U.S. Department
of Transportation
National Highway
Traffic Safety
Administration

NHTSA
www.nhtsa.gov

DOT HS 811 405

October 2010

Advanced Crash Avoidance Technologies (ACAT) Program – Final Report of the Volvo-Ford-UMTRI Project: Safety Impact Methodology for Lane Departure Warning – Method Development And Estimation of Benefits

DISCLAIMER

Technical Documentation Page

1. Report No. DOT HS 811 405	2. Government Accession No.	3. Recipient's Catalog No.

4. Title and Subtitle
Advanced Crash Avoidance Technologies (ACAT) Program – Final Report of the Volvo-Ford-UMTRI Project: Safety Impact Methodology for Lane Departure Warning – Method Development and Estimation of Benefits

5. Report Date
October 2010

6. Performing Organization Code

7. Authors
Gordon, T., Sardar, H., Blower, D., Ljung Aust, M., Bareket, Z., Barnes, M., Blankespoor, A., Isaksson-Hellman, I., Ivarsson, J., Juhas, B., Nobukawa, K., and Theander, H.

8. Performing Organization Report No.

9. Performing Organization Name and Address
Volvo, Ford, UMTRI
Ford Motor Company, 2101 Village Road, PO Box 2053, Dearborn MI 48121-2053

10. Work Unit No. (TRAIS)

11. Contract or Grant No.
DTNH22-06-H-00068

12. Sponsoring Agency Name and Address
USDOT/National Highway Traffic Safety Administration
Office of Human Vehicle Performance Research
1200 New Jersey Avenue, SE,
Washington, DC 20590

13. Type of Report and Period Covered

14. Sponsoring Agency Code

15. Supplementary Notes

16. Abstract
The Volvo-Ford-UMTRI project: Safety Impact Methodology (SIM) for Lane Departure Warning is part of the U.S. Department of Transportation's Advanced Crash Avoidance Technologies (ACAT) program. The project developed a basic analytical framework for estimating safety benefits in the form of computed reductions in US crash numbers assuming the vehicle fleet was fully equipped with Volvo Lane Departure Warning (LDW) systems. Attention was limited to crashes initiated by the lane departure of a light passenger vehicle. The SIM uses computer models, Monte-Carlo methods and extensive batch simulations to fuse together diverse data sources into a Virtual Crash Population. The simulation model incorporates sub-models for the driver, vehicle, environment and technology (DVET model). The simulated population of virtual crashes makes use of historical crash data (NASS GES and Michigan Crash File), naturalistic driving data, objective test data from test track and driving simulator experiments, as well as highway data to populate the environmental sub-model. The vehicle model was based on a representative mid-sized sedan, and test data was used to estimate initial conditions, parameters and parameter ranges in the various sub-models. The driver component of the DVET model includes sensing, information processing and control action modules, and represents the stochastic effects of distraction and delayed driver reaction to a lane departure event. Batch simulations were run in cases where the LDW technology model is active or suppressed. Taken together with estimates of system availability and driver responsiveness, an estimate for the range of safety benefits was developed. The SIM provides detailed indications of how the DVET components are expected to interact in the field, and hence the results provide safety-related information that goes beyond the numerical benefit estimates, considered preliminary in this first analysis. This project was led by the Ford Motor Company.

17. Key Words
ACAT, Lane Departure Warning, active safety, safety benefits, crash avoidance, driver model, simulation

18. Distribution Statement
Document is available from the National Technical Information Service www.ntis.gov

19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 218	22. Price

Form DOT F 1700.7 (8-72) Reproduction of completed page authorized

Acknowledgements

The authors are grateful to the following individuals who made substantial contributions to the technical direction of the project and provided numerous suggestions for improving the format and content of this report.

- M. Blommer, M. Shulman, K. Sullivan, L. Tijerina and R. Wheelock (Ford)
- N. Harish and D. LeBlanc (UMTRI)
- R. Biveby, A. Eugensson, L. Jakobsson and M. Lindman (Volvo).

The authors also recognize and appreciate the efforts of NHTSA staff for their suggestions and comments throughout the project.

Contents

List of Figures

List of Tables

Executive Summary

NHTSA's Advanced Crash Avoidance Technologies (ACAT) program has two broad objectives. The first is to develop a formalized Safety Impact Methodology (SIM), to estimate the ability of advanced technology applications in full vehicle systems to address specific types of motor vehicle crashes. The second objective is to demonstrate how the results of objective tests can be used by the SIM to forecast the safety benefit of a real technology. Both objectives have been achieved in this research and are documented in this report.

The report describes the methods and operation of a SIM developed by researchers from Volvo, Ford and UMTRI (VFU). The SIM applies computer based modeling and simulation along with statistical analysis, and fuses a wide variety of data – from crash, naturalistic driving, driving simulator, test-track and highway sources – to estimate which safety benefits may be generated by introducing a safety technology into the interactions between driver, vehicle and driving environment.

This study focuses on crashes occurring after a subject vehicle exits the travel lane. Rather than trying to reconstruct a set of individual crashes and assess a technology's potential influence on that set, the crash reconstruction in this project has taken place at what can be called a population level. First, a set of crash types associated with inadvertent drift out of lane events has been defined. Next, a set of "Driving Scenarios" (DS's) relevant to these crash types were constructed, using both crash data and other non-crash data sources. Each DS represents the typical range of **D**river, **V**ehicle and **E**nvironment (DVE) conditions preceding one of the identified crash types.

The DS's were defined and represented in a parametric form, so they can be used as input to a **D**river-**V**ehicle-**E**nvironment-**T**echnology (DVET) simulation model. The model has four components; a driver model, a vehicle model, a model of the traffic environment and an ACAT safety technology model. The model components were first calibrated and validated against relevant data, such as test track and naturalistic driving data. Next, the full DVET simulation model was used to run multiple time-stepping simulations of each DS, i.e. the virtual driver would drive the virtual vehicle in the virtual environment, subject to the DVE pre-conditions defined for each DS.

The DVET model time-steps from the starting point of each DS until the DS has run for a certain time interval. This is typically 10 - 20 seconds, the exact value depending on how the simulation evolves. Furthermore, for each new simulation within a DS, the DVE pre-conditions are varied slightly (within the ranges defined for that DS). Each simulation is also run twice, once with and once without the technology. The set of simulations run for each DS thus can be said to explore the outcome of a large number of variations within the general DS pre-crash conditions, both with and without the technology present.

Due to the way that DVET-model and DS are defined, a crash is a possible but not a necessary outcome for each simulation. Whether a crash occurs or not depends on whether the virtual driver is capable of keeping the vehicle in the lane, given that simulation's particular DVE pre-conditions, and whether the virtual driver is supported by the technology or not. The resulting outcome of the simulations can therefore be called a *virtual crash population,* i.e. a large number of pre-crash condition simulations, where some lead to crashes and some do not.

Once all simulation runs are finished, the virtual crash population is split into two groups, one with outcomes where the technology is present and one with outcomes where the technology is absent. These two groups are then used in the system performance analysis. By reference to differences in crash frequency between the two groups, an estimation of technology benefits could be generated. The relative importance of each DS, and hence each underlying crash type, is calibrated by assigning weights to each DS, reflecting their underlying crash type proportions in the real world data.

In summary, at its core, the SIM involves the following key processes:

(i) A sampling scheme for basic driving scenario definition based on crash data, and including highway geometry and environmental conditions;
(ii) A "DVET" simulation model incorporating **D**river, **V**ehicle, **E**nvironment and **T**echnology (safety technology), which is calibrated by track tests and simulator tests and then used to generate an ensemble of virtual conflicts and crashes;
(iii) A sampling scheme based on naturalistic driving data to initialize detailed simulations;
(iv) A mapping from the ensemble of simulated conflicts to real-world crash types and frequencies, including optimized scenario weights, to provide an underlying "virtual crash population"; and
(v) A unifying analytical method to find out which safety benefits can be estimated from any changes that occur in the virtual crash population when the technology is present and supports the driver.

The reason for taking a simulation approach is that for the crash types considered, no simple program of experimental testing can directly predict the performance of relevant safety technologies in the field. There are so many combinations of contributing factors and coincident actions involved, that an unmanageably large number of pre-crash conditions would have to be considered for evaluation. Monte Carlo simulation has therefore formed the core of the predictive component of this research, based on an underlying computational DVET model with inputs from both real world driving and crash data.

One of the main achievements of the study was to fully integrate diverse data sources in a coherent and objective manner, at a sufficient level of detail to run the simulations. Crash data do not contain sufficient detail about early pre-crash conditions, especially relating to driver state and detailed vehicle kinematics in the period immediately preceding the initiation of the crash sequence. To avoid placing excessive emphasis on formalized crash reconstruction, the methodology has been developed so that naturalistic driving data and crash data supply "two legs of the stool", while the third leg is provided by objective testing, including both technical testing of the system performance and human factors testing of driver-system interaction. The three data source legs are then united via the simulations, run over what can be viewed as a very large set of test conditions.

It is worth noting that since the DVET simulations are the core activity of this project, the role of objective testing has been driven towards calibration and validation of DVET simulation model performance, rather than to be used as a direct basis for the benefit assessment.

This approach increases the complexity of the program beyond a simple "perform test and analyze data" format, but the DVET simulation model provides a feasible way to fuse diverse data sources into a single predictive tool. Indeed, it seems inevitable to take such an approach, as no single data resource appears rich enough on its own to provide the desired predictive power.

The SIM developed in this research has been focused on a specific set of technologies. However the methodology adopted is very general, and is especially appropriate for technologies that can involve significant interactions with the driver; this is typical for driver assistance technologies that begin to operate early in the crash phase sequence[1] (Figure I). Alternative approaches may be appropriate in other cases, for example if the technology intervenes only in the crash-imminent phase. The focus of this project on methodology development has led to a number of simplifications in the research; for example a single representative light vehicle (mid-sized sedan) was used in DVET model. Also, there was no explicit modeling of the effects of driver age and skill, even though broad random variations in driver performance, especially reaction times, have been included in the analysis. The methodology can accommodate extensions to multiple vehicle and driver types etc. in the future.

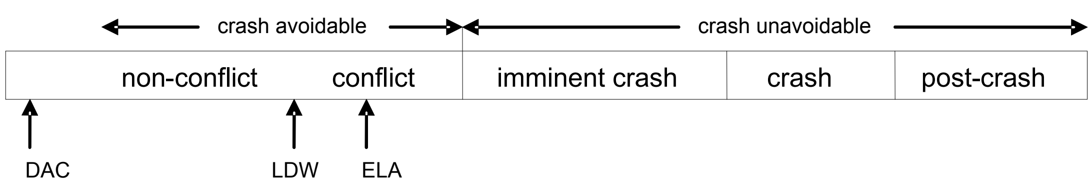

Figure I: NHTSA Crash Phase Timing Diagram

The SIM was developed and applied in the context of a Volvo Lane Departure Warning (LDW) system, aimed at warning the drivers if they are drifting out of the current travel lane. Under such a scenario, LDW supports the driver by generating an audible alert. The system will not take any automatic action to prevent a possible lane departure, and responsibility for the safe operation of the vehicle remains with the driver.

[1] The crash phase sequence (or crash phase timing) diagram was previously formulated by NHTSA in the request for applications (NHTSA, 2006).

Two related technologies have also been included in the study: "Driver Alert Control" (DAC) and "Emergency Lane Assist" (ELA), though full and formal benefits analysis was limited to LDW. DAC is designed to detect degraded control in the driver's lane keeping, to estimate the level of degradation, and to warn the driver when the inferred reduction in alertness or vigilance reaches a certain threshold. While DAC operates early in the normal (non-conflict) phase, ELA operates later, in the full conflict phase as indicated in Figure I. ELA relies on the detection of the host vehicle position with respect to the road lane markings as well as detection of other vehicles (both oncoming and those being overtaken) in the adjacent lanes. If the host vehicle is drifting into an adjacent lane, and the ELA detects a collision course with another vehicle in that lane, the system provides an active steering intervention to return the subject vehicle back into the original travel lane, provided that return to lane has a clear path. Taken together, the three active safety technologies address aspects of potential crashes arising from inadvertent drift out of lane. As mentioned, to contain the scope of the safety benefits estimation and present the methodology as clearly as possible, the performance analyses of these latter two technologies were not formally incorporated into the large-scale benefits estimation process. However, it was thought worthwhile to investigate the feasibility of evaluating such a combined system within the SIM methodology, and this is included in the report.

The VFU partners all made key technical contributions to the research. The technology definition and its implementation in test vehicles and driving simulator were the responsibility of Ford and Volvo. Defining the associated target crash modes and linking them to crash data was jointly conducted by all partners, with UMTRI leading the resulting analysis of crash and naturalistic driving data. Physical tests were conducted on the track at Volvo in Sweden and in the driving simulator (VIRTTEX) at Ford in Michigan. The main focus of the track testing was to validate the physical performance envelope of the safety systems. In the driving simulator, the emphasis was on human factors tests with naïve subjects (e.g. distracted and sleep deprived driver tests with lane departure warning) though again some controlled technical tests were also conducted.

The simulation tools were developed at UMTRI with significant input from Volvo and Ford on the scope and functionalities of the models. Commercial software was used, in the form of CarSim for vehicle simulation, SIMULINK for custom modeling of the driver and the highway conditions, and MATLAB for the data management and job control for batch simulation. MATLAB was also used to compute estimated safety benefits, and this analysis was also led by UMTRI researchers.

Data used to develop the target crash types and to describe crash conditions include the National Automotive Sampling System General Estimates System (NASS GES) and NASS Crashworthiness Data System (NASS CDS), crash data from the State of Michigan that was geo-located, and roadway geometric information from the Highway Performance Monitoring System (HPMS). Crash data obtained from the NASS provided a key resource for the whole project. The GES data were used to estimate crash numbers under different conditions, and CDS was used for confirmation that the codes used for extracting the crashes types from GES broadly captured the crash conditions of interest. Naturalistic data from the RDCW (Road

Departure Crash Warning) Field Operational Test were found to be sufficiently comprehensive for the purposes of populating elements of the scenarios that could not be obtained from the crash data or from test track data.

The input data for the batch simulations were over-sampled from higher risk conditions in the DS's to ensure that the simulation results covered the crashes of interest in an efficient manner. The outcomes from the batch simulations were later adjusted by the relative frequencies of the crashes from GES to ensure that the safety benefits estimates were aligned with real-world data. In the batch simulations, the road types were broadly classified into urban and rural and within each the roads were further classified as 4-lane divided, 4-lane undivided and 2-lane undivided. The crash outcomes were partitioned into four categories based on the type of crash: paved shoulder, adjacent lane – same direction, adjacent lane – opposite direction, and off-highway.

In order to estimate the safety benefits for the specific crash problem, approximately 15000 simulations were run to generate the underlying virtual crash population. By optimizing driving scenario weights it was possible to produce a reasonable degree of fit to the actual (GES coded) crash population, albeit with some error in the locations of first harmful event predicted for off-highway crashes in multi-lane divided highways. The lack of perfect fit is hardly surprising in a first study of this type; however the chosen measure of fit is high – 85% of variations are explained according to the analysis. A further simple test of validity was made by randomly selecting a small number of relevant real-world crash events from NASS CDS and confirming that similar events can be found in the virtual crash data. Therefore, the use of a virtual crash population is at least plausible, and there remains clear scope for improvement in the future, especially through the use of more detailed and geographically diverse highway data.

The results of the batch simulations show that in every driving scenario the effect of LDW is to reduce the estimated crash numbers. This is certainly expected since the system triggers an early reaction from the driver or, if the LDW alert occurs too late to generate a useful response, its effect is at worst neutral.

The most striking result is the relative uniformity across road types and outcome types. There are however some trends in the results. One is towards greater effectiveness of the LDW technology in rural conditions, partly due to the large number of single-vehicle crashes in rural conditions. The effects of a number of adverse factors (curved road, nighttime driving, fatigued driver, and wet roads) were also evaluated. The SIM analysis predicts system effectiveness is higher under the adverse conditions with the exception of *curved road*, where the LDW effectiveness is reduced. This appears to be due to the reduced time available for driver reaction under these conditions, though formal analysis of the effect was not included in the study. Analysis also predicts that while LDW is effective in reducing crashes across most of the speed bands, there is a trend towards reduced LDW effectiveness at the highest speeds.

The initial raw estimate of overall "system effectiveness", i.e. the proportion of crashes reduced in total is 47%, which corresponds to 85,000 crashes annually for the target vehicle type and crash type. However, this raw estimate does not include several important effects such as system availability and driver acceptance and compliance. When these factors are considered, more plausible estimates result.

As can be expected, the overall effectiveness is strongly affected by system availability. The LDW system relies on a camera to register lane markings and compute vehicle positions relative to the lane. The system is only available when this data capture is successful, and hence the probability of the system being available depends on the quality and continuity of the lane markings, any contamination of the road surface (e.g. standing water or repair markings), and also the lighting conditions. System availability was estimated using a combination of published data and test data, and was estimated to vary from a high of over 90% during daylight to a low of under 20% for wet roads and nighttime driving conditions. When system availability was factored into the estimation, the overall system effectiveness was reduced to 33% from the unadjusted value of 47%.

Another factor closely tied to availability is the lower speed threshold of 40mph for the LDW activation – below this speed the system was disabled, which is equivalent to the system being unavailable in this condition. This was found to have a very small effect on the overall system effectiveness – reducing it from 33% to 32%.

The initial raw estimate includes driver response delay times that were randomly sampled from ranges chosen on the basis of the VIRTTEX driving simulator studies. The analysis shows that benefits from LDW are affected by driver reaction time; if the mean delay in driver response is increased by 0.1s, the predicted overall system effectiveness is reduced by three percentage points.

There are several other factors expected to influence the overall safety benefits. Unlike system availability there is insufficient objective data to fully quantify the effects of these additional factors, though where possible the broad magnitude of their influence were estimated. These include:

- Under-reporting of driver fatigue in crash data. The analysis suggests that this could add around 1% to the overall system effectiveness.
- Driver acceptance of the technology and compliance with its alerts. Review of naturalistic data suggests that for approximately half the time, drivers were prepared to maintain the lateral position at which the alert was sounded (or drift even further from center) and didn't feel the need to correct. This has a strong potential influence on the benefits of LDW.
- Long-term driver adaptation to the safety system may change the frequency of driving scenarios that are associated with distraction. No relevant data were collected in this project, and in the absence of clear published data, the effect was assumed to be neutral.

Taking account of the above factors, the final estimate for the effectiveness range is a 13% - 31% crash reduction for the target crash modes; this would equate to between 24,000 and 57,000 fewer accidents annually for the light vehicle class analyzed.

While the above results have inherent value (and no doubt will be refined in the future), it is perhaps the methodology development that is of greatest value as an output of this project, having applicability that goes beyond the evaluation of one specific safety system. The VFU-ACAT research team is not aware of any other implementation of active safety benefits estimation that utilizes such a wide range of data sources and combines that information with detailed computer simulations of crash occurrence and avoidance. In particular it is worth noting the critical use of naturalistic data to refine detail in the driving scenario and provide the necessary kinematics to initiate scenario simulations. The data used for this are not specific to the ACAT system evaluated or even to the crash types considered; this provides for efficient re-use of the supporting data elements across multiple safety systems in the future.

The development and implementation process has highlighted a number of areas where either basic research knowledge or supporting data have significant gaps at present, and it is hoped that the research findings will stimulate fresh efforts to fill these gaps.

From a general perspective, all approaches to the evaluation of active safety technologies face the same challenge of representation – using some form of model (i.e. a test setup) to represent how a scenario unfolds with or without the technology of interest. Regardless of where and in which form an evaluation takes place, the characteristics of the four components needed to perform the evaluation, i.e. Driver, Vehicle, Environment and Technology (DVET), must be represented in a way which is sufficiently faithful to their counterparts in real world crashes. The major strength of the SIM developed here is that it makes full use of pre-existing crash and driving data, and integrates these through large-scale batch simulations, so that multiple factors and complex interactions are allowed to play a role similar to that seen in the real world. It also permits detailed predictions to be made about real-world effectiveness of LDW under different conditions, and it is possible that these predictions can be tested in the future, for example via the SHRP2 Naturalistic Driving Study.

1. Introduction and Overview

This report describes the methods and operation of a Safety Impact Methodology (SIM) developed by researchers from Volvo, Ford and UMTRI (VFU). The SIM applies modeling, simulation and statistical analysis, and fuses a wide variety of data – from crash, naturalistic driving, driving simulator, test track and highway sources. The aim of this report is to provide a description of the SIM methodology applied in a "start to finish" analysis of a specific set of active safety technologies considered by the VFU team.

Specifically, the SIM is developed and applied in the context of a Volvo lane departure warning (LDW) system. Two other technologies have been included in the study: "Driver Alert Control" and "Emergency Lane Assist"; together, the three active safety technologies all address aspects of potential crashes arising from inadvertent drift out of lane, and are introduced in more detail in Section 1.3. However, to contain the scope of the safety benefits estimation and present the methodology as clearly as possible, the performance analyses of these latter two technologies are not formally incorporated into the large-scale benefits estimation process.

In this introductory section we highlight the aims, objectives and broad limitations of the program of research, how the work was carried out, how the three active safety technologies are intended to function, and then outline the overall report structure.

1.1. *Study Aims and Objectives*

The underlying purpose of this research is to address gaps in current knowledge about the performance and likely effectiveness of new and emerging active safety technologies in reducing crash numbers. An analysis method is formulated to estimate safety benefits arising from interactions between driver, system, vehicle and driving environment. The analysis includes a formulation of the conditions that precede relevant crash types in the form of "Driving Scenarios", with simulations linking these relatively normal driving conditions to real-world crash outcomes.

As described in the original request for applications (NHTSA, 2006), the motivation behind the Advanced Crash Avoidance Technologies (ACAT) program has been twofold. The first is to develop a formalized Safety Impact Methodology (SIM) tool to estimate the ability of advanced technology applications in full vehicle systems to address specific motor vehicle crashes. The second objective of the program is to demonstrate how the results of objective tests can be used by the SIM to forecast the safety benefit of a real system. Both objectives have been achieved in this research and documented in this report. As part of the development, test and evaluation program, careful attention has been placed on the crash sequence, or crash phase timing, previously formulated by NHTSA in the request for applications (NHTSA, 2006). This is illustrated in Figure 1.1, together with the phases most relevant to the active technologies considered. Driver Alert Control is expected to influence the exposure of drivers to episodes of diminished vigilance, for example due to sleep deprivation, while driving and hence operates in the earliest phase. Lane Departure Warning and Emergency Lane Assist are relevant to the evolution of vehicle kinematics during a conflict, and are operating in the early and late conflict stages. The imminent crash phase is not expected to be relevant to these technologies, nor are the actual crash and post-crash phases.

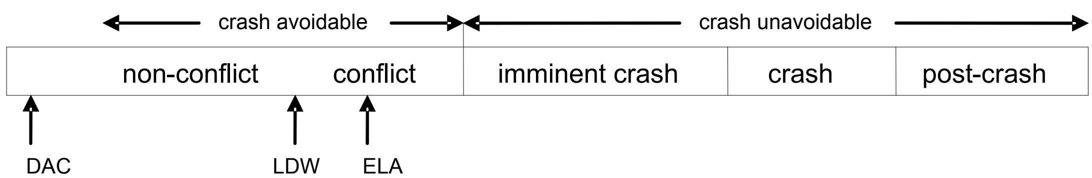

Figure 1.1. NHTSA Crash Phase Timing

The fundamental requirement for the present research has been to develop predictive tools. This is because, for new and emerging technologies, there is no readily available field data to analyze how an active safety technology performs in the real world. Field operational tests (FOT's) have been conducted previously to fill this gap, for example under related NHTSA programs (LeBlanc et al, 2006, 2007), but such programs directly evaluate the specific algorithms and HMI of one particular system design proposal at a time; with the wide range of active safety technologies currently being developed, it is desirable to have additional tools such as the SIM instead of having to perform such large-scale evaluations on each and every separate technology proposal.

One specific aim of the current project has been to avoid any structural dependence of the SIM on FOT data. In fact, the research team did mine pre-existing FOT data for the purpose of including naturalistic driving behavior in the analysis; other sources of high-fidelity naturalistic data might have been used if they had been available. The key point is that, where real-world driving data has been used in this project, the focus has been entirely on utilizing data to characterize the naturalistic driving aspects, not on system function or performance in the field.

In the research it has become clear no simple program of experimental testing can directly predict the performance of the safety technologies in the field. Rather, the role of objective testing has been driven towards the calibration and validation of computer models of the system and its interaction with both the external environment (highway, other vehicles) and the driver of the subject vehicle. There are so many combinations of factors and coincident actions in the types of crashes considered, that a very large number of candidate pre-crash conditions need to be considered. For this reason, Monte Carlo simulation has formed the core of the predictive component of this research, based on an underlying computational model and linked to both real world driving and crash data. This increases the complexity of the program beyond a simple "perform test and analyze data" format, but the simulation model provides a feasible way to fuse diverse data sources into a single predictive tool. Indeed, it seems inevitable to take such an approach, since no single data resource appears rich enough on its own to provide the desired predictive power.

One fundamental concern of the research program has been to represent all major relevant tasks and functions relevant to crashes caused by inadvertent lane departures. As will be described in Section 4, two major components are driver distraction and fatigue. For this reason, a driver model has been developed that includes variable attention switching, as well as information processing that generates variable time delays in the lane-keeping control process. This is a relatively simplified functional model that includes the critical sub-processes of visual capture of boundary points, recognition and classification of those points, memory management, threat assessment and finally the manual control of steering. The model is developed and calibrated to include the major mechanisms of control disruption and recovery during error-inclusive lane keeping; clearly such actions are crucial to representing the effects of the LDW and ELA systems.

One of the main challenges of the study was to make the best use of diverse data sources. As will be described, crash data does not contain sufficient detail about early pre-crash conditions – especially relating to driver state and detailed vehicle kinematics in the preceding driving scenario. Hence, to avoid placing excessive emphasis on formalized crash reconstruction, the methodology has been developed so that naturalistic driving data and crash data supply "two legs of the stool". The third leg is provided by objective testing, including both technical testing of the system performance and human factors testing of driver-system interaction. The three legs are united via simulations, run over a large set of test conditions. The overall analysis concept is developed in Section 2 of this report.

It has not been an explicit goal of this research to provide a universal analysis tool for estimating safety benefits for all vehicle safety systems, or even all active safety systems. The goals of the study are already ambitious, so the team's approach has been to focus on specific technologies and develop an approach that seems most appropriate for those cases. For a different technology an alternative approach might be preferred, especially if that technology operates in a different phase of the crash timing (Figure 1.1). The current systems operate early in the crash phase and involve a high degree of interaction of the system with the driver; the outcome of the event depends critically on how and when the driver reacts. A safety system operating in a later crash phase and offering little driver interaction may be more appropriately assessed by detailed crash reconstruction. The research focus on methodology development has also led to imposing constraints on the number of test cases considered, and has therefore not explicitly included effects of vehicle type, driver age and skill (though broad variations in driver performance, especially reaction times, have been included). For the vehicle model we have adopted a single target vehicle type (a mid-sized sedan class), assuming this to be sufficiently representative of the light passenger vehicles explicitly addressed in the crash data (see Section 4).

While the presented SIM analysis is not claimed to be universal and applicable to all technologies, neither is the analytical method limited to the specific technologies used. At its core, the SIM involves the following key processes

(i) A sampling scheme for basic driving scenario definition based on crash data, and including highway geometry and environmental conditions;
(ii) A "DVET" simulation model incorporating **D**river, **V**ehicle, **E**nvironment and **T**echnology (safety technology), which is calibrated by track tests and simulator tests and then used to generate an ensemble of virtual conflicts and crashes;
(iii) A sampling scheme based on naturalistic driving data to initialize detailed simulations;
(iv) A mapping from the ensemble of simulated conflicts to real-world crash types and frequencies, including optimized scenario weights, to provide an underlying "virtual crash population"; and
(v) A unifying analytical method to find out which safety benefits can be estimated from any changes that occur in the virtual crash population when the technology is present and supports the driver.

Given the very wide scope of the research, the focus here is on methodology development and demonstration more than predicting safety benefits with any high degree of accuracy. Where available data is sparse or imprecise, alternative sources have been determined, or – where this has proved impractical – simplifying assumptions have been made and recommendations put forward for how gaps in knowledge or data might be addressed in future research. These are detailed in Appendix B.

1.2. *Conduct of the Research Program*

The research program has broadly followed the steps shown in Figure 1.2. The technology definition and its implementation in test vehicles and driving simulator were the responsibility of Ford and Volvo. Defining the associated target crash modes and linking them to crash data was jointly conducted by all partners, but with UMTRI leading the resulting analysis of crash and naturalistic driving data.

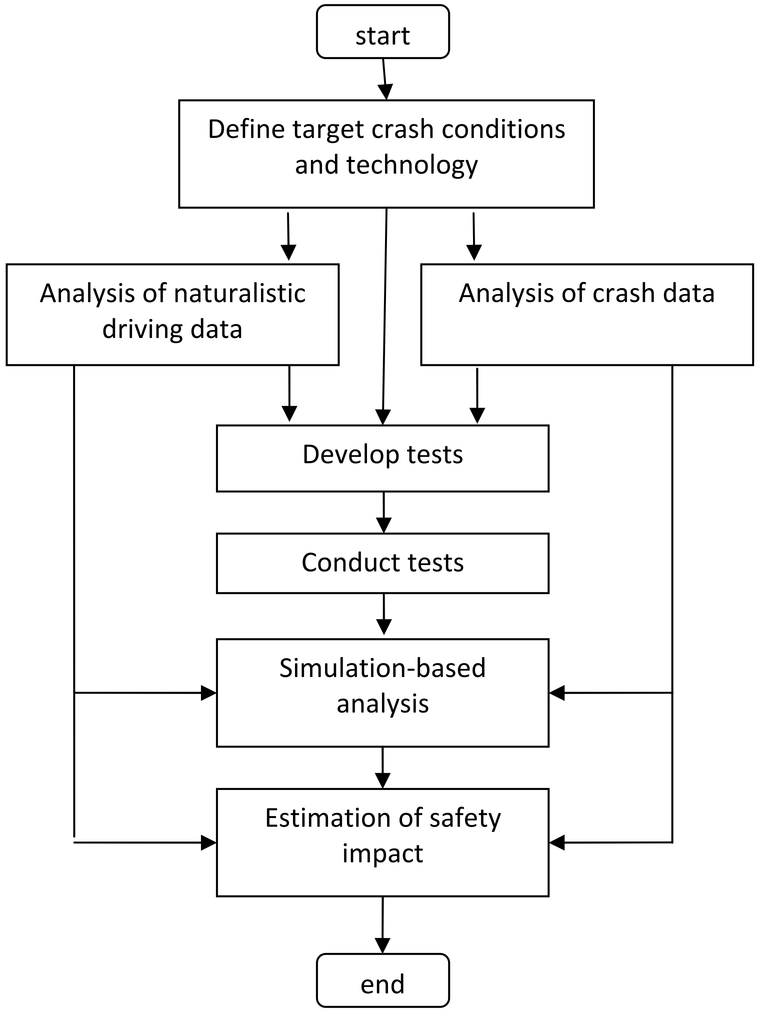

Figure 1.2. General Overview of the Research Program Steps

Physical tests were conducted on the track at Volvo in Sweden and in the driving simulator (VIRTTEX) at Ford in Michigan. The main focus of the track testing was to validate the physical performance envelope of the safety systems. Some track testing with naïve subjects took place at Volvo and at Ford as part of the evaluation of Driver Alert Control (see below). In the driving simulator, the emphasis was on human factors tests with naïve subjects (e.g. distracted and sleep deprived driver tests with lane departure warning) though again some controlled technical tests were included. Further details of the objective testing are presented in Section 5.

The simulation tools were developed at UMTRI, though with significant input from Volvo and Ford on the scope and functionalities of the models. Commercial software was used, in the form of CarSim (Mechanical Simulation Corporation, 2008) for vehicle simulation, SIMULINK (Mathworks, 2009) for custom modeling of the driver and the highway conditions and MATLAB (Mathworks, 2009) for the data management and job control for batch simulation. MATLAB was also used to compute estimated safety benefits, and this analysis was also led by the UMTRI researchers.

The research program broadly followed the flow shown in Figure 1.2, with the formal task structure of Table 1.1. Of course there was some degree of iteration between tasks, as information and tools were developed and assumptions revisited. Throughout the project there was close communication between team members, with regular conference calls and web-based meetings, as well as very frequent email exchanges and occasional live meetings – including informal technical workshops in both Michigan and Sweden.

Table 1.1. Work tasks for the ACAT project

Task No.	Task Name
0	Program management
1	Safety Impact Methodology (SIM) Development
2	Safety Area Characterization and Technology Specification
3	Objective Test Development
4	Objective Test Performance
5	Safety Benefits Development using SIM

Additional meetings and workshops took place in Washington to brief and exchange information with NTHSA staff. During these workshop meetings, most team members were able to drive test vehicles or the driving simulator to experience the systems in an interactive environment. This was found to be extremely valuable, improving the level of shared understanding and communication within the team. Some NHTSA staff were also able to experience some of the system functions, and again this was found to be highly beneficial to the project.

1.3. *Active Safety Technologies*

The overall Volvo system consists of three active safety subsystems, each of which is designed to help the driver avoid crashes in certain specific circumstances: Driver Alert Control (DAC), Lane Departure Warning, (LDW), and Emergency Lane Assist (ELA). Each of these systems is now briefly described, while further details are provided in the context of system simulation.

Driver Alert Control (DAC). Driver Alert Control is designed to detect degraded control in the driver's lane keeping, estimate the level of degradation, and inform the driver of his/her reduced alertness, where reduced alertness is predicted through quality of lane keeping over time. The driver is informed of his/her state so as to support a decision to continue to drive or take remedial steps like a rest break. This system uses an algorithm described in the literature (Birk et al., 2006) to grade the driver's lane keeping variability. When a threshold of lane keeping degradation is exceeded, the driver receives a warning through a Human-Machine Interface (HMI). DAC relies on detection of the vehicle position with respect to the road lane markings. The vehicle position is evaluated making use of a camera system that measures the lateral distance from the camera center line to the left and the right lane markings. DAC is operational at speeds between 65 kph and 180 kph, and if at least one lane marking on either side of the center of the vehicle is present.

DAC was developed with the particular intent of addressing diminished driver vigilance, for example due to sleep deprivation. The aim is to detect and predict diminished vigilance for at least 95% of the driving population at least one minute before normal driving control is lost, based on the assumption that diminished vigilance is a state that appears slowly and remains fairly constant for some duration of time. In order to verify that the degradations in quality of lane keeping identified by the algorithm actually correlate with diminished vigilance in driving, two test track studies were carried out, one at the Hällered proving ground in Sweden and one at the Michigan proving ground in the U.S. A total of seventy-seven (77) sleep-deprived test subjects were asked to drive for two hours. The drive was stopped when drivers were judged to have fallen asleep (based on the vehicle being 50% out of the lane). During the testing, the DAC HMI was suppressed (meaning that no alert was given to the driver even if DAC triggered), in order for any subsequent lane departures to happen in a natural way. All 77 test participants fell asleep during testing, and 76 of 77 drivers did trigger a DAC warning at least 1 minute before the lane departure (though with the alert not issued to driver). These tests indicate that degraded performance detected by the algorithm correlates well with diminished driver vigilance due to sleep deprivation.

Lane Departure Warning (LDW). Lane Departure Warning is aimed at warning the driver if he or she is drifting out of the current travel lane. Under such a scenario, LDW supports the driver by generating a warning. The system will not take any automatic action to prevent a possible lane departure. Responsibility for the safe operation of the vehicle remains with the driver. Like DAC, LDW relies on the detection of the vehicle position with respect to the road lane markings (via a camera system) in order to detect a lane departure. Thus visible lane markings must be present on both left and right hand sides. The system is intended to operate during daytime and nighttime. At night, a minimum of low beam illumination from the host vehicle's headlamps is required, but no other illumination is needed. The speed range of operation is similar to DAC.

Emergency Lane Assist (ELA). ELA relies on the detection of the host vehicle position with respect to the road lane markings as well as detection of other vehicles (both oncoming and those being overtaken) in the adjacent lanes. The host vehicle position is evaluated making use of the same camera system as that used by DAC and LDW systems. Oncoming vehicles in the adjacent lanes are monitored by millimeter-wave radar. If the host vehicle is drifting into the adjacent lane, and the ELA estimates that it is on a collision course with an obstacle in that lane, the system provides an active steering intervention to steer the host vehicle back into the original travel lane under requisite conditions, including the condition that a clear path exists in the original travel lane.

Combined System Operation. Combined operation of DAC, LDW, and ELA subsystems offers potential benefits in both non-conflict and conflict conditions. LDW and DAC continuously operate in parallel and utilize much of the same hardware on the vehicle. DAC assesses the quality of lane control over time and detects diminished vigilance, for example due to sleep deprivation, or other degradation in control based on lane keeping. LDW provides a timely warning in certain instances of lane line crossing. When a lane departure has begun and there is a forward threat in the new lane, ELA is intended to help return the vehicle to the original travel lane if the collision threat is clear and unambiguous. Thus, a sequence of alerts and interventions provide a progression of filters to reduce risk (DAC), and assist the driver through warning (LDW) and active steering (ELA). Figure 1.3 provides an overview of Volvo DAC-LDW-ELA systems combined operation.

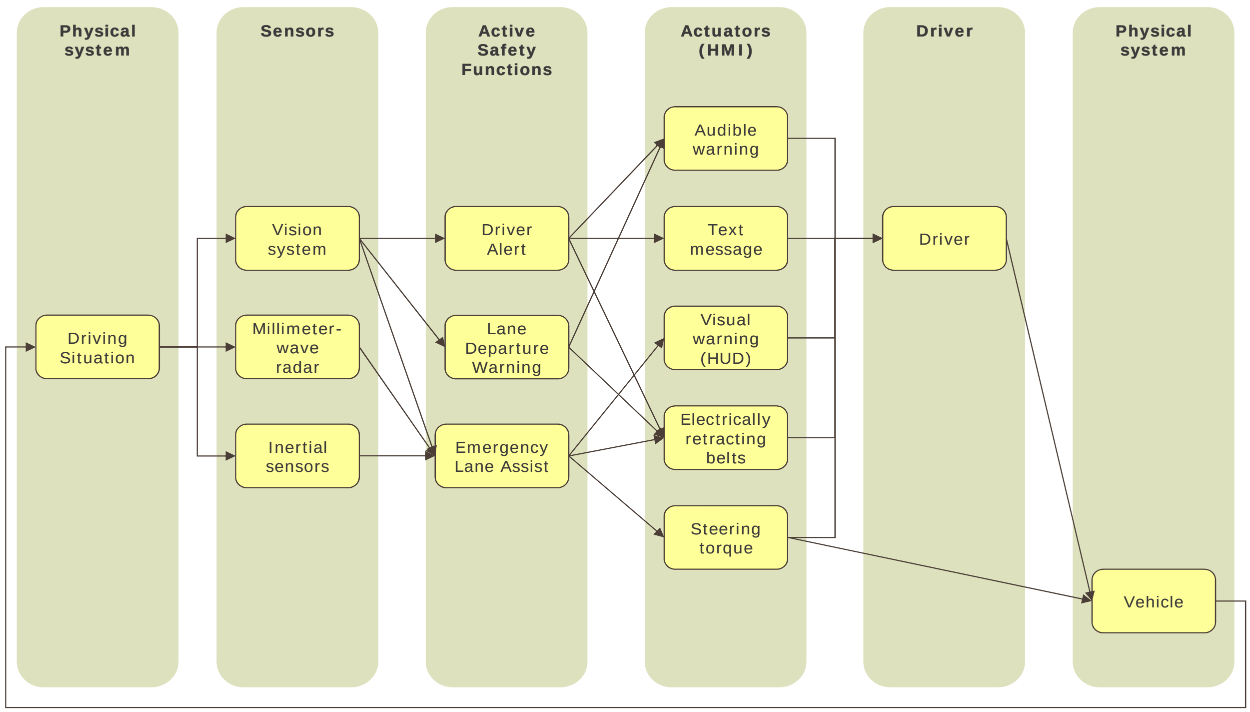

Figure 1.3. System Integration

In this research the full SIM analysis was limited to standalone LDW performance; extending the work to include the other systems, whether alone or in combination, is discussed in Section 10, with some partial analysis presented in Appendix D.

