the safety criterion and the hazards, the method of complete specification and fault-tree analysis becomes a basis for procedural standards for safe design and evaluation.

## SPECIFICATION OF HAZARDS

The hazards used in the fault-tree analysis, which form a basis for design, are specified elsewhere (1). They assume there is platooning, which is a basis of this design. The four hazards are then the possible precursors of a high-delta-V collision involving a platoon or vehicle under automatic control. These hazards are

1. 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 (to be defined).
2. A vehicle not under system control is an unmeasured and unknown distance in front of a platoon or single controlled vehicle.
3. A vehicle is released to manual control before the driver has given a positive indication of readiness.
4. 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, say, 2 sec.

Here, platoon spacing is defined as the safe spacing between platoons according to the criterion of Shladover (6). A preceding vehicle is halted violently (say, with a deceleration of 1.0 g), and the follower must brake to rest without collision.

PATH, Institute of Transportation Studies, University of California-Berkeley, Richmond Field Station, 1301 South 46th Street, Richmond, Calif. 94804.

Manual spacing is that spacing at which drivers feel comfortable and use in normal driving. In the system proposed, these quantities, which depend on the condition of the road, are set by the system controllers.

During the fault-tree analysis, it became apparent that hazard specifications erred. The last one (Item 4 in the list) omitted to say that a vehicle should not be released to manual control while the brakes are being applied. Otherwise the driver may not be able to regain control.

No formulation of the hazards can exclude all possibilities of high-delta-V collisions. There are parts of the road (transition lanes, or TLs) where vehicles are taken from manual to automatic control or released from automatic to manual control. Both manual and automatic vehicles are present here. Thus, automatically controlled vehicles cannot be protected from all errors by manual drivers. Sideswiping or cutting in on an automated vehicle are particular examples. Hitchcock has shown that in any design, a fence must protect vehicles on the automated lane from such accidents (1). There is an exception when accident debris is projected through a gate just in front of a platoon. Vehicles on the transition lane are, however, open to such accidents, just as they would have been without automation. The collisions that may result, if a vehicle in a platoon fails, take place at low delta-Vs (6). Such collisions do not have to be guarded against in the same way as do collisions between vehicles in different platoons. This is the basis on which platooned designs are accepted (6).

If a platoon is fully formed and there is a vehicle failure within it, the ensuing collisions are slight. The entire platoon may then come to rest if a vehicle cannot continue. Provided the wreckage does stop without hitting something else, no occupants will suffer large, injury-provoking decelerations. The fences referred to earlier also have this effect. However, if vehicles are joining or leaving the platoon, the collisions can be more serious. Just how serious has not been made clear. We do not know the frequency of the vehicle faults that cause such accidents. In the current design, some automatic inspection of vehicles entering the system is envisaged. Whether dangerous faults will then be detected is not known.

Within-platoon collisions are not examined here. If there is only a single automated lane, system design can do little to reduce the numbers of such collisions. The time that elapses while vehicles join or leave a platoon can be minimized. If there are multiple lanes, the more serious incidents that occur when a vehicle is joining or leaving a platoon can be reduced. The wish to do so must be overriding. The system is arranged so that a platoon (or a single vehicle) joins another platoon only at the rear, by cutting into position from another lane. Platoons divide in the same way. It will often happen that a vehicle in the center of a platoon wants to leave the automated lanes (ALs). The platoon then must reform. The back of the platoon must make two lane changes in quick succession—which is likely to be uncomfortable. The cure seems worse than the disease.

**SYSTEM CONCEPT**

The system considered here has a single AL on which vehicles run in platoons on a freeway that is also used by manually controlled vehicles. Such systems are discussed by Hitchcock


(1), who shows that, if the hazards are to be respected, the physical configuration is necessarily that which is presented in Figure 1.

The AL is separated from the unconcerned lanes (ULs) by a fence through which vehicles may pass only at brief on-gates and off-gates. The gates are grouped into one logical on-ramp (LONR) and one logical off-ramp (LOFR) per block (of 1 mi or so in length). The last gate in the block is an on-gate. To the right of the AL, there is a TL in some places. It stretches from some distance upstream of the LOFR and LONR to a short distance downstream of the LONR. It may be discontinuous, as it is in the case considered here.

