Intelligent Vehicle Highway System Safety: Specification and Hazard Analysis of a System with Vehicle-Borne Intelligence

A. Hitchcock

Hsu has described the normal operation of a system of automated freeways that minimizes the degree to which the infrastructure is involved in maneuvers. No account is given of procedures on entry and exit or of possible faults. The Partners for Advanced Transit and Highways (PATH) safety program demanded a second example of the process of full specification and fault tree analysis to determine if this process was generally applicable. Accordingly, Hsu's specification has been completed, retaining minimal infrastructure-based involvement in the maneuvers. In this design the safety criterion was adopted that no hazard shall arise unless there are two independent faults. The hazards were those used earlier. Because a hazard is the precursor of a catastrophe, no multivehicle high-relative-speed collisions should occur without three independent faults. A fault tree analysis was carried out in which there were never more than four branches in any line of the tree. It is concluded that it is possible to produce a design that meets these safety criteria and that behaves during normal operation in the same way as the system defined by Hsu. Further, the method of complete specification and fault tree analysis suffices to produce a system of demonstrated safety with a practical allocation of resources and time. A comparison is made of the safety of this vehicle intelligence system and the earlier infrastructure intelligence one. On the basis of experience rather than analysis, it is concluded that it is possible to design a safe automated freeway.

The process of complete specification and hazard analysis has been defined previously in relation to automated freeways by Hitchcock (1). The procedure already has been demonstrated in one example (2–4). That was a case in which the system intelligence was mainly located in the infrastructure. The fault tree analysis was successful in finding errors in the design. There is interest also in evaluating cases in which the intelligence is mainly vehicle based. It is possible that the basic method would be less successful or would need modification in such a case. Therefore a second example was undertaken of complete specification and fault tree analysis, in which the system considered is based on vehicle-mounted intelligence (2,5,6).

First it was necessary to define the hazards. A hazard is defined as a situation in which a further fault can lead to a catastrophe. A catastrophe is defined, as is usual in system safety, as a multicasualty event; in practice this means a high-speed collision involving platoons or vehicles. The hazards restated below are the same as those used in the earlier work (1,2). As in the previous work, the statement of hazards applies to platooned systems. However, in principle they apply to nonplatooned systems also, although small changes to the formal statements are required. Also, for the reasons given by Hitchcock (1), there are barriers, called dividers, between the lanes that prevent vehicles from straying from one lane to another. The hazards do not expressly refer to this, although they can again be construed to apply to the problems that arise if the dividers are absent. For example, if a vehicle strays into another lane, it becomes the "one ahead" in Hazard 1 below. The hazards are as follows:

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 (see below).
2. A vehicle not under system control is an unmeasured and unknown distance in front of a platoon or free agent (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.

A platoon is a succession of any number of vehicles that are closely spaced in the same lane and under coordinated control. In the hazards above, platoon spacing is defined as the safe spacing between platoons according to the criterion of Shladover (7). Manual spacing is the spacing at which drivers feel comfortable and that they use in normal (i.e., not automated) driving. In the system proposed, these quantities, which depend on the condition of the road, are set by the system controllers.

To say that a system meets, or does not meet, a safety criterion is a statement of fact. This paper discusses means by which the validity of such a statement can be demonstrated. To say that a system is safe, on the other hand, is a statement that will mean different things to different people.

California PATH Program, Richmond Field Station, Institute of Transportation Studies, University of California, Berkeley, Berkeley, Calif.

SYSTEM SPECIFICATION

When the work reported here started, a system with the requirements for this work had been partially specified by Hsu et al. (8). This system was used as the basis for a complete specification. A full description is given by Hitchcock (5). Although the full specification is much too lengthy to be reproduced here, in this paper the specification and the design method are described and examples are given.

It is a principle of the design method adopted here that procedures such as checking for faults are integral to the design. To design a system for normal operation and then "bolt on" safety features has been shown by experience in other fields to lead to difficulties. The designer therefore considered what checks and messages were necessary to confirm compliance with expectations at each stage, both with and without a detected fault. In the description that follows, system behavior in normal and fault conditions is considered in parallel. The architecture encompasses both functions.

Each length of lane in the system (called a block, defined below) may operate in one of several modes. In some of these, vehicles are at rest while emergency vehicles under human direction deal with accidents. One is a restart mode. One is the normal operating mode, and the others are degraded ones, in which speeds, for example, may be reduced to facilitate lane changing if a lane is closed ahead (SA mode = slow ahead). Another example is no-entry mode (NE), used when a vehicle in the lane has a communication fault and therefore cannot cooperate if another vehicle is about to enter its lane.

As will be seen, the control system requires that messages be exchanged between vehicles. It is convenient to give such messages names that indicate their function. Such names contain an underline character (_).

Architecture

The system architecture, as described by Varaiya and Shladover (9), is a hierarchy of six levels (see Figure 1). The three lowest form the safety-critical subsystem. As safety demands, the construction is modular, and messages or data pass between modules only at fully specified interfaces. The physical level defines the physical processes (such as tire-road friction) as a result of which control actions are translated into vehicle movement. The regulatory level controls the movement of vehicles: lateral and longitudinal control (i.e., the subsystems in each vehicle that maintain it in its lane and keep it properly spaced from its predecessor) are regulatory-level functions. The selection, from among all messages received, of messages relevant to this vehicle, is another regulatory-level function. The platoon level organizes the maneuvers whose combination enables a vehicle to proceed from its entry to its desired exit. The link level advises on the desired route to achieve the objective. Because the route of each vehicle is affected by the routes of others, link-level functions are roadside based. But link advises only; before a maneuver is initiated, platoon-level procedures ensure that it can be done without hazard. The platoon, regulatory, and physical layers thus comprise the safety-critical subsystem. Functions of the higher levels will not be discussed here.