1.4. *Report Outline*

Section 2 provides a definition and overview of the methodology adopted for safety benefit estimation in this project, and has a strong focus on the techniques used to combine the complementary data sources to create the underlying parameters for the SIM. The overall SIM analysis is structured in the form of a flow chart and an overview of the various stages is presented. Section 2 also motivates and describes the analytical method used for safety benefits estimation.

The flow chart of Section 2 also provides the framework for the main section topics of this report. Section 3 covers the main data sources, and how information is extracted in the analysis via a number of key parameters. Section 4 considers the relevant crash cases from multiple perspectives, most particularly from the technology function and the crash data analysis. Objective testing is covered in Section 5, while the model development – the model formulation, test and refinement – is covered in Section 6. This leads to the parametric formulation of the driving scenarios and the preparation required to establish large-scale batch simulations in Section 7. The performance analysis within individual scenarios is presented in Section 8, and this is integrated into an overall safety benefits estimation process in Section 9, where safety benefits estimates are developed. Finally Section 10 provides the summary and conclusions of this study.

2. Safety Impact Methodology

This section provides a road map for the SIM developed in this project. The purpose of this section is mainly to provide an overview of the activities that take place within the SIM development process. However, some description will be included concerning why alternative approaches were not followed.

The overall structure of the project is quite straightforward. As the purpose of the project is to assess the possible influence of certain safety technologies on an existing crash problem, the first step is to describe how one intends to carry out the assessment, i.e. formulate the research plan. The second step is to generate a detailed description of the specific crashes to be addressed by the technology. This step addresses questions such as: exactly which crashes are of interest in relation to the technologies under evaluation, and how much can be learned about what these crashes look like and why they happen, given existing data sources?

Once the problem definition is clear, it is time to set up the evaluation procedure. In this project, evaluation has taken place through computer based simulation rather than through track or simulator testing, but the basic concept is the same. One takes a group of drivers and vehicles believed to be representative of the drivers and vehicles that experience the crashes of interest in the real world, and make them drive under certain pre-conditions in a chosen test environment. The intent is to replicate interesting pre-crash events in a controlled and safe way, and see whether the technology under evaluation helps drivers respond faster and/or better in the situations created.

The strong point of using simulation for technology evaluation is that rather than having to select one or two pre-crash condition sets (which usually is necessary in a physical or driver simulator setting, otherwise evaluation costs too much and/or takes too long), one can run thousands of simulations. It is therefore possible to explore a very large number of variations in the pre-conditions. This means that questions such as “what if driver response is a little bit slower?” or “what if vehicle speed was a little bit higher?” go from being discussion points to questions which can be answered quantitatively. An important additional strength is that computer simulation completely avoids one problem of physical testing with naïve subjects: test-to-test variations for individual subjects can be very large, especially as adaptation to the test environment and anticipation of events takes place.

The potential weak point of using simulation for technology evaluation is that, since all its components are computer models rather than physical objects, a lot of effort has to be put into making sure that these models behave as their real life counterparts would. From a research point of view, a test setup at a track or in a driving simulator is also a model of the real world. However, verifying, for example, that a computer based driver model steers a vehicle in the same way as a real world driver does takes more effort than verifying the steering behavior of a test track driver. A large portion of this report is therefore dedicated to the issue of verifying the behavior of the computer models used in the simulation.

Once the evaluation environment has been set up, two groups of simulations are run, i.e. one with and one without the technology under evaluation present. By comparing the outcomes of these two groups, a basis for the safety benefit analysis is established. However, just as what is measured at a test track rarely translates directly into a crash count reduction, the simulation outcomes cannot be used directly to predict what difference the technology would make to the real world crash population. Rather, a process has to be developed and applied for transforming the difference in simulation outcomes with and without the technology present into a real world safety benefit estimate.

Figure 2.1 presents a high level picture of the SIM framework adopted during the project. The figure was developed based on a more generic SIM flow chart considered by NHTSA (Carter et al, 2009). As seen, the flow chart relates data sources, case scenarios, objective testing methods and simulation model development leading to a somewhat idealized definition of the candidate scenarios used for analysis. In the figure "scenario" is used broadly and covers both driving in the non-conflict phase as well as the conflict type, crash conditions and crash outcomes. As the methodology is defined however, the key scenario definition is that of the non-conflict *driving scenario*, and the reason for this will be made clear in the following discussion. Once scenarios are defined and represented in a parametric form, the simulation model is completed, and batch data generation is initiated. The batch simulations involve time-stepping simulation across multiple scenarios to map the predicted outcomes with or without the ACAT system active. The results of this are used in the system performance analysis, and by reference to crash frequency data for the scenarios of interest, an estimation of benefits is generated. In total we use real-world data to drive a large-scale computer simulation to create a calibrated "virtual crash population"; re-running the same simulations with the ACAT system active provides a basis for estimating the potential safety impact of the technology.

The processes shown in Figure 2.1 are diverse, as are the data sources and even the domains of expertise required to develop them. It has not been possible to cascade all aspects of the SIM methodology down to the finest level of detail that will be desired in the future to increase the quality of the benefits estimates. However, due to the various timing, data and resource constraints of this project a range of simplifying assumptions have been adopted to allow completion of the various components of the SIM. The major simplifying assumptions are listed in Appendix B. On the other hand it has been possible to create a unified framework for analysis with connections between sub-processes objectively defined. Also, the SIM methodology should be recognized as a framework, not a deterministic algorithm; as we will see, a number of decisions are required in formulating the model and the driving scenarios, and linking these with candidate crash codes, etc. In the next section we describe the various components of Figure 2.1 in greater detail.

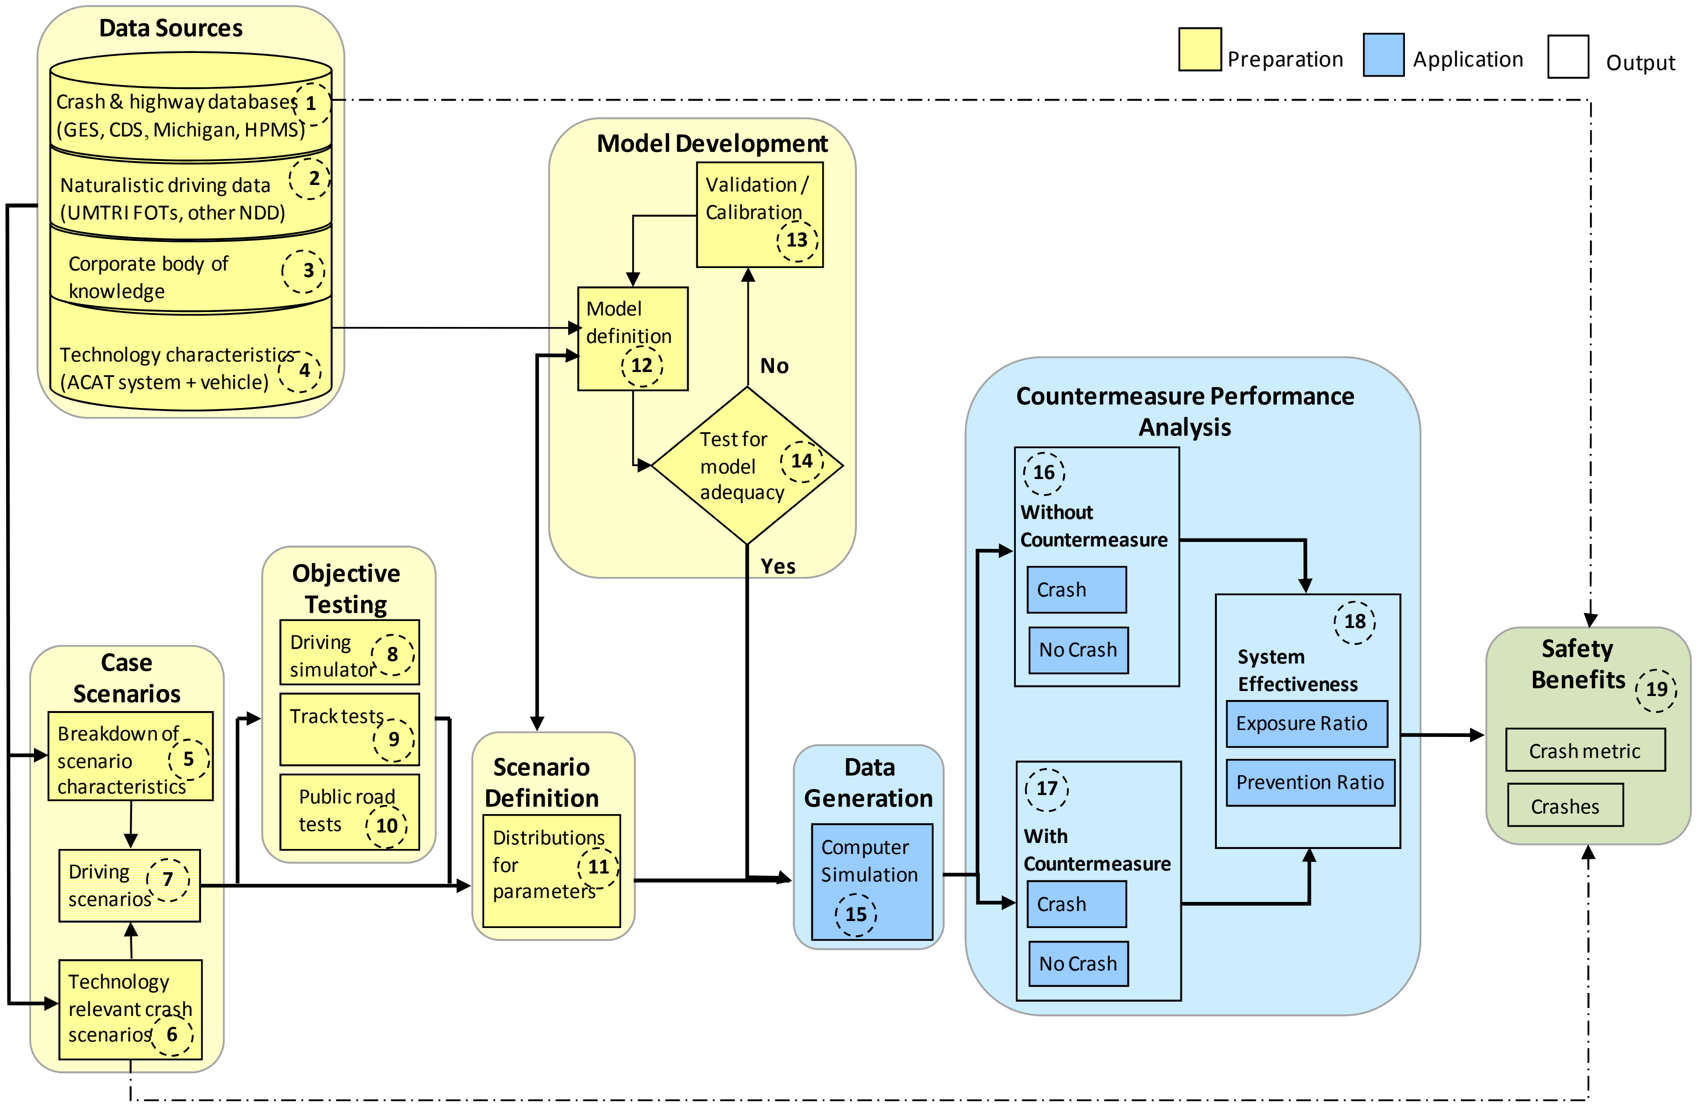

Figure 2.1. Overview of the SIM Components

2.1. *Components of the Safety Impact Methodology*

Analysis starts with a review of relevant data sources, as depicted in blocks 1-4 of Figure 2.1 for the SIM. Crash data obtained from the National Automotive Sampling System (NASS) provide a key resource for the whole project. Scenarios are initially formulated relative to the design intent of the safety system. An informal process of iteration takes place in the research team to establish common ground between this design focus and the relevant coding conditions for particular crash types in the US national crash databases, i.e., the General Estimates System (GES) and the Crashworthiness Data System (CDS) components of NASS. The GES data will ultimately be used to estimate crash numbers under different conditions, and CDS – which contains more detail, including photographs, summary descriptions and scene diagrams – is used for confirmation that the codes are broadly capturing the crash conditions of interest. It is not necessary or even possible to establish a unique match; for example, if the speed range in the crash data is greater than the operating range of the ACAT system, no error is introduced provided that the actual speed range of the crashes is factored into the system availability estimate used in benefits estimation. As a rule however, the closer the match between crash types and technology function, the more direct the benefits estimation will be[2].

The result of analyzing crashes and conflicts relative to the technology functional description is to define the key characteristics of the crashes, identify the supporting GES codes, and the corresponding selection criteria and limits (e.g. speed range) for the candidate driving scenarios – blocks 5-7 in Figure 2.1.

Note that the defined driving scenarios need not lead to any conflict or crash outcomes. In the present study, a distracted driver who wanders in the lane may not exit the lane, in which case the conflict and crash types never emerge. However the driving scenario (distracted or fatigued driver engaged in lane keeping under different road and environmental conditions) is consistent with the recorded crashes and matched to the design intent of the system. All aspects of modeling and testing also need to be consistent with the driving scenarios. In this way, the definition of the driving scenario, as represented in block 11 in Figure 2.1, becomes the cornerstone of the overall SIM analysis; however, to reach this point it is necessary to create and refine the simulation model (model development, blocks 12-14) as well as fill in critical objective data for the system and the driver (objective testing, blocks 8-10).

The model includes components of the driver, vehicle, environment and the ACAT system. It is developed in terms of the key processes identified in the case scenario analysis, so that blocks 12-14 are developed based on the mechanisms identified. There is a minimal level of model complexity implied by these mechanisms, but once that level has been met, the goal has been

[2] We note that if a scenario is completely within the design intent of the system, we might expect the system to address it completely. However, as with the GES coding, on-board sensing is similarly limited in its ability to diagnose the scenario, and the need to reduce the possibility of false warnings will naturally reduce the real-world performance of the countermeasure. In the current analysis, the efficacy of the countermeasure is also limited by the driver's ability to respond quickly and intervene successfully.

to keep the model as simple as possible. Otherwise the task of parameter estimation and validation can become overwhelming, not to mention the difficulty in maintaining software and keeping track of changes to algorithms and system parameters.

To perform the calibration and validation of the model, it is necessary to conduct objective tests, thus creating the necessary reference data (blocks 8-10). In this study, driving simulator tests were used to characterize driver performance, while track tests provided objective data to characterize the performance envelope of the various ACAT systems. It is important to note that the purpose of the objective tests is to characterize the performance envelope of the technologies and to characterize the driver interaction with the safety systems for use in the model, not to directly test the system in "representative conflicts". The test conditions, in terms of vehicle kinematics (e.g. speed) were chosen to be broadly representative of the crash conditions and driving scenarios of interest, but "performance in objective tests" is not intended to map directly to the estimation of safety benefits. As we shall see later, the estimation of safety benefits is based on a large population of simulations designed to explore a wide range of relevant driving conditions and crash outcomes; this is difficult to accomplish with any practicable physical testing program. One can think of the batch simulations in block 15 as a method of amplifying the available data to provide the necessary input for countermeasure performance analysis and overall safety benefits estimation. This "numbers issue" is critical, and is in addition to the more obvious challenge of objectively testing high risk driving conflicts via physical testing; in the current methodology this challenge is largely removed.

Under objective testing, block 10 (public road testing) plays a peripheral but important role in the analysis. The ACAT system availability depends on environmental and highway variables, especially in terms of the visibility of the lane markers to the on-board camera and image processing algorithms. The system availability – percentage of time it can provide a warning or other intervention – plays an important role in the safety benefits estimation; for example if the system is available only 50% of the time, it is clear that the benefits might be roughly halved, the exact figure depending on the co-dependency of system availability and crash risk on highway type and environmental conditions.

Once the model has been developed and calibrated, it is formally connected to the data sources via the parameters needed to populate its subsystems (e.g. vehicle parameters, friction constants, delay times for driver control application) and for the specific conditions needed to initiate a simulation (e.g. lane width, vehicle speed, position in lane). This information is assembled in block 11 of Figure 2.1 (the "analysis cornerstone"). Roughly speaking, the parametric structure is defined for those parameters coded or estimated in crash records ("*beta*" parameters) as well as those that are estimated from naturalistic driving data ("*alpha*" parameters). The beta parameters define the road and environmental conditions that can be directly estimated or determined from the crash data, as well as finer detail conditions of road geometry that may not be directly available in GES, but may be inferred (in this study geo-located crash data in Michigan was used to define actual lateral curvature of sample crash events, as well as lane width and shoulder width). More strictly, the beta parameters are found or estimated from crash data, and hence there is no requirement to estimate frequencies or relative frequencies in normal driving data – that is the key distinction. By contrast, alpha parameters are related to frequencies in naturalistic driving, conditional on the given scenario (beta) conditions. Thus naturalistic driving data is mined for instances of lane keeping in non-dense traffic on a multi-lane rural highway, to establish distributions of speed and other lane-keeping kinematic values. As the research methodology is developed in this report, some simplifying assumptions are introduced so that full multivariate distributions of such variables do not need to be modeled (see especially Section 7).

Sometimes there is a non-trivial choice to be made, whether to treat a variable as alpha-type or beta-type. The approach adopted in this study is to *use beta-types wherever possible*. This is because the use of alpha parameters always puts more emphasis on the predictive power of the simulation model, and that model cannot be expected to reproduce all the complex interactions and dependencies present in the real world of driving. For example, we could ignore the information in the crash data relating to wet or dry roads. We may estimate from naturalistic driving data the relative frequencies of driving in dry and wet conditions and the simulation model can be used to represent the change in surface friction for wet roads. But it is clear that the simulation cannot take account of all the factors that influence the relative frequencies of wet and dry lane departure crashes, including the variations in rainfall patterns in different parts of the country and at different times of the year. It is far more reasonable to treat the surface condition as a beta parameter, especially since GES data readily provides this information. On the other hand, we have access to detailed information about lane positioning in the naturalistic driving data, while we know virtually nothing about this from the crash data; in the case of initial lane position of the vehicle, the only reasonable choice is to use the alpha form of the variable. If in the future all such details were to become available for large numbers of crashes, they would be preferentially chosen to take the beta form.

At this time we reach block 15 in the block diagram, and at this stage all the major steps are completed relative to the problem formulation. The information feeding into block 15 is controlled by a sampling plan that reflects real-world frequencies in crash and driving populations. The sampling plan is described in Section 7, but it is worth mentioning a particular point from the sampling of alpha parameters from naturalistic driving. Severe inadvertent lane departures are relatively rare in real-world driving, and it is therefore wasteful to run simulations uniformly based on real-world vehicle kinematics in lane-keeping; in most cases there is no significant lane departure seen in the simulation and therefore a very large number of simulations must be run, just to create a reference set of candidate crashes. To address this, over-sampling of higher risk conditions is required, and the method for this is outlined in Section 2.3 below, as well as in Section 7 where more specific details of the batch simulation process are described.

Another way to limit the number of simulations to a feasible set is to limit analysis to the most commonly occurring scenarios seen in the crash data. In the current study it was found that most relevant crashes are associated with a relatively small subset of relevant driving scenarios (see Section 4), so in this case there was no need to simulate all possible cases. Therefore in blocks 16-18 systematic selections are made to concentrate on those most common cases. As part of the sampling and randomization plan for blocks 16-18, random number seeds are stored during an initial sweep of “without countermeasure” simulations. Hence, when the simulations are repeated “with countermeasure” the exact same initial conditions and parametric conditions apply; if there is no lane excursion and hence the ACAT system does not become activated, the outcome will be unchanged and no repeat simulation “with countermeasure” is required.

The final block (block 19) in Figure 2.1 then performs five major steps:

- processes trajectory information to compute various components of crash probability via a defined “crash metric” (see Section 2.3 and also Section 7)
- estimates driving scenario relative frequencies (weights) as well as crash metric scale factors by matching simulated crash outcomes to the GES crash data
- estimates relative crash frequencies with and without the simulated countermeasure, and calculates a safety benefits estimate
- overlays the effects of system unavailability due to factors such as sensor limitations and driver non-compliance (switching the system off or ignoring its operation)
- conducts feasible sensitivity analyses based on varied assumptions in uncertain parameters

In the following sections we focus on the underlying analysis method developed in this study, starting in Section 2.2 with an existing analysis method that provides much of the required statistical framework.

2.2. ***Basic Analytical Method for Safety Benefits Estimation***

A simple statistical concept had been developed by Najm and others based on Bayes’ theorem, and using posterior probability estimates (Najm et al, 2006, Pack et al, 2006). Let *E* be a global measure of exposure, e.g. all driving in the USA in vehicle miles per year. Then the current number of crashes (of the type of interest) per year is $N_{wo} = E \cdot P_{wo}(C)$, where *wo* denotes *without* the safety system, and $P_{wo}(C)$ is the associated probability of a crash per vehicle mile traveled. Similarly let the subscript *w* denote the corresponding situation *with* the safety system, i.e. $N_w = E \cdot P_w(C)$.

In this approach S_i (*i* = 1,2,…,*n*) represents a set of mutually exclusive pre-crash scenarios, defined so that all or most relevant driving situations that lead to the crash type of interest can be associated with these scenarios. Typically S_i is referred to as a “conflict type” or “pre-crash scenario”. The crash numbers and probabilities can be resolved into contributions from pre-crash scenarios, e.g. with the safety system:

$$
\begin{aligned}
P_w(C) &= \sum_i P_w(C \mid S_i) \cdot P_w(S_i) \\
&= \sum_i \frac{P_w(C \mid S_i) \cdot P_w(S_i)}{P_{wo}(C \mid S_i) \cdot P_{wo}(S_i)} \cdot P_{wo}(C \mid S_i) \cdot P_{wo}(S_i) \\
&= \sum_i \pi_i \, \varepsilon_i \, P_{wo}(C \cap S_i)
\end{aligned} \tag{2.1}
$$

where $\pi_i = \dfrac{P_w(C \mid S_i)}{P_{wo}(C \mid S_i)}$ is called the *prevention ratio* and $\varepsilon_i = \dfrac{P_w(S_i)}{P_{wo}(S_i)}$ is the *exposure ratio*. It is helpful to simplify notation, in preparation for extending the method in Section 2.3. Hence, in the following the subscript *wo* will become implicit – the probability or crash number is in the normal *without* population, so for example the probability of driving scenario S_i occurring in a randomly chosen driving "event" in the US driving population is written $P(S_i)$ instead of $P_{wo}(S_i)$. The corresponding terms for the *with* population is written with a prime symbol, e.g. $P'(S_i) \equiv P_w(S_i)$.

The final term in equation (2.1) can then be re-written

$$
P(C \cap S_i) = P(S_i \mid C) \cdot P(C) \equiv c_i P(C) \tag{2.2}
$$

where $c_i = P(S_i \mid C)$ is the conditional posterior probability (estimated by the historical relative frequency) of the particular scenario occurring given a crash event, and satisfies $\sum_i c_i = 1$. The analysis of Najm et al. (2006) assumes this is available and that the scenario can be uniquely determined from the post-crash conditions. Thus

$$
P'(C) = \sum_i c_i \pi_i \, \varepsilon_i \, P(C)
$$

and subtracting this from the equivalent but trivial equation for $P(C)$

$$
P(C) = \sum_i c_i \, P(C)
$$

yields the benefit equation in terms of crash numbers reduced:

$$
B = E \cdot (P(C) - P'(C)) = \sum_i c_i (1 - \pi_i \, \varepsilon_i) E \, P(C)
$$

or, since $N = E \, P(C)$

$$
B = \left(\sum_i c_i (1 - \pi_i \, \varepsilon_i) \right) N \tag{2.3}
$$

The expression in parentheses is the “system effectiveness” in the form of a weighted sum over the effectiveness $(1-\pi_i \varepsilon_i)$ within each particular scenario, the weighting being derived from the crash record.

The approach requires that the nature of each scenario is unaffected by the presence of the safety technology; only crash and exposure probabilities are affected. It also assumes that the posterior probabilities $c_i = P(S_i \mid C)$ can be reliably and accurately estimated from the crash record. The first condition is met reasonably well as long as the driving scenario is defined to precede any system intervention, and it is hard to see how a dormant intervention system can change the dynamics of the driving event (unless the safety system has a “training” effect on the driver or otherwise leads to driver adaptation – we return to this general point in Section 11).

However the second condition is overly restrictive for the present study. In general it is not possible to derive all relevant aspects of the driving scenario (DS) from the post-crash information. In this study we rely on naturalistic driving data to supply important “within scenario” information – we do not assume a single representative set of conditions for each scenario. And while we have attempted to minimize the amount of "supplemental" information (i.e. not directly available from crash data) at the level of the scenario set S_i definitions, one exception is for crashes on a multi-lane highway; here we do not know the initial travel lane of the vehicle involved in the lane departure or road departure crash. We could treat this as an alpha parameter (see above) and estimate relative frequencies from naturalistic driving data (either from onboard measurements, or from traffic counts by lane). However, as noted, the basic rule adopted in this study is to “prefer beta variables” where this is possible. In the absence of definitive data, we have assumed the initial travel lane to be correlated to the location of the first harmful event in a crash, and since this is reported, it should be possible to estimate relative frequencies in the initial travel lane to match the crash locations. In any individual case this may be inaccurate, but aggregated over large number of events the proportion of actual crash events initiated from the different travel lanes can be estimated with a greater degree of confidence.

Another limitation with applying the standard formulation in the present study is that simulations do not directly predict crash/no-crash outcome. To do so would require a combined simulation of both the subject vehicle (drifts out of lane) and a potential principal other vehicle that happens to be in the adjacent lane or perhaps is parked on the shoulder. It would also require the allocation of other collision partners, such as poles and trees, in the off-roadway part of the environment. Not only does this increase complexity (and create a corresponding need for a very large number of simulations), it also seems unnecessary. The vehicle trajectory is presumed to be largely a result of driver failure to observe lane markings and presumably there would normally be a failure to detect other vehicles and stationary objects. In this case the trajectory is defined in the absence of such other objects, and a generalized crash metric is applied to the trajectory in terms of a probability of collision given the magnitude or duration of the excursion. The details are developed in Section 7. Here we note that the crash metric contains a small number of uncertain parameters (scaling factors). These parameters are expected to be constant across a number of different scenarios, but the optimal choice of these parameters is confounded with the initial lane chosen in the simulation. Hence, in establishing the scenario weights, it is necessary to take account of outcomes (locations of first harmful event) not just the recorded pre-crash conditions. Thus, because of an initial uncertainty in some of the beta parameters, and also in the crash metric, the resulting set of uncertain parameters are lumped into a combined optimization scheme to estimate the scenario weights. A modified version of the benefits equation (2.3) can then be applied, based on the resulting information for the "without" population.

2.3. *Formulation of the Safety Benefits Analysis*

As mentioned above, and further described in Section 4, individual driving scenarios are not simply considered as uniform or homogenous conditions in the SIM benefits analysis. Alpha and beta type parameters are sampled from databases to provide important *within-scenario* variations in vehicle kinematics and detailed highway conditions such as lane width or horizontal curvature.

2.3.1. **Basic Formulation**

The crash types under consideration are assumed to arise only from a given set of driving scenarios $\{S_1, S_2, \ldots, S_n\}$, and only crashes of the relevant type are considered. Based on a given driving population the "scenario weights" are defined:

$$w_j = P(S_j)$$

i.e. the probability that a randomly chosen moving vehicle is being driven in this condition over a randomly chosen time interval[3]. Given a global measure of exposure *E* (defined as a large set of discrete driving events, e.g. 10 second intervals of driving) the scenario S_i is expected to occur $w_i E$ times. We define a transition probability $T_{ji} = P(C_j \mid S_i)$, i.e. the probability of a crash outcome of type C_j resulting from the driving scenario. The expected number of crashes of type C_j is then

$$N_j = \sum_i T_{ji} w_i E \tag{2.4}$$

and the total population crash number is expected to be

$$N = \sum_j N_j = \sum_i \sum_j T_{ji} w_i E \tag{2.5}$$

[3] The time interval is required to be standardized and suitably small. Note that to quantify exposure in the context of driving it is more convenient to use time rather than distance as the measure of exposure; frequency distributions extracted from naturalistic driving are based on uniform sampling over time, so this is the exposure measure used throughout when driving activities are considered.

The above expressions all relate to the baseline population, without the safety technology. With the technology present, the scenario and transition probabilities may be altered:

$$T'_{ji} = P'(C_j \mid S_i) \quad , \quad w'_j = P'(S_j)$$

where P' denotes the probability in the "with technology" case. Defining the exposure ratio

$$\varepsilon_i = \frac{w'_i}{w_i}$$

the global "expected benefits" in terms of crash numbers reduced is

$$B = N - N' = \sum_i \sum_j (T_{ji} - \varepsilon_i T'_{ji})\, w_i E \qquad (2.6)$$

This is an abstract form of benefits estimate, because it hypothesizes the historical driving population (with associated crash data) to have been retro-fitted with the safety technology, and no other changes. However, in principle, it does provide a standardized estimation of the performance of the safety system, independent of the particular crash types addressed. The fundamental aim of this research is provide a systematic way of estimating the components of the benefits equation using simulation, crash data, naturalistic driving data, as well as objective tests.

We define crash outcome probabilities (or "outcome weights") relative to the chosen crash types,

$$c_j = P(C_j \mid C) = \frac{P(C_j)}{P(C)} = \frac{N_j}{N}$$

and this clearly satisfies $\sum_i c_i = 1$. The probability of a crash (of any relevant type) is given by

$$P(C) = \frac{N}{E}$$

so dividing equation (2.4) by N we find

$$c_j = \sum_j T_{ji} \overline{w}_i \qquad (2.7)$$

where we have introduced the normalized DS weights,

$$\overline{w}_i = w_i / P(C).$$

In the following we will determine a metric for crash probability based on simulation outcomes and in this way the transition probabilities T_{ji} and T'_{ji} can be estimated. The outcome weights are by definition available (estimated) from GES crash data, so – provided the outcomes are defined with sufficient detail to provide redundancy in equations (2.7) the normalized DS weights can be estimated by optimizing $\overline{w}_i$ to give best fit. Since the equations are linear, a suitable optimal solution is given by the generalized inverse matrix $T_{jk}{}^{\dagger}$ (which reduces to the inverse matrix in the special case that *T* is square and nonsingular):

$$\overline{w}_j \approx \sum_k T_{jk}{}^{\dagger} c_k$$

Using this, the above benefits estimator (2.6) may be re-written:

$$B = \sum_{i\,jk} (T_{ji} - \varepsilon_i T'_{ji}) T_{ik}{}^{\dagger} c_k N \tag{2.8}$$

Typically we assume $\varepsilon_i = 1$, meaning that the presence of the technology does not affect the frequency at which the various scenarios occur. In that case $B = \sum_{i\,jk} (T_{ji} - T'_{ji}) T_{ik}{}^{\dagger} c_k N$, and the expression for benefits estimation is seen to be a weighted sum of the differences $T_{ji} - T'_{ji}$.

In analysis of the above type, a *prevention ratio* is normally defined for any given scenario; using the present notation this is

$$\pi_i = \frac{P'(C \mid S_i)}{P(C \mid S_i)} = \frac{\sum_j T'_{ji}}{\sum_j T_{ji}} \tag{2.9}$$

i.e. it is easily found from the transition probabilities as the ratio of (in practice it is estimated from the estimated transition probabilities) but in the present formulation it is a derived quantity, not a fundamental one.

Another derived quantity is the *system effectiveness* defined as the estimated relative reduction in crash numbers within a given class of cases. In the present formulation, the class can refer to outcomes or driving scenarios or both. The relative reduction for outcome *j* and scenario *i* is

$$e_{ji} = \frac{n_{ji} - n'_{ji}}{n_{ji}} = \frac{T_{ji} w_i - T'_{ji} w'_i}{T_{ji} w_i} \tag{2.10}$$

where (see equation (2.5)) $n_{ji} = T_{ji} w_i E$ is the expected crash number in the given case. Summation (separately in the numerator and denominator) provides the *system effectiveness* in any given class (e.g. sum over all *j* to obtain the estimated system effectiveness for scenario *i*). Partial summation can be used to obtain composite system effectiveness values, for example separated into rural and urban driving.

The above formulation is greatly simplified in the case where driving scenarios and crash outcomes are uniquely associated. In that case T_{ji} is a square diagonal matrix and the benefits estimator (2.8) reduces to the simple form presented in equation (2.3)

$$B = \sum_j (1 - \varepsilon_j \pi_j) c_j N$$

The above expressions need to be estimated from available data, and we present a method for this based on both crash and naturalistic data. This requires some further definitions based on a parametric resolution of the above scenario probabilities.

2.3.2. Scenario Parameters

Here and in the next sub-section we consider the role of alpha and beta parameters relative to the above formulation, and also consider how terms in the benefits equation are estimated. The use and meaning of these parameters will become clearer in Section 4, when they are defined in terms of the lane departure crash problem. In Section 9 the application of the method will also be described in a more practical way.

Real world driving conditions that give rise to the relevant driving scenarios are assumed to be described via an underlying joint distribution $P(\alpha, \beta)$; this is the probability density relevant to making a random selection from the target driving population. This probability density is not directly available, but is basic to the analytic formulation.

In terms of underlying probability density functions

$$w_i = P(S_i) = P(\beta \in S_i) = \int_{S_i} P(\alpha, \beta)\, d\alpha\, d\beta$$

Here beta integration covers the subset $\beta \in S_i$. We expect $\sum w_i < 1$, since the scenarios under consideration only cover a subset of all driving situations. Also, as mentioned, these scenario weights are not directly known, and are inferred from crash data and simulations using frequencies of crash outcomes and using a best fit estimator (the pseudo-inverse matrix taking this role in the above equations).

To simplify further, we assume that within each scenario α and β are statistically independent, i.e. given $\beta \in \beta(S_i)$,

$$P(\alpha \mid \beta \in S_i) = f_i(\alpha)$$

independent of the particular value $\beta \in S_i$; then

$$\int f_i(\alpha)\, d\alpha = 1$$

and we reduce the *within-scenario* variations to a set of "alpha only" probability density functions $f_i(\alpha)$. We also obtain

$$f_i(\alpha) = P(\alpha \mid \beta \in S_i) = \frac{P(\alpha \cap \beta \in S_i)}{P(\beta \in S_i)} \equiv \frac{P(\alpha, \beta_i)}{w_i}$$

and hence

$$P(\alpha, \beta_i) = w_i f_i(\alpha) .$$

2.3.3. **Estimation**

The various terms within the benefits equation (2.8) are estimated from data by random sampling from the simulated scenarios S_i. We randomly select an instance for simulation: $\beta_i(k) \in \beta(S_i)$, and results are averaged over many instances *k*. The corresponding instance $\alpha(k)$ for α is randomly sampled (with distribution conditional on the scenario choice). Each simulation provides a trajectory $\xi(k)$ (instance *k* from scenario *i*, the scenario label is now implicit) and a crash probability is extracted from each trajectory, $P(k) = h(\xi(k))$. Here the crash metric *h* includes a small number of parameters used to balance crash probabilities for different sub-metrics based on crash outcome variables. For any specific choice of (α, β) - and here alpha includes random number seeds - a unique $\xi_{\alpha,\beta}$ trajectory results. The crash probability is

$$P(C \mid \alpha, \beta) = h(\xi_{\alpha,\beta})$$

The crash probability averaged across driving scenario S_i is then a weighted integral (or sum) over alpha's

$$P(C \mid S_i) = \int h(\xi_{\alpha,\beta}) f_i(\alpha) d\alpha \tag{2.11}$$

Resolving this further in terms of crash outcomes C_j (we will use the first harmful event in GES data) the crash metric can also be reduced to sub-components

$$P(C_j \mid \alpha, \beta) = h_j(\xi_{\alpha,\beta})$$

The transition probability T_{ji} is then given by

$$T_{ji} = P(C_j \mid S_i) = \int h_j(\xi_{\alpha,\beta}) f_i(\alpha) d\alpha \tag{2.12}$$

This equation is fundamental to the benefits estimation, and the expectation with respect to the multivariate probability density $f_i(\alpha)$ may be estimated by random sampling from the naturalistic driving data. In Section 7 we provide a practical method for achieving this.

2.4. ***Summary***

In this section, an analytical framework for the ACAT study has been formulated. The cornerstone of the methodology is in the defined *driving scenarios*, a set of natural driving conditions under which the safety system is expected to operate and potentially influence crash rates. In most cases the top level scenario definition is to be directly associated with the pre-conditions for crash events represented in the NASS GES database, though the method allows additional freedom to specify scenario components, such as initial travel lane, which are unknown but later inferred at the crash population level. Two classes of parameter have been introduced to allow sampling of model parameters and initial values of dynamic variables from distributions; in this way scenarios are resolved in detail, and batch simulations are to be performed, based on Monte Carlo methods, and results aggregated. In the following sections we turn to the empirical data that supports this analysis.

3. Data Sources

This section describes the various data sources for the Safety Impact Methodology (SIM). The SIM is a forecasting approach, utilizing a high level of modeling and simulation, since the safety technologies under consideration have not yet been widely deployed. Predictive evaluation requires objective data and the most powerful predictor is one that fuses data from all relevant sources and combines them with the most plausible assumptions.

In order to do modeling and simulation of inadvertent lane departures, a large number of data sources are needed. Apart from details on the how and why the crashes of interest take place, any approach that makes critical use of dynamic simulation to predict safety benefits must have a way to characterize vehicle-system-driver interactions in the scenarios and environments of interest. Objective testing is a natural way to characterize these interactions. One must also retrieve data for calibrating and validating the simulation models in other aspects, such as driver and vehicle behavior during normal lane keeping. Furthermore, the typical roadway properties must be identified, along with what can be expected in terms of how the technology will perform on such roads during the type of weather conditions typical for the crashes of interest.

The data sources used for this ACAT project are therefore quite diverse, and the approach leads to a fusing of objective test data, crash data and naturalistic driving data. Regarding the latter two, if the relevant driving scenarios could be adequately defined from crash information alone, the need for naturalistic data is reduced and possibly removed. But when the driving scenarios are not sufficiently detailed in the crash data, information from naturalistic driving is crucial for supplementing the crash data. On the other hand, because the whole crash problem is inherently complex, it is hard to conceive that naturalistic data alone can be sufficient.

