U.S. Department of Transportation

Federal Highway Administration

Precursor Systems Analyses of Automated Highway Systems

An Hypothesized Evolution of an Automated Highway System

DTFH61-93-C-00201

Final Report
August 1994

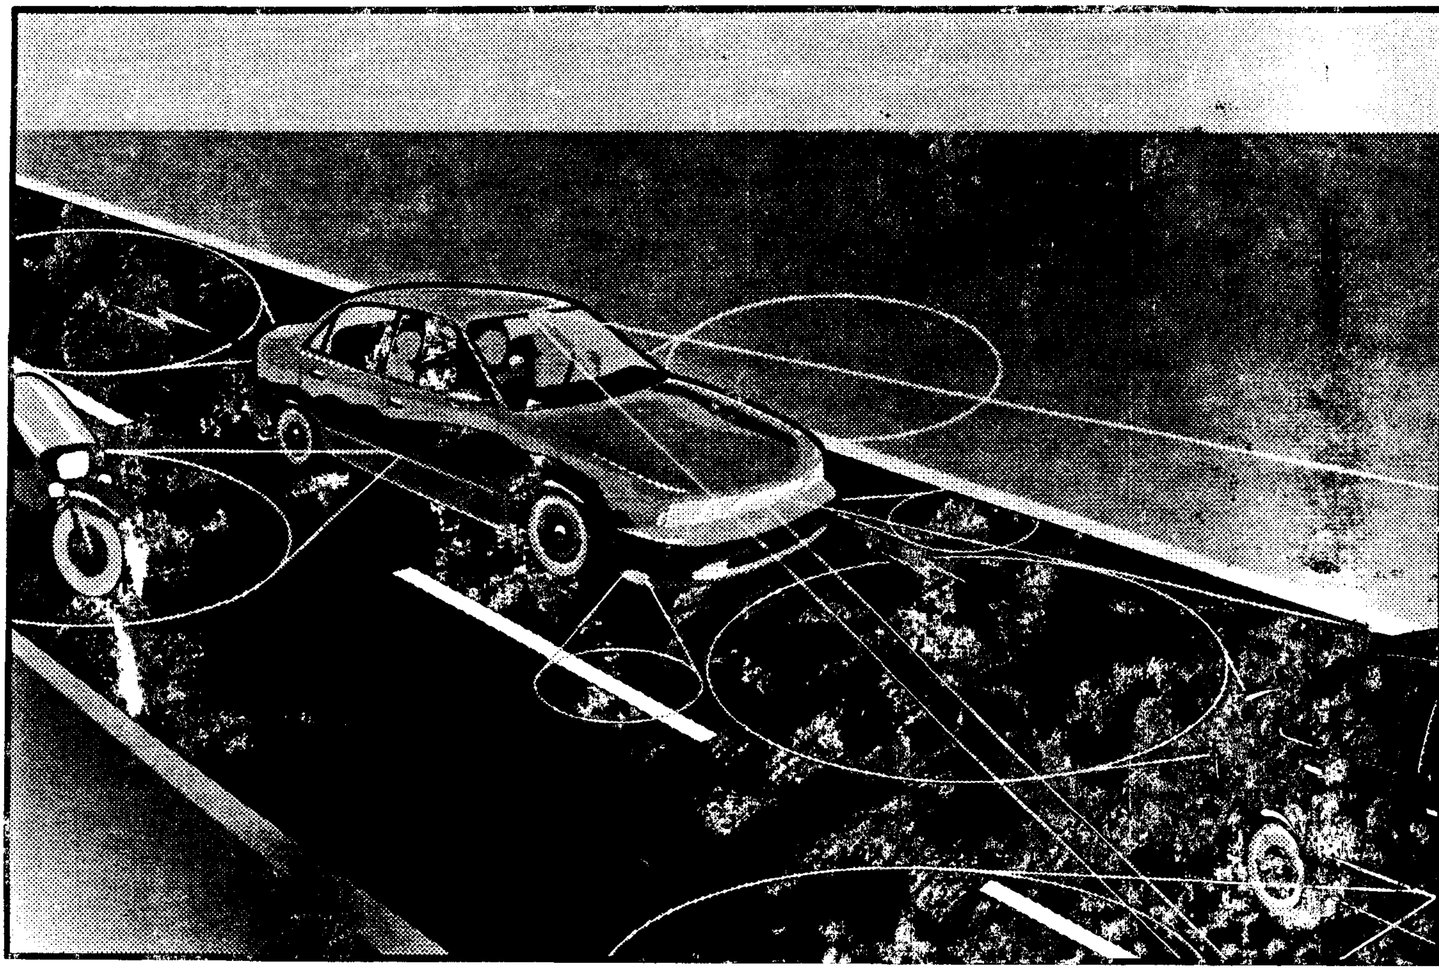

Rockwell International
3370 Miraloma Avenue
Anaheim, CA 92803-4192

EPS.30308.P01.1

Technical Report Documentation Page

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
4. Title and Subtitle Precursor System Analyses of Automated Highway Systems (AHS) An Hypothesized Evolution Of an Automated Highway System Final Report		5. Report Date August 1994 6. Performing Organization Code 8. Performing Organization Report No.
7. Authors Jerry Ward		
9. Performing Organization Name and Address Rockwell International 3370 Miraloma Ave. Anaheim, Ca. 92803-3150		10. Work Unit No. (TRAIS) 11. Contract or Grant No. DTFH61-93-C-00201
12. Sponsoring Agency Name and Address Federal Highway Administration Turner-Fairbanks Highway Research Center 6300 Georgetown Pike McLean, Va. 22101-2296		13. Type of Report and Period Covered Technical Report 14. Sponsoring Agency Code

15. Supplementary Notes

16. Abstract

This document describes and broadly analyzes an evolutionary deployment scenario for an Automated Highway System in which vehicles equipped for automated operation are deployed in mixed traffic with unequipped, manually operated vehicles. A gradual evolution of incremental change, both in scope and capability of the automated features in the vehicle fleet is assumed.

