Title

Development and Performance Evaluation of AVCSS Deployment Sequences to Advance from Today's Driving Environment to Full Automation

Permalink

https://escholarship.org/uc/item/33w2d55j

Authors

Shladover, Steven
VanderWerf, Joel
Miller, Mark A.
et al.

Publication Date

2001-08-01

Development and Performance Evaluation of AVCSS Deployment Sequences to Advance from Today's Driving Environment to Full Automation

Steven Shladover, Joel VanderWerf, Mark A. Miller, Natalia Kourjanskaia, Hariharan Krishnan

CALIFORNIA PARTNERS FOR ADVANCED TRANSIT AND HIGHWAYS

Development and Performance Evaluation of AVCSS Deployment Sequences to Advance from Today's Driving Environment to Full Automation

Steven Shladover

Joel VanderWerf

Mark A. Miller

Natalia Kourjanskaia

ACKNOWLEDGEMENTS

ABSTRACT

This report presents the findings of its investigation into deployment sequences to better understand the paths that could be taken from today's driving environment to vehicle-highway automation. One of the most vexing problems has always been that of determining how to advance from the present-day manually-controlled vehicles to the future fully automated vehicles. Considerable research attention has been devoted to defining the architecture and operating protocols, as well as the technology, of automated highway systems. Rather less attention has been devoted to defining the steps by which we can get there.

Initially, targets of opportunity were identified for accelerating progress toward highway automation, taking account of the operational constraints. Next, after reviewing existing literature on automated highway systems deployment, a set of principles to govern the design of deployment strategies is suggested followed by proposed deployment sequences for automated highway systems, beginning with adaptive cruise control and then adding elements of vehicle-vehicle cooperation and lane protection to build toward automated highway systems capabilities within constraints of technological, human factors and economic feasibility. A general deployment staging sequence is then presented along with example deployment "roadmaps" shown for transit buses, heavy trucks and light-duty passenger vehicles.

Finally, we discuss the findings of our modeling and evaluation work for the beginning stages of a specific deployment sequence for light-duty passenger vehicles in the setting of a single highway lane. This sequence incorporates the use of cooperative adaptive cruise control along with conventional or autonomous adaptive cruise control and manual-driven vehicles. The evaluation assesses the impact of each of these three operational driving modes on traffic flow dynamics and highway capacity as well as of increasing proportions of both autonomous and cooperative adaptive cruise control vehicles relative to manually driven vehicles. Such effects are difficult to estimate from field tests on highways because of their necessarily low market penetration of these vehicles. Our approach uses Monte Carlo simulations based on detailed modeling work to estimate the quantitative effects of varying proportions of vehicle control types on lane capacity and on queue lengths and wait times at on-ramps.

The results of this study can help to provide realistic estimates of the likely effects of the introduction of adaptive cruise control to the vehicle fleet, so that transportation system managers can recognize that the autonomous adaptive cruise control systems now entering the market are unlikely to have significant positive or negative effects on traffic flow. An additional value of studying such systems in this way is that these scenarios can represent the first steps in a deployment sequence leading to an automated highway system. Benefits gained at early stages in this sequence, particularly through the introduction of cooperative adaptive cruise control with priority access to designated (but not necessarily dedicated) lanes can help to provide support for further investment in and development of automated highway systems.

EXECUTIVE SUMMARY

This report constitutes the final deliverable for PATH Project MOU 366 — "Development and Performance Evaluation of AVCSS Deployment Sequences to Advance from Today's Driving Environment to Full Automation". The project has investigated the need for deployment sequences to better understand the paths that could be taken from today's driving environment to vehicle-highway automation. Throughout the history of thinking about highway automation, one of the most vexing problems has always been that of determining how to advance from the present-day manually-controlled vehicles to the future fully automated vehicles. Considerable research attention has been devoted to defining the architecture and operating protocols, as well as the technology, of automated highway systems. Rather less attention has been devoted to defining the steps by which we can get there.

Initially, targets of opportunity are identified for accelerating progress toward highway automation, taking account of the operational constraints that must be addressed. Next, we build upon such opportunity targets and address the most serious challenge to the credibility of highway automation as a potential solution to transportation problems, that is, the lack of convincing deployment strategies. Such a strategy is needed to show how to advance, step by step, from today's transportation system to a future system that includes automated highway systems. After reviewing existing literature on automated highway systems deployment we suggest a set of principles to govern the design of such deployment strategies. A deployment sequence is proposed for automated highway systems, beginning with adaptive cruise control and then adding elements of vehicle-vehicle cooperation and lane protection to build toward automated highway systems capabilities within constraints of technological, human factors and economic feasibility. Finally, some example deployment "roadmaps" are shown for transit buses, heavy trucks and light-duty passenger vehicles.

While vehicle type-specific deployment staging sequences for transit buses, trucks, and light-duty passenger vehicles have differences from each other because the economics and needs of these vehicle classes, they possess many commonalities and these are captured in the general deployment staging sequence "roadmap" shown in Figure ES-1. In each case, certain enabling technologies and AVCSS services may be available now or in the foreseeable future and can

serve as the foundations for introducing additional capabilities. Most of these are being developed as private industrial initiatives, although some additional developments are likely as a product of publicly sponsored programs. Based on the current philosophies and plans of the private industry and USDOT in particular, those building blocks are likely to be very heavily vehicle based rather than infrastructure based. Additional public investments should be targeted at key cooperative system technologies and infrastructure needs that fill the remaining gaps.

In Figure ES-1, we begin on the left with the autonomous safety warning and control assistance systems that are already under active development by the motor vehicle industry and its suppliers throughout the world. These systems will come to the market within the next few years, regardless of whether any public sector or coordinated public-private activities occur. However, advancing beyond these systems will require the active participation of public agencies, working in concert with the industrial developers of the vehicle-based technologies. The shaded blocks in the figure indicate the enabling elements that will in particular need public sponsorship and/or coordination.

Note that there are two parallel development paths near the start, and it is possible that different localities may choose to follow the separate paths. In some locations, vehicle-vehicle communication may be combined with ACC to provide cooperative ACC service anywhere in the highway network before the first protected lane is provided. In other locations, where the infrastructure providers are more proactive, a protected HOV (or truck-only) lane could be provided first, to provide a higher level of HOV (truck) service before vehicle-vehicle communication is available. Regardless of which sequence is followed, these represent important steps toward the first AHS service. If the cooperative ACC and protected lane are augmented with vehicle steering actuation (combined with the lane-tracking sensing function already developed for lane departure warning), there is a possibility for getting close to the first really automated operations on a protected lane. Adding some more intensive vehicle-roadside communication (VRC-2) for condition checking at entry, the means for automatically coordinating entering traffic with the traffic already in the lane, and some enhanced traffic management capabilities for integrating operations with the rest of the traffic system (ATMS+), we have the makings of the first single-lane automated highway system.

FIGURE ES-1 Generic AHS Deployment "Roadmap"

Reaching that first single-lane AHS is really the challenging milestone. Once that has been reached, the public and their political leaders should be able to see and experience the AHS benefits directly, which should stimulate the support needed to extend from that to a comprehensive AHS. That extension needs the addition of capabilities for automatic lane-changing control and control at the link and network levels. As the market penetration of AHS-capable vehicles increases and the utilization of the initial single AHS lane approaches capacity, it should become politically feasible to convert pre-existing mixed-flow lane(s) to AHS, providing the civil infrastructure for the complete multiple-lane AHS at modest incremental cost.

The next step in our investigation of AVCSS deployment sequences involved the modeling and evaluation of these sequences. Due to the significant amount of modeling and simulation work necessary, especially for the case of manually-driven vehicles, we were constrained to focus on only initial steps immediately beyond adaptive cruise control and analyzed the impacts of

cooperative adaptive cruise control in a mixed-lane traffic setting, i.e., with manually-driven vehicles and vehicles under AACC control in a single lane. We used light-duty passenger vehicles as the vehicle type for this analysis.

A set of mathematical models are defined to predict the effects that emerging driver control assistance systems such as adaptive cruise control can have on traffic flow dynamics and capacity. It is important to understand these effects in order to ensure that adaptive cruise control systems are implemented in ways that improve, rather than degrade, traffic conditions. Existing traffic models were not designed for, and are not suitable for, this purpose, so it has been necessary to develop a new family of simulation models incorporating the key elements of driver behavior and control system design that will affect traffic flow dynamics and capacity.

Initially we conducted a review of the literature in the area of adaptive cruise control that have attempted to predict the effects that the introduction of adaptive cruise control vehicles into the traffic stream would have on overall traffic behavior including congestion, safety, emissions, and fuel consumption. Past work has been based on analytical models and simulations since there are not yet sufficient adaptive cruise control vehicles in use on the roadway to perform direct experimental evaluations. Much of this past literature has been deficient in its attempts to completely and accurately describe the impacts of adaptive cruise control on traffic, for reasons including 1. a lack of a realistic model for human car-following behavior and thus no evaluation of the dynamic interactions between adaptive cruise control and manually-driven vehicles, 2. dynamics of human responses not fully understood or explained, 3. results specific to non-U.S. highways, driving behavior, and traffic laws and so may not be totally applicable to the U.S. case, 4. simulation unrepresentative of real traffic with its variety of disturbances, and 5. no prediction of traffic flow volume achievable with adaptive cruise control.

Models were developed, validated, and then used to assess the impact on traffic flow volumes. Such models included 1. vehicle-following logic for each operational mode, i.e., human driving, autonomous adaptive cruise control, and cooperative adaptive cruise control, 2. merging of vehicles entering a limited-access highway from an on-ramp, 3. "free-driving" of a vehicle in

uncongested traffic or with no immediately preceding vehicle to follow, and 4. vehicle dynamic responses to speed change commands.

These models were tested initially for simple vehicle pairs, then for sequences of 20 vehicles with the same type of vehicle-following controller, and finally more complicated cases involving mixtures of controller types and larger numbers of vehicles, with vehicles entering and leaving the flow at on- and off-ramps, respectively. Testing for model validation consisted of showing how a group of 20 vehicles each with the same controller (human, ACC or CACC) responds to severe disturbances. The disturbances are generated by a lead vehicle that follows a velocity trajectory taken from severely congested stop-and-go traffic on I-880 in the San Francisco Bay Area. We allow the following vehicles to follow each other stably before injecting the disturbance. For the human driver model, we see the kind of shock wave disturbance expected in normal traffic with larger disturbances at later times for vehicles further behind the lead vehicle. For the adaptive cruise control case, disturbances are reduced the further out we get from the lead vehicle. While responses are delayed in time, the amplification of disturbances experienced with the human driver control law is not encountered here. Thus reduction in shock wave propagation illustrates one potential improvement to traffic dynamics from use of adaptive cruise control. For the cooperative adaptive cruise control case, the tighter control of this controller manifests itself with a closer match to the velocity trajectory of the lead vehicle from following vehicles as well as a reduction in response delay from one vehicle to another. Therefore, all simulation model validation test cases produced reasonable results.

Results have been shown for the validation cases used to test the models individually, for the capacity estimates for the 100% market penetration cases for each of the three modes of operation: manually driven vehicles, AACC, and CACC, and for the capacity impacts of different combination of market penetrations for AACC and CACC mixed with manually-driven vehicles. For the three 100% market penetration cases, nominal capacity estimates for the manual driving, AACC, and CACC cases were, respectively: 2,050, 2,200, and 4,550 vehicles per hour. The estimate for manually driven vehicles is consistent with current estimates for such driving. Even under the most favorable conditions, with ideal ACC system design and performance, it appears that autonomous ACC can only have a small impact on highway

capacity. Assuming that average ACC users choose a mid-range time gap of 1.4 seconds, highway capacity can be increased by at most 9.5% when the market penetration of autonomous ACC is in the 40% to 80% range. Moreover, diminishing returns set in quickly with respect to the capacity increases from introducing AACC into the traffic stream. Increases in the market penetration of autonomous ACC above 80% can lead to a modest loss of highway capacity, based on ACC users choosing an average time gap for ACC that is somewhat longer than the time gap they use when driving manually. Because of the modest effects of AACC on highway capacity, there does not appear to be any justification for providing AACC vehicles with priority access to special lanes such as HOV lanes. Cooperative ACC systems, using vehicle-vehicle communications to enable closer vehicle following (down to a time gap of 0.5 second), have the potential to produce significant highway capacity increases. Cooperative ACC can represent an important step in a progressive deployment strategy to lead toward highway automation, because it can potentially double the capacity of a highway lane at a high market penetration. The capacity effect is very sensitive to market penetration, which means that it is important to gather as high a proportion as possible of CACC vehicles into the same lane. This provides a strong justification for giving priority access to a special lane for CACC vehicles.

Much more work remains to be done before this can be turned into an action plan. More quantitative studies of the costs and benefits associated with the different stages are needed. The results of this and further analysis will then need to be explored with representatives of the relevant stakeholder groups for verification or refutation. The real action can come only after the stakeholders have been convinced of the genuine benefits that they will gain from each significant step on the road toward AHS.

TABLE OF CONTENTS

1.0   INTRODUCTION

This report constitutes the final deliverable for PATH Project MOU 366 — "Development and Performance Evaluation of AVCSS Deployment Sequences to Advance from Today's Driving Environment to Full Automation". This report consists of our investigation into the need for deployment sequences to better understand the paths that could be taken from today's driving environment to vehicle-highway automation and the modeling and evaluation of such sequences.

The need or a progressive deployment strategy for AHS was recognized by the National Automated Highway System Consortium (NAHSC), but the NAHSC was not able to satisfy that need by the time its work was prematurely terminated by the USDOT. The NAHSC struggled with attempts to define a roadmap toward deployment of automated highway systems (AHS) incorporating a diverse collection of feasible and infeasible elements, and the limitations of its deployment plans were among the primary targets for criticism by the Transportation Research Board (TRB) study committee that reviewed its program (1). This shortcoming of the NAHSC work was one of the principal criticisms by the special TRB committee. Strategies for AHS deployment have been addressed in a variety of papers published in recent years, the limitations of which will be discussed later in this report.

In Japan, under the leadership of the Ministry of Construction, AHSRA has defined a very general framework for the deployment staging of its AHS ("Advanced Cruise-Assist Highway System") program, with successive stages called AHS-i (information), AHS-c (control assistance) and AHS-a (automation) (2). However, this framework does not appear to have been specified in much depth, with no indication of the extent to which these three stages of advanced vehicle control and safety systems (AVCSS) may overlap in time or differ in their application to the different classes of road vehicles.

Research conducted over the past decade has provided strong indications that automated highway systems (AHS) can offer very significant benefits to transportation system operations (3). Principal among these is the potentially very large capacity increase per lane that can be gained when all vehicles in that lane are fully automated. Many of the technical impediments to highway automation have been addressed successfully, and many more are continuing to receive

active research attention. However, one major impediment has remained largely intractable – the absence of a clear strategy or strategies for advancing from today's manually-controlled road vehicle system to the future system in which traffic on some portions of the road network could be fully automated.

Adaptive cruise control (ACC) systems are the first driver control assistance systems to enter the passenger car and heavy-duty truck marketplace having the potential to influence traffic flow characteristics. Such systems can also be viewed as a first step toward future automation of the complete driving function and are thus essential as part of the overall deployment equation and it is important to understand how ACC system designs and operating concepts can be developed to best support such future advances. To evaluate interactions between ACC vehicles and conventionally driven vehicles, it is necessary to use sophisticated and sensitive models of each having sufficient fidelity to capture their sometimes subtle interactions. While it is not practical now to evaluate these interactions in full-scale experiments because of the need for a large ACC vehicle fleet, in the longer term, it would, however, be valuable to perform these experiments. As a first step in developing and evaluating deployment sequences leading toward AHS, this report describes a new set of models for evaluating large-scale interactions occurring when varying proportions of ACC vehicles are mixed with conventionally driven vehicles in a limited-access highway operating environment. To address the longer-term advances in ACC capabilities, models are developed for higher-capability cooperative ACC systems (CACC) as well as the first-generation autonomous ACC.

2.0   DEPLOYMENT OF VEHICLE-HIGHWAY AUTOMATION SYSTEMS: OPPORTUNITIES AND CONSTRAINTS

The central challenge for AHS deployment staging remains the "chicken and egg" problem of which comes first – the vehicle technology or the roadway infrastructure technology? Of course, the real answer is neither, because they need to come in coordination with each other rather than in sequence, and the stakeholders on both the vehicle and infrastructure sides need to work together to coordinate their development work and investments. The technical and institutional considerations, which tend to strongly favor opposing approaches to deployment staging, need to be balanced with each other in seeking a feasible deployment strategy.

2.1   Evolution or Revolution?

In recent years, the deployment staging discussion has often been framed in contrasting terms of "evolution or revolution"(4,5). This false dichotomy forces people to choose between two extreme solutions, neither of which is likely to be successful. The "evolutionary" approach assumes that all the needed technologies will gradually be added to vehicles, which will have steadily increasing degrees of automation until eventually they will be fully automated. This approach cannot work because of the insuperable technical problems associated with automating vehicles that must coexist freely with manually driven vehicles and the human factors constraints that preclude drivers from being effective supervisory controllers of "mostly automated" vehicles (6). The "revolutionary" approach assumes that a completely new and large-scale transportation system will be implemented alongside today's system within a short period of time. Although this approach is easier to accomplish technologically, it faces difficult economic, political and social challenges. In all likelihood, a hybrid approach combining elements of both the "evolutionary" and "revolutionary" approaches will be needed. The underlying technologies in sensors, communications and actuators will gradually be increasing on vehicles, for reasons largely unrelated to the needs of vehicle automation. These can provide some of the building blocks for automation, and can significantly reduce the incremental costs associated with automation, but they are far from sufficient. At the same time, roadway infrastructure modernizations, expansions and additions will continue to be made, even if not at the rate they were made in the 1960s. With technological enhancements of modest cost, these can provide the infrastructure building blocks to help advance toward some automated lanes of limited scope. The core of the deployment staging planning process is identifying the targets of opportunity for combining these vehicle and infrastructure improvements in ways that are mutually supportive.

2.2   Deployment Targets Of Opportunity

Progress toward an automated highway system is not likely to "just happen", but has to be based on deliberately capitalizing on the targets of opportunity that may be available. These may be based on the needs of specialized fleets of vehicles, geographic locations with specific transportation problems and needs, or unique political or technological opportunities, such as:

- Applications where the vehicles and roadway infrastructure are under common ownership and operation (HOV lanes for buses and vanpools, freight movements at terminals and ports, highway maintenance vehicles, public service fleets);

- Special bottlenecks where conventional infrastructure expansions would be particularly costly and difficult (bridges, tunnels);

- Commute corridors with high congestion and some available right of way;

- Specialized vehicle fleets where there is an opportunity to subsidize some incremental vehicle costs (vanpools, public service fleets);

- Public officials interested in promoting transportation and high technology issues;

- Applications where substantial transportation infrastructure investments are already planned and the incremental costs of adding AHS capabilities are therefore modest;

- Locations where the roadway design permits sufficiently tight access controls that it may be possible to carefully limit the interactions between automated vehicles and vehicles that are manually driven or partially automated.