The following data sources have been used in this research:

- Design information and algorithms associated with the safety technology
- Basic scientific knowledge about vehicle dynamics and driving dynamics
- Detailed investigations of crash causation
- Nationally representative crash databases (e.g. NASS GES, CDS, etc.)
- Databases of naturalistic driving (obtained from previous field operational tests)
- Databases of roadway characteristics (e.g. HPMS)
- Objective tests in the form of detailed technical tests of the vehicle and its safety features, typically on the test track
- Objective tests designed to capture typical ranges of human performance where the driver is in the loop, typically on a test track or in a driving simulator.

In particular, Figure 3.1 shows how data sources are used before and during the batch simulations which are at the heart of the ACAT SIM analysis. DVET sub-models require parameters and initial conditions before any dynamic simulation can be run. The driver model parameters are mainly derived from human factors tests in the driving simulator, but also supplemented by naturalistic driving data (dashed line in the figure). The vehicle model is essentially fixed in terms of its parameters, being based on a representative mid-sized passenger car similar to that used in the VIRTTEX driving simulator (see Section 6 for details). Initial values for vehicle kinematics such as speed, yaw rate and lane position are determined by the naturalistic driving data. The environment model uses information such as highway geometry and surface friction, and these are obtained from Michigan highway data (HPMS data); to obtain these very specific data, Michigan crashes are matched to national GES crashes and then sampled. Because the Michigan crashes are geo-located it is possible to join crash data with the corresponding highway data and extract the necessary highway information. Finally the LDW technology model has parameters set from technical testing of the system (see Section 5).

Secondary data sources also exist for running batch simulations: corporate knowledge (especially about vehicle and technology sub-models), prior published research and preliminary simulations all provide important background information. For example a preliminary simulation is needed to complete the definition of vehicle initial conditions, so that steering angle and throttle pedal are pre-set to a stable equilibrium condition. The background information also includes any fixed initial conditions, such as the representation of the driver state: here visual attention is assumed to be initially switched off, even though it will be recovered later in the simulation, either by random attention switching (internal to the driver model) or by an external trigger or warning.

Overall, each of the data sources has a unique and traceable effect on the resulting benefits estimation. The crash data sources are discussed in more detail in Section 3.1, the naturalistic data is covered in Section 3.2, and the HPMS data is discussed in Section 3.3.

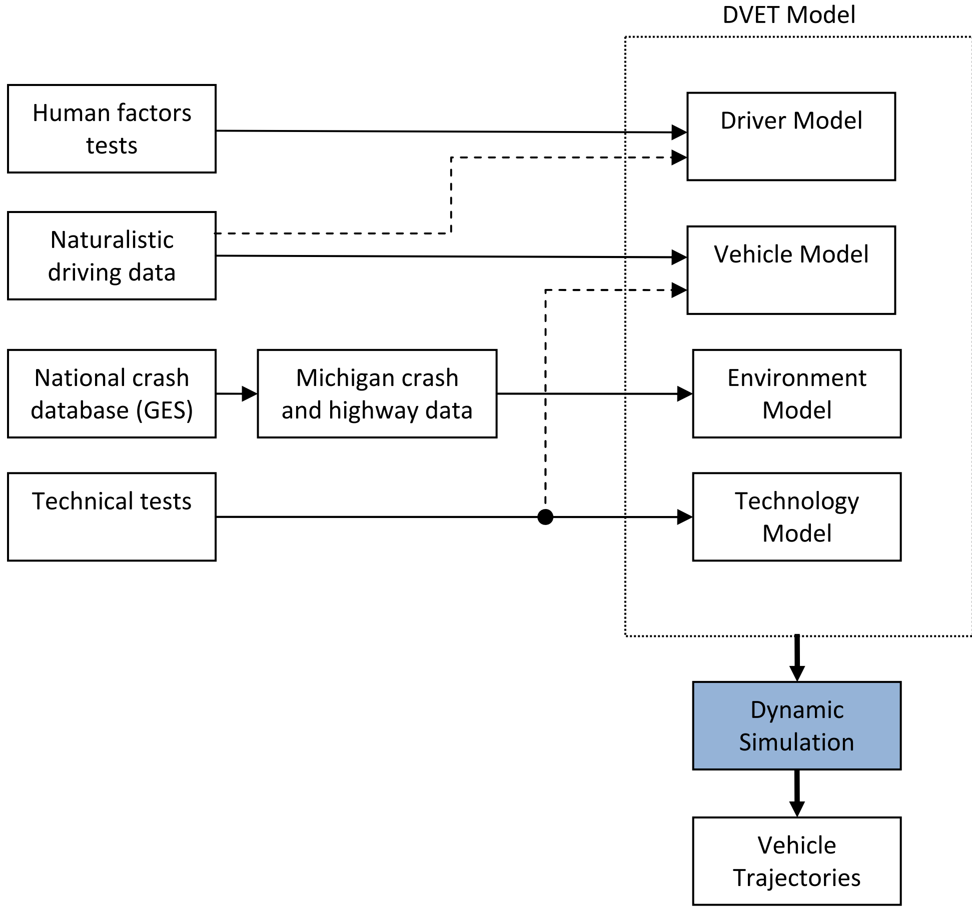

Figure 3.1. Major Data Sources used to Support Batch Simulations

3.1. *Crash data files*

This section describes the crash data sets, including the variables available to identify the target crash types. In addition, the process of establishing the algorithms to identify them is described. Data used directly to develop the target crash types and to describe crash conditions include the National Automotive Sampling System General Estimates System (NASS GES) and NASS Crashworthiness Data System (NASS CDS), crash data from the State of Michigan, and roadway geometric information from the Highway Performance Monitoring System (HPMS).

The **NASS GES** is a nationally-representative sample of police-reported crashes, compiled by the National Center for Statistics and Analysis in the NHTSA. GES is a probability sample of motor vehicle crashes that occurred in the United States. The GES file covers police reported crashes of all severities and all vehicle types. Police accident reports (PARs) are sampled from approximately 400 police jurisdictions within 60 primary sampling units. The sampled PARs are sent to a contractor for coding. The GES data includes a description of the crash environment, each vehicle and driver involved in a crash, and each person involved in a crash. GES data are coded entirely from police reports, without any supplemental investigation. Consequently, the data in GES is limited to what is available on a PAR. GES typically samples about 100,000 motor vehicles involved in 60,000 crashes annually.

The crash data in the GES file was used primarily to identify crash types relevant to the ACAT technologies, to estimate the population size of the relevant crashes, and to characterize at a high level the crash types relevant to the technologies. The GES data include a set of coded variables that captures the sequence from just prior to the initiation of the "crash envelope" to the collision or other harm-inducing event. The crash envelope is defined as extending from the point in which the driver first perceived danger or the vehicle was in an imminent path of collision with another vehicle, animal, or non-motorist to the point at which the driver either has successfully avoided the collision or the collision has occurred. Data elements categorize the vehicle maneuver immediately prior to the critical envelope (in the pre-crash maneuver variable), the event or condition that made the situation critical (critical event), the corrective action taken by the driver, and the stability of the vehicle after the maneuver. There is also an accident type variable that captures the relative position and movement of the vehicles leading to the first harmful event (NHTSA, 2007a). All of these variables appear in the GES and CDS data sets (NHTSA, 2007b, pp 320-335).

Five years of GES data were used in the process of identifying relevant crash types. An analytical file combining GES from 2002 through 2006 was constructed. GES is a sample file, based on a hierarchical, stratified sampling procedure and standard errors are relatively large, particularly for estimates of infrequent crash types. By combining many years of crash data, the estimates are more stable and reliable.

The **NASS CDS** file was also used in the development of the algorithm to identify relevant crash types. CDS is a data system complementary to the GES file. While the GES crash file is a sample of about 60,000 crashes per year covering all vehicle types, to provide nationally-representative estimates of a broad range of crash data, the CDS file contains detailed investigations of roughly 5,000 crashes per year, focused on light vehicles crashworthiness. The CDS investigations go beyond the PARs to include on-site investigation and documentation of the scene, as well as measurements of all crash damage on the vehicles, extensive documentation of crash injuries using hospital and other records, and estimates of the change in velocity (delta *v*) for each vehicle in the crash, where possible. Note that CDS crashes require at least one light vehicle towed from the scene due to damage, so the CDS crashes are, on average, higher in severity than GES crashes.

The in-depth investigations of CDS were used to examine some of the variables and algorithms used to identify crashes relevant to the technologies. Case materials for individual CDS crashes are available on the internet.[4] The materials include the crash summary, scene diagram, scene photos, photos of the vehicles involved, in addition to the coded data elements that are in the

[4]Cases may be accessed at http://www-nass.nhtsa.dot.gov/BIN/NASSCASELIST.EXE/SETXMLFILTER

analytic CDS files. The CDS data uses the same set of pre-crash and crash descriptor variables as the GES file. The CDS case materials were used to develop and refine different selection algorithms and variables used in the GES file. CDS cases that met the selection algorithms developed in GES were reviewed to see if the crashes had the characteristics sought. For example, to determine if cases in which the critical event was loss of control due to speed or other factors were potentially relevant to the technologies under consideration, 40 randomly-selected CDS cases were reviewed to determine if the loss of control occurred prior to lane departure. In almost all cases reviewed, the loss of control preceded lane departure. Similarly, 40 randomly-selected CDS cases were reviewed in which the critical event was lane departure, and in only one was there evidence of loss of control prior to lane departure, even if the vehicle subsequently lost control. Since the detailed case materials for the GES cases are not publicly available, the ability to review more in-depth cases in CDS provides valuable insight into how crashes are classified.

Crash data from **Michigan Crash files** were also used in the development of the driving scenarios for simulation. The GES data do not provide the depth of detail needed for the simulations, particularly in terms of roadway geometries and details such as the presence and composition of shoulders. The Michigan crash file provides a link to such data, because all crashes in Michigan are geolocated (located by latitude and longitude) and those locations can be linked to roadway inventory data. The details of the specific data extracted from the Michigan crash data and the rationale for using the Michigan data will be described later in this report. Here we will simply describe the Michigan crash data files used. The crash data captures the information on all reportable crashes involving a motor vehicle entered by police officers on the Michigan crash report, the UD-10. Reportable crashes involve a motor vehicle in transport on a roadway resulting in a fatality, injury, or property damage of $1,000 or more. Although reporting thresholds do vary, this criteria is reasonably comparable with reporting standards in most other states and thus with crashes in the GES file. Crash data from five years of the Michigan crash file were used, from 2001 through 2005. Many of the variables in the Michigan crash file that describe the crash environment, such as weather, road condition, number of lanes, speed limit, and travel speed, are compatible with similar variables in GES. For others, it was necessary to develop comparable selection criteria. The methods used will be described in the next section.

Data from the Swedish project **FICA** (Factors Influencing the Causation of incidences and Accidents) was also used in the project. The FICA data comprises of 200 crashes occurring in Sweden, from different categories, aggregated in similar situations[5]. Data from FICA was not

[5] The crash analysis method at FICA is called the Driving Reliability and Error Analysis Method (DREAM) (Sandin & Ljung, 2007). DREAM makes it possible to systematically describe and store what can be determined about the causes of a crash. It is sometimes possible to identify causes, collected and managed from police reports, but they are not detailed enough to use as a platform for the development of active safety systems, at least not if the development should be associated with a meaningful evaluation of the potential benefit of the system. DREAM provides a structured way to sort out the reasons attributed for the crash, and classifies them as belonging to a set of categories developed from previous research.

used directly in this project but has contributed with knowledge of how to set values and ranges of values for the different components of the driver behavior model.

3.2. *Naturalistic Driving Data*

Naturalistic data from UMTRI's RDCW (Road Departure Crash Warning) Field Operational Test (FOT) was found to be sufficiently comprehensive in terms of the data for the purposes of populating the relevant elements of the scenarios. The RDCW FOT collected data from 78 drivers distributed evenly by gender and within three age groups ranging from participants in their 20s to participants in their 60s. The total distance traveled was 83,000 miles, covering almost 2500 hours and over 11,000 separate trips spanning a 10-month window that included summer, fall, and winter weather. The drivers used 11 specially instrumented passenger sedans equipped with the safety technologies being evaluated and UMTRI's data acquisition system (LeBlanc et al, 2006).

The RDCW system targeted crashes involving vehicles that drift off the road edge or into occupied adjacent lanes, as well as those involving vehicles traveling too quickly into turns for the driver to maintain control. Included in the RDCW system were two warning functions:

- The first warning function consisted of a lateral drift warning system that was designed to help drivers avoid drifting off the road by providing a set of driver-alert cues when the vehicle was determined to be moving over either dashed or solid lane edge boundaries. The driver was expected to assess the situation and consider steering the vehicle back into the original travel lane if the drift was unintentional. The crashes addressed by such a system are often associated with driver inattention, and fatigue. This system used a camera to observe visual features that delineate lane and road edges, such as painted lane boundaries. Furthermore, a set of onboard radars was used to modulate the warnings when potentially dangerous objects were sensed alongside the edge of the lane or road.
- The second warning function consisted of a curve speed warning system (CSW) that was designed to help drivers slow down to a safe speed before entering an upcoming curve. The desired driver response to a CSW alert was for the driver to consider applying the brakes to slow the vehicle and reduce the lateral acceleration in the curve ahead. The CSW system relied on GPS and a digital map to anticipate the curve location and radius. Measurements of recent driver control actions, such as changing lanes or applying turn signals, were considered in CSW's decision to issue an alert. Both the CSW and the LDW used a set of visual, audible, and haptic cues to alert the driver.

This FOT data proved to contain a rich and detailed source of measurements relevant to lane keeping and drift out of lane events, and while such events were analyzed in the original FOT, the purpose of the data mining carried out here was specific to establishing values and distributions of parameters needed to initialize simulations. Thus the emphasis here was on quantifying baseline driving and variations via key kinematic (vehicle model) variables such as lane deviation and yaw rate. This is further described in Section 4.5.

3.3. *Highway Performance Monitoring System (HPMS) Data*

The HPMS database contains data on the extent, condition, performance, use and operating characteristics of the national highway system (FHWA, 2005). The HPMS is collected annually and submitted by each state department of transportation to the FHWA HPMS group. The database is comprised of universal and sample data. The universal data is a set of basic inventory information. Data fields 1-46, included in the universal set include: unique section id, functional class, rural/urban designation, and traffic volume. The total length of road reported in the universal set should agree with the Certified Public Road Mileage. The sample set, fields 47-98, record more detailed information from a statistical sample of the major functional class system. Sample set data fields include: surface/pavement type, lane width, shoulder type and width, curve and grade classes and speed limit.

Selected data fields from the Michigan HPMS database were spatially joined along the digital map segments for the state public road system. Data from eight counties in southeast Michigan were used because most of the RDCW travel was accumulated in this area. The outcome of this effort was a digital map of SE Michigan road segments that included HPMS data as a part of the spatial database attribute table. While the HPMS database does not include all public road mileage, it does contain approximately 6,747 miles or 77% of the public roads in SE Michigan. The attributes selected include:

- Road Functional Class
- Traffic Volume – AADT
- Curve Radius Group
- Speed Limit
- Lane Width
- Shoulder Type
- Shoulder Width - Left
- Shoulder Width - Right

The roadway descriptive data from HPMS provides a very detailed geometric description of the roadway. When joined to the Michigan crash data, this detail is a significant enhancement to the level of roadway data available in the crash data. For example, while typical crash data distinguishes only straight roads from curved, when the HPMS data are joined to a crash file, the radius of curvature, lane width, and type and width of shoulder can also be used to characterize the scene.

4. Driving Scenarios and Crash Cases

4.1. *Driving Scenarios (DS)*

As described above, the SIM (Safety Impact Methodology) developed by the VFU-team is intended as a computational device to estimate safety benefits of preventive safety systems for the target crashes initiated through lane departures. These crashes include road departure crashes, head-on collisions, sideswipes, and other crash modes. For the analysis conducted here the central criterion is that the driving scenario is one of lane-keeping, and that visual distraction or fatigue can cause an interruption of the driver's lane-keeping control. The scenario must also correspond to a speed range where LDW is operational.

The basic approach is to start by exploring real-world crash circumstances using both analysis of national crash databases and in-depth analysis of recorded crash events in order to understand contributory factors and event sequences, including the role of driver fatigue, distraction and judgment in actual crashes involving lane/road departure. This information is then used to develop a comprehensive set of driving scenarios (DS) which precede the crashes. These DS are not pre-crash scenarios in the sense that a crash inevitably follows DS development. Rather, crashes may or may not result from any given DS as it develops over time (this applies both to real driving and simulations). The DS are thus "coarse-grain" in the sense that they cover a broad parametric range which leads to both crash and non-crash outcomes. Hence, rather than looking at single case accident reconstruction, the aim here is to generate an ensemble of crash/no-crash situations, and then study the extent to which the safety technology changes the proportion of crash/non-crash outcomes in the ensemble.

Following the definition of the DS set, the DS's are then represented in software via a computational model, which time-steps from the parameterized DS and includes random effects (e.g. driver distraction) as well as fixed parameters (e.g. road curvature, lane width, surface friction). Because of the coarse-grained limitation stemming from the recorded crash events, there remains a need to resolve parameters further within each DS – for example to specify precise initial speed, lane position, yaw rate, etc. Such within-scenario frequency distributions are determined using naturalistic driving data. When parameterization is complete, multiple cases are sampled and run in a Monte-Carlo simulation, from which a distribution of outcomes is obtained, as illustrated in Figure 4.1.

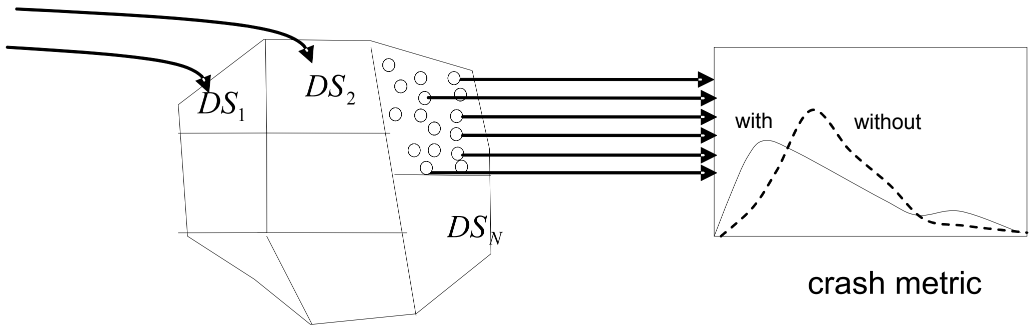

Figure 4.1. Driving Scenario Resolution for Monte Carlo Simulation - With and Without the Specific Safety Technology Being Evaluated.

The advantage of the continuous-time simulation approach is that a very large range of interactions between driver, vehicle, infrastructure and features can be included. The number of simulation runs can be in the tens of thousands, compared to the single or low double digit numbers of objective measurements one could feasibly perform. The simulation approach thus offers a potentially higher degree of representativeness compared to direct objective performance measurement or crash reconstruction based on a small number of actual crashes. This is especially important where interactions and coincidences (e.g. the driver "looks away at just the wrong time") play a substantial role in the event outcome. The main issue to address to ensure these simulations are representative of the real-world crash population is not so much "did we choose the right examples" but "did we represent the mechanisms in a reasonable way".

4.2. *Technology Relevant Crash Cases*

This section describes the method used to identify crash types relevant to the technologies and the process of identifying those crashes in the crash databases. First, the features of the technology that form the framework within which the technologies function are described. The features are described at a level of detail and specificity that can be applied in the crash data. Next, there is a discussion of the approach taken in applying the technology description to the crash data to identify crash types that might be influenced by the technology. An algorithm is developed that identifies a set of crashes that may be addressed by the technology. From this algorithm estimates of the national population of crashes are developed. Next, a set of environmental and roadway factors associated with the crash types are identified in the crash data.

The national crash data, specifically the GES data that are used to identify the general crash types relevant to the technologies, do not include details about the roadway geometry at the crash scenes at a level that is useful for the simulation work. However, specific roadway information is available from the HPMS data in Michigan. To acquire the needed distributions of certain roadway geometrics that are present in the technology relevant crashes, a method is developed to identify the set of technology-relevant crashes in Michigan, and to link those crashes to the HPMS data.

4.2.1. LDW System

The LDW system monitors the vehicle's position in the lane by measuring the distance to lane edge markings. The system issues a warning if it detects a lane excursion. The system has two sensitivity settings, high and low. There is a minimum speed for system operation, and this applies to both settings. LDW is turned to operational mode when speed exceeds 65km/h and is turned to non operational mode when speed drops below 60km/h (40 mph), or when the turn signal is activated. On the low sensitivity setting there are a number of other conditions under which warnings are suppressed:

- The brake pedal is depressed
- The accelerator is quickly depressed (as in passing maneuvers)
- Fast steering inputs occur that indicate an avoidance maneuver

Each of those actions implies that the driver is actively maneuvering the vehicle, and that the lane or road departure is not likely to be inadvertent. In general, conditions such as these might be included in the system model. This is mostly unnecessary in the present study since, in the absence of field data to support driver selection preferences, we will assume a high sensitivity setting. While the driving scenarios are to be selected and simulated in a way that reflects the general intent of lane keeping rather than active driving, only a limited set of variables are available for the selection of crash cases. In general, the lack of any precise information about the driver's use of brake, accelerator and steering controls in actual crashes means that it is possible that some crashes counted as relevant in the GES data will actually be preceded by suppression conditions. To the extent possible (described in the next section), an effort was made to exclude crash cases consistent with the suppression conditions. However, it is recognized that some cases include one or more of the suppression conditions, and this may potentially introduce errors in analysis. We address the issue of simple speed-based suppression conditions via the estimation of system availability – see Section 9. For the more complex ones (e.g. upper limit on steering velocity), this would more appropriately be included in the system model, though for the high sensitivity setting this condition is not relevant. The use of the turn signal remains problematic however, since it is not available in the crash data (no such filter can be applied) and it was not possible to include a naturalistic-based analysis of turn signal use within the scope of this project.

4.2.2. Selection of Crash Conditions Based on GES Data

The GES data include a set of coded variables that span the sequence from just prior to the initiation of the "crash envelope" to the collision or other harm-inducing event. The crash envelope is defined as extending from the point in which the driver recognizes an impending danger or the vehicle was in an imminent path of collision with another vehicle, animal, or non-motorist to the point at which the driver either has successfully avoided the collision or the collision has occurred. Data elements record the vehicle movement prior to the crash envelope (in the pre-crash maneuver variable, named P_CRASH1 in the NHTSA's Statistical Analysis

System (SAS) analysis files), the event or condition that made the situation critical (critical event, named P_CRASH2), the corrective action taken by the driver (P_CRASH3), and the stability of the vehicle after the corrective action (P_CRASH4). There is also an accident type variable that captures the relative position and movement of the vehicles leading to the first harmful event. All of these variables are also included in the CDS data sets (although the variables have different names).

The approach to capturing crash events in GES (and CDS as well) is well-suited to evaluation of crash avoidance technologies. Many other crash data systems focus on the first harmful event, or provide a sequence of events in the crash, which record the series of harmful events. But in crash avoidance research, information about the vehicle state prior to the initiation of the crash sequence and any harmful event is of interest. The technologies under consideration here monitor vehicle position within the lane in normal driving, prior to any crash or conflict. This information is recorded in P_CRASH1, which describes the vehicle's activity prior to the driver's realization of an impending critical event or just prior to impact if the driver took no action or had no time to attempt any evasive maneuvers (NHTSA, 2007b). Similarly, the critical event variable records how the crash sequence was initiated, which is not the same as the first harmful event in a crash. It is the event that precipitated the crash and thus is the event that the crash avoidance technologies are designed to address. Vehicle movement prior to the critical event (P_CRASH1) and critical event for this vehicle's first impact (P_CRASH2), in the GES file are of primary interest in identifying the relevant crash types.[6]

Pre-crash movement (P_CRASH1) includes a comprehensive set of actions or vehicle states prior to the initiation of the crash sequence. Some of the levels may be characterized as active maneuvers, indicating active control and engagement by the driver. These maneuvers include turning, stopping, starting, slowing, passing, making a U-turn, backing, entering or leaving a parked position, and avoidance maneuver. In these cases the resulting crash mode is not likely to be connected with the design intent of the system, and even if a warning were given it is unlikely that the crash mode would be affected by a lane departure warning (or emergency lane assist intervention). In the case of a low sensitivity setting for the LDW system, the active maneuvering would also suppress any alert.

Three pre-crash maneuver code levels are consistent with the types of crashes in which the LDW may be activated: “going straight”, “negotiating a curve,” and “changing lanes”. The LDW functionality described above constrains relevant crashes to those in which the crash sequence was initiated by a lane/road departure of the reference vehicle. The variable P_CRASH2 in GES includes codes that identify crashes for which the critical event was a lane or road departure. There are four code values that identify LDW relevant crashes: "Vehicle traveling over the left

---

[6] This approach is analogous to aspects of the “Universal Description,” in Burgett, et al. (2008). The cited Universal Description uses some of the same variables in classifying crashes into mutually exclusive categories. The approach taken here, however, begins with the pre-crash maneuver. The Universal Description is a useful classification for crash avoidance research. The approach in the current methodology development work is tailored to the specific operations and capabilities of the technologies being evaluated here.

lane line", "Vehicle traveling over the right lane line", "Vehicle off the edge of the road on the left side" and "Vehicle off the edge of the road on the right side".

The identification of target crash types was accomplished primarily by the two variables, Vehicle movement prior to critical event (P_CRASH1) and critical event for this vehicle's first impact (P_CRASH2). However, a number of other variables were included to refine the identification of crashes that might be influenced by LDW. These variables record the number of vehicles in the crash, whether the vehicle was involved in the first harmful event in the crash, the travel speed of the vehicle, and whether the driver was under the influence of alcohol or drugs.

Samples of CDS cases were reviewed during the process of identifying variables and specific code levels of the variables to use in identifying relevant crashes. The purpose of the reviews was to see if the codes selected identified the types of crashes that would be addressed by the technologies. Scene diagrams, scene photos, and the researcher’s text description of the crash are all available for individual crashes. Since CDS uses the same set of variables and definitions as GES, the CDS cases were used as a sample of the types of events that would be selected in GES. The CDS case review confirmed that the algorithm identified the appropriate crashes in GES.

The Volvo LDW technology is currently designed to be active at speeds greater than 60 kph (40 mph). Travel speed and speed limit were both used to identify crashes that took place at speeds at which LDW would be active. The variable for travel speed is based on the police-reported estimate of travel speed. The variable is unknown in about 63 percent of the crashes and the available speed data are of uncertain accuracy. This is not unexpected because it is very difficult for the reporting officer to estimate that value. However, because travel speed is unknown in such a high percentage of crash involvements, the posted speed limit was considered for those cases where travel speed is unknown. After examining the distribution of known travel speeds by speed limit, it was decided to use a posted speed limit of 35 mph or greater as a surrogate for travel speed of at least 40 mph, where travel speed was unknown.

The number of vehicles involved in the crash was used to identify crashes in which one or two vehicles were involved. In addition, only crashes in which the reference vehicle was involved in the first harmful event are included. There are crashes which are initiated by another vehicle or vehicles, and the reference vehicle is involved as a consequence of that earlier event. For example, two other vehicles may collide and one might be forced into the lane of the reference vehicle. Such crashes are excluded. Additionally, cases in which the driver was coded as using alcohol or drugs were excluded due to the uncertainty of driver response in those cases.

Vehicle types were also limited to simplify the modeling of the vehicles and drivers. Light trucks (pickups, sport utility vehicles, and vans) were excluded because of differences in handling properties as well as in the types and proportions of crashes in which they are involved. Only vehicles identified as passenger cars in the BODY_TYP variable, as shown in Table 4.1 below, were included.

Table 4.1. GES Body Type Automobiles Included

Convertible (excludes sun-roof, t-bar)
2-door sedan, hardtop, coupe
3-door/2-door hatchback
4-door sedan, hardtop
5-door/4-door hatchback
Station wagon (excluding van and truck based)
Hatchback, number of doors unknown
3-door coupe
Other automobile type
Unknown automobile type

Based on the general crash characteristics identified above as relevant to the LDW technologies, four dynamically-distinct crash types were identified as relevant to the technologies.

- Single-vehicle road departure
- Prior lane-keeping, lane departure
- Changing lanes, lane departure
- Other lane or road departure, prior lane-keeping or changing lanes

Table 4.2. Rules in GES to Identify Target Crash Involvements

Vehicle	SAS code to identify
Passenger cars	If 1<=body_typ<=9 or body_typ=17;
Scenario	**SAS code to identify**
Single-vehicle road departure	if veh_invl=1 and p_crash1 in(1,14,15) and p_crash2 in(12,13);
Prior lane-keeping, lane departure	else if veh_invl>1 and p_crash1 in(1,14) and p_crash2 in(10,11);
Changing lanes, lane departure	else if veh_invl>1 and p_crash1=15 and p_crash2 in(10,11)
Other lane or road departure, prior lane-keeping or changing lanes	else if p_crash1 in(1,14,15) and 10<=p_crash2<=13;
Filter	**SAS code to identify**
Travel speed of 40 mph or more	if 40<=speed<200 or (speed=999 and 35<=spdlim_h)
Driver not coded for alcohol or drug use	if per_alch ne 2 and per_drug ne 2;
Vehicle in first harmful event	if not (veh_invl>2 and acc_type=98);

Table 4.2 shows the SAS code for the GES crash data, to identify the vehicles, crash types, and exclusions. Note that the final entry in the table is intended to exclude vehicles that are not involved in the first harmful event in the crash. Code level 98 of the accident type variable is used for "other accident type" (i.e., not one of the other crash types in the variable) or for vehicles involved in a crash after the first harmful event occurred elsewhere in the crash. The filter excludes vehicles in any crash with more than 2 vehicles and for which the accident type for that vehicle is 98.

A word of caution is offered about using the GES data, that are coded primarily from police reports. There is no supplementary investigation of the crashes. Therefore, the data are limited to the level of information available on police reports, and in a subsidiary way to what is available to the reporting officers. Police officers are trained in completing the crash form and often have considerable experience in crash reporting. But only the most serious crashes are typically investigated by a crash reconstructionist. Instead, the officer builds the narrative of events from personal observation of the physical evidence, interviews with the involved parties, and sometimes witnesses. As mentioned, the LDW function is suppressed if the turn signal is activated. But there is no direct evidence from the vehicle whether a turn signal was used; nor do any of the codes in P_CRASH1 capture whether the turn signal was activated. Similarly, there is usually little direct evidence about the vehicle's travel speed after the crash, only statements by witnesses and involved parties, and to a less extent inferences from crash data. The movements of the vehicle, as recorded in the database, imply certain things about what the driver did, but it is useful to keep in mind the original source of the information. In particular, GES data are based on police reports that are of uneven quality and completeness and it is important to keep in mind that data elements are usually estimated or reported rather than objectively measured or recorded.

4.3. *Estimates of Relevant Crash Involvements*

The GES crash data files for 2002 through 2006 (the most recent available when the analysis was done) were used to identify relevant crash types and estimate the scope of the crashes the LDW technologies might address. Over that period, an annual average of 6,107,000 passenger cars, as defined in Table 4.1, were involved in police-reported crashes. That total includes all passenger cars involved in crashes, regardless of travel speed, whether the vehicle was involved in the first harmful event, or whether the driver was coded as using alcohol or drugs. However, the exclusions, in terms of travel speed, driver alcohol/drug use, and number of vehicles involved in the crash, reduces the number of relevant crashes, i.e., those crashes that might be affected by the technologies.

Table 4.3 shows the average annual involvements of passenger cars in traffic crashes, estimated from five years of the GES data. The table shows all passenger car crash involvements, split between those at an estimated travel speed of 40 mph or more and all others. Essentially, crashes at the 40+ mph travel speed fall into the domain of the LDW system in the sense that the safety technologies under consideration may be active in that domain. Overall, about half of the passenger car involvements, or 2,985,000 in total, fall into that domain. The crash types relevant to LDW are shown separately in Table 4.3, and account for 241,000 annual involvements. These involvements are 3.9 percent of all passenger vehicle involvements, but 8.1 percent of the 2,985,000 passenger car involvements in crashes with speeds > 40 mph. Crash involvements where the driver was coded for drugs or alcohol are excluded from the target types, as are involvements that happened after the first harmful event. An estimated 3,122,000 passenger cars were involved in a crash at travel speeds less than 40 mph. Since they occurred at speeds below the threshold for activation of the technologies, no crash type detail is shown.

The frequencies for the LDW crash types slightly underestimate the true number of such involvements in the crash population because there are a small number of cases with missing data in the variables used to define the crash types. To account for these crashes with missing data, we assume that the target crash types follow the same distribution in the cases with missing data as the cases where all relevant data are known. This adjustment resulted in a new estimate of the annual number of crash involvements relevant to LDW. Table 4.4 shows the estimated crash involvements for travel speed >= 40mph, adjusting for missing data. The total number of crash involvements in the relevant speed domain does not change, because the speed domain identification included imputed variables, which have no missing data. The new estimate for the total of relevant crash involvements is 245,000. These are 8.2 percent of all crashes in the relevant speed domain.

Table 4.3. Annual Average Involvements of Passenger Cars by Crash Types Relevant to LDW Technology and All Other Crash Types, GES 2002-2006

Travel speed domain	Crash type	Annual Estimate	Percent
Travel speed >= 40 mph	Target crash types:		
	Single vehicle, road departure	102,000	1.7
	Lanekeeping, lane departure	41,000	0.7
	Change lanes, lane departure	80,000	1.3
	Other lane/road departure	19,000	0.3
	Subtotal of target	*241,000*	*3.9*
	Other crash involvements	2,744,000	44.9
	Total crash involvements > 40 mph	2,985,000	48.9
Travel speed < 40 mph		3,122,000	51.1
*Total crash involvements, all speeds		6,107,000	100.0

**minor discrepancies in totals due to rounding. Percentages calculated using unrounded frequencies*

Single-vehicle road departures account for 104,000 annual involvements, the most frequent target crash type, and 3.5 percent of involvements within the relevant speed domain. Crashes in which the vehicle was coded as lane keeping and drifted out of lane account for another 41,000 annual involvements. Crashes in which the vehicle was coded as changing lanes and the critical event was the lane departure account for 81,000 annual involvements (2.7 percent of crashes in the relevant speed domain). Finally, the "other lane/road departure" set, which consists primarily of single-vehicle lane departure crashes in which the vehicle collided with a parked vehicle, accounts for 19,000 crash involvements annually.

Table 4.4. Annual Average Passenger Car Involvements in Crashes with Speed ≥40 MPH and Target Crash Types Relevant to LDW Technology Adjusted for Missing Data

Crash type	Annual frequency	Percent
Target crash types		
Single vehicle, road departure	104,000	3.5
Lanekeeping, lane departure	41,000	1.4
Change lanes, lane departure	81,000	2.7
Other lane/road departure	19,000	0.6
Subtotal of target	*245,000*	*8.2*
Other crash involvements	2,740,000	91.8
*Total crash involvements > 40 mph	2,985,000	100.0

**minor discrepancies in totals due to rounding. Percentages calculated using unrounded frequencies*

4.4. *Crash Conditions Derived from GES Data (β -parameters)*

A number of road, environmental, and driver factors were examined in relation to the target crash types. The purpose was to identify environmental and other factors associated with the crash types to be used in the simulations to specify the driving situation at the time of the crash. Table 4.5 shows the percentage distribution across the target crash types for roadway, roadway alignment, weather, road surface condition, light condition, and driver fatigue. (The roadway variable shown in Table 4.5 was constructed using the variables for number of travel lanes and traffic flow). For comparison purposes, the table also shows the distribution of each of the characteristics for all other, non-target crashes with travel speed over 40mph, that is, all the other crashes in the speed domain for LDW. Thus, the parameters are univariate distributions within a factor. Joint distributions of multiple factors are not provided. This may impact the quality of benefits estimations. Collectively these parameters are referred to as the β parameters in the SIM process.

The distributions of each of the road, environment, and driver factors listed in Table 4.5 vary substantially by target crash type and also typically differ from other crash involvements in the relevant speed domain (speeds > 40 mph). For example, a higher proportion of single-vehicle lane departure crashes occurred on two-lane, two-way undivided roads than any of the other target crash types or the non-target crashes. Single-vehicle lane departure crashes showed relative over-involvement on curved road segments with adverse weather. The lane-keeping, lane departure type occurred more frequently on roads with three or more lanes, especially divided highways. Those crashes occurred less frequently on curved segments than the other target crash types, and more frequently on dry roads in good weather and daylight than any other crash type, including all the non-target crash types. Excluding the lane-keeping lane departure type crash, each of the target crash types involved a greater proportion of fatigued drivers than non-target crashes in the ACAT-relevant speed domain. Only about 0.3 percent of such crashes are coded with a fatigued driver, while the proportion ranges from 8.0 to 9.9 for the single-vehicle road departure, change lanes lane departure, and other relevant crashes types.

The factors in Table 4.5 characterize at a high level the main features of the driving environment at the time of the crash and thus provide parameters for the simulations. The factors shown are referred to as beta parameters, in that they represent relatively stable conditions at the crash site that are captured in police-reported crash data. Alpha parameters, which are discussed in the next section, are more transient and cannot be recovered from the available national crash data. These include factors such as lane position, speed, and driver attention. Distributions of alpha parameters are taken from naturalistic driving data, but the beta parameters can be known with some confidence from the crash and other data. Road type and road curvature are stable and are assumed to be reliably reported. Weather, road surface condition, and light condition are less stable but still considered sufficiently reliable in the crash data. Driver fatigue is difficult to identify in a reliable way, and while both crash and naturalistic data support some form of identification, in this study it was only considered feasible to use the crash data. While many cases of fatigue may be missed, it is assumed that the cases that are identified are true cases of fatigue and it is further assumed that they are representative of fatigue-related crashes in general.

Table 4.6 shows the top 25 detailed driving scenarios for crashes relevant to LDW. These are formed by cross-classifying the target crashes by the factors identified in Table 4.5 as associated with the crashes. To help limit the scenarios to a tractable number, it was decided to combine code levels for some of the factors. No aggregation was done for road type and roadway alignment, but the light condition codes were simplified to daylight and not daylight, and weather and roadway surface condition were combined to form a three-level variable: no adverse weather and dry road; adverse weather and not dry road; and no adverse weather and not dry road. Weather condition and road surface were combined since they strongly interact. Only 0.6 percent of cases were coded with an adverse weather condition but dry road surface, and it was felt that this combination could be ignored.

Table 4.5. Road, Environmental, and Fatigue Factors Associated with Target Crash Types (β parameters), GES 2002-2006

Factor	Target Crash Types				Non-Target Crashes with Speed ≥40 MPH
	Single Vehicle Road Departure	Lane Keeping, Lane Departure	Change Lanes, Lane Departure	Other Relevant	
Roadway					
One lane	0.3	0.3	1.5	1.4	1.7
2 lane, 2-way, undivided	38.9	2.6	34.4	31.2	21.8
3+ lanes, undivided	19.0	26.2	11.9	16.5	26.8
3+ lanes, divided	26.0	56.5	35.5	38.3	31.3
Unknown	15.8	14.4	16.7	12.6	18.4
Roadway alignment					
Straight	73.7	93.0	79.1	77.0	90.0
Curve	26.3	7.0	20.9	23.0	10.0
Road surface condition					
Dry	73.1	85.7	81.3	80.4	76.1
Wet	20.8	13.1	15.5	16.2	19.0
Snow or slush	4.6	0.8	2.0	2.3	2.6
Ice	1.5	0.2	1.0	0.9	2.0
Sand, dirt, or oil	0.0	0.1	0.1	0.2	0.1
Other (Specify)	0.0	0.1	0.1	0.1	0.1
Weather					
No Adverse	81.2	89.8	86.6	86.3	83.6
Rain	13.5	8.4	10.0	10.2	12.1
Sleet	0.5	0.1	0.1	0.2	0.3
Snow	4.0	1.1	1.6	2.1	3.2
Fog	0.4	0.4	0.4	0.7	0.5
Rain And Fog	0.0	0.0	0.6	0.1	0.0
Sleet And Fog	0.0	0.0	0.0	0.0	0.0
Other	0.5	0.2	0.7	0.3	0.3
Light condition					
Daylight	68.4	75.5	61.0	64.0	70.0
Dark	11.1	5.7	15.6	16.6	10.3
Dark But Lighted	16.5	15.7	19.5	15.6	15.6
Dawn	2.4	1.0	1.9	2.0	1.5
Dusk	1.6	2.1	2.1	1.7	2.5
Fatigued					
No	92.0	99.8	91.2	90.1	99.7
Yes	8.0	0.2	8.8	9.9	0.3

Note: all values are percentages within the factor categories.

