

Precursor Systems Analyses of Automated Highway Systems

RESOURCE MATERIALS

An Hypothesized Evolution of an Automated Highway System



U.S. Department of Transportation
Federal Highway Administration

Publication No. FHWA-RD-96-051
January 1996

FOREWORD

This report was a product of the Federal Highway Administration's Automated Highway System (AHS) Precursor Systems Analyses (PSA) studies. The AHS Program is part of the larger Department of Transportation (DOT) Intelligent Transportation Systems (ITS) Program and is a multi-year, multi-phase effort to develop the next major upgrade of our nation's vehicle-highway system.

The PSA studies were part of an initial Analysis Phase of the AHS Program and were initiated to identify the high level issues and risks associated with automated highway systems. Fifteen interdisciplinary contractor teams were selected to conduct these studies. The studies were structured around the following 16 activity areas:

(A) Urban and Rural AHS Comparison, (B) Automated Check-In, (C) Automated Check-Out, (D) Lateral and Longitudinal Control Analysis, (E) Malfunction Management and Analysis, (F) Commercial and Transit AHS Analysis, (G) Comparable Systems Analysis, (H) AHS Roadway Deployment Analysis, (I) Impact of AHS on Surrounding Non-AHS Roadways, (J) AHS Entry/Exit Implementation, (K) AHS Roadway Operational Analysis, (L) Vehicle Operational Analysis, (M) Alternative Propulsion Systems Impact, (N) AHS Safety Issues, (O) Institutional and Societal Aspects, and (P) Preliminary Cost/Benefit Factors Analysis.

To provide diverse perspectives, each of these 16 activity areas was studied by at least three of the contractor teams. Also, two of the contractor teams studied all 16 activity areas to provide a synergistic approach to their analyses. The combination of the individual activity studies and additional study topics resulted in a total of 69 studies. Individual reports, such as this one, have been prepared for each of these studies. In addition, each of the eight contractor teams that studied more than one activity area produced a report that summarized all their findings.

Lyle Saxton
Director, Office of Safety and Traffic Operations
Research
and Development

NOTICE

This document is disseminated under the sponsorship of the Department of Transportation in the interest of information exchange. The United States Government assumes no liability for its contents or use thereof. This report does not constitute a standard, specification, or regulation.

The United States Government does not endorse products or manufacturers. Trade and manufacturers' names appear in this report only because they are considered essential to the object of the document.

CONTENTS

	PAGE
PART I. INTRODUCTION AND SUMMARY	
THE SCENARIO IN BRIEF	1
A NOTIONAL BASIS FOR DESIGN: EMULATE THE HUMAN	5
SUMMARY OF THE PRIMARY ISSUES AND RISKS	6
THE PROSPECTS FOR SYSTEM DEVELOPMENT	6
COST AND PRICES	7
SYSTEM ROBUSTNESS AND FAILSAFETY	7
SAFETY CRITERIA	7
AUTOBRAKE - OBSTACLE DETECTION REQUIREMENTS	7
AUTOBRAKE - ~ CRITERIA	7
AUTOGAP - SENSOR INTERPRETATION PERFORMANCE	8
AUTOGAP - COMMUNICATION <u>WITH</u> THE TRAFFIC	
MANAGEMENT SYSTEM	8
AUTOLANE AND AUTOMATIC CRUISE CONTROL - DRIVER	
ALERTING	8
AUTOPLATOON - DRIVER AND PASSENGER ACCEPTTANCE OF	
PLATOONING PROXIMINY	8
THE PLATOON-DEPLATOON DECISION	8
SUBSEQUENT EVOLUTION - SYSTEM ARCHITECTURE	8
PART II. DISCUSSION AND ANALYSIS OF THE INDIVIDUAL FUNCTIONAL	
STEPS	9
GENERAL PRESENTATION SCHEMA	9
1. THE AUTOMATIC BRAKING FUNCTION (AUTOBRAKE)	10
OPERATIONAL CONCEPT	10
TECHNICAL DESCRIPTION	10
POTENTIAL ~ - MOTIVATION TO PURCHASE	19
DEVELOPMENTAL ISSUES AND RISKS	20
2. THE AUTOMATIC GAP HOLDING FUNCTION (AUTOGAP) -	
INTELLIGENT CRUISE CONTROL	21
OPERATIONAL CONCEPT	22
TECHNICAL DESCRIPTION	22
POTENTIAL BENEFITS—MOVITATION TO PURCHASE	24
DEVELOPMENTAL ISSUES AND RISKS	25
3. THE AUTOMATIC LANE HOLDING FUNCTION (AUTOLANE) -	
AUTOMATIC CRUISE CONTROL	26
OPERATIONAL CONCEPT	26
TECHNICAL DESCRIPTION	27
POTENTIAL BENEFITS	28
DEVELOPMENTAL ISSUES AND RISKS	28
A DIGRESSION: DEPLOYMENT AND THE DYNAMICS OF COST	
BEHAVIOR	30
4. THE SPONTANEOUS PLATOONING FUNCTION (AUTOPLATOON) -	
ADVANCED AUTOMATIC CRUISE CONTROL	33
OPERATIONAL CONCEPT	33
THE PLATOONING-DEPLATOONING DECISION	33
TECHNICAL DESCRIPTION	35

POTENTIAL BENEFITS - MOTIVATION TO PURCHASE	36
DEVELOPMENTAL ISSUES AND RISKS	40
	PAGE
PART III. SUBSEQUENT EVOLUTION OF AHS AND AUTOMATED	
VEHICLE CONTROL	43
AUTOMATIC LANE CHANGE	43
SUPERCUISE	43
SURFACE STREET OPERATION: THE INTEGRATED IVHS	44
POSTSCRIPT	47
APPENDIX A.SPONTANEOUS PLATOONING: THE CALCULATION OF	
PROBABILITIES AND PARAMETERS	

ILLUSTRATIONS

Figure		Page
1.	The Evolution of Automated Vehicle Systems	5
2.	Generic Automated Vehicle Control System Functions	9
3.	Generic Automated Vehicle Control System	9
4.	Emergency Braking Illustration	11
5.	Faster Reactions and Better Brakes Permit Closer Operation	12
6.	Safegap	12
7.	Machine Reaction Times Permit Closer Spacing - Brakes, Not Time Become Dominant	13
8.	The Relationships Among the Gap Between Vehicles, Speed, and Flow	14
9.	The Relationships Among the Gap Between Vehicles, Speed, and Flow	14
10.	Stopping Distance as a Function of Speed and Braking Capability	16
11.	Typical Scene Looking Forward	18
12.	Autobrake Mechanization	19
13.	Autobrake Mechanization: Functional Description of Technical Elements Required	19
14.	Automatic Gap Control	22
15.	Autobrake and Autogap Mechanizations: Functional Description of Technical Elements Required	23
16.	Autobrake and Autogap Mechanization - Intelligent Cruise Control	24
17.	Potential Capacity Benefits from Gap Control	25
18.	Auto Lanehold Mechanization: Functional Description of Technical Elements Required	27
19.	Adding Auto Lanehold: Automated Cruise Control Hands-Off Freeway Cruising - The Mark I AHS	29
20.	An Illustration of the Dynamics of Cost Behavior - I	31
21.	An Illustration of the Dynamics of Cost Behavior - II	31
22.	An Illustration of the Dynamics of Cost Behavior - III	32
23.	Spontaneous Platooning Mechanization: Functional Description of Technical Elements Required	36
24.	Add Spontaneous Platooning: Advanced Automated Cruise Control	37

25.	Density, Spacing, and Platooning Relationships	37
26.	Lane Flow as a Function of Proportion of Vehicles Equipped with the Autoplatoon	39
27.	Variation in Platoon Size Mix with Proportion Equipped	39
28.	Add Automated Lane Change: Fully Automated Freeway Trip	44
29.	TMS-AVCS Integration Fully Automated on Freeways and Traffic Controlled Surface Streets	45
30.	Automated Manual Control - The Safe Vehicle	46

PART I

INTRODUCTION AND SUMMARY

It is frequently assumed that an Automated Highway System (AHS) will require lanes exclusively dedicated to automated operation, and that operation in mixed traffic is either not safe or not practical. Since these assumptions preclude automated vehicles from sharing existing lanes with manual traffic, they imply the introduction of an AHS will require either new construction or diversion of existing lanes from their current use, and at a time when there are likely to be few vehicles equipped to use them.

If these assumptions are indeed valid constraints, they make the initial deployment step a very steep one. First, there would be the need to justify investment in infrastructure well before there are compensating benefits. At least as serious is the prospect of the heat that will be generated by drivers who must continue to endure the unrelieved, or possibly worsened, congestion in full view of this underutilized toad space.

An Automated Highway System is much more likely to be brought into being if the operational concept and the system that supports it are compatible with a gradual and nondisruptive introduction into the existing freeway system operational environment.

Here we describe and broadly analyze an evolutionary scenario in which the vehicles equipped for automated operation are assumed to be capable of safely operating in mixed traffic with unequipped, manually operated vehicles. We assume a gradual evolution of incremental change, both in the scope and capability of the automated features, and in the vehicle fleet as new, equipped vehicles replace old ones.

We believe the scenario described is both technically and operationally feasible, and offers the prospect of an essentially seamless and nondisruptive path to an Automated Highway System - and beyond. In devising this scenario we have attempted to meet two primary criteria:

- The evolution of the system should be technically sensible, with each step building upon previous steps. We do not want Step #8 to obsolete Steps #2 and #3.
- At each step in deployment, there should be reasonable correlation in time and degree between costs and benefits. Investment, either by individuals or by political bodies, is only motivated by perceived benefit.

Our first step has been to hypothesize what is believed to be a sensible technical evolution for an AHS, using our best judgment and knowledge from already available analyses. Since every technical step does not necessarily constitute the basis for a saleable product, we have also defined a sequence of deployment steps for bringing the hypothesized system into operational use. There has been no intent to carry out a detailed system analysis; if the results are judged to be promising, that can be done in a later stage of the program.

In our hypothesized scenario, the initial AHS is designed for operation where our current congestion problems are most obvious: on urban freeways. In addition to the freeway analysis, the scope of the potential usefulness of the particular function for operation in other venues is briefly addressed in materials and other surface streets, interstates, other rural highways, and intersections of various types.

THE SCENARIO IN BRIEF

There are experimental cars on our highways today that are already taking the first steps toward an AHS: these are cars with their Cruise Control units augmented by a forward looking sensor that detects when one might be getting too close to the car in front and lets up on the throttle perhaps even down-shifts. This adds a kind of minimum gap function to the normal speed-hold function to produce what is being called *Intelligent Cruise Control*.

We can raise the IQ of this Intelligent Cruise Control (ICC) by letting this same sensor - actually a somewhat better one - also provide automatic emergency braking. Now our much smarter ICC would be capable of two functions, automatic emergency braking (Autobrake) and automatic use of both throttle and brake to maintain a safe gap behind the car in front (Autogap).

It is purely a conjecture, but we suspect the reason the first versions of Intelligent Cruise Control are of the low IQ variety - that is, without Autobrake, using throttle control only - is, first, that the more capable system does represent a more difficult technical problem, and, second, the specter of product liability is a much greater concern.

We will henceforth use Intelligent Cruise Control (ICC) to refer to the smarter version, including both the Autobrake and Autogap functions. From the driver's perspective it will operate much like Cruise Control today, except that when his or her car overtakes another vehicle, the system will automatically revert from holding speed to holding fixed-but-driver-adjustable distance behind that vehicle. Thus at some moderate transition distance the speed controls become position controls, varying the gap between the vehicles in response to the driver's desires. Like Cruise Control today, a tap on the brakes or the throttle gives drive-train control back to the driver.

The autobrake feature operates at all times, invisible to the driver. In fact if the driver is alert enough to never have a lapse in concentration and lucky enough to never have someone going slower cut directly in front of him, the system will never be actuated, because the driver's braking will always occur before the point that emergency response is needed. In addition to sensing when emergency braking is needed, the system will also be designed to recognize when it is not needed, such as in sharp turns when objects may suddenly appear dead ahead, but are not threats to safety.

We anticipate that these systems will incorporate a complete self-testing and self-diagnostic capability; and will be designed to be completely fail safe and fail soft. We believe this to be possible, but it will require extensive design effort and verification testing. We suspect that the major portion of the developmental engineering and much of the testing will be focused on providing and verifying these features.

Given this self-check capability; these systems create no special requirements for entering the freeway; the driver behaves as he or she does today. If there had been a malfunction in the system, the self-test would have already indicated it, and the driver would know that auto-operation was precluded until the system was repaired. The system will have been designed to fail softly if a malfunction occurred during operation.

There are at least two powerful motivations for Intelligent Cruise Control. The first is safety. Artificial sensors don't get tired or fall asleep, and their attention doesn't wander. lapses in attention are the largest source human accidents; automated Systems could virtually eliminate these. This one trait of almost eternal vigilance is worth a large investment itself.

The second advantage, one that should have great appeal to drivers who have to spend long hours on the

freeway, is relief from the constant brake-throttle-brake jockeying needed to hold position in heavy traffic.

There is a third benefit that may or may not turn out to be important: the potential for increased throughput. Automated controls are potentially very fast, it is reasonable to expect that automated systems could react to emergencies on the order of five to ten times faster than a human driver. This capability offers the potential to increase the effective throughput of the freeway, because it enables vehicles to drive closer together safely.

There will be acceptance problems to overcome. For example, because of the potential for shorter headways, some drivers will perceive themselves to be the victims of uncomfortably close tailgating. This new driving behavior will need to be made acceptable to fellow drivers. Simple techniques may be enough. Perhaps a small green light visible in a rear view mirror could tell drivers that the vehicle behind them has automatic braking. Information programs could help drivers understand what is going on. Time and gradualism should help.

We visualize this introduction of Intelligent Cruise Control (ICC) as the first step toward the Automated Highway System. It will require no modification of the infrastructure, and equipped vehicles will operate in mixed traffic with unequipped vehicles. The primary motivations for purchase are improved safety and more relaxed driving in freeway traffic.

The next major function is automatic lane holding (Autolane). Autolane will require on-board sensors that can accurately determine vehicle position relative to the roadway lane. This implies the need for sensing the road itself, perhaps by tracking the white lines, or perhaps guiding on embedded markers such as the magnetic "nails" being used in the PATH program experiments. There may be other techniques. We conjecture that an operational system will combine several such methods to enhance robustness and failsoftness.

We also conjecture that lane modifications such as painting brighter lane demarcations or installing magnetic nails will be sufficiently inexpensive and practical that all lanes in a freeway can be modified. This would enable drivers with equipped vehicles to operate in the lane that matches his or her speed preference.

It is our current opinion that Autolane will be used only in conjunction with Autogap and Autobrake, not alone. With this combination, the driver would manually enter the freeway as he or she does today, drive to the lane desired, and activate these modes, perhaps with a single control. Now the vehicle is under completely automatic control.

The vehicle would remain in this mode until the driver indicates his intent to take control by turning the system control back to "Off", and actually disengaging the system by grasping the steering wheel and introducing some small manual input, overriding the automatic steering.

We suggest that a submode of Autolane is automatic lane departure warning. This could be used during manual control to prevent accidental drifting out of the lane, a significant safety feature for the sleepy driver.

These new capabilities need not be restricted to just freeway use. Since the infrastructure modifications needed to support the Autolane feature should be relatively inexpensive and disruptive, its use could readily be expanded to interstates and many lesser rural highways. The payoff would be a marked improvement in both safety and in driver convenience. It should be a particular boon on long intercity trips.

We have now attained the "Mark I" Automated Highway System, which we might appropriately call Automatic Cruise Control (ACC). It will require modification of lanes to support the Automatic lane Holding. It will operate in mixed traffic, sharing these lanes with unequipped vehicles. The primary motivation for purchase is the fully automatic, hands-off cruising on urban freeways and intercity trips.

Avoid about safety criteria is appropriate. There is frequently a feeling that if it is automatic it must be made 100 percent safe in all circumstances. That is an impossible criteria for man or machine. It seems much more rational to set the standard in relation to what happens today; for example, we could that the automated system must be twice as safe as manual control, or five times as safe, but not 100 percent safe. The perfect is the enemy of the good.

The next step in our hypothesized evolution is the augmentation of Automatic Cruise Control with the capability to platoon. We have hypothesized an approach in which the platooning takes place spontaneously. In this mode, the individual vehicles under full Automatic Cruise Control (Autobrake plus Autogap plus Autolane) will be sensing and deducing the state of the traffic: is there plenty of capacity in the lane or is more capacity needed to accommodate additional vehicles desiring entry into the lane? If a vehicle senses the need for more capacity, and if it is next to a vehicle also under autocontrol, it will spontaneously move up to platoon with that vehicle. This creates more space in the lane for other vehicles. Possible techniques for implementing this capability are discussed in the Autoplatooning Section in Part II.

Autoplatooning is envisioned to also work in reverse: deplatooning occurs when there is no longer need for it to increase lane capacity. It also occurs when the driver signals that he wants to assume manual control, probably in order to move out of the lane and exit the freeway.

The sole purpose of planning is to increase the effective capacity of the freeway. It is dubious that people will like it for themselves, but purchase may be motivated by the recognition that it enhances the prospects of a free-flowing freeway, or by partial government subsidy. ACCs equipped with this feature. If absolutely necessary it could be by decree.

This evolution would continue into the fixture, although our crystal ball begins to get much cloudier. Probably the next step is to provide automatic lane change. This enables the entire freeway trip to be automated: entry into the freeway, movement to the desired lane, then subsequent exit from the freeway. This is fairly ambitious technically, but in a decade should appear much less so. Fully automating the freeway portion of a trip may have considerable appeal to our growing contingent of elderly drivers.

An additional possibility is very high speed intercity travel. It seems likely that this step would require the dedication of special lanes, but whether they would require permanent physical separation to the normal speed lanes is an open issue.

Somewhere in this evolution the step will be taken that we believe will be the true watershed in surface transportation: the integration of the Automated Vehicle with the Traffic Management System, the Smart Car with the Smart Road. This will begin to happen when we first start to introduce the automated vehicle to surface streets. The first steps may be undramatic, but they will open the door to the fully automated origin-to-destination vehicle.

Once the automated system can permit hands-off operation on surface streets, we can begin to think in terms of total trip automation, probably with the basic navigation coming from the vehicle's own guidance system.

Somewhere in the more advanced stages of this evolution of the Automated Vehicle we can expect another important metamorphosis. We will review the steps so far to better lay out the overall logic.