17. Key Words	18. Distribution Supplement

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price

Form DOT F 1700.7 (8-72)

An Hypothesized Evolution of an Automated Highway System

August 1994

Contract no.: DTFH61-93-C-00201

Prepared for: U.S. Department of Transportation
Federal Highway Administration

Prepared by:

Rockwell International

Autonetics Electronic Systems Division
Defense Systems
3370 Miraloma Avenue
P. O. Box 4192
Anaheim, California 92803-4192

CONTENTS

Page

PART I. INTRODUCTION AND SUMMARY 1
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
Costs and Prices 7
System Robustness and Failsafety 7
Safety Criteria 7
Autobrake – Obstacle Detection Requirements 7
Autobrake – Braking Criteria 7
Autogap – Sensor Interpretation Performance 8
Autogap – Communication with the Traffic Management System 8
Autolane and Automatic Cruise Control – Driver Alerting 8
Autoplatoon – Driver and Passenger Acceptance of Platooning Proximity 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 Benefits – 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 – Motivation 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

CONTENTS

Page

PART III. SUBSEQUENT EVOLUTION OF AHS AND AUTOMATED VEHICLE CONTROL 43

AUTOMATIC LANE CHANGE 43

SUPERCRUISE 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 Descritpion 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 Techmical Elements Required	27
19.	Adding Auto Lanehold: Automated Cruise Control Hands-Off Freeway Cruising – <u>The Mark I AHS</u>	29
20.	An Illustration of the Dynamics of Cost Behavior – I	31
21.	An Illustration of the Dynamics of Cost Behacior – 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 Feature	39
27.	Variation in Platoon Size Mix with Proportion Equipped	39
28.	Add Automated Lane Change: Fully Automated Freway Trip	44
29.	TMS-AVCS Integration Fully Automated on Freeways and Traffic Controlled Surface Streets	45
30.	Automated Manual Control – The <u>Safe</u> 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 road 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: arterials 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 a 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 of human accidents; automated systems could virtually eliminate these. This one trait of almost eternal vigilance is worth a large investment in 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 nondisruptive, 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.

A word 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 specify 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 platooning 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 of ACCs equipped with this feature. If absolutely necessary it could be by decree.

This evolution would continue into the future, 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

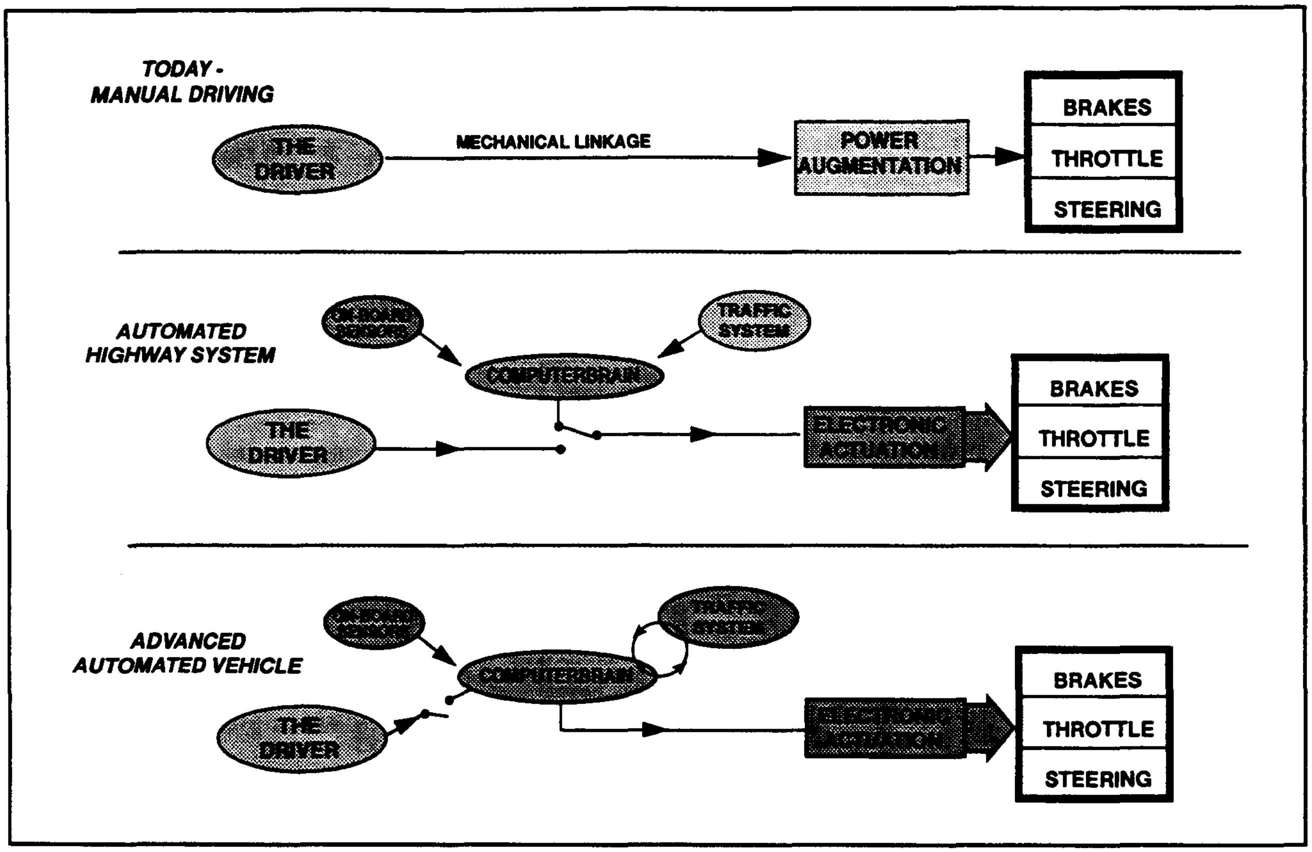

Figure 1. The Evolution of Automated Vehicle Systems

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 curves 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 far 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 won't 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 function is with vision-based sensors, radar, ladar, 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 entirely, 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-our-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 surface 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 Mazda[1] 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.