Roadway and Vehicles

The automated freeway has several lanes, divided into blocks, perhaps 10 km (6 mi) long. They are separated by dividers that contain gaps, called gates, through which lane changes are made. On the sending side of each gate (the side from which vehicles leave) there is a "turning point"—an electromagnet or other signalling device that, when activated, gives a signal to the lateral control system of a vehicle to begin a turn, if that has been commanded. If the turning point is inactive, the vehicle will not commence a lane change. The gate advises vehicles changing lanes of the position in the block of the gate, appropriate speed limits, and the system mode of this lane. The presence of a vehicle on the receiving side will inhibit turning-point activation. Rerouting instructions may also be passed here. In fault conditions, some lanes may operate in degraded modes in which speed is reduced and lane changing may be restricted. At the commencement of each block is a marker that is used by vehicles crossing it to zero a vehicle-mounted odometer. The discrete lateral control references enable distance along the lane to be measured to ±1 m or better. Every vehicle thus contains a register containing its current absolute position to within 1 m.

Each vehicle is equipped with a longitudinal control system, a lateral control system, a self-monitor, a communication system, and an overall controller:

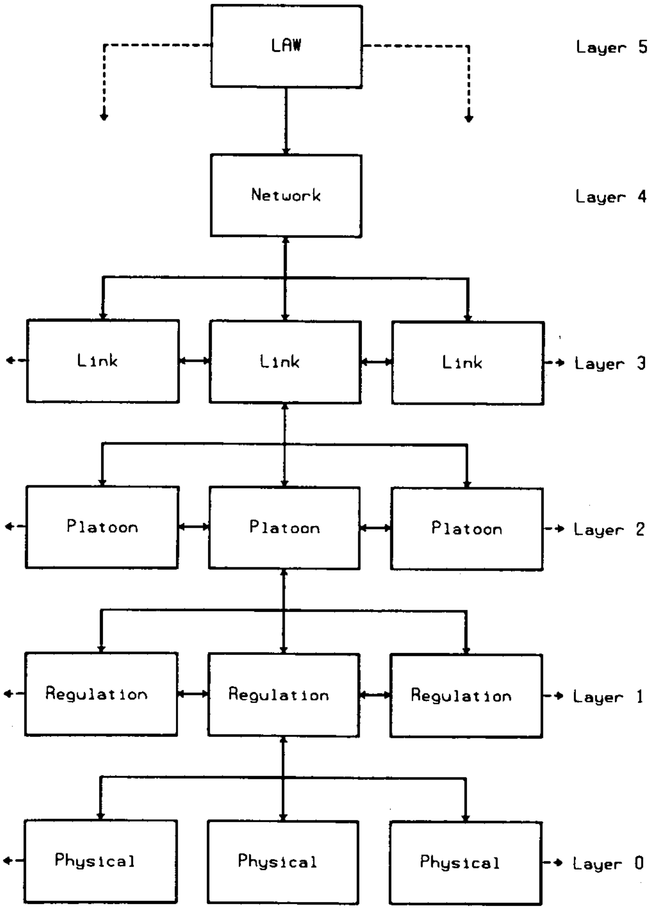

FIGURE 1 IVHS control architecture [after Varaiya and Shladover (9)].

1. The longitudinal control system has a forward sensor that can determine the distance and relative velocity of the vehicle ahead in the same lane within a maximum range. This range exceeds the maximum interplatoon spacing. If the vehicle is in platoon, the longitudinal control system will maintain the vehicle at a fixed distance (about 1 m) from the vehicle ahead. If the vehicle is a platoon leader or a free agent, the vehicle will proceed at a constant speed [between 90 and 110 km/hr (55 and 70 mph) in normal conditions], except that it will not approach closer than platoon spacing from the vehicle ahead.

The odometer provides a means for a vehicle to identify its position to others. In forming platoons, for example, the initial action is for a platoon leader to send the message request_merge (see below) to the vehicle ahead in the same lane. But the follower has no identifying name for the one ahead, and the message may be received by several vehicles. By including the odometer reading, lane number, and block number, the intended recipient can be identified.

2. The lateral control system will keep a vehicle on track in an automated lane. If a "turn" instruction has been given, it will execute a lane change at the turning point.

3. The self-monitor will detect faults in the other subsystems, including the mechanical parts of the vehicle and conditions such as shortage of fuel. If a fault is detected, the control system's objective changes. Instead of trying to follow a route designed to reach a stated destination, the vehicle will leave the automated lanes as soon as possible.

4. The communication system will transmit and receive messages from other vehicles ahead, behind, and to either side, and also messages broadcast to or from the system. These have to be eight separate independent functions: thus, for example, failure of rearward reception does not imply loss of rear transmission or of system reception. If any two of the eight functions are combined in one device, it is relatively simple to show that a hazard will arise following the single fault of its failure. This violates the safety criterion.

5. The overall controller, called supervisor, determines when it is time to join or quit a platoon, change lanes, and so on. One of its features is a busy flag, which is set while a vehicle engages in a maneuver. If the busy flag is set, the supervisor will usually deny a request from another vehicle to engage in another maneuver.

Maneuvers

Hsu et al. (*8*) describe three basic maneuvers that enable the system to perform its functions when there are no faults. In the merge maneuver, the leader of a following platoon contacts the one ahead, and, if their sizes are appropriate and the busy flag is reset, the two platoons will merge to one. In split, the reverse happens. A vehicle requests that the platoon split into two. By means of at most two such split maneuvers a vehicle in a platoon can become a free agent. Only a free agent may initiate the change-lane maneuver.