In each case, it is necessary for the vehicle and infrastructure stakeholders to be interested in gaining the benefits of highway automation, willing and able to work with each other, and able to perceive tangible rewards that exceed their risks. The risks can be minimized if the deployment can be staged in steps, each of which clearly has positive net benefits for each decision maker, so that they do not need to be worried about whether they can achieve a more uncertain later step in order to recover their investments.

2.3 Operational Constraints To Deployment Of AVCSS/AHS Capabilities

The three key operational issues that determine the difficulty of implementing automation capabilities on vehicles and that must be considered carefully when designing automation deployment sequences are:

a. Physical protection of the driving environment

b. Minimum capabilities of vehicles that are permitted to coexist with automated vehicles in that environment

c. Driver role(s) in control of the automated or partially automated vehicles.

If these characteristics are most severely constrained, the technical problems of implementing automation become simpler, but if these characteristics are left relatively unconstrained those technical problems can become insuperable. However, the most constrained cases also present challenges in the economic, operational and political domains, which must be traded off against the technical challenges.

2.3.1 Physical Protection Of The Driving Environment

Beginning with the simplest case, and then proceeding to more complicated cases, we have:

1. Single lane, protected by barriers and fences – This represents the highest level of protection for the automated vehicles, preventing the intrusion of foreign objects, unequipped vehicles, animals and pedestrians. It also prevents other vehicles from passing or cutting in front of the automated vehicles. With this level of protection, automated operations could be based on use of vehicle-follower longitudinal control and lane-following lateral control, without the need for special obstacle detection capabilities.

2. Single lane plus shoulder, protected by barriers and fences – This provides almost as much as protection as the previous case, but the shoulder could provide an opportunity for an aggressive driver to "weave" between automated vehicles, introducing the possibility of side-swipe or cut-in hazards. This would impose significantly increased sensing requirements on the automated vehicles, but in compensation would provide better access to failed vehicles by emergency response crews as well as an opportunity for bypassing a failed vehicle without halting all traffic.

3. Multiple lanes, protected by barriers and fences – This introduces the need to implement lane-changing automation, as well as protection against more frequent sideswipe or cut-in hazard conditions.

4. Elimination of fences on the barriers – In any of the above cases, eliminating the fences from the top of the barriers substantially reduces the protection from road debris and animal and pedestrian intrusions. This would impose significant obstacle detection capabilities on the automated vehicles, which appear to substantially exceed the state of the art in sensing technology.

5. No physical protection – This imposes the least infrastructure constraint, but represents

the maximum complexity of operating environment for the technology.

2.3.2   Minimum Capabilities Of Vehicles That Can Coexist With Automated Vehicles

Beginning once again with the simplest case and proceeding to the more complicated cases:

1. Other fully automated vehicles – This is the easiest to achieve technically, because all vehicles would then have comparable, predictable, performance and full communication capabilities.

2. Vehicles with two-way data communications – These vehicles can exchange information about their current behavior and possibly also their future intentions, whether they be manually or automatically driven. This facilitates coordination and anticipation of problems, but leaves the unpredictability of driver behavior as a remaining challenge.

3. Vehicles with transponders or tags – There are several possible levels of capability here, but the key element is that they can respond to sensors or communication devices on the automated vehicle in a distinctive fashion (identifying themselves as vehicles, perhaps also by type or performance class, and at higher levels with some information about current behaviors such as speed, acceleration, or failure condition flags). The information that they are able to provide reduces the uncertainties about the driving environment for the equipped vehicle.

4. Vehicles with adaptive cruise control (ACC) – These vehicles can provide a somewhat more predictable longitudinal behavior than unequipped manually-driven vehicles. That improvement is relatively small, but it could be increased in the case of vehicles with cooperative ACC, combining the basic ACC with inter-vehicle communications for coordination of maneuvering.

5. Any vehicle – This is the most difficult case because it provides no assistance to the automated vehicle and leaves the driving environment as unstructured and unpredictable as any public road today. This is also the easiest case from the political and institutional perspective because it requires no limitations on the movements of any vehicles.

A promising approach for unlocking the vehicle/infrastructure "chicken/egg" conundrum is to focus on combinations of driving environment restrictions and vehicle capability-based access restrictions that minimize safety concerns. Both of these restrictions are intended to reduce the complexity of the driving environment encountered by the automated systems. As the driving

environment becomes more complex, it becomes increasingly difficult to design and implement the automation system, which also means that less of the driving tasks can be transferred from the driver to the system without jeopardizing safety.

For example, since it is likely to be politically difficult to restrict a protected lane to fully automated vehicles when the population of such vehicles is extremely small, those lanes could perhaps also be made available for use by manually-driven vehicles equipped with suitable communication devices so that the automated vehicles would be able to communicate with them. Similarly, if a fully protected lane is not initially achievable in some locations, perhaps usage of the partly protected lane(s) would be restricted to vehicles having intermediate capabilities, but neither fully automated nor totally unequipped. The feasibility of these intermediate solutions must be considered on the basis of the third key characteristic, the role(s) that the driver plays in control of the vehicle.

2.3.3   Driver Role(s) In Control of the Automated or Partially Automated Vehicles

Considerations of human factors limit the range of feasible driver roles. The roles of the driver and the automation system will need to be defined so that, when combined, all of the essential safety-critical functions are performed at least as well as they are today. The roles of the driver and the automation system would then be the complements of each other. Obviously, the current condition of the driver maintaining complete responsibility for the control of the vehicle remains as a viable alternative. However, in addition to this, it is also possible to have:

1. Adaptive Cruise Control – Automatic control of vehicle speed and spacing to forward vehicle(s) in well-structured traffic flow, while driver retains steering control, hazard detection and response and stop-and-go control responsibilities.
2. Driver control assisted with hazard warnings – A variety of hazard warning systems currently under development could give the driver auditory, haptic, kinesthetic, or possibly visual cues about hazards in the driving environment, and the driver would retain responsibility for taking corrective actions.
3. Automated driving with driver responding only to take-over warnings – Under normal

conditions, the vehicle would be driven automatically. However, when the automated system encountered a failure or a driving condition beyond its capabilities, it would alert the driver to the need to intervene and would then cede control of the vehicle to the driver as soon as he interacts with the steering wheel or pedals. (The inverse case of a system intervening to protect the driver when he is controlling the vehicle [a kind of automated "guardian angel"] poses such daunting technical challenges that it should not be regarded as a realistic precursor to completely automated driving.)

4. Completely automated driving – The system would be designed with sufficient fault detection and malfunction management to handle all driving conditions, including bringing the vehicle to a safe stop as the ultimate fall-back condition. In this case, the driver would have no role and would not need to remain conscious during vehicle operations. This is certainly the most technically challenging to implement, while providing the maximum of task relief to the driver.

Note that this listing does not include any "partial" or "shared" control cases between ACC and fully automated driving. That is because of concerns about the difficulty of designing such systems to be easy to learn and safe to implement. These intermediate cases encounter two fundamental problems:

- Not being able to ensure that the driver has a correct mental model of the system's capabilities and limitations;
- Not leaving the driver with a sufficient workload to retain the attentiveness necessary to quickly detect any and all hazardous conditions that the system is not designed to handle (5).

2.4   Deployment Staging "Roadmap"

Based on the considerations that have been discussed already, hypothetical deployment staging sequences have been developed for transit buses, freight vehicles (primarily trucks) and light-duty passenger vehicles. These sequences are somewhat different from each other because the economics and needs of these classes of vehicles differ considerably, but their common essence is captured in the general "roadmap" of Figure 1 (Vehicle-specific sequences are presented in Chapter 3). In each case, certain enabling technologies and AVCSS services may be

available now or in the foreseeable future and can serve as the foundations for introducing additional capabilities. Most of these are being developed as private industrial initiatives, although some additional developments are likely as a product of publicly sponsored programs such as the USDOT's Intelligent Vehicle Initiative (IVI) in the United States or the ASV or AHSRA programs in Japan. Based on the current philosophies and plans of the private industry and USDOT in particular, those building blocks are likely to be very heavily vehicle based rather than infrastructure based. Additional public investments should be targeted at key cooperative system technologies and infrastructure needs that fill the remaining gaps.

At the left of Figure 1, we begin with the autonomous safety warning and control assistance systems that are already under active development by the motor vehicle industry and its suppliers throughout the world. These systems will come to the market within the next few years, regardless of whether any public sector or coordinated public-private activities occur.

However, advancing beyond these systems will require the active participation of public agencies, working in concert with the industrial developers of the vehicle-based technologies. The shaded blocks in Figure 1 indicate the enabling elements that will in particular need public sponsorship and/or coordination.

Note that there are two parallel development paths near the start, and it is possible that different localities may choose to follow the separate paths. In some locations, vehicle-vehicle communication may be combined with ACC to provide cooperative ACC service anywhere in the highway network before the first protected lane is provided. In other locations, where the infrastructure providers are more proactive, a protected HOV (or truck-only) lane could be provided first, to provide a higher level of HOV (truck) service before vehicle-vehicle communication is available. Regardless of which sequence is followed, these represent important steps toward the first AHS service. If the cooperative ACC and protected lane are augmented with vehicle steering actuation (combined with the lane-tracking sensing function already developed for lane departure warning), there is a possibility for getting close to the first really automated operations on a protected lane. Adding some more intensive vehicle-roadside

FIGURE 1 Generic AHS deployment "Roadmap"

communication (VRC-2) for condition checking at entry, the means for automatically coordinating entering traffic with the traffic already in the lane, and some enhanced traffic management capabilities for integrating operations with the rest of the traffic system (ATMS+), we have the makings of the first single-lane automated highway system.

Reaching that first single-lane AHS is really the challenging milestone. Once that has been reached, the public and their political leaders should be able to see and experience the AHS benefits directly, which should stimulate the support needed to extend from that to a comprehensive AHS. That extension needs the addition of capabilities for automatic lane-changing control and control at the link and network levels. As the market penetration of AHS-capable vehicles increases and the utilization of the initial single AHS lane approaches capacity,

it should become politically feasible to convert pre-existing mixed-flow lane(s) to AHS, providing the civil infrastructure for the complete multiple-lane AHS at modest incremental cost.

3.0 PROGRESSIVE DEPLOYMENT STRATEGIES LEADING TOWARD AN AUTOMATED HIGHWAY SYSTEM (AHS)

This chapter builds on the foundation established in chapter 2, and extends that work in more depth and detail. Following a review of the existing literature on AHS deployment, the first focus is on definition of some basic principles that should be applied to the synthesis of deployment strategies. Next, the three operational constraints on deployment of AHS that were previously introduced are used as the basis for synthesizing operational concepts for some "pre-AHS" (intermediate-term) systems. The incremental benefits that can be gained at each step are then explored, to try to ensure that they are sufficient to justify the investments needed to attain that step. Finally, deployment "roadmaps" are hypothesized for transit buses, heavy trucks and passenger cars to provide some tangible examples to illustrate the similarities and differences in how the use of automation technology can progress for each of these types of road vehicles.

3.1 Prior Literature on AHS Deployment

In the early literature on AHS, little attention was devoted to defining a sequence of steps by which one could advance to AHS. This should not be too surprising, because until the 1960s, and in most cases the 1970s, it was still relatively easy to build entirely new highways. That was the era of major Interstate highway construction in the United States, and in that environment it was easy to conceptualize the construction of another completely new (automated) highway. However, in subsequent years the economic, political and environmental landscape has changed to the extent that construction of a completely new highway would be difficult to envision in virtually any major urban area in the U.S. Attention, instead, needs to be focused on expansions of existing highways that can be accomplished largely within existing rights of way.

Given the difficulty of constructing completely new highways for AHS, some investigators became interested in the concept of operating automated vehicles freely mixed within the conventional traffic flow on existing highways (7,8). If this concept were technically feasible, the infrastructure deployment challenges could be ignored and all attention could be focused on the vehicle. However, the technical challenges that must be surmounted in order to make such a

concept safer than today's highway travel are so overwhelming that this is probably not a realistic target within the dawning century.

The amount of research attention focused on AHS in the U.S. began to increase substantially in 1993, motivated by the USDOT sponsorship of the AHS "Precursor Systems Analysis" projects. These projects, combined with the already-existing PATH program of AHS research, and the subsequent work of the NAHSC, led to the publication of several papers addressing AHS deployment staging. Tsao (4,9) identified a set of challenges to AHS deployment and then proposed a detailed sequence of stages toward AHS, beginning with an automated shuttle van service that would operate in mixed traffic under the supervision of a driver, who would be responsible for handling failures or conditions that exceed the capabilities of the automation system. Unfortunately, even this initial stage is unlikely to be feasible because of the severely limited ability of humans to execute this type of supervisory control (5). Several additional treatments of the deployment issues appeared in the book about AHS edited by Ioannou (3), which either fell prey to the same misconception about human capabilities or else blithely assumed away the technical challenges to automating driving in the extremely complicated mixed traffic environment.

Two different systematic frameworks for thinking about AHS deployment problems were proposed by Tsao (10) and Koopman and Bayouth (11). Tsao (10) chose an "axiomatic" approach based on explicitly stating assumptions and then deriving inferences and requirements from those assumptions. The difficulty here was in the lack of a sufficiently robust, universally applicable, set of assumptions that could provide a solid enough foundation for the subsequent conclusions. Koopman and Bayouth (11), in contrast, defined a set of four "capability building blocks" that they considered to be orthogonal to each other and then defined the 16 combinations of these as the map of possible deployment options. However, the building blocks that they selected were not comprehensive (not spanning the full range of useful alternatives), nor were they truly orthogonal and independent of each other. Both of these papers showed the limitations of trying to distill the AHS deployment problem down to an equivalent logical or mathematical problem for which an algorithmic solution can be found.

In a later paper, Bayouth and Koopman (12) approached the deployment problem from a functional perspective, defining the technical capabilities that could be combined to lead to ever-increasing levels of vehicle capabilities. This provided some useful insights into deployment issues, although the authors unfortunately chose to assign priority to autonomous in-vehicle functions, treating inter-vehicle communications and infrastructure support as later additions rather than integral elements in the improvement of system capabilities. Coming from a different direction, McKendree (13) tried to estimate the time-scales for development and growth of AHS and for an alternative concept in which automated vehicles would mix freely in the conventional highway traffic stream. This used simple system dynamic models to represent many assumptions about conscious deployment decisions for AHS and about the general advance of technology for the automated mixed traffic vehicles. With different choices of numerical values to represent those largely arbitrary assumptions, substantially different results could have been derived.

3.2   Principles to Govern AHS Deployment Strategies

Successful AHS deployment will require a set of favorable decisions by a host of stakeholders in both the public and private sectors. Because AHS deployment requires investments in both roadway infrastructure and vehicles, the essential stakeholders include a wide range of players, beginning with state, regional and local governments, then crossing over to vehicle OEMs and their suppliers, and finally to the purchasers and users of vehicles (not only automobiles, but also buses and trucks), as well as various other special interests such as the insurance industry, environmentalists, and the legal community. Given this complexity, a successful AHS deployment strategy is likely to need to adhere to principles such as these:

a.   Define a deployment staging "roadmap" in which each development stage produces benefits that will exceed its costs for all of the key decision-makers so that they will have continuing positive interest in investing in the program.

b.   Seek targets of opportunity to combine vehicle and infrastructure elements and services to produce added value that would not otherwise be achievable.

c.   Seek synergies across "vehicle platforms", particularly in cases in which services for one platform can be used to establish an infrastructure framework and elements that can facilitate the advance of services on other platforms having larger vehicle populations.

d.   Build on the work that has been and is still being done on pieces of the system (industry AVCSS product developments, prior PATH and NAHSC research, ATMS and ATIS services, IVI, etc.) and avoid duplicating those efforts, so that resources are used most efficiently.

e.   Encourage commercial spin-off products at every stage, to sustain the interest of private sector participants.

f.   Apply a systems approach, so that all systems and services can fit together compatibly within an architectural framework that should be developed as early as possible, but should also be sufficiently flexible to adapt as technologies mature.

g.   Maintain the visibility of the goal of developing a truly automated highway system on dedicated lanes (in order to be able to gain the capacity benefits), and try to ensure that activities contribute toward reaching that goal rather than diverging on tangents.

The first two of these principles are particularly influential in the design of a deployment staging sequence. The dominant consideration for the staging sequence must be that the incremental benefits at each stage must exceed the perceived costs and risks required to reach that stage. If this does not hold, then some or all of the stakeholders are likely to hesitate to move forward to that stage, because they cannot base their plans on uncertain benefits further in the future.

Although not explicitly stated in the above list, it is also essential that each stage be feasible in technical and human factors terms. Some of the deployment-staging ideas that have been proposed in the past have been based on some very dubious assumptions about human and technological capabilities. The feasibility considerations need to be based on proper treatment of the three key operational issues that were described in Chapter 2:

i.   Physical protection of the driving environment from intrusions of debris, animals, pedestrians or other vehicles.

ii.   Minimum capabilities of the "other" vehicles that are permitted to coexist with the automated vehicles.

iii.   Driver role(s) in control of the automated or partially automated vehicles.

3.3 Operational Concepts for "Pre-AHS" Systems and Their Benefits

Although much research attention has been devoted to the definition of operational concepts for (fully automated) AHS, rather less attention has been devoted to the definition of the systems that might precede it in deployment. The U.S. National Architecture for ITS identifies a series of advanced vehicle control and safety system (AVCSS) user services (14). Similarly, the AHSRA program in Japan has defined its set of analogous user services (15). Under the auspices of the NAHSC, an exhaustive listing of possible "pre-AHS" and AHS user services was also defined (16). While these listings of possible services can serve as a useful starting point, they fall short of identifying the crucial sequencing by which we can advance from today's entirely manual driving to the future fully automated driving (AHS).

The first steps toward AHS are relatively easy to identify, because these are autonomous systems that are already under development by vehicle OEMs and their suppliers throughout the world: adaptive cruise control (ACC), forward collision warning, and road (lane) departure warning systems. All of these will be on the market within the next few years, regardless of any action or inaction on the part of public authorities. The ACC will enhance driving comfort and convenience, while the two warning systems should provide some safety benefits. These systems may later be joined by side collision ("blind spot") warning systems as well. The warning systems could possibly be extended from audible or visual warning to haptic feedback to drivers (through pressure on pedals or torques applied to the steering wheel).