[1]SAE Paper 901484, Okuno et al, Towards Autonomous Cruising on Highways, Mazda, August 1990

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 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 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.

Costs 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 function 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 by 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 far 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 "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.

THE PLATOON-DEPLATOON DECISION.

The primary operational-technical 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 subfunctions that must be carried out to effect automatic controls of almost any 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.

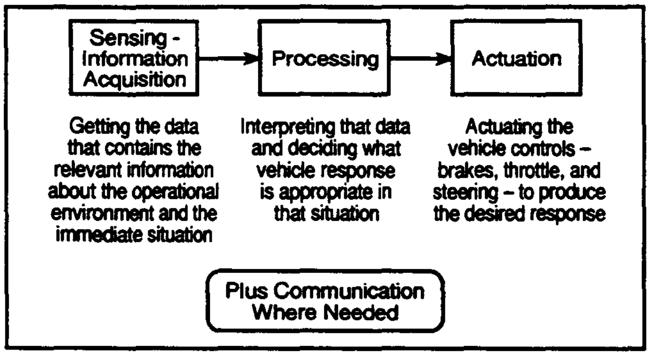

Figure 2. Generic Automated Vehicle Control System Functions

Figure 3 is a finer grain decomposition of these three subfunctions; we will use this more descriptive depiction of system subfunctions 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.

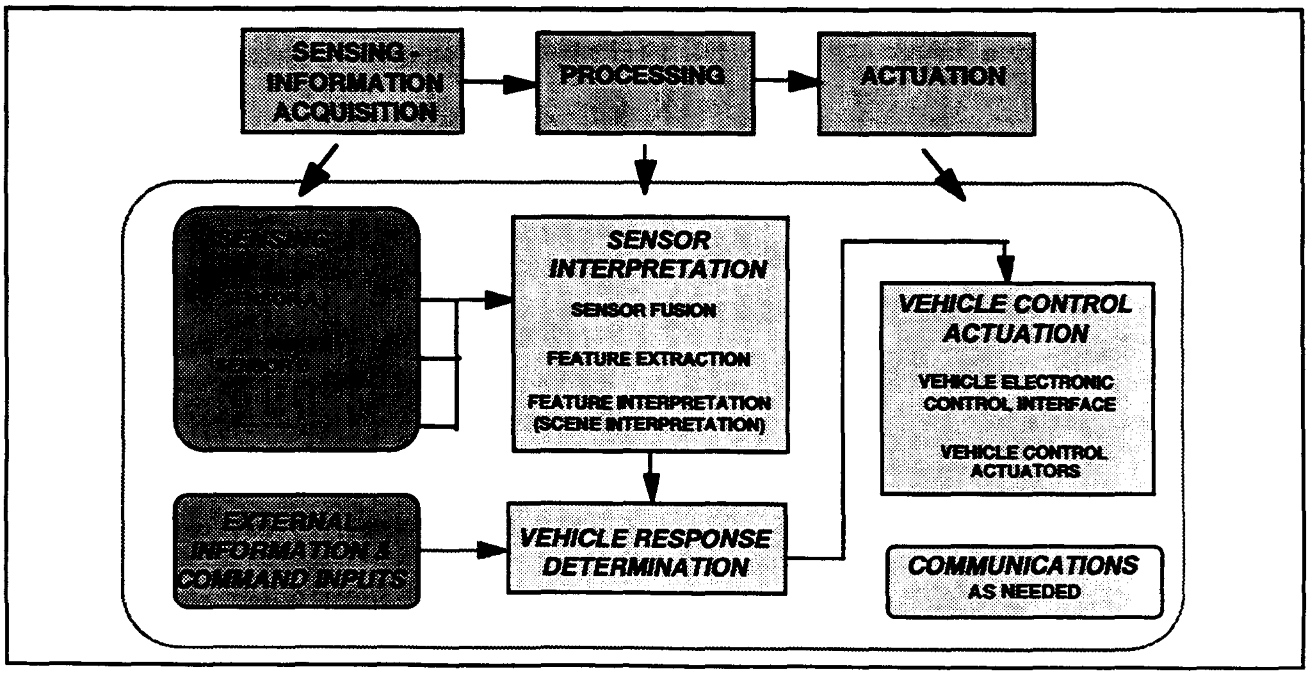

Figure 3. Generic Automated Vehicle Control System

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 rear-enders. 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