Beginning with Autobrake and Autogap, we have modified the vehicle so that the steering and drive-train controls are actuated by electrical signals, not muscle: we have replaced drive-by-muscle with drive-by-wire. With drive-by-wire, the driver, by the manipulation of his controls, is sending an electrical signal to the actuators, not a mechanical movement.

We have then progressively taken the control of these actuation signals away from the human eye-brain system and substituted, function-by-function, a "computerbrain" system fed by artificial sensors. We have built into that computerbrain all the responses to keep the vehicle safe in a very wide variety of driving situations.

The next logical step is to use the computerbrain system to insure safety during manual driving. Now when the driver wants to manually control the vehicle he does so not by sending electrical signals to the steering and drive-train controls, but by sending his signals to the computerbrain system - he drives *through* the artificial system, not *around* it. The vehicle follows his gross steering signals, but doesn't permit him to do anything unsafe. This will be a boon for the elderly portion of our population - and for the people who share the road with them.

This broad pattern of evolution is illustrated in Figure 1. We can be almost sure, however, that it will not happen in the three neat steps we have laid

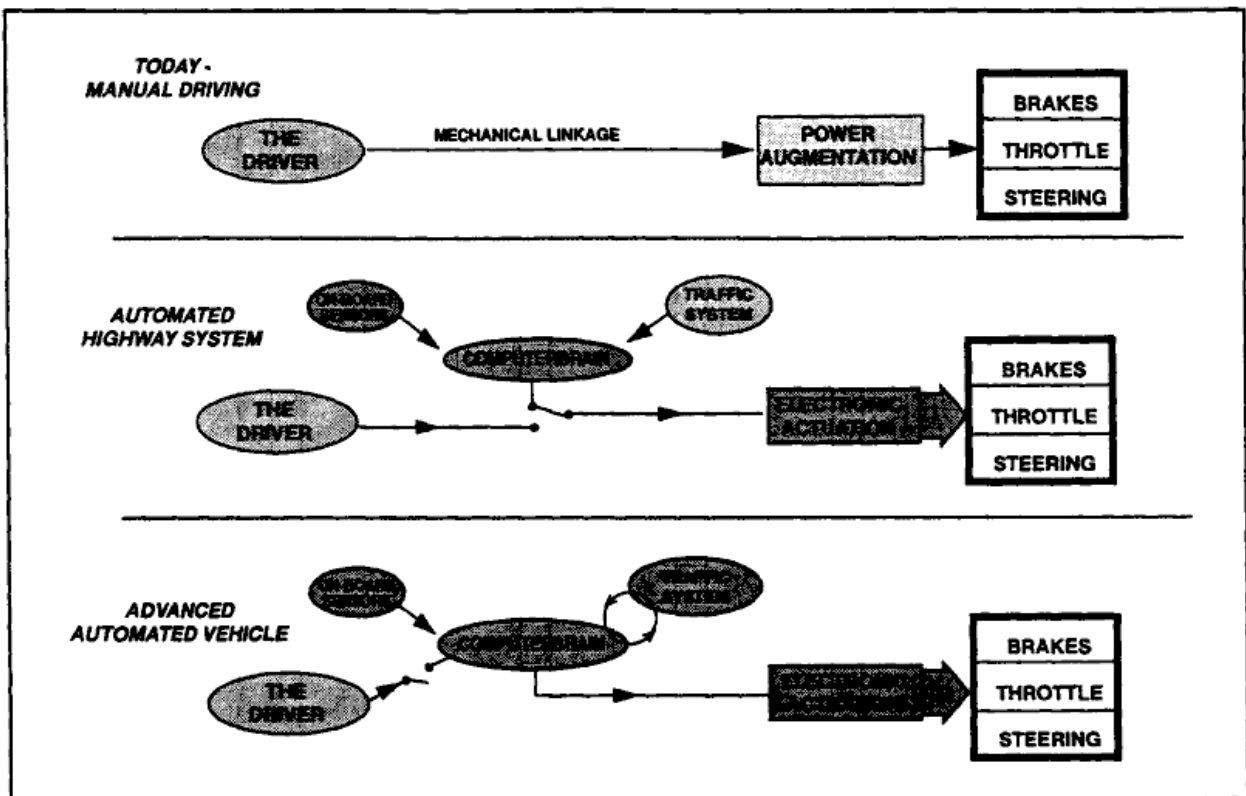


Figure 1. The Evolution of Automated Vehicle Systems

out; rather, like everything else, it will happen in bits and pieces, in small, almost experimental steps. For example, once we are comfortable with Intelligent Cruise Control - fully debugged by millions of hours of operational experience - we may well see some new cars whose throttles are inputs to the ICC unit, not the engine.

All of these steps are treated in more detail later in Parts II and III.

A NOTIONAL BASIS FOR DESIGN: EMULATE THE HUMAN

Consider how a human drives in a freeway lane. He or she gets most of their information on the state of traffic around their vehicle through their own eyes, and they keep track of the location of their vehicle on the road the same way. They use all the visual cues available: lane markers, "road ~~drives~~ ahead", et cetera. They look most closely at the relatively few vehicles that could endanger their own vehicle, like the one in front of them if it stops suddenly, or the one in the next lane if its turn signal is on or if it decides to pull in front of them.

Second, the driver is continuously interpreting all this information. The driver decides what is a safe following distance. He has some "rules" for deciding, like it is safe to follow closer at low speeds than at high speeds. In wet weather the braking is not as good, so the gap - the following distance - should be larger. And so on. (It is apparent that not all drivers subscribe to the same rules.)

Drivers also have rules for taking action. If he is gaining slowly on the car in front, he backs off slightly on the accelerator (or thinks about passing). If he is closing a bit faster, he applies gentle braking. If he has glanced down to retune his radio, and looks back up to see the car in front of him suddenly closer and closing rapidly, he slams on his brakes and stimulates his adrenaline.

If we set out to automate this "gap control" function, and do it entirely on-board the vehicle, then emulating this pattern of sensing and interpretation, decision, and action seems a sensible approach to mechanization.

The human eye, along with the brain that interprets its signals, is ~r more capable and versatile than any artificial sensor-plus-interpreter that could be built today. But gap holding does not require all that capability; nearly all that is needed is an unambiguous measure of the distance to the car ahead and to the few vehicles in the immediate vicinity that might move into the intervening gap, thus becoming the new target.

The point is that if we deal with just a few functions at a time, it is only necessary to emulate very specific and limited functions of the eye-brain, not duplicate its total capability.

This is not intended to imply that the job is easy, nor that we wont want even more information for more capable systems in the future. Even now, just to get unambiguous range, the system must be smart enough to not unintentionally measure the distance to the car in the next lane because the road is curved. It should also be able to recognize when another vehicle (or motorcycle) starts moving into the lane. Or to react appropriately when the car ahead moves out of the lane. The actual range measurements can be made by man-built sensors with greater accuracy than a human can estimate it; the more complex part is interpreting the sensor outputs.

It is still to be determined if the best way to handle this general sensing flinction is with vision-based sensors, radar, radar, or some combination. And the best approach to interpretation of the output of these sensors is also yet to be defined, and is, in fact, the most difficult technical challenge facing these new systems.

The decision rules that govern vehicle responses are easier to visualize. Compared to "expert Systems,, that can substitute for human decision making in general areas like law and medicine, the area of our interest is comparatively limited: it seems well within current capability to devise a set of rules, probably couched in "fuzzy" terms, that adequately decide when and how vigorously to put on brakes or change the throttle setting.

So while the versatility of the eye or the data processing capability of the brain will never be matched in its entirety, for very specific functions it may be possible to actually improve upon it. We can measure distances and rates of closure with great precision. We can have "eyes-in-the-back-of-head" - and wherever else we want them. In time we should be able to see through fog.

This perspective on system mechanization is consistent with the notion of incremental evolution just described. The start is with simple functions. As sensor-interpretation capabilities grow, and our ability to define behavioral rules for more complex control situations improve, the scope of automation could be expanded a step at a time within the same basic mechanization framework. As the need for information broadens beyond that which can be acquired by vehicle-mounted sensors, then communication links to infrastructure intelligence can be introduced; we have conjectured that this will occur when auto-operation on suffice streets is introduced.

While the "eyes-and-brains" of the system we have been discussing are the key challenge, it is obvious that it will also be necessary to replace human muscle in the operation of the vehicle's throttle, brakes and steering: we will require vehicles designed for drive-by-wire. This is relatively straight-forward engineering.

The idea of an autonomous vehicle is not new, and appears to have been the primary focus of automated vehicle research in Japan for at least the last few years. A 1990 paper from Mazda¹ lays out quite logically the philosophy and the broad mechanization.

SUMMARY OF THE PRIMARY ISSUES AND RISKS

THE PROSPECTS FOR SYSTEM DEVELOPMENT.

If no one is willing to develop the system and offer it for sale, then all is for naught.

In our judgment, there are two areas of risks that could seriously inhibit private sector investment in the program. The first is product liability risk: we can be fairly sure that fault for all future rear-enders will be shifted from drivers to the deeper pockets of the manufacturers of the ICC system. This prospect must give a potential supplier of the equipment considerable pause for thought.

The second is the uncertainty in predicting the cost of development and validation. Costs are very much a flinching of the technical difficulty of the tasks being undertaken. We view the technical problem of devising a sensor and sensor interpretation system that can provide the performance desired, keep false alarms to an acceptably low level, and incorporate the self-verification and failsafe features required as one that will require considerable innovation and ingenuity. And it will be necessary to produce the product at a cost that will sell. Not only are these development costs hard to predict with confidence, but in this case the difficulty is compounded by the ever-present possibility of unanticipated regulatory mandates.

We visualize that a primary objective of public sector participation in the program is to reduce such

risks to proportions that do not inhibit private sector investment.

¹SAE Paper 90~484, Okuno et al, Towards Autonomous Cruising on Highways Mazda, August 1990
COST AND PRICES.

This is just an explicit statement of the risk just discussed. Unless prices are low enough to attract buyers, the potential benefits of the system will not be realized. Prices may be purely a flinching of actual costs, or they may be reduced by possible subsidies that may be justified by the public benefits the system brings.

SYSTEM ROBUSTNESS AND FAILSAFETY

This is an issue, perhaps more properly considered a risk, that is common to all functions.

As noted, an automatic self-test and self-diagnostic system would be an essential part of any system design. We believe that the engineering development of this Integrity Verification Subsystem and the testing required to insure its proper functioning will be a major element of development cost.

SAFETY CRITERIA

Another issue is the choice of safety criteria. The choice made will drive both the performance and the cost of the system. As we have noted, there is an almost a knee-jerk reaction to demand perfection, which is, of course, not attainable. We suggest a more rational criteria might be set by comparison to what is now obtained with human drivers. Perhaps something like twice as good, or even five times as good. The final criteria should not be selected until we have a better understanding of the cost of attainment.

The remaining issues have features that are unique to individual functions.

AUTOBRAKE - OBSTACLE DETECTION REQUIREMENTS.

A serious performance issue arises in determining the obstacle detection requirements for the initial systems. It will be easier and cheaper to build a system that just tracks vehicles in one's own lane out to, say, 100 feet than to simultaneously track vehicles in adjacent lanes, or see out to 300 feet. There is the issue of the size of the objects that must be detected (and broadly identified?): motorcycles?, bicycles?, people: adults, children?, dogs and cats?, trash cans?.

Another dimension of this issue is the level of permissible degradation with weather: fog, rain, snow...

The choices made here could have a substantial impact on both system capability and system cost, and it would be wise to defer final decisions until the trades are thoroughly understood.

AUTOBRAKE- BRAKING CRITERIA

This design issue is a particular of the safety criteria issue: the determination of the criteria for a safe following distance. This involves assumptions about relative braking capabilities, road surfaces, reaction times, and a few other secondary variables.

The criterion chosen implicitly determines the trade between safety and capacity.

The system could be designed so that it is adjustable, so the initial choice could be readjusted as operational experience became available.

AUTOGAP - SENSOR INTERPRETATION PERFORMANCE.

There is the general issue of just how good the sensor and sensor interpretation system have to be in assessing the tactical driving situation. We want the system to be able to behave as a prudent and defensive driver would. How close do we need to come, and how do we articulate the requirement? It is our suspicion that we will discover that we ultimately need a sensor interpretation system that is fir better than we will settle for in initial systems:
sensor interpretation is at the heart of both system performance capability and system safety

This issue applies in differing degrees to autobrake and the other functions. It will leap in importance when we reach automatic lane change, which has not been analyzed here.

A specific detail: do we need to be able to read turn-signals. If the answer is yes, then a radar only system will be inadequate.

AUTOGAP - COMMUNICATION WITH THE TRAFFIC MANAGEMENT SYSTEM.

We initially assumed that Intelligent Cruise Control is completely autonomous. There is the issue of whether a communication link to the Traffic Management System is worth the cost. It may be desirable to permit speed commands directly to the vehicle's **Computerbrain** rather than to the driver as we do today

AUTOLANE AND AUTOMATIC CRUISE CONTROL -DRIVER ALERTING.

Automatic Cruise Control removes the motivation - if not the need - for driver vigilance. The implication is that the driver may not only take longer to react to an emergency or an unusual situation, but may also be in a mental state that is not conducive to an immediately rational response. This problem will require attention.

If communication with the TMS was not made a requirement for ICC, it will have to be reexamined here.

AUTOPLATOON -DRIVER AND PASSENGER ACCEPTANCE OF PLATOONING PROXIMITY.

We do not know if or how quickly people can adjust to the very close spacing of vehicles at speed that platooning involves. The transition itself may be disturbing. This is a significant human factors issue.

PLATOON-DEPLATOON DECISION.

The primary operational4technical issue is the way in which this decision is made.

SUBSEQUENT EVOLUTION - SYSTEM ARCHITECTURE.

Heretofore there has been minimal interaction between the essentially autonomous AHS and the Traffic Management System - just the ability for the TMS to issue occasional stream speed commands, and receive emergency transmissions from vehicles in distress. The AHS architecture was almost entirely independent from the infrastructure system architecture.

This is no longer true when automated operation is extended to surface streets: now an intimate integration of the two systems is required. This should be the primary focus of the 20 year version in the on-going architecture studies. This is the true IVHS system.

PART II

DISCUSSION AND ANALYSIS OF THE INDIVIDUAL FUNCTIONAL STEPS

GENERAL PRESENTATION SCHEMA

We will begin the discussion of each functional step with a brief description of the Operational Concept envisioned. This is a statement of the broad goals of the function, and the general operational approach selected to attain them.

The next section is the Technical Description of the system. The presentation logic derives from a generic decomposition of the automatic control function. This is illustrated in Figure 2, which shows the three classic system sub functions that must be carried out to effect automatic controls of almost any

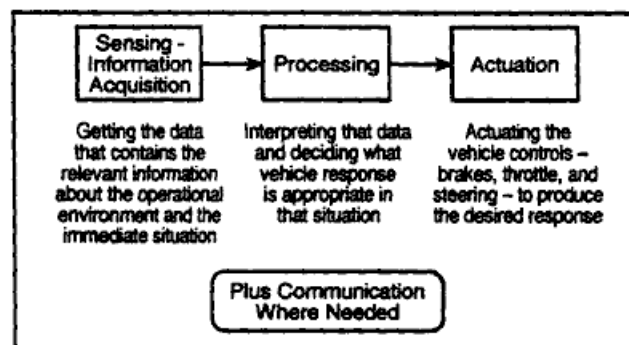


Figure 2. Generic Automated Vehicle Control System Functions

kind. These apply whether we are thinking of just one function, like automatic braking, or the total three-axis vehicle control to be described later.

Figure 3 is a finer grain decomposition of these three sub functions; we will use this more descriptive depiction of system sub functions as our guide for describing the system at each stage and for illustrating evolution over time.

Contrary to the advice to Alice from the Red Queen, the best place to begin thinking about the design of the system is not in the beginning, but in the middle: the first logical step is to determine the information we will need to make the system operate as we want it to.

In turn, the information needs are driven by how we want the automatically controlled vehicle to respond to various driving situations to accomplish the goals identified in the Operational Concept. The information is that needed to recognize those situations and thus trigger the desired vehicle responses. Thus the most logical starting point is the Vehicle Response Determination step in Figure 3.

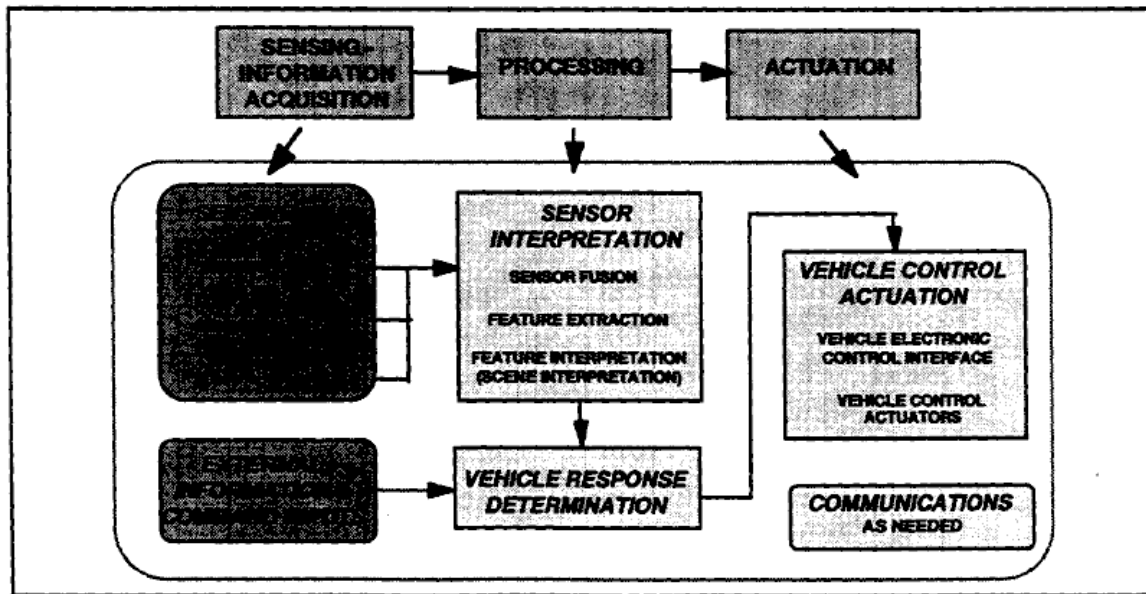


Figure 3. Generic Automated Vehicle Control System

We discuss this in the section labeled *Desired Vehicle Response - Information Needs*.