The steps that might follow these initial systems are not so easy to define. Some have suggested that lane departure warning might be extended to lane tracking control, but this seems to be an extremely dubious step because of its implications for driver alertness. If the driver is relieved of responsibility for steering to track the lane, there is a serious likelihood of the driver losing overall attentiveness to the driving task and to the hazards that could arise (5). Moreover, the availability of automatic steering control would likely tempt drivers to combine that with conventional or adaptive cruise control, for a simulacrum of fully automated driving. The safety implications of this could be devastating, because the seemingly automated vehicle would in fact have no capability to deal with driving hazards such as obstacles in the roadway, aggressive maneuvering by other vehicles, poorly marked roads, etc., etc. One cannot expect the typical

driver to appreciate all the limitations of the new technologies on his vehicle, particularly in the context of his own limitations of attentiveness when relieved of much of the driving task.

In trying to advance toward AHS, it appears to be more productive to build on the longitudinal control capabilities of ACC and try to provide benefits beyond the comfort and convenience offered by the basic autonomous ACC. Additional benefits could be associated with safety improvements or improvements to lane capacity (translating into congestion relief and time savings for travelers). These benefits could be gained by augmenting the ACC with vehicle-vehicle and perhaps vehicle-roadside communication (leading to cooperative ACC – CACC) and with access to special protected highway lanes.

3.3.1   Cooperative ACC (CACC)

Vehicle-vehicle communication can provide an ACC system with more and better information about the vehicle it is following. This could include precise speed information, acceleration, fault warnings, warnings of forward hazards, maximum braking capability, and current braking capability (based on local tire/road coefficient of friction). With information of this type, the ACC controller can better anticipate problems, enabling it to be safer, smoother and more "natural" in response, and potentially to operate at a smaller time gap (i.e., closer to the preceding vehicle).

The "communications" could be as primitive as equipping the back of the preceding vehicle with a tag that could be clearly distinguished by the ACC radar, or it could extend to a complete two-way wireless communication system. The tag (designed to reflect the incident radiation from the ACC radar) could be implemented on all vehicles relatively quickly and cheaply by incorporating it in the annual license plate renewal sticker.

The two-way communication system would involve more time and expense, particularly if it needs to have access to the outputs of various sensors and subsystems throughout the vehicle. The incremental expense of capabilities such as these will inevitably decrease over time as more vehicles are equipped with in-vehicle networks such as CAN (Controller Area Network) and sensors for collision warning and improved active braking control. However, it is not reasonable

to expect vehicle owners to select an option such as the two-way communication system unless they, themselves, receive a clear and direct benefit from it (i.e., beyond serving as a better target for somebody else's CACC system). These benefits could be associated with traveler information systems, in-vehicle multimedia systems or perhaps with enhanced access to special highway lanes (as described below).

3.3.2   Protected Highway Lanes

A variety of levels of "protection" can be applied to highway lanes to restrict access and minimize hazards. Higher levels of protection can support higher levels of system performance, but they also tend to cost more and be more difficult to implement politically. Some possible levels of protection, ranging from least effective to most effective, include:

a.   "virtual barriers" using paint stripes or other forms of delineation – these are used to mark HOV lanes in many places, and represent a legal boundary beyond which only certain vehicles are authorized to operate. However, they provide no physical protection against scofflaw drivers' vehicles, debris or animals.

b.   Concrete "Jersey barriers" – these are used to separate HOV lanes in some locations, and generally provide protection against inadvertent vehicle intrusions and some debris intrusions. However, they are sometimes breached by vehicles, and can be violated easily by crash debris and animals.

c.   Concrete barriers with fences on top – these can protect against the large majority of vehicle, debris and animal intrusions. If the "protected" facility is wider than a single lane width, it is still possible for vehicles within that facility to have adverse interactions if drivers overtake or change lanes too aggressively or sideswipe other vehicles.

d.   Single lane behind concrete barriers with fences – this provides the maximum protection for the vehicles, producing the simplest and most predictable driving environment.

As the level of protection increases, the hazards to which the vehicles are exposed decrease and it becomes possible to place more reliance on automated system capabilities. The reduction in hazards should make it possible to reduce the need for driver vigilance, enabling vehicles to

operate closer together without compromising safety, which, in turn, is where improvements in lane throughput and travel time savings can become evident.

The pacing factor for the deployment of AHS is the availability of sufficient protected lanes. In those locations where space is available to build new protected lanes or where the political environment is favorable to the addition of roadway infrastructure, the AHS facilities could be provided readily and the stimulus would thereby be applied for encouraging the growth of the AHS-capable vehicle population. In cases where the "empty lane syndrome" creates political difficulties while there are still relatively few AHS-capable vehicles, it may be necessary to go through many of the intermediate steps outlined below. However, most of those steps can be skipped where it is politically feasible to deploy an AHS lane before all the vehicles are available. This could be done, for example, in the form of a busway that is only open to AHS-equipped buses at the start, but is then made available to suitably equipped vanpools and carpools.

3.3.3 Potential Steps Beyond ACC

Building on the foregoing elements, some steps that could be devised to lead from the first generation of ACC and collision warning systems toward AHS include:

1. Cooperative ACC in mixed traffic – The cooperative elements provide better ability to identify and track equipped target vehicles, improving the comfort and convenience of ACC and its ability to distinguish stopped vehicles (bringing it closer to having a true forward collision warning capability). The improved performance can only be achieved when the vehicle ahead is equipped with the cooperative tag or communication system, so the increase in effectiveness of this system is strictly limited by the market penetration of the cooperative elements on other vehicles.

2. Cooperative ACC in a *separated lane*, mixed with both cooperative and uncooperative vehicles – If the cooperative ACC vehicles and cooperating (communication-enabled) vehicles concentrate in a separated lane, rather than being spread across all highway lanes, the probability of being able to use the cooperative capability is greatly increased. This could initially be done by providing these vehicles access to under-utilized HOV

lanes. At this stage, even the "virtual barrier" is an adequate separation to demarcate the protected lane, because the only purpose of the lane restriction is to increase the percentage of cooperating vehicles. The user benefits will be relatively small (improved ACC performance and access to a previously-inaccessible lane), but the costs will also be negligible.

3. Cooperative ACC in a separated lane, *restricted to other cooperating vehicles* – As the *local* market penetration of cooperative (communication-enabled) vehicles increases, it becomes politically feasible to dedicate a lane to their use, prohibiting non-cooperative vehicles from this lane. On a freeway with four lanes in each direction, this could correspond to a local market penetration of 25% (or even less if one assumes that the equipped vehicles are used by longer-distance commuters and therefore might be driven more than non-equipped vehicles). It is important to emphasize the local nature of this market penetration condition. There is no need to wait for the national or even the regional market penetration to reach a threshold value in order for this to become practical. Rather, it is the market penetration of vehicles using the specific highway facility that should be considered. The marketing of equipped vehicles could be concentrated in the locations with the greatest need for, as well as ability to pay for, the AHS-precursor capabilities. Once all vehicles in the lane are in communication with each other, they can operate somewhat more safely and smoothly, and possibly at a smaller average spacing, but the effect of this on lane capacity should be minor.

4. Cooperative ACC and other cooperating vehicles in a *protected*, separated lane – By providing physical protection (barriers topped by fences) to the separated lane, the safety of operations should be increased, particularly by excluding road debris and crashing or scofflaw vehicles, and the level of stress on the drivers should be reduced.

5. Protected lane *restricted* to cooperative ACC vehicles – As the local market penetration of cooperative ACC vehicles continues to increase, the protected lane can be restricted to vehicles with that full capability, so that longitudinal control behavior is consistent and predictable. Safety, smoothness of driving (via elimination of shock waves), and some consequent small energy/environmental improvements should be the additional benefits gained at this stage. Recent research has shown that proper design of the vehicle-follower

control law can produce asymptotically stable traffic flow, without shock waves, by eliminating the driver behavior that causes shock waves (17).

6.   Addition of *automatic steering control* to CACC in *single* fenced, protected lane – This is the step when the driver's role changes most significantly to that of an observer or possibly a general supervisor rather than a direct participant in the driving task. At this point, the driver can no longer be depended upon to identify hazards or failures because his attentiveness cannot be assured. Therefore, the driving environment within the protected lane needs to be made as safe and predictable as possible, through physical protection (barriers and fences) and carefully limited interactions between vehicles (cooperative longitudinal motions only, no space for passing or sideswiping maneuvers). This step should probably be taken initially on a "pipeline" road facility with no intermediate entry or exit points in order to avoid the complications introduced by entry and exit maneuvers. It could be seen as the first truly automated driving step, albeit under strictly limited conditions. Drivers could gain the full comfort and convenience and safety benefits of AHS, for limited travel duration, but not any additional throughput increases because there would not yet be a means for inserting larger volumes of vehicles into the traffic stream.

7.   Addition of *automatic entry and exit coordination* for access to and from the single automated lane – Add vehicle-roadside communication and vehicle synchronization capabilities at the entry and exit points to and from the automated lane, so that the longitudinal maneuvering of the entering vehicles can be coordinated with the automated mainline operations and drivers can be relieved of the task of synchronizing their vehicles with the mainline traffic gaps. These communication and control systems also need to be integrated with the traffic management system for the rest of the roadway network to ensure that new bottlenecks are not created in the vicinity of the entries and exits. At this point, we have reached the stage of a basic single-lane AHS, which could be extended to an arbitrarily long highway. In the absence of completely automatic lane-changing control, the entry and exit maneuvers would still present some possible challenges and hazards to drivers, and it may therefore not be possible to gain the full capacity advantages of AHS. However, the throughput increase should be substantial at this stage, enabling real reductions in congestion and travel time.

8. Addition of *lane-changing control* – With the addition of automated lane changing, it becomes possible to extend automated operations from a single-lane to a multi-lane protected facility and to provide for fully automated entry and exit maneuvers. At this stage, the full lane or link capacity benefits of AHS can be gained on the protected facility.

9. Addition of *link and network layer control* – With the further addition of these capabilities, it becomes possible to develop a connected AHS network, rather than just a single AHS line, so that travelers can take an automated trip of significant length, involving transfers among different highways, without needing to resume manual control.

The above sequence assumes that CACC would precede the availability of protected lanes. However, there are already some highways in the U.S. equipped with segregated, protected HOV lanes. In locations such as these, it could become possible to begin with the segregated, protected lanes and then add cooperative capabilities later. This depends on the motivation of the HOV operator to support the equipping of the buses and employer-based vanpools that use the HOV facility with ACC and potentially newer and more capable technologies to maximize the utility of the HOV facility.

In some locations, highway expansions involving new separated lanes are already being planned. With minor incremental costs, these facilities could be made much more hospitable toward automated vehicle operations, representing wonderful early targets of opportunity for advancing toward AHS. In cases such as these, it is not necessary to proceed through all of the above steps in sequence, but it could be possible to start at step 5 or 6 and then advance through steps 7 and 8 simultaneously, before proceeding to step 9 when more facilities are developed and connected.

3.3.4 Benefits
The investments required to advance through each of the above steps need to be justified based on the benefits to be gained by both the public and private sector investors. Some of these benefits, such as driver comfort and convenience, are subjective and difficult to quantify. Others, such as safety improvements, are challenging to quantify in advance because of their strong dependence on specifics of system design and implementation. The benefits associated with lane

throughput, which translate directly into savings in travel time and reductions in congestion, are somewhat more tractable.

A preliminary estimate of the throughput effects of mixing automated with manual vehicles, based on a highly simplified model, has already been reported (18). However, that analysis did not address the mixing of both autonomous and cooperative longitudinal controlled vehicles with manual vehicles, nor did it consider the driver role with these partial automation systems. The driver role can become a more severe constraint than the technology in defining the separations that should be maintained between vehicles.

Under normal highway traffic conditions, the maximum stable traffic flow is about 2200 vehicles per lane per hour, which corresponds to an average headway of 3600/2200 = 1.63 seconds. If we assume an operating speed of 30 m/s (67 mph) and vehicle length of 5 m, the average time gap between vehicles is then 1.63 – 5/30 = 1.47 s. This average includes the extra-long gaps that will inevitably occur when vehicles change lanes and when a slow vehicle trails behind a fast vehicle, so it must also include a substantial fraction of gaps significantly smaller than 1.47 s.

There has been a vigorous debate about what minimum time gap should be permitted by standards for autonomous ACC systems. The current international agreement establishes a minimum setting of 1.0 s, based primarily on the limited ability of drivers to react sufficiently quickly to anomalous conditions (failures or driving conditions that exceed the capabilities of the ACC system). This means that even if the entire population of vehicles in a highway lane were equipped with ACC, all choosing the minimum setting, the theoretical capacity of that lane could not exceed 3600/(1.0 + 5/30) = 3086 at a speed of 30 m/s. Allowing a margin of 25% for extra gaps to accommodate lane changing reduces this capacity to 2469 vehicles per hour. However, since many drivers will not be comfortable choosing the minimum setting, given the limited response capabilities of these systems, it is unlikely that autonomous ACC could produce any appreciable increase in lane throughput, even at a 100% market penetration.

Increases in lane throughput beyond today's are only likely to be achievable through the addition of cooperative vehicle-following capabilities (cooperative ACC, or CACC), combined with

separation from non-cooperative vehicles and perhaps in combination with physical protection of the driving environment from hazards. These make it possible to provide tighter coupling between vehicles so that drivers may feel secure selecting a shorter minimum setting than the 1.0 s time gap for autonomous ACC. The communication link also improves the safety of the vehicle-following system by providing the following vehicle with near-instantaneous warning of any problem (failure, forward hazard) experienced by the leading vehicle. This can be used to justify permitting shorter time gaps than 1.0 s, but the time gaps can still not be decreased dramatically as long as the driver must be counted on as the "backup" safety supervisor to handle the hazards that the system cannot handle.

The requirement for driver supervision is likely to be maintained for any operations that are not in fully protected lanes, because the hazards that can then be encountered are so diverse and complicated that sensors cannot be expected to identify them successfully within the foreseeable future. When the lane is fully protected, it could become possible to disengage the driver from even the supervisory responsibilities, and then the minimum separation between vehicles could be determined entirely by technological limitations.

It is not possible to determine *a priori* what time gap would be acceptable to drivers of CACC vehicles. If, for example, they could operate at a time gap of 0.8 s rather than 1.0 s, the previous lane-changing-derated capacity estimate of 2469 could be increased to 2980 vehicles per lane per hour, an increase of about 20%. The actual capacity that is achievable will ultimately depend on the portion of the driving population that feels comfortable choosing each part of the available range of vehicle-following time gaps. Unless a substantial percentage of the drivers select the closest setting, the overall lane capacity will probably not increase appreciably. The significant capacity increases that lead to congestion reductions and travel-time savings cannot be gained until the longitudinal control of the vehicles is completely automated, removed from the capricious decision making of individual drivers.

## 3.4   Deployment Staging Sequences by Vehicle Category

The generic deployment sequence outlined in the preceding section does not account for the different economic and operational factors that distinguish the major categories of road vehicles.

These factors are likely to mean that automation will not progress at the same pace, nor necessarily even in the same sequence, for these different categories. In general, we should expect the heavy vehicle categories (both transit buses and commercial trucks) to be earlier adopters than light-duty vehicles. There are three primary reasons for this:

- When the automation technology is still new, it will be relatively expensive. The fixed cost of the technology will be similar for all classes of vehicles, but it will represent a much smaller fraction of the total vehicle cost for the heavy vehicles, which can cost from five to ten times as much as a passenger car;

- The improvements in productivity produced by the automation technology can be translated directly into an economic benefit by the heavy vehicle owners, which strengthens their justification for investing in the technology. The productivity effect is particularly strong because driver labor saving translates directly into dollars for them, while it does not have the same impact for most private automobile owners;

- Heavy vehicles in fleet operations are driven by professional drivers and maintained by professional maintenance staffs, in contrast to private passenger cars. This means that drivers can be trained to understand limitations of the technologies that are implemented before they are robust to all possible operating conditions. The new components and subsystems can also be tested regularly in an orderly preventive maintenance program, leaving them less vulnerable to the neglect that is generally suffered by private passenger car systems.

Transit buses have the additional advantage that in many cases their owners and operators are public agencies that may also own and operate some of the roadway infrastructure on which they operate (busways, HOV lanes). This means that the vehicle and infrastructure decision-making is consolidated, eliminating a major potential institutional hurdle and short-circuiting the normal "chicken and egg" problem of who invests first. Investments in AHS technologies for transit buses can be extended to vanpool fleets, and the special roadway infrastructure could be made available for use by private vanpool and carpool users as well, providing a logical sequence for expanding the population of user vehicles.

Sample AHS deployment staging 'roadmaps' for transit, trucking and private automobile users are illustrated in Figures 2-4. These sequences have much in common with each other and with the generic deployment roadmap (Figure 1) that was shown in Chapter 2, although they differ in some specific details peculiar to their respective needs. More importantly, they differ in time scales, with transit probably being the fastest and the private automobile the slowest. In these figures, the rectangular boxes represent enabling technologies that are needed, while the ovals represent user services that make use of the technologies. Time proceeds from left to right, and the shaded elements are those that cannot be expected to develop "on their own" as part of a natural market progression (or evolution). These are the key elements that will need public investment and/or the cooperation of public/private partnerships in order to become a reality. The unshaded elements are assumed to become available through normal product development and marketing processes. The blocks below the dashed line in each figure represent supporting capabilities that are likely to become available, but are not integral to the progress toward AHS.
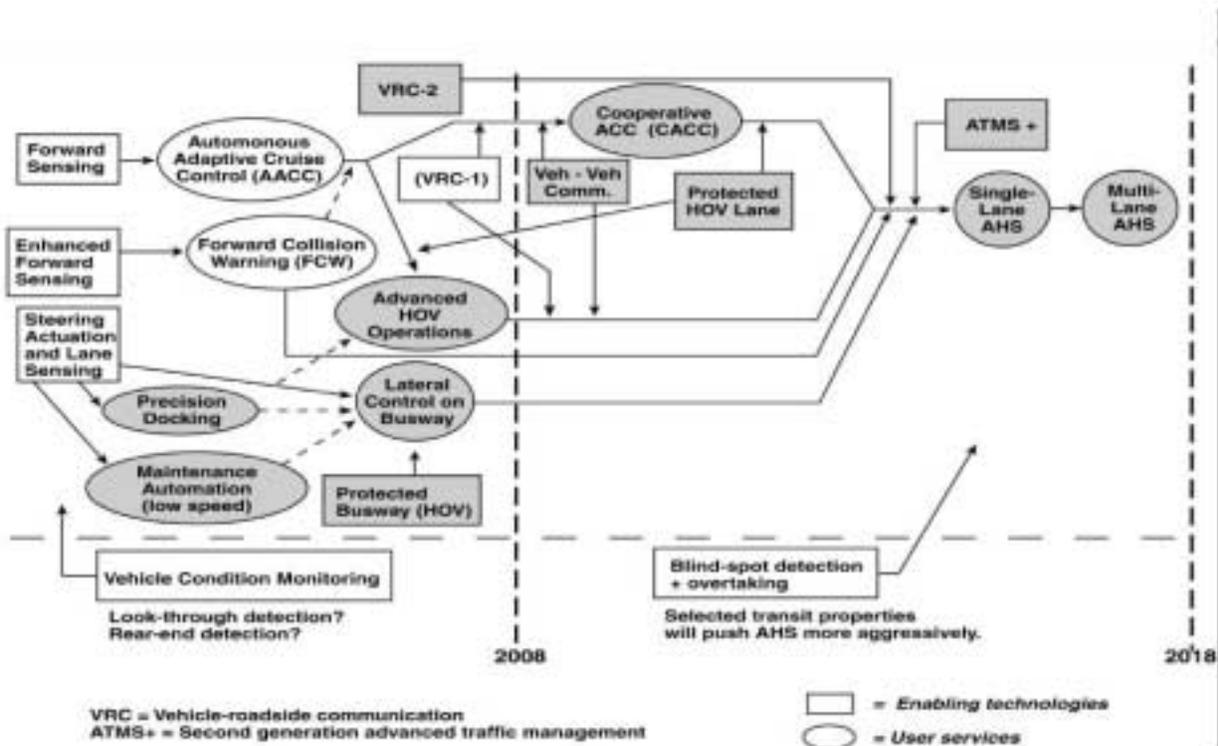
FIGURE 2 Deployment staging for transit (buses, vanpools).