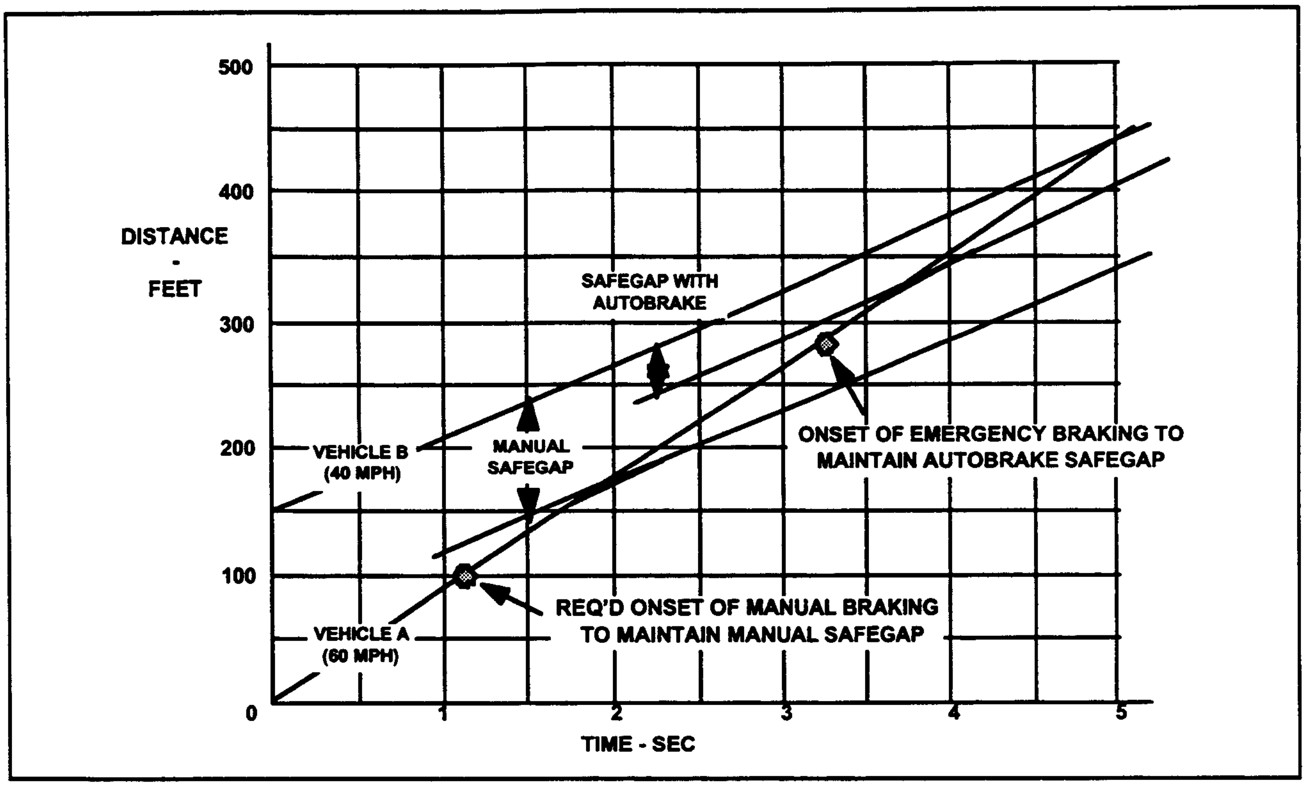

Figure 4. Emergency Braking Illustration

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.