Knowing what information is needed to produce the desired vehicle behavior leads us to the second design step: identifying how that information can be obtained. This drives the choice of sensors and the sensor interpretation subsystems. This becomes the second part of our Technical Description section, *Sensing and Sensor Interpretation*.

Third, we address specifically the information best obtained from outside the vehicle, and the communication requirements it implies. The communications needed are primarily dictated by the location of the various system elements and the nature of the data flow between them. This is the *External Information & command Inputs- The Need for Communications*.

Next, we give passing attention to *Vehicle Control Actuation*, which is relatively straightforward engineering, responding to the requirements for precision and speed that will be developed in the future during the detailed system design process.

In the last section of the Technical Description we illustrate the overall *System Mechanization*.

After the Technical Description, the next section addresses the potential benefits of the system, and therefore the factors that might motivate its purchase by potential users.

Last, we summarize the issues and risks we foresee in its development and deployment.

We begin with the Autobrake function.

1. THE AUTOMATIC BRAKING FUNCTION (AUTOBRAKE)

OPERATIONAL CONCEPT.

The purpose of the Autobrake function is to improve safety by reducing the number of accidents. We do not expect that Autobrake would be offered for sale as an individual system, but as one of the functions embodied in the mature Intelligent Cruise Control system.

We envision this system as primarily a back-up to the human driver; it would be actuated only if a lapse in attention or a sudden intrusion into his or her path created a dangerous situation. Autobrake would operate by constantly sensing the distance to the vehicle in front, and if the distance and relative velocities were such as to indicate immediate danger of collision, it would automatically apply the brakes.

Because the reaction time for the autobrake system is much faster than for the human driver, the vehicle equipped with Autobrake can operate safely much closer to the vehicle in front than should the manually controlled vehicle. While, in time, some drivers may take advantage of this extra margin of safety to operate closer than they had before Autobrake was available, most will probably maintain old habits and try to maintain the following distance they perceive to be safe under their own control -with normal human reaction times. The consequence of this behavior is that Autobrake becomes essentially invisible to the driver, because it is actuated only when the driver has failed to react as he or she should.

This point is illustrated in Figure 4, which shows the trajectories of Vehicle A at 60 mph is overtaking Vehicle B at 40 mph. For the driver who maintains the following distance he or she is accustomed to, the onset of manual braking should occur several seconds ahead of the onset of automatic braking needed to maintain an equally safe Autobrake following distance.

TECHNICAL DESCRIPTION.

We now turn to a technical description and analysis of Autobrake. The overview summary of this section is presented later in Figure 13.

While the treatment here is sketchy, it presents a higher level of detail that most readers desire, so they will be forgiven if they skip directly to - POTENTIAL BENEFITS - MOTIVATION TO PURCHASE

Desired Vehicle Responses - Information Needs.

We have defined four operational situations that will collectively influence the requirements placed on the individual subsystems. The first is the




Figure 4. Emergency Braking Illustration

common situation of one vehicle following another: we want to avoid collision if the first vehicle brakes abruptly. The second emergency braking situation can occur if a slower vehicle intrudes into the lane. The third, and probably the most demanding on system requirements, is the extreme case of overtaking: a vehicle is stopped - essentially parked - in one's lane. The fourth is intermediate to number one and number three: the car we are following rear-ends a stopped vehicle. We will address these in turn.

Normal following. We have defined a term "safegap" to refer to the distance between vehicles at which it is possible to avoid collision if the lead car suddenly decelerates at its maximum capability. "Safegap", by our definition, assumes both vehicles are traveling at the same speed; if they are not, the gap that is safe is, of course, differs from "safegap" by some function of the rate of closure between the vehicles.

This safegap distance is primarily a function of the velocity of the vehicles, the magnitude of the deceleration of the lead vehicle (Vehicle A), the time delay of the following Vehicle B in apply its brakes after Vehicle A brakes, and the magnitude of Vehicle B's deceleration. Safegap defines the minimum safe following distance of Vehicle B behind Vehicle A. If B is closing on A, it is necessary to start braking soon enough that when the rate of closure has dropped to zero, the actual distance separating the vehicles is not less than safegap.

Figure 5 illustrates the effect of system reaction time on collision velocity as a function of the original spacing between vehicles. Safegap is defined by the distance at which collision velocity goes to zero; in this illustration about 130 feet if the system time lag is 1 second, and about 50 feet if the system time lag is one-tenth that value.

Figure 5 suggests that the algorithms used to calculate safe separation distances will need to account for the effect of system time delays. This probably can be included as a predetermined value.

Figure 6 shows that safegap is a strong, nearly linear function of velocity, and again emphasizes the importance of system reaction time. Thus the safe separation algorithm will clearly require vehicle velocity as an input.

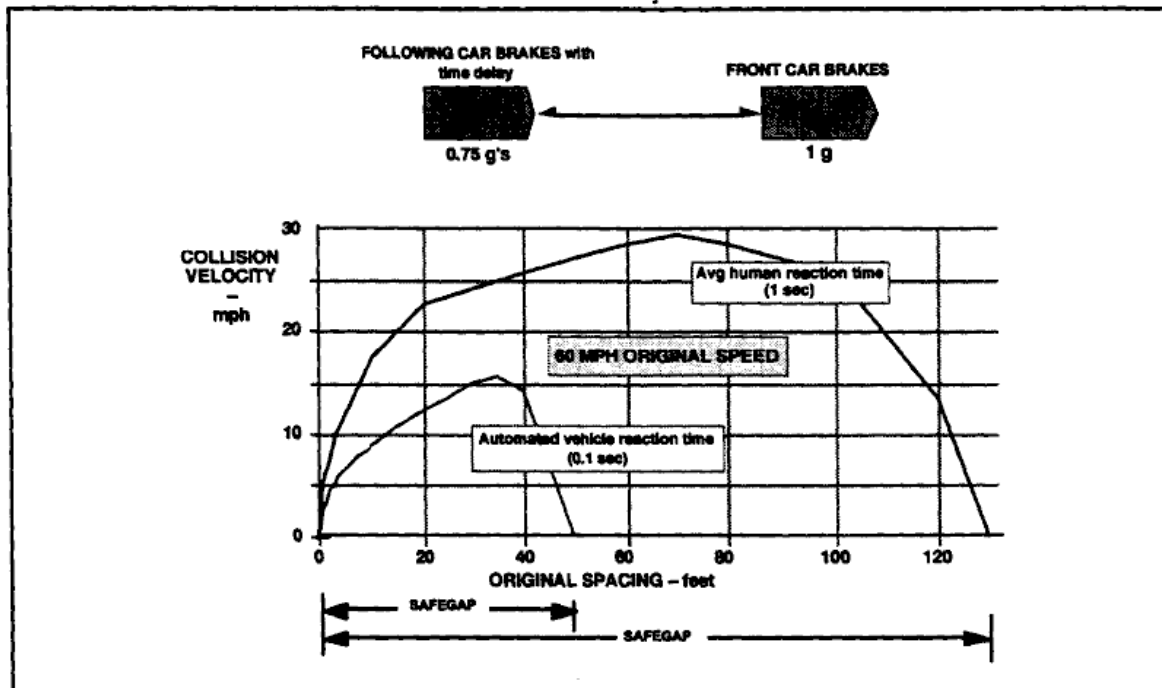


figure 5. Faster Reactions and Better Brakes Permit Closer Operation

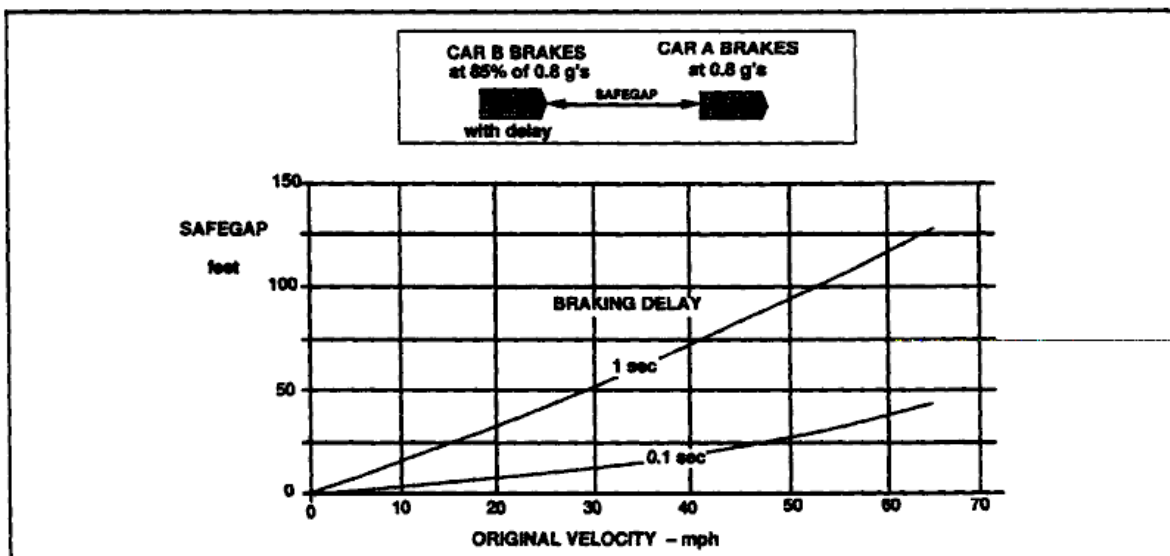


Figure 6. Safegap

Figure 7 shows how crucially important the braking capability of the following car is in setting the criteria for safegap. The time lag is held constant at 0.1 seconds, and following car's assumed braking is varied. The figure also shows the lane flow that corresponds to several different values of safegap, and thus the sensitivity of system productivity performance to the criteria finally chosen for determining safegap.

At this point we suspect that it will be necessary to measure actual braking capability in roughly real

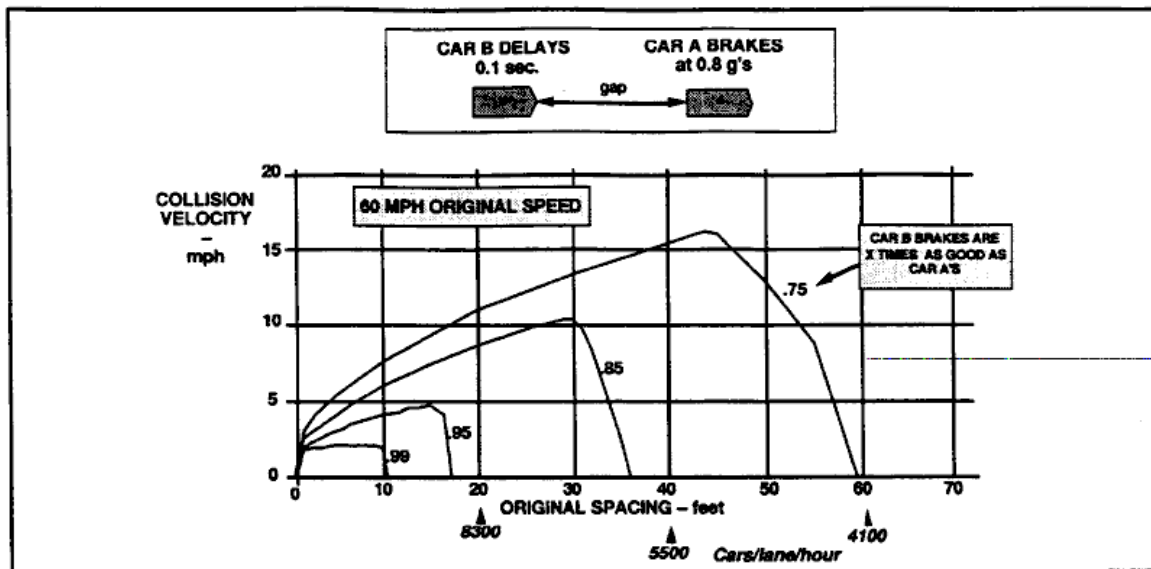


Figure 7. Machine Reaction Times Permit Closer Spacing - Brakes, Not Time, Become Dominant

time. This opinion is based on the fact that a deterioration in braking (wet surface, ice...) causes the safegap to increase substantially; even though both vehicles experience the same road surface. For example, if the braking criteria is based on the assumption that Car B's brakes are always eight-tenths as good as Car A's, and Car B brakes were actually capable of stopping at $0.8 g$ ($25.6 \text{ ft per sec}^2$), then a safegap of 39 feet would be deemed safe. But if actual braking capability of both cars were cut in half by, say, a rain-slick highway, then the required safegap rises to 69 feet. If the system did not revise the safegap algorithm as a function of actual capability, then this deterioration could permit an unsafe cruise condition.

Figure 8 shows the relationships among the three variables of lane flow, space between vehicles (gap), and speed. These are the same relationships that define the familiar highway capacity curves, except we have used average spacing between vehicles as an alternative to the more familiar variable of vehicle density.

Some orientation points are shown on Figure 8. The Highway Capacity Manual shows that the maximum lane flow of 2000 cars per hour occurs at about 30 mph; this corresponds to an average spacing between vehicles of a little over 60 ft. The line shown is the safety advice of the California Highway Patrol: maintain a minimum spacing of one car length for every 10 mph, i.e., leave 6 car lengths at 60 mph. Maximum flows of about 2700 cars per lane at 55 mph have been observed on California highways; such driver behavior is seen to be fairly consistent with the CHP's guidelines. The comparison between the driver behavior reflected in the Highway Capacity Manual and the more recent observations in California show there is a lot of flexibility and variability in peoples' driving habits, which presumably reflects their judgment as to what is safe.

Figure 9 is a repeat of the framework of Figure 8 that shows some illustrative operating points for several assumed levels of autobrake performance: these are the spacing and flows that would obtain at the various speeds if the vehicles cruised at the

minimum safe spacing (safegap) permitted by the design braking criterion. The upper shaded region is based on a braking criterion in which the lead vehicle suddenly brakes at 1 g (32 ft per sec^2) and the second vehicle brakes at 0.8 g's ($25.6 \text{ ft per sec}^2$). The upper edge of the region reflects a 0.2 sec delay and the lower a 0.4 sec delay. At 60 mph this shows the spacing is somewhere between 50 to 60 feet, and the corresponding flow between 4000 and nearly 5000 vehicles per lane per hour.

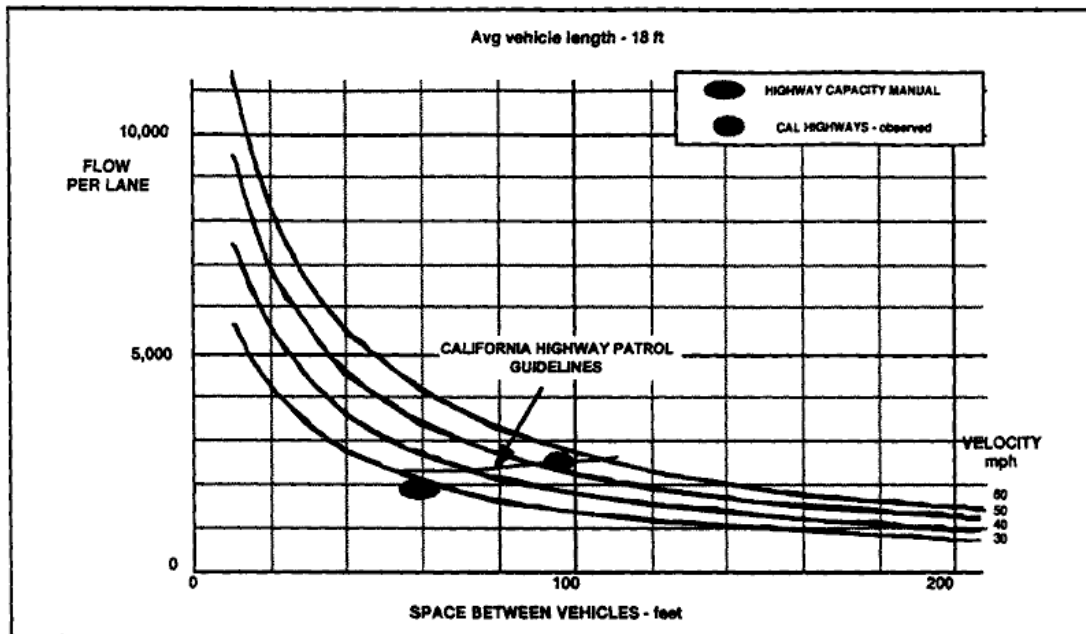


Figure 8. The Relationships Among the Gap Between Vehicles, Speed, and Flow




Figure 9. The Relationships Among the Gap Between Vehicles, Speed, and Flow

The lower shaded region reflects the autobrake behavior if the braking criterion is changed by reducing the braking of the second vehicle to 0.5 g's (16 ft per sec^2). It is obvious that the choice of braking criteria has substantial leverage on the maximum flow potential of Autobrake-equipped vehicles.

Based on these data, sensor(s) capable of measuring distance from about 10-15 feet out to 120-150 feet would be adequate for normal vehicle following. As will be shown, however, this is not adequate to handle the stopped vehicle situation.

Precision in measuring distance is not necessary; as a horseback guess, a one-sigma of 3-5 percent should be adequate.

Overtaking and Intrusion into the Lane Assume the situation in which the vehicle with Autobrake is closing on the vehicle ahead. If the driver fails to do so, the Autobrake system should begin braking so that when the rate of closure drops to zero, the distance separating the two vehicles is not less than the safegap corresponding to that speed.

Figure 4, shown earlier, illustrates that it takes time to slow the following vehicle to match the speed of the lead vehicle. The time required depends on the degree of braking. It is probably not desirable to wait until it requires maximum braking, so some lesser value more comfortable to passengers should be set into the Autobrake system as a default value.

The choice of this braking level for autobraking is a trade-off between passenger comfort and frequency of actuation: decreasing the braking severity shortens the timelag between when the driver should have braked and when Autobrake intercedes, so it brakes more often but more gently. (Note that it is not necessary that the driver maintain the human driver safegap; with any realistic braking criterion and the Autogap system operating, the vehicle is safe at the smaller gap. Figure 4 was drawn under the assumption that the human driver would not change his habits, and if alert would brake to maintain the gap to which he or she was accustomed.)

In the case of a vehicle intrusion from the next lane, a more dangerous situation could arise. There are several possibilities. First, the intruder may move in at less than safegap, but pulling away. Probably no response is appropriate here; just wait until the actual distance has widened to safegap. If the rate of opening is too slow or the distance too close (both would have to be defined, possibly in the language of fuzzy logic), then mild braking might be appropriate.