Table 4.6. Top 25 Driving Scenarios for Target Crash Types, GES 2002-2006

Scenario ID	Road type	Roadway alignment	Weather & road surface	Light condition	Driver fatigued	Percent
1	2 or more lanes, divided	Straight	Not adverse, dry	Daylight	No	19.7
2	2 or more lanes, undivided	Straight	Not adverse, dry	Daylight	No	9.9
3	2 or more lanes, divided	Straight	Not adverse, dry	Not daylight	No	9.1
4	2-lane, 2-way undivided	Straight	Not adverse, dry	Daylight	No	8.6
5	2-lane, 2-way undivided	Curve	Not adverse, dry	Daylight	No	5.8
6	2-lane, 2-way undivided	Straight	Not adverse, dry	Not daylight	No	5.2
7	2-lane, 2-way undivided	Curve	Not adverse, dry	Not daylight	No	4.1
8	2 or more lanes, undivided	Straight	Not adverse, dry	Not daylight	No	3.1
9	2 or more lanes, divided	Straight	Adverse, not dry	Daylight	No	2.5
10	2 or more lanes, divided	Curve	Not adverse, dry	Daylight	No	2.4
11	2-lane, 2-way undivided	Straight	Not adverse, dry	Daylight	Yes	1.9
12	2 or more lanes, divided	Straight	Adverse, not dry	Not daylight	No	1.9
13	2 or more lanes, divided	Straight	Not adverse, dry	Not daylight	Yes	1.8
14	2-lane, 2-way undivided	Curve	Adverse, not dry	Daylight	No	1.8
15	2 or more lanes, divided	Straight	Not adverse, dry	Daylight	Yes	1.7
16	2 or more lanes, divided	Curve	Not adverse, dry	Not daylight	No	1.6
17	2-lane, 2-way undivided	Straight	Adverse, not dry	Daylight	No	1.4
18	2 or more lanes, divided	Straight	Not adverse, not dry	Daylight	No	1.2
19	2 or more lanes, undivided	Straight	Adverse, Not dry	Daylight	No	1.2
20	2-lane, 2-way undivided	Straight	Not adverse, dry	Not daylight	Yes	1.2
21	2-lane, 2-way undivided	Straight	Adverse, not dry	Not daylight	No	1.0
22	2-lane, 2-way undivided	Curve	Adverse, not dry	Not daylight	No	1.0
23	2 or more lanes, undivided	Curve	Not adverse, dry	Daylight	No	1.0
24	2-lane, 2-way undivided	Straight	Not adverse, not dry	Daylight	No	0.8
25	2-lane, 2-way undivided	Curve	Not adverse, not dry	Daylight	No	0.7
	TOTAL					90.5

The twenty-five scenarios shown in Table 4.6 account for 90.5 percent of target crashes (based on NASS GES number estimates). The top 36 scenarios account for 96.6 percent of target crashes with non-missing data on all variables. Missing data was primarily a problem with the roadway variables; 12.3 percent of the cases could not be classified into one of the crash types because of missing data.

Detailed roadway geometric data are needed to describe the characteristics of the roadway at the point of the crash. This information is not available in the GES data, which only distinguishes curved from straight roads, the number of travel lanes, and a few other details. The GES data do not include other information of interest for lane/road departure crashes, such as lane widths, lane markings, radius of curvature, shoulder width, and shoulder type. This more detailed information is available in roadway inventory files, such as the Highway Performance Monitoring System (HPMS) data on the road system. GES crashes are not geo-located (i.e., located using a standard geographic reference system such as longitude and latitude), so it is not possible to link the GES crashes to a roadway inventory file.

Accordingly, it is necessary to obtain descriptions of the roadway geometry for the target crash types from some other source. Roadway inventory files on the road system in Michigan provide a convenient source. All police-reported crashes in Michigan are geolocated, which allows the crash sites to be linked to the HPMS roadway files and the detailed roadway information they contain. A combination of the State of Michigan digital map database and the Michigan HPMS inventory were used to represent the road system.

The Michigan Department of Transportation annually completes an inventory of public roads for the FHWA HPMS program. The data items establish the characteristics of the road and include items such as: functional class, urban/rural designation, volume and, on a sampled basis, road geometry. The HPMS inventory data was appended to the road segment attributes of the State of Michigan digital base map using ERSI GIS software. While the HPMS inventory includes some 98 data fields, not all fields are collected on the inventoried system. A portion of the fields are collected on a sample basis. The sampled segments have the most detailed geometric data. It was necessary to use crashes that occurred on the sampled segments to obtain the most detailed geometric data.

The use of the Michigan crash data and roadway files carries with it the assumption that crash types and associated characteristics identified in Michigan are reasonably consistent with the national crash picture as represented by GES. If the target crash types can be identified in the Michigan crash data, and if the road system in Michigan is reasonably comparable to the national road system, then the Michigan crashes can be used as a sample of crashes to characterize the distribution of roadway geometries for the target crash types. Note that the assumption is not that Michigan is "typical" of the national experience in terms of weather severity or the physical condition of the roads, which may influence the performance and effectiveness of the technologies. The distributions of weather condition and road surface condition (dry/not dry) are part of the high-level driving scenarios derived in GES and their relative frequencies are taken from the national GES crash data. The Michigan geolocated crashes are used as a sample to obtain the fine detail available, such as shoulder widths, curve transitions, and shoulder types, all of which are assumed to be sufficiently typical of all states.

As will be described below, the Michigan crash file was analyzed to identify the set of crashes that are as close as possible to the target crash types identified in the GES file. In addition, selected standard characteristics of the Michigan road system were compared with national distributions. The purpose of these comparisons was to demonstrate that the Michigan roadway system is reasonably comparable with the national road system. For example, when compared to the national system, the distribution of functional roadway class in the Michigan road system is similar to the distribution of functional roadway class nationally (see Figure 4.2 below). Accordingly, the crash and road data were then fused to develop variables for simulation input. While not as complete as desired, linking these two sources does provide credible inputs of the necessary detail to represent the crashes and the road characteristics.

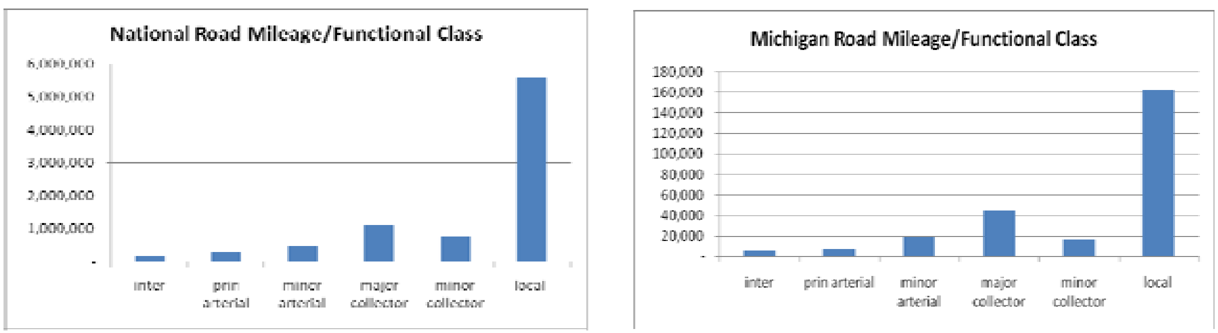

Figure 4.2. Distribution of Road Miles by Functional Class, National Road System and Michigan Road System.

Every effort was made to reproduce the target crash types from the GES analysis in the Michigan crash data. Five years of the Michigan crash file were used for the purpose, from 2001 through 2005. Only passenger cars were selected for the analysis and filters similar to those used for GES were used to identify and exclude travel speeds less than 40 mph and drivers coded as using either drugs or alcohol. The Michigan crash data includes a "prior action" variable that captures what the driver was doing just prior to the crash. The variable is comparable to the GES variable for pre-crash maneuver (P_CRASH1). Though the Michigan variable includes more possible actions, it was possible to match the GES variable coding. Michigan does not code a "critical event" variable, but there is a detailed sequence of events variable that was used to identify crashes precipitated by lane or road excursion. Table 4.7 shows the results of the attempt to reproduce the target crash selection algorithm in the Michigan data, along with the comparable distribution in the GES file. Overall, the distributions are quite similar, and the target crash types account for a similar percentage of all crash types in the LDW speed domain. The results are consistent with the conclusion that the effort to identify comparable crash types in the Michigan crash data was reasonably successful.

Table 4.7. Distribution of Crash Types Relevant to LDW Technology in Michigan and GES Michigan, 2001-2005; GES 2002-2006

Crash type	Michigan		GES
	Annual frequency	Percent	Percent
Target crash types			
Single vehicle, road departure	39,652	3.5	3.4
Lanekeeping, lane departure	11,814	1.0	1.4
Change lanes, lane departure	15,923	1.4	2.7
Other lane/road departure	1,282	0.1	0.6
Subtotal of target	*68,671*	*6.1*	*8.1*
Other crash involvements	1,062,973	93.9	91.9
Total crash involvements > 40 mph	1,131,644	100.0	100.0

The more detailed crash conditions from the GES data, as shown in Table 4.6, were also reproduced in the Michigan crash data. Michigan codes variables for number of lanes and traffic way flow that are identical to those in GES, so the road type variable could be reproduced. Light condition, weather, and roadway surface condition coding are also identical between the two files. Michigan also captures driver fatigue as part of a multiple-response driver condition variable. Unfortunately, Michigan does not discriminate curved from straight roads in the coded crash data, so it was necessary to link in this information from the roadway files, which was done by linking in the needed information from the HPMS data. In the end, between the coded crash data and the information about roadway curvature from the HPMS data, it was possible to identify detailed crash conditions relevant to the LDW in the Michigan data that are comparable to those in the GES crash data.

The combined Michigan crash and roadway data can reproduce the crash conditions shown in Table 4.6 and the distributions of roadway geometries. The steps involved resulted in a substantial reduction in the number of cases. Of the original 68,671 crash involvements identified, 66,595 had valid longitude/latitudes. The crashes were then spatially located, joining points that fell within 75 feet of an HPMS road segment on which the crash occurred. This step was carried out to help eliminate those crashes located in parking lots or minor streets. Using the 75 feet location criteria resulted in 51,239 of the initial 66,595 crashes being processed. The most detailed roadway geometric information is available only for sampled segments; 10,441 of the crashes fell along sampled road segments. Figure 4.3 below illustrates the process to fuse the Michigan crash and HPMS road data.

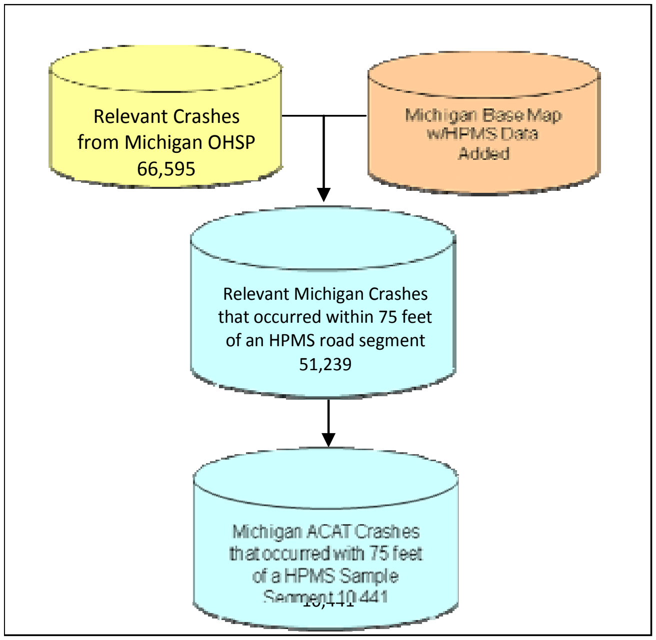

Figure 4.3. Fusing Michigan Crash and Road Data

Table 4.8 shows the specific roadway information that is linked to the crash data by this method from the sample segments.

Table 4.8. Roadway Characteristics from Michigan Sampled Segments

Road Functional Class	Lane Width
Traffic Volume – AADT	Shoulder Type
Curve Radius Group	Shoulder Width - Left
Speed Limit	Shoulder Width - Right

Certain critical roadway geometric features, such as roadway radius of curvature, lane width, shoulder type, and shoulder width vary between roads in urban and in rural areas. Moreover, the sampled roadway segments, which have the most detailed geometric information, over-represent urban areas. Unless the urban/rural difference is accounted for, the sampled segments would be biased toward the distribution of geometrics in urban areas. Accordingly, a method was developed to match the crash conditions on the sampled segments in Michigan to national estimates from GES, accounting for the urban/rural split.

GES and the Michigan crash data both include a variable that codes the population of an area, although in different ways. Table 4.9 shows the code levels used to record the population of the crash area in each file:

Table 4.9. Population Categorization in GES and Michigan Crash Data

GES	Michigan
<25,000	<5,000
25,000 to 50,000	5,000 to 50,000
50,000 to 100,000	50,000 to 200,000
>100,000	>200,000

Though the two scales do not match well, they include a common boundary at 50,000, which was adopted as the split point between rural and urban.

Table 4.10 shows the count of target crashes on sampled segments in the Michigan data by driving scenario and split by the urban/rural dimension. Each cell can be mapped directly to the same cell in the national GES data. The count in the matrix shows the number of crash involvements in the Michigan data that are available to provide “seed” parameters for the simulations.

Table 4.10. Top 25 Crash Cases for Target Crash Types on Sampled Roadway Segments, Michigan 2001-2005

Road type	Roadway alignment	Weather & road surface	Light condition	Driver fatigued	Rural	Urban
2 or more lanes, divided	Straight	Not adverse, dry	Daylight	No	117	1,731
2 or more lanes, undivided	Straight	Not adverse, dry	Daylight	No	76	1,274
2 or more lanes, divided	Straight	Not adverse, dry	Not daylight	No	45	683
2-lane, 2-way undivided	Straight	Not adverse, dry	Daylight	No	104	389
2-lane, 2-way undivided	Curve	Not adverse, dry	Daylight	No	242	14
2-lane, 2-way undivided	Straight	Not adverse, dry	Not daylight	No	64	219
2-lane, 2-way undivided	Curve	Not adverse, dry	Not daylight	No	122	9
2 or more lanes, undivided	Straight	Not adverse, dry	Not daylight	No	17	338
2 or more lanes, divided	Straight	Adverse, not dry	Daylight	No	34	342
2 or more lanes, divided	Curve	Not adverse, dry	Daylight	No	152	143
2-lane, 2-way undivided	Straight	Not adverse, dry	Daylight	Yes	7	25
2 or more lanes, divided	Straight	Adverse, not dry	Not daylight	No	29	395
2 or more lanes, divided	Straight	Not adverse, dry	Not daylight	Yes	16	71
2-lane, 2-way undivided	Curve	Adverse, not dry	Daylight	No	54	2
2 or more lanes, divided	Straight	Not adverse, dry	Daylight	Yes	16	56
2 or more lanes, divided	Curve	Not adverse, dry	Not daylight	No	64	45
2-lane, 2-way undivided	Straight	Adverse, not dry	Daylight	No	42	121
2 or more lanes, divided	Straight	Not adverse, not dry	Daylight	No	8	182
2 or more lanes, undivided	Straight	Adverse, Not dry	Daylight	No	12	216
2-lane, 2-way undivided	Straight	Not adverse, dry	Not daylight	Yes	5	27
2-lane, 2-way undivided	Straight	Adverse, not dry	Not daylight	No	36	185
2-lane, 2-way undivided	Curve	Adverse, not dry	Not daylight	No	60	3
2 or more lanes, undivided	Curve	Not adverse, dry	Daylight	No	22	26
2-lane, 2-way undivided	Straight	Not adverse, not dry	Daylight	No	22	67
2-lane, 2-way undivided	Curve	Not adverse, not dry	Daylight	No	37	1

4.5. *Definition of Within-Scenario Parameters*

As noted earlier, many parameters needed for the detailed representation and simulation of any given driving scenario are not directly available from GES crash data. For example, information about vehicle kinematics such as specific speeds, yaw rates, lane positions, etc, just prior to a lane departure or road departure crash is not available in the database. Hence these data must be derived from other sources, specifically from naturalistic driving databases. Additionally, specific characteristics of the safety technologies are obtained more appropriately from objective testing. Collectively, these parameters extracted from the naturalistic databases and the objective testing were labeled as alpha (α) parameters. Table 4.11 shows a listing of the typical parameters that are represented in driving scenarios along with their type and sources.

Given the probabilistic nature of the crash occurrences, we definitely expect that a proportion – even a large proportion – of the outcomes from simulation will result in safe recovery and no resulting crash. In this case, safety benefits can only be assessed from the cases that do lead to a predicted crash, and the simulations resulting in "no crash" are not directly useful for benefits estimation. For this reason, sampling from the naturalistic behaviors may be inefficient. However, without the very detailed information on kinematics (and, more fundamentally, on driver awareness and responsiveness), there appears to be no reasonable alternative. Further, in Section 7, we will introduce an "alpha sampling scheme" that greatly reduces this inefficiency.

Another efficiency saving is that distributions of alpha parameters need be derived only for driving scenarios relevant to the safety technology, and not the entire universe of driving behavior. Therefore, the data mining from naturalistic databases was carried out conditional on the scenario conditions (β parameters) from Section 4.4; α parameters are extracted for each of the scenarios in Table 4.6, with filters applied for road type, roadway alignment, weather and road surface conditions and lighting. Driver fatigue level was not directly available in the naturalistic driving data, and although progress was made in developing such a parameter (see Section 10) the simple choice was made to exclude this factor from the filter conditions used.

Another key filter in the naturalistic data mining is to ensure that the distributions exclude cases where the dominant crash risk is not relevant to the technology, for example in hard braking events or in situations of high traffic density. The RDCW database includes a traffic density parameter that has been used to filter out high traffic situations. Additional filters for high traffic density are determined by excluding cases with a combination of repetitive braking and acceleration, inappropriate following distances, or vehicle speeds not commensurate with roadway types (e.g. too slow for freeway driving). Again, traffic situations involving turning at intersections were removed from the analysis. Applying these filters, distributions of any number of alpha parameters are obtained.

Table 4.11. Driving Scenario Parameters and Initial States

	Type	Source
External parameters		
surface conditions		
road class	β	GES (NASS GES crash data)
curvature	β	MICH (Michigan crash data)
lanes and widths	β	HPMS (Highway Performance Monitoring System)
lane markings	β	NAT (naturalistic driving data – RDCW)
grade	β	HPMS
friction/surface condition	β	GES
cross-slope	β	MICH
shoulder width and surface	β	MICH
off-road slope and friction	β	MICH
off-road obstacles	β	HPMS
environment		
weather/visibility	β	GES/CDS (NASS GES and CDS databases)
traffic condition	β	HPMS
location of other vehicles	α	NAT
motion of other vehicles	α	NAT
External initial states		
disturbances		
aerodynamic drag	α	VP (Vehicle Parameters: manufacturer data)
road induced yaw moment	[α]	NAT [assumed zero in this study]
Vehicle parameters		
type, size, load	α	VP
suspension and steering	α	VP
systems (ABS, ESC, …)	α	VP
safety system level	α	VP [ESC absent; with/without LDW]
vision limits (e.g. A pillar)	[α]	VP [no limits imposed in this study]
Vehicle initial states		
CG position	α	NAT
speed	α	NAT
yaw angle	α	NAT
yaw rate	α	NAT
lane position	α	NAT
lateral velocity	α	NAT
steer angle	α	NAT
throttle position	α	NAT
Driver parameters		
steering control gain	α	HMI (objective testing with naïve subjects)
throttle control gain	[α]	HMI/NAT [nominal – constant speed reference]
min/max preview times	α	NAT
steer/brake bias	[α]	HMI/NAT [nominal – bias always towards steer]
time delay parameters	α	HMI/NAT
fatigued driver	β	GES
Driver initial states		
task switching state	α	HMI
attention/delay states	α	HMI
lane boundary error state	α	HMI
steer torque disturbance state	α	NAT
ACAT system parameters		
LDW availability	α	TECH (objective testing of the technology)
LDW initial state	α	NAT
LDW delay/error	α	TECH

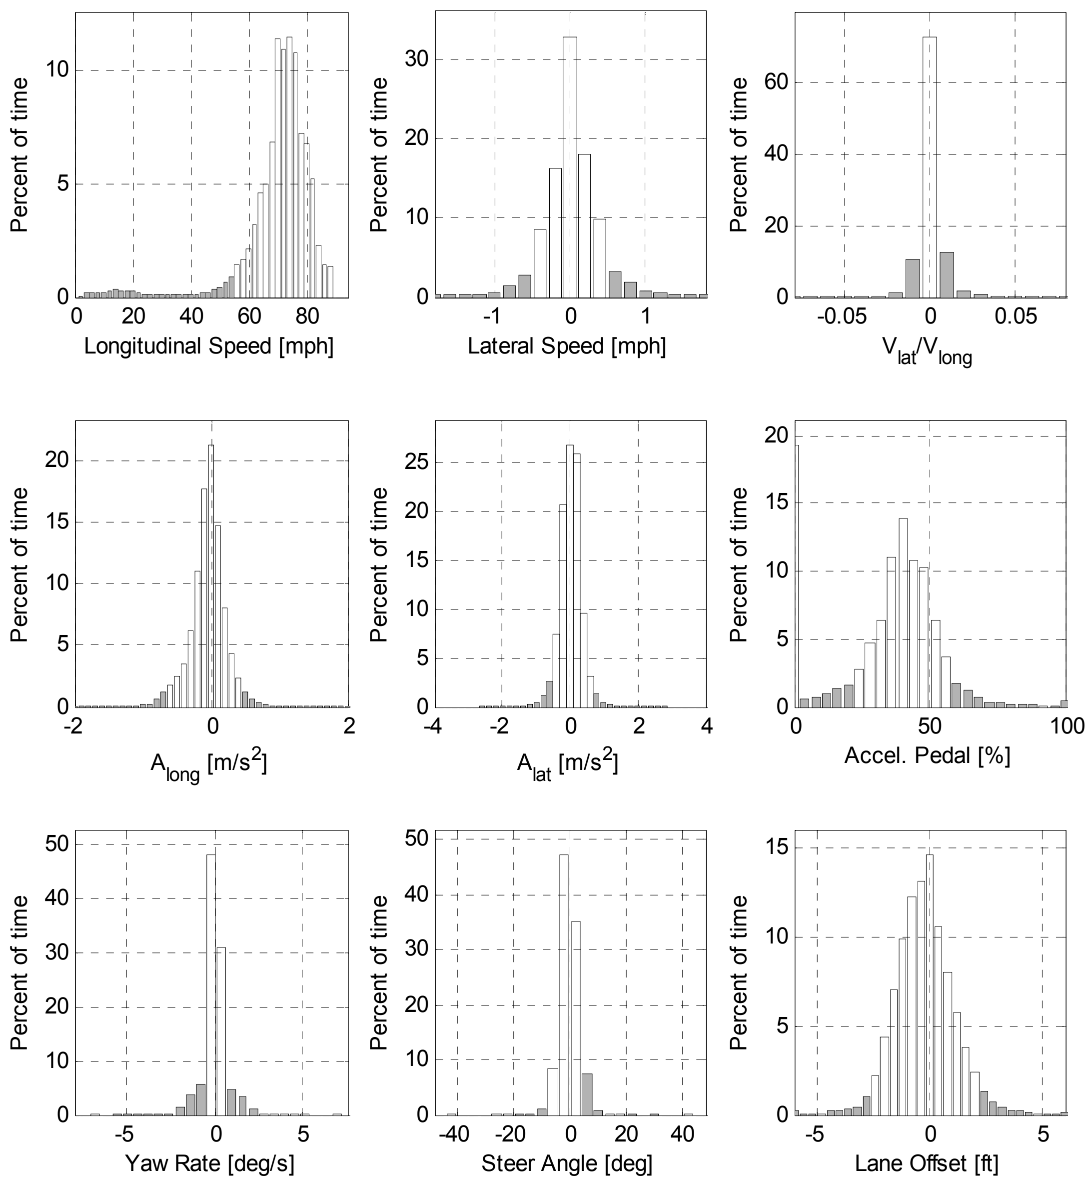

Figure 4.4. Distributions of Alpha Parameters for Driving Scenario 1

Figure 4.4 shows the distributions of some key alpha parameters, obtained from the RDCW FOT data, for the first driving scenario in Table 4.6. This scenario consists mainly of driving on limited access roads that are two or more lanes with divided medians (e.g. freeways & interstates), no adverse weather conditions, straight and dry roads, with the crashes occurring during daylight hours.

Naturally enough the (multivariate) distributions are not the same for different driving scenarios. For example, since scenarios 1 and 10 differ in roadway alignment (straight and curved) the distributions of the yaw rate and steering angles are noticeably different - those for scenario 10 display double peaks (for curves to the left and right) compared to the unimodal distributions of Figure 4.4.

To apply the information gained from filtering and sampling the naturalistic driving data requires a specific sampling plan. Although the distributions shown in Figure 4.4 are univariate, they arise from aggregation within a truly multivariate probability density function, since we naturally expect variables such as speed, yaw rate and lateral acceleration to be statistically dependent. When random samples are drawn from the multivariable α distribution in any major scenario, it is possible to avoid the statistical dependency by sampling specific events from the naturalistic driving database. This is again taken up in detail in Section 7.

It is not only α variables that need to be randomly sampled within the SIM analysis plan. Based on examples of Michigan crash locations from each scenario, it has been possible to extract geographical coordinates and then use HPMS data to define actual lane and shoulder widths. Similarly, the simple road curvature variable (straight/curved) used in the high level scenario definition is supplanted by sampling actual highway geometries. Thus random sampling is used to define within-scenario parametric values for simulation, not only for the α variables but also in creating necessary detail for the β parameters. Again see Section 7 for details.

5. Objective testing

Objective testing consists of functional testing of the vehicle and its safety technology (technical testing), as well as controlled tests involving driver interaction (HMI testing) with or without the safety technologies active. Testing performed of the ACAT technologies cover the blocks 8-10 from the SIM flowchart in Figure 2.1; driving simulator tests (block 8), vehicle tests on track (block 9) and vehicle tests on public roads (block 10).

In the approach chosen and developed by the VFU-team, it is important to point out that the data from the objective tests is not used directly in the safety benefits estimation. Rather it is used to establish parameter values for the operational capabilities and limitations of components in the computational Driver-Vehicle-Environment-Technology (DVET) simulation model.

Regardless of whether the testing is performed in a driving simulator, on a test track or on public roads, the objective tests can be said to come in two forms. One is in the form of detailed technical tests of the vehicle and its safety features. The other is objective tests designed to capture typical ranges of human performance where the driver is in the loop. This is illustrated further in Figure 5.1.

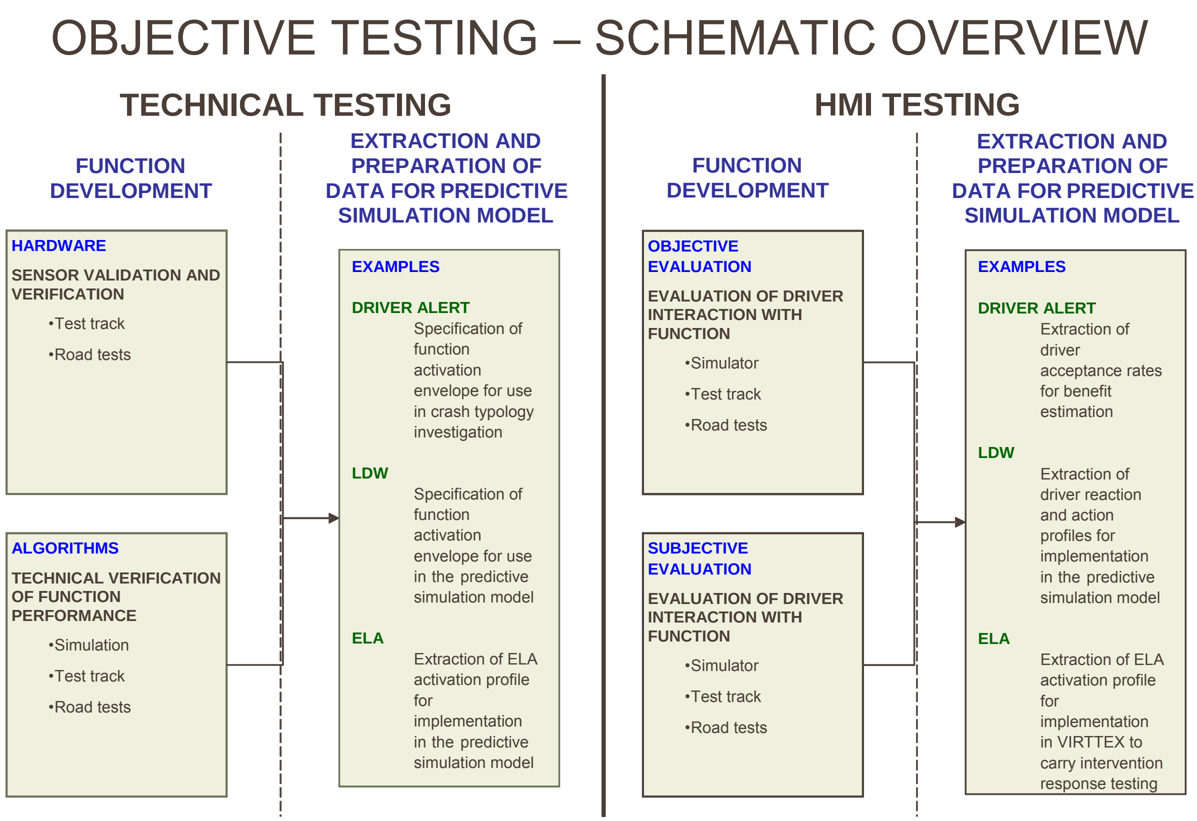

Figure 5.1. Objective testing of ACAT technologies

The technical tests are used to calibrate and validate performance of the predictive simulation model relative to the vehicle technologies, while the HMI tests are used in a similar way to calibrate the model for the driver performance. In the HMI tests, there is a great deal of variability in performance, so the model will attempt to capture a range of driving behaviors rather than just specific values in single recorded events.

In relation to the ACAT technologies, this means technology performance in the computational model can be adjusted to how the technology performs when fitted on a real vehicle driving on real roads under various conditions. In relation to the driver model, driver performance can be adjusted to how real drivers in the real world behave under various conditions. Objective testing is therefore important in the sense that without it, the ACAT technologies would operate under only idealized conditions which exist in the computational model (such as continuous perfect information on where the lane boundaries are, no precipitation, etc), and the driver model would not capture the behavior of real drivers.

Since the main focus of this report is on how the VFU model has been developed for Lane Departure Warning (LDW), the following section will focus on how the two types of objective tests have been carried out in relation to LDW and where the results have been incorporated into the model. Following that description, a shorter section on the objective testing performed with Driver Alert Control (DAC) and Emergency Lane Assist (ELA) will follow.

5.1. *Lane Departure Warning (LDW)*

Lane Departure Warning (LDW) works in a later phase in the crash sequence than DAC. In this phase the model is active and uses the system performance for the LDW to estimate the preliminary safety benefit.

5.1.1. LDW Driving Simulator Tests

The purpose of this testing was to assess drivers' responses to imminent lane departure events and determine which HMI gives the best situational performance and acceptance. Two main studies were carried out:

VIRTTEX Study 1

The Ford VIRTTEX LDW sleep deprived driver study (hereafter referred to as *Study 1*) was designed to measure a sleep deprived driver's reactions to, and acceptance of, different HMI solutions in lane departure situations. The study consisted of simulated interstate driving at 60-70mph (95-110kph). Driving conditions were night-time with moderate traffic volume, the simulated road had two lanes each way separated by a median, and the driver was instructed to stay in the right lane for the entire drive. The total duration of the drive was 3 hours, and drivers were sleep deprived for approximately 23 hours prior to starting their drive.

At random intervals, a “yaw deviation" was introduced by sending a steering command to the simulated vehicle while suppressing the effects of the command on the simulator motion platform.

The purpose of this maneuver was to make the simulated vehicle drift out of the lane in order to activate the LDW, without drivers noticing (i.e. feeling the vehicle movement), thus increasing the number of situations that would activate LDW.

VIRTTEX Study 2

The Ford VIRTTEX LDW distracted-driver study (hereafter referred to as *Study 2*) was designed to test distracted and alert driver reactions to, and acceptance of, different LDW HMIs. The study simulated interstate driving at 60-70mph (95-110kph). Driving conditions were daytime with moderate traffic volume, the simulated road had two lanes each way separated by a median, and the driver was instructed to stay in the right lane for the entire drive. To avoid any "mental preparation" for the generated surprise lane deviation, drivers were told an alternate reason for the study – it was described as "Lane-keeping Aid Study."

Drivers were instructed to perform a secondary task, so that when an LDW alert is introduced, its effectiveness and acceptability could be assessed. That secondary task was to read aloud a series of 6 numbers sequentially displayed on a screen low in the center console near the passenger seat. Each digit came on every 0.5 s with a display time of 0.3 s and a blank period of 0.2 s between displayed digits (so the total task time is approximately 3 seconds). The time between distraction events was uniformly distributed from 15-45 seconds. The total duration of the drive was 20 minutes, and each subject experienced 40 distraction events. During 16 of these, a "yaw deviation" (as described above) was introduced, to make the simulated vehicle drift out of the lane and activate the LDW without drivers noticing. These sixteen true positive events were then used in the following estimation and validation process.

From these studies, driver response data was taken and used to calibrate the performance of the driver model used in the simulations in two ways. One is in terms of driver reaction times to lane departures and lane departure warnings. The other is in terms of drivers' steering response when they have detected a lane departure. Data from alert, distracted and sleep deprived subjects were extracted and used for calibration, based on the assumption that reaction times and steering performance may be different for these groups.

Data on acceptance and usability was not used directly in the model, but rather to gain an understanding of drivers' willingness to comply with LDW warnings. If some drivers are not willing to respond to the warnings provided, the benefit of having the technology in the vehicle is reduced. A low willingness to comply would therefore affect the overall benefit analysis. In the testing however, drivers did to a large extent comply with the warnings, and acceptance was high. For a further discussion of this topic, see Section 9.

5.1.2. LDW Track Tests – Assessing True Positive Performance

The aim of the true positive performance testing was to quantify the distribution of lateral warning distance (Activation profile) of a Lane Departure Warning feature. Eight tests, chosen as a combination of longitudinal and lateral velocity and direction (right / left) were performed by a test driver on high speed test track. The results of the eight test scenarios is given as warning distance in relation to the longitudinal and lateral velocities measured for the specific test. Warning distance is the distance between the inner edge of the lane marker (the side facing the middle of the lane) and the outer edge of the lead tire. The number turns negative when outer edge of the lead tire is outside the inner edge of the lane marker. In all tests the numbers were positive. Data from this testing was used to calibrate the performance of the LDW implementation used in the simulations, to ensure that the warnings in the simulation are issued at the same time as warnings would be given on a real road in a real vehicle.

5.1.3. LDW Track Tests – Assessing HMI Acceptance and Usability

Acceptance and usability testing of display, warning sound and on/off buttons and activation/deactivation speed has been performed on a test track. A clinic with 23 Volvo Cars employees was run to investigate differences in driver's acceptance and attribute rating of two different modality implementations of the LDW function. The two modalities were a haptic belt (four rapid pulls on a motorized seat belt) and a monaural audio warning signal.

The test simulated a situation with a potential for unintended lane drift. The task was to drive the car within a lane on the test field marked by lane markings at a speed of approximately 50 km/h, and after a while slowly let the car slide over the lane marking to activate a lane departure warning from the LDW system.

Half of the test group got a haptic warning using the seat belt, while the other half got an auditory warning. All test participants experienced both warnings. Results were given with regard to the warning implementations for preference and acceptance.

As with the simulator tests of acceptance and usability, this data on acceptance was not used directly in the model, but rather to gain an understanding of drivers' willingness to comply with LDW warnings in relation to the safety benefit estimate for the technology. For a further discussion of this topic, see Section 9.

5.1.4. LDW Public Road Tests - Assessing False Positive Performance and Availability

The aim of false positive performance testing was to quantify the rate of false warnings of the LDW feature. System availability testing was performed to calculate the percentage of time during the test drive that the LDW system was active and operable relative to the time driven within the operational scope of the function.

Selected test drivers drove the test vehicles on public roads in several different countries to determine the false alarm rate and availability of LDW. The reason for including driving in several countries is that Volvo Cars is a global company, and the technologies it deploys should function in all markets, despite possible variations in road designs and standards of maintenance. More than 2000 km of data where LDW was operational (not disabled due to too low speed) was collected. All the LDW events (warnings) were classified as either true positive or false positives.

The false positive performance (number of false warnings divided by the time driven within the operational scope of the function[7]) was calculated for all the 18 combinations of road type, weather and lightning conditions classified as follows:

Road Type:
- interstate, freeways and Expressways
- rural arterials
- urban arterials

Light condition:
- day
- night

Weather condition
- clear
- light rain
- light snow

The availability was calculated as the time LDW was available divided by the time driven within the operational scope of the function. Details are provided in Section 9.4.

The data on false warnings was not used directly in the model, but rather to gain an understanding of drivers' willingness to comply with LDW warnings. If drivers receive a large number of false warnings, their willingness to respond to the warnings may be lowered, and the benefit of having the technology in the vehicle would thus be reduced. In this testing, the number of false alarms was within the number deemed acceptable. For a further discussion of this topic, see Section 9.

Data on availability was used in the benefits assessment. Basically, for each simulated scenario, the number of successful lane departure recoveries classified as being the result of a lane departure warning can be reduced in proportion to the estimated availability of LDW in that scenario. For further discussion of this topic, see Section 9.

5.2. *Driver Alert Control (DAC)*

Driver Alert Control (DAC) is intended to help the driver in the pre-conflict phase of a crash sequence. The results of the technical and HMI tests are used to develop an estimate of the DAC benefit. The testing results indicate that drivers who receive a DAC warning will be motivated to do something about their diminshed vigilance, for example by taking a rest break and sleep for a few minutes. If this is the case, there is no need to extract data from the DAC testing to the simulations, because drivers will take action before the driving situation enters a conflict phase. This means that the DAC technology basically can be treated as a filter in the SIM tool; by having

[7] LDW is turned to operational mode when speed exceeds 65km/h and is turned to non operational mode when speed drops below 60km/h

the technology in the vehicle the number of drivers experiencing an unintended lane departure due to diminished vigilance, for example due to sleep deprivation, will be reduced to a substantial degree. Analysis of the DAC system is discussed in more detail in Appendix D.

