Chapter 8

Conclusions

This chapter opens with a brief discussion of the automated highway system as a possible solution to current traffic problems and a short evaluation of the reinforcement schemes as an automata learning process. The results and contributions of this dissertation are listed in the second section. The chapter concludes with the recommendations for future research.

8.1 Discussion

The automated highway system envisioned by researchers and institutions around the world suggests a complex structure that requires extensive research efforts. Among many issues, the realization of AHS depends mainly upon the realization of intelligent vehicle control methods. No matter what design approach is taken for vehicle control, there are two important issues that will 'haunt' the researchers. The first problem in front of the AHS is a social issue: user acceptance. Without the user acceptance, the aim of automating the highway system is bound to fail. We will briefly discuss the possible future of the AHS deployment in the light of the current influence of the technology on society. The second problem is the technical feasibility of the designed systems. While designing the intelligent path planner described in this work, the transition from abstraction to implementation was one of our concerns. We will therefore briefly discuss the technical issues that hinder the realization of an automated highway system.

We designed our intelligent vehicle controller using an AI technique called the learning automata. The adaptation of the automata to existing conditions is provided by the reinforcement schemes mentioned in previous chapters. We will thus conclude our discussion with a few ideas on reinforcement schemes.

8.1.1 AHS as a Social Issue

In the last five years, we have witnessed the birth of the automated highway system. Highway automation efforts showed themselves in various forms of technological gadgets (e.g., automatic vehicle identification, electronic toll collection, emergency signaling devices, video enforcement systems, and visual vehicle detection systems [ITSW96, TrafT96]) as far as the 'user' is concerned[1]. The first question in the mind of those familiar with the computation and

---

[1] The author of this dissertation has traveled more than 40,000 miles on US interstate highways while working on design of intelligent vehicle and highway systems, and on the evaluation tools for intelligent transportation systems.

communication aspects of AHS (and ITS) is the 'information content' of such a system. Most of the projects for full automation of the highway system envision total knowledge about the highway and the vehicles.

Present AHS and ITS applications such as traffic monitoring, electronic toll collection and license plate recognition are already employing identification/classification techniques that can be used to generate information about the drivers/owners of the vehicles. Human factors research on AHS systems shows that drivers would not want the infrastructure to know everything about their actions, locations, and daily routines. This problem is called the "big brother syndrome." The objections to usage of this type of data include fear that the security of the data will be compromised, and overall discomfort with the modern trend toward routine collection of personal data. Initial system analyses indicated that check-in procedures envisioned for AHS implementation may include "processing of financial, medical, or driving records to verify potential users' qualifications" [Delco94].

The information content of the data collected for AHS purposes is very important for the user acceptance of the automated highway system. While the collection of driver data may be perceived as an invasion of privacy, some of the problems may be solved by choosing the least invasive AHS implementation with minimal data storage. Such issues are important because they may influence the design of the communication and processing functions of the (longitudinal) control system for a hierarchical infrastructure [Delco94]. For the autonomous vehicle approach that will most probably include some form of visual control system, similar problems exist.

With the current technological progress and its applications to daily life, the 'invasion of privacy' is becoming an important issue. A recent example is the implementation of a global computer network, a.k.a. Internet or World Wide Web. More and more individuals are starting to use the network in their daily routine, and in return, they are providing the WWW server, the Internet service provider or the routing system with some form of information. This could be login data, an e-mail message, or a server log of the visited pages with the address of the client machine. In most cases, it is possible for a computer user to monitor the activities of another from his console. A similar and more common example for today is the use of credit cards. Every time a charge card is used, the corresponding computer system records the time, date, location and type of the transaction which can later be sold to third parties. Yet, many are comfortable with the accessibility that a charge card or an online computer system provides (to all interested parties).

The use of AHS or ITS systems such as automatic toll collection or automatic vehicle identification (or ATIS, ATMS, AVL systems in the near future) is no different than technologies mentioned above. A trade-off between comfort and privacy exists, and will exist in the future. The "evolution" of the AHS will stop at a level where the majority of the users are no longer willing to trade their privacy with the additional comfort the system may bring. As AHS becomes more visible, its complexity may lead to misunderstandings and misrepresentations that

---

The reflections in this section are the results of his AHS experience as a 'researcher' and his driving experience as a 'user' of the present highway system.

may jeopardize public acceptance [Batelle94]. Considering the recent mistakes made with the latest technological revolution, the Internet, it is obvious that a balanced and accurate picture of AHS must be presented to the public.

Recent investigations addressing the user acceptance of AHS, and legal and technical privacy issues in AHS technologies can be found in [Stevenson95], [Wright95], [Garfinkel96], [McGurrin96], and [Wigan96].