The second situation is more dangerous: the intruder is slower so the rate of closure is positive, and the distance is already near or less than safegap. In such a case, passenger comfort is sacrificed to more powerful braking, possibly the maximum the vehicle is capable of.

It is unfortunately possible that the intruding vehicle can be too close and too slow for collision to be avoided even with full emergency braking. In this case the only remaining option is a manual lane change to go around the object. It is not envisioned that Autobrake can cope with this situation, which is left entirely to the manual control by the driver.

Stopped Vehicle. The extreme of this situation is the stopped object. Some 70 percent of rear-enders are caused by running into a stationary vehicle, a testament to the frailty of the human attention span. Collision avoidance requires that the sensor detect the stopped vehicle in time to brake to a complete stop, or to provide time for the driver to move into another lane.

Figure 10 illustrate the distance required to brake to a complete stop as a function of one's own velocity and braking capability. While these calculations are simple, their interpretation to arrive at sensible sensor range requirements is not. And it may turn out that the maximum practical capability of

the sensor system will dictate constraints on the region of safe system operation. For example, if the maximum range measuring capability of the sensor is, say, 250 feet, then vehicle speed should be kept below about 55 mph if the road surface will only provide 0.4 g braking capability, but permit 75 mph speeds on an 0.8 g surface. (Here we have found a second reason to measure braking capability).




Figure 10. stopping Distance as a Function of Speed and Braking Capability

It should be kept in mind in selecting specific values for these various requirements that the option of changing lanes is often going to be available. The practicality of this maneuver depends both on the state of traffic and the time to alert the driver and have him or her carry out the maneuver. Figure 10 shows the total time available for this warning-and-lane-change maneuver for a few illustrative cases; we have no judgment as to how much time is enough.

The Pile Up.

Assume we are following a vehicle that also has Autobrake. The car in front of it stops suddenly. The Autobrake on the car we are following brakes almost immediately and therefore stops safely. Our own vehicle detects the change in range rate, and our braking begins almost simultaneously, and we, too, come to a safe stop.

This scenario can be more complex if we are following a manually controlled vehicle not equipped with Autobrake. Assume the vehicle he is following stops suddenly. There are several possible outcomes. First, if the driver of the vehicle in front of us is alert and immediately brakes, then we detect his braking and also brake. In this case the outcome is happy.

If, however, he is inattentive and fails to brake in time, he collides with the stopping (stopped?) vehicle. If our sensors had been able to see around him to the vehicle ahead, we could have started braking when we detected the initial braking, and avoided joining the collision. If, however, the vehicle directly in front of us shielded the one in front of it, we would not brake until the range rate of the front car changed, which might not be until he actually collided with the stopped vehicle. Now the car in front

of us decelerates at 5 or 10 or more g's - much faster than could be achieved through braking. If our safe following distance were based on assumptions about the front car's braking capability; say 1 g, then there is inadequate space to permit stopping. The result is a three-vehicle pile up.

This last scenario is very close to having a stopped vehicle dropped from the sky in our path. It poses a real dilemma. Of course, if all vehicles have Autobrake or if all drivers without it are alert, the problem goes away. But if the driver ahead of us is not alert, and does rear-end the car ahead, then we are following too close. At the moment we see four choices: (1) accept the risk - it won't happen often; (2) open up our normal following distance to be completely safe - this causes too large a penalty on flow capacity; and anyway cars will keep cutting in to fill the gap; (3) hope that our sensor can see two cars ahead often enough to maybe detect the original slowing; or (4) retrofit all vehicles with a small radio frequency emitter aimed backward that is actuated by hard braking to warn following cars (or some such scheme).

These various operational scenarios deserve considerably more analysis than is presented here, and their various kinematics are being analyzed in other AHS Precursor Analyses.

Sensing and Sensor Interpretation.

“Scene” Sensing... Sensor and sensor interpretation requirements are intertwined, and the nature of the sensor at least partially dictates the sensor interpretation task. Further, some of the interpretation task may go on inside the box labeled "Sensor" in one case, and in the box labeled "Sensor Interpretation" in another. For example, the typical radar will internally interpret its signals to calculate range, but for stereo ranging the calculations to convert angles into ranges occur outside the "sensors". In the discussion following we will not be concerned about this separation.

We concluded above that we will require continuous measurement of distance and rate of closure on the vehicle in our lane. From Figure 9 it appears that sensing over a range from around 10-15 feet out to 120-150 feet would be adequate for normal cruising, but that the stopped-vehicle threat pushes this range out to 300 feet, or even more. This later may turn out to be impractical of attainment, requiring constraints on operating regimes.

Precision is not required, accuracy of about 10 percent is probably adequate.

More analysis may show that continuous measurement of the distance to the nearest vehicles in adjacent lanes can reduce system time lags in reacting to an intrusion. Whether this is a desirable system function depends on a weighing of the kinematic gains against any additional system complexity. This increases the demands on the sensor and sensor interpretation subsystems, but probably only modestly.

In order to obtain these data, it is obviously necessary to distinguish between the vehicle in one's own lane from those in adjacent lanes, as well as quickly recognize an intrusion into one's lane; Figure 11 illustrates this operational situation. The interpretation of such situations will probably require recognition of lane markers. If tracking the lane markers is, in fact, necessary; it imposes an additional requirement on both the sensor and the sensor interpretation subsystems.

In addition to recognizing the operational situations and extracting the data discussed above, the sensor-sensor interpretation system must be able to differentiate between legitimate threats and false ones. There are two situations of concern. One is when the vehicle is turning, so the sensors sweep across objects or vehicles that will not remain in one's own path. The second is when a vehicle crosses

the path first enough to insure being clear before one's vehicle reaches the crossing point. Both of these situations suggest discrimination based on the angular velocity of the objects relative to one's own vehicle. So measuring, or being able to extract, relative angular velocity is added to the list of requirements. (This may help identifying an about-to-intrude vehicle.)

Last, there might be some advantage if the sensor and sensor interpretation system could crudely distinguish among different classes of vehicles: motorcycles or 18-wheelers. Because their braking capability varies, this knowledge would offer the option of tailoring the safe gap to the type of vehicle. This is not a requirement, perhaps only a mild desirement: we confess to being unsure of its usefulness.




Figure 11. Typical Scene Looking Forward

We describe the sensor-plus-interpretation process as real time"; actually we mean it will have to be very fast, with the optimum speed defined by the trade between the performance of the system and the cost of computing power. Here is a prime opportunity to use ingenuity in choice of techniques as a substitute for sheer computing power.

In addition to the requirement for speed is the requirement for validity: there is almost no latitude for misinterpretation of the scene. There will be little tolerance for calling for braking when it is not desired, and even less for failure to call for braking when it is desired. This requirement for an extremely low false alarm rate may have a substantial influence on the choice of sensors (more likely: sensor combinations).

*Speedometer...*The velocity on ones own vehicle is a key variable: this can be obtained from the vehicles speedometer system.

*Longitudinal Accelerometer...*It will be necessary to obtain at least a crude measure of current braking capability, which is a variable primarily because road conditions vary geographically and change with the weather. This braking capability might be deduced in roughly real time by correlating longitudinal deceleration with degree of braking every time the brakes are used. It may develop that a longitudinal acceleration deceleration feedback loop will improve internal system dynamic performance, strengthening the case for an accelerometer.

There are two other possibilities that might be considered if the accelerometer approach fails. First, this information could possibly be supplied by the infrastructure, in which case communication to send and receive it will be needed. Second, the possibility of rain or snow sensors be added to the vehicle, and the braking automatically be degraded when they occur.

External Information & Command Inputs- The Need for Communications.

We see Autobrake as completely autonomous, with no reason for communication with the infrastructure.

Vehicle Control Actuation.

Obviously, the autobrake function requires brake control only. The need is for reliable, fail-safe, fast, and accurate control; the technology to provide it is well in hand.

System Mechanization.

Figure 12 depicts the basic mechanization of the Autobrake System, which infers the high-level architecture of the system.

While the figure is generally self-explanatory one point should be noted: given that control by wire for the brake system is available, it is feasible to also have the driver's input be an electrical signal, and also operate through this channel. We have elected, however, to continue to show the driver exerting manual control. Our reasoning is that if Autobrake is sold alone as an early product, it is unlikely that there would be sufficient faith in electronic controls at that time to warrant abandonment of the tried-and-true. We do show the other alternative in Figure 16, which depicts Autobrake plus Autogap. This is, of course, arbitrary: the system would function just as well if manual control continued to be a separate system.

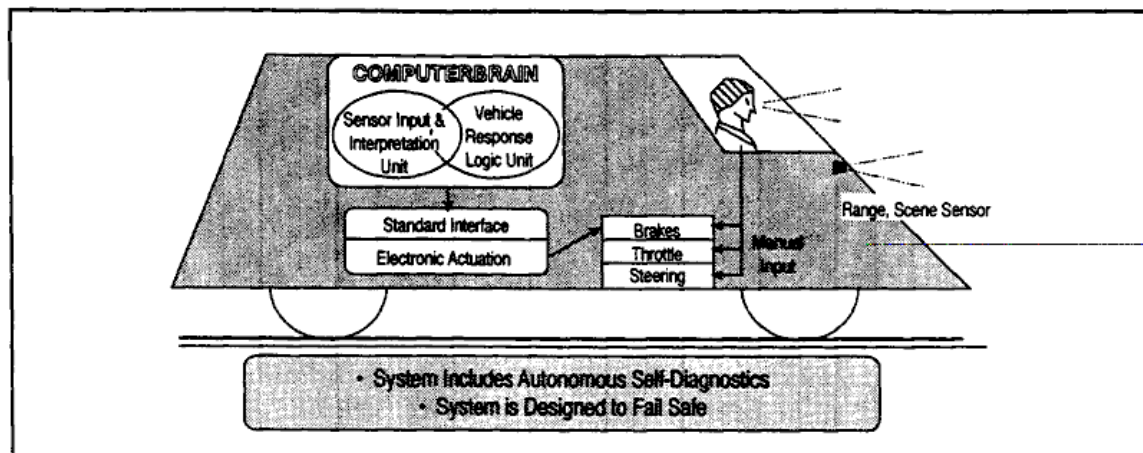


Figure 12. Autobrake mechanization

As already noted, the major points made here are summarized in Figure 13. We now turn to a discussion of the potential benefits of the system that might be a motivation to purchase.

POTENTIAL BENEFITS - MOTIVATION TO PURCHASE.

Safety.

The only motivation for the purchase of an Autobrake system is improved safety: Twenty-three percent of all freeway accidents are rear-enders, and the Minnesota DOT reports that half the urban

freeway accidents in that state are rear-enders. In common with airbags and Anti skid braking systems,

TECHNICAL FUNCTIONS	AUTOBRAKE SYSTEM
DESIRED VEHICLE RESPONSES - INFORMATION NEEDS	<ul style="list-style-type: none"> • DISTANCE TO AND RELATIVE VELOCITY OF NEAREST VEHICLE-SIZE OBJECTS IDENTIFIED BY THE LANE THEY OCCUPY. • OWN VEHICLE VELOCITY. • CRUDE MEASURE OF REAL-TIME BRAKING CAPABILITY. • BASIS TO DISCRIMINATE NON-THREATENING OBJECTS.
SENSING AND SENSOR INTERPRETATION	<ul style="list-style-type: none"> • FIELD-OF-VIEW ROUGHLY FORWARD QUADRANT. "SEES" OBJECTS AND LANE BOUNDARIES. • INCLUDES FUSION OF MULTIPLE SENSORS.. • INTERPRETS SENSOR SIGNALS TO PROVIDE VARIABLES DESCRIBED. • ESSENTIALLY REAL-TIME PROCESSING. VERY LOW FALSE ALARM AND FAILURE-RATE. • SPEEDOMETER & LONGITUDINAL ACCELEROMETER.
EXTERNAL INFORMATION, COMMAND INPUTS - COMMUNICATION	NOT REQUIRED
VEHICLE CONTROL ACTUATION	ELECTRONIC BRAKING NEEDED. (MANUAL BRAKING BY DRIVER MAY BE THROUGH CONVENTIONAL MECHANICAL INPUT, OR THROUGH ELECTRICAL SIGNAL)

Figure 13. Autobrake mechanization: Functional description of Technical Elements Required

Autobrake serves no function during normal driving, and acts only in emergency or unusual driving situations. Better safety offers rewards through improved peace of mind, and it may translate into monetary savings as well through reduced insurance rates.

Such a system has two attributes that enable it to improve safety beyond the capabilities of human drivers. First, its attention doesn't lapse. Some 70 percent of rear-enders on freeways are driving into stopped vehicles, most of which are the result of driver inattention. For the automatic system, almost the only 'inattention' is caused by system failure, which at worst will be far less frequent than human lapse. Further, the driver would be informed of a nonfunctional system. It should be possible to reduce the risk of dangerous failure to essentially zero. (But we can't eliminate the accident caused by a stopped vehicle that is hidden by an intervening truck or until a hill is crested; if there is inadequate braking space the only other option is a manual lane change.)

Second, as already noted, an automatic system's reactions will be very fast, probably on the order of five to ten times faster than the average human driver. And because it is measuring the distance and rate of closure on the lead vehicle with much better accuracy than most drivers - and probably with better built-in "judgment" as to stopping distance requirements - it can safely tolerate closer vehicle proximities than are safe for human drivers. Thus it can still prevent collisions in situations that are already past the point of collision-avoidance by the driver. For example, at 60 mph the roughly 1 second in reaction time saved by automatic controls implies 88 feet of safe space is gained

It is probably less important, but worth noting that the system requires no learning on the part of the human. In the vast majority of situations the human driver will initiate braking well before the Autobrake system would react. The Autobrake system would therefore require no change in normal driver habits, and for the driver who was both careful and lucky he might never know it was there.

On the other side of the coin, the system will not protect against all accidents. The human will remain much better at recognizing and assessing unusual driving situations, particularly those involving lateral motions. While one would have to be careful in claiming superiority for our eyes over artificial

sensors in all dimensions of performance -estimating distance, for example - there is no question that our brains are fir better at image interpretation than any automatic system we can now conceive. The automated system will not recognize and react properly to situations it is not capable of "seeing" or programmed to interpret.

Usefulness of the Freeway

We see no reason that the Autobrake function need be restricted to freeway use only. As long as the system can properly discriminate threats from non-threats during turns, the system could be activated at all times.

DEVELOPMENTAL ISSUES AND RISKS

We view the development of Autobrake alone as less probable than that of a joint development of the Autobrake and Autogap functions - Intelligent Cruise Control - so we delay discussion of the issue until that section of the report.

Robustness and Failsafe

An automatic self-test and self-diagnostic system would be an essential part of any system design.

Robustness and failsoftness will require redundancy, particularly in the sensing system. Redundant sensors also afford the opportunity for a more capable overall system. For example, a system using both forward looking radar and passive, vision-based sensors can detect more features about the operating environment than could either type of sensor alone: they are complementary in their detection and measurement techniques.

We believe that the testing required to insure proper functioning of the Integrity Verification Subsystem - the self-test and self-diagnostic system -will fir exceed that required to verify performance. This aspect of system design will be a major item of development cost.

Sensor Development.

This is a demanding area of development, largely because of the cost and the reliability issues; the performance required is well within the state-of-the-art, but comes today at too high a cost for this application.

There is one performance issue. All of the objects that it would be desirable to avoid are not vehicles; in the future we may be concerned with neighborhood driving ,where such things as children and pets are serious problems. Today we depend entirely on driver reaction for their safety, and clearly it would be desirable if Autobrake could supplement that capability just as it is designed to do for vehicles.

The issue is whether such a capability should be required. It is a substantial technical complication, because driving up the capability to detect such threats also drives up the propensity for false alarms. The issue deserves serious attention, but our intuition is to not impose this requirement for the initial systems, and depend on the improving state-of-the-art over time to bring it about later.

Sensor Data Processing - Sensor Interpretation.

In our judgment, this is the longest pole in the technical development tent.

Sensors only convert the pattern of electromagnetic radiation from the operational scene into electrical signals. These signals then have to be interpreted: the features of the scene have to be extracted and identified with real-world objects. The data from multiple sensors have to be combined and compared ("fused"). Critical measurements have to be derived. All of this must be done essentially instantaneously and continuously. There is very little tolerance for errors: braking when it is unnecessary or not braking when it is are both very undesirable events.

While the situation being interpreted is more constrained than that encountered in the military target recognition and identification problem, the requirements for validity are probably more stringent. Overall, this sensor interpretation problem is far from simple, and as it is currently done, requires a great deal of computing power. There is little doubt that it can be done adequately for this application, but it will take substantial effort. It is at the edge of the state-of-the-art.

Potential for Retrofit.

Without more careful study, it would appear that the key questions are finding a reasonable location for the sensors, and providing electronically controlled braking. Clever designers may find ways.

Human Factor Issues.

Being essentially invisible to the driver, the only issue is the design of the displays of system status

Costs.

We have little basis at this time for projecting probable costs of either development or production. We conjecture, however, that the three major items of development cost will be

- (1) The devising the data processing schema for sensor interpretation,
- (2) The design and testing of the Integrity Verification System - self-test and self-diagnostics, and
- (3) Performance and reliability testing.

The generic cost issue is discussed later under A DIGRESSION: DEPLOYMENT AND THE DYNAMICS OF COST BEHAVIOR

2. THE AUTOMATIC GAP HOLDING FUNCTION (AUTOGAP) - INTELLIGENT CRUISE CONTROL

Autogap combines naturally with Autobrake, using the same sensors and its drive-by-wire braking capability, but for fully modulated control, not just responses to special situations. It adds throttle control, so that we now have full automation of the vehicle drive-train.

As already noted, we envision that Autobrake will operate at an times, even when Autogap is Off.

OPERATIONAL CONCEPT.

It seems reasonable to assume that Autogap will be combined with the Cruise Control system. The new Cruise Control - now the *intelligent* Cruise Control system - functions as an ordinary cruise control system when the leading vehicle is further than some transition distance (probably around 150 feet at 60 mph), but reverts to gap control when the lead vehicle is inside of this distance. The driver's hand control adjusts speed when nothing is closer than the transition distance, but begins to adjust the following distance to any vehicle that is nearer. He can manually close the distance all the way down to the safegap corresponding to automated operation (the Autobrake safegap). The geometry of operation is illustrated in Figure 14.

Autogap has three primary goals: to further increase safety by preventing Cruise Control from overrunning the leading vehicle, to reduce driver work load on crowded freeways, and to offer the prospect of an increase in effective lane capacity

TECHNICAL DESCRIPTION.

The technical features for Autogap are summarized in Figure 15, which also repeats the main points associated with Autobrake. These Autogap features are discussed below, using the format introduced earlier.

