

Functional Evolution of an Automated Highway System for Incremental Deployment

TRB Paper Number: 981060

Michelle Bayouth
Research Engineer
Carnegie Mellon University
Robotics Institute

5000 Forbes Avenue
Pittsburgh, PA 15213 USA
(412) 268-7680
(412) 268-5571 (fax)
mbayouth+@ri.cmu.edu

Philip Koopman
Assistant Professor
Carnegie Mellon University
Department of Electrical & Computer Engineering
and the Institute for Complex Engineered Systems
5000 Forbes Avenue
Pittsburgh, PA 15213 USA
(412) 268-5225
(412) 268-6353 (fax)
koopman@cmu.edu

ABSTRACT

A combination of market forces, cost constraints, and other factors necessitate incremental evolution of a fully automated highway system (AHS) rather than instantaneous deployment. Thus, an understanding of the interdependencies among required AHS functional capabilities is essential for planning. This paper proposes a set of three AHS functional evolution reference models that include essential as well as supplemental functions. The reference models include lateral motion handling, longitudinal motion handling, obstacle handling, and selected infrastructure support functions. This family of three models is used to present the needs of baseline autonomous tactical vehicle operation, the benefits of adding inter-vehicle communications, and the benefits of adding infrastructure support. The reference models reveal a critical need for vehicle motion prediction capability, and suggest that both communications and infrastructure support are beneficial but not mandatory for achieving an AHS. Furthermore, there appear to be a number of safety and efficiency benefits that can be realized with only partial automation and in some cases no automation. These results could help set priorities and guide strategies for incremental introduction of AHS technology into vehicles and roadways.

Keywords: AHS functional evolution; incremental deployment

ACKNOWLEDGMENT

This work was sponsored by USDOT under Cooperative Agreement Number DTFH61-94-X-00001 as part of the National Automated Highway System Consortium.

INTRODUCTION

For many years, scientists and engineers have envisioned building an automated highway system (AHS) to increase both the safety and efficiency of the nation's highways. In such a system, the vehicles become driving robots, capable of sensing and reacting to the surrounding environment while the driver is free to do other tasks. Automating the vehicle has significant potential advantages: it can reduce accidents caused by driver error and can potentially increase traffic-carrying capacity and fuel economy by eliminating human driver inefficiencies. Automating the vehicle also presents difficulties: shifts in legal liability, issues of technical feasibility, and questions of political and social acceptance make the design of an AHS highly constrained, and often subject to heated debate.

There is not yet consensus on exactly what policies and configurations will be used in the operation of a fully deployed AHS. However, it is clear that an automated system will require a number of common functions such as the ability to stay in a traffic lane and to avoid collisions. Furthermore, a number of cost, technical, social, and customer constraints make it seem likely that any deployment of AHS will need to be an incremental one, as opposed to a fielding of a completely automated system from the beginning (*National Automated Highway System Consortium Stakeholder Workshop #3*, Minneapolis Minnesota, September 18-20, 1996, unpublished).

This paper addresses the technical evolution of the AHS by asking the following questions: what common functions are all automated vehicles required to have? and what are the evolutionary dependencies among these functions? The answers to these questions form a proposed family of three reference models for AHS evolutionary development. In this context, the reference models illustrate a reasonably straightforward way to implement the incremental introduction of technology into the vehicle-highway system. It is not necessarily the only possible way, but rather a general representation that encompasses essential issues. These reference models may be useful for both planning and for providing common ground for detailed discussions about technological dependencies and

tradeoffs. These models do not include control policies (*e.g.*, whether vehicles follow each other at short distances in “platoons”), but rather focus on the common underlying technical functions (such as maintaining safe headway spacing among a group of vehicles).

PREVIOUS WORK

In order to deploy AHS capabilities, the uncertainties in the research and development of new technology must be managed well. Additionally, it is impractical to introduce fully automated vehicles on all highways instantaneously. Incremental deployment, then, is a significant issue, and several alternative strategies have been proposed. One strategy advocates the deployment of fully automated vehicles on dedicated lanes, but restricts the deployment to heavily used roadway segments equipped with special-purpose AHS guidance infrastructure.⁽¹⁾ Another strategy is to introduce AHS capabilities onto mass transit vehicles for use on existing High Occupancy Vehicle (HOV) lanes, subject to the supervision of a safety driver.⁽²⁾ A third general strategy involves gradually increasing the degree of automation of new and refitted vehicles over time, with both AHS vehicles and manually driven vehicles sharing essentially all interstate highways.⁽³⁾⁽⁴⁾

This paper does not assume that any one of the above deployment strategies will be implemented. Rather, it presents the set of functions and sequencing constraints that are likely to be involved in deploying an AHS. Whichever deployment strategy is used, the system will need to contain some subset of the reference models’ functionality to be considered a partial AHS. And, no matter the deployment strategy selected, substantially all of the functions will need to be implemented to achieve a complete AHS.

This paper begins by presenting the baseline functional evolution reference model of an autonomous robotic vehicle, assuming that inter-vehicle communications are not universally available. An expanded reference model is then presented that includes the use of inter-vehicle communications, and is used to illustrate functions that are enhanced or enabled for the first time. Finally, a fully elaborated reference model is presented that adds communications with roadside intelligence (highway infrastructure support), enhancing and enabling even more functions. These three reference models illustrate important technical dependencies, highlight the effects of communication and infrastructure support on the deployment of autonomous vehicle systems, and can be used as a roadmap for tracing the development and potential usage of the functions inherent to an AHS.

FUNCTIONAL EVOLUTION

The evolution of the AHS is broken into three diagrammatic reference models: vehicle automation, the addition of inter-vehicle communications, and the addition of infrastructure support. This is done to delineate the evolutionary linkages among functions beginning with the vehicle itself and then adding system-wide applications. This delineation does not suggest that complete vehicle automation must come before communications or infrastructure-based systems, but rather is used to separate development issues and clearly define how the addition of system-wide functions eases (and does not ease) the development of core vehicle systems.