The special elements particular to transit buses in Figure 2 are the early use of lane sensing and steering actuation to provide precision docking, maintenance automation and then lateral control on busways. These, combined with the possible early availability of a protected busway/HOV lane, give the transit buses the head start over the other categories of vehicles.
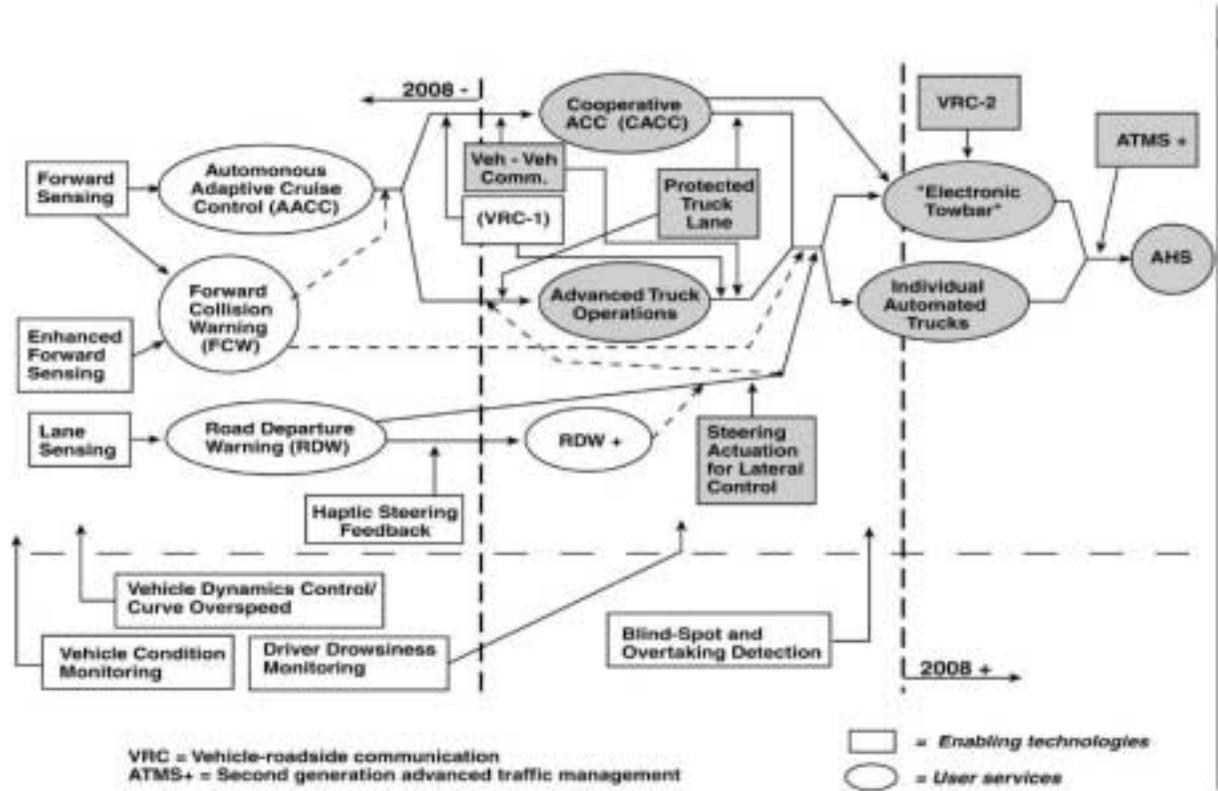
FIGURE 3 Deployment staging for heavy trucks.

The 'roadmaps' for the heavy trucks, Figure 3, and light-duty passenger vehicles, Figure 4, are quite similar to each other, but the trucks have an opportunity to reach higher levels of automation sooner *if* their stakeholders are genuinely interested in getting there. They also have the peculiar opportunity to adopt the "electronic towbar" to automatically link pairs of trucks while sharing the roadway with normal manual traffic (19). The milestone years of 2008 and 2018 are indicated on the figures to give a general sense of when some of the indicated capabilities could be achieved if there is sufficient stakeholder interest to move forward. Note that Figure 3 shows a large uncertainty about the year 2008, indicating the range between what

could be accomplished if the trucking industry stakeholders express a strong positive interest in use of automation to improve their operations and what is likely if they remain passive.
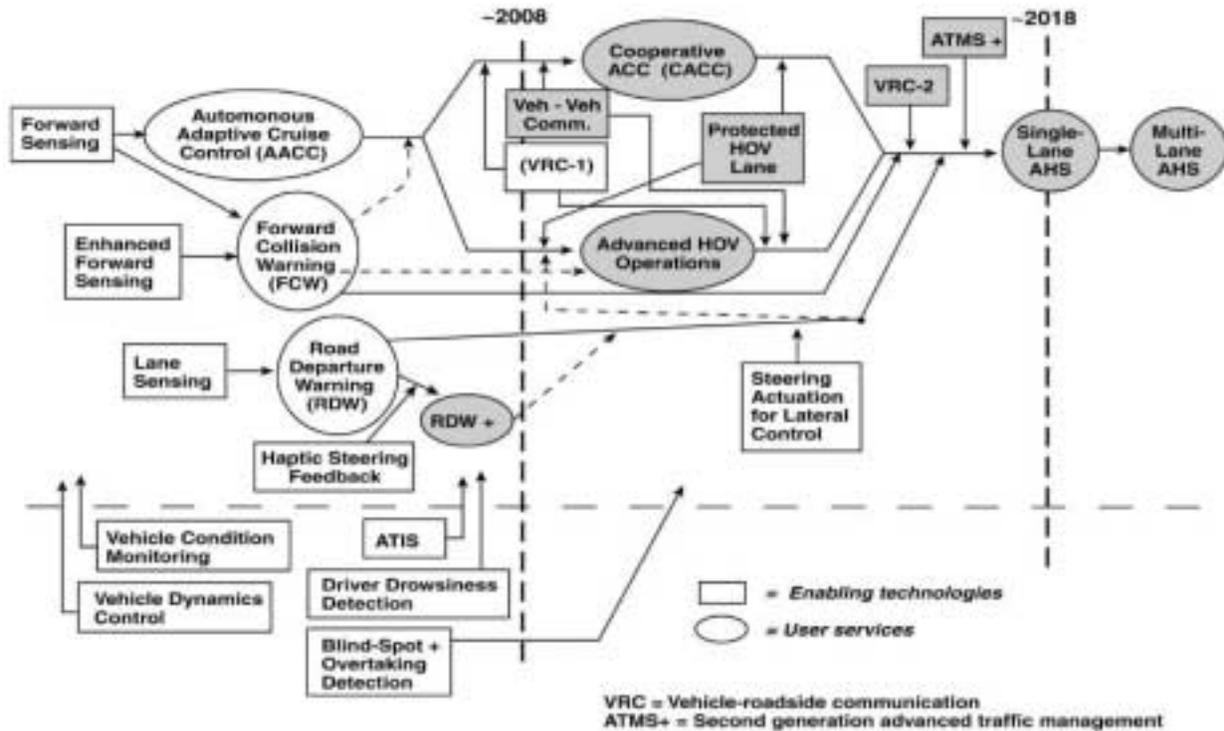
FIGURE 4 Deployment staging for light duty vehicles.

4.0   MODELING AND EVALUATING THE EFFECTS OF DRIVER CONTROL ASSISTANCE SYSTEMS ON TRAFFIC

The next step in our investigation of AVCSS deployment sequences to advance from today's driving environment to full vehicle-highway automation involved their modeling and evaluation. Due to the significant amount of modeling and simulation work necessary, especially for the case of manually-driven vehicles, we were constrained to focus on only initial steps immediately beyond adaptive cruise control and analyzed the impacts of cooperative adaptive cruise control in a mixed-lane traffic setting, i.e., with manually-driven vehicles and vehicles under AACC control (See sections 3.3.1 and 3.3.3) in a single lane. We used light-duty passenger vehicles as the vehicle type for this analysis.

Adaptive cruise control (ACC) systems are the first driver control assistance systems to enter the marketplace having the potential to influence traffic flow characteristics. These systems use forward ranging sensors (typically millimeter wave radar or infrared laser) to measure the clearance (and sometimes the closing speed) to a forward vehicle to modify the commanded speed of the equipped vehicle's cruise control system. The ACC system's control logic substitutes for driver decision making in determining the desired vehicle speed at each time instant. This has profound implications for traffic flow theory, which has historically been founded on the dynamics produced by the aggregate decision making of the driving population.

ACC systems have recently been introduced into passenger car markets and they are also being provided in heavy-duty trucks. It is important to understand their effects on traffic dynamics early in their development so that if they are discovered to inadvertently create problems their designs can be adjusted before adverse traffic effects are widely manifested. ACC systems can also be viewed as a first step toward future automation of the complete driving function. Thus, it is important to understand how ACC system designs and operating concepts can be developed to best support such future advances.

To evaluate interactions between ACC vehicles and conventionally driven vehicles, it is necessary to use sophisticated and sensitive models of each having sufficient fidelity to capture their sometimes subtle interactions. It is not practical now to evaluate these interactions in full-scale experiments because such experiments would require a large ACC vehicle fleet, not yet widely available. In the longer term, it would, however, be valuable to perform these experiments. This report describes a new set of models for evaluating large-scale interactions occurring when varying proportions of ACC vehicles are mixed with conventionally driven vehicles in a limited-access highway operating environment. To address the longer-term advances in ACC capabilities, models are developed for higher-capability cooperative ACC systems (CACC) as well as the first-generation autonomous ACC.

Existing traffic simulation models incorporate simple representations of drivers' car-following behavior. However, those models do not represent car following behavior at the level of detail

needed to predict expected traffic flow changes when ACC vehicles are mixed into the traffic stream. Hence, it has been necessary to adopt a new state-of-the-art driver behavior model (20). Its calibration and testing for use in this project are described in this report's Human Driver Model section. Subsequent sections describe the ACC and CACC models used in our simulations. The ACC model is based on a simple, first-generation ACC system with a fixed time gap setting of 1.4 seconds (s), representing a middle performance range for the ACC systems that have been tested and are entering the market (21, 22). The CACC model represents a more sophisticated form of ACC, in which the preceding vehicle can communicate actively with the following vehicle so that their speed changes can be coordinated, permitting significantly closer vehicle following (time gap of 0.5 s).

4.1   Review of the Literature

The literature is well supplied with papers attempting to predict the effects that the introduction of ACC vehicles into the traffic stream will have on overall traffic behavior (congestion, safety, emissions, fuel consumption, etc.). These papers have been based on analytical models and/or simulations since there are not yet enough ACC vehicles to perform a direct experimental evaluation. Another body of literature addresses the prediction of the analogous effects for fully-automated driving in an automated highway system (AHS), but these are not addressed here.

Rao and Varaiya (23) based their analysis on a modification of a simulation originally developed for AHS evaluation, which led to assumptions that are not really appropriate for evaluation of ACC (clustering of vehicles into platoons and use of a transition lane for entering and exiting an ACC lane). Since they lacked a realistic model of human car-following behavior, they were not able to evaluate dynamic interactions between ACC and manually-driven vehicles.

Chang and Lai (24) did detailed simulations to represent anticipated conditions on Taiwan highways, using their own ACC control law design (with separation dependent on velocity and velocity squared). The human drivers' choices of following distance were assumed to follow the Greenshields linear dependence on speed, but the dynamics of human responses were not explained fully. They showed expected traffic flow effects at market penetrations of 30, 50, 80

and 100% ACC, with a nonlinear increase in traffic flow capacity, leading to a maximum of 2750 vehicles/lane/hour with all ACC vehicles.

Zwaneveld and van Arem (25) reviewed a variety of papers to synthesize existing predictions of traffic effects of ACC, concluding that the general tendency was to show little potential for congestion reduction, but some potential to improve traffic flow stability.

The MIXIC model (26) was developed to study effects of ACC on traffic flow. Different ACC market penetrations were studied, concentrating on traffic flow characteristics, such as traffic speed, time gaps, and shock waves. They concluded ACC systems could contribute to more stable traffic flow without sacrificing capacity. However, a deterioration of traffic performance was found at higher levels of traffic demand: speed collapsed, especially in the left lane. This tool was calibrated for Dutch highways and Dutch driving behavior, and incorporated European traffic laws.

Ioannou and Bose (27) investigated the ACC car's influence on traffic flow and stability. They compared the 100% manual, 100% ACC, and mixed traffic cases, on a straight one-lane road. Traffic was simulated with two kinds of disturbances: smooth and rapid acceleration of lead car. They concluded that ACC vehicles smooth traffic flow by filtering the response of rapidly accelerating lead vehicles. Since only two disturbances were used for the lead vehicle, their simulation does not represent real traffic with its variety of disturbances.

Yokota, et.al. (28) performed detailed simulations to explore the effects ACC could have in overcoming bottlenecks produced by the decelerations of typical human drivers encountering a significant highway grade change (-2.6% to +0.3%). They represented human driver behavior using Koshi's car-following model to reproduce existing shock wave conditions, then introduced 20%, 40% and 100% ACC vehicles to the traffic stream to determine their traffic stability effects. Even at the 20% rate, ACC vehicles significantly reduced the shock wave, while at higher percentages they almost completely eliminated it. They also simulated an urban scenario with ACC vehicles representing 10%, 30% and 100% of traffic flow, with ACC time gaps of 0.9 s and 1.6 s. Trip times and delays were the chosen measures of effectiveness, rather than overall

lane capacity. Yokota et al (29) also attempted to make some simple aggregate calculations of ACC's effects on traffic speed, using the fundamental traffic flow-velocity diagram, but without benefit of any explicit modeling of ACC or human driver dynamic responses.

Cremer et.al. (30) simultaneously compared one macroscopic and four microscopic models for predicting the traffic flow effects of ACC (time gap of 1.2 s). Results showed little change in traffic flow at ACC market penetrations of 40% and 70%, but indicated that either shorter or longer ACC time gaps could adjust those traffic flows up or down respectively.

Darbha and Rajagopal (31) explored traffic flow stability issues that need to be considered when designing ACC controllers when all vehicles are equipped with ACC, but with no representation of manually driven vehicles and no prediction of traffic flow volume achievable with ACC.

Minderhoud and Bovy (32) developed detailed simulations of a variety of ACC operating concepts for use on Dutch motorways, emphasizing the effects observed when entering traffic disturbs on-ramp flow. The simulation cases were run with ever-increasing traffic flow rate generation on the mainline highway and on-ramp, well past the flow rates that could be sustained in steady state, so that queuing effects could not be evaluated realistically. Results showed that the ACC with time gap setting of 1.2 s would leave traffic flow capacity essentially unchanged from the baseline manual driving conditions, while time gaps of 1.0 s and 1.4 s could cause noticeable, yet small, increases and decreases in flow capacity respectively at the highest market penetrations. It was necessary to reduce the time gap to 0.8 s to generate capacity increases in the 10% range at market penetrations of 50% or higher.

Godbole et.al., (33) analyzed the behavior of two ACC designs (with and without active braking) to determine their ability to respond to the kind of transient disturbances occurring in highly congested freeway driving. All vehicles were assumed to be ACC-equipped except for the lead vehicle that causes the disturbance, so effects of different market penetrations of ACC could not be evaluated. For the assumed ACC time gap of 1.0 s, the achievable capacity was estimated to be approximately 2,800 vehicles per lane per hour.

4.2   Method of Approach

For the study reported upon here, it was necessary to develop and then validate several mathematical models, including:

- Vehicle-following logic for each operational mode:
  - Human driving,
  - Autonomous adaptive cruise control (AACC), and
  - Cooperative adaptive cruise control (CACC), both with and without a cooperating preceding vehicle.
- Merging of vehicles entering a limited-access highway from an on-ramp.
- "Free driving" of a vehicle in uncongested traffic or with no immediately preceding vehicle to follow.
- Vehicle dynamic response to speed change commands.

These models were tested initially for simple vehicle pairs (one following another), then for sequences of 20 vehicles with the same type of vehicle-following controller, and finally more complicated cases involving mixtures of controller types and larger numbers of vehicles, with vehicles entering and leaving the flow at interchanges. In this section of the report, these models and their test cases are described in general terms, while subsequent sections address the three classes of vehicle-following logic in more depth, with results of the validation test simulations.

4.2.1   Vehicle Following Logic

Each of the four operational modes, human, AACC, and CACC with and without communication, has its own following logic. In all cases, however, the free-driving acceleration is limited to +/- 2 m/s/s and no vehicle ever accelerates to a velocity greater than its desired velocity. The variables included for vehicle following are the current distance between vehicles, speed of both the preceding and following vehicles, and the length of the vehicle (all vehicles are assumed to be of the same length).