Desired Vehicle Response - Information Needs

Outside the transition distance illustrated in Figure 14, the system operates to hold constant velocity; inside the transition distance it operates to hold zero rate of closure except when the driver is adjusting the following distance. If the driver wants to change the gap, his control action signals for some passenger-comfortable rate of deceleration or acceleration that overrides the zero rate of closure command until the control is released.

The bare-bones logic does not sound complex, but there are many aspects that will require more careful thinking, and ultimately extensive testing. By and large, the automated system will be required to behave much as a careful and defensive driver would. There are many different circumstances that we encounter in manual driving that the prudent driver almost instinctively reacts to in order to lessen the chances of unpleasant surprises: we would like to build into our automated vehicle response logic the same prudence.

For, example, we frequently observe vehicles using their turn signals to request entry into a crowded freeway lane, and, perhaps surprisingly, we see vehicles responding by slowing down to permit entry. To emulate this behavior automatically will require sensors that can read turn signals.

Sensing and Sensor Interpretation.

If vision-based sensors capable of reading turn signals are used for Autobrake, then the Autobrake sensing and sensor interpretation system will also support Autogap. The interpretation system will, of course, require some adaptation.

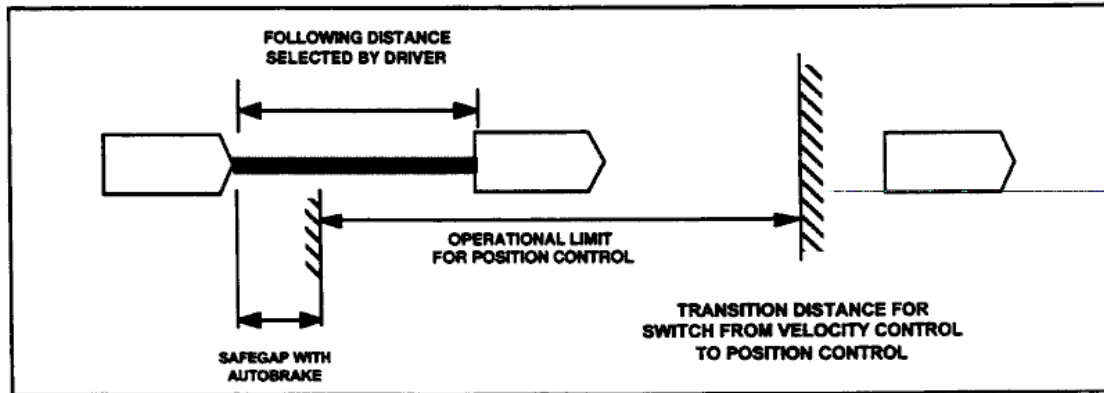


Figure 14. Automatic Gap Control

TECHNICAL FUNCTIONS	AUTOBRAKE	AUTOGAP
DESIRED VEHICLE RESPONSES - INFORMATION NEEDS	DISTANCE, RATE OF CLOSURE OF FORWARD QUADRANT VEHICLE-SIZE OBJECTS IN OWN & ADJACENT LANES DISCRIMINATION OF NON-THREAT OBJECTS. OWN VELOCITY. REAL-TIME BRAKING CAPABILITY	ESSENTIALLY THE SAME PROBABLY NEED TO DETECT TURN SIGNALS OF NEAREST VEHICLES.
SENSING AND SENSOR INTERPRETATION	FIELD OF VIEW ROUGHLY FORWARD QUADRANT. • ESSENTIALLY REAL-TIME PROCESSING. VERY LOW FALSE ALARM AND SPEEDOMETER &	ESSENTIALLY THE SAME FAILURE-TO-ACT RATES. LONGITUDINAL ACCELEROMETER
EXTERNAL INFORMATION, COMMAND INPUTS - COMMUNICATION	NOT REQUIRED	NOT REQ'D. IF COMMUNICATION NEEDED, WILL BE TMS-TO-VEHICLE
VEHICLE CONTROL ACTUATION	ELECTRONIC BRAKING	ELECTRONIC BRAKES AND THROTTLE

Figure 15. Autobrake and Autogap Mechanization's: Functional Description of Technical Elements Required

It is more important for Autogap than Autobrake for the system to be able to continuously track vehicles in adjacent lanes. If not included in Autobrake, it should probably be added.

External Information & command Inputs - The Need for Communications.

It is our opinion that since a vehicle with ICC is still a driver controlled vehicle, there is no need to transmit traffic commands directly to the automated system, and that the currently used visual signals to the driver are adequate. This would leave ICC as a completely autonomous function, not requiring communication with the infrastructure.

This is an arguable issue.

Vehicle Control Actuation.

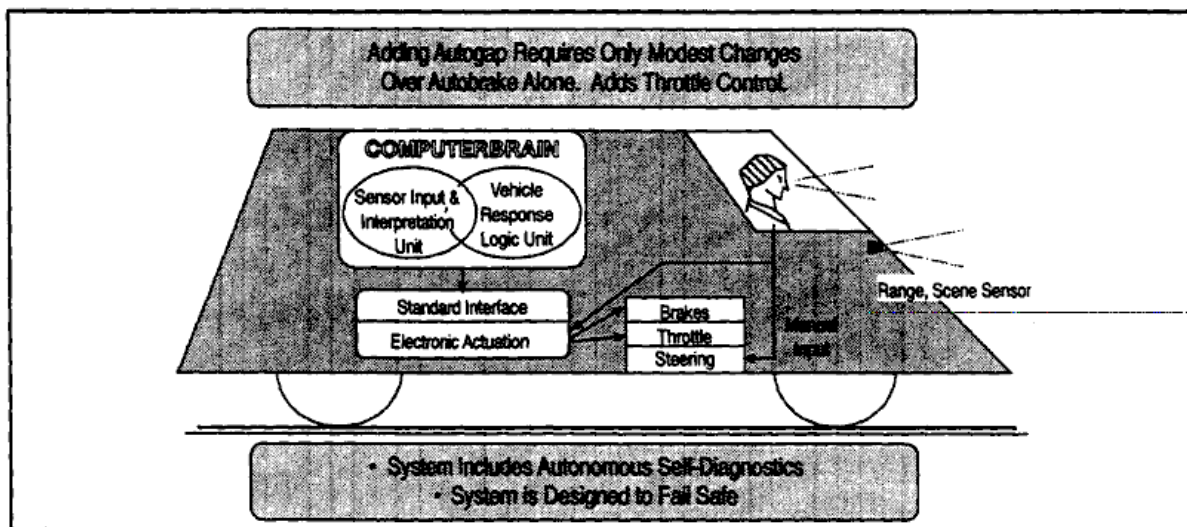
The throttle control probably needs to be more sensitively modulated than some cruise control systems currently in use.

It would be sensible to initially design the automated brake system of Autobrake for the level of modulation needed for Autogap (and superior Cruise Control). This will clearly happen if; as we conjecture, Autobrake is only offered for sale when incorporated into Autogap.

System Mechanization.

Figure 16 depicts the mechanization for Intelligent Cruise Control. Comparing it to Figure 12, it can be seen that the changes at the top level shown in the figures are slight. Electronic throttle control has been added, and we now show the driver's control being effected through the electronic system rather than through mechanical linkage.

We now turn to a discussion of the potential benefits of the system that might be a motivation to purchase.



**Figure 16. Autobrake and autogap Mechanization
Intelligent Cruise Control**

POTENTIAL BENEFITS - MOTIVATION TO PURCHASE

While Autobrake offered only 'invisible' safety, the combination of Autobrake and Autogap - Intelligent Cruise Control offers both invisible safety and the obvious safety of not letting Cruise Control, in a moment of inattention, drive one into the vehicle ahead.

Driver Convenience.

The Autogap addition to Cruise Control also offers relief to the driver from the constant attention and drive-train control actions required to maintain one's position in a string of traffic. This speed-up-slow-down is particularly irksome on a crowded freeway operating on the unstable side of the highway

capacity curves - an all too frequent event. We conjecture that this aspect of ICC may be the most powerful enticement for purchase of the system.

Capacity Improvement.

While Autogap (because of the Autobrake feature) offers the driver the option of decreasing the following distance down to the safegap corresponding to automatic braking, its questionable whether many drivers will exercise this option because such close following distances will probably be uncomfortable, at least initially. It is therefore dubious safegap will contribute to an appreciable increase in effective lane capacity when first introduced.

But we notice that many drivers seem to not be disturbed by driving very close to the car in front (a number of these frequently follow the author). There is, in fact, an advantage to close following: shortening the gap makes it more difficult for other cars to cut in front of one. This is motivation for many to drive closer than they would otherwise chose. It may be that in time more and more drivers will opt for closer following distances, particularly as the freeway becomes more crowded and as they gain confidence in the automated system. Figure 17 illustrates the potential impact on ~~lane~~ ^{flow} as a function of the number of people who chose to operate at Autobrake-Safegap instead of Manual-Safegap. The values chosen are only intended to be illustrative of possibilities, and tend to represent extremes in driving distance choices.

Usefulness of the Freeway

We see no reason that the Autogap function need be restricted to freeway use only. Operation on interstates and any relatively straight rural highways does not appear to pose particular problems.




Figure. 17. Potential capacity Benefits from Gap Control

There is a problem with going around corners. If the system is to be used on surface streets, it will be necessary to provide logic to prevent unwanted behavior, either by adapting the system to the maneuver in some way, or deactivating it. This point is noted as an issue in the next second.

DEVELOPMENTAL ISSUES AND RISKS

Prospects for Development

It is generally assumed that the final products that will implement AHS will be supplied by the private sector. The issue addressed here is whether the private investment can be attracted for their development.

The calculus for private sector funding is conceptually simple: there must be a reasonable prospect of future profitable return. Investment is inhibited if risk is high, or payoff a long time in coming. The cost of capital influences the acceptable time horizon.

We do not believe that marketing risk is high. If products along the lines of those described here can be offered for sale at reasonable prices, then the benefits just outlined should be sufficient to motivate purchase.

There are, however, two areas of risks that could inhibit a successful program. The first is product liability risk: there is a very real prospect that fault for many future rear-enders will be shifted by sympathetic jurors from drivers to the manufacturers of the ICC system.

The second is the uncertainty in predicting the cost of development and validation. Costs are very much a function of the technical difficulty of the tasks being undertaken. We view the technical problem of devising a sensor and sensor interpretation system that can provide the performance desired, keep false alarms to an acceptably low level, and incorporate the self-verification and failsafe features required as one that will require considerable innovation and ingenuity. And it must still result in a product at a cost that will sell. Not only are these development costs hard to predict with confidence, but

in this case the difficulty is compounded by the prospect waiting in the wings of unanticipated regulatory mandates.

The public sector can play an important role in reducing these risks, and it may require the joint development now being started to produce a successful AHS. One might suggest that this should be the fundamental public sector strategy: to reduce the risks and lower the obstacles to the point that private funds could be reasonably expected to carry the program.

Robustness and Failsafety

As discussed under Autobrake, an automatic self-test and self-diagnostic system would be an essential part of any system design. Autogap does *not* as far as we can see, raise any new issues.

Operation on Surface Streets.

As already noted, the sharp turns associated with surface streets may pose problems we have not adequately thought through. They will certainly complicate the vehicle response logic, and may create requirements for additional information from the sensor/sensor-interpretation subsystem.

Potential for Retrofit.

As the system gets more comprehensive, the complexity of retrofit goes up. But we will have to be further along in design to have a basis for judgment

Human Factor Issues.

The issue of operation in other venues than *freeways* rises a human factors issue. If the system is versatile enough to be left on most of the time, it is important that the driver is adequately warned to turn it off in situations where it might give unwanted responses. This implies either that the driver must remember to turn it *off*, or the system is able to recognize such situations, neutralize itself, and warn the driver that action has been taken.

Costs.

The observations about Autobrake generally apply here; we know too little to estimate costs. If Autobrake has already been developed, the incremental cost of adding Autogap should be small. The total costs should be minimized, however, by concurrent development.

3. THE AUTOMATIC LANE HOLDING FUNCTION (AUTOLANE) AUTOMATIC CRUISE CONTROL

OPERATIONAL CONCEPT.

The purpose of the automatic lane hold function (Autolane) is to maintain the vehicle in the center of the lane once the vehicle has been driven manually into the lane and the Autolane system actuated. It continues to hold lane until the driver disengages the system.

Disengagement will occur at the volition of the driver. It will be necessary to configure the system

to prevent accidental disengagement that could go unnoticed, but at the same time not preclude very quick resumption of manual control if some emergency warrants it. Some possibilities for meeting these criteria are described under the Human Factors section of Issues and Risks.

We consider it unlikely that Autolane will be offered for sale alone (although it could well be preceded by an Automatic lane Deviation Warning system that does not INCLUDE automatic steering). The reason for this conjecture is our difficulty in seeing the benefit to a driver just because he doesn't have to steer, but does have to maintain sufficient vigilance to continue operating the brake and throttle.

It seems much more likely that Autolane will be offered for sale only as a part of the complete package of Autobrake plus Autogap plus Autolane, a combination that could appropriately be called Automatic Cruise Control (ACC). With ACC the driver can manually drive to the desired lane for cruising, and turn *complete control of the vehicle over to the automatic system.*

Thus the primary goal of Autolane is to offer the potential for fully automated cruise. It will also improve safety by preventing vehicles from inadvertently leaving their lane. As already noted, submode of Autolane could easily be Lane Departure Warning.

TECHNICAL DESCRIPTION.

The technical requirements hypothesized here are summarized in Figure 18 and discussed in the following.

Desired Vehicle Response - Information Needs

The vehicle steers to stay on the centerline of the lane.

The information required for Autolane is the position of the vehicle relative to the center of the lane. It may be desirable to also measure the attitude of the vehicle relative to the centerline.

Sensing and Sensor Interpretation.

We have conjectured that one of the requirements imposed on the Autobrake-Autogap sensors is to establish which vehicles are in which lane, the implication of which is that they track the lane boundaries. Whether this information will also fulfill the needs of the Autolane system is yet to be determined. If not, then specialized sensors will be required for detecting the lane boundary lines.

Other techniques are available. The PATH program at UC-Berkeley tracks magnetic nails" buried along the centerline of the road. Visible or RF reflectors have been suggested. There may be others.

Sensor interpretation is a less demanding problem than for Autobrake and Autogap, although it is not trivial. Almost surely an operational system will not depend on one measurement technique alone, so data fusion and reconciliation will be a requirement.

As with Autobrake and Autogap, processing must be very fast with a very low probability of persistent or bias errors.

At least some automatic steering schemes use a knowledge of road geometry beyond that derived from the on-board sensors. There are a variety of ways to acquire that information, but we make no

attempt to cover them here.

The output of the sensor and interpretation system will be an error signal telling how far the vehicle is from the centerline, plus the rate of change of that signal.

TECHNICAL FUNCTIONS	AUTO LANEHOLD AND AUTOMATED CRUISE CONTROL
DESIRED VEHICLE REPONSES - INFORMATION NEEDS	STEER TO CENTER OF LANE NEED TO KNOW POSITION OF VEHICLE RELATIVE TO LANE, MAYBE RATE OF DEVIATION AND ATTITUDE
SENSING AND SENSOR INTERPRET ATION	REAL-TIME SENSING OF LANE POSITION. PROBABLY NEED MULTIPLE SENSORS FOR FAILSAFETY & ROBUSTNESS, SO NEED SENSOR FUSION. (MAYBE ALREADY AVAILABLE FROM AUTOBRAKE, AUTOGAP SENSORS.)
EXTERNAL INFORMATION COMMAND INPUTS - COMMUNICATION	AUTO CRUISE CONTROL NEEDS TRAFFIC COMMANDS DIRECTLY TO VEHICLE PROBABLY NEED VEHICLE-TO-TMS FOR EMERGENCY REPORTING.
VEHICLE CONTROL ACTUATION	ADDS ELECTRONIC STEERING (MANUAL CONTROL EITHER MECHANICAL OR ELECTRONIC.)

Figure 18. Auto Lanehold Mechanization: Functional description of Technical Elements Required

*External Information ~ Command Inputs - The Need
for Communications.*

All traffic commands will go directly to the vehicle, introducing the almost unequivocal need for communication from the Traffic Management System (TMS) to the vehicle. The "deadman problem" may also require signals from the TMS; this is discussed later.

Additionally, there will probably be a need for malfunction alerts from the vehicle, so two way communication becomes a highly probable requirement.

Vehicle Control Actuation.

Autolane brings in the third element of vehicle control, and introduces the need for electronically actuated steering.

There is a serendipitous byproduct to a drive-by-wire capability now that we have electronic steering in all three vehicle controls. Now it becomes technically feasible to begin considering new or substantially modified approaches to manual control. We are still using today techniques of control that were derived from the original need to use human muscle to steer, brake, and control the throttle. With that constraint gone, the manual controls could be designed differently if there appeared to be advantages in doing so, and if failsafety concerns could be dealt with.

Infrastructure Modifications

The point has already been made that fiducial marks of some kind will be necessary to permit orientation of the vehicle relative to the lane. It is believed that these will be simple enough that they can be installed relatively inexpensively and with little disruption.

System Mechanization.

The system mechanization is shown in Figure 19. It can be seen that the basic architecture is unchanged, just incremental modifications. Given standardized interface specs, there is substantial freedom for technical innovation in the individual subfunctions.

POTENTIAL BENEFITS - MOTIVATION TO PURCHASE.

Driver Convenience.

We conjecture that Automatic Cruise Control would be a boon to anyone whose commute involves even a moderately long freeway segment, or anyone who does much intercity traveling. ACC could take much of the strain and boredom out of long trips.

While the primary motivation for purchase might be driver relief, ACC should also offer another increment in improved safety: the essential elimination of wandering out of the lane or off the road for any travel on equipped highways. The sleepy driver threat is essentially nullified.

Usefulness of the Freeway

We have already alluded to the fact that ACC should be usable on any freeway or highway that has the lane markings and other modifications required to accommodate Autolane. But it is improbable that ACC could be used on surface streets without integration with the Traffic Management System.

DEVELOPMENTAL ISSUES AND RISKS.