8.1.2 Technical Feasibility of AHS

The technology needed to create an intelligent vehicle control system is already available, although still expensive for full implementation. The term 'technology' here refers to our ability to implement the necessary components for vehicle control. On the other hand, vehicle control problems related to AHS have been vigorously investigated by many researchers, as discussed in Chapter 2. However, the research on sensing and communication methodologies for AHS is relatively new, and many problems remain. These problems are solvable, yet they require significant research effort for safe and successful implementation of AHS. Several observations and ideas related to technical feasibility of the AHS implementation are discussed in this section.

Since the vehicles in an automated highway will be autonomous, they will require on-board sensing equipment to gather information about their surroundings. Sensing is vital for AHS approaches that do not consider hierarchical control structures where most of the information is relayed from the higher levels of the hierarchy. Non-hierarchical control approaches, on the other hand, do require simple communication systems. No matter what the sensing and communication requirements are for a specific realization of AHS, several issues such as headway measurements, lateral sensing, local communications, and sensor fusion must be addressed.

The most important issue for autonomous vehicle navigation are the sensor capabilities. The question of "can we trust our observations?" needs to be answered. For example, the headway measurement is the heart of longitudinal vehicle control. There are several different technologies such as sonar, microwave radar and optic sensors for headway sensing, but a single method is not reliable for all environment conditions. Many researchers were able to implement longitudinal and lateral control methods [Lui89, Özgüner95, Özgüner96, Pomerlau96, Weber96], however operability under various environmental conditions is still an important design problem. The solution seems to be the sensor fusion. Although more reliable observations are possible by combining distinct sensor information, finding an intelligent way of 'fuse' sensors is very difficult. Furthermore, an autonomous vehicle must have redundant sensors, and an additional back-up mechanism to guarantee safe operation [Singer95, Agogino96]. Thus, the complexity and the cost are the two main factors in autonomous/intelligent vehicle implementation.

Current commercially available sensing technologies include radar-based headway sensors for adaptive cruise control, global positioning sensors for intelligent navigation, and vision systems for active driver support for lane keeping [ITSW96, TrafT96]. Magnetic sensors for lane keeping, adaptive vision systems for lateral and longitudinal vehicle control, laser radar for headway control, and proximity sensors for lane departure warning are currently used in research vehicles [Özgüner96, Pomerlau96, Tan96, Weber96, Yanagisawa92, Construction96].

The research on sensing technologies is gaining momentum due to the ongoing research on AHS, and due to the fact that most of these technologies find quick commercial applications. The research on sensing technologies for AHS was initially driven by the robotics field; during the last few years more research has been carried out specifically for intelligent vehicle applications. The sensor capabilities for autonomous navigation will increase with the investigation of new alternatives and the progress of the previous technologies.

Communication with the roadside and other vehicles is another important factor for AHS implementation. Again, a reliable communication system is vital for a hierarchical control approach. A hierarchical control structure requires three types of communications: communications for control, maneuvers, and navigation information. For the first type, high bandwidth is required. Furthermore, the loss of data must be minimized to maintain stability of platoons, and to assure collision-free lane changes [Foreman95]. Communicated data must be "trustworthy." A real-life analogy to the problem of reliability of the received information is to wait at the intersection until another car approaching from left makes a right turn although its right turn signal flashes while approaching the intersection. This type of behavior is a clever one on the part of a human driver and it is useful in avoiding collisions, but it is against everything AHS tries to achieve. An automated system repeatedly checking every bit of information against errors cannot be successful in enhancing the traffic flow. On a greater scale, a hierarchical automated highway system must optimize traffic flow by coordinating vehicle actions in an intelligent way. The need to increase traffic flow theoretically increases the risk of accidents. The reason for this is a decreasing possibility of recovering from a wrong decision due to a smaller "reaction window."

Commercial communication systems for AHS and ITS include new radar detectors with road hazard and emergency vehicle signals, traffic information systems, fleet-management systems, in-vehicle mobile offices and on-line services, MAYDAY devices and emergency services, and electronic toll collection [ITSW96, TrafT96]. On the other hand, vehicles currently used for research purposes are equipped with 2.5GHz band radio transceivers, vehicle-to-vehicle radio modem communications [Construction96, PATH96]. Ultra wide band communication is also considered as an alternative [James96]. The network architecture for vehicle-to-vehicle communications is currently being defined [Bolla96, Fuji96], and related problems such as interference and scattering are being investigated [Gavan96, Zoratti95].