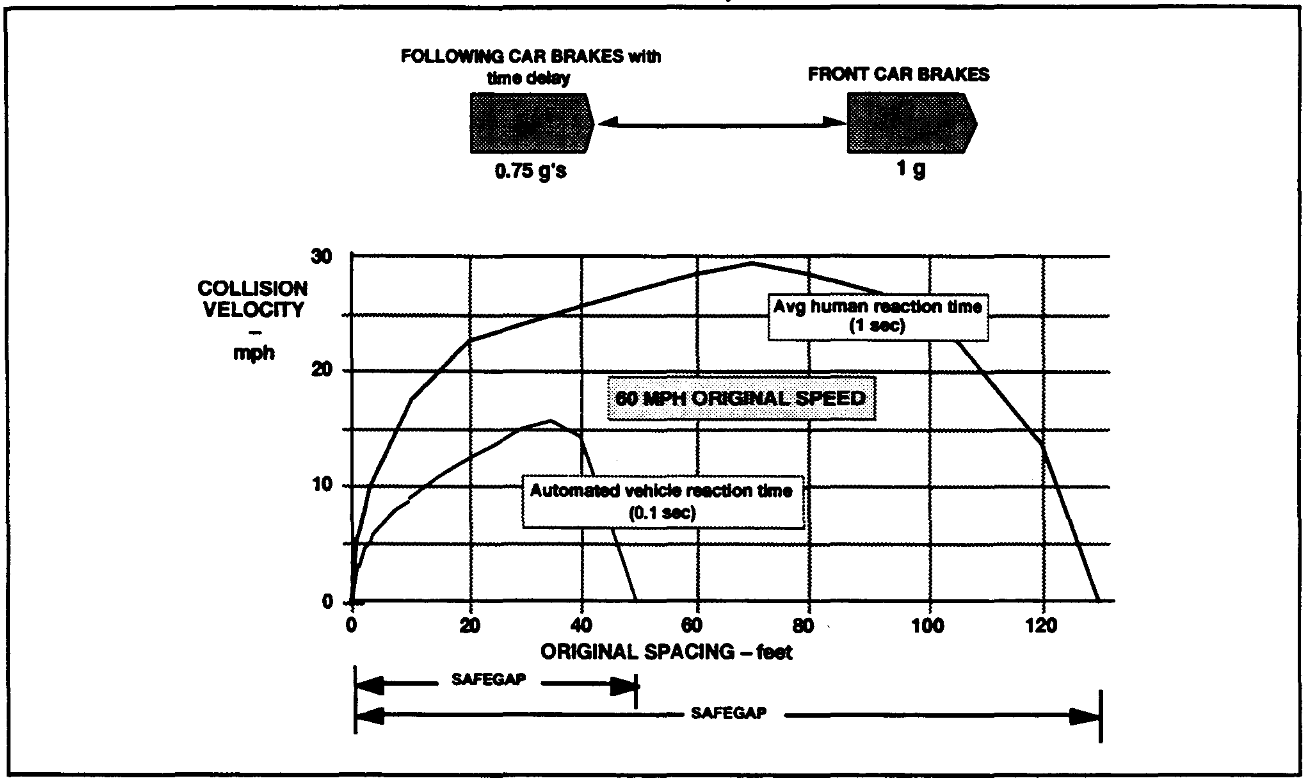

Figure 5. Faster Reactions and Better Brakes Permit Closer Operation

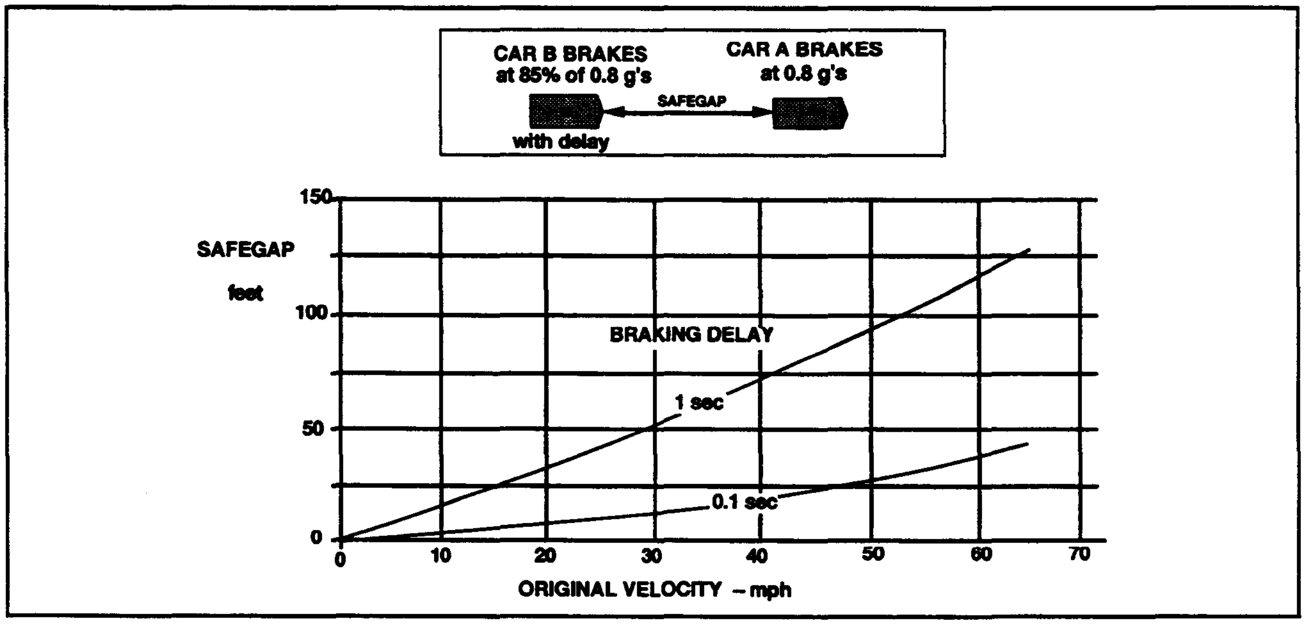

Figure 6. Safegap

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

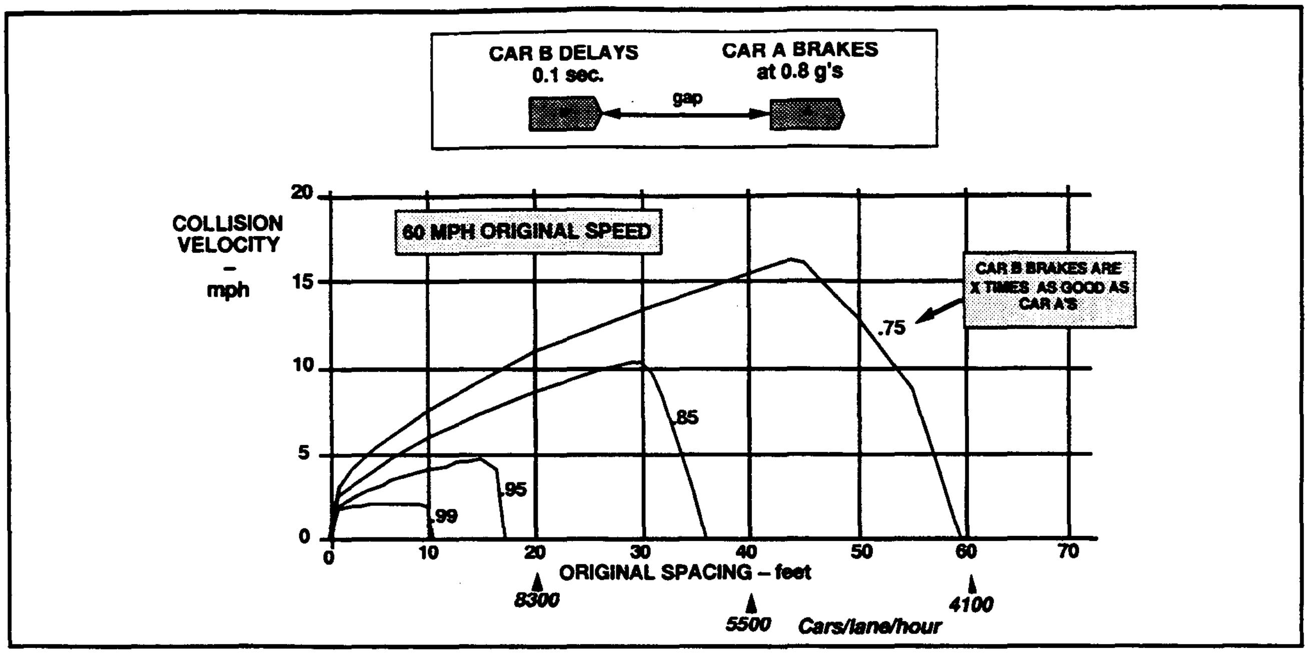

Figure 7. Machine Reaction Times Permit Closer Spacing – Brakes, Not Time, Become Dominant

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's (25.6 ft per sec2), 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 carlengths 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 sec2) and the second vehicle brakes at 0.8 g's (25.6 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.