The TL is instrumented with vehicle position detectors (VPDs). There are also VPDs on the AL in the neighborhood of the LONR and LOFR. Vehicles are taken under the control of the system on the TL and are under control as they are passed through the gates. Vehicles enter the AL only at the rear of a platoon. Along the length of both TL and AL runs a lateral guidance reference. This reference defines the proper course to vehicles' later controls. Close to the gates there are turning points marked on the TL. Turning points act as reference points in lane changing. Each vehicle bears a lateral and a longitudinal control system that keeps it on track and property spaced within a platoon. These systems contain sensors that can detect a vehicle ahead within a defined minimum range called sensor range. The speed at which platoon spacing equals sensor range is called sensor range speed.

All these conditions have been shown to be necessary for safety (1). In this design, the TL also contains identifiers and chicanes. At identifiers, vehicles that wish to enter the system identify themselves. At chicanes, their claim to have an operative control system is verified. Appropriate control signals are sent. It is checked that the vehicle accelerates, decelerates, or steers to the left or the right. The chicanes are, of course, much less severe than those on racing circuits—the occupants may not notice the test.

Varaiya and Shladover have described a possible control architecture of an automated freeway—it is shown, with the addition of a top layer (law), in Figure 2 (2). In this case, the link layer is concerned to organize the formation of platoons



**FIGURE 1** Layout of AL and TL.

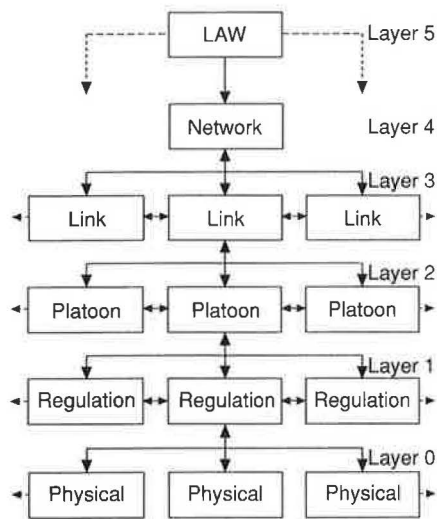


FIGURE 2 IVHS control architecture (2).

and selection of gates for entry and exit. It also organizes the movement of platoons on AL and TL so that there are appropriate gaps when lane change is desired. The link layer thus determines system capacity. The link layer is not part of the S-CS, and the maneuvers it recommends are not started until the platoon level has verified that they can be executed safely. The verification uses information from the VPDs. Control and interpretation of VPD signals are carried out at the regulatory level. The platoon and regulatory and physical layers compose the S-CS and are localized within each block.

When some fault occurs, the system in one block is put into a degraded mode. The fault may arise because the presence of a faulty vehicle in the system is detected. Alternatively, the roadside system detects an internal error. The law requires that vehicles with faulty (or no) controls do not enter the system. Entrance tests try to enforce this. However, faults may develop after entry. Deceit ("hacking") is also possible. In the simplest degraded modes, speeds on the AL are reduced to sensor-range speed. In one of these modes, vehicles may be required to exit the AL at the off-gates. In another, vehicles may continue on the AL the next block, where there may be no restrictions. If operation in these modes is not safe, further degradation is necessary. All vehicles will be brought to rest. In these conditions human intervention is necessary. The system controllers should be able to direct unusual operations. These include automated backing up on the AL so as to clear the way for removal of casualties or debris. Reversion to normal is also under human supervision.

## INITIAL DESIGN CONSIDERATIONS

No previous descriptions of an automated freeway system have covered such events as joining and leaving the automated lanes, nor do any cater for fault conditions. The purpose here is to design a system on which a method of safety analysis is to be tested. The hazards are therefore defined first. In ad-

dition, the designer has three concepts that shape much of the design. These are

- The system should verify that each event demanded has in fact occurred. If it has not, suitable action should be defined. This ensures the completeness of the specification.