Progress in the electronics filed, especially in wireless communications, will positively affect the deployment of AHS[2]. AHS is actually a combination of many engineering fields. For example, the design of the ITS/AHS technologies are carried out by the government organizations in conjunction with the Institute of Electrical an Electronics Engineers (IEEE), the Society of

[2] AHS will use emerging technologies as soon as they become available during its evolution to fully automated highways. For example, drivers with cellular phones are already acting as 'probes' in traffic via emergency phone lines. Such an information source was not available ten years ago.

Automotive Engineers (SAE), the Institute of Transportation Engineers (ITE), and the American Society for Testing and Materials (ASTM) [ITSst96]. National AHS Consortium consists of many aerospace and automotive companies as well as civil, mechanical, electrical engineering, and computer science departments of several educational institutions. AHS is a multi-disciplinary research filed. In addition to technologies mentioned above, its realization depends on many engineering applications such as systems and control, fiber optics, automotive design, data mining, data scheduling, human factors. It slowly creates its own niche in many related fields of engineering boosting the research efforts and resources.

Automotive companies are hoping to make AHS technologies such as cruise control and lane departure warnings commercially available in the next five years. Lane tracking and vision enhancement applications are expected to be available in ten years [Cogan96]. It is interesting to note that the current AHS research has already found answers to the problems related to these applications. The transition from research to commercial availability is only a matter of cost reduction. Considering the rate of technological progress in our century, the future of the AHS technologies looks bright.

Returning to our assumptions on the vehicle capabilities, it is fair to say that the assumptions made on sensing and communications for autonomous vehicles are feasible, but not widely available.

8.1.3 Reinforcement Learning

The intelligent vehicle path controller discussed in this work is based on learning automata techniques. Learning stochastic automata are capable of adapting to changing environments by using reinforcement schemes described in Chapters 3, 4 and 6.

Vehicle control in the context of AHS is a very complex problem. Automated vehicles are part of a dynamic environment where the control decisions must be made very quickly and by coordinating with other vehicles. In an environment where there are multiple fast-moving vehicles, making the right decision to avoid collisions and optimize the traffic flow is difficult.

The inherent complexity of the AHS has forced researchers to look for new methods capable of creating flexible and powerful high level control systems. The artificial intelligence methods mentioned in Section 1.2.1 are the results of this need. Precursor System Analyses for AHS concluded that these techniques, especially knowledge-based systems and learning methods, can make a strong contribution to AHS especially in the areas of malfunction detection and fault tolerance, high level control of vehicle actions, overall traffic management, and sensor fusion [Rayteon94]. The work presented here attempts to make contributions to three of these problems except the traffic management.

“In any eventual AHS, there will be widely varying capabilities among the possibly millions of AHS vehicles because there will be a mixture of new and old models, plus there will be widely varying levels of maintenance. Moreover, the world is complex and unexpected things can happen, such as a moose entering the highway, or a refrigerator falling off of a truck?” [Rayteon94] It is clear from this argument that the AHS environment will be more chaotic than structured. It will not be easily predicted, nor analyzed. In the case of mixed traffic (human drivers and automated vehicles coexisting on the same highway), the problem grows enormously. As a result, an automated vehicle must be able to deal effectively with a large number of possible traffic situations.

The last statement above emphasizes the importance of learning and adaptation. What is needed is a system that can handle situations unforeseen by the designers. For this type of capability, the choice of artificial intelligence technologies is obvious. While expert systems, and knowledge-based planning [Forbes95] are effective, learning methods are capable of discovering new situations and optimal responses using simulated environments [Sukthankar96].

One of the most important characteristics of the reinforcement learning is the choice between the actions that are known to return positive responses and the actions with unknown performance. For example, a learning automata starting with *a priori* action probability values needs time to evaluate the actions: it learns from its mistakes.

This process can be examined from two different points of view [Ashby60]. It can be thought of simply as an attempt to find the best solution. In case of failure, the attempt is not rewarded (or is penalized). This view of the learning process is not very attractive. On the other hand, the trial and error method is invaluable in gathering information necessary to achieve optimal action. Choosing non-optimal actions to gain information improves the long-term performance while resulting in temporary penalty responses from the environment. Learning automata is one of the successful examples of such methods.

Another important aspect of the reinforcement learning is that the actions performed by an automaton influence the environment response that the automaton will receive in the future. Rather than receiving an independently chosen set of action-response pairs, the agent has some form of control over what response it will receive and complete control over what actions will be generated in response. “In addition to making it difficult to make distributional statements about the inputs to the agent, this degree of control makes it possible for what seem like small experiments to cause the agent (automaton) to discover an entirely new part of its environment” [Kaelbling94].