Each of these reference models is further segregated by three functional categories: lateral motion handling, longitudinal motion handling, and obstacle handling. The last reference model adds a fourth category, infrastructure support, in which roadside-based assistance systems are presented. Note that these categories are deliberately chosen to separate technically relevant deployment functions. Lateral and longitudinal motion consider vehicle automation in each of these domains while also handling the effects of surrounding vehicles. Obstacle handling is separated from the core of the vehicle maneuvering functions because of the unique technical issues surrounding obstacle detection and vehicle response. Infrastructure assistance provides unique capabilities that are separate from the basic automation of the vehicle itself.

The three diagrammatic reference models also depict two important relationships: functional dependencies, and functional beneficiaries. In each of these diagrams, the presence of a solid arrow from one function to the next indicates a functional dependency, where the first function must be technically viable for the next function to be. (This is not to say that the first function must be marketable; this is a separate issue.) For example, lane departure warning cannot be achieved without first being able to detect the lane boundaries. A solid arrow therefore links the two functions.

A dashed arrow indicates a functional beneficiary, where the technical development of the first function provides some benefit to that of the second, however the viability of the first is not required for the second to work. For example, the technologies developed for vehicle and vehicle motion detection will likely speed the development of obstacle detection and relevance determination; the unique issues associated with obstacle handling require a separate development effort however, that is not directly linked to vehicle detection.

What is important to note is how these arrows change from one diagram to the next: the addition of communications and infrastructure-based systems alleviates the technical development of vehicle-based systems in some cases yet in others provides little benefit. This provides useful insight into where the larger AHS system can make the in-vehicle technology simpler, but more importantly delineates where it does not do so. This is an important point in the use of these diagrams as a decision-making tool for directing the development efforts of AHS-like systems.

Each of these diagrammatic reference models is presented with cognizance of time-to-market issues. Within the best judgment of the authors, those functions displayed further to the right on the diagrams are likely to be available after those depicted on the left. This is true both within each category and across categories.

Baseline Reference Model: Vehicle Automation

The baseline reference model presents the evolution of vehicle automation capabilities in terms of lateral motion handling, longitudinal motion handling, and obstacle handling. Within each of these categories, Figure 1 shows the key technical functions of an AHS and their dependencies. This diagram presents the core in-vehicle functions that rely neither on inter-vehicle communication nor on the infrastructure for any support in autonomous operations.

While there are clear sequences of functions within each category, there are also cross-functional dependencies that increase in complexity with proximity to the end-state of full tactical driving capability (*i.e.*, fully autonomous operation). The various functional capabilities and their relationships are discussed below.

Lateral Motion Handling

The lateral (side-to-side) motion of the vehicle has a number of different functions, from vehicle-centric maneuvers such as lane keeping to those involving merging in heavy traffic. First, if the vehicle is to stay within the lane, it needs to know where the lane boundaries are. **Lane detection** is currently achieved through a number of different technologies, including a vision system(5), magnetic nails buried in the roadway which are then sensed by the vehicle(6), or a radar-reflective stripe (unpublished work at Ohio State University).

With the advent of lane detection capability, the system can then detect where it is within the lane, leading to **lane departure warnings** when a vehicle strays out of the lane unintentionally. This is an attractive function to have available for incremental deployment as 31% of all highway fatalities are a result of single-vehicle, run-off-road accidents.(7) Marketable systems might be in the form of warnings which alert the driver when a lane change is attempted without prior activation of a turn signal, or might involve a driver-based model that adapts to the characteristic driving patterns of the driver.

Once a vehicle knows the lane boundaries and is able to determine its own position within the lane, **lane keeping** also becomes possible. Lane keeping is the ability of the vehicle to drive down the center of the lane, taking into account upcoming curves and the required pre-steering used to maneuver into them gracefully. A vision system(5) uses a preview area of the lane image for this purpose. Magnetic markers and radar reflective strips could encode upcoming curve information in the roadway, or rely on navigational position information and an electronic map in order to alter the vehicles' trajectories appropriately.

Simple lane changing is the ability of the vehicle to move smoothly out of one lane and into another in light traffic conditions. The technical requirements for such a system include side-looking sensors that detect a gap, and the ability to cross between adjacent lanes and begin lane keeping in the new lane. Such a system could be considered “simple” if it changes lanes only in the absence of nearby vehicles, thus being assured of no risk of collision during the lane changing operation. Simple lane changing requires elementary, side-looking vehicle detection. (**Vehicle and vehicle motion detection**, a precursor functional dependency, is discussed at length in the following section.)

Lane changing in traffic is a more complicated task. The vehicle must not only sense whether a suitable gap exists in traffic, but reason as to whether the gap will continue to exist throughout the maneuver execution. (This requires the use of vehicle motion prediction, a function discussed in the following section.) The vehicle will have to employ rear- and side-looking sensors to accurately detect and track vehicles in adjacent lanes, and tracking/prediction models to determine whether those vehicles have a closing rate that might reduce the gap space unacceptably or otherwise make the maneuver less safe. The vehicle must also be concerned with traffic two lanes over, as a vehicle from the far lane might also merge into the same gap while the vehicle is executing the maneuver, creating a risk of lateral collision.

Merging is a special case of lane changing in traffic in that the lane change must be executed within a fixed space. By the time the lane ends, the vehicle must successfully merge or come to a stop and wait for traffic to clear. This function requires the ability to change lanes in heavy traffic as well as satisfy the constraint of changing lanes within a fixed distance.