5.2.1. DAC Driving Simulator Tests

Driving simulator testing is not necessary for tuning DAC. HMI and compliance-related aspects for DAC require methods where long durations and a large spectrum of drivers are covered.

5.2.2. DAC Track Tests - Assessing True Positive Performance

The aim of the true positive performance testing was to estimate the probability of the Driver Alert Control warning a sleep deprived driver before lane departure. Tests were run at the Michigan proving ground in the US in 2005, and then at the Hällered proving ground in Sweden in 2006, using seven different vehicles and 77 test subjects. All the tests took place between midnight and 05:00 AM and the test subjects were sleep deprived before the tests. All the 77 test participants' vehicles ended up at some point 50% (or more) out of the lane due to falling asleep and 76 of the 77 drivers did receive a DAC warning at least 1 minute before the lane departure.

5.2.3. DAC Track Tests - Assessing HMI Acceptance and Usability

To determine driver's evaluation of DAC warnings, a test track HMI clinic with sleep deprived drivers was performed. In this clinic all but one of the drivers received a DAC warning during the drive. When asked to give feedback on how they perceived this warning (questionnaire study) a large majority of the drivers felt that the feedback was useful. They also reported that the feedback influenced how they drove and that it made them more awake. Some of the drivers were surprised by the feedback from DAC. The drivers did not perceive the feedback as annoying or frightening.

This data on Acceptance and Usability was not used directly in the model, but rather to gain an understanding of drivers' willingness to comply with DAC warnings in relation to the safety benefit estimate for the technology.

5.2.4. DAC Public Road Tests - Assessing False Positive Performance

The aim of false positive performance testing was to quantify the system detection of degraded lane keeping when the driver is, in fact, not distracted nor has diminished vigilance. This fact does not directly influence the outcome of the SIM-tool, nonetheless it is a part of the technical testing, since excessive false positives may influence a driver's confidence in the system and potentially result in the driver deciding to turn the system off.

Tests were carried out on pre-defined routes on public roads in three countries, using four vehicles and 105 test subjects. The number of false warnings occurring during the tests, within the total of 282 hours of driving, forms the basis for the calculation of the false warning rate (FP, i.e. false warnings per hour).

The data on false warnings was not used directly in the model, but rather to gain an understanding of drivers' willingness to comply with DAC warnings. If drivers receive a large number of false warnings, their willingness to respond to the warnings may be lowered, and the benefit of having the technology in the vehicle would thus be reduced. In this testing, the number of false alarms was below the number deemed acceptable.

5.2.5. DAC Public Road Tests - Assessing Availability

System availability testing is performed to define the percentage of time during the test drive that the system is active and operable relative to the total drive time. With respect to functionality the DAC is turned to operational mode when speed exceeds 65 km/h and is turned to non-operational mode when speed drops below 60 km/h.

More than 43 hours of driving data was recorded, with a total of 24 hours at speeds above 65 km/h by eight test drivers on public roads. The DAC availability performance metric calculated as the percentage of time that DAC is active and operable relative to the total drive time was calculated for the same combination of road types, weather and light conditions as shown above for LDW.

However, since DAC benefits are not estimated in this project, no detailed results are presented. Data on availability was used in the benefits assessment. Basically, for each simulated scenario, the applicability of DAC to the initial scenario was assessed (i.e. is fatigue (as indicated by GES) a factor which influences the occurrence of the scenario?). In principle, for each applicable scenario, the estimated effectiveness of DAC can be reduced in proportion to the estimated availability of DAC for that scenario.

5.3. *Emergency Lane Assist (ELA)*

If the host vehicle is drifting into the adjacent lane, and the ELA estimates that it is on a collision course with an obstacle in that lane, the system provides an active steering intervention to steer the host vehicle back into the original travel lane under certain conditions (e.g., a clear path in the original travel lane).

Because ELA is a technology in an early stage of development, testing has so far been limited to tentative studies of its true positive performance on a test track, and driving simulator studies of how its performance capability may be reduced as a result of driver interference.

5.3.1. ELA Driving Simulator Tests – Assessing the effects of driver interference during intervention

Driving simulator testing was carried out to assess how the function performs if driver interferes with the steering input from the technology during an intervention. The technology was implemented in Ford’s driving simulator VIRTTEX. Testing with 5 non-naïve subjects was carried out. These were instructed to induce mild, moderate and heavy interference (through grip on the steering wheel) during intervention.

5.3.2. ELA Track Tests - Assessing True Positive Performance

Preliminary testing of the technology's ability to return the host vehicle into the original lane when facing an imminent collision threat was carried out at Volvo's Hällered test track. The testing was performed by non-naïve subjects, and used a stationary balloon car as the potential collision object.

5.3.3. Extraction/preparation of ELA data for the simulations

In principle, if drivers who receive an ELA intervention do not interfere with the steering input from the technology, the safety benefit will be determined by two things; its availability in relevant driving scenarios, and how quickly it can return the host vehicle to the original lane in relation to the time available before a collision with objects/vehicles in the adjacent lane occurs. This means that the ELA technology does not have to be included in the simulations. Rather it can be treated as a filter in post processing of simulation results. By assessing how many of the trajectories which come out of the simulation would require ELA intervention, and which of these are covered by ELA given a certain level of performance and availability, the potential reduction in number of drivers risking a collision with an object or vehicle in the adjacent lane as the result of an unintended lane departure can be estimated. If drivers do interfere with ELA, it is necessary to first gain an understanding of how well the technology performs under those conditions (i.e. to which extent the performance capability of the technology is reduced due to driver interference) before a relevant benefit estimate can be made, but the general principle still applies. The ELA system is discussed in more detail in Appendix D.

6. Model Development & Validation

This section describes the model development and validation and addresses blocks 12-14 of the SIM flowchart in Figure 2.1. The preceding blocks (1-11) have focused on the formulation of the underlying methodology, the definition of the types of crashes that could potentially be prevented by the active safety technology and the modeling of the active safety technology itself, as well as data preparation via existing databases and targeted objective testing; the following blocks, from 15 onwards, deal with the application of the model and supporting data, starting with the setting up of batch simulations in block 15 through to the analysis of output data in the remaining blocks. Since the model is pivotal in the SIM benefits analysis, it is described in some detail in this section.

We consider here the definition, implementation and validation of the computational model used in this study – the complete DVET model – via the four components of driver, vehicle, environment and safety technology.

To be able to control the vehicle and to respond appropriately during the identified driving scenarios (DS's), the driver model has to have certain capabilities. It must be 1) able to keep the lane (i.e. find out where the lane is going and steer accordingly), but 2) also be capable of losing track of where the lane is (i.e. of being distracted from the lane keeping task), so lane departures may occur. Also, it should 3) be able to return to the lane of travel once a lane departure has occurred and the driver model becomes aware of this fact.

The vehicle model should be able to respond realistically to driver control inputs, including for example loss of traction/skidding when steering is excessive. The safety technology model must issue warnings at the same time and in the same place (here, lateral position) as a real world vehicle implementation of it would. Finally, the environment model should mimic the properties of the real world roads where the crashes occur, in terms of for example road curvature and lane width.

Development and validation of the first three components (driver model, vehicle model, safety technology model) are described here in Section 6, while the environment model is covered in section 7.1.

Figure 6.1 shows the model architecture, with connections marked "R" for road data, "V" for vehicle motion data and "L" for the LDW audible and/or haptic alert. The input conditions for the driver model are dynamic variables and their initial states at the start of each simulation are set by the alpha and beta parameters. In the following sections we describe in turn the vehicle model, the LDW system model and the driver model. Note that the vehicle motion outputs are split into the position, velocity and acceleration information, and it is only the position and velocity data that are used by the LDW system model.

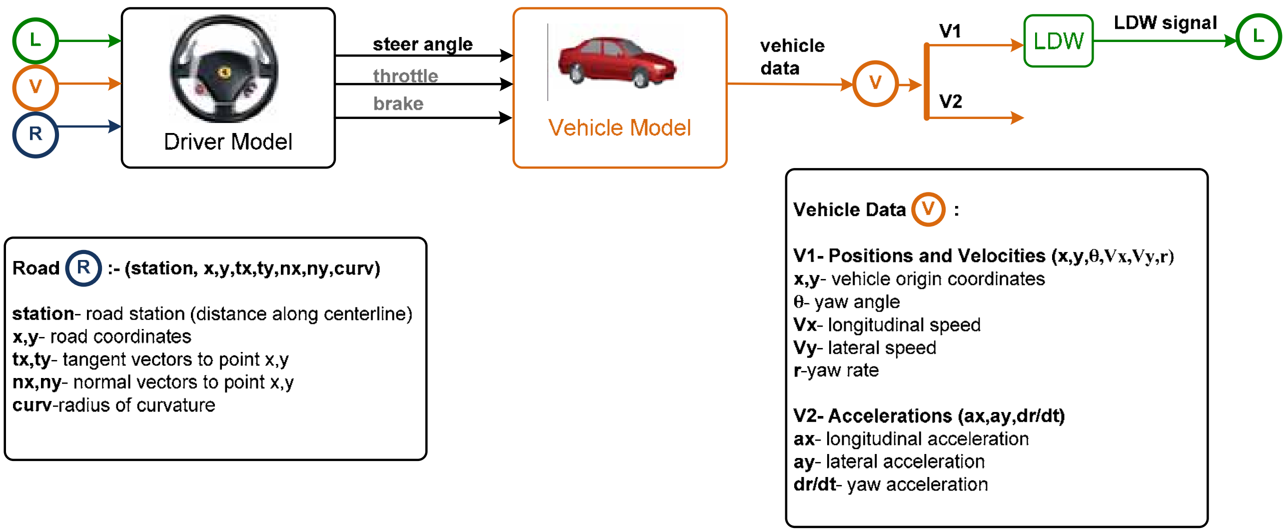

Figure 6.1. Architecture of the Coupled Vehicle-Highway-Driver Model

Model validations were performed using physical (track-based) measurements of the vehicle and LDW system for comparison purposes. Further test data was obtained from the Ford Motor Company VIRTTEX (Virtual Test Track Experience) driving simulator, and used for both vehicle and driver model development and validation.

Sections 6.1-6.2 describe the structure modeling of the vehicle and the LDW safety system respectively, including parameter estimation and key validation results. Section 6.3 details the structure of the driver model, while Section 6.4 describes how the key driver model parameters are estimated.

6.1. *Vehicle Model*

As shown in Figure 6.2, the vehicle model represents the response of a real vehicle when subject to given driver inputs and road conditions. Although it is primarily the steering input that is used by the driver in the model, throttle and brake are also potentially used; throttle in particular is used to maintain desired speed, and in any case the preference has been to keep the model simple, but at the same time general. The effect of the environment is limited to the friction and slope of the road, and in the batch simulations only flat roads are used. (As noted below some cross slope was introduced to assist model validation). Friction does vary between the paved road surface and the off-highway regions, and this will be considered further in Section 7 when we describe the connection between driving scenarios and the simulation model in preparation for conducting batch simulations. The unpaved (off-road) region is also presumed flat, and this will certainly reduce fidelity of simulation in events that include road departure.

The vehicle selected for the simulation was a mid-sized sedan. It was not considered feasible or worthwhile to try to model and simulate a wide range of light vehicles, so therefore the size and dynamic performance of the mid-sized sedan was assumed to be representative of the population of light vehicles considered in the analysis (see Table 4.1). The vehicle model was required to have sufficient fidelity to reproduce representative vehicle response for a given set of driving inputs, and show acceptable sensitivity to environmental changes such as road friction. The model needed to have a flexible representation in software to connect modules for the driver, environment and safety technology functions. A further requirement is for high computational speed so that large numbers of batch simulations may be run.

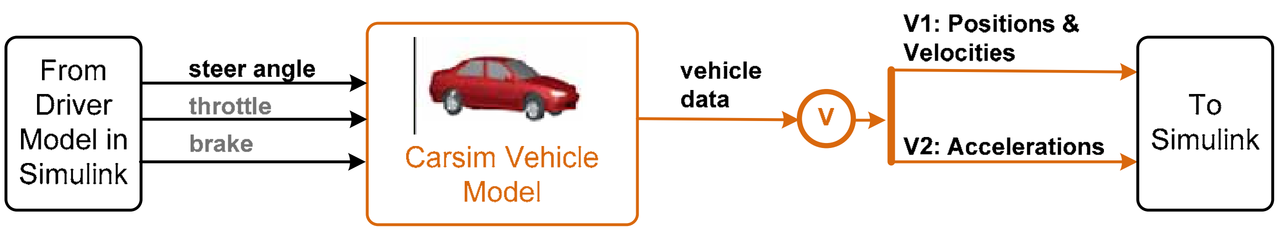

Figure 6.2. Architecture of the Vehicle Model

The desired vehicle model was implemented using Mechanical Simulation Corporation's CarSim software for modeling vehicle dynamics. CarSim was chosen because it has been built on decades of research in characterizing vehicles and reproducing their behavior with mathematical models; it is widely used in the automotive industry, can be coupled very easily to other sub-models using Mathworks' SIMULINK®, and is known to run efficiently. It is based on standard multi-body dynamics principles (e.g. Sayers and Riley 1996) and the software includes well validated parameter sets. CarSim's vehicle model parameters are also very adaptable (e.g. wheelbase, sprung mass and tire characteristics are selectable should vehicle characteristics be changed in the future) and the user can define and modify a large number of vehicle components, as seen, for example, in Figure 6.3 (a) and (b).

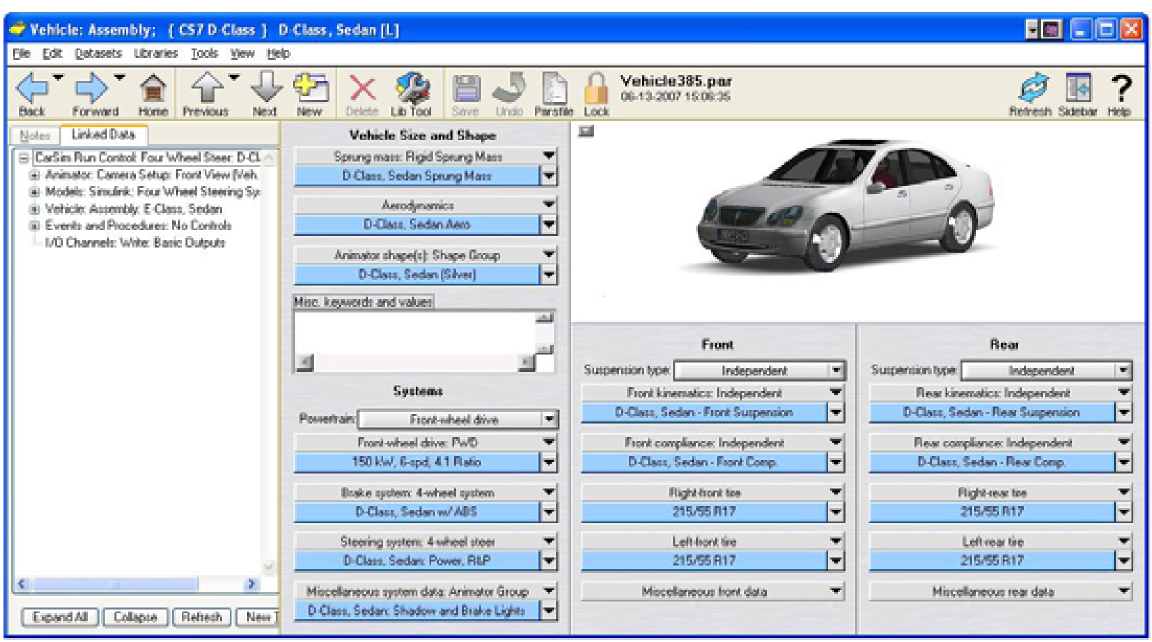

Figure 6.3(a). CarSim Screenshot: Vehicle Assembly

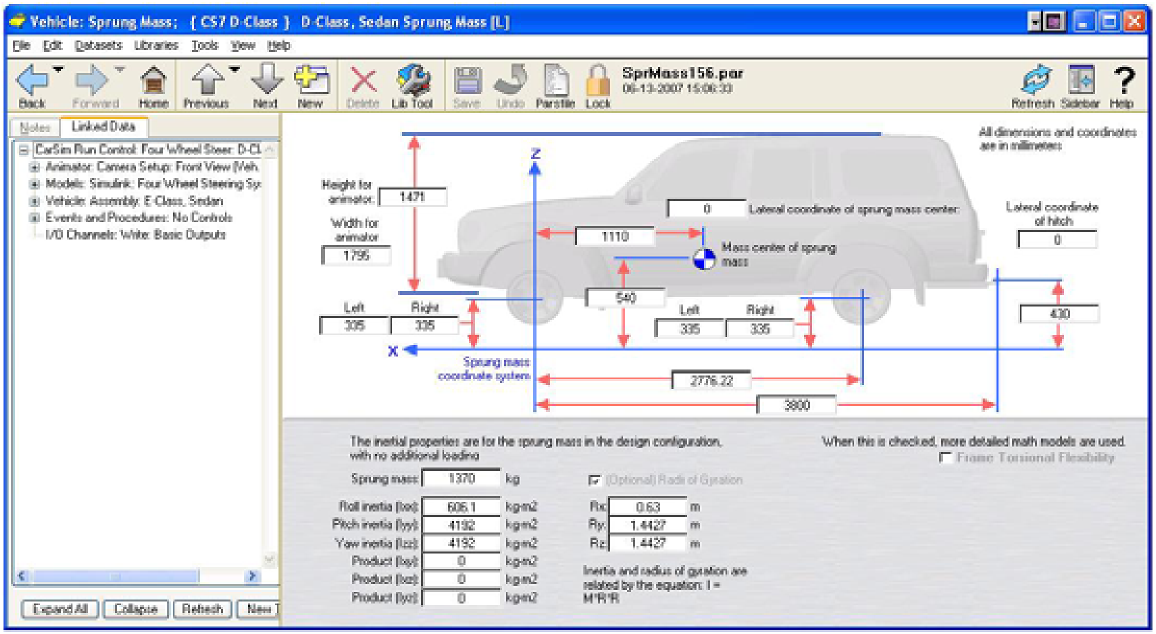

Figure 6.3(b). CarSim Screenshot: Vehicle Sprung Mass Parameters.

As shown in Figure 6.4, the vehicle model is embedded in an associated SIMULINK model that deals with many aspects of the inputs and outputs. The CarSim vehicle model is implemented in Simulink as a standardized *S-function* sub-system block. The simulation of the whole system takes place in a fully synchronized manner, with Simulink and CarSim exchanging data throughout the simulation.

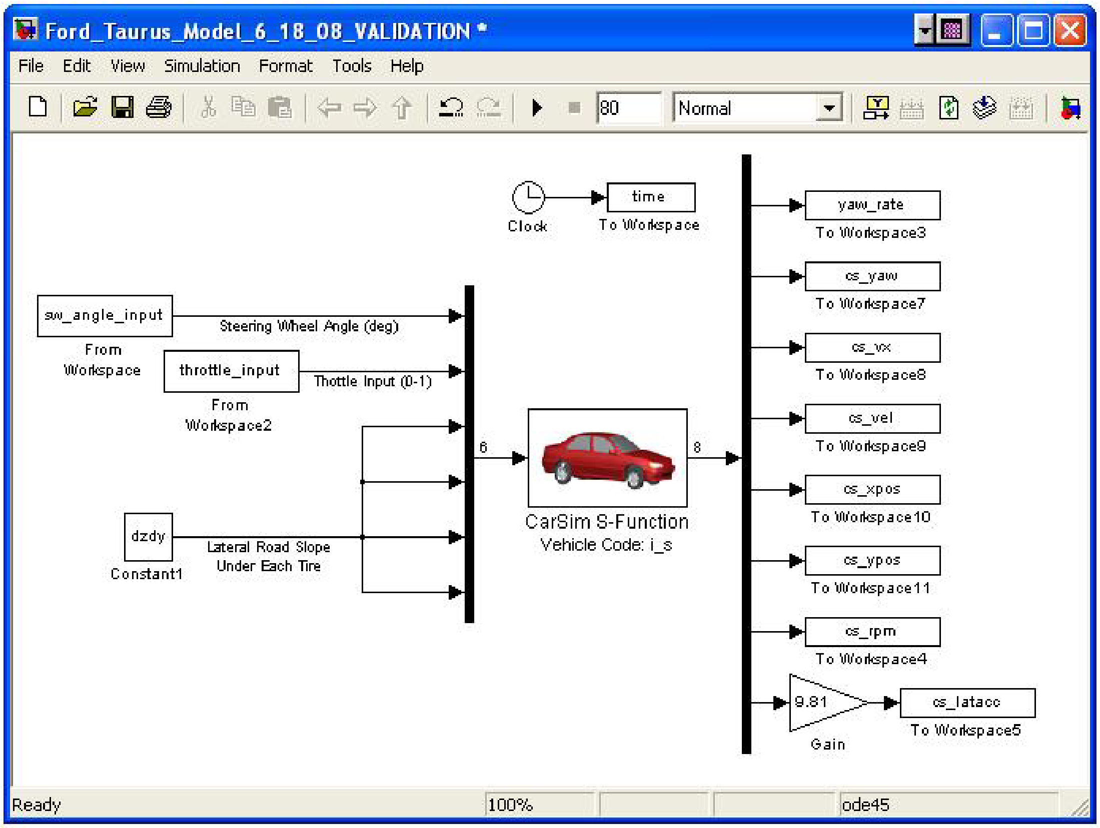

Figure 6.4. Simulink Diagram Used During CarSim Vehicle Model Validation

CarSim is an established multi-body vehicle simulation code, and has been validated against vehicle data many times in the past; thus any generic and representative set of vehicle parameters would probably be adequate for the present study. However, since it is important to compare test data and simulation results for cases when the LDW system and the driver are interacting with the vehicle, it was considered worthwhile to select vehicle parameters that are most representative of the test vehicle used. Since the greater degree of parameter estimation and validation took place using the VIRTTEX tests, the focus was placed on establishing a link with the vehicle model used in the driving simulator, a 2000 Ford Taurus. Though not reported here, we note that the Taurus model has itself been extensively validated by Ford based on objective test track data.

In the source VIRTTEX data, relevant time histories of many vehicle inputs and responses were available, and hence it was possible to compute the response of both the CarSim model for the same driver inputs and compare it to that of VIRTTEX. For the validation, time histories of the steering wheel angle and throttle position were chosen as inputs and the outputs included the yaw angle, velocity, and the global Y-coordinate. The simulation road of the VIRTTEX simulation included a crown, and so this was also implemented in the CarSim vehicle model.

Multiple vehicle model parameter lists for the VIRTTEX simulator were available, which allowed many of the CarSim variables to be directly set to match those of the VIRTTEX model. Parameters that could not be exactly matched (e.g. tires), were set to those of an internal CarSim model (D-

Class Sedan) representing a typical mid-sized sedan, and then further tuned to match the VIRTTEX results.

At the start of each VIRTTEX test the vehicle was at rest and the driver was asked to accelerate the vehicle to roughly 65 mph. This time period of relatively high longitudinal acceleration was used to tune the powertrain parameters of the CarSim vehicle model; from the velocity and engine speed outputs the engine torque curves were set, together with the transmission gear ratios and the transmission shift map required by the CarSim model. Figure 6.5 shows the velocity responses of the tuned CarSim model and the VIRTTEX system, and these are closely matched for the same throttle input. Figure 6.6 shows the corresponding throttle input. In these validations the performance of the driver is largely irrelevant, since the comparison is simply between how the two vehicle models respond, given a common set of inputs.

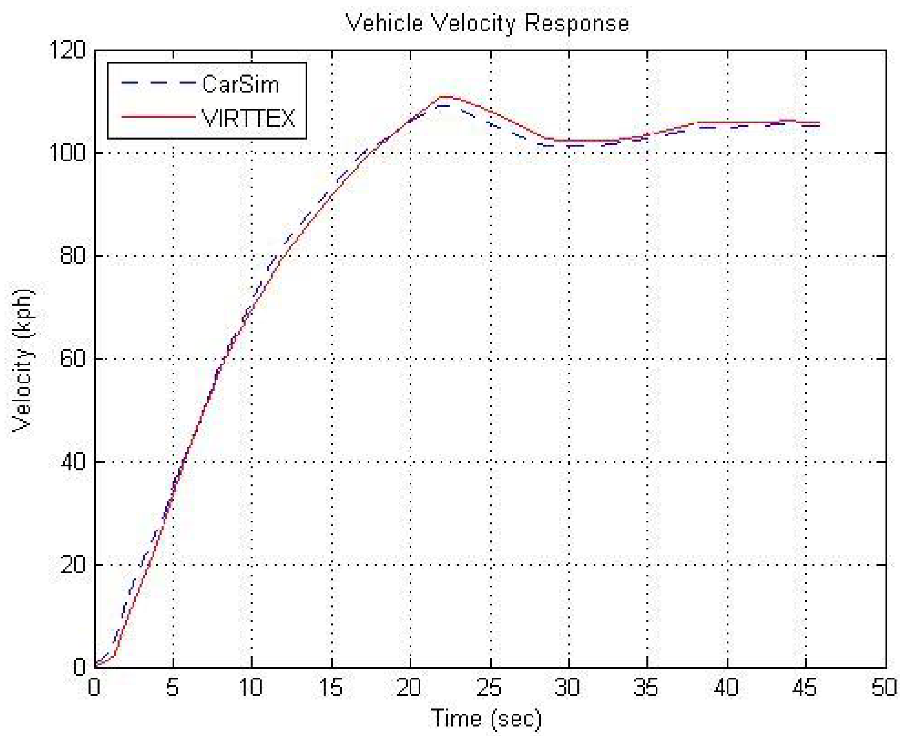

Figure 6.5. Vehicle Model Velocity Response to the Same Throttle Input

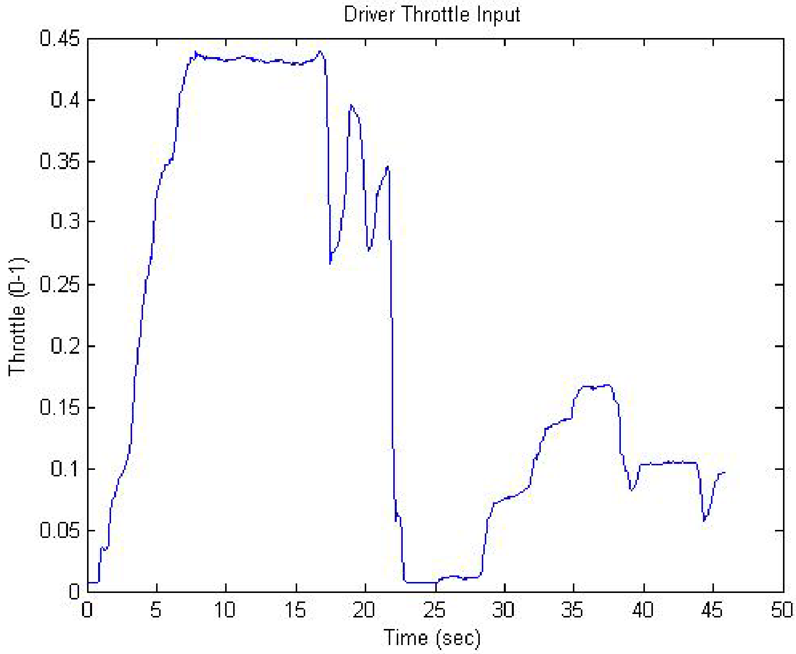

Figure 6.6. Throttle Input From the Driver During the Period of High Longitudinal Acceleration

The virtual road used in the VIRTTEX simulator tests is primarily straight but there are a few gentle curves which can be distinguished when studying the vehicle yaw angle time history. In order to validate the lateral dynamics of the CarSim vehicle model it was beneficial to study the vehicle responses around periods of transient lateral accelerations. An 80 second period of the VIRTTEX simulation data during which the road makes a gentle left followed by a right turn was chosen for this validation. The driver inputs during this time period are given in Figure 6.7.

The steering system, suspension, and tires of the CarSim model are of high importance when studying the lateral dynamics of the vehicle. Many parameters of the steering system were supplied but there was again limited information concerning the suspension stiffness, suspension damping, and also the tire model. For validation, the relationship between the displacement of the rack and the rotation of the front wheels and also the power steering boost curve were the parameters tuned to ensure the vehicle responses matched. The variables of interest during this validation study were the velocity, yaw angle, and the lateral distance (Y coordinate) of the vehicle relative to the road centerline. Responses of the CarSim model and VIRTTEX are shown in Figures 6.8 - 6.10. It is apparent that the vehicle model responses for these three variables are very similar and at this point we have high confidence that the simulation model will accurately follow the response of VIRTTEX, at least in the range of longitudinal and lateral accelerations used in the tests.

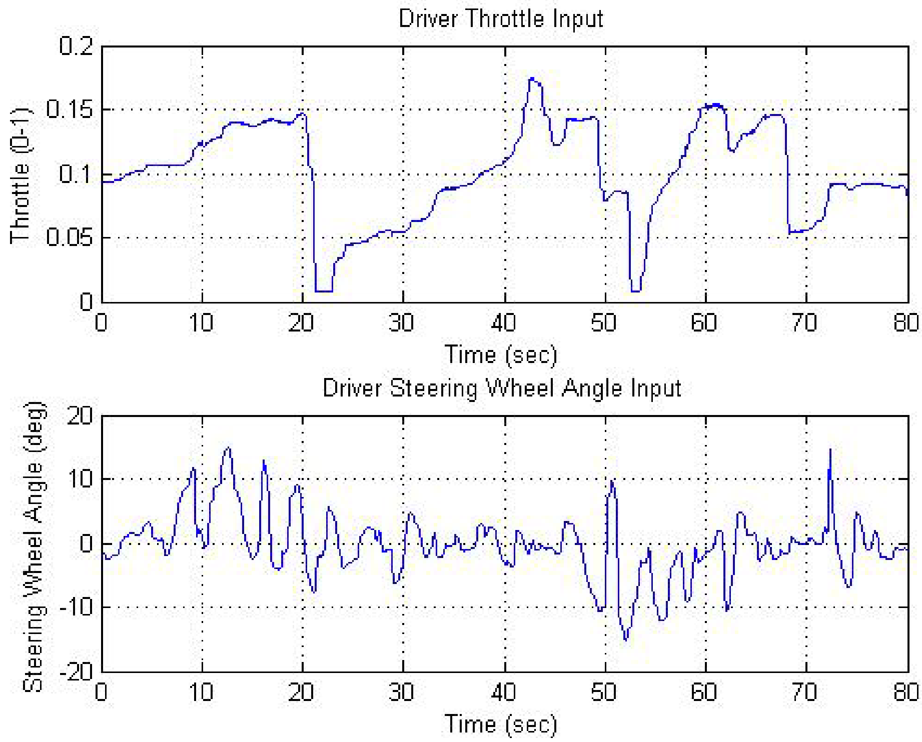

Figure 6.7. Driver Inputs During a Period of Transient Lateral Acceleration.

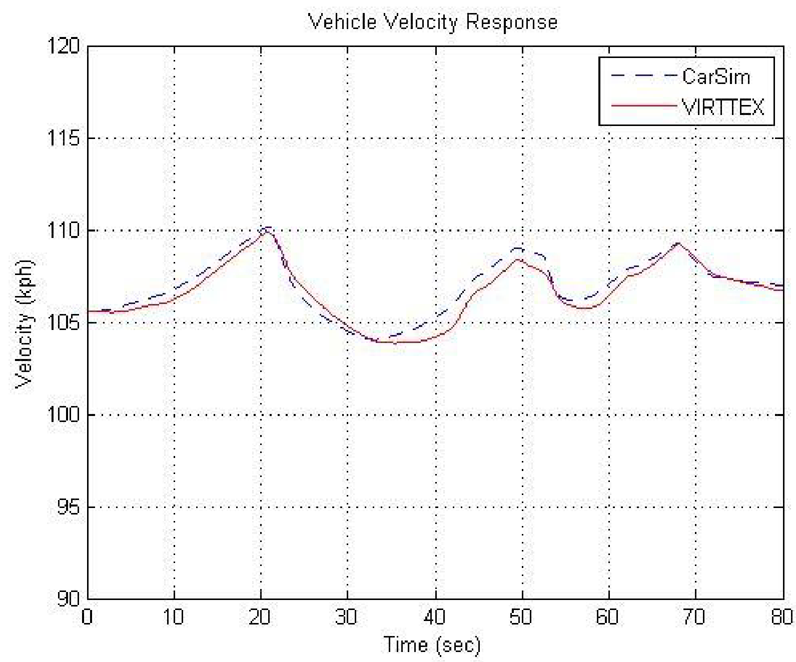

Figure 6.8. Vehicle Model Velocity Responses to Identical Driver inputs

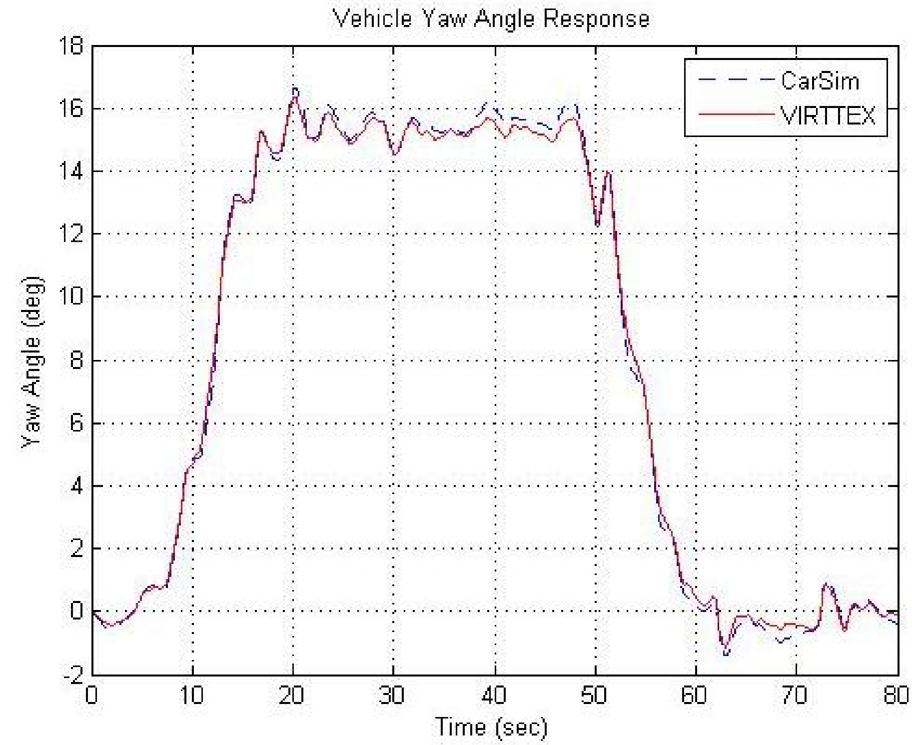

Figure 6.9. Vehicle Model Yaw Angle Responses to Identical Driver Inputs

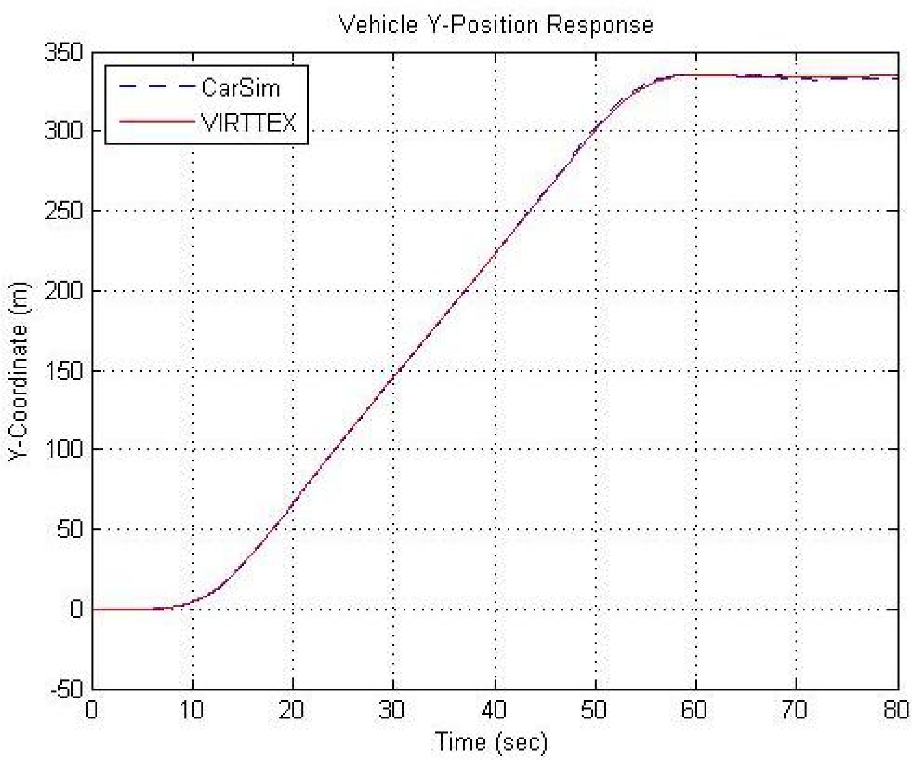

Figure 6.10. Vehicle Model Y-coordinate Responses to Identical Driver Inputs.

6.2. ***LDW System Model***

A generic model of Volvo Cars' Lane Departure Warning (LDW) system was developed for implementation in the SIM. LDW uses a camera mounted between the windscreen and the rear-view mirror to monitor the vehicle's position between the road lane markings. Figure 6.1 shows how the LDW sub-system is connected to the vehicle and driver model. The road information and vehicle position and velocities are used to represent the view seen by the camera and in particular to determine the distance between the front tires and the lane markers. As we now describe, this information is used together with trigger conditions to issue an alert to the driver.

Figure 6.11 shows a simple schematic of the system concept. The system alerts the driver with a warning sound if the car crosses a lane marking, provided there is no clear indication that the lane excursion was intentional. In the driving scenarios considered, the driver model never has any intention to leave the lane, and driving conditions are not used that would deliberately suppress the LDW signal (turn signal, braking or active steering effort); thus in all driving scenarios where it is switched on, the system will issue a warning signal once the trigger conditions are met. Note that LDW will not take any automatic action to prevent a potential lane departure, and responsibility for the safe operation of the vehicle remains with the driver.