Figure 2 is based on the report of Hsu et al. (*8*) and shows the sequence of events in the merge maneuver in the present work. The appearance of an underline character (_) indicated a message name. As the figure shows, the following platoon leader, B, initiates the merge by sending a message, request_merge, to its predecessor, A. (In fact, the message is received by the rearmost member, C, of A's platoon, and passed forward by the within-platoon communications.) If A is busy, or if the resulting platoon would exceed the maximum size, A rejects the invitation by sending nack_request_merge. Otherwise A sends (via C) ack_request_merge. As it passes the message on, C sets the busy flag. B then accelerates and joins on to C. When it does so, it sends confirm_merge, again through C. Platoon leader A then changes its regularly transmitted control data message to take account of the increased platoon size. This is received successively by C and B, and as each receives the new control data, it resets its busy flag, so that it is ready to start a new maneuver.

Hsu et al. (*8*) do not distinguish the role of C, and it sets no busy flag. The change was made to deal with fault con-

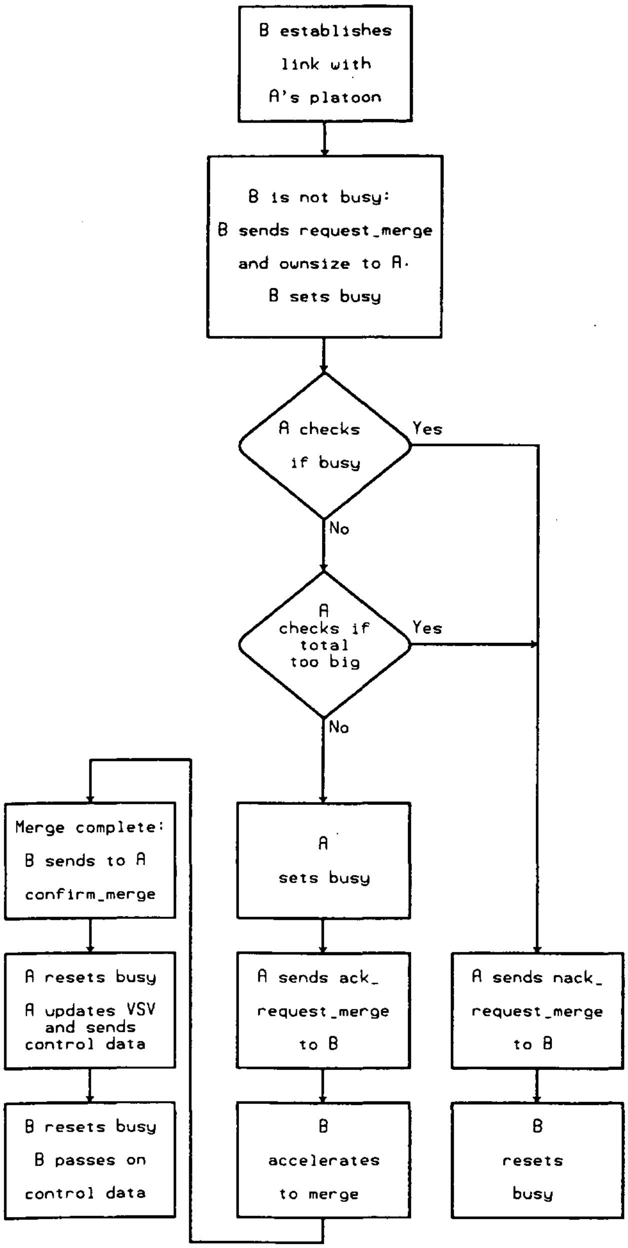
B establishes link with A's platoon
B is not busy: B sends request_merge and ownsize to A. B sets busy
A checks if busy
Yes
No
A checks if total too big
Yes
No
A sets busy
A sends ack_ request_merge to B
B accelerates to merge
Merge complete: B sends to A confirm_merge
A resets busy A updates VSV and sends control data
B resets busy B passes on control data
A sends nack_ request_merge to B
B resets busy

FIGURE 2 Merge [after Hsu et al. (*8*)].

ditions: C's busy flag is significant only in fault conditions. Similar changes have been made to the other maneuver protocols described by Hsu et al. (8). These will not be detailed here. They identify the last vehicle as part of a platoon that is engaged in a maneuver. In the case of the change-lane maneuver, arrangements are made to break off the maneuver before completion, if a forced-split maneuver (see below) is called.

Faulty Vehicles

A treatment of fault conditions was devised for this work. When the internal monitor detects a fault, or the vehicle is advised of one by the system, a flag is set in supervisor. Different flags are set for different faults. Supervisor will thereafter arrange that the vehicle (if it will still move) is driven out of the system as quickly as possible. This is necessary, as may be seen from the following argument. If a pair of unrelated faults occurs on one vehicle, a hazard may arise. Also, if a vehicle with a fault starts to interact with another, and no special action is taken, a hazard may ensue. In this system, therefore, if two faulty vehicles start to interact (because neither has yet exited), all vehicles in the relevant lanes are automatically halted. This is clearly undesirable, although not unsafe, and to avoid it vehicles with one fault should exit as soon as possible.