4.2.2 Free Driving

When a vehicle has no vehicle ahead of it or has at least 100 m of clearance to the preceding vehicle, the controller (human, ACC, or CACC) is in the free driving state. In this state, the controller attempts to maintain a desired velocity, assigned when the vehicle entered the simulation, drawn from a normal distribution with a mean of 28.9 m/s (65 mph) and a standard deviation of 4.4 m/s (10 mph). Each type of controller uses an error-based control law:

$$u_2(t) = -k_f(\dot{x}_2(t) - v_{2d}(t))$$

Where: $v_{2d}(t)$ is desired speed of the second vehicle

         $k_f = 0.4$

In addition, acceleration is limited to +/- 2 m/s/s.

4.2.3 Vehicle Model

The detail level of our models is constrained by the need to simulate hundreds of vehicles at one time. We decided to use a unit mass dynamic model: $\ddot{x}(t) = u_2(t)$. We also include a first-order lag between the controller command and the vehicle's response. The flow of data among components of the vehicle in the simulation is as follows:

1. Sensors output range and range-rate to the controller.
2. Controller outputs acceleration command to the first-order lag.
3. Lag outputs actual acceleration to vehicle mass (dynamics).
4. Dynamics integrates and outputs velocity.
5. Velocities of all vehicles are integrated to update their positions, which can then be read again by sensors.

Sensors

Sensors are assumed to be perfect, except in the case of a CACC vehicle following a non-CACC vehicle, in which case the control algorithm's sensitivity to small errors in range-rate justifies the added computational expense of simulated range-rate noise.

Braking capabilities

In our simulation, hard braking is unlikely. However, to meet safety standards, our controllers must be able to handle such situations. We use a distribution of light-vehicle braking capabilities from the NAHSC (34), assuming dry pavement, as shown in Figure 5. Engine braking is not modeled separately. Brake dynamics are not explicitly modeled, although the first-order lag applies to deceleration as well as to acceleration.

Lag

The lag between controller command u and vehicle response a is a first order lag with the equation:

$$a' = (u - a)/Tc$$

The choice of time constant Tc turned out to be critical for the CACC controller. We found that a time constant of 0.05 s provided a realistic representation of vehicle response without requiring an excessively short simulation time step.

4.2.4   Test Cases

We show how a platoon of 20 vehicles each with the same controller (human, ACC or CACC) responds to severe disturbances. The disturbances are generated by a lead vehicle that follows a velocity trajectory taken from severely congested stop-and-go traffic on I-880 in the San Francisco Bay Area. We allow the following vehicles to follow each other stably before injecting the disturbance. For each run we plot the velocities of vehicles 0 (the leader), 5, 10, 15, and 20, and the trajectories of all 21 vehicles.
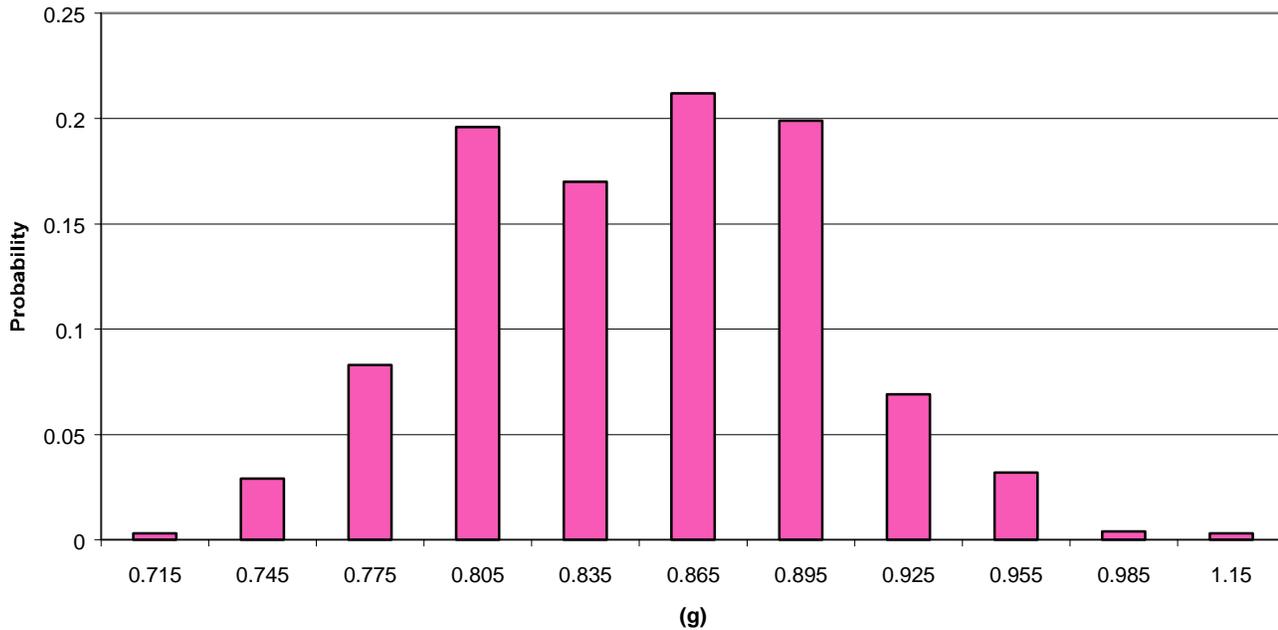
FIGURE 5 Weighted distribution of maximum deceleration capability.

4.3   Human Driver Model

We considered several existing models. MIXIC (26) has only one control region, and the parameters for that region are tuned for stable following behavior. It did not perform realistically in the other kinds of interactions that occur in our simulations. Fancher's model (35) has more control regions but has some problems with the transition from free driving to following. We settled on the model of Song-Delorme (20) because it has more room for tuning, and incorporates recent research on perception (36).

Desired time gap is assumed to be normally distributed, with a mean of 1.1s. This value was taken from a statistical analysis of the UMTRI ACC FOT baseline case human driving data (21). The model has many other characteristics (six discrete control regions, and for each region a gain factor), which have not been fully tuned from microscopic driving data. However, macroscopic observations of the simulation outputs appear to match standard sources describing aggregate traffic characteristics such as the Highway Capacity Manual (37).
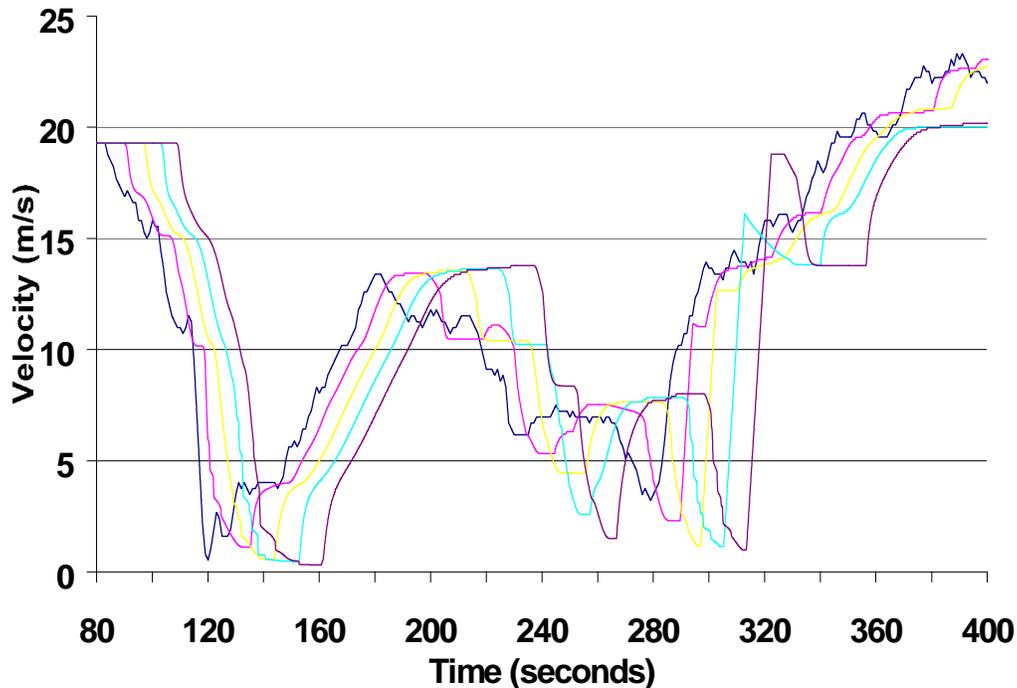
**FIGURE 6 Velocity trajectories for lead vehicle and human followers
#5, 10, 15, and 20.**

Figure 6 shows the velocity trajectories for four sample vehicles of a 20-vehicle sequence using the human driver model, all of which are following the lead vehicle trajectory previously described from the I-880 database. The plotted results clearly show the kind of shock wave propagation expected in normal traffic response, with larger disturbances at later times for the vehicles further behind the leader. The higher-frequency disturbances of the lead vehicle velocity are also clearly damped out by the filtering effects of the drivers' perceptual limitations and response lags. Figure 7 shows a portion of the same simulation from 110 seconds to 170 seconds, in a distance-vs.-time plot, which clearly shows the compression effect of a typical shock wave.
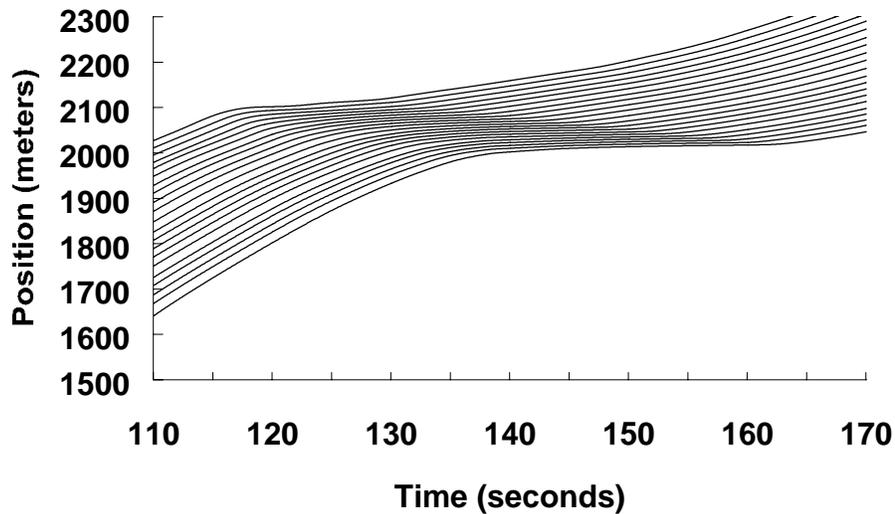
FIGURE 7 Human vehicle following trajectories.

4.4   Autonomous ACC Model

The purpose of ACC is to regulate the range between vehicles to a user-selected value, and to adjust the vehicle speed to the speed of downstream traffic. It is mainly designed to increase comfort and convenience for the drivers. The driver is responsible for safety and can override the ACC system at any time. The driver takes over control of the vehicle in situations demanding more braking authority than the ACC system can provide. It usually uses sensor measurements of range and range-rate to regulate the spacing between the preceding and following vehicles to a driver-desired value, r_d.

The first controller considered was ACC with braking proposed by Godbole (33). Here, three surfaces divide the range and range-rate space into five regions, with a different mode of ACC operation in each. Controller design is based on the notion of the safety surface, which is defined as a function of braking capabilities of the preceding and following vehicles, range, range-rate, driver time delay, and current vehicle speed. The desired vehicle behavior is to maintain a safe distance from the preceding vehicle.

Safety distance and hence desired following distance depends heavily on the braking capabilities of both vehicles, which, without communication between vehicles, are unknown. To avoid collisions the algorithm assumes the preceding vehicle has the maximum deceleration capability, which raises the safety and desired following distance to maximum values. Various studies have shown that the desired car-following distance depends on driver personality (21). Therefore, we chose desired range to be a function of the driver-selected constant time gap. When we changed the definition of desired distance we lost one of the three surfaces of ACC design, and could not use this controller anymore. Inspecting its five regions we noticed that three of them, corresponding to vehicle following, utilize linear control laws with deceleration limits. The vehicle was in the other two regions in the more extreme cases, while it was driving alone, or the situation required the maximum braking capability of the ACC. The three linear control laws similarly use range and range-rate and differ in their gains and maximum allowed deceleration. Since we did not have a safety surface to separate those regions we decided to merge them and use the same controller in the whole area.

Finally, the ACC controller that we adopted is linear with limits on deceleration and acceleration. The control variable is the acceleration command.

$$u_2(t) = k_1 * (\dot{x}_1(t) - \dot{x}_2(t)) + k_2 * (r(t) - r_d(t))$$

$$u_2(t) \leq a_{max}$$
$$u_2(t) \geq d_{max}$$
$$r_d(t) = t_g * \dot{x}_2(t)$$

Where: $r_d(t)$ = desired distance between vehicles

$r(t)$ = current distance between vehicles

t_g = desired time gap

$\dot{x}_1(t)$ = speed of the preceding vehicle

$\dot{x}_2(t)$ = speed of the following vehicle

$a_{\max}$ = maximum allowed acceleration

$d_{\max}$ = maximum allowed deceleration

$k_1 > 0$ and $k_2 > 0$ are gains, which are tuned for the ACC performance

Since ACC is a comfortable system we bounded its acceleration capability by a "comfortable value" of 2 m/s/s (38), while the deceleration capability has the limit of -3 m/s/s.

We chose a time gap of 1.4 s representing a mid-point among the settings that are being provided in commercially available ACC systems, which typically range from 1.0 to 2.0 s.

This ACC controller is structurally simpler than the controllers that are being commercially implemented. Because it has been designed to ensure string stability of vehicle following, its response to traffic disturbances is equivalent to the best that could be achieved with real vehicles. The results derived using this model therefore likely provide higher capacity and smoother damping of shock waves than would be achieved in practice.

Figure 8 shows the velocity trajectories for four sample vehicles of a 20-vehicle sequence using the ACC model, following the same lead vehicle trajectory as in Figure 6. Because this ACC control law was designed specifically to provide string stability (unlike some that have been developed elsewhere), disturbances are clearly damped out the further we get from the lead vehicle. The responses are obviously delayed in time, but we do not encounter the amplification of disturbances experienced with the human driver. Absence of shock wave propagation can be seen clearly in Figure 9, showing one potential improvement to traffic dynamics from use of ACC.
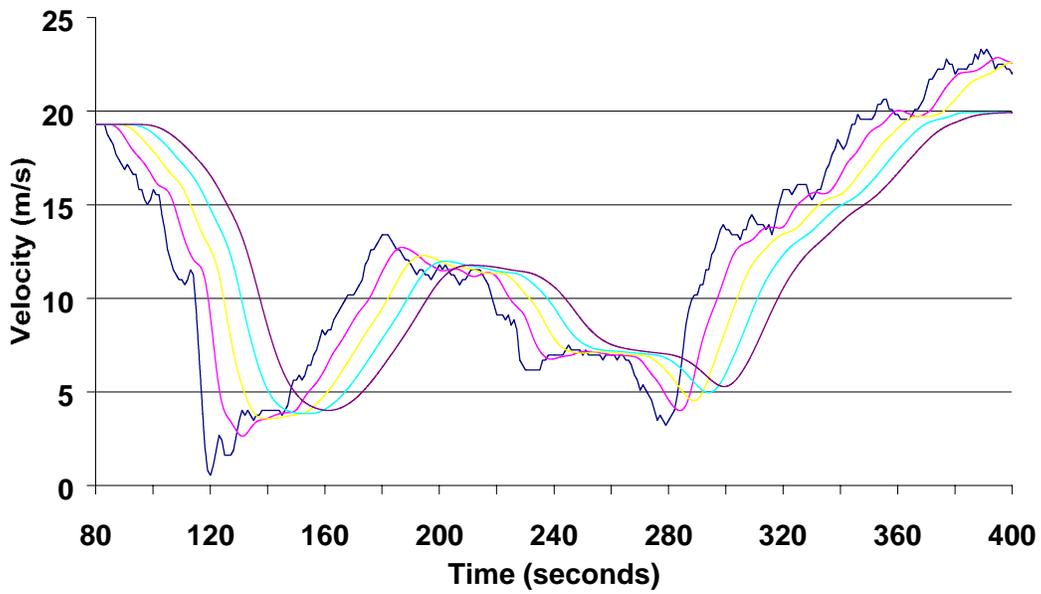
**FIGURE 8 Velocity trajectories of lead vehicle and AACC followers
#5, 10, 15, and 20.**
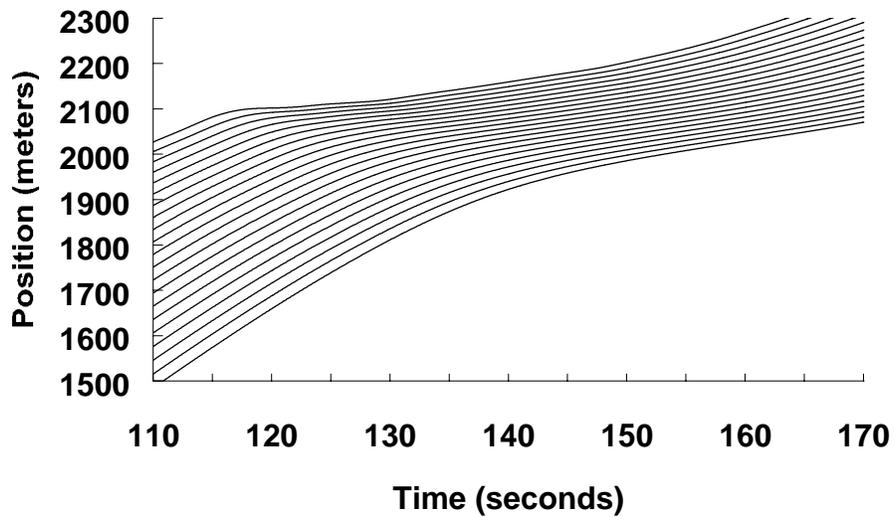
FIGURE 9 AACC vehicle following trajectories.

4.5 Cooperative ACC (CACC) Model

Cooperative ACC (CACC) systems add vehicle-to-vehicle communication to the capabilities of ACC systems. We restricted our model to one-way, "hop-by-hop" communication. That is, messages are transmitted from a vehicle to the upstream vehicle in the same lane, providing the simplest possible cooperative architecture and the easiest to implement. This form of cooperation can be implemented using simple line-of-sight wireless communication technologies and does not require any complicated addressing schemes or infrastructure.

We have assumed a perfectly reliable and instantaneous form of transmitting three double-precision floating-point numbers per time-step (50 ms). Estimates of transmission delays for this amount of data with existing technology fall well below our simulation time-step and no data needs to be relayed past the receiving vehicle; it is simply used by the controller to produce its control outputs.

The data we chose to transmit are:
- velocity of lead vehicle,
- acceleration of lead vehicle, and
- braking capability of lead vehicle, i.e., its maximum deceleration.