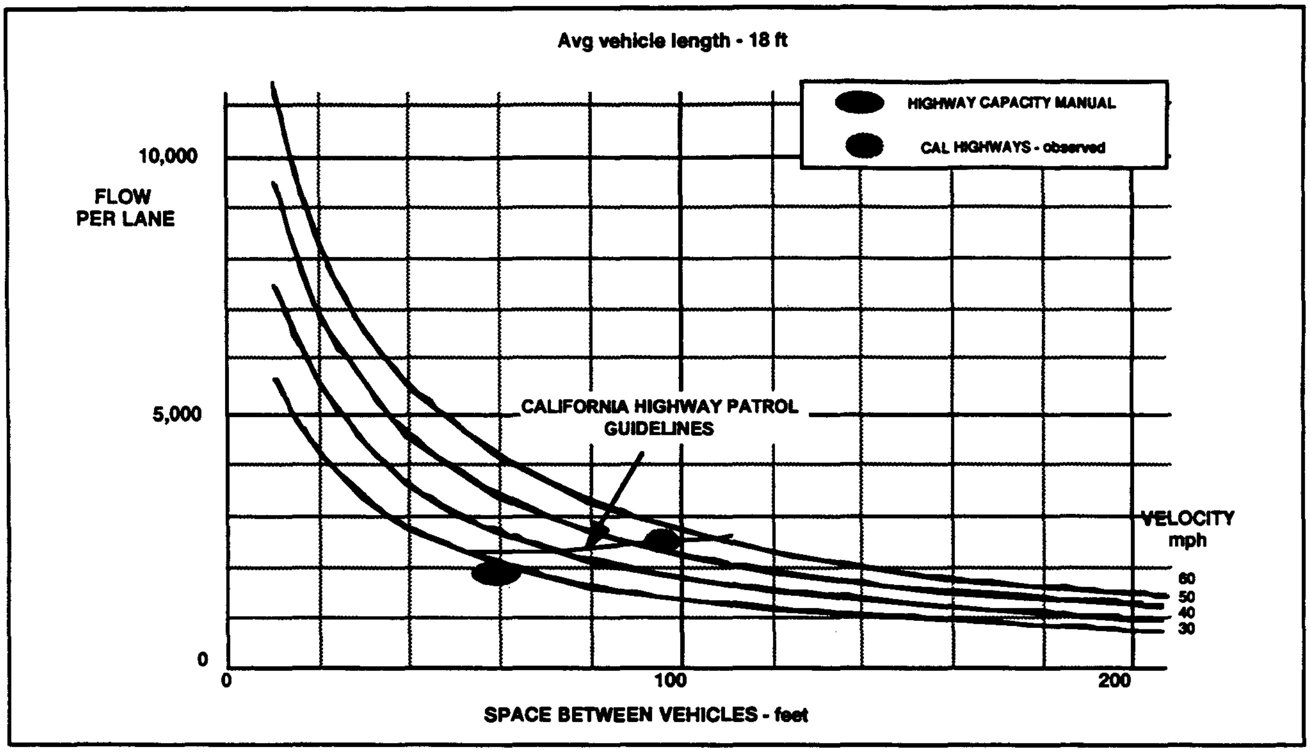

Figure 8. The Relationships Among the Gap Between Vehicles, Speed, and Flow

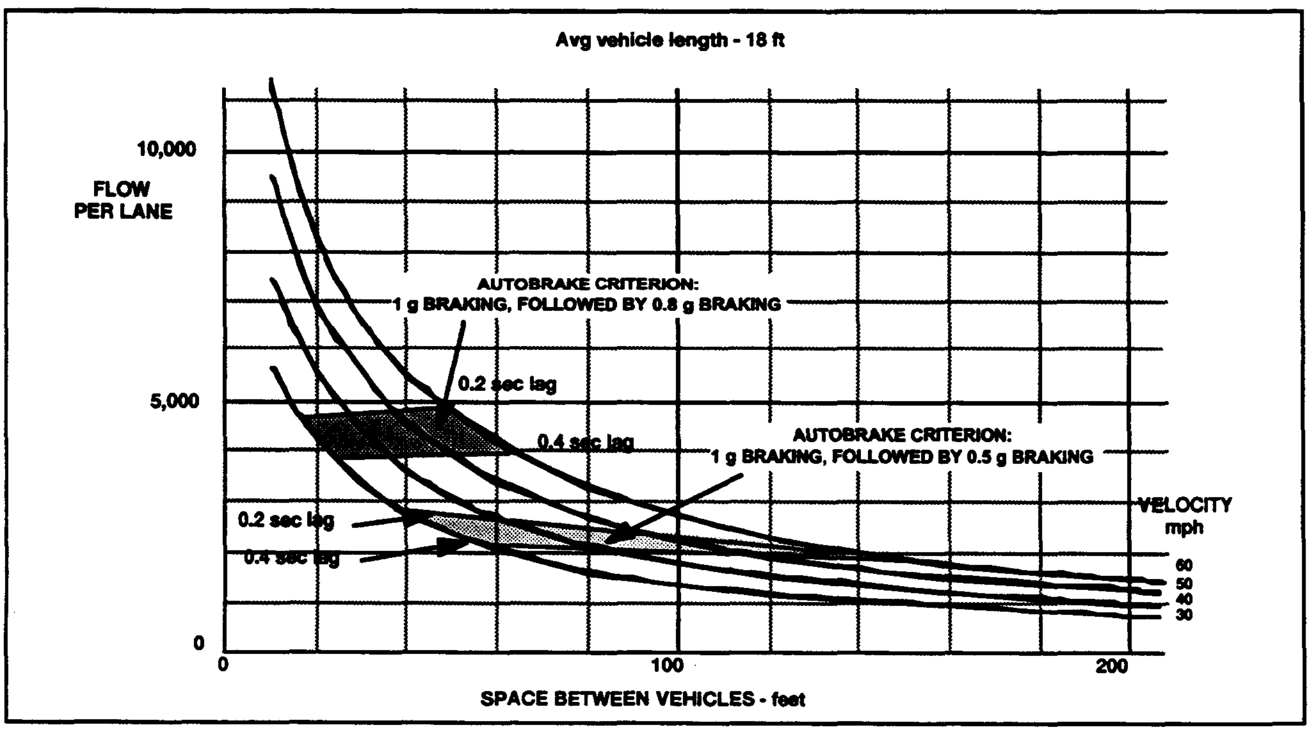

Figure 9. The Relationships Among the Gap Between Vehicles, Speed, and Flow

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 sec2). 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 to 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

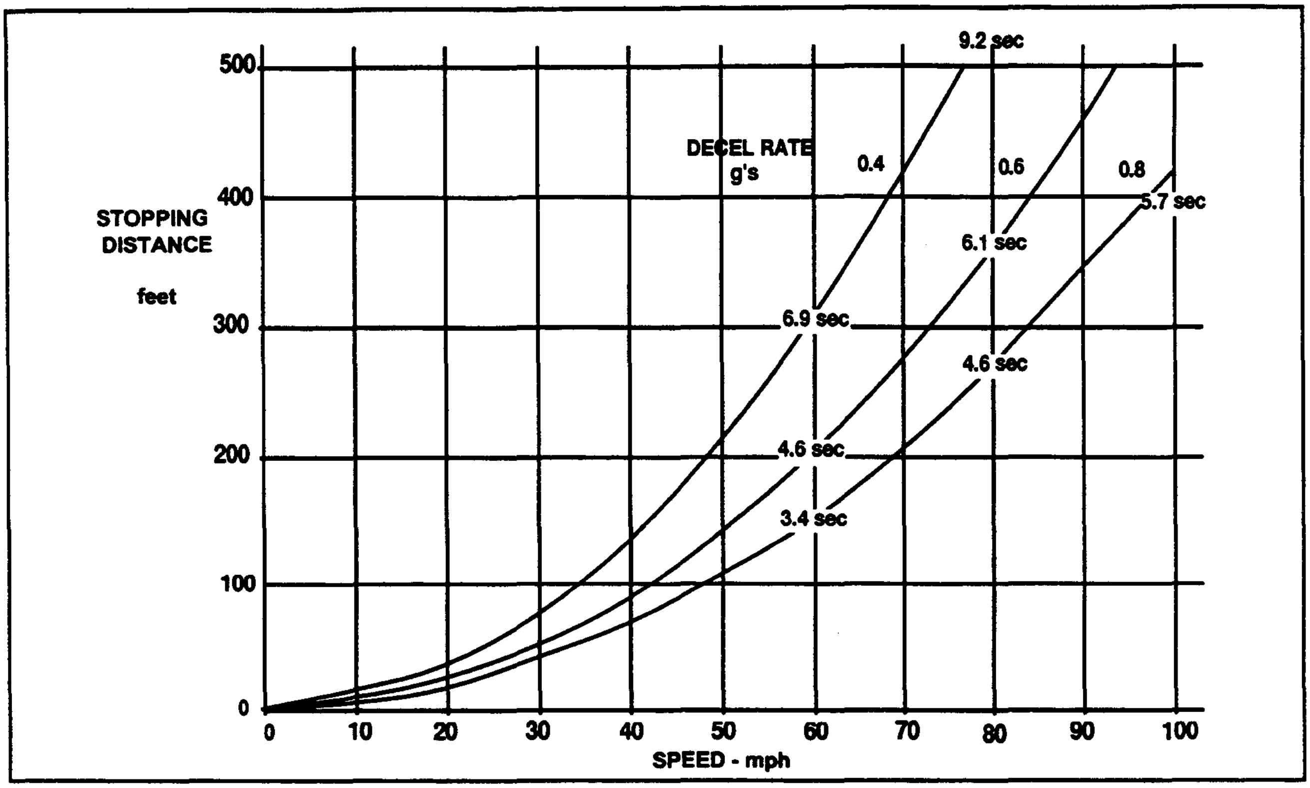

Figure 10. Stopping Distance as a Function of Speed and Braking Capability

0.8 g surface. (Here we have found a second reason to measure 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 fast 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 identify 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 safegap to the type of vehicle. This is not a requirement, perhaps only a mild desirement: we confess to being unsure of its usefulness.