To extract a faulty vehicle from the system, first it must be extracted from its platoon. The extraction must be done immediately because a faulty vehicle in platoon can destroy the control arrangements for the whole platoon. If the platoon is engaged in another maneuver, the faulty vehicle will cause the maneuver to proceed to an unsatisfactory conclusion or prevent it from concluding at all. The forced-split maneuver is therefore called by a faulty vehicle in a platoon. In relevant cases it will also be called by the vehicle it has ceased to communicate with (see "probes" below). It is necessary that this vehicle call the maneuver: when communication has failed, each half must act independently. Forced split has absolute priority. If there is another maneuver in progress it is either broken off or interrupted. The busy flag in the last vehicle of a platoon ensures that a maneuver in progress is not forgotten and remains accessible. Forced split also differs from the usual split maneuver in its termination. The message confirm_for_split is sent, but no reliance is placed on its receipt. After a maximum of two forced splits a faulty vehicle becomes a free agent. Thereafter it will behave as though it were busy if a message inviting participation in a maneuver is received. Depending on the nature of the fault, the infrastructure-based part of the system will take other precautions by putting selected lanes in some blocks into degraded modes; speeds may be reduced, the lane may be cleared, or entry may be forbidden. Full details of this part of the design are not given here; the essence is that lane changing by the faulty vehicle is made both hazard free and easier. Once it is a free agent, a faulty vehicle must change lanes successively until it exits.

To do this a fifth maneuver protocol, emergency change, is required. The standard change-lane protocol may place the changing vehicle close to other vehicles or involve it in merge maneuvers with other platoons, which may be unsatisfactory and is likely to delay further changes. In the emergency-change protocol the faulty vehicle is shepherded between other platoons, if there are any in the neighborhood. Thus, even if the faulty vehicle's speed is reduced, a faster vehicle in another lane cannot interfere with the lane change.

The working of the emergency-change maneuver is described in Figure 3. As will be seen, the faulty vehicle engages the cooperation of four platoon leaders in other lanes who prevent any other vehicle from entering the change-lane area that the faulty vehicle will use. In this case, there is no need to prepare for the possibility that another emergency change or forced split will arise while the emergency-change maneuver is in progress. If another emergency-change or a forced-split is called, there is an immediate call to the system, followed by shutting down of motion in the affected lanes. Thereafter, operations will be under the control of a human on the spot (presumably a member of the highway patrol).

Probes

For most functions, on-vehicle testing will confirm or refute the presence of a fault. From time to time a situation will arise in which it is clear that there is a fault, but it is not clear in which vehicle it lies. The roadside parts of the system keep records of both faulty vehicles in their areas and "suspect" ones. (A vehicle is suspect if it is one of several that an event has shown may be faulty.) If a suspect vehicle is involved in similar incidents twice, it is declared faulty, and the other suspect vehicles are acquitted. However, there are two circumstances in which "fault" action is required, although the self-monitor cannot detect a fault. To cover these, probes are employed.

The platoon leader's probe tests the forward sensor. A platoon leader may "think" it sees a vehicle ahead either because there is indeed one present or because its forward sensor is faulty. Periodically therefore, if it is not otherwise busy, a platoon leader will send a message, indicating that it can see a vehicle ahead of it at a particular odometer reading or that it sees nothing. If it receives an unexpected reply, or no reply when one is expected, it will try again. If again there is a logical clash, the platoon leader will usually declare itself faulty. However, it will not do so if a check with the system reveals that there is a vehicle nearby that may have lost one of its communication subsystems.

The in-platoon probe deals with the situation in which a vehicle in a platoon develops a fault in communication with the vehicle behind or ahead of it. The faulty vehicle has its fault diagnosed by its self-monitor. However, the other vehicle in the noncommunicating pair must also start a forced split and must have a means of knowing that it must do so. The in-platoon probe works by means of acknowledging the control data message passed down through the platoon. Uniquely this ack_control message is not passed on to the platoon leader. The logic is shown in Figure 4.

Entry and Exit

Entry and exit are change-lane maneuvers or (for the exit of a faulty vehicle) an emergency-change maneuver. It is envisaged that there are separate on- and off-ramps. The on-ramp