8.2 Results and Contributions

The results of the work described in this dissertation can be divided into two groups: (a) learning automata and reinforcement schemes, and (b) intelligent vehicle control and path planning. This section lists the significant results of the completed work. The contributions are indicated together with the associated results.

The application of learning automata techniques to intelligent vehicle control for automated highway systems provided the following conclusions about learning automata, reinforcement schemes, and games of automata:

- The reinforcement scheme is the heart of the learning process, and its characteristics define the behavior of the automata. However, besides the reinforcement scheme, the definition of the reward-penalty structure (*i.e.*, the teacher) is vital in the success of the learning automata application. The automata cannot be expected to learn the best actions with a defective teacher structure. The significance of the reward-penalty structure is emphasized in environments with multiple teachers. We have spent a significant amount of time to design a function that intelligently combines multiple environment responses.

- The lack of applications using continuous environment response is well understood. The problem of combining such responses as well as the difficulty in finding a meaningful representation of the several factors with the same measure hinders the use of S-model environments. There are no known applications that use this model.
- The use of learning algorithms in our application is based on the assumption that the physical environment around an autonomous vehicle can be associated with a stationary automata environment (or multiple switching automata environments). We assume that the automata adaptation to the environment is faster than the changes in the physical environment. Although the current computational technology is sufficient for such an assumption, we defined two new reinforcement schemes for our purposes. These new algorithms have convergence rates up to four times faster than previously designed counterparts, especially in situations frequently occurring in AHS applications.
- One of these new reinforcement schemes has been previously known. However, its convergence characteristics have not been proven due to the difficulty in analytical modeling the learning process. By using nonlinear methods and a special condition resulting from our design, we were able to prove optimality of the algorithm.
- The second algorithm is also the result of our attempts to improve the convergence characteristics of the learning automaton. It is an extension of a previously designed nonlinear reinforcement scheme, but performs better. This algorithm is proven to be absolutely expedient.

Our work on the design and analysis of an intelligent vehicle control using learning automata enables us to conclude the following:

- The intelligent vehicle path controller consisting of two stochastic learning automata is capable of adapting to its dynamic environment by choosing the best (lateral and longitudinal) actions to be sent to the lower level of control.
- This non-model based approach would be especially useful in situations wherein the complete knowledge about the overall traffic flow is not provided by higher level of the control hierarchy (if such levels exists at all). Instead of trying to foresee all possible traffic situations, we take the approach of defining a mechanism that can make intelligent decisions based on local sensor information, keeping in mind the fact that the initial phases of the AHS will include non-automated vehicles as well as intelligent vehicles capable of communicating with others.

- The method is capable of capturing the overall dynamics of the system that include the vehicle, the driver and the roadway. Definitions of the learning and sensor parameters determine the behavior of each vehicle.
- Initial results indicated that capabilities of the vehicle sensors must be extended. As presented in Chapter 5, these extensions are not too tasking, although some may be relatively expensive to realize. A simple rate-feedback method for the headway distance measurements can overcome the oscillatory behavior in longitudinal control.
- Simulations of intelligent vehicles also indicated the need for additional information sources besides the local sensors. No matter what control structure is used for AHS deployment, some form of communications must exists between vehicles. Although visual clues can be used to coordinate lateral actions, the lane changing capabilities of the intelligent vehicles as well as the safety of the actions increase with local vehicle-to-vehicle communications. We have found that, in order to avoid pinch situations, vehicles may coordinate their lane changing actions by simply sending an ‘intention’ signal to neighboring vehicles.
- Our attempt to design an intelligent path controller inadvertently extended, to some degree, to other levels of vehicle control. For example, we have found that if a higher level of control/decision mechanism provides desired lane information, many local solutions (that are not globally optimal) may be solved to optimize overall traffic flow. Furthermore, the capabilities of the lower layers of control are significantly important when designing the intelligent path controller. Although we did not consider feedback from the control layer that actually carries out the vehicle actions, the design parameters of the lower layer affects the design parameters of the path planning layer. All levels of control are interconnected and cannot be treated as a single free-running entity.
- There is a trade-off between what the automated vehicle can accomplish and how simple the sensing/information system is. The more global the information content of the decision mechanism is, the more the vehicle can accomplish autonomously. It is our belief that neither a fully decentralized control method, nor a completely hierarchical system can solve the problem completely. While a decentralized system will require some form of local communications, a hierarchical system will suffer from being unable to foresee all conflicts between vehicle paths.
- Our approach to the vehicle path control problem can be viewed as an autonomous vehicle approach. Automated vehicles are making their own decisions on which action to take by using information from their own sensors and limited communication systems. Although each vehicle is capable of avoiding collisions, the combined actions of multiple vehicles are not always optimal. Previous research efforts in autonomous vehicles discussed in Section 1.2.1 have not considered multiple vehicle interactions. In this work, we also attempted to find a structured methodology for the interactions of multiple autonomous vehicles.
- The method of evaluating possible environment state transitions based on associated automata environments enabled us to define additional decision mechanisms we called ‘flags.’