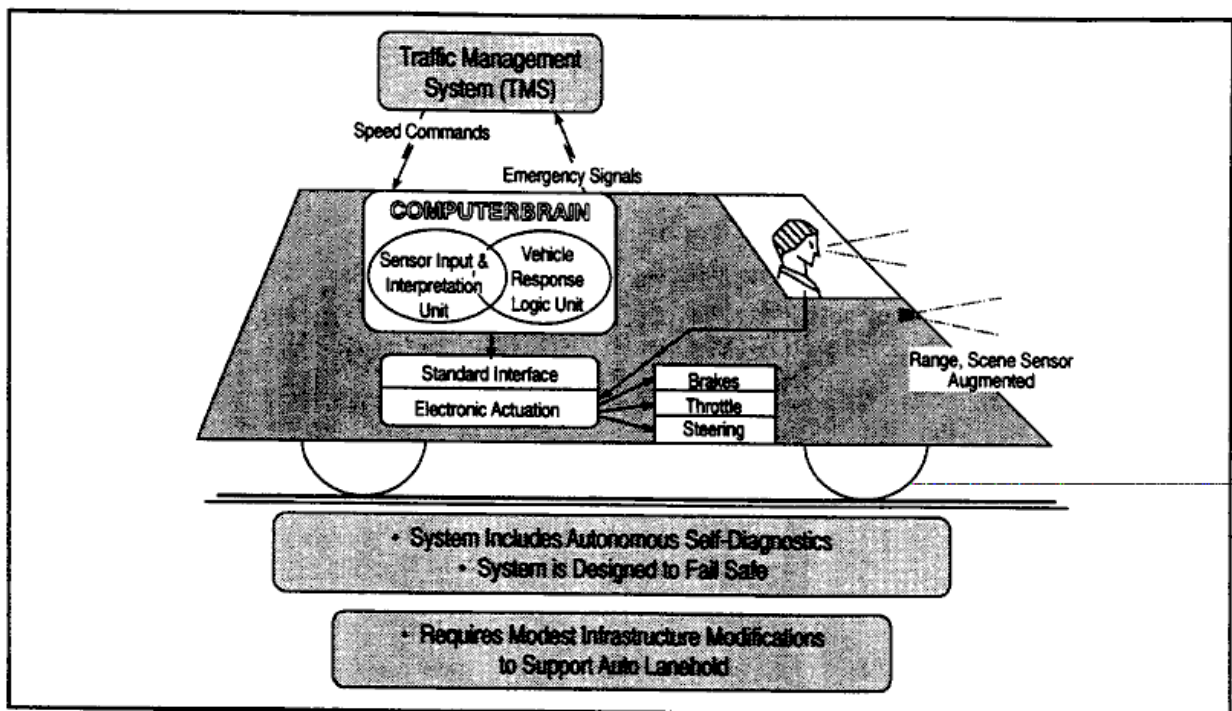
Prospects for development.

The prospects are quite good if the ICC appears to be fulfilling its promise. The liability hurdle has been passed with the introduction of Intelligent Cruise Control, and the most difficult part of the sensor interpretation problem apparently solved.

Robustness and Failsafety

As already noted, an automatic self-test and selfdiagnostic system would be an essential part of any system design. The addition of Autolane and its integration into a comprehensive Automatic Cruise Control system does reemphasize the need for a carefully designed and comprehensively tested Integrity Verification Subsystem.

The failsafety-failsoftly problem is now more complex than with only Intelligent Cruise Control, and will require even more careful consideration; this is discussed under Human Factors, following.



**Figure 19. adding Auto Lanehold: Automated Cruise Control
Hands-Off Freeway Cruising**
THE MARK I AHS

Vehicle Behavior Algorithms.

It is our belief that there will be no important interactions between the steering mode and the drive-train control mode, so that the addition of the lateral control channel will not dictate the need for change in the Intelligent Cruise Control algorithms.

Other System Areas.

Electronically actuated steering is already in extensive test, and its addition should not pose unusual problems. More sensors are required, but they are well within the state-of-practice.

Potential for Retrofit.

The prospects look dimmer and dimmer.

Human Factor Issues.

ACC does introduce new human factors considerations that go beyond the proper design of controls and displays. With only Intelligent Cruise Control, the driver was still required to watch the road and the general driving situation in order to steer; ICC merely relieved him from the frequent braking and acceleration otherwise required. While this may have permitted some relaxation in vigilance, the driver was still in a position to take immediate manual control of the vehicle in the event of an emergency

Automatic Cruise Control removes the motivation - if not the need - for driver vigilance. The implication is that the driver may not only take longer to react to an emergency or an unusual situation, but may also be in a mental state that is not conducive to an immediately rational response. This problem will require attention.

It has already been noted that it will be necessary to configure the system to prevent accidental disengagement that could go unnoticed, but at the same time not preclude very quick resumption of manual control if some emergency warrants it. The temptation to have the driver reassume steering control by just turning the steering wheel slightly is precluded by the first constraint, yet substituting too complex a procedure violates the second.

Perhaps something along the following lines. Small manual inputs to the steering wheel do not disengage the system unless the driver has already set the "Off" switch. But a vigorous steering wheel input immediately gives control to the driver. Obviously it would take considerable testing to validate this or any other approach, and obviously there are many variants of such schemes that will require evaluation.

A DIGRESSION: DEPLOYMENT AND THE DYNAMICS OF COST BEHAVIOR

Thus far we have hypothesized two major deployment steps. The first was the introduction of Intelligent Cruise Control (ICC), which included a full-time Autobrake system and an integration of conventional cruise control with the new Autogap function.

The second step was the addition of the automatic lane holding function, Autolane, producing a system capable of fully automated, hands-off cruise on any (freeway) highway, or road modified to support Autolane. We called this system Automatic Cruise Control (ACC). It is, in fact, the Mark I version of the Automated Highway System.

The rate of deployment of these systems will hinge very importantly on their costs, but development is not far enough along to be able to estimate these costs with any confidence. Neither do we have any useful intuition about the demand curve - the relationship of number of sales to prices. But it is possible to illustrate at least one aspect of the problem: the change in costs as deployment progresses.

The central player is the so-called "experience curve". The experience curve describes the predictable reduction in production costs as a function of the number of units produced. It is a phenomenon that has been observed in many industries, and has been well established by empirical data. It is a consequence of both learning to make the production process more efficient, and incremental improvements in technology and design to make it easier.

These data reflect the fact that unit costs typically vary as an exponential of the total number of units built. Specifically, it implies that production costs reduce by some fixed percentage every time the total number of units built is doubled. As an example, a 90% curve implies that the second unit produced would cost 10% less than the first, the 20th would cost 10% less than the 10th, and the 200,000th would cost 10% less than the 100,000th.

The experience curve is central to the pricing strategy of many companies, and it is reasonable to think that it will be extremely important to the future of the AHS, as the following illustrates.

Figure 20 is a purely conjectural assumption about the rate of diffusion of ICC and ACC as a function of time: it shows the hypothesized cumulative number of equipped vehicles over the first ten years after initial introduction. In drawing these curves we have assumed that once the fully automated

capability of ACC is available here shown five years after the introduction of ICC, ACC will begin to almost fully displace ICC.

These curves are in no sense a forecast, but a mechanism to illustrate the behavior of costs over time. within broad limits, this behavior is relatively insensitive to the particular values shown: we could double them or halve them without affecting the basic phenomena we wish to illustrate.

The top of Figure 21 repeats the cumulative sales chart of Figure 20. The plots at the bottom show the production costs of one unit at the beginning of each year of production (we assume production equals sales). In constructing Figure 19 we have arbitrarily assumed that the first production model of an ICC will cost \$10,000, and that unit costs thereafter will decrease along a 90% learning curve.

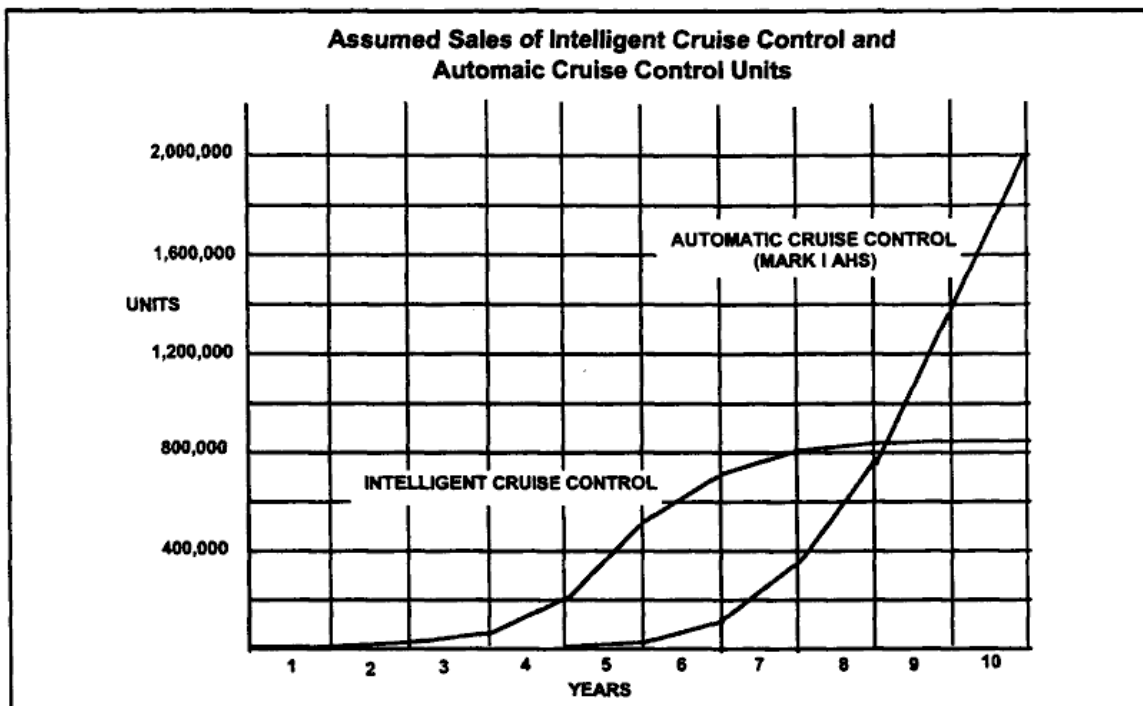


figure 20. An Illustration of the dynamics of Cost Behavior -




Figure 21. An Illustration of the dynamics of Cost Behavior - II

Since ACC builds on ICC, we have also assumed that the initial production model of ACC - here shown five years later - will be less expensive at \$7500. We have used these assumptions to calculate how costs will vary over time as deployment proceeds along the paths assumed in Figure 20.

As can be seen, costs drop very rapidly in the first year, but as more units are produced, the doubling time increases so costs come down more slowly. But they come down significantly: over the total of 872,000 ICC units produced the unit production costs drop from \$10,000 to \$1,250 for the 90% experience curve assumed. Over the 2,000,000 ACC units assumed, the unit production costs drop from \$7,500 to \$827.

For pricing purposes, the cumulative average costs of production may be more relevant than the unit by unit costs. In Figure 22 we show the average unit production costs as a function of the number of units produced.

Figure 22 also illustrates a second, and more important, point: the extreme sensitivity of costs to the slope of the experience curve. For the largely solid state technology that we anticipate for AVCS and ACC, slopes of 85 percent or lower are a reasonable expectation.

We will go through a hypothetical - and simple minded - pricing exercise for ICC. Let us assume the manufacturer believes that an 85% experience curve can be achieved, that he can sell at least 500,000 units, that he wants to recover \$150,000,000 investment in development by the 500,000 unit, and that he wants a 10% margin over these total costs to cover marketing and other associated overhead. From Figure 22, the cumulative average unit costs at the 500,000th unit are about \$600, and the unit share of the development costs are \$300. This gives a total cost of \$900, which with a 10% markup gives a sales price of \$990. Sales at this price beyond the 500,000 level become profitable, particularly if the \$150 million now recovered represented all of his development costs; in that case the \$990 includes a gross margin of 33%.




Figure 22. An illustration of the Dynamics of cost Behavior - III

Will 500,000 people buy ICC at this price in the first five years? We think yes. Are our assumptions about experience curve slopes and initial production costs reasonable? We don't know. Can ICC - self-check, failsafety and all really be developed and thoroughly tested for \$150,000,000? Probably not. But we know what to begin to think about.

4. THE SPONTANEOUS PLATOONING FUNCTION (AUTOPLATOON) -ADVANCED AUTOMATIC CRUISE CONTROL

OPERATIONAL CONCEPT.

The operational concept put forward here differs substantially from the notion of platoons of vehicles forming up off-freeway, going through some check procedure, entering the freeway and then proceeding as a unit to their common off-ramp.

The operational concept hypothesized here, based on the notion of spontaneous platooning, was first articulated to us by Dr. Loren Clare of the Rockwell International Science Center. His basic idea was that of a self-organizing system, a system which spontaneously reorganized itself to continuously adapt to its changing operating environment.

We begin at the beginning: the purpose of platooning is to increase the effective capacity of the freeway. It accomplishes this by reducing the normal separation between some vehicles from perhaps 50 to 100 feet down to 2 or 3 feet, thus freeing up space that can accommodate additional vehicles. Put in mathematical terms, it increases the density of vehicles without the necessity of changing speed to maintain safety.

(A glance back to Figure 7 illustrates that cruising with a spacing of 2-3 feet is essentially as safe

in terms of longitudinal collision velocity as the safegap at the other end of the curve.)

We would like to platoon only when the additional freeway capacity it provides is actually needed. We therefore would like to initiate platooning Only when the density of vehicles approaches the capacity limit (maximum flow accommodation) of the freeway at the existing speed. When that point is approached the need for platooning is signaled, and some of the vehicles equipped for platooning that are adjacent to other equipped vehicles automatically move to platoon.

As vehicles desire to leave the lane, the drivers signal their intent to resume manual control and the vehicles automatically deplatoon, permitting the driver to take over and drive out of the lane. As vehicles leave the lane, the need for platooning decreases, allowing additional deplatooning.

The principles of spontaneous platooning have been simulated extensively, and have clearly demonstrated the ability of vehicles to self-organize into platoons as the traffic flow demand increases, and to deplatoon when demand decreases.

The trick is how to make the decision to platoon or to deplatoon in response to changes in the demand. Some of the possible techniques are briefly discussed in a special section below: The PlatooningDeplatooning Decision.

THE PLATOONING-DEPLATOONING DECISION.

One can think of three basic approaches to making this decision.

The fundamental assumption behind the first two approaches outlined here is that as more vehicles enter the lane, vehicles drive closer together until some minimum spacing is reached, at which point some drivers begin to slow down. The result, of course, is that the whole traffic stream is forced to slow correspondingly. This suggests two possibilities for inferring the level of congestion:

- First, measuring the gaps between vehicles to directly measure vehicle density, and use this value to determine when the platooning process should begin. The advantage of this approach, if we can make it work, is that it forestalls the need for traffic to slow.
- Alternatively, wait until speed change actually occurs as a sure signal that maximum capacity at the original speed has been exceeded. This latter approach has the advantage of being much simpler to implement than the density approach.

Every equipped vehicle is measuring the gap in front, and is able to communicate to any adjacent equipped vehicle. One operational scheme to exploit this basic idea of measuring the gap to infer incipient congestion is to have each vehicle communicate the gap it is measuring up and down the line of traffic, flagging those readings from vehicles in the midst of lane-changing or platooning-deplatooning maneuvers. There would be some limit on how far the sampling should extend from each car - two vehicles?, four vehicles?, ten vehicles?. Thus each vehicle would have a value for the local vehicle density, with which it could use the equivalent of the highway capacity curves to make the platooning deplatooning decision.

There are at least two problems with this scheme. First, when only a small proportion of the

vehicles are equipped, the sample will be truncated when the communication hits an unequipped vehicle. Perhaps it doesn't matter, because there is little need for precision, and, as will be shown, the gain from platooning is small until about half the vehicles have the spontaneous platooning feature added to ACC (which we call Advanced ACC).

The second difficulty is that the gap in front of each vehicle represents the choice of the driver when he or she adjusted Autogap, and may or may not reflect the level of congestion. The seriousness of this problem depends on how most drivers react to other vehicles entering the lane in front of them. If vehicles repeatedly enter - in itself a signal that there is a need for more capacity - some drivers at least close up the gap to discourage new entrants in front of them. If most drivers react this way, then the gaps may become a reasonable reflection of the state of congestion and demand.

The second approach noted above, in which changing speed signals the need for platooning, was to let high density manifest itself through decreasing stream speed, and let this be the trigger to initiate platooning. This is much less ambiguous than the approximate-density-measurement approach, and certainly easier to implement. This is the technique used in the simulations described in the Agre-Clare paper. It works.

²J. Agre and L. Clare, "Spontaneous Platooning: A Self-organizing Approach to Improve Flow Capacity", Presented at the Third Annual IVHS America Conference, Washington D.C., April 1993.

It is probable that some combination will finally evolve as the preferred approach. If each vehicle was counting the rate of intrusions - or the surrogate, the small speed loss if the car was forced to slow because of the intrusion - it might be possible to infer a combination of conditions that are a reliable guide to platooning. The wide variety of possibilities will require much more in-depth thinking and simulation.

There is also a basically different approach to the problem: rather than leave platooning to some form of distributed decision made by the individual cars, it could be turned over to the Traffic Management system (TMS).

For example, when the TMS observed congestion rising, it could broadcast a start-platooning signal to all vehicles in that segment of the freeway. It may not be required for all eligible vehicles to platoon, so some scheme to do it in steps might be preferable. For example, have one-fourth of all platoon-capable systems sensitive to A-signals, one-fourth to B-signals, and so on. This time TMS sends a B-signal, so all B-vehicles automatically platoon with the nearest equipped vehicle, whether it be A, B, C, or D. This results in something less than half the equipped vehicles joining in two-vehicle platoons; the B-B combinations and the Bs not next to an equipped vehicle keep the total below one-half. Still more congestion causes the TMS to send both B and D signals. And so on. This scheme removes the spontaneity from the decision, but retains the autonomous action in forming and dissolving platoons. (It also may be an important tool of velocity-flow control to cope with flow interruptions, but that is another subject.)

Once the basic decision is turned over to the Traffic Management System, one can think of other variations. One possibility is to postpone, and possibly forestall, the need for platooning. A signal broadcast into the segment of the freeway that was beginning to congest, for example, could take gap control away from the driver, and have all equipped cars close the gap to safe gap - the minimum safe gap at that speed with the automatic system operating. As noted before, this would reduce the gaps at 60 mph from a driver-chosen 100 or so feet to less than half that. Figure 12 characterized the potential gain

from this approach.

TECHNICAL DESCRIPTION.

Figure 23 summarizes the technical profile for Autoplatooning. Briefly, there are only three changes from the requirements for Automatic Cruise Control:

1. Vehicle-to-vehicle communication is required between adjacent vehicles in the same lane.
2. The range of accurate distance and rate of closure measurements is extended to close distances: down to a foot or so.
3. The vehicle response algorithms must be extended to control platoon formation, cruise in the platooned state, and deplatooning. If the algorithms require information beyond that available from currently-defined on-board measurements, then a source for that information will also be needed.