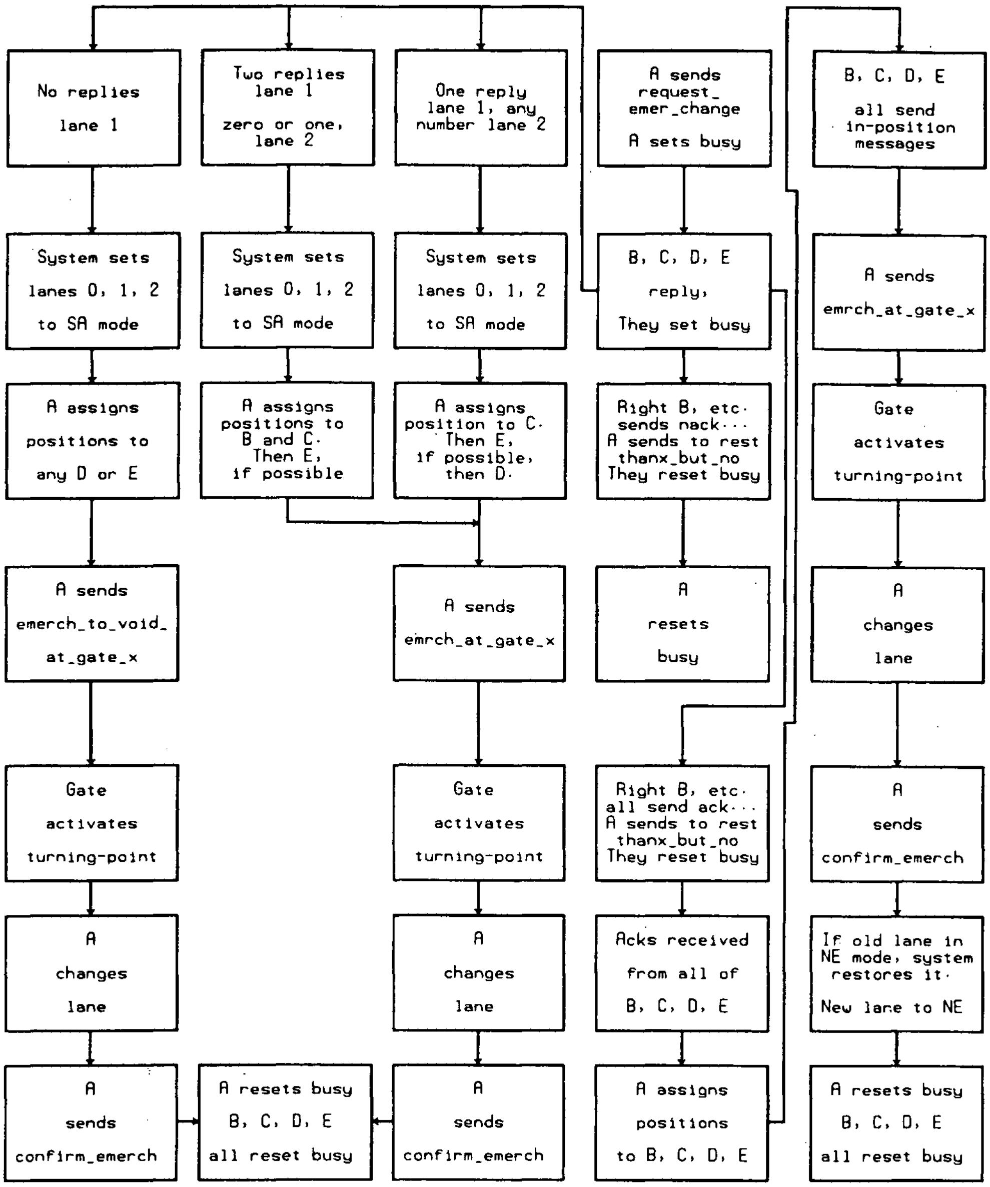
No replies lane 1
System sets lanes 0, 1, 2 to SA mode
A assigns positions to any D or E
A sends emerch_to_void_at_gate_x
Gate activates turning-point
A changes lane
A sends confirm_emerch
Two replies lane 1 zero or one, lane 2
System sets lanes 0, 1, 2 to SA mode
A assigns positions to B and C. Then E, if possible
A resets busy B, C, D, E all reset busy
One reply lane 1, any number lane 2
System sets lanes 0, 1, 2 to SA mode
A assigns position to C. Then E, if possible, then D.
A sends emrch_at_gate_x
Gate activates turning-point
A changes lane
A sends confirm_emerch
A sends request_emer_change A sets busy
B, C, D, E reply, They set busy
Right B, etc. sends nack... A sends to rest thanx_but_no They reset busy
A resets busy
Right B, etc. all send ack... A sends to rest thanx_but_no They reset busy
Acks received from all of B, C, D, E
A assigns positions to B, C, D, E
B, C, D, E all send in-position messages
A sends emrch_at_gate_x
Gate activates turning-point
A changes lane
A sends confirm_emerch
If old lane in NE mode, system restores it. New lane to NE
A resets busy B, C, D, E all reset busy

FIGURE 3 Emergency change lane. Change is to Lane 1: Platoons B and C are in Lane 1, B leading; Lane 2 is adjacent to Lane 1: Platoons D and E are in Lane 2, D leading. SA = slow ahead, NE = no entry.

is equipped with a separate exit, so that a vehicle that is refused entry can return to the manual lanes. The off-ramp is equipped with a "dormitory"—an area where a vehicle whose driver does not resume manual control on exit can be parked. On entry, a vehicle declares itself to be carrying control equipment that its self-monitor declares to be in working order. The driver must position the vehicle, relative to the traffic on the freeway, so that the change-lane maneuver is possible within the limited distance to the gate. On exit, a series of messages both before and after the change are sent to the driver and the vehicle reminding the driver of the need to resume manual control. A positive reply is needed before automatic control is released. This means, if a reply is delayed, that the vehicle must be brought to rest in the dormitory before it leaves the physical limits of the system.

DESIGN CHOICES

In the earlier work of Hitchcock (*2*) it was observed that, once the basic concepts were fixed, there seemed to be very few choices about the way the system design was put together. In the current study, there did not seem to be any design choices of importance at all, except the selection of four rather than one or two "shepherding" platoons in the emergency-change maneuver. However, it is less possible to be certain about this. The work of Hsu et al. (*8*) was taken as a basis, and there may be more choices within the area covered by that work. It is also not certain that there are no alternatives to the rule that a faulty vehicle should leave as soon as possible. Even if it is certain, it may be that there are other maneuvers besides forced-split and emergency-change, as de-