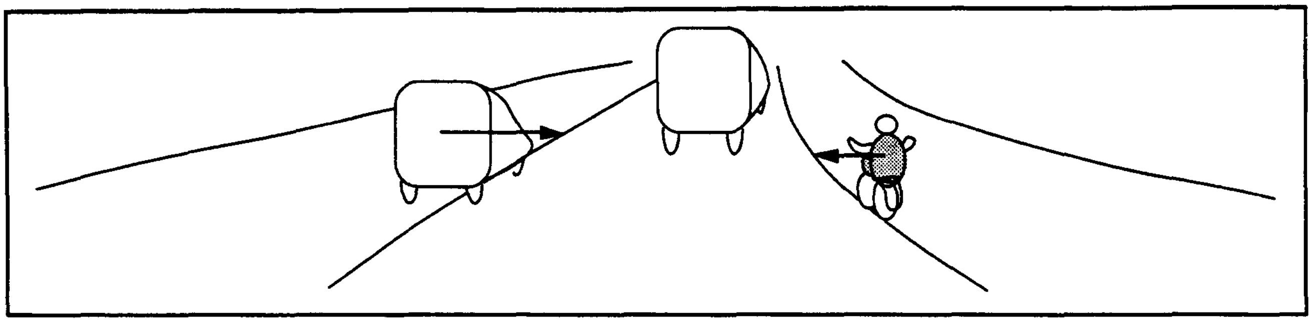

Figure 11. Typical Scene Looking Forward

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 one's 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, failsafe, 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 a 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

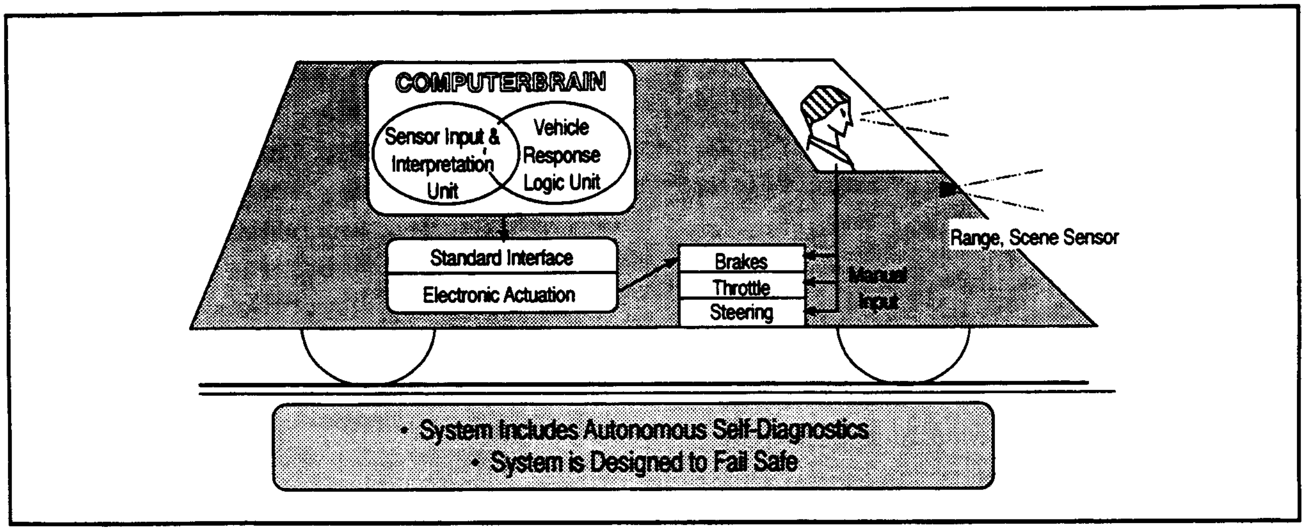

Figure 12. Autobrake Mechanization

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.

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	• 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	• FIELD-OF-VIEW ROUGHLY FORWARD QUARDRANT. "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-TO-ACT RATES. • 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 ELECTICAL 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 worse will be far less frequent than human lapse. Further, the driver would be informed of a non-functional 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 far 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 Off 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.

Prospects for Development.

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 Failsafety.

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 far 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 all 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

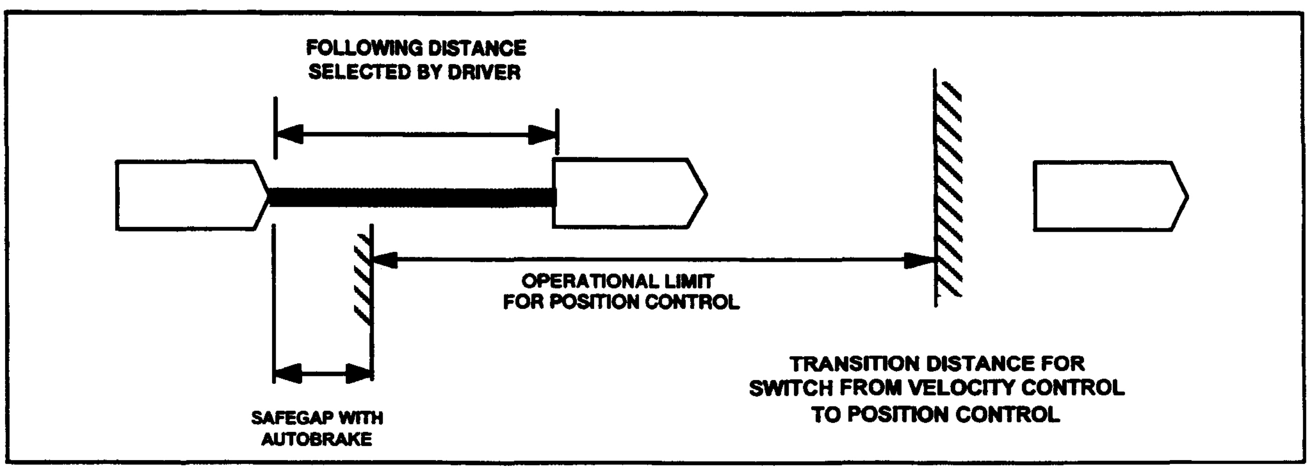

Figure 14. Automatic Gap Control