It is even more probable than with ACC that communication from the Traffic Management System to the vehicles will be needed; the circumstances are discussed further in the following.

The individual elements in the system mechanization are discussed in somewhat more detail below.

Desired Vehicle Responses - Information Needs.

The vehicle response controller must now include the dynamics of the platoon formation maneuver, the steady-state platooning during cruise and during speed changes, and the deplatooning maneuver.

Possibly the most demanding new requirement is providing the basis for the decision as to when platooning or deplatooning is desirable; this has already been discussed.

Sensing and Sensor Interpretation.

The same information required for Autogap is required for platooning, but the range of interest must now be broadened down to platooning distances. Further, platooning demands more precision in position control than Autogap, so the continuity and increment thresholds in measurement must be commensurate.

It is possible that platooning will require new sensing, but we do not now foresee that need.

External Information ~ Command Inputs - The Need for Communications.

Both platooning and the platooning deplatooning maneuvers require coordination of control actions between the involved vehicles, so communication between these vehicles is required.

Further, in order to count the vehicles in a platoon, and perhaps to pass on information relevant to the spontaneous platooning decision, it will probably be necessary to have a relay capability and a means of identifying the source of original message.

Last, if the platoon-deplatoon decision is relegated to the TMS, communication from the TMS to

the vehicles will be required.

As with ACC, it should be possible for the Traffic Management System to vary the stream speed by communicating commands to the ACC-equipped vehicles, and for communication from vehicle-to-TMS for May Day signals if the driver fails to retake command of the vehicle.

Vehicle Control accuation

These are the same as for ACC, except the precision of control during platooning is more demanding. This requires more vernier control than with ACC.

TECHNICAL FUNCTIONS	SPONTANEOUS PLATOONING
DESIRED VEHICLE RESPONSES - INFORMATION NEEDS	LIKE AUTOGAP, DEGREE OF BRAKING AND ACCELERATION - BUT FOR MORE VERNIER CONTROL THAN AUTOGAP REQUIRES. MUST CONTROL DYNAMICS OF PLATOON FORMATION AND DEPLATOONING.
SENSING AND SENSOR INTERPRETATION	SAME AS AUTOGAP, BUT PERHAPS BETTER ACCURACY AND MORE SENSITIVE INTERPRETATION AT CLOSE RANGE.
EXTERNAL INFORMATION, COMMAND INPUTS - COMMUNICATION	. VEHICLE-TO-Infrastructure SAME AS AUTO CRUISE CONTROL . VEH-TO-VEH REQ'D TO COORDINATE CONTROL ACTIONS 'WITH VEHICLES FORE AND AFT.
VEHICLE CONTROL ACTUATION	VERNIER CONTROL OF BRAKES AND THROTTLE

Figure 23. Spontaneous Platooning Mechanization: functional Description of Technical Elements Required

System Mechanization.

The mechanization architecture is illustrated in Figure 24; the only change is the addition of Vehicle-to-Vehicle communication.

POTENTIAL BENEFITS - MOTIVATION TO PURCHASE

Capacity.

The basic reason for adding a platooning capability to an ACC system is to increase the effective capacity of the freeway. This, in a sense is a substitute for the expense and environmental impact of additional concrete. As is discussed below, this fundamental motivation may or may not be persuasive to the individual driver.

Figure 25 illustrates the fundamental phenomenon that motivates platooning. Without changing a given safe spacing, platooning increases vehicle density. Since flow is the product of density and speed, platooning increases the flow potential at any given level of safety (almost - see the discussion of safety following).

As can be seen in the figure, the maximum lane capacity depends on both the size of the platoons and the spacing between them; the influence of platoon size is increasingly important at larger spacings. As we have already noted, the spacing that is safe depends on the speed of the system and the relative braking capability of the vehicles involved.

Platoon size is an analytically more complex issue. As already described, platooning can only occur when both vehicles are equipped with the Autoplatoon feature. Under the assumption that Autoplatoon

vehicles are randomly scattered among noAutoplatoon vehicles) the level of platooning is very low when few vehicles are equipped, and rises as higher levels of equipping increases the probability of contiguous equipped vehicles. As more vehicles can platoon, both the average and the maximum platoon sizes increase.

Under the spontaneous platooning concept, the size of platoons is also strongly influenced by the proportion of vehicles of vehicles entering the lane relative to those leaving at each ramp location. Platoon size is thus determined by the proportion of vehicles with the Autoplatoon feature, and the

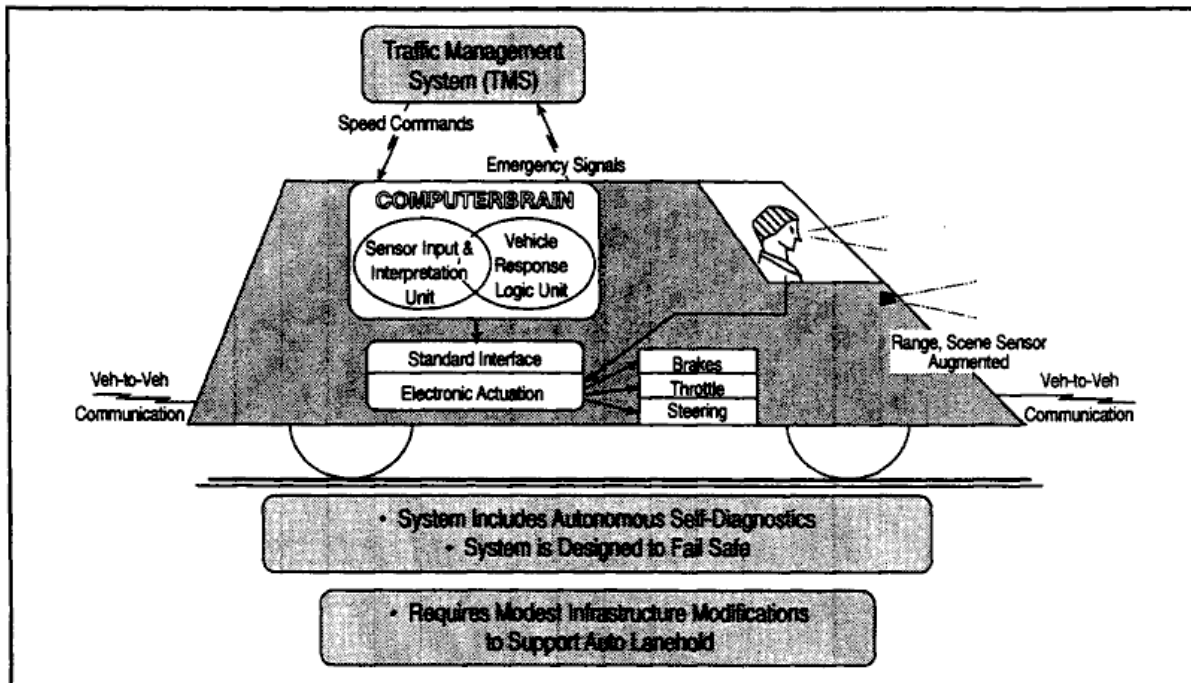


Figure 24. Add Spontaneous Platooning: Advanced Automated Cruise Control




Figure 25. Densityspacing, and Platooning Relationships

dynamic balance between the number vehicles entering the lane and potentially adding to platoon size, and the number leaving which force them to break apart. When these latter numbers are in balance, very few platoons will ever get above two or three vehicles. When more vehicles are entering than are leaving, platoons will get larger unless they are constrained to some specified maximum.

Figure 26 shows the variation in lane flow as a function of the percent of vehicles equipped with Autoplatoon. These curves include both the effects of decreasing average "safegap" and increasing average platoon size with the proportion of equipped vehicles. The maximum platoon size is constrained to 4 vehicles. The speed effect is mitigated by the fact that safegap increases with speed, and maximum flow is apparently peaking around 80 mph.

For the conditions shown here, the proportion of cars exiting at each ramp has little impact on maximum flow, but, as will be shown later, it does impact the distribution of platoon sizes.

One of the simplifying assumptions in the computations of Figure 26 was the omission of any extra lane space to permit entry and exit maneuvers. It is not clear to us today whether extra space is really needed, but prior calculations - more conservative on this issue - would suggest that this omission results in overstating the maximum flow by as much as 15 percent.

Figure 27 shows the distribution of platoon sizes for two different levels of exiting. As expected, more exiting breaks up more of the larger platoons, thus changing the distribution of sizes, but because the broken 4-vehicle platoons add to the number of smaller platoons, the overall impact on flow is dampened.

While we believe the trends shown in Figure 27 are valid, we have less faith in the specific numbers because of simplifying assumptions made in the calculations. These computations and the major assumptions are described in Appendix A.

Driver Motivations.

The addition of the spontaneous platooning feature to Automatic Cruise Control offers no new relief from driving chores. Further, it will take some getting used to: people are not accustomed to driving only three feet or so from the next vehicle.

The effective capacity increase just discussed is, in itself, a dubious motivation. The argument can be that, in fact, drivers will perceive only negatives from platooning, that the primary beneficiaries of this extra capacity are the additional cars coming into the lane, not the drivers already there who are doing the platooning.

There is an alternate perspective, based on the notion that the consequences of platooning is less to make room for more cars than it is to maintain reasonable speeds in spite of more cars entering the lane. In this perspective the platoon is directly benefiting from his or her vehicles' action through reduced congestion. It is true that he or she is not the only beneficiary which may dilute the motivation to be the one who does the platooning.

The difficulty with this latter perspective is that increasing effective freeway capacity is unlikely, by itself, to reduce congestion, simply because growth in demand will probably outstrip growth in capacity. If demand outstrips capacity, then congestion itself allocates freeway space, and cars will continue to enter a freeway until traffic is already severely slowed - no matter what the capacity level.

An actual reduction in congesting will require some form of demand management, like more stringent metering or road pricing. The additional capacity means that the queues do not have to be as long or the price hurdle that permits entry to the freeway doesn't have to be as high because it is not

³J Ward, "The Contribution of Platooning to Increased Freeway Capacity", Presented at the First Annual IVHS America Conference, Reston, VA, March 1991.

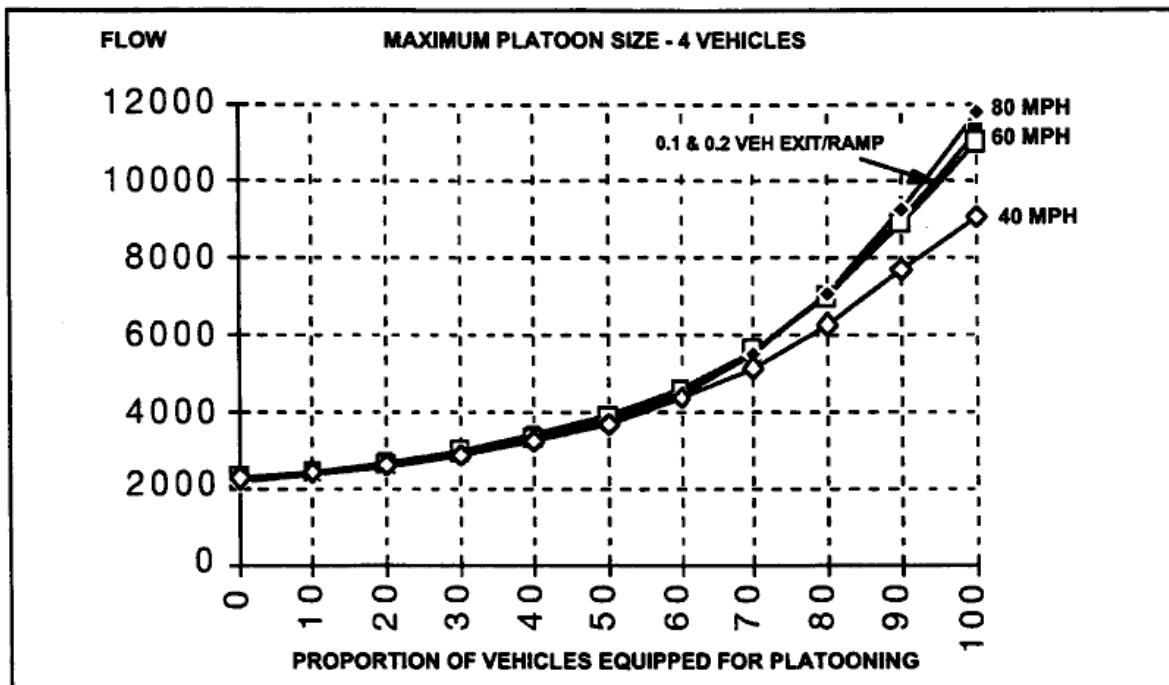


Figure 26. 1-8n. Flow 838 Function of Proportion of vehicles Equipped with the Autoplatoon Feature

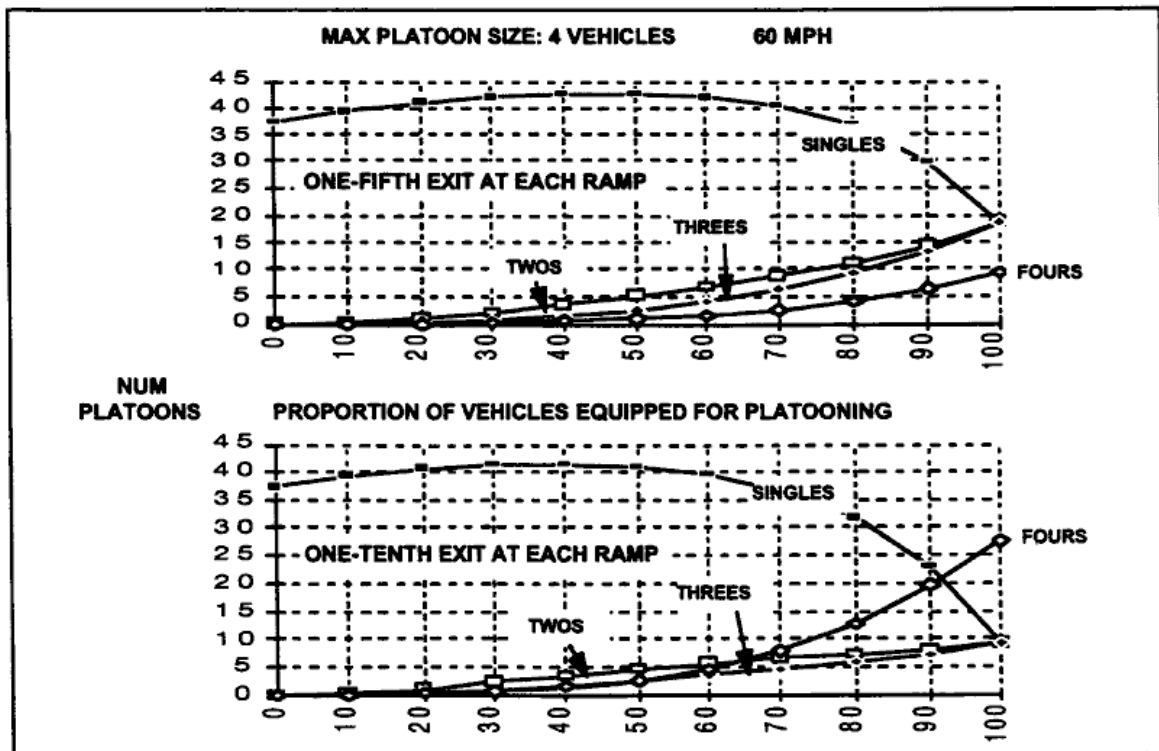


Figure 27. Variation In Platoon Size Mix with Proportion Equipped

necessary to discourage as many vehicles. But it is unclear whether drivers who would prefer not to platoon will ultimately perceive the benefits of doing so anyway.

The clearest beneficiaries are the public at large who don't have to tolerate or pay for new freeways because of the effective capacity gains of those already in place. It is just not clear how the individual driver will view the system.

Safety.

Spontaneous platooning introduces some penalty to safety. There are two sources. The first is the obvious one of driving close: while the collision velocity between the two adjacent vehicles is too low to be unsafe, there is always the possibility, however small, that lateral forces - or a lateral collision from another vehicle - could cause a general pile-up; this is an obvious incentive to keep the size of platoons small.

Second, there is some new risk added from the closing-to-platoon and the opening-to-deplatoon maneuvers. Figure 4 shows that while being close is safe and being beyond safegap is safe, in between is not. To perform these maneuvers vehicles will have to cross through these unsafe spacings. While system design will be such as to reduce the risk to almost zero, the "almost" may be small, but it won't be zero.

There is a trade between safety and capacity in this concept. We believe that platooning can be mechanized to actually improve net safety and still significantly increase effective capacity. But evaluating this trade is beyond the scope of this paper.

Summary. Withal, it appears that the motivation to equip one's vehicle with the ~~the~~ Advanced Automatic Cruise Control system capable of spontaneous platooning is less than it has been for either ICC or ACC. One can think of two approaches for increasing that motivation. First, it may be rational to pay people to equip their cars with the ACC with Autoplatoon, rather than just ACC, since it substitutes for more road surface.

Second, it may be justifiable to require that all ACC systems sold beyond some date include Autoplatoon. This would cut the Gordian knot on motivation to equip lot platooning, just as the scheme based on TMS control of the platooning decision cuts the Gordian knot on who, when, and where platooning takes place.

Usefulness of the Freeway

Platooning should be useful on some crowded interstates. The issue of its practice on surface streets will require more investigation; at the moment it appears dubious.

DEVELOPMENTAL ISSUES AND RISKS.

Prospects for development.

The problem here is sorting out who benefits, who pays. The ambiguity in motivation to purchase will ease doubt on the market potential in the eyes of private developers. And if the Traffic Management System is made part of the decision to platoon, then there is a need for compatibility with a public sector owned and operated system. We cannot predict with confidence that this enhancement to ACC will be

viewed as a good use for private capital.

Since, however, platooning is a substitute for additional concrete, there is justification for either public subsidy or even public mandate.

Robustness for Development

As always, the Integrity Verification Subsystem -the self-test and self-diagnostic system - will be an essential and integral part of the system.

Vehicle Behavior Algorithms.

Working out the best way to handle the platooning decision process is one new element in the development. The second is the control of the dynamics of the maneuvers themselves, including steady state cruise in the platooned position.

These classify more as work that will have to be done than technical risks. *Human Factor Issues.* There is nothing here that pushes the state of the art.