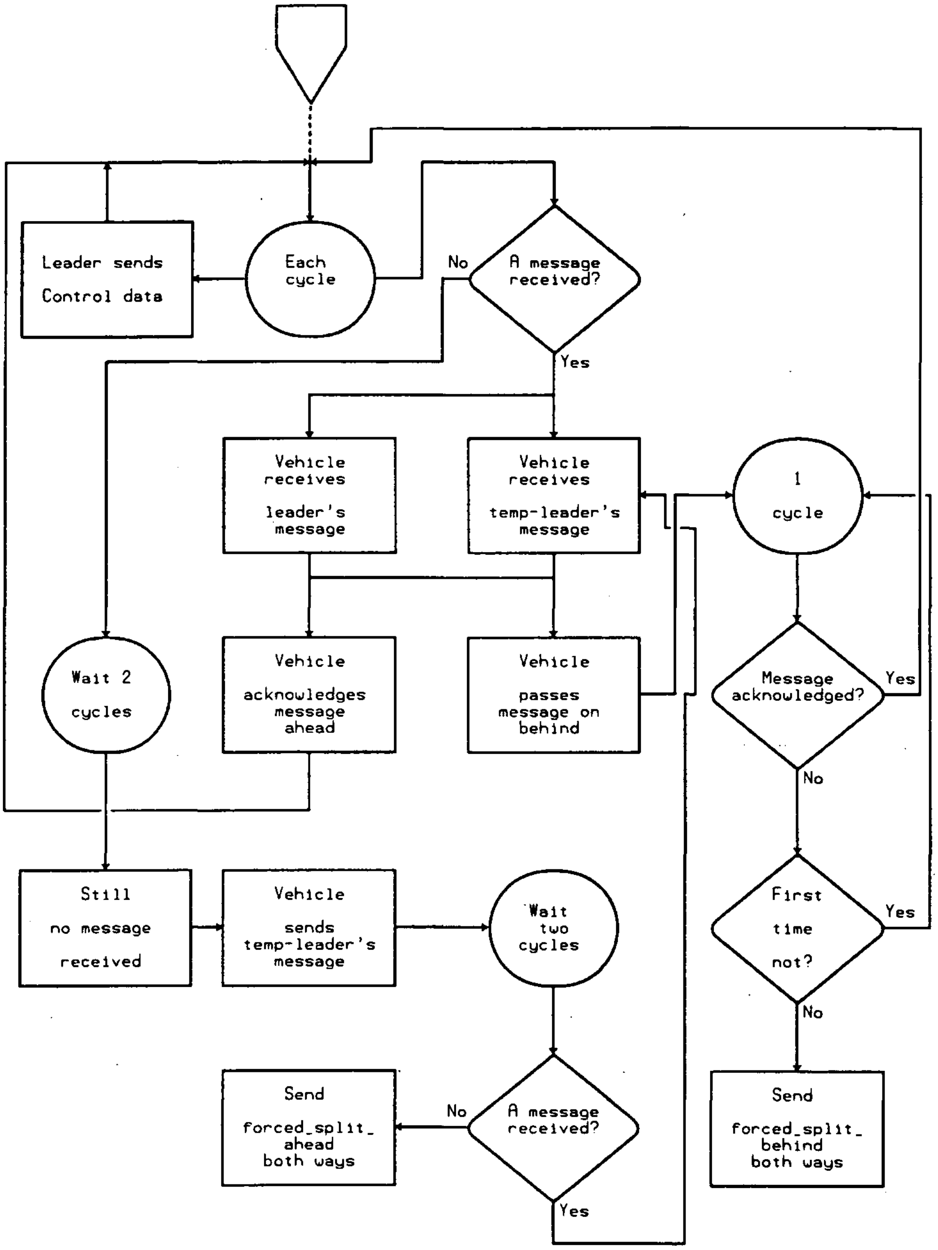
Leader sends
Control data
Each cycle
No
A message received?
Yes
Vehicle receives leader's message
Vehicle receives temp-leader's message
1 cycle
Wait 2 cycles
Vehicle acknowledges message ahead
Vehicle passes message on behind
Message acknowledged?
Yes
No
First time not?
Yes
No
Still no message received
Vehicle sends temp-leader's message
Wait two cycles
Send forced_split_ ahead both ways
No
A message received?
Yes
Send forced_split_ behind both ways

FIGURE 4 In-platoon probe. Probe operates in a cycle; there is no starting point. The figure is easier to understand by starting at the point indicated.

scribed here, that could be used to achieve this. The author did not conceive of alternatives, but this does not mean that there are none. It does, however, once again appear that the safety criteria suffice to restrict the acceptable designs to a very small number.

FAULT TREE ANALYSIS

In a fault tree analysis, each hazard 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 where the proposition is one of the following:

1. "This" (A, B, C, etc.) can happen as a result of a single fault in a vehicle or other system component. In this case a design error has been found.
2. "This" implies that two simultaneous faults have occurred.
3. "This" implies that there has been a computer error (the computers are assumed to be so redundant that this implies two simultaneous faults).

4. "This" is a proposition and is not possible (e.g., involves reversal of gravity).
5. "This" is a proposition that implies that there has been systematic failure to maintain the infrastructure.

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 the investigator is the designer.

Nevertheless, in both specification and analysis, the process has been carried out with formal rigor. Each module in the design (there are about 230) has been specified in a standard form and is also included in some 15 pages of flowchart drawings. This form shares many features with the forms used for module specifications in formal-method computer languages. The specification language used here, however, is not based on formal axioms. The complete formal specification has been stated and discussed previously (*5*). In the fault tree analysis, similar rigor has been employed: there are some 60 elements in the tree, and the arguments in each have been recorded precisely (*6*). Both reports are long and complicated, and no attempt is made to summarize them here.

DESIGN FAULTS AND RESULTS OF THE FAULT TREE ANALYSIS

In no case did a branch of the fault tree contain more than three elements, which means that the fault tree analysis is practical. The alternative of arguing forward from possible combinations of faults would have involved millions of times more cases, if not billions or trillions.

The analysis detected faults of two kinds. There were some potential violations of the safety criteria, strictly construed, which were known about before the fault tree analysis was started. That the analysis detected all these is reassuring but should not convince anyone that it will detect all faults. Some faults that arose from genuine mistakes in the design process, and therefore were not foreseen, were also detected. This should increase confidence that the method will detect all errors but does not prove it.