- There will be a strong temptation for some people to attempt to beat the system. Such people might send signals that falsely allege that a vehicle has automatic controls. They might suppress data indicating a fault. This can bring economic benefit to—and gratify the ego of—the driver. If this is done, however, it is possible that things will go wrong. This can be a source of catastrophe.

- Misinterpretation or nonreceipt of transmitted messages is a potential source of hazard. Noise can cause this, so it must be kept to a minimum. A strict discipline of sequenced transmissions is desirable. Each should identify the vehicle referred to, that is, which one is transmitting or being addressed.

The second of these considerations suggests that the output of vehicle-borne intelligence is untrustworthy. The intelligence should be emphasized in the infrastructure. This has therefore been decided on for the present design. Thus, much of the transmission between units will involve the infrastructure, and there will probably be economic advantage in eliminating any vehicle-vehicle communication. The third consideration supports this decision.

These decisions are basic to this design. They reflect unproven assumptions. It is certainly possible to construct system designs based on vehicle-borne intelligence—the author has done so. It may be possible to avoid the deleterious effects on hacking and noise by other means, or perhaps hacking and noise are less significant than assumed. Nevertheless, in this case, infrastructure-mounted intelligence is emphasized.

It is possible to arrange that a message to or from a vehicle is possible only if the vehicle is at one particular spot. This is necessary at the identifiers, for there is no other way in which the system can identify the vehicle seeking entry. It may also be desirable at chicanes. In general, however, the message identifies the vehicle being referred to. To ensure receipt it should be heard over the whole length of the block. Transmission and reception are therefore accomplished by way of a line of devices operating in parallel.

## Functions of Architectural Levels

The regulatory layer must maintain each vehicle on track and in motion. By control of VPD signals, it provides performance data that enable an independent check on vehicle behavior. The platoon layer must initiate maneuvers, such as platoon formation or lane change, advised by the link layer. It does this only after verifying, in the light of data from the regulatory layer, that it is safe to do so. The platoon layer also analyzes the performance and position data provided by the regulatory level. It checks for a hazardous condition. If there is an incipient hazard, the platoon layer will revise maximum speeds. The platoon layer also passes data to the link level describing the position and speed of platoons and solo vehicles.

The link layer is outside the S-CS. It advises in maneuvers, including platoon and vehicle speed changes, that will

- Enable vehicles to leave at their selected exit,
- Form platoons as appropriate,
- Organize the pattern of traffic near gates in order to make gaps, and
- Enable changes from TL to AL and the reverse.

The link level thus has an optimizing function. Link maximizes the achieved capacity of the system so as to meet demand. However, it only advises on actions. The platoon levels mediate these for safety—that is, the local controllers at the platoon level check that each maneuver requested by the district controllers at the link level violates no hazards. Only then does the platoon level initiate the maneuver. This is achieved in part by sending a maximum speed to the vehicle-borne regulatory level. The regulatory level determines the speed of the vehicle it controls as follows:

1. The platoon level's maximum speed will not be exceeded, whatever the other rules say.
2. The speed of a vehicle in platoon is determined by the control algorithm, sensor readings, and other data.
3. The speed of a platoon leader or solo vehicle is the target speed, set by the link layer.

**Alternative Designs**

A very considerable superstructure has been built on initial considerations that are not without merit, but should not be overriding. The resulting design is not unsound, but the concentration on infrastructure intelligence and communication is extreme. Other designs are possible and would be arrived at if there were different initial considerations (for instance, that vehicle owners should pay directly for their benefits). It is therefore planned to make a parallel demonstration of the specification and fault-tree analysis technique, starting from the partial design of Hsu et al. (7). This is based on the idea that the greatest possible amount of intelligence should be vehicle-borne.