Other lateral motion handling functions that do not reside on this direct dependency path include lateral collision warning and stability augmentation. **Lateral collision warning** provides a way to warn of adjacent vehicles encroaching on lateral safety buffer space (for example, an adjacent vehicle that is attempting an inappropriate lane change, or has lost lane keeping ability). Lateral collision warning differs from lane departure warning: the former is concerned with other vehicles encroaching upon the automated vehicle's safety buffer space, while the latter is concerned with the vehicle departing from its own lane unexpectedly.

Stability augmentation is a function wherein the vehicle gently resists changing lanes when a turn signal is not used. The steering wheel resists the motion of the driver, making it feel as if the vehicle is driving with its wheels in ruts, helping keep the vehicle in its lane. This type of function might also prove useful in high crosswind situations, where the vehicle assists the driver in stabilizing a position within the center of the lane.

Longitudinal Motion Handling

The longitudinal (front-to-back) motion of the vehicle also has a variety of functions which range from simplistic in-vehicle handling to tactical driving within a congested traffic scene. **Speed keeping** is the most elementary function within this category, involving the maintenance of a constant travel speed. It is widely deployed in the form of "cruise control." **Headway keeping**, also known as adaptive cruise control, is a function which adapts the speed of the vehicle to match that of a lead vehicle while maintaining a safe distance. Headway keeping is currently being deployed on a limited scale in foreign markets.

Headway keeping, like all of the advanced functions in this category, depends upon reliable **vehicle detection & vehicle motion detection**. This is the ability to ascertain fundamental information about surrounding vehicles and their behavior. This capability will likely evolve from simple look-ahead functions to include look-behind and look-to-the-side as well. The term "look" is used loosely in this context, and refers

to an ability to obtain information about surrounding areas in a particular direction. It does not mean to imply that vision-based systems must necessarily be used; indeed radar, lidar, and sonar systems may prove far more useful than vision systems, especially in reduced visibility situations such as rain and fog.

Given the ability to detect vehicle and vehicle motion, **longitudinal collision warning** becomes possible. This function will probably have two distinct sub-functions: forward- and rearward-looking warning. In the forward-looking case, the system warns the driver if the vehicle is closing too fast on a lead vehicle. In the rearward-looking case, it is concerned with whether a following vehicle is closing too fast. These warning functions, in and of themselves, rely on the driver to respond to the hazard. As technology develops, this capability can lead to several safety enhancements: avoidance via signaling, avoidance via hard braking, and avoidance via simple swerve. These three functions are not integral to the main deployment path, however they provide interim benefits until the time when more advanced automation functions are enabled. **Avoidance via hard braking** is where the vehicle brakes to avoid an accident with a lead vehicle or obstacle; **avoidance via simple swerve** is where the vehicle swerves into the shoulder or into a clear, adjacent lane; **avoidance via signaling** is where the vehicle sends a warning signal to a following vehicle. Avoidance via signaling may use rapidly flashing brake lights to gain the attention of the rear vehicle, or may employ communications as will be discussed later.

Vehicle motion prediction is the heart of advanced in-vehicle intelligence. This function, which is critical to the development of all high-level fully automated systems, predicts the traffic scene forward in time and will likely employ driver models, vehicle capability models, probabilistic analyses, and a fundamental understanding of driving “rules”. Vehicle motion prediction enables the vehicle to determine the safest maneuver amongst the numerous available possibilities, and has the potential to predict an unsafe situation before it unfolds and can be sensed.

Notice that vehicle motion prediction is the only function which provides retroactive functionality to already-deployed capabilities. Both lateral and longitudinal collision warning can be significantly enhanced with additional in-vehicle intelligence. It is also clear that all higher-order fully automated systems depend on this capability in order to function. For this reason, vehicle motion prediction is identified as a critical link in the deployment of AHS and is shown in black in order to highlight its importance.

Cordial driving, the next step in maneuvering, is best described as driving which accommodates the desires of other vehicles in a friendly way. For example, if a vehicle is in a merge lane with its turn signal activated, manual drivers tend to (but do not always) give way so that vehicle can merge. Cordial driving by an automated or semi-automated vehicle can be enabled once the vehicle can discern the intentions of surrounding vehicles. In Figure 1, the assumption is that communications are not universally and dependably available; therefore, the vehicle must infer intention by vehicle location, vehicle motion, and vehicle signaling. (Signaling could be augmented with, for example, a radio transmission but such signaling is not assumed to be universally available.) This type of inference is, therefore, a large part of vehicle motion prediction.

Avoidance via lane change in traffic is a capability that enables the vehicle to change lanes given a hazardous situation, such as the sudden and unexpected slowdown of a preceding vehicle. This capability is a reasonable and attractive alternative to hard braking, as it may avoid secondary and tertiary accidents that could ensue when braking is used. In order to perform this function safely in traffic, the vehicle system must be acutely aware of the operating environment, and predict the behavior of surrounding vehicles with enough accuracy that the lane change can be done safely and efficiently. Note that lane changing in traffic, a lateral motion handling capability, must be supported in order for this function to be available.

Finally, **full tactical driving** is introduced. Tactical driving is the ultimate form of full automation within an individual vehicle. At this stage in the deployment process, the vehicle not only tracks and reacts to other vehicles, but also proactively plans out series of maneuvers which are executed to achieve a goal or goals. For example, if the vehicle (or the driver) desires to increase speed and the vehicle is “boxed in,” the tactical capability of the vehicle will enable it to plan a way out of the box in order to achieve its goal with a combination of speed changes (including perhaps temporarily slowing down) and lane changes.

Obstacle Handling

Obstacle avoidance capabilities reduce or eliminate safety hazards caused by obstacles on the automated highway roadway. This includes rocks, vegetation, dropped vehicle parts, disabled vehicles, and animals such as deer.