There were six foreseen departures from the satisfaction of the safety criterion. The first four of these (1–4 below) do not seem to be readily remediable. They are clearly rare. A quantitative analysis might show that they are so rare that they can be ignored. However, data that would enable such a quantitative analysis were not collected for this study. They are listed below because they cast light on the inherent safety problems of this system, which may well also be present in most or all similar systems.

These foreseen hazards are as follows:

1. It is not usually a catastrophe if some malfunction of a vehicle in a platoon causes vehicles in the platoon to crush together. The collisions occur at low relative speeds, and the dividers ensure that the vehicles come to rest without striking anything else. In a low-speed collision, also, the lateral control systems should keep the vehicles on lane. If such an accident occurs at a gate, however, some wrecked vehicles, perhaps with occupants, could trespass onto another lane and be struck by another platoon. There have to be gates in the dividers because all vehicles must change lanes. This class of accident seems unavoidable; in fact, it will be rare. It requires (a) that a low-speed collision occur; (b) that it occur in a small space within a gate; and (c) that it damage lateral control systems sufficiently to cause them to malfunction. If the lateral control systems have some robustness, malfunction following a low-speed collision should be rare. But until the detailed design of the system is known one cannot say how rare.

A parallel situation arises if debris from an accident between manually controlled vehicles on parallel lanes, or a dropped load, protrudes onto the automated lanes. Alternatively, the debris from an accident on manual lanes can be so massive that it breaches the divider. In these cases there would have been an accident in the left-hand lanes even if there had been no automation; the accident is in no sense caused by faulty automation. However, the greater density of traffic in the automated lanes can mean that the number of casualties is greater than it would have been.

2. During the merge-and-split maneuvers, Hazard 1 (see above) is necessarily violated for a brief period. Platoons in automated lanes are separated by less than platoon spacing.

3. An object dropping from above onto the automated lanes presents a situation in which there would have been an accident in the absence of automation; however where there is increased density of traffic on the automated lanes there is an increase in the number of casualties.

4. If a manually controlled vehicle (illegally) enters the automated lanes, unpredictable behavior by the driver can give rise to hazards. It is a weakness of any design in which the intelligence is concentrated in the vehicles that a "rogue" cannot be tracked. In this case, the fact that the rogue will not respond to messages means that it is not detectable by sending messages. Hazards can then arise in the change-lane maneuver. Further, if the forward sensor requires a responder, the responder may be absent from a rogue. In such a case the rogue may not be detected at full range by the forward probe. There are no easy means of predicting the frequency of this occurrence or of its propensity to result in accidents. Much will depend on the probability that rogue drivers can be apprehended.

5. A fifth foreseen hazard is a design error in the original specification of Hsu et al. (*8*). Because this design is based on the work of Hsu et al. it was not thought right to remove it, even though it meant that the safety criterion would be violated. The change-lane maneuver permits a vehicle to enter just ahead of a platoon, which is hazardous because, if the entering vehicle strikes the gate, a high-speed collision results. This is readily remedied by requiring that the platoon in the receiving lane drop back a full platoon spacing.

6. The merge-and-split functions involve situations in which Hazard 1 arises. However, the merge operation is not necessary, provided the restriction that only free agents can change lanes is relaxed. If a vehicle always changes lane either into a large gap or immediately to the rear of a platoon, there is no hazard in that lane. If it leaves a lane from the center of a platoon, there is no hazard in that lane. However, because this work was based on that of Hsu et al. (*8*), these changes were not made in this work.

The fault tree analysis detected all these errors.

Some other errors were also detected, which were real errors. The fact that the fault tree picked them up demonstrates that the method is effective. It also demonstrates that mistakes in design are easily made and that the verification and validation process is indeed necessary to ensure that design meets safety criteria.

The following are design errors that were not foreseen but were detected by the fault tree:

1. If the fault of "loss of the forward sensor" is present at entry, it will not be detected immediately, and hazards can follow. The fault can be corrected by including a check on the forward sensor as part of the entry procedure.
2. If a vehicle is moving slowly and another wishes to change into the lane in which it is moving, the slow mover may not respond to the initial message (request_change_lane) because it is too far away. Nevertheless there is a possibility that the "changer" will catch up with it and change lanes too closely. As the design stands, a lane containing a vehicle that cannot communicate laterally is made a no-entry lane. Lanes with vehicles that are slow moving should also be treated in this manner.
3. If two vehicles each wish to change lanes into the same lane, but from opposite sides, and their speeds are ill matched, a hazard can arise. When the messages request_change_lane were sent, the vehicles were too far apart to need to reply; however, it is possible that at the time of change they are too close.
4. On exit, if a vehicle has a fault in its forward sensor and there is a manually controlled vehicle ahead of it, the separation of the two is unknown. This violates Hazard 2. Any danger can be avoided by causing such a vehicle to leave at a very low speed and to maintain such a low speed until manual control is resumed.

SUMMARY

A particular design to demonstrate the techniques of complete specification and fault tree analysis has been examined. These techniques have been successful here and in another example (*2*). It seems reasonable to conclude that they are suitable methods for examining the safety of all automated highway design concepts.

In the system considered here, lane changes were made only by free agents, and in the process of becoming a free agent, a split maneuver involves a situation in which Hazard 1 arises. This was done because the ideas of Hsu et al. (*8*) were being followed. An alternative is to require that lane changes take place by a vehicle's leaving a platoon from any position and either joining another at the rear or entering a large gap.