**Development of Design**

Returning to the initial design, the considerations rehearsed so far lead to the concept of the iterator as the roadside component of the regulatory layer. An iterator is a control and communication computer controlling a number of similar elements. An iterator communicates with each in element in rotation. Iterators on ALs and TLs address each solo vehicle or platoon member, thus stimulating a reply, in strict rotation, so as to avoid message overlay. The vehicle receives information about its maximum speed along with other data as are for within-platoon control (speed of platoon leader). The vehicle responds with an account of its speed, distance to the vehicle ahead (if it is detectable), and lateral displacement. The fact of reply ensures that the communication equipment is in order. This information is passed to roadside state vectors (RSVs), one per vehicle in a block. RSVs are data records held in asynchronous data stores. These stores are accessible

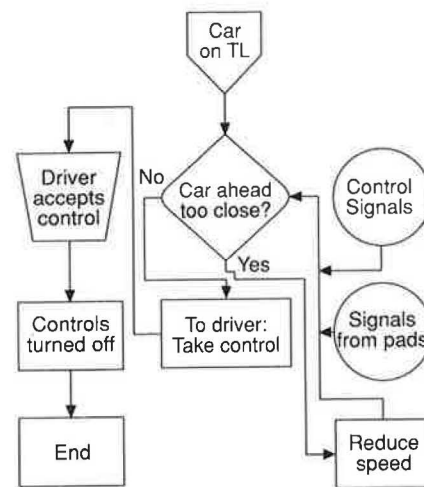
by the platoon level. Messages are sent to platoon level, if appropriate. Thus, the continuous independent check on vehicle control system behavior can be made.

The other part of the roadside regulatory system is the monitoring via the VPDs of all vehicles on the TL and on the instrumented part of the AL. Each VPD is monitored in sequence. Vehicles (including unconcerned vehicles) are tracked using the traces they leave on the line of VPDs. Vehicle position and measured speed are communicated to the RSV. Again, messages are sent to platoon level if there is an unusual feature, such as too little space ahead of a platoon. Thus, the rest of the information needed in order to monitor vehicle control systems is gathered. In addition, there is a special RSV for each gate. If it is set, passage through the gate is barred because a vehicle is present on the other side. These RSVs, too, are controlled as part of the logical process described earlier.

The design can now be built up, module by module. Initially the required normal behavior of vehicles, drivers, and system is defined. This behavior is largely determined by the fact that each vehicle must reach its destination. As each maneuver is proposed, the system must have a means of checking that it has been carried out. The design must contain an alternative safe procedure if the maneuver is not executed.

An example may make this clear. Figures 3 and 4 show part of the flow diagram starting with a solo vehicle that must resume manual control and complete its journey. To avoid hazard, the driver must indicate readiness to resume control. A message is sent. If the driver replies, control is passed, and the iterator ceases to communicate with the vehicle. All is well. But, according to the design method, the question is asked, "What if the driver does not reply?"

The TL is not continuous—the vehicle cannot remain in motion on it indefinitely. What should happen is not certain. A choice must be made by the designer. The present system tries to reinsert the vehicle into the AL. The decision point is the last on-gate. Here, there may be a gap for entry at the rear of the platoon the vehicle has just left. The VPD signal is checked to ensure safety before instructing the vehicle to



**FIGURE 3** Normal exit.



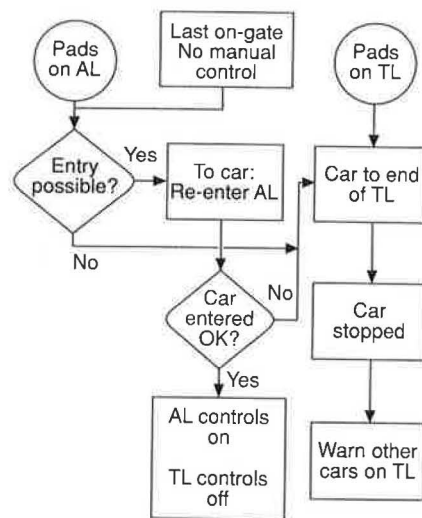


FIGURE 4 Driver does not take over.

enter. The VPDs are then checked again. (The case for which the vehicle enters is not followed further.) Entry may be unsafe because of vehicles on the AL. The gate may be closed because of operation in a degraded mode. The check may reveal that the vehicle did not enter when invited. Further checks are therefore needed to ensure that the vehicle comes to rest (or resumes manual control) before it reaches the end of the TL.