One way to reduce the need for obstacle avoidance is to implement obstacle exclusion. To a limited degree this is already deployed with fencing and highway department maintenance of the interstate highway system. Obstacle exclusion can significantly reduce the frequency of obstacles on the roadway, but it seems implausible that any exclusion method can be 100% effective. Thus, more sophisticated forms of obstacle handling will be needed in most foreseeable end-state AHS implementations. Note that obstacle exclusion, as an infrastructure-based system, is depicted on a later diagram that introduces infrastructure functionality.

Obstacle detection and threat determination is a much more difficult task than vehicle detection due to the technical difficulties of sensing obstacles and identifying whether those obstacles present a threat. A mylar toy balloon, for example, is easy to sense but would preferably not trigger a severe braking maneuver that could risk minor injury to vehicle occupants. On the other hand, deer present a much greater threat but are more difficult to sense and track. It seems likely that obstacle detection and threat determination will be available well after vehicle detection and collision avoidance functions are fielded. Notice that a dotted arrow links the two functions as there may be value to the research done for the detection of vehicles in the development of obstacle detection capability.

Given that obstacles can be detected and properly categorized, **simple obstacle warning** may be possible, whereby the driver is warned of hazards in and around the roadway. This function can provide obstacle information for the driver to use, or it could be coupled with the **avoidance via braking** capability. In that avoidance via braking was fielded previously as a longitudinal motion handling function, it is likely that it could be marketed as an obstacle avoidance function as soon as reliable warning systems are available. For this reason, the avoidance via braking box appears directly above the obstacle warning function box on Figure 1.

Obstacle motion detection & prediction may be a particularly difficult capability to develop. Unlike vehicles, which are physically constrained in realizable maneuvers, obstacles may not behave in readily predictable ways. Deer may run into the road and stop abruptly. Loose tires can bounce randomly, depending on road surface, tire wear, and angle of incidence. Mylar balloons, plastic bags, and tarps are at the whim of the wind and thermal pressures created by hot roadways. None of these are easy to quantify, and as such,

create very difficult detection and prediction problems. Notice that vehicle motion detection and vehicle motion prediction might provide benefit to the development of this capability.

Assuming that obstacle motion prediction can be achieved, the vehicle can achieve **avoidance via lane change**. This becomes a unique capability that depends heavily on good prediction, not only of the motion of the obstacle, but also of how other vehicles will react to that obstacle.

Baseline Reference Model: Discussion

Figure 1 presents several important items. First, it delineates the evolution of functions within an automated vehicle system, beginning with warning functions and evolving into full vehicle control. The functional dependencies, denoted by solid arrows, identify the critical linkages between functions and clearly define which systems must be mature in order for further capability to be technically feasible.

A major finding in the process of identifying these linkages is the importance of vehicle motion prediction. This capability is critical to the successful implementation of all higher-order full automation functions including lateral, longitudinal, and obstacle handling tasks. It has the added distinction of providing recursive benefits to several already-deployed systems, notable in that no other function does this. The criticality of this link has not been previously appreciated in the development of AHS subsystems, with efforts generally concentrated in the development of sensors and actuators that manage the lower-level automation functions. This reference model serves to identify and highlight this function's importance.

Communications Reference Model

Figure 2 identifies changes to the functional evolution depicted in Figure 1 when inter-vehicle communication links are introduced. Perhaps the most striking thing about Figure 2 is that there are only two additional functions introduced, indicating that communications are not a necessary element to achieve fully automated tactical driving capabilities. On the other hand, there are many functions highlighted to indicate that they may significantly benefit from reliable inter-vehicle communications.

First, consider those existing functions from Figure 1 that benefit from communications being introduced. Numerous functions gain from having “pre-maneuver announcement information,” which

provides greater insight into the future motion of surrounding vehicles, and “inter-vehicle negotiations,” which allow vehicles to gain assurances from surrounding vehicles. These existing functions are highlighted in Figure 2, and each of these will be discussed in turn.

Consider a **simple lane-changing** maneuver. With the addition of communications, a vehicle may “negotiate” with surrounding vehicles to gain assurances that the gap it senses will remain available throughout the maneuver. This may allow simple lane changing to occur in a heavy traffic situation where two adjacent lanes are moving at about the same speed, a significant increase in capability over lane changing in the absence of surrounding vehicles. Even in the case where vehicle-to-vehicle negotiations are not employed, knowing that a local vehicle is moving into a gap near to your vehicle provides additional information about the ever-changing traffic scene, and enables a vehicle to begin readjusting its headway in anticipation of the new situation.

Lane changing in traffic may also benefit from pre-maneuver announcements. If traffic in adjacent lanes is moving at significantly different speeds, communications enhance maneuver safety by providing intent information before the maneuver can be sensed, and are particularly useful when a maneuver may not be anticipated by the vehicle motion prediction models. So too, **merging** benefits from knowing a vehicle’s exact intention, as opposed to presuming intention. Vehicles that communicate can send assurances that a gap will be maintained, or even created, so that other vehicles can merge into the mainline traffic stream.

A number of longitudinal motion functions also benefit from the introduction of communications. If following vehicles can respond based on the knowledge that a lead vehicle is about to respond to an incident, safety margins are increased over the situation where the vehicle must first sense the incident before responding to it. **Avoidance via hard braking, avoidance via lane change, and avoidance via shoulder swerve** can improve system safety by announcing maneuvers before executing them, providing a small but possibly significant amount of additional time for the surrounding vehicles to proactively respond. Additionally, swerving and lane changing information enables surrounding vehicles to respond proactively and create a gap, as opposed to reactively, in a variety of situations. In the **avoidance via signaling** case, a trailing vehicle may be provided with information that the closing rate is too high, or a vehicle that wants to pass might signal its desire for a leading vehicle move out of the lane.