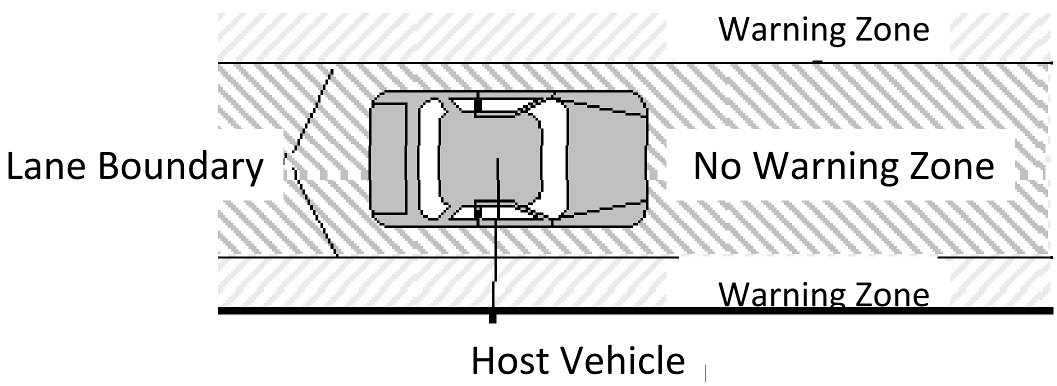

Figure 6.11. LDW Warning Zones

The actual LDW system has two sensitivity modes (high and low) which are driver-selectable. In the low sensitivity setting the warning will occur when the outside edge of the front tire contacts the outside edge of the lane marking. For the high setting, the warning is designed to occur close to where the outside edge of the front tire contacts the inside edge of the lane marking. In this study we implemented a version of the LDW with the high sensitivity. The algorithm controlling the warning is a function of the lane position along with the vehicle's lateral velocity with respect to the lane markings. To model this is relatively simple, as shown in Figures 6.12 and 6.13, which show the LDW architecture and the realization of the LDW functionality in Simulink. The inputs are the distances of the outside of the front tires from the inside of the lane marker (in each case a positive value means the tire is inside the lane boundary) obtained in a prior calculation via lane geometry, vehicle position and wheelbase. If either of these is less than a tolerance distance (called *ldw_tol* in Figure 6.13) a lane excursion is flagged; if the tolerance is zero, an LDW alert is generated as the tire touches the stripe. The model also includes a small processing delay (*Transport Delay*) and a final switch block disables the resulting output if the vehicle speed is below a critical threshold. The combination of position-based trigger and processing delay means that the location of the front wheels relative to the lane boundary when the audible alert is sounded does depend slightly on the lateral velocity of the vehicle relative to the lane.

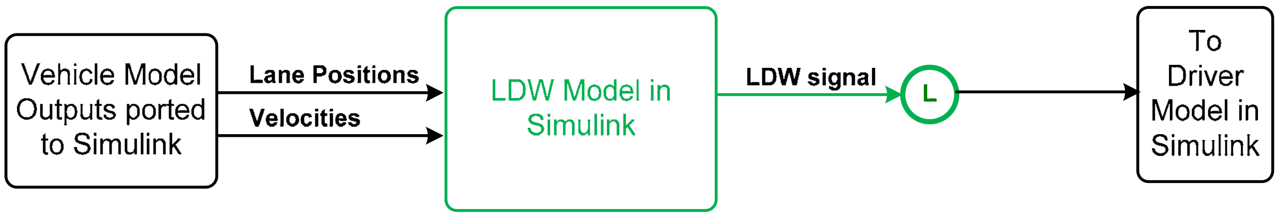

Figure 6.12. Architecture of the LDW Model

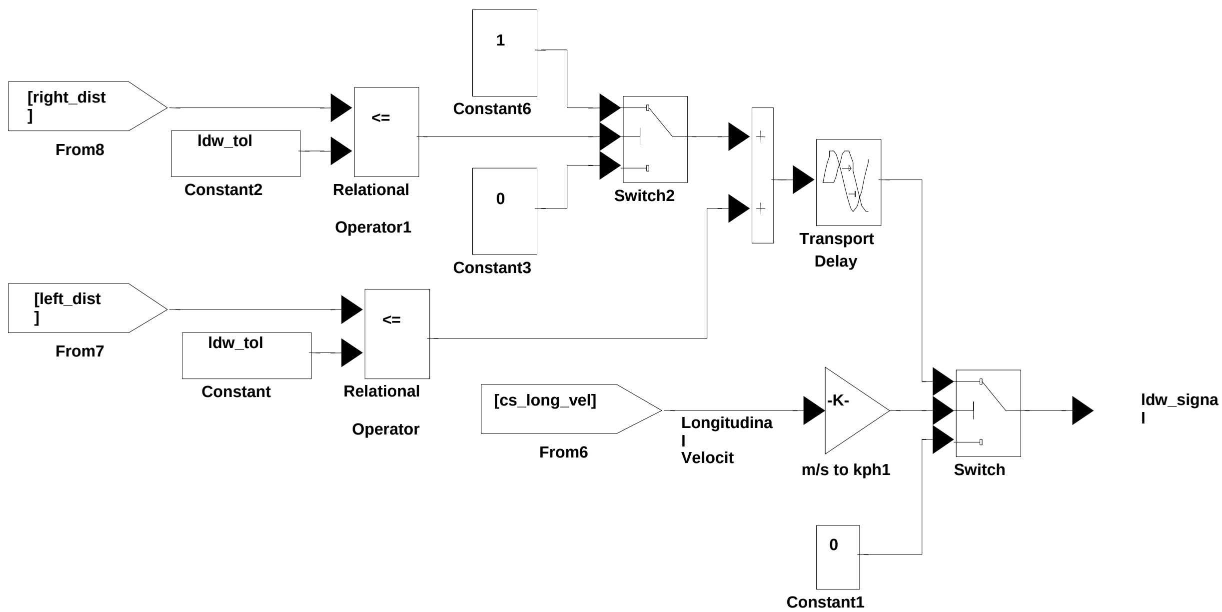

Figure 6.13. LDW System Sub-Model in Simulink

The exact control algorithm of the LDW system involves sensitive proprietary information, and therefore a generic model was developed based on measurements conducted in VIRTTEX. The VIRTTEX test study included a time history of the LDW warning signals, and this was used for detailed model calibration. The warning signal is coded with a value of 0 when there was no warning, a value of 1 when the LDW system indicated a deviation to the left, or a value of 2 when the system indicated a deviation to the right. From time histories of the VIRTTEX test data (LDW proprietary), LDW warning events were compared to the model, it was possible to estimate best fit values for tolerance distance (74mm) and transport delay (5ms). Table 6.1 shows summary timing errors if the transport delay is not used. With these fitted parameters included in the model, it was found that out of a total of 53 true positive LDW activation signals in VIRTTEX, the developed model reproduced the 53 actual warnings and produced only one additional activation (when the actual system was not activated). Also, for all of the warnings recorded, the difference between the warning activation times of the two signals was always less than 0.02 seconds. It was concluded that the simplified SIMULINK system model was sufficiently representative of the (VIRTTEX implemented) LDW system to be used with confidence in the simulation study. Note that no specific details of lane marker detectability or estimation confidence was included in the simulation model.

Table 6.1. Example of LDW model validation results (tolerance = 76mm, transport lag = 0)

Subject	Total Number of Warnings	Average Time Difference Between Warnings
1	17	0.001 sec. (SIM Model Reacted Sooner)
2	15	0.008 sec. (SIM Model Reacted Sooner)
3	8	0.006 sec. (SIM Model Reacted Sooner)
4	13	0.005 sec. (SIM Model Reacted Sooner)

The simulated LDW system was integrated with the CarSim vehicle model. Given the vehicle's center of mass coordinates, wheel track, tire width, yaw angle, and the distance from the front axle to the center of mass, the coordinates of the outside edge of each front tire are calculated relative to the lane boundary. If either of these distances is less than the previously calculated tolerance (74 mm) then the LDW signal is set to 1 (left deviation) or 2 (right deviation), this signal being subject to the 5ms time delay ahead of the warning being sounded.

A simple simulation was conducted to verify the integrated LDW/CarSim model performed as planned. This simulation was conducted in a similar fashion to those in the VIRTTEX simulator except that a driver model was used instead of actual human subjects. The vehicle was given an initial velocity of 100 km/h on a long straight segment of road. A directional (rightward or leftward) yaw angle error of one degree was imposed to produce simple lane departures. The simulations ran for eight seconds and an example LDW signal output is shown in Figure 6.14. Lane deviations to the right and left were detected during this time period. The path of the vehicle can be seen in Figure 6.15 while the vehicle location during the two times of LDW activation can be seen in Figures 6.16 and 6.17. In this visualization, the two white strips (not shown to scale) symbolize the lane markings and the system was designed to activate when the outside of the tire was 74 mm away from the inside edge of the lane marking.

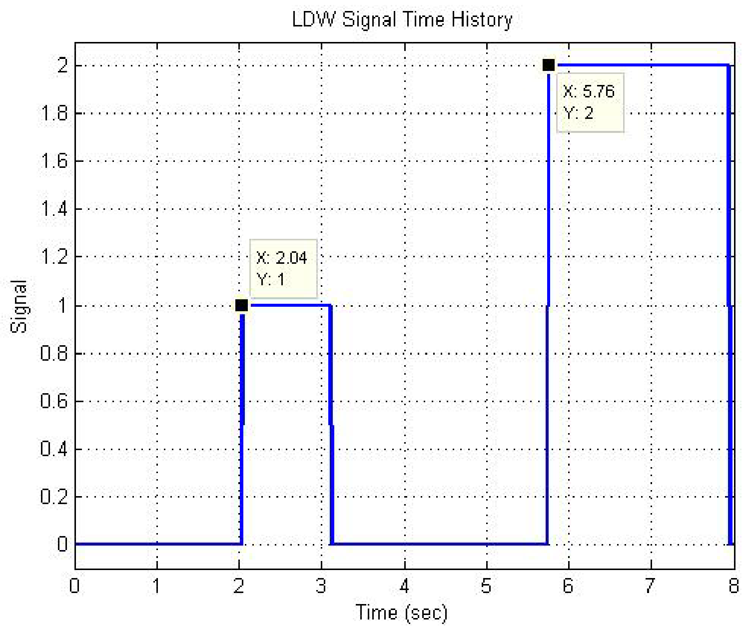

Figure 6.14. Time Periods of LDW Signal Activation [0: no warning, 1: left side warning, 2: right side warning]

Figure 6.15. Screenshot from CarSim's Animator Function Showing the Vehicle Path

Figure 6.16. Vehicle Location when the First Lane Departure is Detected (2.04 sec)

Figure 6.17. Approximate Vehicle Location when the Second Lane Departure is Detected

6.3. *Driver Model*

As previously mentioned, in order to control the vehicle and to respond appropriately during inadvertent lane departures, the driver model has to be capable of three things. First, it must be able to keep the lane, i.e. find out where the lane is going and steer the vehicle accordingly. This includes setting up a procedure for recognition of lane boundaries, a way of calculating whether (and how much) it is necessary to steer in order to stay in the lane, and then communicating the desired steering to the vehicle model.

Second, the driver model must be capable of losing track of where the lane is, i.e. of being so distracted from the primary lane keeping task that lane departures may occur. In this model, this has been implemented through a form of visual distraction. Selecting visual distraction to disrupt the lane keeping task is consistent with, for example, Dingus and Klauer (2008), who report that visual distraction seems to be a primary contributor to real-world crashes and near-miss events. Although the model has the capability of including non-driving workload and cognitive distraction, this is noted as a feature for future development and was not applied in the present project.

Third, the model should be able to return the vehicle to the lane of travel once a lane departure has occurred and the driver model becomes aware of this fact (i.e. becomes alert again), and this return should be performed in the same way a real driver would when discovering s/he's drifting out of the lane.

This driver model thus represents the closed-loop lane-keeping dynamics of the driver and vehicle, including episodes of visual distraction, and recovery to the travel lane after an excursion. According to Figure 6.1, the virtual driver is aware of the vehicle kinematics (V), road geometry and potentially other surface conditions (R), as well as any LDW alert (L). Driver visual attention may be switched between secondary activities and the forward road scene, and various processes direct and activate this switching process. Attention switching is treated as being internal to the driver model.

Driver intention is fixed throughout – the aim is to remain in the original lane – so while attention switching may lead to a lane excursion, there is no particular trigger or decision point to cause the vehicle to leave the desired travel lane. Once attention is directed to the forward road scene, the lane keeping dynamics are essentially continuous. The model comprises three basic elements in a modular structure: input acquisition, information processing including memory management, and control application. Inputs from outside the driver model, shown in Figure 6.1, come from the vehicle (V), the road (R) and the LDW system (L). An additional input is generated within the driver block, but external to the lane-keeping control components; this represents distraction-related activities such as performing visual-manual non-driving tasks, or looking at objects unrelated to lane keeping (the model is not specific; the cause could be driving-related such as checking road signs, or completely unrelated to driving). This additional internal source of input to the driver model therefore corresponds to visual attention switching, such that when attention is directed away from the road, no new visual information is available from the road. During periods of inattention, driver control actions are suspended, and so for example the steering wheel angle is held constant. Information processing is based around a stored "image" of the world scene, represented in the model as a tabular array of lane boundary points and geometric/logical information about those points (see below). As new visual information is acquired, the tabulated information is updated and organized. Information processing includes several sub-functions, as described below, and these take a finite time to execute. Therefore there is a time delay implicit in the model; the information available to support control action is slightly "out of date" when it is applied, and the magnitude of the delay automatically varies based on the details of how much new information is being processed. In particular, in the case where attention is switched away from the road, this information is not updated, the time delay can be several seconds and the stored information may become inaccurate over time, at least until visual attention is restored. Control output is in the form of steering, brake and throttle actions, with steering as the principal action of relevance in the model. The control is based around feedback of errors, so for example when the vehicle drifts towards the points making up one of the lane boundaries, this is detected and the steering is applied to reduce the error.

Figure 6.18 shows the sub-systems within the driver model in more detail. A random attention switching signal is presented to the driver attention block. In the absence of any other demands the output of Driver Attention will follow this demand to carry out non-driving tasks and divert visual attention from the lane keeping task. A lane departure warning signal (together with a number of "attention grabbing" conditions – see Section 7) may override this attention switching, leading to a filtered attention enable/disable signal (E). When E=0 attention is diverted from the roadway and both road scene acquisition and control application are both interrupted. When visual attention to the roadway is restored (E = 1) new data is obtained via the basic image processing, and application of steering, throttle and brakes is again responsive to the information processing block. Note that simulations are initiated with the attention switching signal in the visually distracted state; the vehicle is within the lane boundaries at this time, so E=0 and a drift out of lane is feasible.

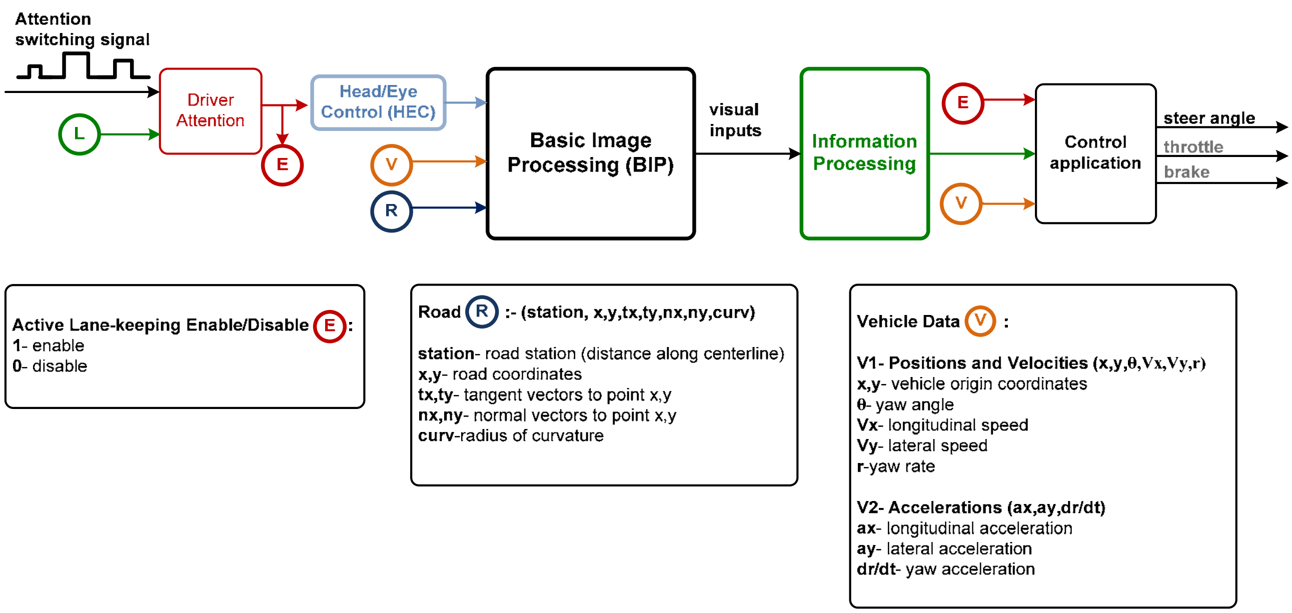

Figure 6.18. Driver Model: Overall Architecture

Figure 6.19 depicts the Information Processing sub-system in more detail. We consider all the sub-systems functions in detail in the following sections, but it is useful to provide an overview here. Visual inputs are in the form of geometric information about the road, represented in the form of sampled points on the lane boundary. This data is placed temporarily in a visual memory block (a memory buffer labeled *visual field* in the following section). A house-keeping function moves data to a working memory that is then used as a common data store for a number of parallel processes. Object recognition and classification selects any entries in the working memory that are new and unclassified; with some limits on the number of points that can be processed at any one time, it decides on relevance and classifies the boundary point as being "left of vehicle path" or "right of vehicle path". The classification is needed so that threats can be assessed and control intervention applied. The threat assessment block evaluates a yaw rate error criterion for each recognized boundary point and decides (and flags) which are most critical. Finally the essential kinematic information for these critical points – range, azimuth, etc. – is provided to the control application block (Figure 6.18). At this point the controller may or may not make corrections, depending on the level of criticality.

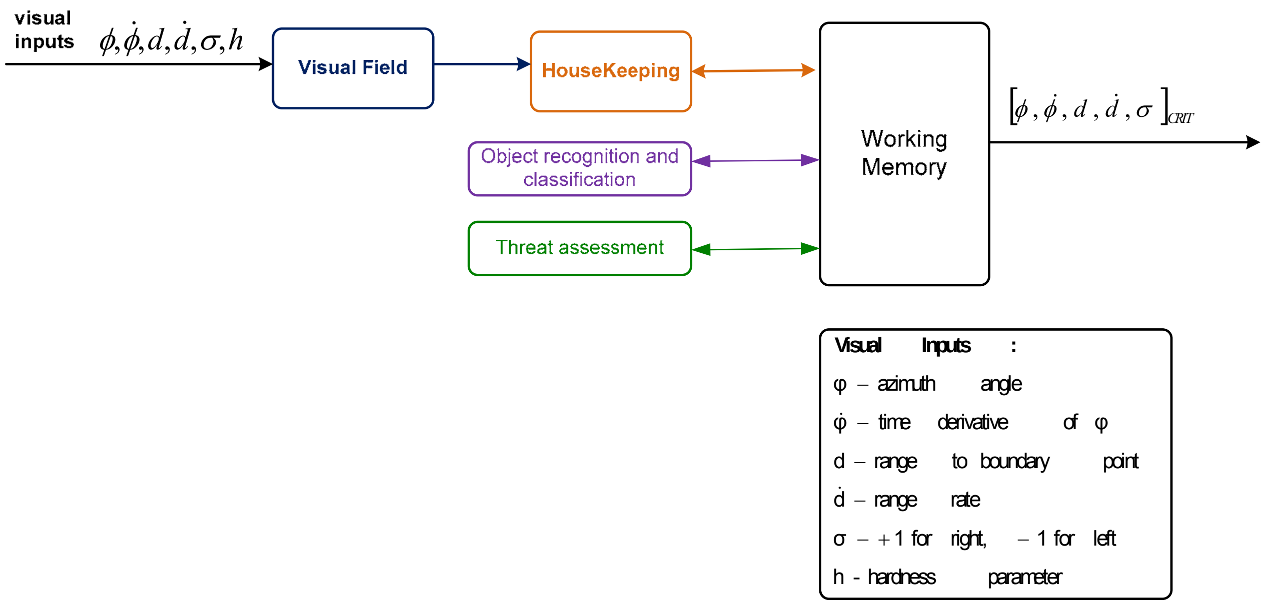

Figure 6.19. Driver Model: Information Processing

One more process is worth mentioning: over time, as new information is received into working memory, its capacity may become saturated, so the house-keeping function removes oldest or least-used data when necessary. We also note that the visual input and working memory data include two parameters σ and *h* which are used in the threat assessment. σ is simply a flag to identify whether the points should be classified "left" or "right", while *h* is a continuous variable that defines criticality in terms of whether the boundary is perceived as "hard" or "soft". In general $0 \le h \le 1$ would be higher for a concrete barrier (e.g. $h=1$) than for a lane edge marker adjacent to an empty paved shoulder (perhaps $h=0.2$); however, for simplicity in this project, we make no such distinctions and treat all boundary points with equal perceived hardness, *h*=1.[8]

The above describes the general driver model architecture, and outlines how each sub-system functions within it. Each block that performs information processing is assumed to operate in "discrete time", i.e. it takes a fixed time to execute a single cycle of processing. This is typically true of computer systems, and is a simplifying assumption for the biologic information processing in the human brain. The common time constant T_s for these processes (time required to execute one cycle of processing) is then related to the overall processing delay of the driver model, and this point is picked up again in Section 6.4.7. We now provided more detail of the data structure and algorithms used in each of the blocks within the driver model.

[8] For computational efficiency the parameters (σ,h) are included in the visual field data, however they are not made available for threat assessment or control until the object recognition and classification has processed the boundary points in question.

6.3.1. **Visual Field and Working Memory (VF and WM)**

As described above, the various sub-processes move data and operate on the information stored in the visual field (*VF*) and working memory (*WM*) – see Figure 6.19. These two memory blocks share the same general structure of a rectangular array, each with *N* rows, one row for each boundary point. *VF* is an *N × 6* array, where each of the *N* rows holds specific information regarding a given boundary point and 6 properties are associated with each boundary point. The information in each row takes the form $[\phi,\dot{\phi},d,\dot{d},\sigma,h]$,where:

- ϕ is the azimuth angle to the point relative to the vehicle heading
- $\dot{\phi}$ is the time derivative of ϕ
- d is the distance to the point/object
- $\dot{d}$ is the time derivative of d
- σ is the actual (input) signature of the point (=+1 for points on the right lane boundary, -1 for points on the left boundary; note that for curved roads a right-hand boundary point might have a negative yaw angle so this information is needed, in addition to the azimuth angle, to allow the driver to control the vehicle)
- *h* is a "hardness" parameter that is provided for future expansion and development, used to define whether the driver perceives collision with the object to be critical or not. To contain model complexity, we currently regard all lane boundary points with *h* = 1, i.e. the driver model does not recognize different threat levels arising from boundary object type when lane keeping.

To capture a reasonable resolution of the lane boundaries, while maintaining efficiency in the simulation 10 points per lane boundary were found to be acceptable. These are equally spaced between the minimum and maximum preview distances defined by the preview times $T_{\min}$ and $T_{\max}$ given in Table 6.2, hence with a typical separation of less than 5 meters. Thus a minimum value *N* = 20 rows is required for storage in *VF*.

The working memory, *WM,* is represented in similar fashion, as an *N × 9* array; again each row is associated with an individual boundary point copied from the visual field, but with an additional 3 elements of information appended during information processing. For convenience in programming, both *VF* and *WM* contain the same number of rows (i.e. the same value of *N* is used). The working memory uses both stored and refreshed information, so a greater memory capacity is required. For this reason we have chosen *N* = 40 rows in each of the arrays.
Then, as the simulation runs, only half of the *VF* memory is utilized, but this redundancy makes no difference to the results obtained. The additional 3 information elements in *WM* are symbolized as [*age*, *flag*, *star*]:

- *age* this normally remains at zero, but if the object is not matched to *VF* data at any processing step it is incremented by +1. If it is not refreshed within a certain number of processing steps then it is lost from *WM*. Age is also used to flush out excessive points to keep *WM* within a working limit of 40 rows.

flag is initially zero, but object recognition and classification (ORC) sets it to 1, once it is deemed to be recognized. Unless *flag*=1, the unrecognized object cannot be used in threat assessment and hence for inducing a control response.

star is an additional flag, set by threat assessment (*TA*) to 1 for the two points of maximum threat (one for the right boundary, one for the left). Only these "star points" are used to generate a control response.

Note that under most conditions the maximum age (maximum number of time-steps) associated with *age* has little or no effect on performance, but to be specific an upper limit of 10 processing steps was chosen. This corresponds to a maximum storage time in *WM* equal to $10 \times T_s$ seconds, where T_s is the fundamental processing time step in the driver model's information processor (see below). With a maximum value of 100ms for T_s, this corresponds to a maximum retention time of 1 sec for dynamic tracking of boundary objects. The value is somewhat arbitrary, yet has minimal influence on driver model performance.

6.3.2. Head and Eye Control (HEC) and Basic Image Processing (BIP)

In the driver model shown in Figure 6.18, *HEC* performs a the basic switching function to enable "image capture", while *BIP* converts external information about the road – specifically the lane boundary – into an internal representation, i.e. the visual field. Although conceptually different (and shown separately in Figure 6.18) in the model these two entities work together as a single sub-model: if enabled by signal *E* in Figure 6.18, the function is to refresh the visual field at regular intervals with new or updated information about lane boundary points relative to the driver's eye point. In the model no actual image processing takes place, and the function is mainly to perform coordinate transformations to convert absolute position and velocity data for the road and vehicle into relevant range, range-rate, azimuth and azimuth rate variables; in other words the two blocks together represent a perception module in the driver model. Intrinsic information about each point acquired (σ and *h*) is also copied. Note that the HEC/BIP block is assumed to run continuously and without any intrinsic delay. This is for simplicity, because any internally generated distraction signal is not available for separate data capture, and any HEC delay could only be estimated. It is also assumed that time constants for *BIP* are small and can easily be absorbed into the other information processing blocks (see below). On the other hand, when forward visual attention is triggered by an event such as the LDW alert, there is an additional response delay T_{LDW} assumed and estimated (see Section 6.4 and also Appendix C).

6.3.3. Housekeeping (HK)

The functions of *HK* are:

(a) to copy information from the visual field to working memory,
(b) to determine associations of old data points with new ones. When range and azimuth values are sufficiently close, the new kinematic values are used as updates to previous ones and no further object recognition and classification is needed. This reduces the overall perception time for previously recognized objects, and
(c) to sort records according to priority, with nearby and new information rated most highly, and distant or time-expired data ranked lowest. When necessary, to stay within memory

bounds (at most 40 points), the lowest priority records are removed. Note that the driver model does not really track 40 independent target objects; it merely tracks up to 40 defining points on two objects, namely the left and right lane boundaries.

6.3.4. Object Recognition and Classification (ORC)

New objects in the visual field have *flag* set to zero, which represents the fact that they are "unrecognized" by the driver. This has a specific meaning in the context of the model: an unrecognized point is not available for threat assessment or control response. In this way the model can genuinely "look but not see". At each time step, the nearest *n* such objects are processed by *ORC* and determined to be either left or right lane boundary points. In the model used, *n*=2, i.e. at each time step the driver model can only recognize two new lane boundary points.

Thus, in a new situation, the model takes multiple time steps to build an internal picture of the boundary points; in reality this might happen if the driver looks away during a yaw disturbance, so that on returning to the road scene it has changed sufficiently to require re-evaluation by the driver. Or a driver becoming alert after a micro-sleep may similarly wake with minimal stored information (see Section 6.6). The more points that need to be recaptured, the longer the re-evaluation process takes. Although in real driving the ORC process is extremely complex, in the model it takes only a simple set of commands to establish the nearest *n* new but unrecognized objects. Of course this is a great simplification, requiring that all objects are equally easy or difficult to identify in real-time. It is interesting to think that for visually complex road scenes, or in situations where the lane boundaries are ambiguous, ORC may have to process points that are irrelevant to vehicle control, in which case ORC will waste effort and delay the useful recognition of left and right boundary points. Such a process may be relevant to the problem of inadvertent lane departures, though in this study no such ambiguous or irrelevant points are presented to the driver model.

6.3.5. Threat Assessment (TA)

Threat assessment scans all objects in the working memory and prioritizes two points – one on the left boundary and one on the right – those which have maximum associated threat. It clears any existing star-point flags and replaces them with new points. Considering the right boundary first, and from the perspective of the drivers head point, the maximum vehicle yaw rate r_i to avoid conflict with boundary point *i* is given by (Gordon and Magnuski, 2006)

$$r_i = \frac{2U \sin\phi_i}{d_i} \tag{6.1}$$

which is proportional to speed *U* and depends on the distance d_i and azimuth ϕ_i variables mentioned above; this defines the *critical yaw rate* for that point. The right "star point" is the one that requires the maximum degree of correction (or lowest margin of error) i.e. it has the minimum critical yaw rate. Similarly on the left boundary, the star point is the one with maximum critical yaw rate. Again the algorithm is extremely simple, though recent analysis of naturalistic driving data suggests that yaw rate error is a reasonably convincing measure of lane-keeping control feedback

[Gordon et al, 2009]. Note that the star points are recognized from the perspective of the driver's eye position, while in the next block, allowance is made for the offset between the driver's eyes and the relevant front corner of the vehicle.

6.3.6. Control Application (CA)

The *CA* module applies steering and speed correction based on the star points identified in *TA*. If the star points are not identified (for example, during an extended eyes off road time no working memory points are initially matched to the visual field) then no control action is possible, and steer action is inhibited. The critical yaw rate for the right lane boundary, r_{cR}, is identified as the critical yaw rate (6.1) for the identified right star point. Similarly the critical yaw rate for the left boundary, r_{cL}, is the critical yaw rate of the left star point. In CA, the vehicle yaw rate is compared to these two critical yaw rates. The existence of a lane boundary conflict is tested by the following inequalities:

$$r_{cL} < r < r_{cR} \tag{6.2}$$

Here *r* is the vehicle yaw rate sensed by the driver. If both inequalities hold, then no control action is required, and none is applied. In this case the steering angular velocity is set to zero, and current steer angle is maintained. On the other hand, if one of the inequalities is invalid, a conflict exists and steering is applied to reduce the conflict. For example, if $r > r_{cR}$, there exists a conflict with the right lane boundary and the steering rate is set negative according to the proportional control law

$$\dot{\delta} = -K(r - r_{cR}) \tag{6.3}$$

where *K* is the steering control gain and $\dot{\delta}$ is the commanded steering velocity. Note that at any one time, CA can only respond to one critical yaw rate, however many target points exist in the field of view.

The steering gain *K* in equation (6.3) need not be constant across all speeds – indeed, at low speed the yaw rate tends to be low and a high gain is needed, while at high speed we expect a high gain will generate instability. Fortunately there is a simple way to introduce speed dependence prior to parameter tuning, based on the mechanical operation of the steering. According to basic cornering mechanics of a vehicle corning at low speed (Gillespie, 1992)

$$\delta_f = L \times (\text{path curvature})$$

where *L* is the wheelbase of the vehicle and δ_f is the steering angle measured at the front wheels. If ρ is the steering system ratio, and *r* is the yaw rate this implies

$$\delta = \rho L U^{-1} r \tag{6.4}$$

where *U* is vehicle speed and δ is the steering angle measured at the steering wheel. This equation is highly simplified since it applies only in the steady-state and for a "neutral steer" vehicle. In reality most passenger cars are "under steering" which means that at higher speeds a larger steering angle is required than stated in (6.4). However, equation (6.4) does provide a basic way to implement reasonable speed dependence in the steering control law (6.3). According to equation (6.4), if a small change in steering angle $\Delta\delta$ is applied during a maneuver, it roughly generates change in yaw velocity Δr :

$$\Delta\delta = \rho L U^{-1} \Delta r$$

If the change in steer takes place over some characteristic time τ , then the mean steering rate required is $\dot{\delta} = \tau^{-1}\Delta\delta$. Here τ is a simple control parameter (in units of time) related to the speed of steering response; while in an extreme case it might be limited by human neuro-muscular response, it is more realistically related to driver adaptation to the vehicle steering system – a manually tuned parameter that generates adequate speed of response from the vehicle and without any likelihood of causing control instability. For a yaw rate error $\Delta r = r - r_{cR}$ as in equation (6.3) we estimate the mean steering rate to be

$$\dot{\delta} = \tau^{-1} \rho L U^{-1} (r - r_{cR})$$

The steering gain *K* in equation (6.3) is then proportional to the inverse of a characteristic steering response time τ according to the equation

$$K = (\rho L U^{-1}) \tau^{-1} \tag{6.5}$$

In this way the steering control gain is defined via the driver parameter τ that can now reasonably be assumed constant across a wide speed range. Since vehicles typically understeer, the control law (6.3), (6.5) tends to be conservative at high speed, and this is consistent with normal driving behavior.

Two further refinements are included in the *CA* sub-system: *anticipation* and *offset*. Anticipation means that the distance and azimuth angle are both estimated a short time T_{DA} in the future based on current azimuth angle (or vehicle position) and yaw rate (or vehicle velocity):

$$\hat{\phi} = \phi + T_{DA}\,\dot{\phi}\,,\ \ \hat{d} = d + T_{DA}\,\dot{d} \tag{6.6}$$

Here T_{DA} is the driver anticipation time, chosen to offset the major time lags in the operation of the driver model. The second refinement is to implement offset: azimuth and distance are adjusted further to allow for the offset between the driver's head position and the relevant front corner of the vehicle. The geometry of this transformation is elementary, but it is clearly important to make this allowance for the dimensions of the vehicle.

Finally in *CA* we include an additional dynamic response to represent the effect of mechanical action, incorporating a pure delay time (so muscular response is not instantaneous) and a first order filter (so there is some frequency limit in the response). This is mainly for completeness, as the real driver cannot execute steering movements instantaneously. For simplicity we use a single time constant, T_c , and the mechanical action in *CA* is represented by the transfer function

$$G_{CA}(s) = \frac{e^{-T_c s}}{1 + T_c s} \tag{6.7}$$

In this study we do not attempt to characterize the muscular/bio-mechanical response with any realism, so a very small value is chosen: $T_c = 20ms$. This corresponds to a very high frequency bandwidth of 50Hz, so the control response is undoubtedly faster than any real driver can achieve. On the other hand the precise value is not very important in this study, since the mechanical action delay is necessarily confounded with the information processing lags used elsewhere in the driver model; similarly the assumption of equality in filter and delay time constants is unlikely to be significant. In the future a more detailed characterization of biomechanical response may be included.

This completes the basic description of the driver model. But before we consider parameter estimation and model validation, we take a more detailed look at the inherent time delays in the model.

6.3.7. Processing Delays and Subsystem Timing

As mentioned above, all information processing operates in discrete-time. With the exception of HEC/BIP, each sub-system carries out their information processing in time T_s per processing cycle, and during this cycle they are not available to refresh the input data. They also work in parallel, so that as new data is provided into the visual field, the first subsystem to process the data is HK, and other sub-systems must wait for the new data to appear in the working memory. Of course the other blocks are not idle, but are effectively working with "old data" carried into working memory in a previous *HK* cycle. The results of HK updating the working memory become available to the other sub-systems after a period equal to this fundamental time constant T_s. In the worst case all blocks must perform their processing steps before new information reaches the control output component of CA. The precise processing time-delay between visual detection and control action, if any, for a particular road boundary point depends is not deterministic. If the working memory is empty or contains obsolete information (e.g. after a prolonged period of visual distraction) the delay can be as long as $4T_s + 2T_c$. However, under steady conditions when the working memory is being continuously update this delay is reduced to a value closer to the following expression.

$$T_{delay} \approx 2T_s + 2T_c \tag{6.8}$$

From this, and the parameter values deduced in Appendix C, the information processing times in the model vary between approximately 100ms and 400ms, the actual value being non-deterministic. Response to an LDW alert incurs addition delay of between 300ms and 800ms in the model, so that overall reaction times from LDW signal to steering response can be anywhere between 0.4 and 1.2 seconds. These values are consistent with the VIRTTEX driving simulator studies, as described in Appendix C..

6.4. *Driver Model Parameter Estimation and Validation*

The key parameters for the driver model are listed in Table 6.2. The estimation of their values was based on a number of test conditions that involved specific DVET simulations and comparisons with objective test data, the principal comparison being with human factors tests with naïve subjects in the VIRTTEX driving simulator (VIRTTEX Studies 1 and 2 of Section 5.1). It was possible to define all driver model parameters to have common units of time. Here we provide a brief summary of the parameter estimation process; for further details see Appendix C.

Starting with the first entry in Table 6.2, parameter T_s (the underlying time step common to the driver model information processing modules) was allowed to take a wide range of values: 50, 100, 150, and 200 ms; these four cases were carried through to the point where performance comparisons could be made between simulation and VIRTTEX results. At that point it was found that the range 50-100 ms gave reasonable results but that larger processing times did not.

Table 6.2. Summary of Key Driver Model Parameter Values – See Appendix C for Further Details

Parameter	Value [sec]	Source
T_s	0.05-0.1 (uniform distribution)	Matching simulations to data during transient recovery from a yaw disturbance
T_c	0.02	Hypothetical value –connected to the choice of T_s so small value chosen
τ	0.22-0.3	Table C.4: apply linear interpolation based on the chosen value of T_s
T_{DA}	$T_{DA} = 2T_s + 0.22$	Equation (C.2)
$T_{\min}$	1	Chosen to be substantially less than $T_{\max}$ but otherwise arbitrary
$T_{\max}$	2.0-2.4	Table C.5: apply linear interpolation based on the chosen value of T_s
T_{LDW}	0.4-0.8 (distracted) 0.3-0.4 (fatigued)	Distributions of simulation and VIRTTEX reaction times

As seen in equation (6.8) parameter T_c is summed with the information processing delay T_s, so it was not possible to estimate T_c independently; hence it was fixed at a relatively short value, 20ms, leading to a typical overall delay of 40ms in the motor response part of the model.

The next parameter, τ, is a measure of the amplitude of steering response – the steering angular velocity is proportional to τ^{-1} - equation (6.5). On the basis that a driver typically adapts the control response to ensure both stability and responsiveness (instability results if τ is too small, and an unresponsive driver model results if τ is too large) it was possible to estimate the optimal values for τ as a function of T_s (see MacAdam, 2003 for a description of this type of adaptation by the human driver). As would be expected, τ increases with driver delay time; τ =220ms for T_s= 50ms, τ =300ms for T_s=100ms, and linear interpolation is used between these values of T_s.

T_{DA} in equation (6.6) is a measure of the driver anticipation used to improve stability (in control theory terms it provides phase lead to the feedback loop) and it is simply set to cancel the estimated minimum time delay in information processing and motor response in the model.

$T_{\min}$ and $T_{\max}$ define the preview limits of the driver model during lane keeping. $T_{\min}$ is not a sensitive parameter, as the driver model includes a scanning process that normally prefers larger preview times to maintain lane keeping quality and stability. By contrast $T_{\max}$ has a significant effect on lane keeping performance, particularly when recovering from a disturbance. Similar to the case with parameter tau, the driver model is to adapt preview time to ensure stable and effective path tracking via linear interpolation between T_s and $T_{\max}$ (the larger value of $T_{\max}$ corresponding to the larger value of T_s).

Finally, T_{LDW} is a reaction delay for the driver response to an LDW warning. Note that it is a measure of delay internal to the model, and is additive with the other time delay components mentioned above. In Appendix C, a number of overall reaction times are estimated from video and from time histories of the driver response in VIRTTEX. From the other known delays in the driver model, the contribution from T_{LDW} was estimated, and validations are shown in Figures C.9 and C.10. It is interesting to note that delays for distracted drivers were found to be longer on average than for sleep-deprived drivers. We can understand this based on two lines of reasoning. First, it has been reported in the scientific literature that, contrary to popular belief, a drowsy driver has a typical reaction time that is not slowed by the drowsy state (Horne and Reyner, 2001). The sleep-deprived driver may not respond at all to an alert, but if a response is generated its timing remains typical of the alert driver. In fact, in Appendix C we find that the quality of response is also not reduced, and in the cases analyzed the settling time of driver-vehicle-environment system in the simulator is actually faster for the sleep-deprived driver. Second, the distracted drivers had a specific reason to delay attention switching, namely the priority of completing a number reading task – numeric values were lost to the subject when visual attention was fully redirected to the roadway. In fact this research project had little information available on the "task inertia" associated with typical distraction tasks in real-world driving, so a direct comparison with the number reading task could not be made. Overall, because of the relative uncertainty on the T_{LDW} range of values, in Section 9 we include a sensitivity analysis for the estimated effect of T_{LDW} on safety benefits.