By proceeding in this manner throughout operation, the design has been completely specified. In this case, there was a choice to be made (what should be done with a vehicle whose driver fails to resume control?). A different system design would have resulted if some other choice had been made. This happens at a few other places in this system design.

## DESIGN CHOICES

Each time a choice arises, an alternative choice will lead to an alternative design. The number of such choices is small, because in most cases it is quite clear what should be done if a maneuver is wrongly executed. Within the assumptions here, the only major choices are the preceding one and the following:

- Where should platoons be formed or dissipated? Platooning on the TL assists capacity, since more vehicles can pass a gate at once. However, on the TL the platoons are exposed to casualty-causing accidents involving unconcerned vehicles. These accidents would have happened anyway. The consequences are more serious because a platoon is involved. Large exit platoons may have difficulty in dissipating before the end of the TL is reached. Changing lanes in a platoon may place difficult demands on the vehicle-borne control systems and reduce reliability.

- In the design chosen, complexity is introduced by counting vehicles on the AL as they move from block to block. This may not be necessary: an intruder should be detected by the gate VPDs, and a lost vehicle should make its presence

known long before it is missed at the end of a block. The redundancy introduced by counting may not, therefore, add to overall system safety.

- There are many possible ways of identifying the vehicle referred to in a message. In the design selected, the designer found it convenient to regard vehicles as passing through a series of modes as they passed tests at the chicane, joined platoons, entered the AL, and so on. This gave rise to a complex shifting ID system to which there are many alternatives.

- Further choices arise if some functions of the VPDs, such as checking on safe spacing from the platoon ahead, are transferred to the vehicle. This requires that a sensor can be constructed that has the ability to detect a vehicle in the same lane, distinguishing it from vehicles in other lanes.

## FAULT-TREE ANALYSIS

In a fault-tree analysis, each hazard (a predecessor of catastrophe) is considered in turn. One asks, "How could this arise?" The answer will take on a form such as "If *A* happens, or *B* happens, or *C* happens, . . ." One then asks "How could *A* arise?" "If *AA* happens or *AB* happens, . . ." The process of identifying precursors continues. Mathematically, "*A* happens," "*B* happens," . . . are logical propositions, and "and," "or," and "not" are Boolean operators. Sooner or later one arrives at the point at which the proposition is one of the following:

- This can happen as a result of a single fault in a vehicle or other system component. In this case design error has been found.
- This implies that two simultaneous faults have occurred.
- There has been a computer or a communication error (the computers and communication equipment are assumed to be so redundant that this implies two simultaneous faults).
- The proposition is not possible (e.g., it involves reversal of gravity).
- The proposition implies that there has been an inadequacy in maintenance.

In each of the last four cases there is no breach of the safety criterion on this branch of the tree.

A fault tree clearly involves subjective elements. It is always possible that the investigator will fail to realize one of the ways in which a situation could arise. This becomes more likely when, as in this case, the investigator is the designer.

Nevertheless, in both specification and analysis, the process has been carried out with formal rigor. Besides its inclusion in some 20 pages of flowchart drawings, each module in the design (there are about 120) has been specified in a standard form. This form shares many features with the forms used for module specifications in formal-method computer languages such as Z or OBJ-3. The specification language used here, however, is not based on formal axioms. The complete formal specification is stated and discussed elsewhere (4). In the fault-tree analysis, similar rigor has been used—there are some 50 elements in the tree, and the arguments in each have been recorded precisely (5). Both reports are long and complicated, and no attempt is made to summarize them here.

## RESULTS OF FAULT-TREE ANALYSIS

Four design faults were found:

1. On the uninstrumented part of the AL between gates, a following platoon may gain slightly on its predecessor. No mechanism is provided to correct this. In any one block the effect is trivial, but it could accumulate and cause a hazard.

2. Care is taken to check that a vehicle joining the AL does so only at the rear of a platoon or into a large gap. However, no check is made on the vehicle's speed when it enters the AL. If this is grossly mismatched with the platoon speeds, a hazard can arise.

3. If a vehicle develops a fault, it is detected and the driver is invited to resume manual control as soon as this can be made possible. No special precautions, however, are taken before it does so in order to keep other vehicles away from the danger a faulty vehicle presents. This can lead to hazards.