Communications can also provide more detailed and timely information to the **vehicle motion prediction** algorithm. The less that this function has to extract intention, the more appropriate and faster the responses can be. If intent of other vehicles is well known, **cordial driving**, **avoidance via lane change in traffic**, and **full tactical driving** are all simplified.

In addition to existing functions benefiting from the addition of communication, two new functions emerge. Prior to this point, the sharing of braking information has been promoted as a way to provide faster reaction times in emergency situations. **Coordinated braking** might use communications to provide smoother headway keeping in non-emergency situations, such as stop-and-go or slow-and-go traffic. In some AHS alternatives, coordinated braking could be used to permit reduced headway and thus denser traffic; such uses, however, would require an extremely reliable communications system to avoid collisions in emergency situations.

Shared obstacle information might also provide tangible safety benefits, even to manually driven vehicles. A lead vehicle that detects an obstacle might accurately pinpoint its position and transmit that information to the following vehicles. Without communications, a following vehicle might not have time to detect and respond to the obstacle, given that the lead vehicle may be blocking the line-of-site until it swerves out of the way.

Communication Reference Model: Discussion

As shown in Figure 2, the addition of inter-vehicle communications enhances a number of core, in-vehicle functions and enables several new functions. These enhancements and additions, over time, will provide increasingly greater system-wide effects as more vehicles acquire communications capabilities and are able to share information. Two factors influence this reality: the evolutionary nature of these technology-based systems, and the long life span of automobiles. In that a mixture of automation and communication capabilities will be found on any given roadway due to a mixed-aged vehicle fleet, all vehicles must be built to handle all other vehicle types and ages. This is why none of the dependency arrows in Figure 2 change from those presented in Figure 1. The addition of vehicle-to-vehicle communications enables refined estimations of

the upcoming traffic situation, however the expected presence of non-communicating vehicles in traffic precludes using communication to alleviate the core in-vehicle requirements.

In addition to the effects and requirements of an evolutionary deployment scheme, there is an additional and important caution about considering an AHS that relies upon communications. The obvious communications medium is radio, but radio has a number of potential reliability problems including correct and timely frequency allocation among multiple proximate vehicles; interference from other sources such as commercial radio transmitters, illegally boosted citizen band radio transmitters, military transmitters or malicious jamming; and interference from natural sources such as lightning strikes. Because of these issues it would not be prudent to rely upon radio communications in a safety-critical role unless reliability is demonstrated across a wide range of environmental conditions or a cost-effective alternative to radio communications can be found. Instead, communications should be treated as information to optimize performance of an AHS, not as a part of safety-critical control loops.

Even if inter-vehicle communications are deemed to be not wide-spread or reliable enough for safety-critical functions, the use of communications when available has the potential to increase both individual and global system safety by providing an additional source of detailed information.

Infrastructure Assistance Reference Model

Figure 3 identifies how infrastructure-based capabilities might affect the vehicle evolution. Five new functional boxes are added, and four existing functional boxes from Figure 2 are highlighted to indicate that they either benefit from infrastructure support or could be used in a roadside-based application.

The four existing functions which can benefit by or be adapted for use in the infrastructure are merging, vehicle motion prediction, obstacle detection and relevance determination, and shared obstacle information. The **merging** function may benefit from having infrastructure sensors providing gap information, especially on blind merges where the vehicle's own sensors may be limited. Although current highway designs avoid blind merges, many older highways have extremely difficult merge points that are the cause of both congestion and safety problems. Infrastructure-based merge assistance will require infrastructure-based **vehicle motion prediction** capability in order to provide appropriate gap information to equipped vehicles. This will not

obviate the need for vehicle-based vehicle motion prediction, but will use the developed capability in an additional capacity.

Obstacle detection and relevance determination, as an infrastructure-based system, can either complement or completely replace vehicle-based systems. Reasonable arguments can be made for either alternative: infrastructure-based systems can be upgraded more readily as new technology becomes available; vehicle-based systems are applicable and improve safety on all roadways. It is likely that obstacle detection will have a place in both infrastructure- and vehicle-based systems, given the safety benefits to be gained by its implementation.

Sharing obstacle information can also have widespread benefit by sending obstacle information to the infrastructure and maintenance personnel in addition to surrounding vehicles. With this capability, every equipped vehicle becomes a scout, providing information about hazardous situations and relaying that information to the infrastructure for quick removal and to surrounding vehicles for early warning. This capability also improves safety in light density traffic, where vehicle-to-vehicle communication may not occur because of large gaps in the traffic stream.

Five new functions also are enabled by the addition of infrastructure-based intelligence. As mentioned previously, **obstacle exclusion** may obviate the need for obstacle avoidance capability within the vehicle by preventing obstacles from being on the roadway in the first place. For a policy of exclusion to be effective, all vehicles would need to be inspected for loose loads and vehicle parts and made reliable enough that they would not suffer failure. As this seems improbable to an effective enough level of safety, obstacle exclusion is posed as a measure to reduce the occurrence of obstacles that is coupled with additional obstacle handling functions.

Real-time **traffic advisories**, which are communicated to the vehicles by the roadside, can aid the vehicle and driver in determining an appropriate route given existing and historical traffic data. In this system, the driver and/or the vehicle decides the route given the information presented. **Route optimization** involves a more organized effort, where the driver provides his origin, destination, and desired arrival time to the infrastructure, which in turn provides the optimal route to the driver and vehicle.

Lane use optimization is a fourth infrastructure-based system. This allows the infrastructure to determine the fastest moving lane, aggregated over many miles, and dynamically assign vehicles to move into

that lane to even out traffic flow. Even more intricate is **system-wide traffic flow optimization**, which dynamically assigns routes and lane choices in an effort to maximize the system throughput given the current system-wide demand.