Speed and lane flags are used to solve the conflict situations arising from the multiple teacher responses and vehicle interactions

- Although our method of evaluating the physical environment's state changes are based on the learning automata environment, similar methods can also be used with other decision mechanisms. By formal descriptions of the decision/control procedure, transition diagrams similar to those given in Chapter 7 can be created to analyze the highway situations.

8.3 Recommendations

AHS research, especially high level vehicle path control, is a relatively new research area, and there are an extensive number of questions to be answered. There are multiple possibilities for extending the work discussed in this dissertation. Here, we will emphasize some of these possibilities that seem relatively important for realization of intelligent vehicle control. Our recommendations for future research efforts on the subject are the following:

- Simulation:
  Simulation is an indispensable way of testing the effectiveness of vehicle control methods since the system becomes highly complex due to the presence of large number of vehicles. Furthermore, the ideas investigated in this work cannot be tested with real vehicles, partly due to the reasons listed in Section 2.5, but mainly because of the futuristic character of the problem. The overall system is neither continuous nor discrete-time; the combination of vehicle dynamics and the computational modules forms a hybrid system whose design and analysis is inherently difficult. For the work described here, we simplified the vehicle dynamics drastically since the main purpose of this research was the study of learning automata techniques for vehicle control. More realistic vehicle dynamics must be incorporated with the computational system, and the study of the interaction between the discrete and continuous parts of the overall system needs to be carried out for a precise control system simulation.
- Sensor modeling:
  For the purpose of studying the learning automata as an intelligent vehicle controller, we assumed almost perfect sensor models. Our sensor module subroutines in the simulation incorporate small percentage of measurement error. Beyond the addition of the measurement noise, we assumed that the sensors defined in this work can be implemented using the current technology. More realistic models of the sensors based on existing radar, laser radar, infrared and sonar techniques may be created to further the simulations' ability to imitate real-life situations. The issues of sensor degradation and back-up methods also have to be answered for a fault-tolerant autonomous navigation applications.
- High level vehicle path control:
  Our study of vehicle path planner indicated that this level of vehicle control cannot be designed nor analyzed without considering lower and higher level of the control methodology. Especially for the hierarchical control structure, feedback from the regulation layer and the

information sent from the link layer of the hierarchy must be taken into account. This study of the planning layer did not consider a possible 'unable to comply' signal from the low level control modules. Furthermore, we assumed that global information such as the desired lane and speed are available. Previous and current research on AHS has not yet answered such issues. Lane and speed assignments by a higher layer in the control hierarchy may prove to be crucial for obtaining optimal traffic flow.

· Vehicle interactions:

The problem of coordination between autonomous vehicles is inherently addressed by a centralized control approach. However, decentralized AHS methods such as autonomous vehicle approach [Bayouth96] or cooperative approach [McKendree96] must still find an answer to this problem. In case of mixed traffic which may be a possibility in the future of AHS, the problem is far more serious. Analysis and design methods described in this work address the vehicle interaction by simplifying multiple vehicle scenarios (situations) into multiple two-vehicle interactions by using additional simplifications on the relative position of the vehicles (the matrices in Section 7.4 and Appendix D). This is possible because of our definition of sensors and the reward-penalty structure used for the learning automata. The idea of analyzing multiple vehicle interactions can be applied to other vehicle control methodologies provided that the "states" of the physical environment, and the "decision environment" resulting from these are carefully defined. More detailed and complex definitions for highway scenarios may be necessary.

· Learning automata:

The idea behind the adaptive vehicle controller presented here is very simple: two automata synchronously adapt to their (teacher) environment. During our study, we have observed that the automata interactions and multiple teacher learning processes in multi-teacher environments are relatively untouched areas of research. The intelligent controller may be extended to multiple automata associated with a single teacher (*i.e.*, sensor module). In this case, the analysis of the interactions between automata will be more difficult. However, translating the problem of combining the environment response to the problem of combining the multiple actions may have its advantages. The application of the S-model environment then may prove to be relatively easier, bringing the advantage of continuously mapping the sensed physical environment to automata environment. Yet, the idea of combining automata actions have not been investigated previously.