7. Scenario Implementation and Batch Simulation

This section describes how scenarios are formalized for large-scale simulation and hence provide the source data for performance analysis and safety benefits estimation. Referring back to Figure 2.1 this corresponds to blocks 11 (Scenario Definition) and 15 (Data Generation). So far we have set the underlying parameter ranges for the driver and vehicle components of each simulation. Now, the environmental and highway conditions, as well as the detailed initial kinematics, are to be defined for each simulation. The basic approach was described in Section 4.5, and now we give more detail on how the precise within-scenario parametric data is selected.

We focus our attention on the case where the LDW system is absent; for comparisons to be made with the LDW enabled case, the method is to switch a single variable in the simulation model and hence introduce the effect of the system. After enabling LDW in the model, all parameter choices (deterministic or selected via a pseudo-random number generator) are exactly replicated from the baseline (LDW-absent) condition.

7.1. *Parameterization of Driving Scenarios and Beta Sampling*

In Sections 2 and 4, we described how a driving scenario is a simplified representation of a real-world driving situation which may lead to crashes previously identified as relevant to the safety technology. The common aspect of the driving scenarios is one of lane-keeping, in a regime of speed and other conditions that are consistent with the operation of the Volvo LDW system. This means that vehicle speed should be sufficiently high, no active lane change or turning maneuvers are being executed, and the highway type will normally support lane markings of sufficient quality to be “visible” to the optical lane tracker (so, for example, road departure crashes on unpaved roads are not considered in the analysis). The driver is assumed to be nominally engaged in maintaining the current lane, though inadvertent departure from that lane is possible, due to deterioration in tracking performance by the driver – typically though some form of distraction, attention wandering or fatigue. If inadvertent lane departure does occur, the driver may return safely to the lane as attention is regained; or there may be a collision with an object outside of the original lane. Other outcomes, such as the vehicle drifting off the highway and coming to rest safely, are not explicitly considered in this study; this is because there are no data to reliably count how often such outcomes occur, even though such counts would be relevant to understanding influential factors and validating the virtual crash population used.

Data input for the safety benefits analysis is in the form of a large set of simulations based on the defined driving scenarios; each simulation assumes deterministic conditions defined via parameters of the DS, as well as the random sampling methods described below for creating detailed initial conditions, model parameters, highway and environmental factors within the DS. The within-scenario data are required to define the pre-conditions in sufficient detail to actually perform the simulation; for example a general “curved road” cannot be used for running a simulation – we need the particular geometry of an actual road, defined by the coordinates of the road centerline. We call any such a simulation – including pre-conditions and outcomes – a *virtual driving event*. Many such simulations are run, based on the particular driving scenario, to form a base set of virtual driving events; suitably weighted, these are used for estimating safety benefits. Note that we randomly select the conditions for these virtual driving events based on actual driving behavior (for alpha variables) and actual crash cases (for beta variables) – see Section 4 for an initial description of the alpha and beta parameters used in this study.

The high level parameters that define a major driving scenario are formalized as a set of DS options given in Table 7.1. As mentioned in Section 4, we introduce an extra parameter for region (urban or rural) to assist in sampling specific highway and off highway conditions. In principle this presents around

$$3 \times 2 \times 2 \times 3 \times 2 \times 2 \times 2 = 288$$

possible combinations, one factor for each column in the table (the actual number is slightly lower as "driving lane" has only one level for the 2-lane, 2-way road type). However not all combinations need to be used when we restrict attention to the dominant conditions under which relevant crashes occur; the top 25 basic scenarios (shown previously in Table 4.6 and aggregated over urban/rural regions) amount to 91% of all relevant crash types. This restriction is helpful in restricting the number of simulations actually conducted.

Table 7.1. Driving Scenario Parameter Sets

Road type	Driving lane	Roadway alignment	Weather and Road Surface Condition	Light condition	Fatigue	Region
2 or more lanes, divided	left	straight	not adverse, dry	daylight	0	Urban
2 or more lanes, undivided	right	curve	adverse, not dry	not daylight	1	Rural
2-lane, 2-way undivided			not adverse, not dry			

As mentioned above, the high level concept of a driving scenario is realized more specifically as a series of computer simulations (virtual driving events) that represent the interaction between DVET elements (driver, vehicle, environment and – when active – the safety technology). To achieve this, a simulation model has been formulated, tested and validated as described in Section 6. Any particular simulation is populated by a number of parameters relating to the model components (e.g. relating to driver reaction time), fixed conditions (e.g. lane width) and initial conditions (e.g. vehicle initial speed).

Again referring back to Section 4, GES crash data are not sufficiently detailed to fully characterize the highway environment - further specific details are needed to support simulation and analysis. In particular, road (lateral) curvature and lane width are critical for conducting simulation, as well as other geometric features, such as shoulder type and width need to be defined for each virtual driving event. Michigan crash data are used to supply broadly representative values for road geometry, lane width etc., and these road characteristics are matched to the corresponding fields in the GES database. As a further step to increase the fidelity of the match, scenarios are stratified according to an urban/rural classification; although this does not directly influence the conduct of simulations, it is intended to reduce any bias that might arise from simple random sampling of the Highway Performance Monitoring System (HPMS) data. (Sampled HPMS road segments are not uniformly spread over the highway system, and the mix of urban and rural locations in Michigan need not be representative of the entire USA).

To determine suitably representative and detailed highway geometries for curved roads (i.e., to provide a definitive track for the road center-line) a small library of highway segments has been created. For the three Road types in Table 7.1, each stratified by urban and rural categorization, a sample of Michigan crashes are geo-located and the crash site inspected using aerial photographs to check whether the associated segment is locally straight or curved. Use is made of segments in HPMS data where Michigan crashes have occurred, and for which there is access to a wide range of highway variables such as lane and shoulder widths. Curved segments are digitized (centerline points are picked out and a curve fitted to these data) and stored; sampling continues until ten examples are obtained in each set, making a small library of 60 digitized road segments. Figure 7.1 shows an example from the library of a "curved" road segment; clearly this would be poorly represented by a segment of uniform curvature for example. On the other hand, for road segments designated "straight" in Table 7.1, the simulations are conducted using idealized straight roads, so no sampling of actual roads is deemed necessary. Vertical geometry (grade, cross-slope, crown etc.) is not considered, i.e. all road surfaces are represented as being flat.

Figure 7.1. Example Road Geometry Sampled from Michigan Crash Locations.

Three specific data tables were developed as part of this project for use in the sampling of driving scenarios based on crash and associated highway data; these are referred to as follows:

1. ***Table_GES_13K***: comprises approximately 13,000 records obtained during the ACAT project from the General Estimates System for ACAT-relevant crashes. These data were the source for Tables 4.5 and 4.6, and also provide information about the crash outcomes (to be used in the safety benefits analysis – Section 9). The associated driving scenario parameters are those of Table 7.1 with the exception of Travel Lane. Individual cases from the GES table are not used directly in the analysis, but the aggregated national estimates are used to establish the weights for each driving scenario in the safety benefit analysis. Event outcomes – accident type and first harmful event are also used in the benefits analysis to refine the driving scenario weights.
2. ***Table_Mich_51K***: comprises approximately 51,000 individual crash records for relevant crashes in Michigan, consisting of crashes that took place on road segments identified in the Michigan HPMS inventory as part of the ACAT project. These data include latitude and longitude coordinates, so the precise highway geometry (for crashes occurring on curved road segments) can be derived. As mentioned, sampling from this table is used to generate the library of 60 segment geometries described above. This data table is also subset for use in the next data table.
3. ***Table_Mich_10K***: comprises approximately 10,000 records from Table_Mich_51K, at locations where the HPMS segment is part of the "sampled" HPMS database created for the ACAT project. This subset of HPMS-registered segments contains greater detail about the highway, such as lane width, shoulder type and width. Lane width is required for each simulation, and this database is sampled during batch simulation, so each simulation has a unique record from this database associated with it. Note that the actual highway geometry is not matched case-by-case, because currently the determination of these geometric factors requires manual input (for centerline digitization); as stated, the corresponding library of 60 "typical" curved roads is sampled to provide a definition of the highway centerline.

Selecting from known crash conditions means choosing the major classes of beta variables. The detailed sampling within these classes ("beta sampling") makes use of the above tables in the following series of steps:

1. The initial *major scenarios* $\{S_1, S_{2,} \cdots, S_n\}$ are defined by the rows in Table 4.6 (n = 25). Each major scenario has 2 or 4 alterative versions of urban/rural and left/right lane according to the code:

    alt_code = 1: rural, right lane
    alt_code = 2: urban, right lane
    alt_code = 3: rural, left lane (where applicable)
    alt_code = 4: urban, left lane (where applicable)

    A fixed number of baseline simulations are run from each alternative version of each major scenario. The initial travel lane is required for the 13 scenarios that have at least two travel lanes per direction on the road (and as stated we assume precisely 2 lanes in each direction of travel). In all, with 13 multi-lane cases in Table 4.6, plus 12 with a single travel lane, and taking account of the 2 urban/rural choices, there are (13×2+12)×2=76 scenarios requiring simulation. At the level of the major scenario (including their alternative versions) a systematic program of simulation is adopted, with n_s simulations run for each of the 76 cases.

To limit the overall simulation time, the number of simulations was restricted to a few tens of thousands in total. This is based on the fact that virtual driving event simulations typically take around 20 seconds to complete. This applies to both the physical computation time and the duration of the simulated event, since on a high end PC (Dell Optiplex, Intel quad core processor, 2.67 GHz, 4 GB RAM) the simulations run approximately at a real-time rate. In this case a batch of 10,000 simulations takes just over 2 days to complete, assuming there are no interruptions. We have set

$$n_s = 200$$

i.e. 200 cases per scenario (including the available alt code for each major scenario); with 76 such scenarios and a requirement to replicate simulations 2 times each (with and without LDW enabled) we obtain an estimate of the total number of simulations:

$$N_s = n_s \times 76 \times 2 = 30{,}400$$

In fact a number of short-cuts allow this number to be reduced in practice (see below) but there remains a requirement for around 20,000 simulations, taking a little over 4 days of continuous running to complete.

2. For each major scenario assign a run identifier and use this to set the seed of a random number generator (RAND function in Matlab) – the seed is then used for following randomized selections from with the beta classes, and also for alpha sampling.

3. For each major scenario, randomly select a record from the corresponding subset in Table_Mich_10K. This record directly determines specific highway conditions for simulation and crash metric analysis. If the scenario is for a straight road, no further information is needed from these tables in order to conduct a simulation.

4. If S_i is for a curved road, also assign a randomly selected digitized road segment from the library of 60 digitized road segments, corresponding to the particular road type.

The above completes the description of how beta parameters are systematically selected at the level of a major scenario and sub-sampled at the level of individual crash locations found in the Michigan crash data. We now turn to the detailed initial conditions used for the vehicle and the driver.

7.2. *Sampling of Alpha Parameters*

As described in Section 4, alpha parameters provide essential *within- scenario* information, sufficient to conduct a simulation. Two types of α parameters are identified:

- ALPHA1: vehicle kinematic variables used for initializing the simulation, obtained directly by sampling from naturalistic driving data
- ALPHA2: vehicle and driver model parameters obtained via random sampling from pre-determined distributions.

Additional parameters for the safety system are obtained from objective testing, but these are essentially fixed throughout simulations, and are not part of the sampling process; see Table 7.2.

Table 7.2. Vehicle Parameters and Initial Kinematic State Variables.

		Type
External disturbances		
	aerodynamic drag	α 2/fixed
	road induced yaw moment	α 2/fixed
Vehicle parameters		
	type, size, load	α 2/fixed
	suspension and steering	α 2/fixed
	systems (ABS, ESC, …)	α 2/fixed
	safety system level	α 2/fixed
Vehicle initial states		
	CG position	α 1
	speed	α 1
	yaw angle	α 1
	yaw rate	α 1
	lane position	α 1
	lateral velocity	α 1
	steer angle	α 1
	throttle position	α 1
Driver parameters		
	steering gain	α 2
	speed control gain	α 2/fixed
	min/max preview times	α 2/fixed
	steer/brake bias	α 2/fixed
	time delay parameters	α 2
	fatigued driver	α 2
Driver initial states		
	task switching state	α 2
	attention/delay states	α 2
	lane boundary error state	α 2/fixed
	steer torque disturbance state	α 2/fixed

Sampling for ALPHA1 is conditional on the major scenario types of the previous section. Ideally this would include the “alternative version” designations leading to 76 cases, though for the present study the alpha sampling was only made conditional on the 25 major scenarios. The various ALPHA1 variables are expected to be correlated – for example the yaw velocity, lane position and lateral velocity within the lane are expected to be correlated as drivers make typical corrections during lane keeping. Rather than try to model the corresponding multivariate distribution, we instead sample directly from measured RDCW naturalistic driving data (LeBlanc et al 2006); in this way any inherent correlations in the vehicle kinematics are preserved but without explicit statistical modeling. The results of sampling are used in estimation (Section 2.3.2) where expectations relative to the multivariate probability density $f_i(\alpha)$ may be estimated by random sampling from events in the driving database. For example, if *n* samples are chosen randomly based on the driving data relevant to major scenario S_i , the transition probability in equation (2.12) may be approximated by

$$T_{ji} = \int h_j(\xi_{\alpha,\beta}) f_i(\alpha) d\alpha \approx \frac{\sum_k h_j(\xi(k))}{n} \tag{7.1}$$

where *k* =1,..,*n* is the sample number and $\xi(k)$ is the *k*th trajectory obtained by simulation from the major scenario. However, as discussed in Section 2, this random choice of driving events is wasteful: under most conditions the vehicle is being driven in a stable condition within the lane and, even with a distracted driver model in control, adequate lane keeping results in the great majority of virtual driving events, so most cases will not lead to a lane departure event and possible LDW warning. To address this practical issue of simulation efficiency, a stratified sampling scheme has been devised based on a single "control variable" formulated to be closely correlated to excursion probability within any single event. The idea is to partition the α_1 parameter space into regions with known frequencies in the driving population, and, by over-sampling the regions with high probability of lane excursion, increase the number of LDW-relevant driving events. Knowing the naturalistic frequencies, corrections to the above equation can then be made to compensate the oversampling bias.

In this work we adopt the "inverse time to lane crossing" (ITTLC) as the control variable for α_1 sampling. ITTLC is the reciprocal of the estimated time to lane crossing given the instantaneous position and lateral velocity of the subject vehicle. Note that other measures may be used; the requirement is that, when averaged over other factors, such as driver model attention switching, the probability of an excursion increases as the control variable – as measured in the naturalistic driving data – also increases. Thus the reciprocal is used to provide the desired positive correlation. Since the initial conditions of the virtual driving event are always to precede an actual lane departure, the predicted time to lane crossing is strictly positive and hence ITTLC values always produce a finite result. Figure 7.2 shows the simple algorithm for calculating ITTLC, where V_y is the lateral velocity of the vehicle relative to the lane markings and *s* is the distance between the outside of the relevant front wheel and the inside of the lane marker. To represent both left and right lane departures in a consistent manner, ITTLC is considered negative for departures to the left, and the sign of the lateral velocity is used to decide which direction is relevant.

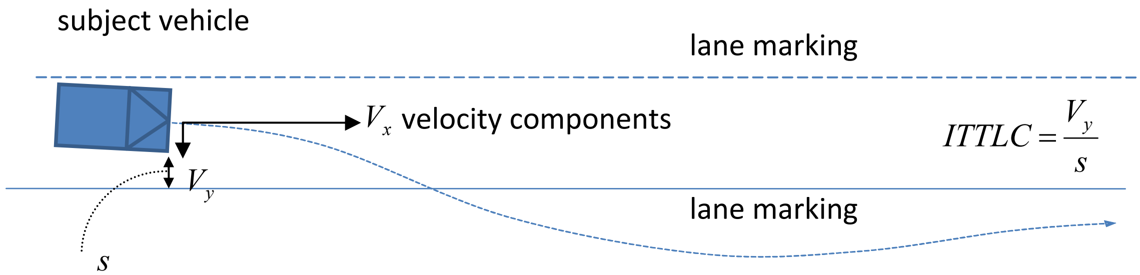

Figure 7.2. Inverse Time to Lane Crossing (ITTLC). Note: *s* is the distance between the outside of the leading front tire and the inner edge of the nearest lane line marker, shown here for lateral drift towards the right lane edge.

To support the analysis, distributions of ITTLC were derived from the RDCW database for each of the 25 major scenarios. The central portion of the ITTLC distribution has a roughly Gaussian form – see Figure 7.3. In the figure, histograms and cumulative frequencies are shown for two comparable scenarios (Scenarios 1 and 10 - see Table 4.6). Figure (b) is for a curved road and, as might be expected, shows a slightly broader distribution than the straight road case (a). The distribution tails are not Gaussian in form, and are not expected to be since ITTLC was truncated by the condition that the vehicle be located within the lane, no closer than 1cm to the nearest lane marker (since the driving scenario captured must always precede any potential LDW alert).

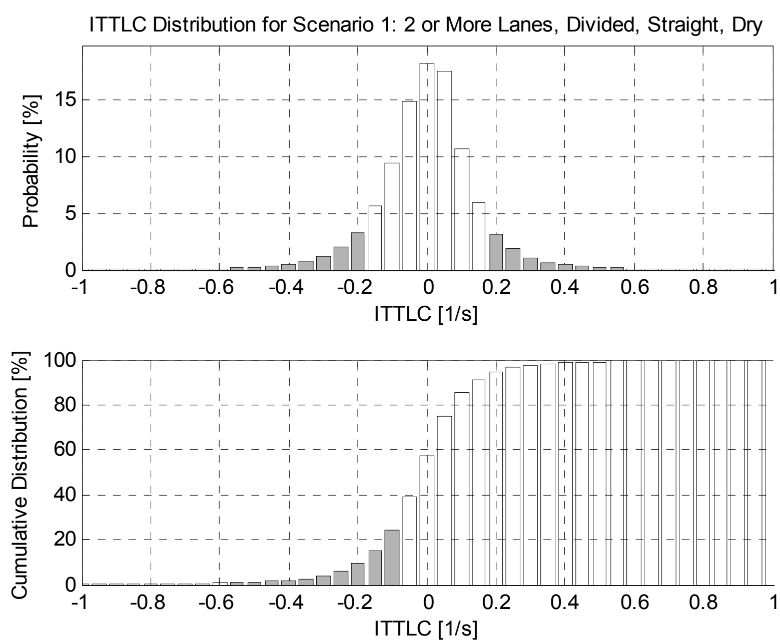

Figure 7.3(a). Probability Density and Cumulative Probability Distributions of ITTLC for Major Scenario 1 (RDCW data)

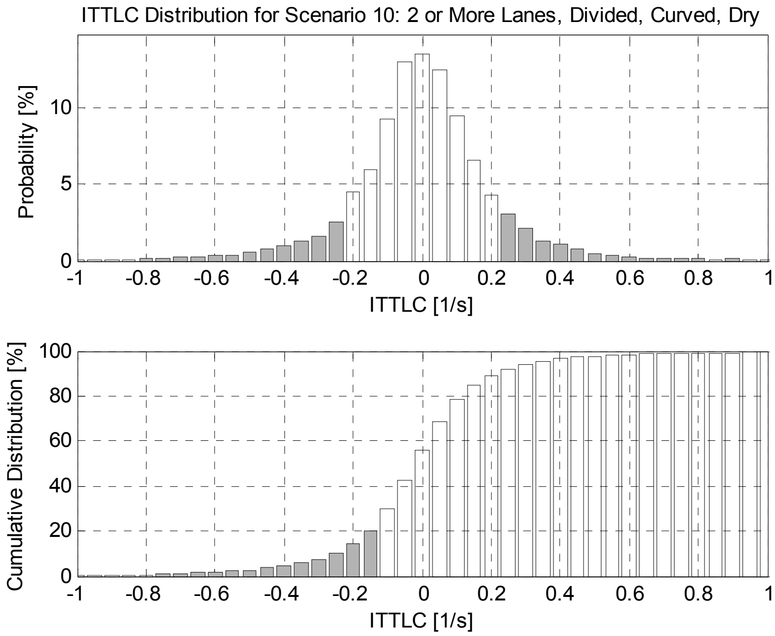

Figure 7.3(b). Probability Density and Cumulative Probability Distributions of ITTLC for Major Scenario 10.

For the purpose of stratified random sampling, the central portion (around 95 - 98% of the total) of the ITTLC distribution was divided into bins of uniform width, and with relative frequencies determined from the relevant driving conditions in the RDCW database. In this study all scenarios used the same 10 bins with ITTLC ranges {(-0.5,-0.4), (-0.4,-0.3), …, (0.4,0.5)}. Thus even in the most severe case there was at least 2 seconds of predicted time in lane before the outer edge of the front wheel made contact with the lane marker. Table 7.3 shows the corresponding relative frequencies as percentages, grouped by bin, for the same scenarios as in Figure 7.3. Again, as expected, tail frequencies are higher for curved roads (Scenario 10) than for straight roads (Scenario 1).

Table 7.3. Relative frequencies for ITTLC bins –Scenarios 1 and 10.

ITTLC range	Percentage Relative Frequencies by Bin									
min	-0.5	-0.4	-0.3	-0.2	-0.1	0	0.1	0.2	0.3	0.4
max	-0.4	-0.3	-0.2	-0.1	0	0.1	0.2	0.3	0.4	0.5
Scenario 1	0.7	1.6	4.2	11.8	33.9	29.3	12.6	3.9	1.4	0.6
Scenario 10	1.5	2.8	5.7	12.9	28.9	23.2	14.4	6.3	2.9	1.3

With 10 bins and 25 major scenarios, it was considered adequate to choose 5 independent cases per bin. The resulting 1250 cases each provided the variables for *vehicle initial states* in Table 7.2. To maximize the independence of the sample events chosen, the 50 cases in any given scenario avoided duplication of individual driver and trip combinations wherever possible, and the number of times any individual driver was represented was also kept to a minimum. Multiple cases from a single driver/trip combination were only occasionally required, this being for scenarios with low levels of exposure (e.g. curve events at night in bad weather).

Of course, with many hundreds of simulations being made for each scenario, each set of initial conditions was to be used more than once. However other independent random selections (for β and α_2 parameters) mean that there will rarely, if ever, be simulations recorded with the exact same trajectory outcomes. Note, due to limitations in the coding of the RDCW data, no cases were explicitly coded as “fatigued”. This driver factor was therefore ignored in the sampling process; for the 4 scenarios coded as “driver fatigued” in Table 4.6 the corresponding “non-fatigued” event data was used (so for example Scenario 11 is considered identical to Scenario 4 for the purpose of α_1 sampling). In the future, assuming it is possible to decide whether certain episodes of driving are subject to fatigue, further stratification can be achieved.

In summary, α_1 sampling consists of the following steps. For each scenario S_i, the RDCW database is sampled, as described above, across all drivers to determine a frequency distribution of inverse time to lane crossing, and a sample of 50 cases is created for each scenario. During batch simulation of a given major scenario, a pseudo-random number generator is used to select one of the 50 cases with uniform probability. The stored data table of sampled kinematics is then used to initiate simulation based on the real-world measurements of vehicle states.

By comparison, α_2 sampling is straightforward. External disturbances were not introduced (aerodynamic drag is included as part of the vehicle physical modeling but does not influence the yawing motion of the vehicle) and vehicle parameters were set to represent a typical mid-sized sedan (see Section 6). Driver parameters were sampled based on uniform distributions of key parameters (again see Section 6) with particular emphasis on the basic discrete-time operating rate (driver model clock speed) T_s and LDW response time (when appropriate). Steering gain was adjusted to the driver time delay based on overall stability and in the case of a fatigued driver the LDW reaction time was adjusted and situation awareness degraded, again as described in Section 6. Compared to α_1 sampling, the supporting data for α_2 is incomplete and simplified, but it is consistent with the tests conducted and at the very minimum serves to demonstrate the methodology. In the future it is expected that more extensive HMI testing will be conducted and events from naturalistic data used to create a more complete parametric representation of driver actions, especially to include a representation of “emergency response” behavior.

We now return to the question of how α_1 combinations based on uniform sampling from ITTLC bins affects the safety benefits analysis, and how the resulting bias is compensated.

7.3. ***Estimation of Transition Probabilities***

The relative frequencies within each ITTLC bin are to be used as within-scenario weights for the benefits calculation described in Section 2. There we estimate crash metric expectations based on the naturalistic distribution of the alpha parameters. In particular the conditional transition probability from driving scenario S_i to crash outcome C_j is expressed as a weighted sum (or integral) of the form

$$T_{ji} = P(C_j \mid S_i) = \int h_j(\xi_{\alpha,\beta}) f_i(\alpha) d\alpha \tag{7.2}$$

where $h_j(\xi_{\alpha,\beta})$ is the crash probability estimated from the simulated vehicle trajectory $\xi_{\alpha,\beta}$ and $f_i(\alpha)$ is the multivariable probability density function with respect to α variables in the population of driving under scenario S_i. Multivariate integration over α is reduced to a univariate summation: we stratify sampling from the ten ITTLC bins and obtain expectations based on the relative frequencies in those bins. In this way we over-sample from the more extreme cases and remove the resulting bias by incorporating the measured ITTLC frequencies. Let the relative frequencies within each bin be given by ϕ_b (*b* =1,2, …,10). The above expression for T_{ji} is then estimated as a weighted sum of expectations within each bin: assuming n_b instances in bin *b* the estimate is

$$T_{ji} = \int h_j(\xi_{\alpha,\beta}) f_i(\alpha) d\alpha \approx \sum_{b,k} \frac{\phi_b}{n_b} h_j\big(\xi(b,k)\big) \tag{7.3}$$

where now the index $k = 1,2,\cdots,n_b$ refers to the instance of simulation within bin *b*. In principle this equation can be used whatever the strategy for sampling between bins. But now we see the importance of maintaining all n_b values strictly positive, and this is most easily achieved by performing uniform random sampling of the 50 sample events across all 10 bins. This is very different from implementing random sampling across the population; as mentioned, the effect here is to over-represent the more extreme cases in the virtual crash population, but weight the outcomes which are weighted via ϕ_a which diminishes towards the distribution tails. Equation (7.3) is a fundamental part of the SIM methodology, providing an efficient method to amplify crash risk in simulations, but without introducing systematic bias.

7.4. *Simulation Batch Control*

Batch simulations are run in groups of 6 related road types, as in Table 7.4. Within each category (such as *Type B Rural*) the probability of a crash as a function of lane or road excursion is assumed uniform, but between categories it may change. This grouping is important for establishing parameters in the crash metric (see Section 7.5) and therefore it is convenient to run simulations and process trajectories (compute crash metrics) according to these groupings.

Table 7.4. Road Types for Common Crash Metric Parameters

Road Category (urban or rural)	Description (total travel lanes)	Scenario Numbers
A	4 Lane Divided Highway	1,3,9,10,12,13,15,16,18
B	4 Lane Undivided Highway	2,8,19,23
C	2 Lane Undivided Highway	4,5,6,7,11,14,17,20,21,22,24,25

Care is required to set repeatable seeds for randomization of the initial conditions, so that repeat runs (scenarios with LDW enabled) are possible with no change in any other conditions. Figure 7.4 shows the overall batch processing activity. The scenario and event numbers were combined to provide random number seeds for selection of α and β variables, as well as the visual attention variable in the driver model. A short (4 second) initial run is executed to allow certain low-level variables such as wheel speed and throttle pedal angle to settle to steady-state values. These are then used to complete the initial conditions in the simulated vehicle and driver so the actual run can be performed.

In the above it was mentioned that around 30,400 simulations are required when executing 200 simulations per major scenario (including urban/rural alternatives, and also left/right travel lane alternatives). This number is reduced to around 20,000 when the following shortcuts are taken account of:

- If the baseline simulation (without LDW) does not generate a lane excursion, the variation in travel lane is irrelevant, as is the effect of LDW being active – in this event no further simulation with LDW active or changed travel lane is made;
- Even if there is a lane excursion, if the vehicle stays on the paved surface, its trajectory can be shifted one lane to the left without any change in simulation conditions; provided the shifted trajectory also remains on the paved surface the trajectory is unchanged and no separate simulation with the left lane start is needed. Of course the LDW enabled case does need to be run, and because the trajectory may intersect different parts of the road when shifted one lane width the crash metric does need to calculated for all four instances (right lane, left lane, right lane with LDW, left lane with LDW) but only 2 simulations are needed (right lane, right lane with LDW).

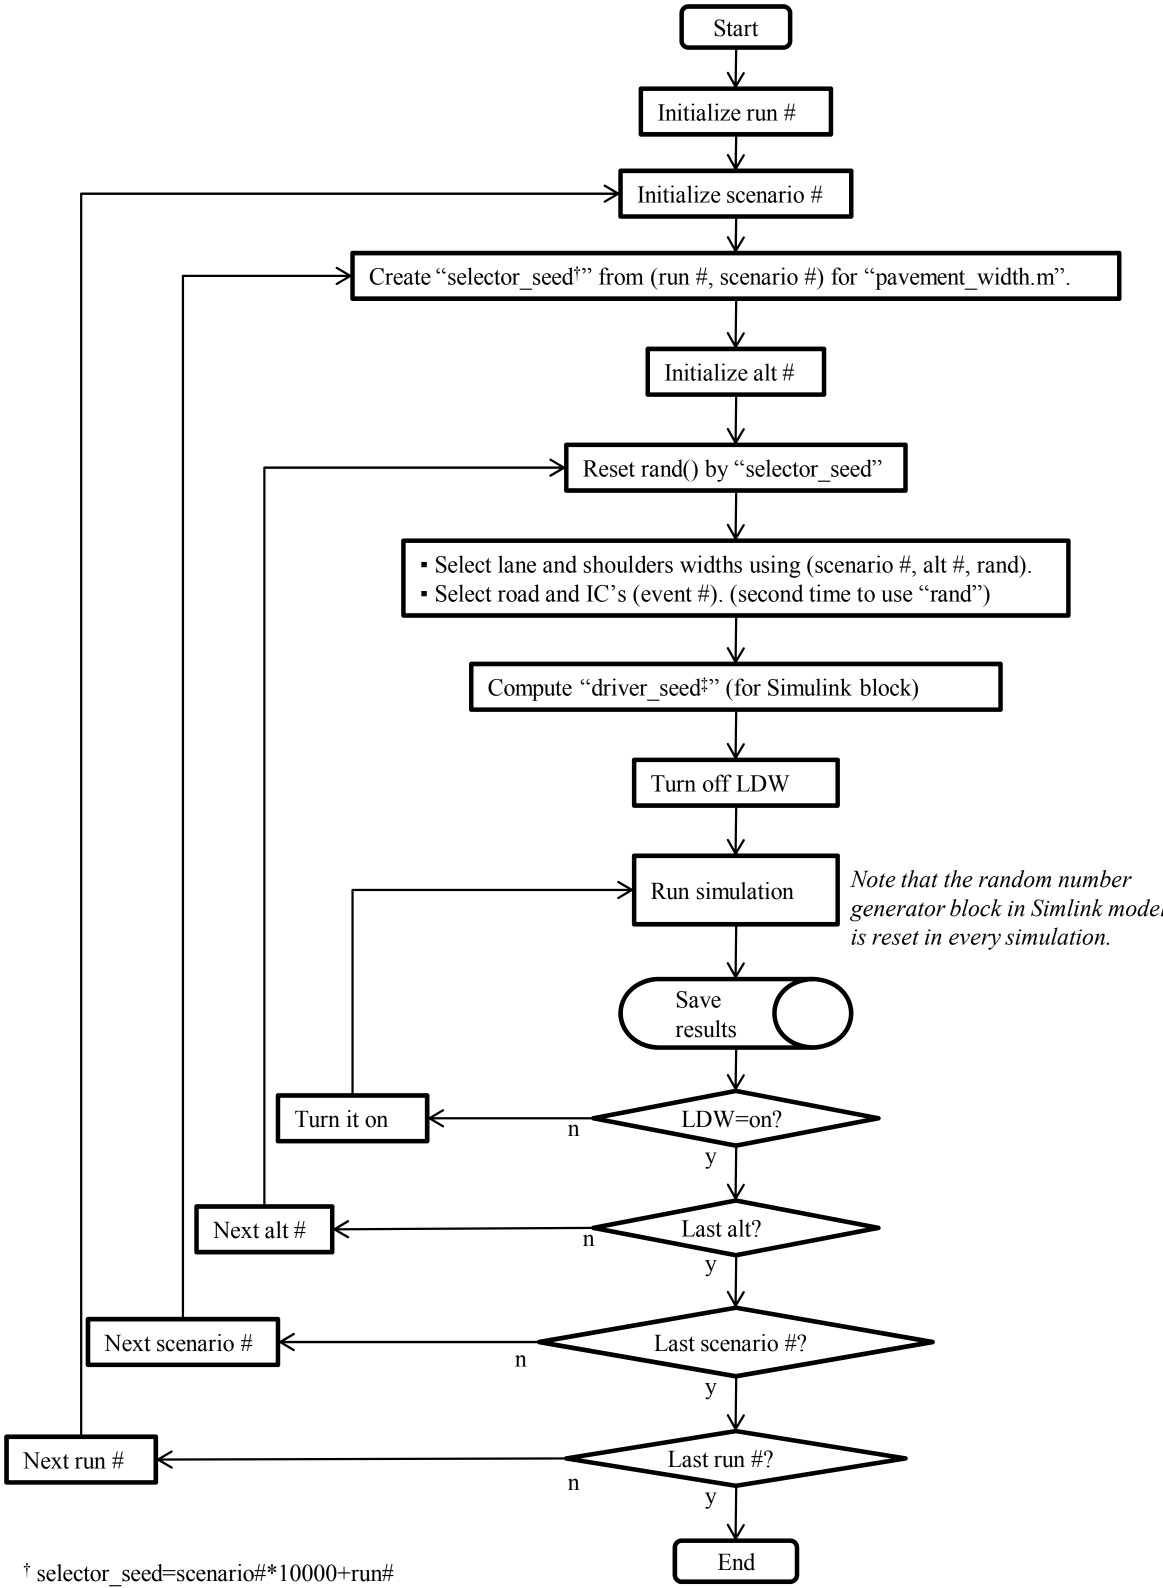

† selector_seed=scenario#*10000+run#

‡ driver_seed=scenario#*100000+event#*1000+roadtype#*100+track#

Figure 7.4. Flow Chart for Batch Run Control (# denotes run counter or identifier, and IC denotes initial condition)**.**

On the other hand there is no possibility to coincide urban and rural cases, since the initial condition sampling is from different data sets in these cases. The batch run procedure is summarized in Figure 7.4 above. An associated file structure has been set up so that all key details of the simulation can be recovered from the supporting data tables. For example the file:

Run_1_S_1_E_11_R_0_0_C_4711652_Alt_3_LDW_1

corresponds to

- Run number 1 from …
- Scenario 1
- Event 11 from the list if 50 ITTLC coded events for Scenario 1
- Road number 0_0 (actually a straight road so null code) from the library of 60 road geometries
- Crash id: 4711652 from the Michigan crash data (Table_Mich_10K)
- Alt_code = 3: rural, left lane
- LDW = 1 (enabled)

A few further details of performing batch simulations are summarized here:

- The surface friction is set to μ=1.0 within dry paved zones, μ=0.25 for off-road, and μ=0.6 for "wet pavement" conditions. These values are broadly representative of the conditions defined in the literature (e.g. Savkoor, 1991), though each type of surface type and zone covers a wide range of possibilities (e.g. the "wet pavement" condition may include standing water or even snow/ice, while the road surface may be concrete or asphalt, so a high degree of simplification has been introduced at this point)
- Maximum visual range for the driver model is $d_{\max}$ = 120 m distance in daylight conditions and $d_{\max}$ = 80 m for nighttime (and again the model has greatly simplified real world conditions which may involve street lighting, glare etc.)
- Driver attention is "fully reset" – meaning the driver attention to the roadway switches on and stays on for the remainder of the simulation - under the following conditions:
    - LDW alert received
    - front tire goes off-road
    - steering angular velocity exceeds a given threshold
    - attention is on at the same time the front wheel has intruded at least 30cm into the adjacent zone (Figure 7.5)

  The point here is that if something very noticeable happens, the driver is expected to stay attentive throughout the remainder of the event: these include "feeling" the vehicle has left the road, noticing a significant lane deviation, or having to make a large correction to the vehicle path. The assumption is that the driver will not lapse immediately into a distracted or inattentive state. Except for "LDW alert received", these reset conditions apply to all simulations – whether LDW is active or not.

- Simulations typically represent 5-20 seconds of simulated driving, in addition to the previously mentioned initial 4 seconds of driving used to achieve steady-state conditions. (See Table 7.5 for various timeout conditions to end a simulation run).

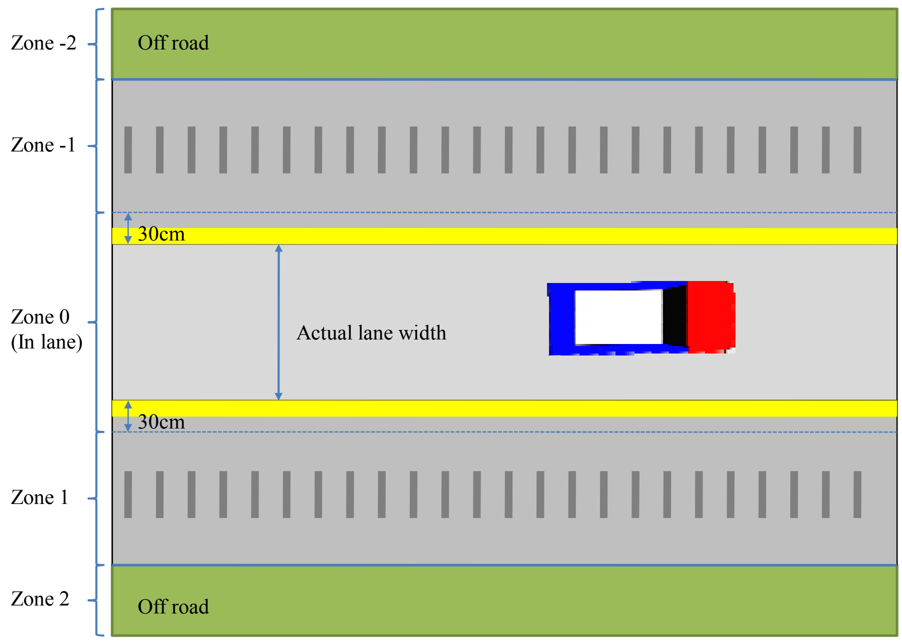

Figure 7.5. Zones for Simulation Control and Crash Metric Analysis.

As will be described in more detail in the next section, the probability of a crash is estimated from the vehicle trajectory across the various zones shown in Figure 7.5. Zone 1 is outside of the initial driving lane, and intrusion here carries some crash probability; in particular this depends on whether the subject vehicle is intruding into an adjacent travel lane, into a lane for oncoming traffic, or onto a paved shoulder or parking lane. Zone 2, which is off the highway, is treated separately, and various weighting factors are used to estimate crash probabilities from related crash data. However, for simulation we only need to determine the physical zone boundaries and compute and store trajectories. Later, in the safety benefits estimation, the nature of the intrusions is taken into account.