Infrastructure Support Reference Model: Discussion

The introduction of infrastructure support provides a number of system-wide benefits that cannot be achieved with in-vehicle intelligence alone. Most notably flow and route optimization can only be handled with coordination between vehicles and the infrastructure. Additionally, merging, obstacle detection, and the sharing of obstacle information can be improved with the addition of infrastructure-based systems.

Notice that the introduction of infrastructure-based systems changes the dependency arrows into the merging function. Merging becomes a beneficiary to lane changing in traffic without being dependent upon it for implementation; the arrow between the two functions is changed to a dashed arrow to reflect this. A new dependency arrow is drawn from vehicle motion prediction, as an infrastructure-based system, to the merging function. In that automated merging may actually become available sooner with infrastructure support than would have been possible with vehicle-based systems alone, this function is moved leftward on Figure 3 relative to Figures 1 and 2.

DISCUSSION

These three-staged reference models are oriented toward evolving in-vehicle systems from warning to vehicle control with an appreciation for the effects of inter-vehicle communication and infrastructure-based assistance. There is a risk in presenting these systems with the “ultimate end-state” of full automation on highways presumed: the real goal of this work is roadway safety and efficiency. While these models provide insight into the elements of evolutionary deployment of AHS, there appear to be additional potential for safety and efficiency gains without the use of automation.

Note in Figure 3 that all of the route and system optimization functions are independent of vehicle automation. It is possible that the existing highway system might be improved with the development of infrastructure-exclusive systems coupled with a simplistic infrastructure-to-vehicle communication system and a

willing driver. Merging assistance, for example, could be provided to the driver with advice on when to accelerate into traffic at difficult entrance points with limited visibility.

Likewise, inter-vehicle communication is a mechanism that has not been considered without being coupled to vehicle automation functions. Maneuver signaling information, shared between vehicles, could be presented in a head-up display format for drivers with limited peripheral vision. A non-freeway application might employ an infrastructure-to-vehicle communication link that presents real-time stoplight information in an easily visible location for these impaired yet licensed drivers.

CONCLUSIONS

Three successively more comprehensive reference models have been presented to depict how an AHS might evolve using incremental introduction of functionality. These models indicate precedence constraints on the introduction of capabilities and depict how the introduction of inter-vehicle communication and infrastructure support can increase the efficacy of an AHS. The models also demonstrate that addition of additional capabilities in most cases does not supplant previously introduced in-vehicle functionality.

A significant result of developing these reference models is the discovery that vehicle motion prediction capability is critical to the development and implementation of all higher-order automation functions. Most AHS analyses to date have emphasized a need to improve sensors and actuators; the results presented here suggest that effort must also be applied to the development of prediction algorithms for use in both vehicles and infrastructure support functions.

The first reference model suggests that full tactical driving capability requires neither inter-vehicle communication nor roadside-to-vehicle communication. However, if robust communication mechanisms can be provided, communications might be used to significantly improve the quality of maneuvering and collision avoidance capabilities.

With the exception of merging, adding communication capabilities does not obviate the need for any vehicle-based functions when these vehicles are incrementally introduced onto an existing roadway system. The presence of non-AHS-equipped and older-but-equipped vehicles in an incremental deployment scheme necessitates the development of in-vehicle systems that can function effectively without relying on communication.

Even though an end-state of full automation is desirable, the reality is that an incremental deployment will be necessary. As shown by the reference models, there appear to be a number of obtainable safety and efficiency benefits at intermediate deployment points utilizing partial automation and in some cases no automation. Using these models as a roadmap may help plan evolutionary approaches to creating an AHS that satisfy both technical constraints and a need to provide value before fully automated operation is achieved.

REFERENCES

- (1) Shladover, S.E., "Potential Freeway Capacity Effects of Advanced Vehicle Control Systems," *Proceedings of the 2nd International Conference on Applications of Advanced Technologies in Transportation Engineering*, American Society of Civil Engineers, New York, NY, 1991, pp. 213-217.
- (2) Tsao, J., "Constraints on Initial AHS Deployment and the Concept Definition of a Shuttle Service for AHS Debut," *IVHS Journal*, 2(2), Gordon and Breach Science Publishers, Langhorne, PA, 1995, pp. 159-173.
- (3) Bayouth, M. and Thorpe, C., "An AHS Concept Based on an Autonomous Vehicle Architecture," *Proceedings of The Third World Congress on Intelligent Transportation Systems*, ITS America, Washington, D.C., 1996.
- (4) Ward, J., "Step by Step to an Automated Highway System—And Beyond." In: (Ioannou, P., ed.) *Automated Highway Systems*, Plenum Press, New York, NY, 1997, pp. 73-91.
- (5) Pomerleau, D., and Jochem, T. "Rapidly Adapting Machine Vision for Automated Vehicle Steering," *IEEE Expert*, 11(2), IEEE Computer Society Press, Los Alamitos, CA, 1996, pp. 19-27.
- (6) Shladover, S.E., "The California PATH Program: a state approach to IVHS research," (92C018) *Proceedings 1992 International Congress on Transportation Electronics*, Society of Automotive Engineers, Warrendale, PA, 1992, pp. 329-338.
- (7) Carnegie Mellon University & Calspan Corporation, Run-Off-Road Collision Avoidance Using IVHS Countermeasures: Task 1 Final Interim Report, Volume 1: Technical Findings, Carnegie Mellon University, Pittsburgh, PA, October 28, 1994.