These, along with range and range-rate from the sensors, represent all the information needed to enable the following vehicle to closely match the behavior of the lead vehicle. This information would not be available to an autonomous vehicle or else would only be available with less accuracy (lead vehicle velocity).

The CACC model evaluated in this study is intended to increase highway capacity by minimizing time gaps between vehicles, while maintaining the typical ACC goals of increasing driving comfort and convenience and preserving driving safety.

The CACC, like the ACC system, uses an error-based controller attempting to maintain three invariants:
- $\ddot{x}_1(t) = \ddot{x}_2(t)$

- $\dot{x}_1(t) = \dot{x}_2(t)$

- $r = r_d$.

The desired gap, r_d, is not a fixed time gap but rather is chosen so that if the front vehicle brakes at its maximum rate and the follower reacts immediately by braking at its own maximum rate, they will not collide. If the braking capability of the following vehicle exceeds that of the preceding vehicle this gap could be quite small, so the follower does not try to follow at a closer time gap than 0.5 s. At this short a time gap, the system designer cannot assume that a normal driver will always be able to respond to anomalous behavior by the preceding vehicle, but such anomalous behaviors are much less likely to occur because the preceding vehicle would not be manually driven but would be operating under ACC control.

The control variable is the CACC commanded acceleration.

$$u_2(t) = k_0 \ddot{x}_1(t) + k_1 (\dot{x}_1(t) - \dot{x}_2(t)) + k_2 (r(t) - r_d(t))$$

$$u_2(t) \le a_{\max} \qquad u_2(t) \ge d_2$$

Where: $r_d(t)$ = desired distance between vehicles

$r(t)$ = current distance between vehicles

$\ddot{x}_1(t)$ = acceleration of preceding vehicle

$\dot{x}_1(t)$ = speed of preceding vehicle

$\dot{x}_2(t)$ = speed of following vehicle

$a_{\max}$ = maximum allowed acceleration

d_2 = braking capability of following vehicle

$k_0 = 1$, $k_1 > 0$ and $k_2 > 0$ are gains, which are the same as for ACC

Desired safe range r_d (t) is defined as a maximum among safe following distance, following distance with 0.5 s time gap, and a minimum allowed distance, chosen to be 2m.

$$r_d(t) = \max(r_{safe}(t), r_{0.5s}(t), r_{\min})$$

$$r_{safe}(t) = \frac{\dot{x}_1^2(t)}{2}(\frac{1}{d_2} - \frac{1}{d_1}) + \delta * \dot{x}_1^2(t)$$

$$r_{0.5s}(t) = 0.5 * \dot{x}_2(t)$$

$$r_{min} = 2$$

Where: $\delta = 20ms$ = communication delay

             d_1 = braking capability of preceding vehicle

             d_2 = braking capability of following vehicle

When a CACC-equipped vehicle follows a non-CACC vehicle, an alternative control scheme must be used. We chose to use the same control laws, but with estimated values for the preceding vehicle substituted for communicated values:

- velocity replaced with noisy velocity,
- acceleration estimated from noisy velocity, and
- braking capability assumed to be the maximum (over all vehicles)

The radar noise is modeled by an error factor uniformly distributed between –20% and +20% of the actual value and sampled every time step. In addition, the minimum desired time gap is raised from 0.5 seconds to 1.4 seconds, to match ACC.

Figure 10 shows the velocity trajectories for four sample vehicles of a 20-vehicle sequence using the CACC model, following the same lead vehicle trajectory as in Figure 6. The tighter control of the CACC controller is evident in the closer match to the velocity trajectory of the leader and the reduction in the response delay from one vehicle to another. While this tight control increases the accuracy of vehicle following and lane capacity, its close matching of velocities also raises the potential for reduced ride quality (particularly if the vehicle being followed is driven erratically) unless the controller is carefully designed to include acceleration and jerk limits. Figure 11 clearly shows both the precision and closeness of the vehicle following achieved here, which is clearly and significantly different from the traditional traffic characteristic shown in Figure 7.
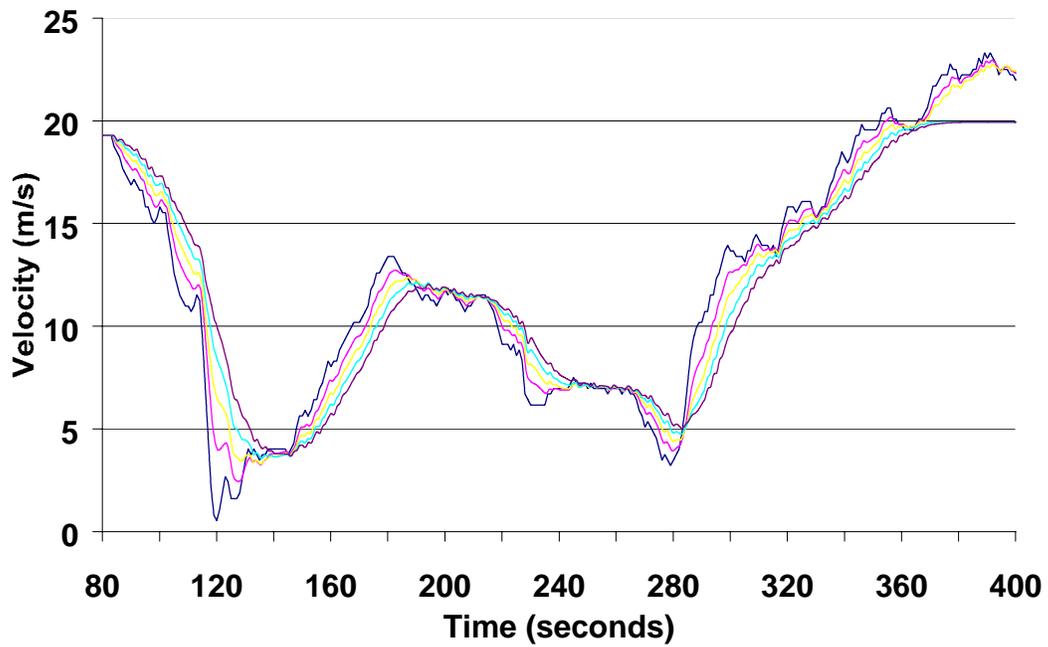
FIGURE 10 Velocity trajectories of lead vehicle and CACC followers #5, 10, 15, and 20.
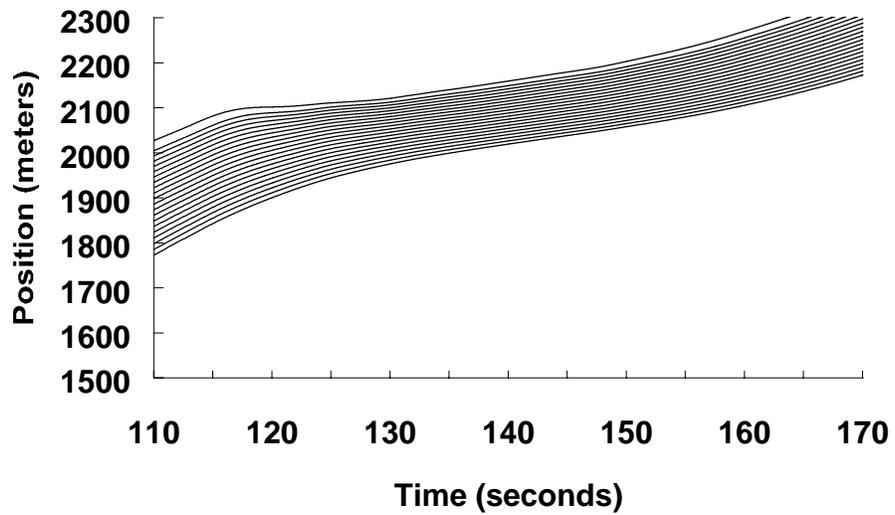
FIGURE 11 CACC vehicle following trajectories.

4.6   Merging Model

We assume that merging is under human control in all cases, since automated control of merging is unlikely to be available in the near future. The driver at the head of the merge queue waits until the gap between passing vehicles is of an acceptable length. When such a gap appears, the driver merges into it at the velocity of traffic. We do not model the lateral movement of the merging vehicle.

Our model for gap acceptance is taken from the dissertation of K. Ahmed (39). This model has been calibrated with driving data and including it in our simulation produces realistic-looking traffic patterns.

Figure 12 shows the trajectories of human-driven vehicles before and after a merge point. Traffic is near capacity and initially moving at about 20 m/s. Merging vehicles enter just after 500 m. The horizontal line represents queued vehicles. Merging vehicles can be identified by a curve that lies entirely above this line. Gaps at the bottom of the figure are due to Monte Carlo simulation of vehicle arrival times.

Several forms of shock waves are evident. First, at around 615 s, a merge causes a disturbance to propagate forward, eventually dissipating. Second, a vehicle queued since about 630 s accepts a gap after three vehicles in the main stream have passed the merge point. The resulting disturbance propagates backwards hundreds of meters over the next few minutes, producing the distinctive shockwave seen in Figure 12. Third, after a period of vehicles queuing but not merging, several vehicles merge into relatively small gaps beginning at about 750 s. However, disturbances dissipate quickly because upstream traffic density here is below the critical threshold for unstable traffic flow.
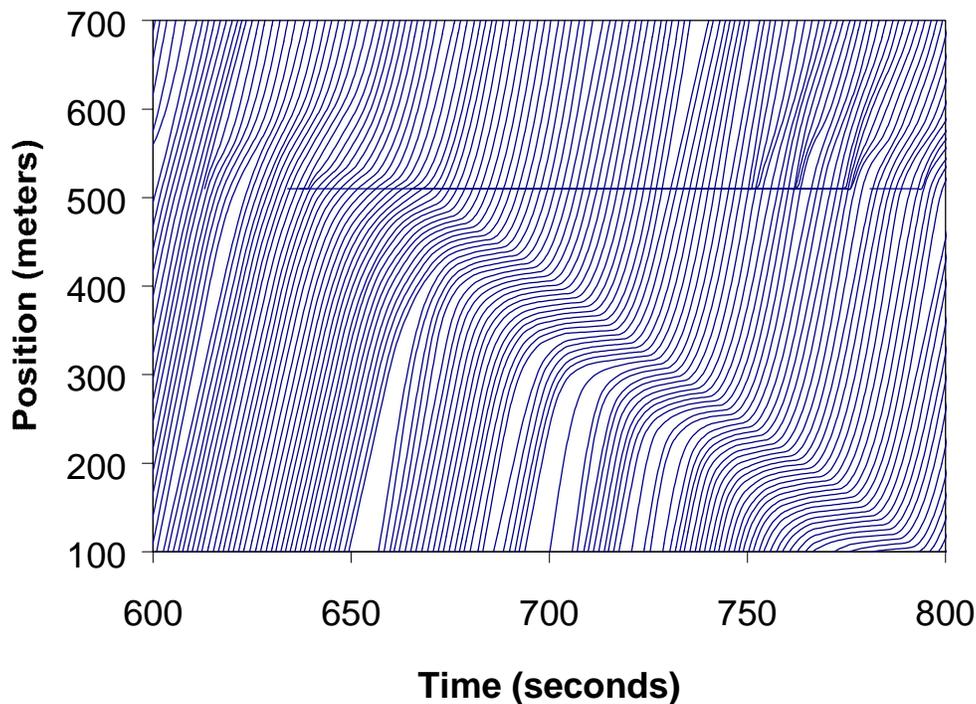
FIGURE 12 Trajectories of human-driven vehicles near merge point.

5.0   EVALUATING THE EFFECTS OF DRIVER CONTROL ASSISTANCE SYSTEMS ON HIGHWAY TRAFFIC FLOW CAPACITY

This section provides quantitative analyses of the lane capacity that could be achievable at several stages in the AHS deployment sequence that was proposed in Section 3. These capacity estimates are important factors in determining the benefits that can be gained at each deployment stage. It is important to understand these benefits because they will serve as the justification for the investments needed to advance to each deployment stage, and if the benefits are insufficient the investments are unlikely to be made. At the same time, these analyses can also help shed light on the effects that ACC vehicles are likely to have on general highway traffic flow as their market penetration increases. This is not yet a well-understood issue, as existing literature on ACC modeling studies provides a wide range of estimates of these effects, ranging from substantial increases in capacity and smoothing of traffic flows to substantial decreases in capacity and worsening of traffic flow instabilities.

5.1   Cases To Be Evaluated

The diversity of possible operating conditions is very wide, so it is important to focus the study on a manageable set of cases that can nevertheless shed light on the likely effects under a wider variety of conditions. The vehicle capabilities that have been chosen for evaluation are:

a) Vehicles driven by "normal" human drivers, represented using a state-of-the-art model of car-following behavior;

b) Vehicles whose speed is controlled by a relatively simple, but high performance, autonomous ACC system, with a driver-selected time gap setting of 1.4 seconds between consecutive vehicles;

c) Vehicles whose speed is controlled by a more advanced cooperative ACC system, using vehicle-vehicle communications to enable operations with a time gap setting of only 0.5 seconds between consecutive vehicles.

The mathematical models used to represent these three classes of vehicles and their calibration and validation were described in detail in Section 4. The reasoning behind the selection of the ACC systems' operating characteristics was as follows:

The autonomous ACC (AACC) system was intended to represent a typical first-generation product such as those now entering the market. The time gap setting of 1.4 seconds is typical of the middle range setting on such vehicles, which are typically adjustable for time gaps from 1.0 to 2.0 seconds. Younger, more aggressive drivers tend to favor the shorter time gap settings, while older, more conservative drivers tend to favor the longer settings (40). The value of 1.4 seconds was chosen as a compromise between these extremes. This ACC system has the capability of automatic acceleration and deceleration, with a maximum deceleration rate of 0.3 g. If larger decelerations are needed to avoid a crash, the driver must intervene (but this kind of emergency condition was not modeled in this study).

The cooperative ACC (CACC) system was intended to represent a significantly more advanced product, in which the equipped vehicle's speed control system could receive wireless communication of the speed, acceleration, and fault conditions of a similarly-equipped preceding

vehicle. When the similarly-equipped vehicle is immediately ahead, it becomes possible to reduce the operating time gap to 0.5 seconds. This shorter time gap was chosen to take advantage of the improved ability of the vehicle to match speed changes of its predecessor. At this short a time gap, the ACC system also needs to have a high enough level of reliability and fault tolerance that it does not need to depend on the driver's manual intervention to avoid hazardous conditions (because the driver would not be capable of intervening quickly enough). When the cooperative ACC vehicle drives behind a vehicle that is not similarly equipped, it cannot operate at the reduced time gap, but must fall back to the larger time gap of the autonomous ACC system.

The analyses were initially conducted for the distinct cases of 100% manually driven vehicles, 100% autonomous ACC and 100% cooperative ACC in order to verify the reasonableness of the results under these simplest cases. Once those cases were completed, mixed vehicle populations were then analyzed, in all feasible multiples of 20% of each vehicle type. These are the cases that will demonstrate the advantages and disadvantages that are obtained with each increase in market penetration of autonomous and cooperative ACC vehicles.

5.2   Simulation Approach

Our goal is to estimate the capacity of a highway as a function of the proportions of AACC, CACC, and human-driven vehicles. Having modeled the three types of controllers described in Section 4, the main difficulties are defining and measuring capacity and generating realistic streams of traffic.

5.2.1   Defining Highway Capacity in Simulation

We define the capacity of our simulated highway to be the maximum rate of flow that it can sustain indefinitely. The following sections explain how we approach and detect this maximum, and how we estimate, in a finite length simulation run, that a certain flow level is indefinitely sustainable.

5.2.2 Measuring Capacity

Our original approach to measuring highway capacity in each case was to simulate a 16 km section of single-lane highway with on- and off- ramps at nodes separated by 1.6 km intervals. The traffic volume at the beginning of this section would be chosen to be well below a conservative estimate of capacity. At each successive node, we would increment traffic flow by a small number of entering vehicles per hour.

For realism, the increment was to be the net effect of injecting a certain number of vehicles through the on-ramp and removing a smaller number of vehicles through the off-ramp. The numbers added and removed would have to be chosen small enough that minor disturbances due to merging settled out without affecting the neighboring merge processes 1.6 km upstream and downstream. In effect, our road would have the same flow properties as a more realistic road with ramps separated by larger distances, but we would not incur the extra computational overhead of simulating long sections of highway with stable flow.

The initial flow and the increments would be chosen so that the demand at the end of the 16 km section was in excess of any reasonable estimate of capacity. This would cause shock waves and queuing at the ramps. By recognizing which node caused these effects, we hoped to measure, to the precision of the chosen increment, the capacity for the given specification of control types and other parameters.

The difficulty with this approach turned out to be the recognition problem. When a shock wave was observed, how could one be sure which node was its source or, if it was caused by vehicle interactions on the link between two nodes, which link? Further, if a shock wave propagates upstream from one node, all subsequent upstream observations are affected; one can say little about where capacity is reached.

To avoid these difficulties, we decided to isolate each node on its own highway, in its own simulation, and model the incoming upstream traffic stochastically. This technique is described in detail below.

5.2.3 Highway Layout

We consider a single protected highway lane, with a single interchange consisting of an off-ramp followed immediately by an on-ramp. We do not model any interaction between exiting and entering traffic. The highway segments adjacent to the merge are 500 m long in the upstream direction and 200 m long in the downstream direction.

Our constraints in choosing these two distances were run time and realism. Shorter road segments require less computation, but we needed to have enough of a window that minor disturbances dissipated before reaching the end of the road. Test runs on longer highways showed that, under sub-capacity flow conditions, disturbances caused by the merge point usually did not propagate beyond 500 m upstream or 200 m downstream. (Near capacity, however, shock waves propagate for much greater distances.)

5.2.4 Traffic Generation

We model two sources of traffic: mainstream traffic entering our highway section from the upstream direction, and merging traffic entering by way of the on ramp.

Mainstream traffic

Our goal in this case is to generate realistically stable traffic, as if vehicles were entering from an upstream segment. To do this, vehicles are generated according to a Poisson distribution without being placed on the road. They are kept in a virtual queue so that when they eventually enter the road, they do so in the order of their creation. The queue in this case is an artifice used to generate steady-state traffic; it does not represent waiting vehicles.

Conditions for a vehicle entering from the queue are as follows. As described in Section 4, the inputs to each controller include range and range-rate to the leading vehicle. The controller of the vehicle at the head of the queue is given the range from the point of entry to the vehicle immediately downstream (typically, the previously entering vehicle). The controller assumes a range-rate of zero, because the vehicle will enter with the same velocity as the downstream vehicle.