The particular design tested is not being suggested as a contender for construction. It is therefore not necessary to reiterate the cycle and attempt to correct the errors. Indeed, the detection of the errors is evidence of the effectiveness of fault tree analysis. However, correction of the unforeseen faults is readily done, as has been shown. It is reasonable to hope that the foreseen ones are rare enough to be ignored, but this remains to be demonstrated.

The following conclusions are made:

1. A specification has been made of an automated freeway. A set of hazards and a safety criterion also have been specified. A number of foreseen circumstances in which the safety criterion is not met seem likely to be rare enough to be accepted.
2. Fault tree analysis was practical. No branch of the tree had more than four elements. Faults were detected, but it is reasonably clear that they can be corrected.
3. It is therefore possible to construct an automated freeway that meets these safety criteria.

SAFETY OF ALTERNATIVE CONCEPTS

The relative merits, from a safety viewpoint, of alternative automated-freeway design concepts will now be considered. The consideration is based on the experience gained in the work reported here plus that of earlier parallel study (*3,4*). That was a study involving a specification and fault tree analysis of a system with a single automated lane and intelligence concentrated in the infrastructure.

The infrastructure-based design (*2,3*) relied heavily, perhaps too heavily, on the integrity of communications. This case is an extreme in the infrastructure basis of intelligence. The control data are passed from vehicle to vehicle in a platoon via the roadside. Any interruption to communication lasting several seconds could produce large disturbances in platoon motion, although the delay would have to be long and other independent disturbances would have to occur also before a hazard arose. How long and how severe are not yet established. But certainly, the more intelligence is concentrated in the infrastructure, the more important the integrity of road-vehicle communication becomes to safety.

Some vehicle-based intelligence is required to keep a vehicle hazard free if communications (vehicle-vehicle or vehicle-road) are less than totally reliable. However, this is not the only need. The system discussed here relies heavily on the performance of the forward sensor. This device, as specified, can be relied on to detect the vehicle ahead of one in the same lane at distances of more than a platoon spacing—perhaps up to 200 m. Roads curve, both vertically and horizontally. At their edges, roads have obstacles to vision. Whether clear sight lines can be guaranteed at this range (the sensor is unlikely to be above the level of the hood) can be doubted. Even more doubtful is the ability to guarantee that an image is that of a vehicle in the same lane. The forward sensor, as specified in this system concept, is not obviously readily implemented.

Alternatives to the use of a forward sensor are possible, In the system discussed in this report we control a vehicle that has lost its forward sensor by using vehicle-vehicle communication. In an earlier report (*2*) the same effect was achieved by the use of vehicle position detectors and a good deal of roadside logic. However, the discussion has turned full circle: it started with an examination of alternatives to totally reliable communications.

There is not necessarily an impasse. A balance of integrity between sensors and communications and a balance of intelligence between the roadside and the vehicle seem likely to offer opportunities to resolve the problem. However, investigation of the safety of such systems would rely heavily on the ability to predict reliability in quantitative terms. One would also need to be able to predict quantitatively the consequences of the loss of reliability, which would involve better understanding of the frequency with which critical configurations occurred in practical operation.

It is concluded, extending the result here, that it is possible to design an automated freeway that meets a required standard of safety. The immediate result here, however, is that it is possible to design an automated freeway that satisfies the safety criteria discussed in this paper.

ACKNOWLEDGMENTS

This work was performed as part of the program of California PATH of the University of California, in cooperation with California's Business, Transportation and Housing Agency and Department of Transportation and FHWA and NHTSA, U.S. Department of Transportation.

REFERENCES

1. A. Hitchcock. Intelligent Vehicle Highway System Safety: Problems of Requirement Specification and Hazard Analysis. In *Transportation Research Record 1318*, TRB, National Research Council, Washington, D.C., 1992, pp. 98–103.
2. A. Hitchcock. *Methods of Analysis of IVHS Safety. Part II: Main Report*. PATH Research Report UCB-ITS-PRR-92-14. University of California, Berkeley, 1992.
3. A. Hitchcock. *A Specification of an Automated Freeway*. PATH Research Report UCB-ITS-PRR-91-13. University of California, Berkeley, 1991.
4. A. Hitchcock. *Fault Tree Analysis of an Automated Freeway*. PATH Research Report UCB-ITS-PRR-91-14. University of California, Berkeley, 1991.
5. A. Hitchcock. *A Specification of an Automated Freeway with Vehicle-Borne Intelligence*. PATH Research Report UCB-ITS-PRR-92-16. University of California, Berkeley, 1992.
6. A. Hitchcock. *Fault Tree Analysis of an Automated Freeway with Vehicle-Borne Intelligence*. PATH Research Report UCB-ITS-PRR-92-15. University of California, Berkeley, 1992.
7. S. E. Shladover. *Operation of Automated Guideway Transit Vehicles in Dynamically Reconfigured Platoons*. Urban Mass Transit Administration Reports UMTA-MA-06-0085-79-1, UMTA-MA-06-0085-79-2, and UMTA-MA-06-0085-79-3, Springfield, Va., 1979.
8. A. Hsu, F. Eskafi, S. Sachs, and P. Varaiya. *The Design of Platoon Maneuver Protocols for IVHS*. PATH Research Report UCB-ITS-PRR-91-6. University of California, Berkeley, 1991.
9. P. Varaiya and S. E. Shladover. *Sketch of an IVHS System Architecture*. PATH Research Report UCB-ITS-PRR-91-3. University of California, Berkeley, 1991.

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

Publication of this paper sponsored by Committee on Intelligent Vehicle Highway Systems.