Figure 1: Baseline Autonomous Reference Model

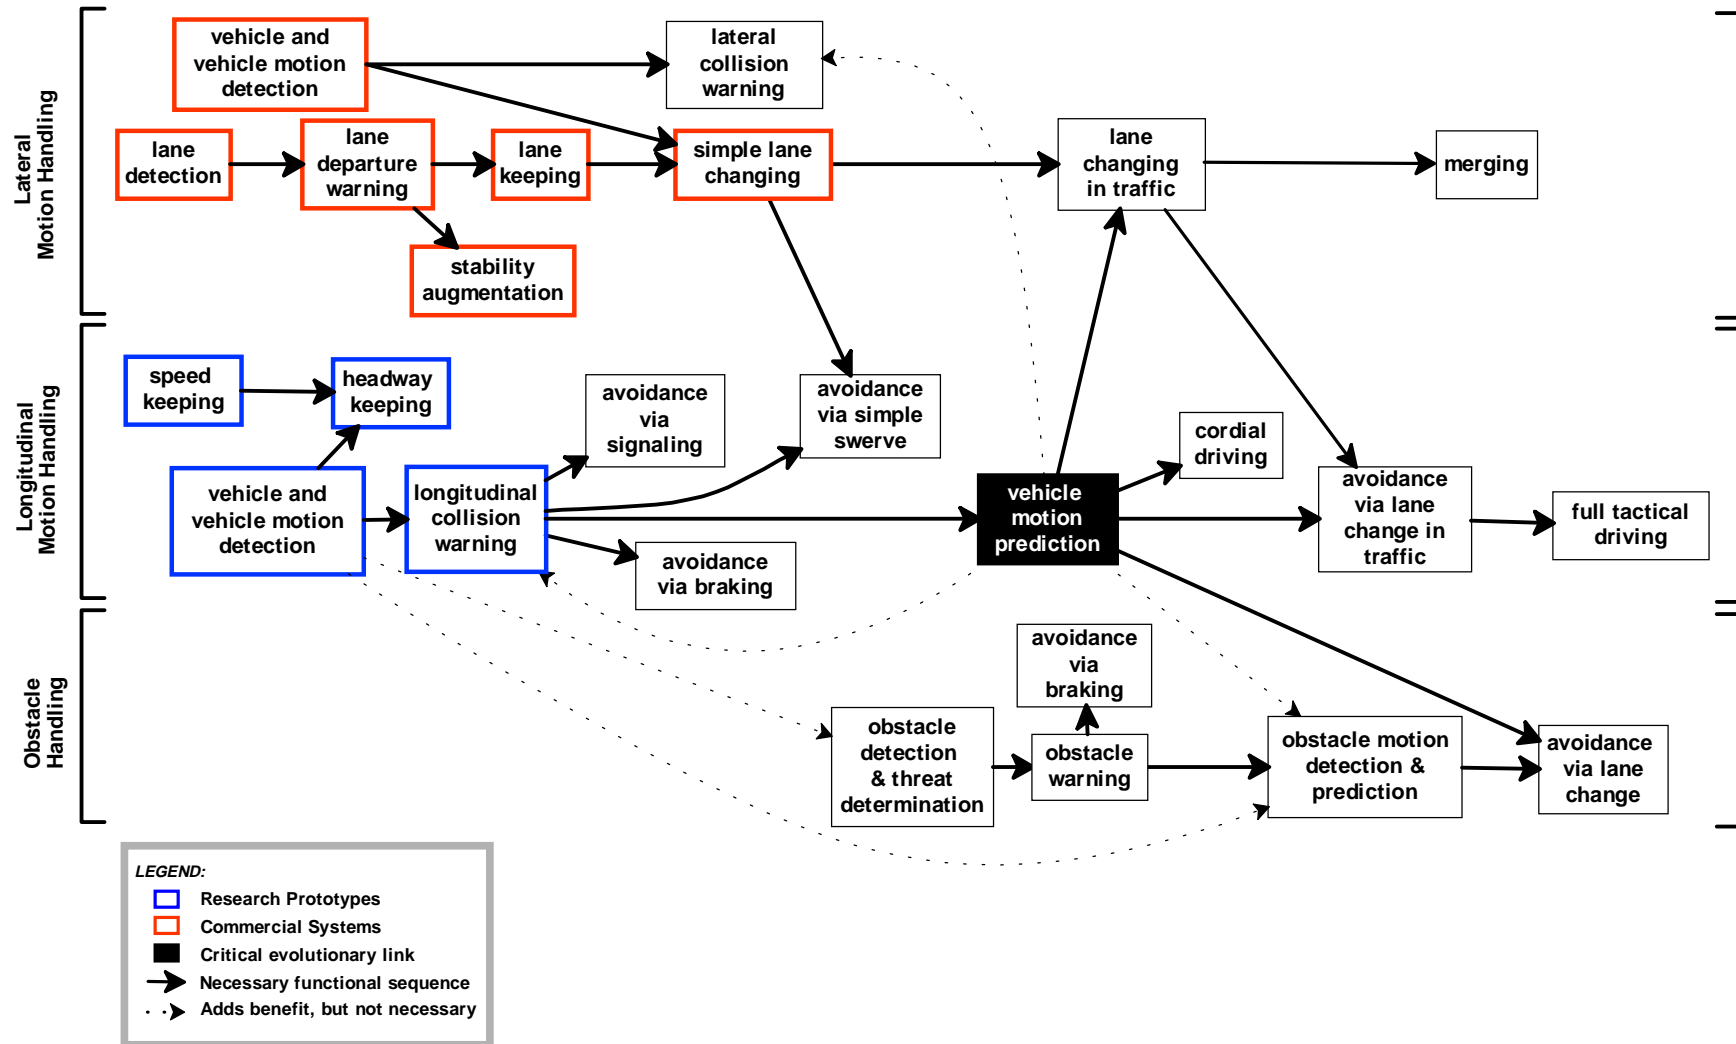


Figure 2: Baseline Autonomous Reference Model with