As the simulation proceeds, the controller outputs an acceleration command based on the inputs. When this command is greater than or equal to zero, the vehicle is allowed to enter the road. Because of the conditions we have imposed and the design of our controllers, the controller is immediately in its steady-state vehicle following mode.

Merging traffic

For merging traffic, our goal is for entering vehicles to inject disturbances that are typical of merge points on real highways. We cannot use the same model as for mainstream traffic, because in that model vehicles only enter when they can do so without disturbance. We used a merging model based on the dissertation of K. Ahmed (39) and described in Section 4. In this case, the merging vehicle is aware of not only the vehicle immediately downstream of the merge point, but also the vehicle immediately upstream, if any, allowing the merge model to consider the entire gap into which the vehicle must merge. On merging, the vehicle assumes a velocity equal to the average of its neighbors. As in the mainstream case, vehicles are generated off the roadway and placed in a queue until they enter. However, in this case the queue is not virtual, but really represents vehicles waiting to enter the roadway.

After the merge, the ranges between the merging vehicle and the upstream and downstream vehicles may be small enough to cause one of the controllers to command a deceleration, which may in turn cause a shock wave. The resulting traffic patterns appear to be realistic (See Section 4).

5.2.5   Monte-Carlo Generation of Vehicles

Each variable in our simulation is sampled from a distinct generator. The variables are described in Table 1.

TABLE 1 Simulation variables.

Variable	Distribution	Mean and Std. Deviation
Vehicle inter-arrival times	Exponential (note that this value determines arrival in queue, not on the road)	Based on flow rate settings
Vehicle destination (off-ramp or beyond)	Weighted uniform, based on flow rate settings	NA
Desired speed	Gaussian, bounded by fixed interval	29 m/s, std dev = 4.5 m/s
Control type	Weighted uniform, based on control type ratios	NA
Desired time gap	Human: Gaussian, bounded by fixed interval AACC: constant CACC: constant	1.1 sec, std dev = 0.15 sec 1.4 sec 0.5 sec if following CACC 1.4 sec otherwise
Braking capacity	Empirical distribution	NA

We bounded the desired speed distribution by a fixed interval (minimum of 20 m/s, maximum of 40 m/s) because of our single lane layout. One slow vehicle will reduce the speed of traffic on the lane for the remainder of the simulation. In effect, we are assuming that minimum speed laws are enforced.

Detailed explanations of these model characteristics can be found in Section 4.

We used antithetic sampling to test the dependence of simulation outputs on particular values of the seeds and on characteristics of the random number sequence. This test involves two runs of the simulation, differing only in that in one case the basic random sequence of numbers in the interval (0, 1) is passed through the function (1-x) before generating the sequences used by the simulation. Our use of this technique is based on (41).

Each run of the simulation used the same seeds for each generator in order to confine the sources of variability between runs to the issue we are evaluating (the market penetrations of the two

types of ACC). So, for example, the sequence of desired speeds of vehicles generated at the ramp is the same in all (non-antithetic) cases.

### 5.2.6 Estimating Capacity

Using our simplified highway layout to estimate capacity requires that we run, for each proportion of control types, a set of simulations at varying levels of traffic demand. We decided to hold fixed the net flow at the ramp, with 200 vehicles/hr entering and 100 vehicles/hr exiting. Traffic flow on the mainstream varies in increments of 50 vehicles/hr.

A simulation run outputs an overflow message and stops when either of the two sources has 50 or more vehicles queued. If the overflow occurs at the ramp, the message indicates that the mainstream is too congested for vehicles to merge into the mainstream at the rate at which they are generated. If the overflow occurs at the "virtual" queue at the mainstream source, the message indicates either that a disturbance has propagated back to the beginning of the road and would presumably continue upstream indefinitely, or that stably flowing traffic cannot exist at the specified rate. In either case, we can conclude that capacity has been exceeded.

Barring such messages, the run continues for 90 minutes of simulation time, after which we declare the simulated highway to be at or below capacity. The 90 minute period is large enough that the number of vehicles queued (at most 50 in each of two queues) is small in proportion to the total number of vehicles passing through the simulation (over 3000 in most cases). Further, the number of vehicles on the road at any one time is small in proportion to the total number of vehicles generated. Therefore, simply adding the initial upstream flow and the net increment introduced at the interchange yields a good estimate of capacity. Another estimate is obtained by measuring the rate at which vehicles reach the end of the road.

A variety of diagnostics were calculated for each case to ensure good understanding of the traffic flow characteristics. These include: downstream traffic speed distribution, queue length and queue time distributions, and vehicle trip times and speeds.

To reduce computation time, we start the sequence of runs at a mainstream flow level that is expected to be well above the capacity for the given proportion of control types, and reduce this flow by 50 vehicles per hour in successive simulations. We continue the sequence until one simulation is below capacity. Figure 13 illustrates this procedure in the manual driving case; note the horizontal axis ('Flow Rate') runs from high to low, in the same order as our sequence of runs from left to right.
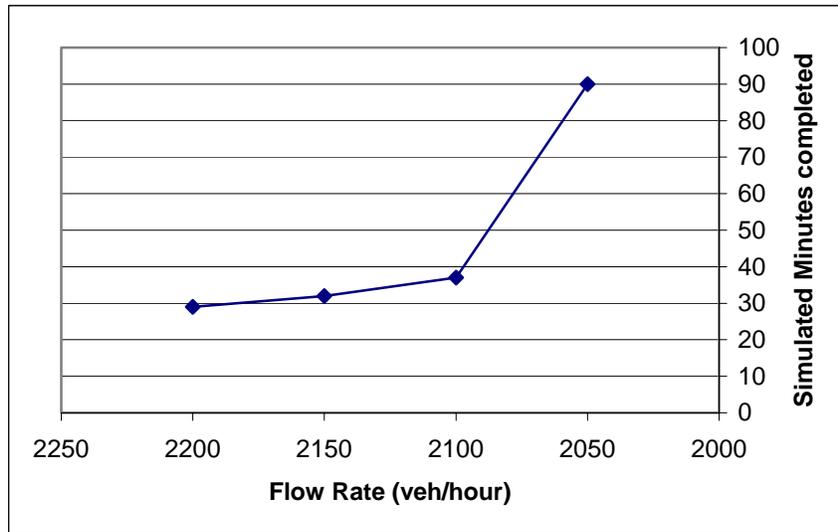
FIGURE 13 Time before source overflow.

Each of these four simulation runs generates 100% manually driven vehicles; plots for other control types show a similar pattern. The first three runs (for flow rates of 2200, 2150 and 2100 vehicles/hr) were terminated early because of queue overflows. The final run was stopped at 90 minutes without a queue overflow. In fact, queue lengths for the upstream and merging traffic rarely exceeded 5 vehicles over this 90 minute run. We conclude from this set of runs that capacity for this case is between 2050 and 2100 vehicles per hour.

5.2.7   Selection of Simulation Cases

In order to maximize our computational resources, we decided to study cases in which the percentage of each control type is a multiple of 20. This level of granularity is enough to show trends without too many runs.

5.3   Results

The flows and other diagnostics for each sub-capacity run (the first run not terminated by an overflow message) for each control type proportion are summarized in Table 2.

TABLE 2 Simulation results.

Control type proportions			Nominal flow (v/h)	Actual flow (v/h)	Wait time (seconds)				Ramp queue length			
						percentile				percentile		
Man'l	AACC	CACC			mean	90th	95th	100th	mean	90th	95th	100th
100	0	0	2050	2050	21.2	76	111	205	1.1	5	7	10
80	20	0	2150	2147	26.9	81	134	214	1.4	5	7	11
60	40	0	2250	2237	106.1	251	265	354	5.5	15	16	21
40	60	0	2250	2241	86.5	222	254	298	4.5	13	16	21
20	80	0	2250	2245	102.1	314	394	528	5.3	13	22	34
0	100	0	2200	2191	85.5	221	339	402	4.5	13	19	22
80	0	20	2200	2192	53.5	144	172	284	2.8	8	11	15
60	20	20	2350	2340	148.6	266	284	353	7.7	14	17	21
40	40	20	2400	2393	168.7	379	463	657	8.9	22	29	41
20	60	20	2500	2477	336.5	567	659	748	17.2	31	33	36
0	80	20	2450	2440	42.5	145	156	232	2.2	8	11	15
60	0	40	2500	2487	134.5	254	280	451	7.0	14	16	24
40	20	40	2600	2592	113.9	302	382	521	5.9	17	21	32
20	40	40	2750	2733	161.2	293	404	640	8.4	16	22	29
0	60	40	2800	2795	76.6	226	261	374	4.2	13	14	20
40	0	60	2800	2793	55.5	183	226	349	3.1	10	14	18
20	20	60	3050	3046	384.4	583	660	741	19.5	31	32	36
0	40	60	3200	3222	88.1	198	247	363	4.6	12	14	22
20	0	80	3400	3416	153.4	390	434	596	8.1	20	24	31
0	20	80	3850	3843	400.4	695	757	831	20.3	36	39	43
0	0	100	4550	4562	227.6	372	393	642	12.2	21	24	31

Nominal flow is not contingent on simulation behavior but is based on the input settings for upstream traffic and ramp entry rates. This value determines, in a probabilistic way, the inter-arrival times between vehicles.

Actual flow is the number of vehicles passing 200 m downstream of the junction per hour of simulation, measured after a 120 second simulation warm-up (to eliminate start-up transients

from the simulation statistics). Actual flow differs slightly from nominal flow because of the probabilistic behavior, because some queued vehicles never enter the road, and because of vehicle interactions (which depend on control type proportions). However, the closeness of the actual flow to the nominal flow supports the interpretation that these are good representations of achievable capacity.

Wait times are sampled when a vehicle enters the road from the entry ramp queue; queue lengths are sampled at 1second intervals. Means and percentiles refer to the distributions of these samples.

The high wait times in cases with a high proportion of CACC vehicles are due to the small gaps between successive (and therefore communicating) CACC vehicles, which makes merging difficult, under our assumption of manually controlled, uncoordinated merging. This observation highlights the need for some kind of cooperative merging system when CACC reaches high market penetrations.

Note that in cases with more human drivers, the high end of the wait time distribution tends to be somewhat more elongated than in the cases with more ACC control. This can be seen by looking at successive ratios of mean wait time, 90th percentile wait time, 95th percentile wait time, and so on. A possible explanation is that ACC vehicles tend to stabilize flow to some extent, making merging more predictable, if slow.

Some of the information in this table is depicted in Figures 14 and 15. Figure 14 shows the effects of two gradual transitions, one from the purely manual case to the purely AACC case, and the other from the purely manual case to the purely CACC case. Figure 15 displays the same AACC results with the vertical axis magnified for detail.
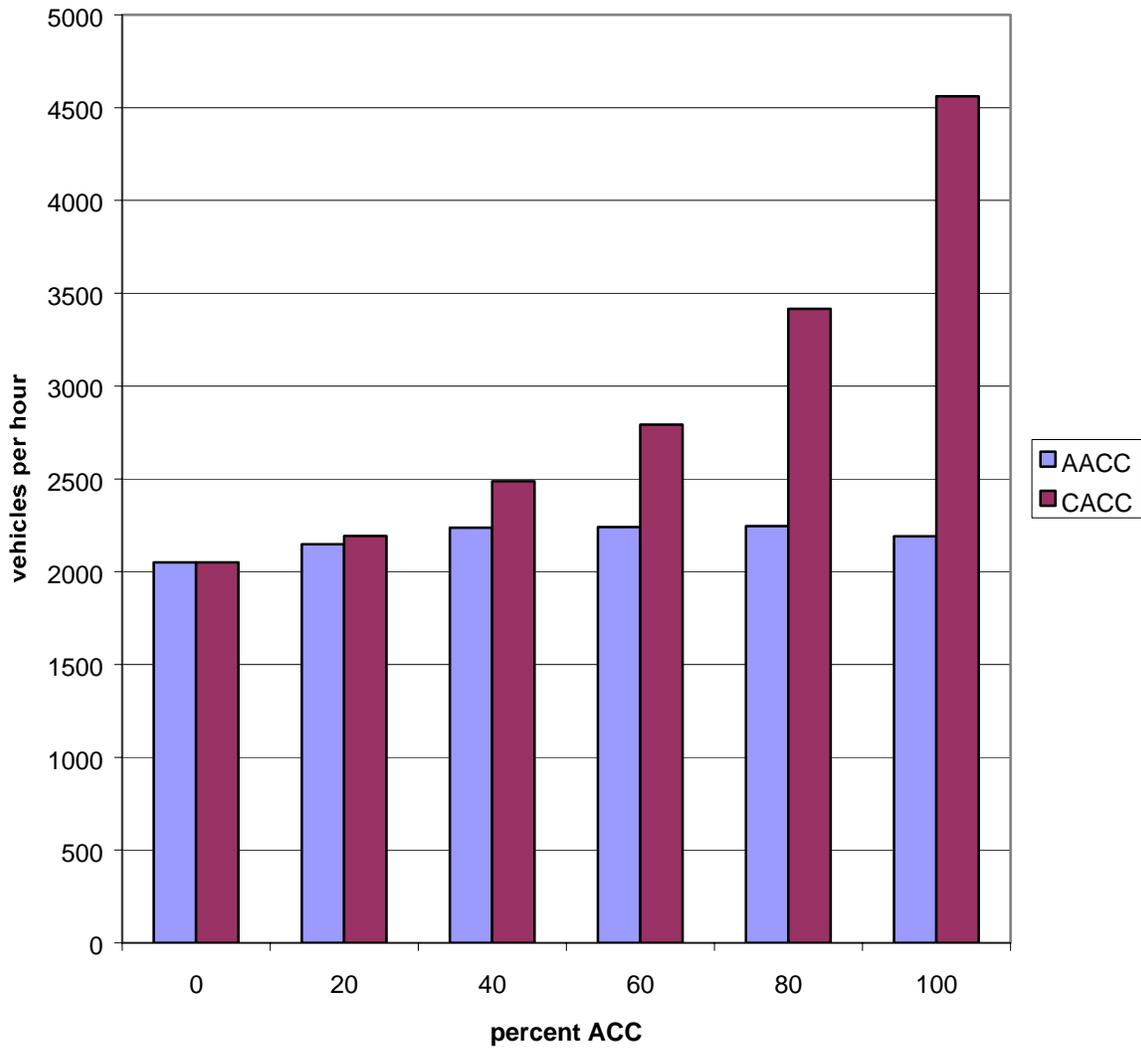
**FIGURE 14 Effect of increasing proportion of one ACC type**
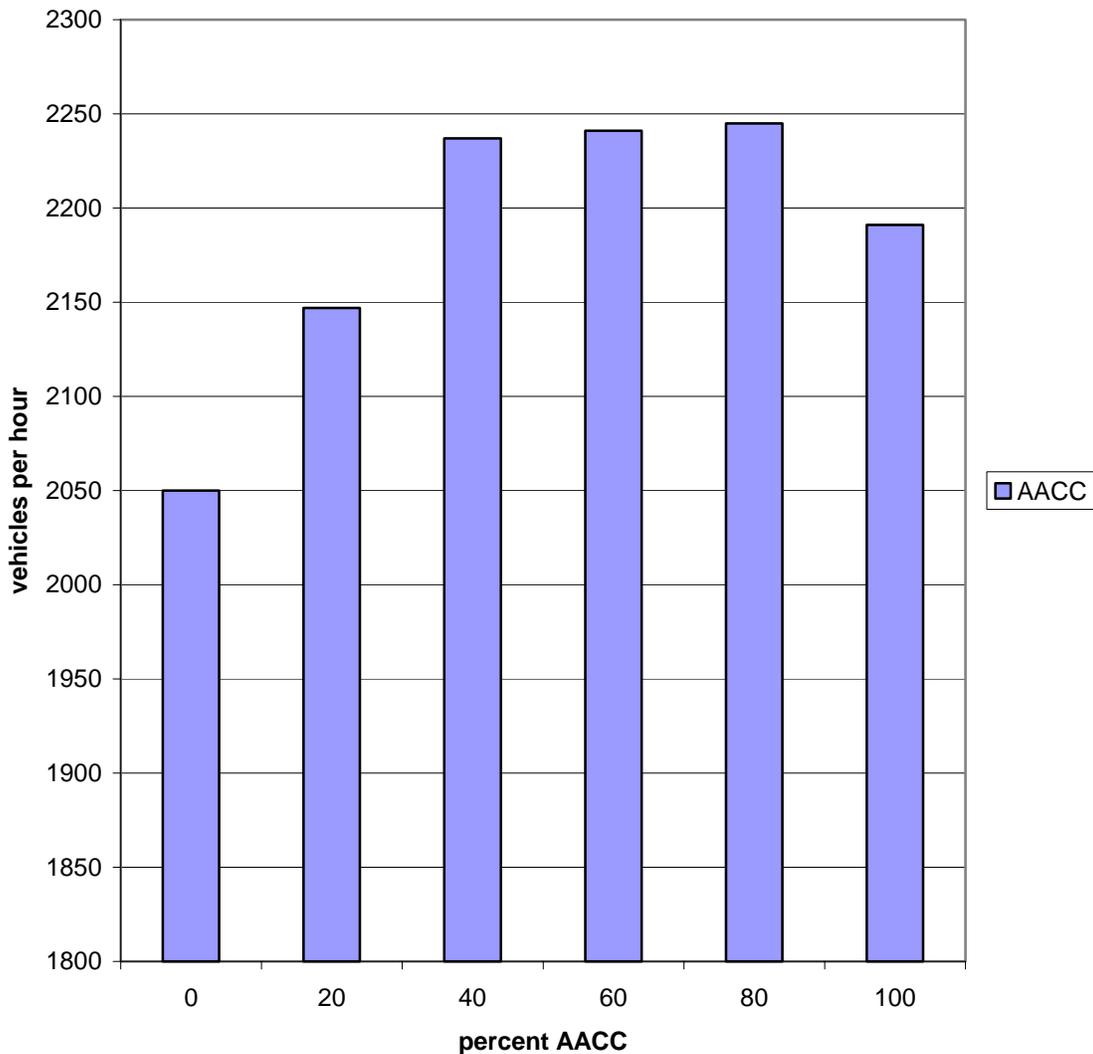
FIGURE 15 Effect of increasing proportion of AACC vehicles with no CACC vehicles

Note that neither of the curves in Figures 14 and 15 appears to be linear. In the CACC case, the upward concavity can be explained by the mode switching in the CACC controller. If the CACC is following another CACC vehicle, it establishes communication and will attempt to maintain a time gap of 0.5 seconds. Otherwise, there is no possibility of communication and the desired gap is 1.4 seconds, just as for AACC. As the fraction p_C of CACC vehicles increases, the chance that any given CACC vehicle will be following another increases linearly. Therefore, the expected fraction of *communicating* CACC vehicles among all vehicles increases quadratically in p_C.