Potential for Retrofit.

Highly unlikely.

The primary one is the acceptance of the close proximities platooning implies. This is clearly daunting now: how much it will continue to be after a few years with ACC remains to be seen.

PART III SUBSEQUENT EVOLUTION OF AHS AND

AUTOMATED VEHICLE CONTROL

The evolution hypothesized so far carries us to fully automated cruise in a freeway or highway lane. Entry to and exit from the lane are under manual control, as is all operation on surface streets. In this Part III, we sketch some possibilities for subsequent evolution.

AUTOMATIC LANE CHANGE

We conjecture that the next step in evolution is the addition of automatic lane change on freeways and highways.

There are at least two justifications for the function. The first is safety: we have noted before that there are occasions when even full emergency braking cannot prevent collisions, leaving the only other option to pull out of the lane to go around the object. Performing this maneuver automatically will be considerably faster than depending on driver reaction.

The second justification is the aging of the population and the desire to make driving safer for the older driver. Having automatic lane change opens the possibility of having the complete freeway trip mechanized, so the driver need only drive the car to the on-ramp, and take control again after exiting the freeway. The driver would have identified the desired exit lane, and the vehicle would navigate to it using either on-board navigation or by detection of identifying signals from the road itself.

The mechanization of automatic lane change is, like all those hypothesized so far ~ autonomous, and insofar as we can envision now, requires no communication with the Traffic Management System other than that already available with Automated Cruise Control. The vehicle will require more on-board information gathering, however, because it will be necessary to monitor conditions in the receiving lane to insure safety. The top level mechanization is shown in Figure 28.

Next Steps. It is unclear what might happen next. We can think of two not-mutually-exclusive possibilities. The first is a dramatic increase in highway cruise speeds. The other is to extend automatic control to surface streets. This latter is a very important step - or series of steps - that could open up a whole new world for urban transportation.

SUPERCruise

The motivation for increased intercity cruise speed is reduced travel time. Higher speed will increase the trip length at which air becomes the dominant mode, and increase the demand for highway travel both by this modal shift effect and by inducing new travel. The impact will be similar for the movement of high value freight.

The advent of higher highway speeds are likely to induce changes in the competitive air and rail modes. In heavily traveled corridors, such a move on the highways may well precipitate new high speed ground systems that would serve both freight and passengers. It is beyond the scope of this study to examine these consequences, or the net impact on demand for highway travel.

Markedly increased speeds introduce two new technical and operational problems. First, it would necessitate a substantial increase in the range of the sensors in the Intelligent Cruise Control system. Second, it opens new issues about the modus operandi for controlling access to the high speed lane. It emphasizes the need to protect the lane against foreign intrusions of any kind, and may pose the need for physical barriers.

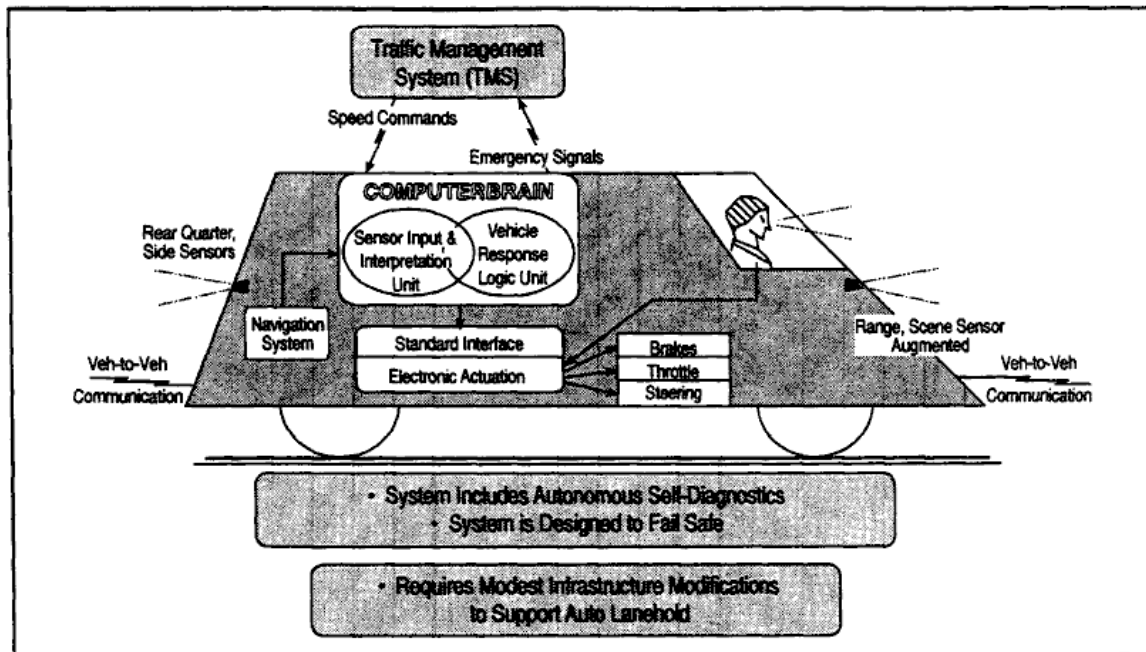


Figure 28. Add automated Lane change: Fully Automated Freeway Trip

SURFACE STREET OPERATION: "THE INTEGRATED IVHS"

In our opinion, this step will be as significant to the future of ground transport as the original introduction of automatic controls.

Freeways and highways basically constitute one long "link" on which the phenomenon of primary interest is the longitudinal spacing between vehicles. The important variable is the spacing between vehicles, a variable most easily measured and manipulated from on board the individual vehicle. This leads naturally to the essentially autonomous mechanization's described.

On surface streets, the primary interest shifts to the "nodes", the intersections. Links are short, and particularly in high density regions the behavior on the links is largely dictated by the conditions at the node. The node, not the link, is the dominate influence on surface street traffic behavior.

This shift from link-behavior-dominated phenomena to node-behavior-dominated phenomena leads to a different mechanization for an automated vehicle system. Now two functions are needed to effectively control flow at an intersection. First, we continue to need to insure safe separation between vehicles in a vehicle stream; we are already getting the necessary information from the vehicles' Automated Vehicle Control system, and it seems sensible to continue to carry out this safe-separation function with an on-board system.

The second function is the macrocontrol of the traffic streams through the intersection, in essence the allocation of time-slots to each of the desired travel directions to optimize total throughput and perhaps bias some flow directions to prevent jam densities in other parts of an area-wide system. The information to perform this function is beyond the reach of any individual vehicle, but is that already being collected and used this way by the Traffic Management System (at least, in advanced systems). It only makes sense to continue with this functional allocation, but rather than send traffic commands to a driver, we will send them to the now much more versatile "Computerbrain" on board the vehicle. The driver is no longer in the loop.

We have made a fundamental change in the AHS architecture from essentially autonomous vehicles behaving in a self-organizing system, to an integration of the Smart Street with the Smart Car. The Advanced Traffic Management System (ATMS) provides macro-commands, and the on-board Automated Vehicle System executes them and provides safe micro-control. Figure 29 illustrates the notion of this new integration.

Who navigates? There are three options. The first is an on-board navigation system that is heading the vehicle toward the destination selected by the "driver-passenger", but reacting to congestion information furnished by the ATMS in route selection.

The second option is guidance by an even-more-capable ATMS: the "driver-passenger" informs the system of his or her desired destination and the ATMS furnishes the necessary commands to the vehicle.

The third is the "driver", whose control inputs go to the wheels, but modulated and modified to prevent any unsafe actions. It is quite likely that the nature of the drivers controls in this mode of operation will have changed dramatically, since there is no need for anything more than 'turn-left-at-the-next-corner' type of inputs. Figure 30 reflects this very important step in the metamorphosis of the Automated Vehicle System.

This is all, of course, distant conjecture. But two points are clear. The first is that once we move to surface streets, we are on the path to the driverless vehicle, with all its implications.

The second is that AHS - or AVCS - cannot be ignored in the development of an IVHS architecture. While there was little interplay when the system was confined to the freeway; when the first steps toward total-trip-automation are taken, the AVCS system becomes inextricably intertwined with the other elements of IVHS. In fact, this will become the heart of IVHS.

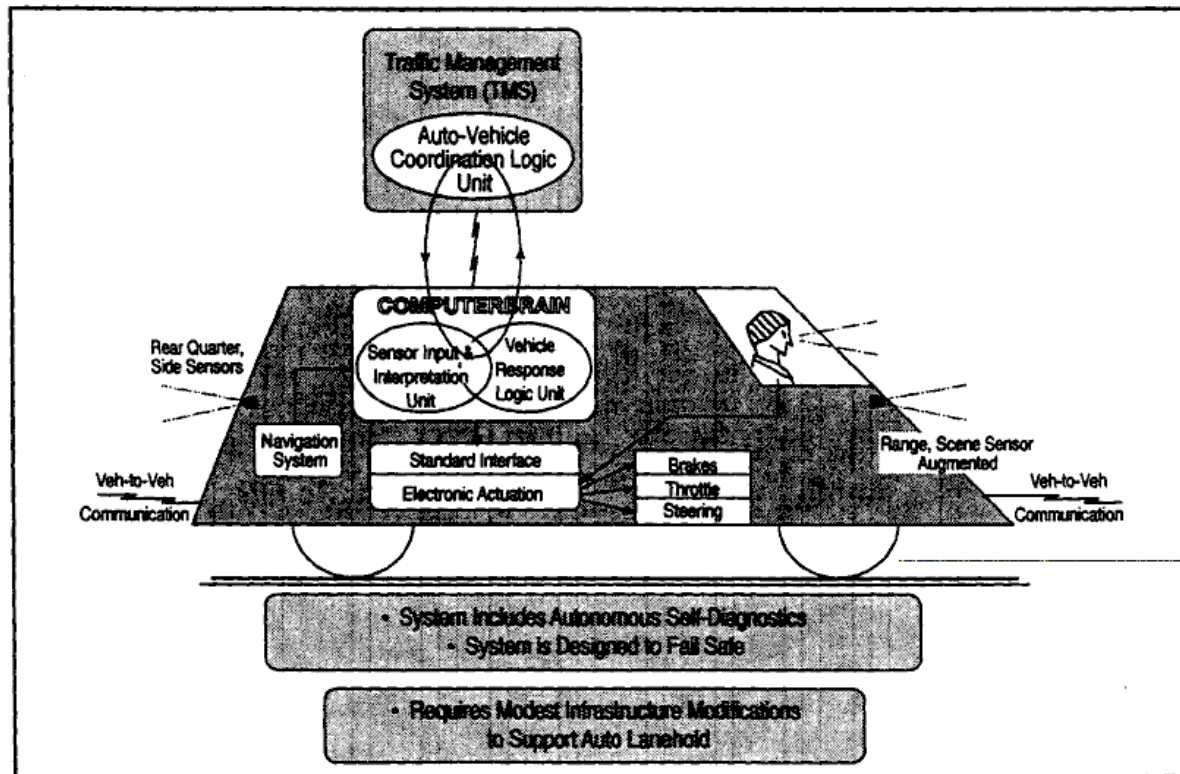


figure 29. TM-AVCS Integration Fully Automated on Freeways and traffic Controlled Surface Streets

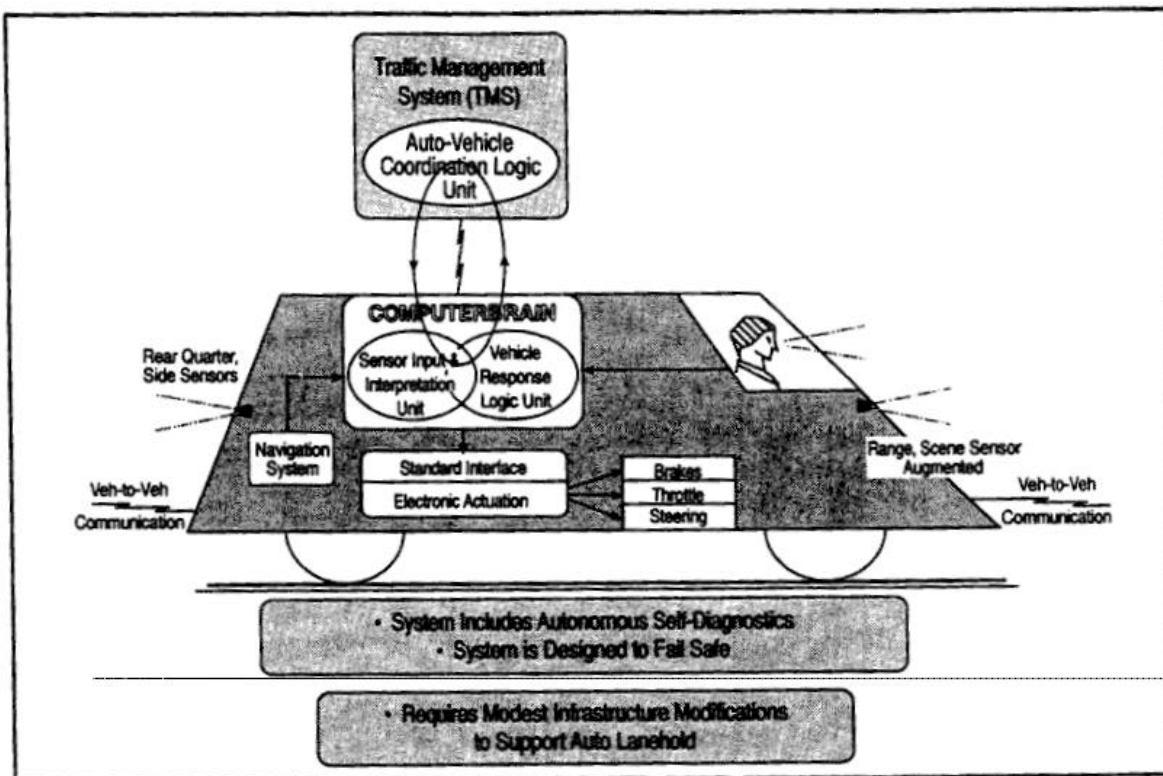


Figure 30. Automated Manual Control
The Safe Vehicle

POSTSCRIPT

There is an almost seamless path to a future for road transport that is almost beyond our imagination today. The technology is either in hand, or within stretching distance. The window of opportunity is open.

Appendix A

SPONTANEOUS PLATOONING: the CALCULATION OF PROBABILITIES AND PARAMETERS

Vehicles can only platoon if both are equipped with the requisite systems. Since it will be many years before all vehicles are equipped, we will have equipped vehicles operating in mixed traffic with unequipped vehicles. It was, therefore, desired to calculate the probability of platoons of various sizes being able to form in such a mixed traffic environment.

Here we describe the calculations carried out to estimate this effect. The problem is complex and dynamic, and detailed simulation like carried out by Agre and Clare (Reference 3) is the preferred approach to analyzing the problem. Such simulation is well beyond the scope of this study.

We assume that equipped and unequipped vehicles are mixed at random. We also assume that there are always some vehicles preparing to exit the freeway at each ramp that, if they are part of a platoon, must break up or break off from that platoon.

The spread sheet used in the calculations is attached. We will comment on the points not considered self-explanatory.

CL is average vehicle length in feet.

DECELA is the maximum deceleration capability of the lead vehicle in g's, and DECELB is that of the following vehicle. Here we assumed it was 85 percent of that of the lead vehicle. The DELAY is the braking delay of the following vehicle. It was varied linearly with proportion of Autoplaton-equipped vehicles, as shown in the sample spread sheet, also attached. These parameters were used to calculate SAFEGAP, the average spacing between vehicles. This average value was a surrogate for the mix of manual safegaps where the following vehicle was unequipped, and auto-safegaps where the following vehicle was equipped.

p is the probability that a vehicle is equipped with Autoplaton, and q that it is not.

The next four variables, PSINGLE... PQUAD, are the probability that a given vehicle will be part of a platoon of size n (i.e., SINGLE, DOUBLE, TRIPLE, QUAD). A "platoon" of size one is a single vehicle.

The variables PPIATn give the probability (the proportion) of platoons of size n.

DISTn (DISTi throughDIST4) give the total distance in the lane taken up by a platoon of size n. As can be seen, it assumes 3 feet between vehicles and SAFEGAP between platoons.

NVEHICLES calculates the vehicle density - the total number of vehicles in one mile.

The next four variables give the number of platoons of a given size. At this point they are artificial, calculated only out of curiosity, because they have no constraint on size from exiting vehicles that would prevent their build-up. Note that on the sample, when all vehicles are equipped, all the platoons are the maximum size. Without the constraint on platoon size built-in, there would be just one solid platoon of infinite size.

We now introduce constraints. With exiting vehicles there are a number of considerations. First, the proportion of vehicles that are exiting is given by the ramp spacing R_s divided by the average trip length (TL). We ignore trying to treat trip length as a distribution.

Second, all the exiting vehicles are not Autoplaton-equipped; only those that are equipped will affect the platoons. This is accounted for in the calculations.

Third, for 3 and 4-vehicle platoons, the probability that the exiting car comes from the outside vehicles in the platoon is, in general, different from the probability that it is an interior vehicle. This probability of an exterior vehicle is given by $2K/n$, and that of an interior vehicle by $1-2K/n$, where n is greater than 2. K is a factor to adjust for the impact of average trip length: if trips are long and vehicles have been entering regularly, the odds are high that the exiting vehicle - which has been in the lane for a long time - has become an interior vehicle, in which case K is nearly zero. If trip lengths are short, then there is a much better chance that the exiting vehicle could be an outside vehicle, in which case k is close to one. We could only guess at the relationship, and we ultimately elected to leave K at 1 for the results shown.

The last consideration is the distribution of exiting vehicles among the platoons. The probabilities we have calculated tell us how many vehicles (probably) are leaving platoons of a given size, but they do not tell us the distribution among them. Largely to simplify the problem - which was already complex enough for this brief analysis - we assumed that each exiting vehicle came from a different platoon. If the choice of values produced more exiting vehicles from platoons of size n than there were of size n , we zeroed out that size platoon.

This last assumption biases the results toward more smaller platoons and less larger ones, so will distort the distribution of platoon sizes. But as noted, it does not seem to materially affect the estimates of maximum flow potential, our key interest here.

Last, if the space required to support the entry and exit maneuvers is greater than normal safe gaps, then some provision should be made for it. Prior analyses lead us to think that it is unlikely to be larger than 15 percent, and may be negligible. This is a problem in itself, since the net space required depends on the time the maneuvers take, the number of vehicles involved, and conditions in the next lane. We did not attempt to accommodate it here.