4. When a vehicle is released from a platoon or admitted to the TL on its way out of the system, its release is controlled so that its separation from the vehicle in front is safe. Controls also ensure that it is not moving much faster than its predecessor. Thereafter its distance from preceding vehicles is controlled to a safe spacing (though an unconcerned vehicle can always cut in). However, at the moment of release no check is made on its speed relative to its predecessor. This too can lead to hazards.

## DISCUSSION OF RESULTS

The design described in this paper suffers from errors. They can readily be remedied should there be a serious intent to develop it. This is not very important, since there is no intent to develop this particular system. More important, at the system level considered here, it does appear to be possible to produce a design that avoids the hazards. This conclusion stands on the basis that controls and sensors can be instantiated that conform to what is specified.

The method of analysis chosen here is detailed, complete specification followed by fault-tree analysis. The example suggests that this approach is sufficient to ensure and verify conformity to safety criteria. The subsequent stage of a quantified hazard and risk analysis is, plausibly, also sufficient to ensure conformity to safety criteria of the whole system. There is more work to be done before these claims can be pronounced valid. That stage may be reached, however. These techniques could then become the basis for standards for design and evaluation procedures against stated safety criteria.

If regarded as an exemplar for standards, however, there are some serious flaws in the present demonstration. First, the designer and the analyst are the same individual. Proper management of system safety, as described in many guides in nonhighway fields (8), requires parallel and independent development of design and safety analysis.

Next, the design method should lead to a complete specification. No check has been made of this. In the work of Hsu

et al., formal methods are used to demonstrate completeness for a part of the system (7). Whether these methods can be extended to the whole system, or even to the S-CS, is still to be investigated. But some independent validation of the completeness of the design concept is necessary before the present work can be regarded as exemplary.

It would be most satisfactory if verification and validation could be done by formal methods, so that completeness would be proven mathematically. If the fault tree also could be proved to be complete, it would be even better. However, this does not yet seem practical.

The choice of hazards constrains the number of possible system designs of an automated freeway. The constraint seems to be severe and the number, to be small. This parallels the earlier result (1), which was restricted to the physical layout.

## ACKNOWLEDGMENTS

This work was performed as part of the PATH Program of the University of California, in cooperation with 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 also like to thank S. E. Shladover and P. Varaiya for helpful discussions.

## REFERENCES

1. A. Hitchcock. Intelligent Vehicle-Highway System Safety: Problems of Specification and Hazard Analysis. In *Transportation Research Record 1318*, TRB, National Research Council, Washington, D.C., 1991.
2. P. Varaiya and S. E. Shladover. *A Sketch of an IVHS System Architecture*. Research Report UCB-ITS-PRR-91-3. PATH, University of California, Berkeley, 1991.
3. N. H. Roberts. *Fault Tree Handbook*. NUREG-0492. U.S. Nuclear Regulatory Commission, Springfield, Va., 1981.
4. A. Hitchcock. *A Specification of an Automated Freeway*. Research Report. PATH, University of California, Berkeley (in preparation).
5. A. Hitchcock. *Fault Tree Analysis of an Automated Freeway*. Research Report. PATH, University of California, Berkeley (in preparation).
6. S. E. Shladover. *Operation of Automated Guideway Transit Vehicles in Dynamically Reconfigured Platoons*. UMTA-MA-06-0085-79-1, 2 & 3. U.S. Department of Transportation, University of California, 1979.
7. A. Hsu, F. Eskafi, S. Sachs, and P. Varaiya. *The Design of Platoon Maneuver Protocols for IVHS*. Research Report UCB-ITS-PRR-91-6. PATH, University of California, Berkeley, 1991.
8. *Procedures for Safety-Critical Software*. DO 178A. Radiotechnical Commission for Aeronautics. Washington, D.C., 1986.

*The contents of this report reflect the views of the author, who is responsible for the facts and the accuracy of the data presented herein. The contents do not necessarily reflect the official views or policies of the state of California. This report does not constitute a standard, specification, or regulation.*

*Publication of this paper sponsored by Task Force on Advanced Vehicle and Highway Technologies.*