Inter-vehicle Communication

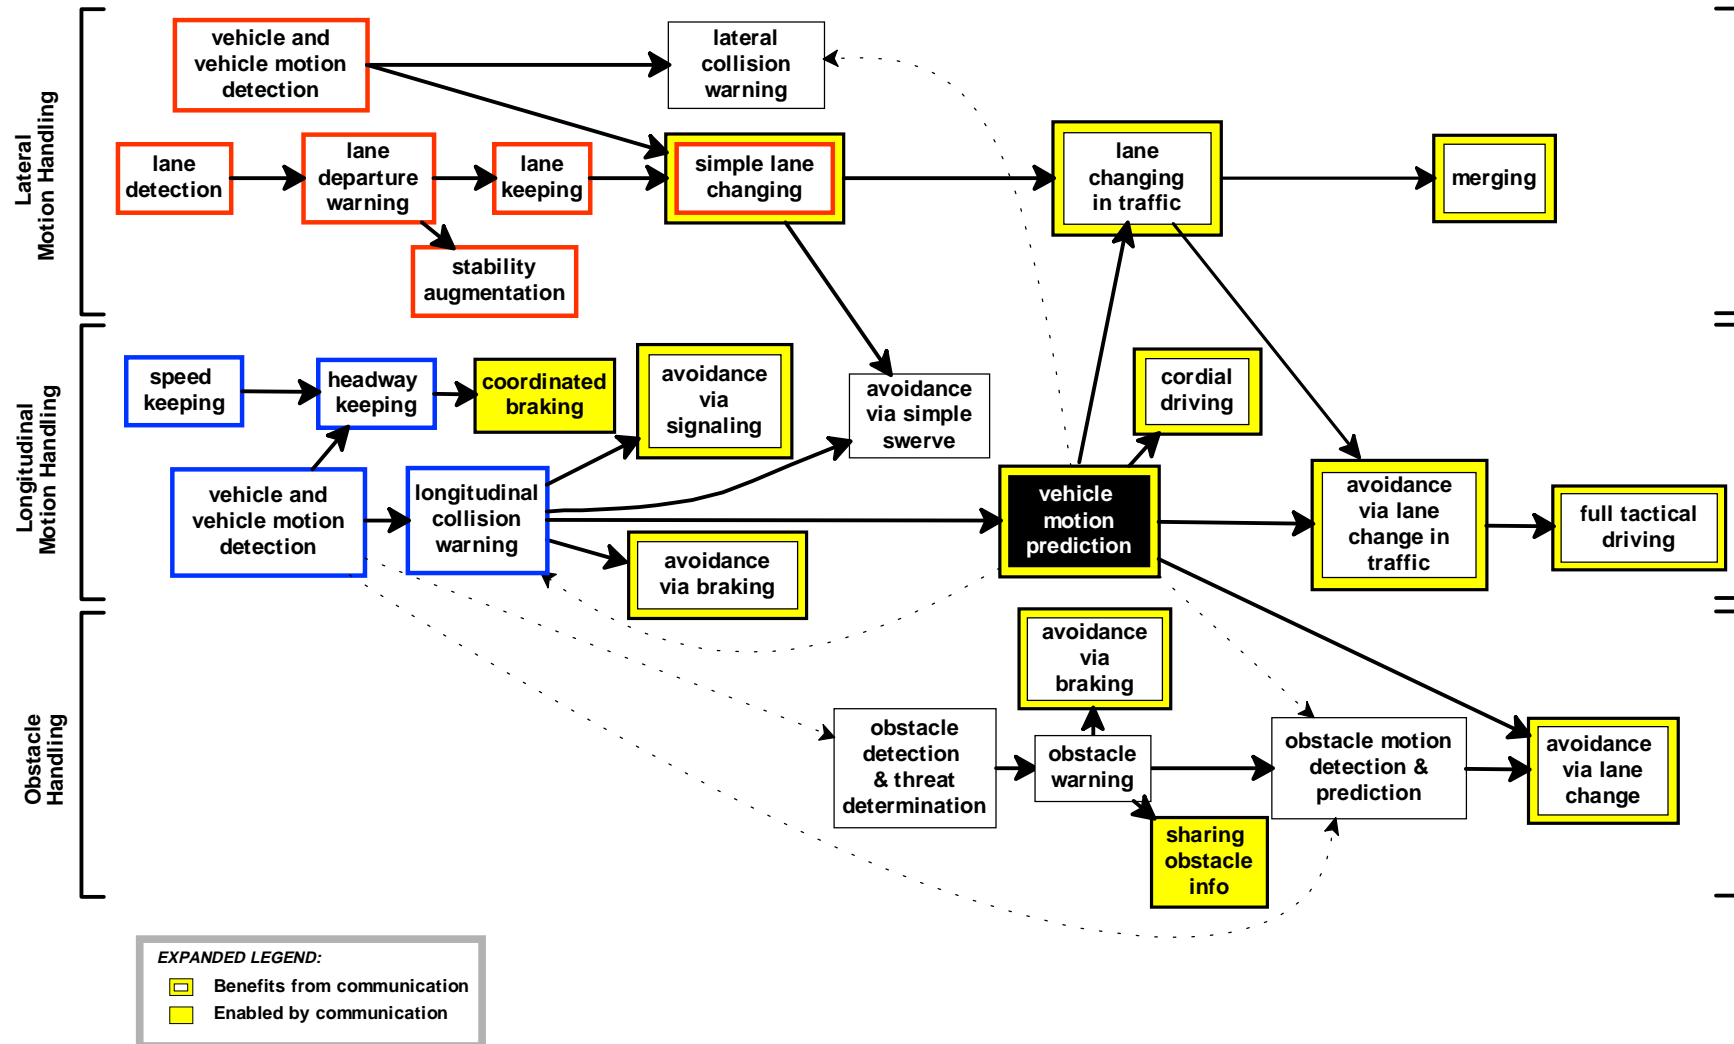


Figure 3: Baseline Autonomous Reference Model with Infrastructure Support