The AACC curve seems to peak at 40% through 80% market penetration. A mixture of AACC vehicles with manually driven vehicles improves flow because they tend to smooth out disturbances. The decrease after 80% is probably because the set time gap of the AACC is 1.4 seconds, whereas the mean desired time gap for manual driving is 1.1 seconds. Without manual drivers (in the 100% AACC case), no additional advantage is gained by the smoothing effect.

Results for all cases are summarized in Figure 16.
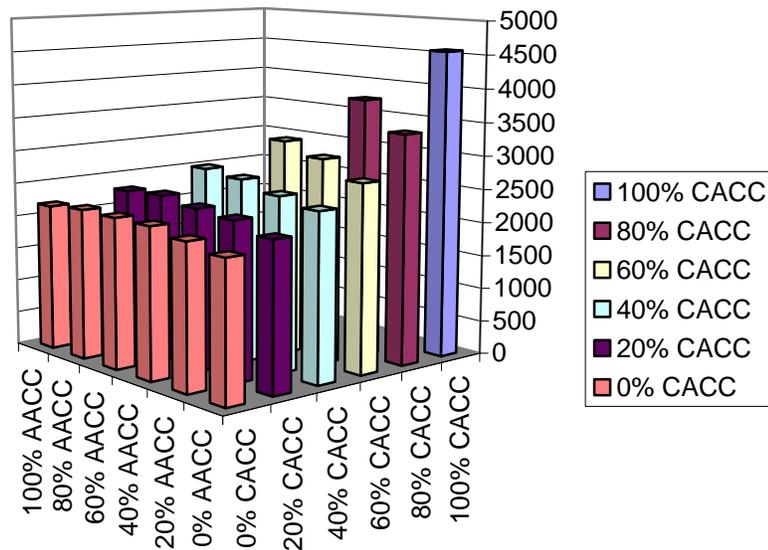
FIGURE 16   Estimated capacities for all control type proportions.

For the three 100% market penetration cases (Table 2), nominal capacity estimates for the manual driving, AACC, and CACC cases were, respectively: 2,050, 2,200, and 4,550 vehicles per hour. Recall that for the AACC case a time gap of 1.4 s was used and in the CACC case, a minimum of 0.5 s time gap was used. The estimate for manually driven vehicles is consistent with current estimates for such driving. The AACC estimate emphasizes the point that adaptive

cruise control is not really designed to increase throughput, whereas for CACC we see a substantial increase in capacity. However, if alternative time gap values were used in lieu of 1.4 s, substantial variability in capacity estimates would be seen (See Figure 17). The 1.4 s value was selected since it represents the middle performance range for the ACC systems that have been tested and are entering the market. In practice, the time gap setting will be selected by each individual ACC driver, and it is very unlikely that the mean value chosen across the entire driving population would be near either of the extreme values (1.0 or 2.0 seconds).
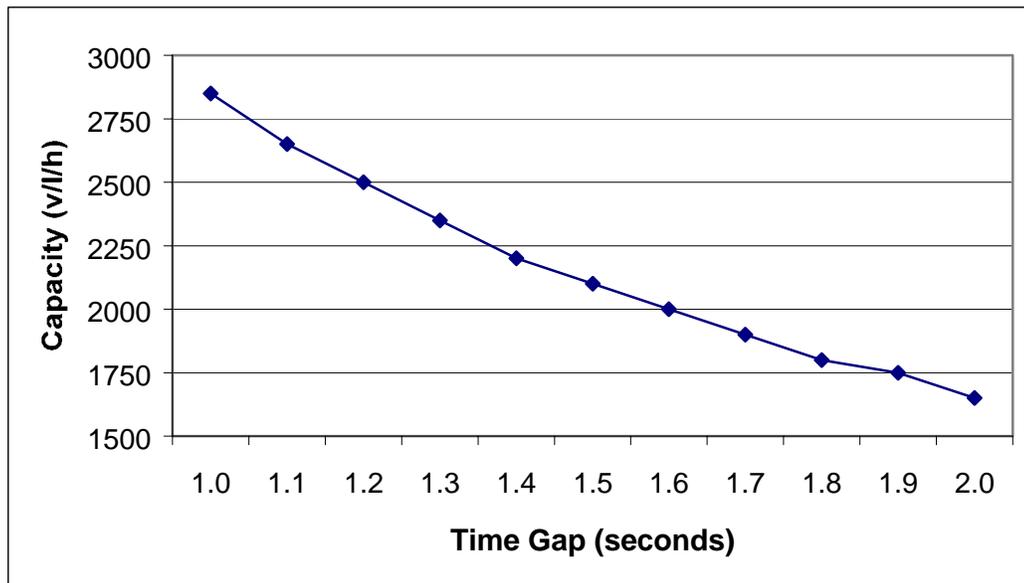
FIGURE 17 Capacity estimates for AACC with different time gap values.

5.4   Comparison with Similar Studies in the Literature

Based on our review of the literature conducted earlier in this study and reported on in Section 4.1, we can observe numerous differences among the assumptions and parameter values used in our study and previous work that lead to different findings and generally make the job of comparing the results across the different studies more difficult. These differences include the distribution of vehicle types in the traffic stream (including heavy trucks or not), the car-following dynamics and nominal time gap maintained by the normal human drivers, the time gap setting chosen for the ACC, the dynamics of the ACC vehicle following algorithm, and the number of lanes of traffic simulated.

The results presented here are intended to represent a "best case" for the capacity impacts of ACC, being based on use of a very stable, high-performance car following controller in a traffic stream with no trucks, and a simulation approach designed to push capacity up to the boundary at which traffic disturbances cause the flow to break down. Therefore, it is not surprising that these results show somewhat higher lane capacity increases for the AACC cases than prior reported studies (30, 32). However, it is very significant that even with these conditions the estimated capacity increases with AACC remain quite modest (at best, less than 10%).

6.0   CONCLUSIONS AND FUTURE WORK

It is possible to define credible progressive deployment sequences that can advance from today's manual control of vehicles to the future fully automated vehicles, without falling into the trap of the "evolutionary/revolutionary" dichotomy. This requires identifying a sequence of steps that can be taken by both vehicle and roadway infrastructure stakeholders with risks that are more than balanced by the expected rewards at each step.

Of all the issues surrounding highway automation, the least well understood to date has been the sequence of steps by which an automated highway system can be developed and implemented. This involves a complicated mix of technological, human factors, economic and political factors. This report has set out to define the most important deployment issues and indicate how they can be addressed in general. Based on consideration of the real constraints to progress, a generic sequence of steps is suggested, and the different ways in which those could be applied to transit buses, heavy trucks and light-duty passenger vehicles have been shown schematically.

This report has also presented models to predict the effects that new driver control assistance systems such as adaptive cruise control can have on traffic flow dynamics and highway capacity. Results have been shown for the validation cases used to test the models individually, for the capacity estimates for the 100% market penetration cases for each of the three modes of operation: manually driven vehicles, AACC, and CACC, and for the capacity impacts of

different combination of market penetrations for AACC and CACC mixed with manually-driven vehicles.

Several significant conclusions can be drawn from these results, with implications for the introduction of ACC into highway traffic and for the longer-term advances toward highway automation:

1. Even under the most favorable conditions, with ideal ACC system design and performance, it appears that autonomous ACC can only have a small impact on highway capacity. Assuming that average ACC users choose a mid-range time gap of 1.4 seconds, highway capacity can be increased by at most 9.5% when the market penetration of autonomous ACC is in the 40% to 80% range.

2. Diminishing returns set in quickly with respect to the capacity increases from introducing AACC into the traffic stream. The increase in lane capacity in advancing from market penetration of 0% AACC to 20% AACC is greater than that from advancing from 20% to 40%, and after 40% there are virtually no capacity increases.

3. Increases in the market penetration of autonomous ACC above 80% can lead to a modest loss of highway capacity, based on ACC users choosing an average time gap for ACC that is somewhat longer than the time gap they use when driving manually.

4. Because of the modest effects of AACC on highway capacity, there does not appear to be any justification for providing AACC vehicles with priority access to special lanes such as HOV lanes. In fact, the tendency of (well designed) AACC to attenuate shock waves in traffic tends to argue in favor of distributing the AACC vehicles throughout all lanes.

5. Cooperative ACC systems, using vehicle-vehicle communications to enable closer vehicle following (down to a time gap of 0.5 second), have the potential to produce significant highway capacity increases. The capacity increases quadratically with CACC market

penetration, based on the fact that the reduced time gaps are only achievable between *pairs* of vehicles that are equipped with CACC.

6. Cooperative ACC can represent an important step in a progressive deployment strategy to lead toward highway automation, because it can potentially double the capacity of a highway lane at a high market penetration. The capacity effect is very sensitive to market penetration, which means that it is important to gather as high a proportion as possible of CACC vehicles into the same lane. This provides a strong justification for giving priority access to a special lane for CACC vehicles. For example, a four-lane freeway occupied entirely by manually driven vehicles could accommodate 8,200 vehicles per hour based on the results shown here. However, if one of those lanes were devoted entirely to CACC vehicles, it could accommodate over 4,500 vehicles per hour by itself, and combined with the other three conventional lanes the overall capacity of the freeway could be increased to 10,650.

Much more work remains to be done before this can be turned into an action plan. More quantitative studies of the costs and benefits associated with the different stages are needed and are either currently underway or being planned for. Several topics for future work are of interest:

1. Using an AACC controller with the same mean time gap as manual drivers (1.1 seconds). Whereas our current study compared controllers operating under parameters we feel are realistic based on expected usage of AACC, using the lower time gap setting would allow more direct comparison of the *potential* capabilities of each control type.

2. Applying our simulation software and capacity estimation methodology to parameters and assumptions used in the other studies of ACC effects on traffic.

3. Studying analogous scenarios with multilane traffic, using the lane-change models of (*4*).

4. Comparisons of results derived using various driver models, including updated versions of the Song-Delorme model (20), and revisions of the MIXIC model (26).

Related work underway at PATH involving further validation and tuning of the driver model includes 1. the freeway performance measurement system (PeMS) data for macroscopic traffic behavior[1], 2. planned for field tests for microscopic driver behavior[2], and 3. data from the Berkeley Highway Laboratory covering 2.7 miles of freeway on I-80[3]. The results of this further analysis will then need to be explored with representatives of the relevant stakeholder groups for verification or refutation. The real action can come only after the stakeholders have been convinced of the genuine benefits that they will gain from each significant step on the road toward AHS.

REFERENCES

1.   "National Automated Highway System Research Program: A Review" TRB Special Report 253, Transportation Research Board, Washington, DC 1998.

2.   Tokuyama, Hideo and Shigenobu Kawasaki, "AHS Strategy in Japan: Theory of Evolution", Fourth World Congress on Intelligent Transport Systems, Berlin, Germany, October 1997.

3.   Shladover, Steven E., "Why We Should Develop a Truly Automated Highway System", Paper No. 980641, *Transportation Research Board Annual Meeting*, Washington, D.C., Jan. 1998.

4.   Ioannou, Petros A. (Editor), Automated Highway Systems, New York: Plenum Press, 1997.

5.   Tsao, H.-S. Jacob, "Stage Definition for AHS Deployment and an AHS Evolutionary Scenario", IVHS Journal, Vol. 2(4), 1995, pp. 359-382 and Commentary and Response, pp. 387-391.

6.   Hogan, Robert M., "Impact of Physical Disengagement on Driver Alertness: Implications for Precursors of a Fully Automated Highway System", IEEE Intelligent Transportation Systems Conference, ITSC '97, Boston, MA, Nov. 1997.

---

[1] This data collects, filters, processes, aggregates, and examines loop detector data in California.
[2] This data will be collected through the use of a Ford Taurus test vehicle on loan from Caltrans. It will collect range, range-rate, gap acceptance, and lane changing behavior.
[3] This data is used to improve basis traffic sensor technologies, to obtain high-quality and real-world surveillance data, and to develop, improve, and validate traffic simulation and prediction models.

7.   Ulmer, Berthold, "Autonomous Automated Driving in Real Traffic", *Proceedings of First World Congress on Applications of Transport Telematics and Intelligent Vehicle-Highway Systems,* ERTICO, Paris, Nov. 1994, pp. 2118-2125.

8.   Bayouth, Michelle, and Charles Thorpe, "An AHS Concept Based on an Autonomous Vehicle Architecture", *Third World Congress on Intelligent Transport Systems*, ITS America, Orlando, FL, Oct. 1996.

9.   Tsao, Jacob, "Constraints on Initial AHS Deployment and the Concept Definition of a Shuttle Service for AHS Debut", *IVHS Journal*, Volume 2, No. 2, 1995, pp. 159-173.

10.   Tsao, Jacob, "Axiomatic Approach to Developing Partial Automation Concepts for Deployment of Automated Highway Systems and Partial Invocation of Vision-Based Lane-Keeping and Adaptive Cruise Control", *Transportation Research Record No. 1651*, 1998, pp. 74- 79.

11.   Koopman, Philip, and Michelle Bayouth, "Orthogonal Capability Building Blocks for Flexible AHS Deployment", *ITS Journal*, Volume 4, No. 1-2, 1998, pp. 1-19.

12.   Bayouth, Michelle, and Philip Koopman, "Functional Evolution of Automated Highway System for Incremental Deployment", *Transportation Research Record No. 1651*, 1998, pp. 80 – 88.

13.   McKendree, Thomas, "Modeling National Deployment of an Automated Highway System", *ITS America Annual Meeting*, Jun. 1997, Washington, D.C.

14.   "National ITS Program Plan", First Edition, ITS America and USDOT, Mar. 1995.

15.   Hosaka, Akio, Yasuhiro Nakanomori and Hideki Ameniya, "A Vision of AHS Development Based on the Principal Requirements", *Fifth World Congress on Intelligent Transport Systems,* VERTIS, Seoul, Korea, Nov. 1998.

16.   Unpublished material

17.   Liang, Chi-Ying, and Huei Peng, "Design and Simulations of a Traffic-Friendly Adaptive Cruise Control Algorithm", *Proceedings of the ASME Dynamic Systems and Control Division*, 1998 International Mechanical Engineering Congress and Exposition, Anaheim, CA, Nov. 1998.

18.   Miller, Mark, and Steven Shladover, "Analysis of Throughput Achievable with Automated and Manual Vehicles Sharing a Lane", *Proceedings of ITS America Annual Meeting*, Washington, DC, Jun. 1997.

19. Schulze, Matthias, "CHAUFFEUR – The European Way Towards an Automated Highway System", *Fourth World Congress on Intelligent Transport Systems*, ERTICO, Berlin, Oct. 1997.

20. Bongsob Song and Delphine Delorme, "Human Driver Model for SmartAHS based on Cognitive and Control Approaches," *Proceedings of the 10th Annual Meeting of ITS America*, Boston, May 2000.

21. Fancher, P. et al., *Intelligent Cruise Control Field Operational Test (Final Report),* UMTRI, 1998.

22. Richardson, Mike, et al., "Adaptive Cruise Control Systems: A Case Study", *Proceedings of the 6th World Congress on Intelligent Transport Systems,* Toronto, Canada, November 1999.

23. Rao, B.S.Y. and Pravin Varaiya, "Flow Benefits of Autonomous Intelligent Cruise Control in Mixed Manual and Automated Traffic", *Transportation Research Record No. 1408,* pp. 36-43, 1993.

24. Chang, Tang-Hsien and I-Shyen Lai, "Analysis of Characteristics of Mixed Traffic Flow of Autopilot Vehicles and Manual Vehicles", *Transportation Research Part C: Emerging Technologies,* Volume 5C, No. 6, December 1997, pp. 333-348.

25. Zwaneveld, Peter and Bart van Arem, "Traffic Effects of Automated Vehicle Guidance Systems", *Fifth World Congress on Intelligent Transport Systems,* Seoul, Korea, October 1998.

26. Arem, Bart van, et al., "An Assessment of the Impact of Autonomous Intelligent Cruise Control", *TNO report INRO-VVG,* Netherlands, March, 1996.

27. Bose, Arnab and Petros Ioannou, "Analysis of Traffic Flow with Mixed Manual and Semi-Automated Vehicles", *California PATH Research Report UCB-ITS-PRR-99-14,* 1999.

28. Yokota, Toshiyuki, et al., "A Study of AHS Effects on Traffic Flow at Bottlenecks", *Fifth World Congress on Intelligent Transport Systems,* Seoul, Korea, October 1998.

29. Yokota, Toshiyuki, et al., "Evaluation of AHS Effect on Mean Speed by Static Method", *Fifth World Congress on Intelligent Transport Systems,* Seoul, Korea, October 1998.

30. Cremer, M., et al., "Investigating the Impact of AICC Concepts on Traffic Flow Quality", *Fifth World Congress on Intelligent Transport Systems,* Seoul, Korea, October 1998.

31.   Darbha, S. and K.R. Rajagopal, "Intelligent Cruise Control Systems and Traffic Flow Stability", *Transportation Research Part C: Emerging Technologies,* Volume 7C, No. 6, December 1999, pp. 329 – 352.

32.   Minderhoud, Michiel, et al., "Impact of Intelligent Cruise Control in Motorway Capacity", *Transportation Research Record No. 1679,* 1999, pp. 1 – 9.

33.   Godbole, Datta N., et al., "Breaking the Highway Capacity Barrier: Adaptive Cruise Control-Based Concept", *Transportation Research Record No. 1679,* 1999, pp. 148 – 157.

34.   *National Automated Highway System Consortium C3 Interim Report*, Phase 1, October 1996-March 1998.

35.   Fancher, P. et al., "Initial Study toward a Methodical Approach for the Engineering of Driver Assistance Technology," UMTRI-99-11, April 1999.

36.   Hoffman, E.R. and R. G. Mortimer, "Scaling of Relative Velocity between Vehicles," *Accident Analysis and Prevention.*, Vol. 28, No. 4, pp. 415-421, 1996.

37.   *Highway Capacity Manual SpecialReport 209*, TRB National Research Council, 1985.

38.   Hoberock L.L., "A Survey of Longitudinal Acceleration Comfort Studies in Ground Transportation Vehicles", *Transactions of the ASME,* 1977.

39.   Ahmed, K. I., Modeling Driver's Acceleration and Lane Change Behavior, dissertation, MIT, 1999.

40.   Sayer, J.R., et al., "An Experimental Design for Studying How Driver Characteristics Influence Headway Control", *Proceedings of the IEEE ITSC '97 Meeting*, Boston, November 1997.

41.   Fishman, George S., *Concepts and Methods in Discrete Event Digital Simulation*, John Wiley & Sons, 1973.