Each simulation is designed to represent a single potential lane or road departure "event", not a series of such events, so it was considered important to stop simulation once the potential event is over. This is partly related to the "full reset" of driver attention – in most cases we expect the driver to recover control and restore stability in the lane keeping (albeit with some risk of a crash due to the inadvertent lane excursion). To cover this and other conditions, such as "non-event", the following stopping conditions were implemented

- The vehicle remains continuously within a particular zone for more than a pre-defined time period (timeout condition – see Table 7.5)

- A lateral deviation occurs more than 10m from the paved boundary (typically resulting from a loss of control on the low μ surface)
- The furthest vision point of the driver is within 10 m of the road segment end (typically around 1km distance from the start)

Table 7.5. Timeout Conditions for Simulation End

Driving Zone	Time Limits
Lane (No lane departure occurs)	10s
Lane (After recovering from lane departure)	5s
Paved zone (shoulder, parking space, etc)	30s
Off-road	30s

7.5. *Crash Metric for Lane Departure Events*

In this study we do not represent the detailed locations of potential collision objects. Instead, we propose a distance-based measure of crash probability associated with any given trajectory. According to figure 7.6, we assume a basic crash risk model that associates a crash metric with the lateral or longitudinal distance traversed at various locations outside the desired lane - with crash risk associated with other vehicle, fixed objects (and uneven ground causing further deviations or rollover etc.).

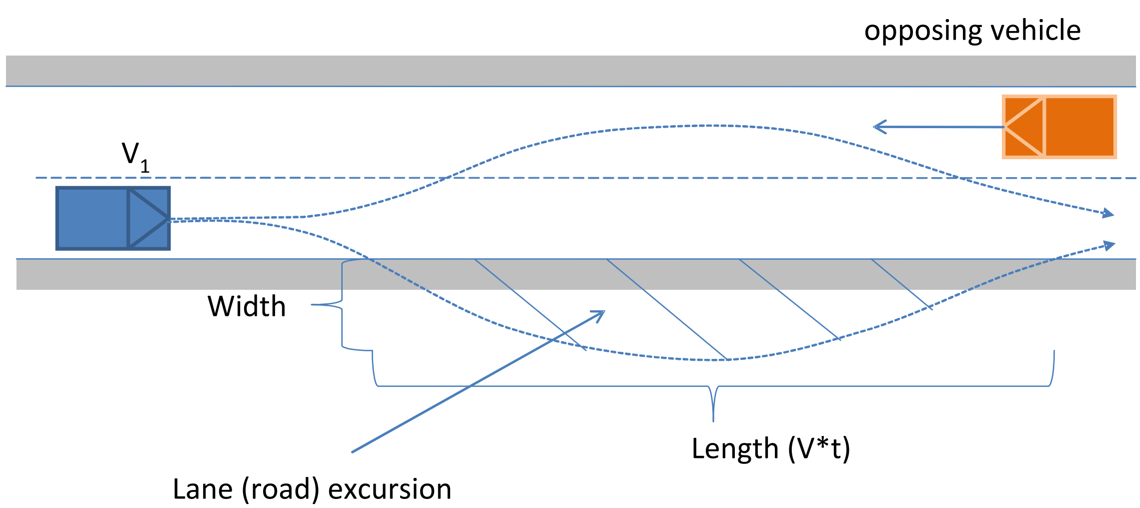

Figure 7.6. Crash Risk Related to Trajectory Output From Simulation.

There are four basic cases to consider, when the subject vehicle drifts into a “conflict zone”:

(a) the shoulder or parking lane, but still on the paved highway
(b) a neighboring lane (traffic in the same direction)
(c) a neighboring lane (traffic in the opposite direction)
(d) off the paved highway (median or right side)

There are many options one could adopt, and for simplicity we adopt one based on

- the distance traveled in the wrong lane for cases (a)-(c)

- the lateral excursion from the highway in case (d).

The simple logic behind this approach – apart from the need for simplicity and the need to avoid a highly detailed representation of the environment – is that the lane deviation is unplanned and there should exist a uniform risk of colliding with a fixed or moving object that is simply proportional to "exposure" – for fixed objects this is distance based (how far does the subject vehicle travel on the shoulder) rather than time based. For collision with other moving vehicles, the subject vehicle speed is certainly a factor, as is the relative speed between the subject vehicle and the POV, but again for simplicity we use a distance-based metric.

In all cases we use the equation

crash probability = $k \times$(crash metric).

where there are four potentially distinct values to be estimated (k_a, k_b, k_c, k_d) corresponding to the four cases. The estimates of these parameters are based on error minimization relative to actual crash numbers, a procedure that is formulated in Section 9.1. Since we do not know absolute exposure values in this study, the *k* values are actually only *proportional* to crash probabilities, so we may normalize so that the maximum of the four parameters equals 1:

$$\max(k_a, k_b, k_c, k_d) = 1$$

In cases (a)-(c) we take distance traveled in the relevant lane as the crash metric, based on a somewhat arbitrary "boundary layer" of 30 cm (12 inches approx.) - the outside of at least one front tire must move at least this far into the conflict zone for the encroachment to "count" in the distance metric. This is because of the maneuvering room typically available when encroachment into an adjacent lane or shoulder takes place. Certainly the parameter value might be better chosen according to some independent analysis, but a zero boundary layer is surely less realistic, so a somewhat arbitrary value has been selected.

For (d), trajectory excursions outside of the pavement (i.e. further out than the edge of the shoulder), we use the maximum road excursion (MRE) as the metric. This normalized metric increases linearly within the clear zone (for road departure) to a maximum at the edge of the clear zone – see Figure 7.7. Based on the β parameter selection (from sampled HPMS data *Table_Mich_10K*) the clear zone is estimated (see below).

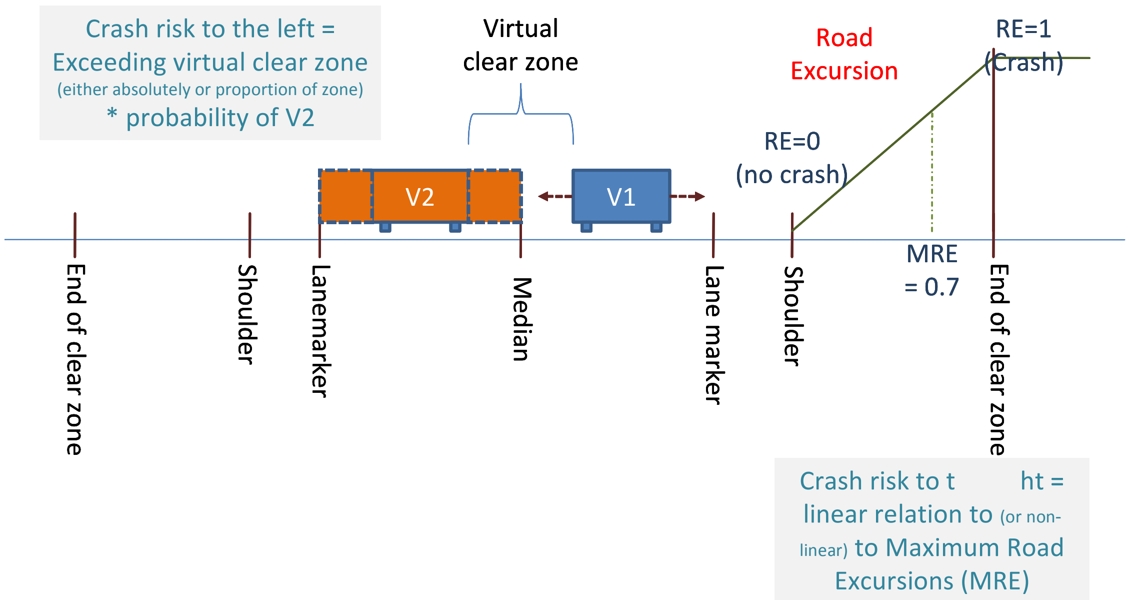

Figure 7.7. Crash Metric Definition as a Function of Maximum Road Excursion (V1 denotes subject vehicle, V2 a principal other vehicle; MRE corresponds to the edge of the clear zone, and crash probability does not increase beyond this point).

The 30 cm "boundary layer" for lane excursion shown in Figure 7.5 is implemented for encroachment into any particular paved zone when computing the distance traveled in that lane, while no such boundary layer condition is applied to the off-road case. Note that when the simulation is run, the fine detail of where different lanes and shoulders sit on the paved surface are considered irrelevant – the driver, vehicle and LDW system are only set to respond to the lane markings from the initial travel lane, and the edges of the pavement. This is a major simplification, since one consequence of this is that a driver inadvertently drifting onto a paved shoulder has the same degree of recovery intent as when drifting into a lane of oncoming traffic. In spite of this simplifying assumption, the model does indeed represent a range of response behaviors, from mild corrections to recover the target lane to severe corrections that induce vehicle instability.

To determine the four components of the crash metric, the lane and shoulder widths, plus clear-zone limits are to be determined. These are found from the crash locations provided in Table_Mich_10K. Each simulation records the crash event it was based on, so in post-processing the same records are accessed to determine these three dimensions. The lane and shoulder (right side and median side if applicable) are directly provided in the table, and indeed these values were used to determine pavement limits in the simulation. Unfortunately, even the sampled HPMS data does not record the actual clear-zone limits, so an alternative approach was needed to estimate this dimension at the chosen crash location.

The AASHTO Roadside Design Guide (AASHTO, 2002) defines the clear zone as "a variable distance from the traveled way, depending on design speed, annual average daily traffic (AADT), and embankment slope rate and direction". In 1989 the Roadside Design Guide was issued by AASHTO and adopted by MDOT as a guide. Using the MDOT Road Design Manual the project team used the AADT and speed data available in the Michigan HPMS to infer a representative clear zone value for simulations. Table 7.6 presents the range of the clear zone values for slopes from 1:6 fill to 1:3 cut grouped by AADT and speed. These values are indicated in section 7.01.11 of the manual. The manual also recommends that for practicality, and to provide a consistent cross section, designers may limit clear zones to 30 feet. However, where there is a high probability or history of crash events, clear zones exceeding 30 feet may be indicated. In the absence of physical data in each case the minimum value in the range was taken as the actual clear zone for the site. Note that all values were converted to meters in the post-processing step.

Table 7.6. Clear Zone Distances

Design Speed(mph)	AADT Volume Groups			
	under 750	750-1500	1500-6000	over 6000
	Clear Zones (ft)			
40	7 - 10	10-12	12-16	14-16
45-50	8-14	12-20	12-26	14-28
55	8 - 18	10-24	14-30	16-32
60	10-24	12-32	14-40	20-44
65-70	10-26	12-36	16-42	22-46

In summary, within this section we have described the main activities involved in running batch simulations. In Section 8 we explore the basic trends in lane departure events and in particular the influence of scenario-related factors on the individual crash metric components. We also focus attention on specific example simulations, the aim being to confirm the overall reasonableness of the individual crash simulations and the predicted LDW effects. Following on, in Section 9, we turn attention to the virtual crash population as a whole, estimating benefits and analyzing the major trends in population estimates.

8. Countermeasure Performance Analysis

The previous section focused on how scenarios are rendered in the model and batch simulations initialized and run. In this and the following section, results are presented from the batch simulations and the safety benefits analysis. This section deals with the basic description of the virtual crash population as well as the major effects of the LDW system. Section 8.1 considers the simulations conducted in the baseline (non LDW) case, so that crash metric components can be determined, and some simple trends can be assessed for reasonableness. Section 8.2 considers simulation output in terms of the LDW effects, detailing the "kinematic benefits" (reduction in lane excursions associated with different regions) implied by the trajectory post-processing. Then Section 8.3 focuses on more detail on some specific simulation examples to highlight how benefits are delivered at the level of the individual simulation. In this section we restrict attention to the level of individual scenarios (and individual simulations) and defer greater synthesis until Section 9, when the overall safety benefits estimation will be considered.

8.1. *Baseline Population*

As described in Section 4, the virtual crash population is expected to comprise 76 scenarios; and with a target of 200 simulations per scenario, this gives 15,200 simulations in the baseline case (no LDW system present). Table 8.1 provides a summary of the actual numbers run in the baseline population. This covers the major road types (labeled Type A, B and C in the following), the urban/rural alternatives as well as the travel lane alternatives in the 4 lane case (2 lanes in either direction).

With the increased sampling of high ITTLC conditions we see the proportion of cases with significant lane deviations is high, averaging 66% of all simulations. Of course this proportion also depends on the degree of driver distraction in the model, and this has not been validated in the population (indeed, for the same reasons ITTLC has been elevated, it is reasonable to expect that efficiency demands again require simulations to be dominated by higher risk cases; however, in this aspect we do not yet have population statistics on visual distraction to populate a bin-wise oversampling as for ITTLC).

Table 8.1. Summary of simulation numbers in the baseline (LDW absent) population

Number of Baseline Simulations	Type A Rural (4 lane divided)	Type A Urban	Type B Rural (4 lane undivided)	Type B Urban	Type C Rural (2 lane undivided)	Type C Urban	Combined
Total	3600	3600	1600	1600	2362	2361	15123
Number with lane excursion	2482	2494	877	903	1636	1615	10007
Percentage with lane excursion	68.9	69.3	54.8	56.4	69.3	68.4	66.2

Figure 8.1 shows a typical example of the ITTLC sampling used in simulation. Here the bins in the ranges (-0.5, -0.4) and (0.4, 0.5) both show a degree of oversampling that was not explicitly part of the initialization scheme. In fact the time to lane crossing at the start of the simulation was often reduced (so ITTLC was increased) due to a smaller lane-width between the simulation and the naturalistic condition. Although largely random, the effect was to increase the number of instances for which ITTLC exceeded 0.4 (and in some cases exceeded 0.5 – the left and right-most bins shown in Figure 8.1 include some data from outside the nominal range ±0.5 s^{-1}). This unplanned increase in initial ITTLC was not considered serious enough to justify re-running the entire set of batch simulations.

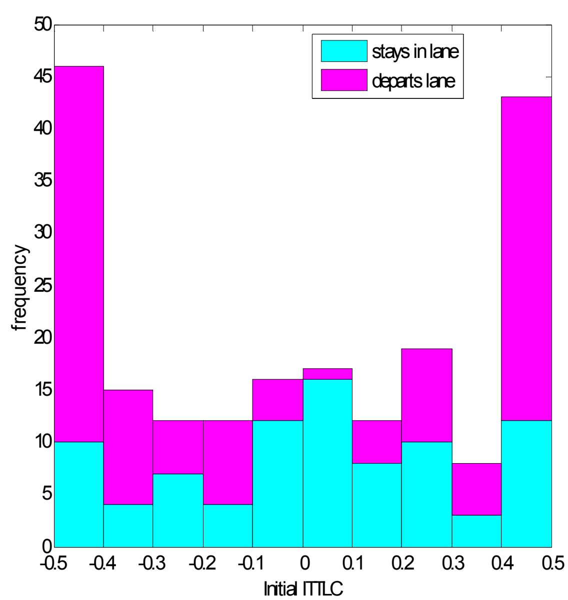

Figure 8.1. Distribution of Initial ITTLC Values (Scenario 1, four lane divided highway, rural locations, right travel lane). Note the extreme bins actually cover a wider ITTLC range than shown

Figure 8.2 presents distributions of crash metrics from Scenario 1 (straight road, 4 lane rural highway, divided). Each sub-plot corresponds to one component of the crash metric defined in Section 7. At this stage we have not established any calibration to estimate the scale factors (k_a, k_b, k_c, k_d) so here the individual components are shown with unit scaling. Each sub-plot has two distributions, the solid line corresponding to when the subject vehicle is in the right lane, and the dashed line for the left lane. Only events with significant lane excursions are included – those that do not generate any lane excursion metric at all are excluded. Paved shoulder distributions are quite similar, as might be expected, as are the crash metrics for the adjacent lane in the same direction. Neither lane generates any crashes for opposing traffic (given the condition that the highway is divided such an outcome is not possible in the simulated environment) and for road departure (which can be left or right side) again the frequencies are comparable. Care should be taken not to compare the on-road crash metrics (upper plots) with the off highway plot, since the definitions are quite different and at this point there has been no calibration against real crash data.

Figure 8.3, corresponding to a 4 lane undivided rural highway is similar, except that collisions in the opposing travel lane are much more frequent when the initial travel lane is on the left, as would be expected for crashes arising from inadvertent lateral drift.

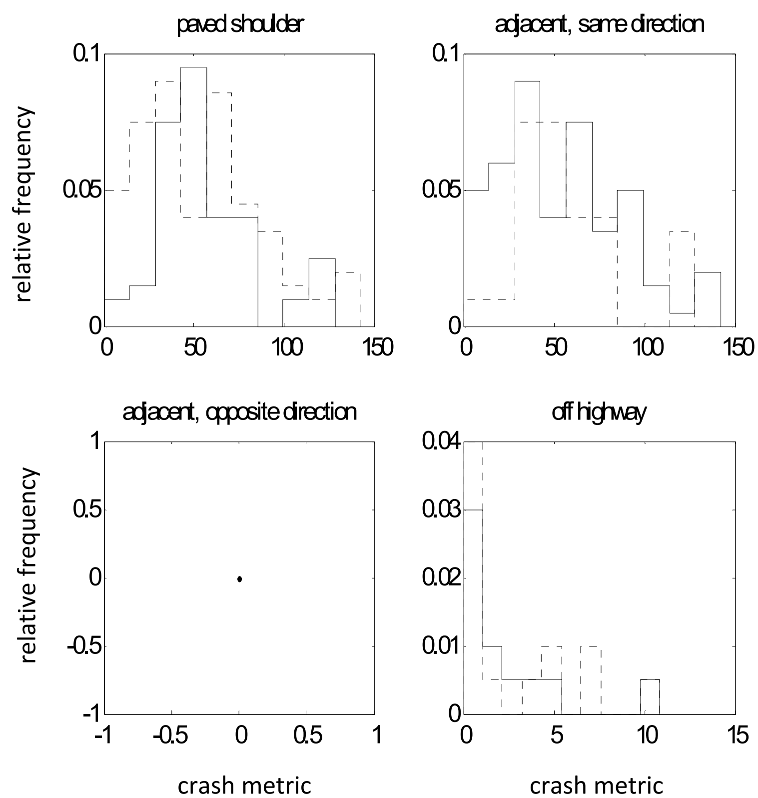

Figure 8.2. Relative Frequencies of Crash Metric Components for Scenario 1 (Rural case, solid=right travel lane, dashed=left lane). Note that crashes with opposing traffic are excluded in this scenario.

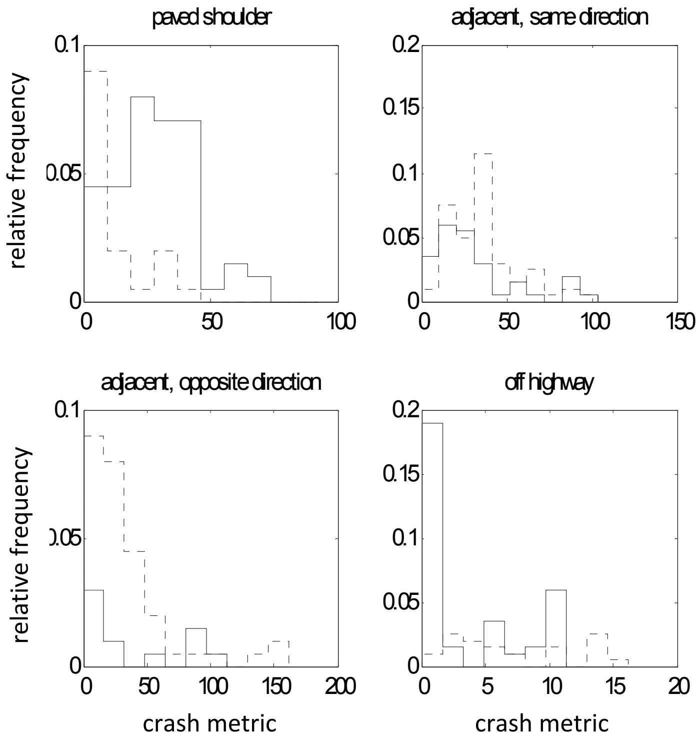

Figure 8.3. Relative Frequencies of Crash Metric Components for Scenario 2 (Rural case, solid=right travel lane, dashed=left lane). Crashes with opposing traffic are expected since the highway is undivided in this scenario.

Having briefly reviewed the types of crashes occurring in the virtual driving population, we now consider this same population alongside the LDW-enabled population.

8.2. *Comparison of Trajectories and Crash Metrics for With-LDW and Without-LDW Simulations*

Here again we restrict attention to illustrative examples and defer the analysis of the full population until Section 9, at which point the scenario weights and crash metric scale factors will be established. Figure 8.4 is again taken from Scenario 1 (right travel lane) and again we consider rural conditions (typically a rural freeway or other limited access highway). The trajectories are shown in a light color so that lane boundaries and other features may be seen. 200 trajectories are shown in each sub-plot. The upper plot is the case with LDW unavailable and there is a significant dispersion from the lane (recall, we have amplified the lane departure risk to increase the efficiency of estimating benefits). The trajectory marker is based at the center of the front axle of the vehicle, so by the time the curve meets the lane boundary, the front wheels are symmetrical about the lane marker.

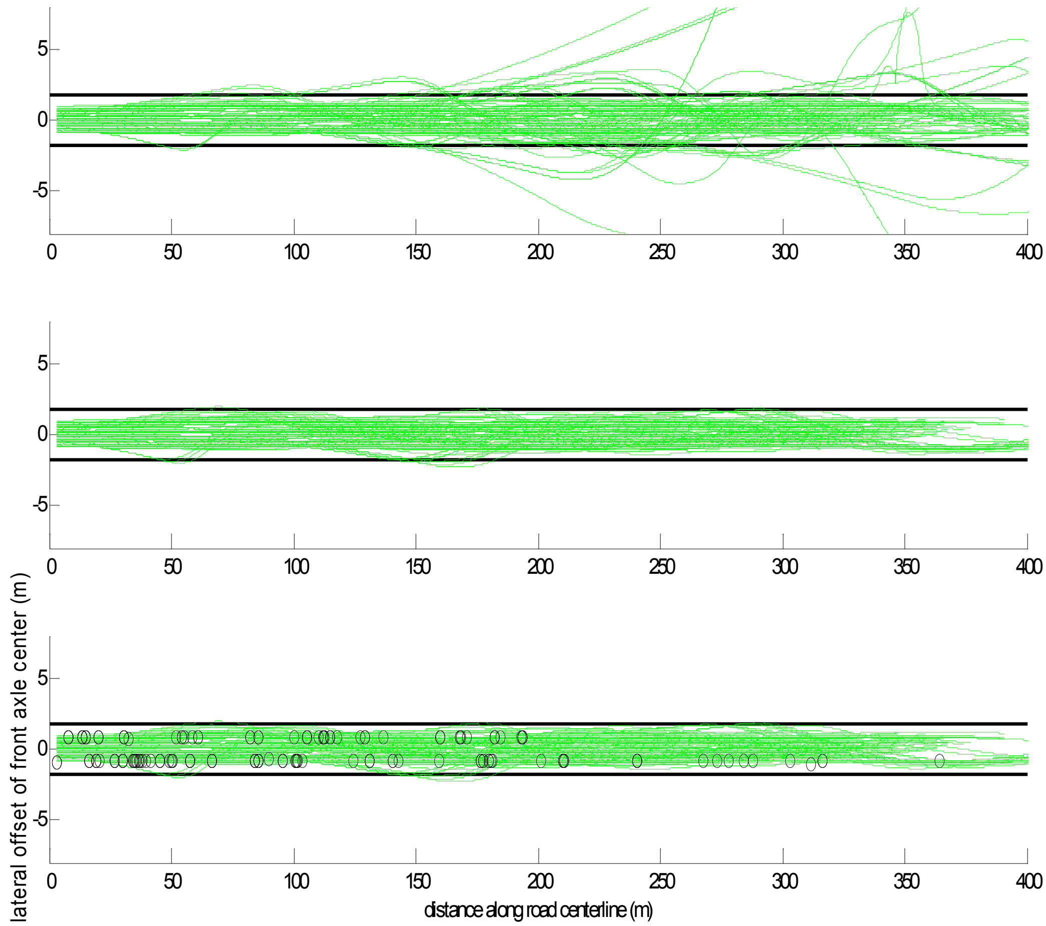

Figure 8.4. Effect of LDW on Lane-keeping. Trajectories shown for Scenario 1, Rural, left travel lane(upper plot: no LDW, center plot with LDW, lower plot also shows the location where the LDW alert is issued)

Even in the upper plot the driver is often capable of recovering directional control and returning to the lane. In a few cases there is a severe steering correction and a loss of lateral stability (recall, the simulated subject vehicle is not fitted with an electronic stability control system). A much reduced dispersion is seen for the LDW-enabled case (center) and it is seen that most alerts were issued during the first 200 meters, not surprisingly since it is the initial conditions that are primarily responsible for the elevated risk of lane departure. For reference, the mean speed for these trajectories is 32 m/s, or approximately 72 mph.

The crash metric has been specifically formulated to quantify the magnitude of the effect of the kind of trajectory change seen in Figure 8.4. For the comparisons to be valid, a condition was imposed that the LDW-enabled and LDW-disabled trajectories were compared over the same time intervals; if the simulation times are not the same, the longer simulation is truncated to match the end time of the shorter simulation, and this is done before the crash metrics are calculated.

In Section 4.1, Figure 4.1 shows an illustrative representation of how the ACAT system is expected to reduce the crash metric within any particular scenario. We are now in a position to present some specific examples based on the computational modeling and simulation. The plots have a similar form to those shown above comparing left and right travel lanes. Here we restrict attention to the right lane only and again consider Scenario 1 (Figure 8.5) and Scenario 2 (Figure 8.6) both for the rural case. The dashed line (with LDW) is almost uniformly lower than the solid line (without LDW).

In the next section we give attention to two specific examples, informally addressing the question: are these simulations “reasonably realistic” and hence worthy of taking forward to generate a safety benefits estimate.

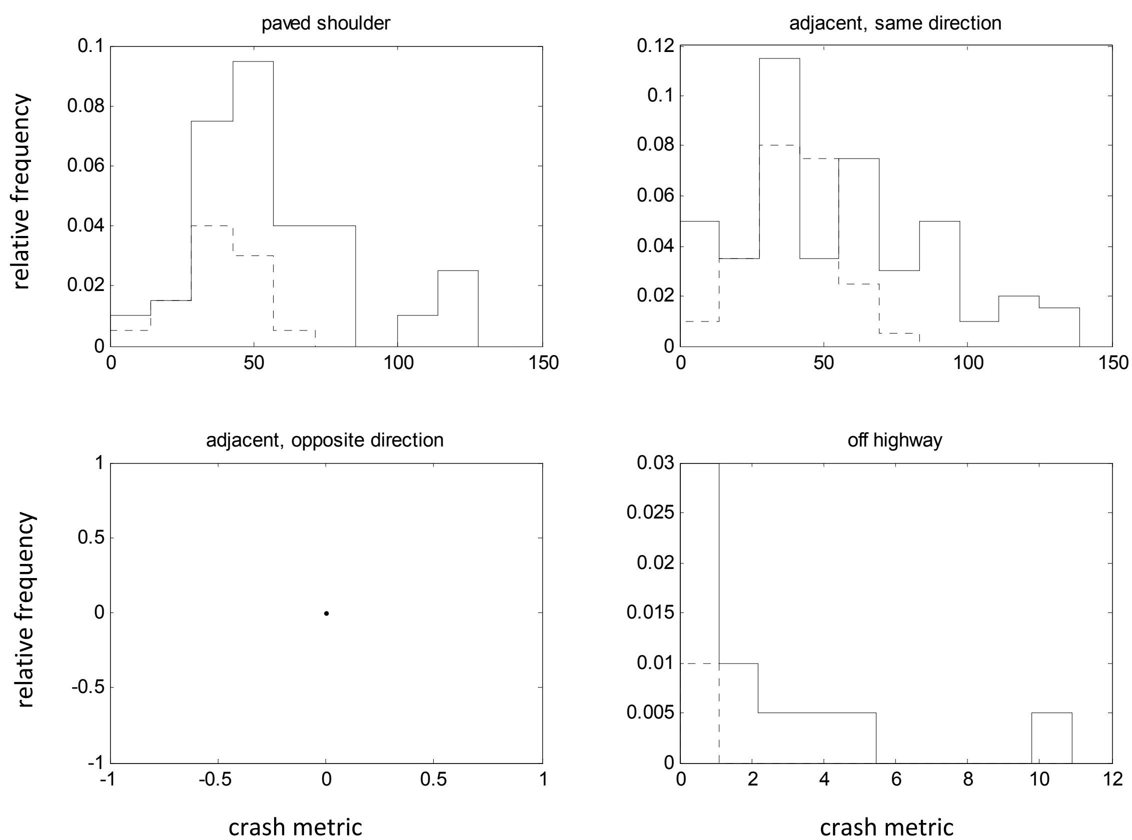

Figure 8.5 Effect of LDW on Frequency Distributions for Component Crash Metrics (Right lane, Scenario 1, Rural). Solid: without LDW, dashed: with LDW.

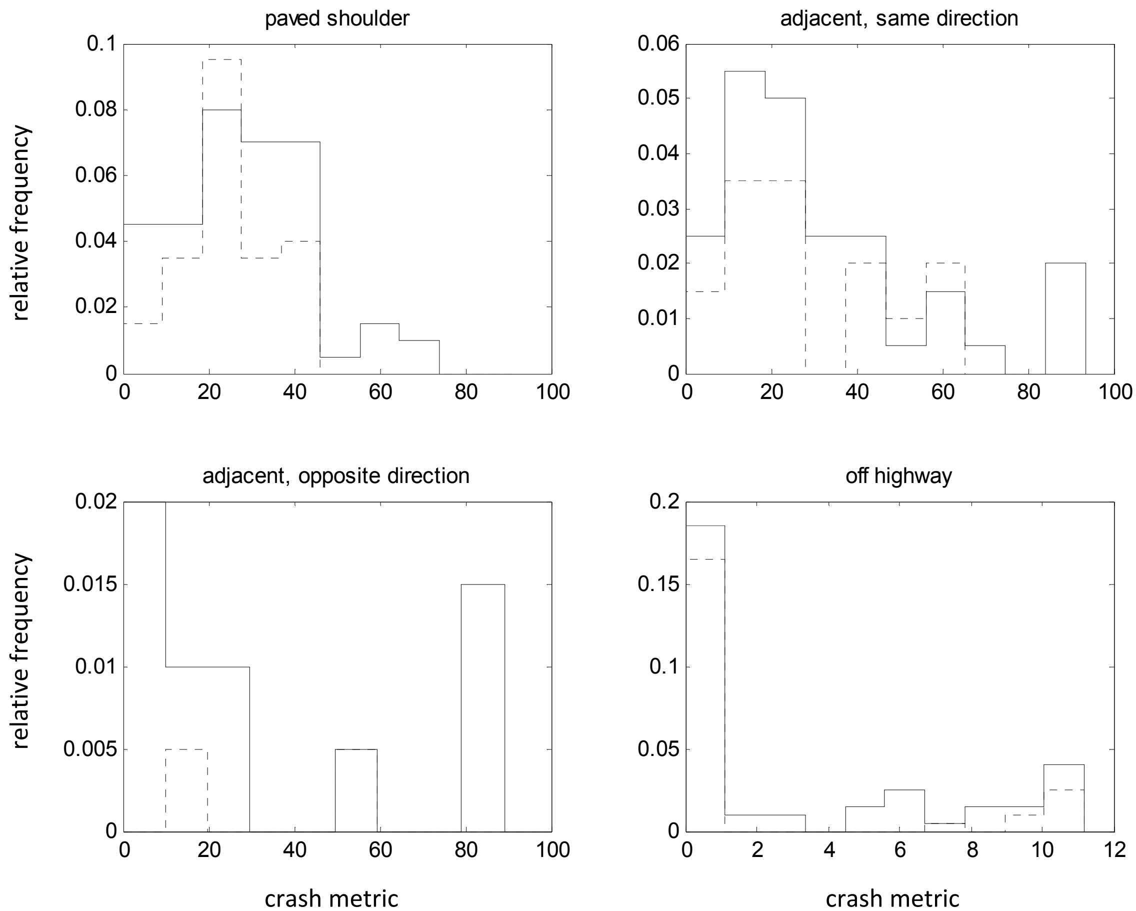

Figure 8.6 Effect of LDW on Frequency Distributions for Component Crash Metrics (Right lane, Scenario 2, Rural). Solid: without LDW, dashed: with LDW.

8.3. *Case Studies*

A small number of sample cases were extracted from CDS for comparison with SIM batch processing results. CDS cases belonging to driving scenarios 1 and 10 from the crash events table were selected for comparison purposes. These driving scenarios are basically identical in all respects, except for the fact that DS 1 is for straight roads, and DS 10 is for curved roads. The intent of this exercise is not to do crash reconstruction, but rather to illustrate the point that the SIM approach is fairly comprehensive and is able to capture a broad spectrum of single vehicle road departure crash events with sufficient clarity for benefits estimation.

8.3.1. CDS Crash for Driving Scenario 1 - Sample Comparison

Table 8.2. Example CDS Case for Driving Scenario 1

Case	2006-073-044
Researcher's narrative	This crash occurred during early morning daylight hours on a clear day on a rural, dry, straight, level, bituminous roadway with 2 northbound lanes divided from the 2 southbound lanes by a depressed grass median. Off the right edge of the road is a narrow level gravel shoulder, then a grass foreslope that goes down into a farmer's field. Vehicle 1, a 1995 Buick Regal occupied by the driver and five passengers, was northbound in lane one, departed the roadway to the right, braked and entered a clockwise yaw, travelled down the embankment, tripped over, rolled 4 quarter turns and came to rest on its wheels. V1 was equipped with dual frontal airbags, which did not deploy. V1 was towed from the crash scene due to damage. The restrained driver received minor laceration but was not treated. The unrestrained middle front seat passenger was not injured. The restrained right front seat passenger was treated and released with minor abrasion. The unrestrained left second seat passenger, who was completely ejected, received minor abrasion and contusion but was not treated. The unrestrained middle second seat passenger, who was completely ejected and found under rear tire, was fatal in ER with internal chest and abdominal injury. The restrained left second seat passenger was treated and released with minor abrasion and contusion.
Additional information:	Posted limit is 55 mph. Travel speed unknown. Driver distractions coded unknown. Attempted to steer back to left once off the road.

Figure 8.7 shows a CDS file diagram of the crash. Note that for straight roads, HPMS road segments were not used since there is no need to estimate road curvature. Figures 8.8 and 8.9 show the approach tracks from the road and off-road. It is apparent that there is a very small shoulder and the road edge dips down rather sharply past the shoulder.

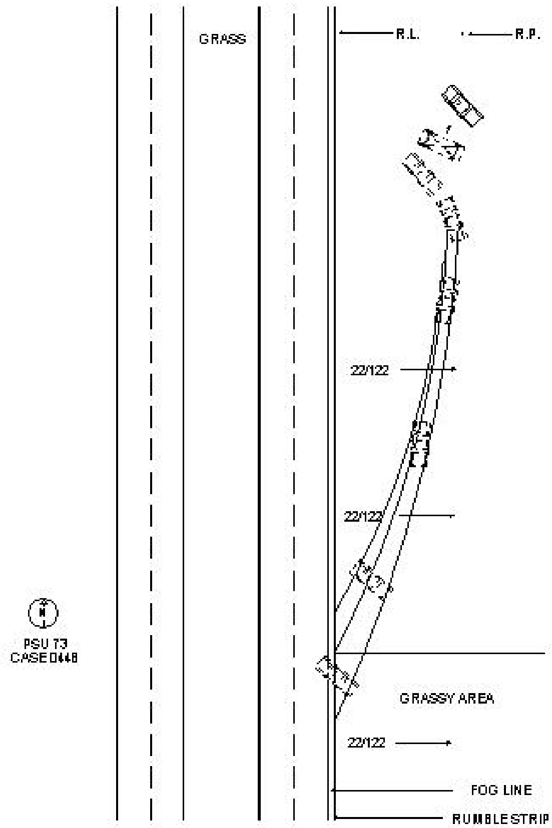

Figure 8.7. Crash Event Diagram for CDS Case 2006-073-044 (DS 1)

For comparison with simulated data, the initial conditions for the run were selected to be relatively close to the researcher's description, road geometry and vehicle longitudinal speed. There were a couple of simulation runs that appeared to closely match the CDS crash description. For example purposes, one of the runs: Run_118_S_1_E_45_R_0_0_C_5632694_Alt_1_LDW_0 has been selected for comparison. The sample run has an ITTLC of 0.358 s^{-1} which corresponds to a vehicle veering to the right and the initial velocity of the vehicle was 58 mph. A variation of this run with LDW ON was also evaluated.

Figure 8.8. Approach Tracks for CDS Case 2006-073-044 (DS 1)

Figure 8.9. Approach Offroad for CDS Case 2006-073-044 (DS 1)

Figures 8.10-8.11 show the results from the batch simulations for the case with LDW OFF. Figure 8.10 shows the trajectory in road coordinates, while Figure 8.11 shows the attention switching by the driver as a function of time. Figure 8.10 shows (approximately) where the vehicle departed the road, and in conjunction with the off-road features that are shown in Figures 8.8, 8.9, it is clear that the vehicle could not return back into the lane. As indicated earlier, the point of this comparison is not to do crash reconstruction, but to illustrate that real world crashes are (and can be) captured quite well using the SIM approach. In terms of the simulation metrics, these are computed the first time the vehicle leaves the road, and subsequent simulation maneuvers that might show the vehicle back in the lane are not considered primarily due to the absence of off-roadway features in the SIM model and their impact on the vehicle trajectory.

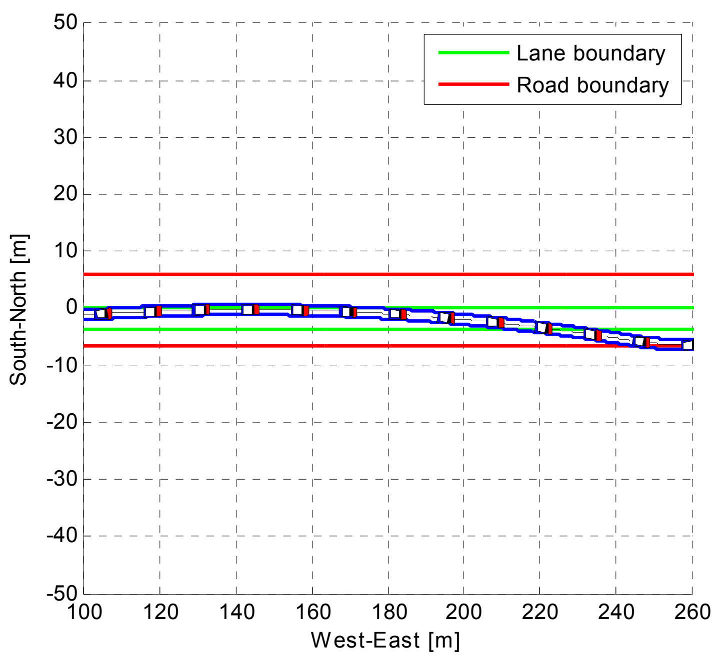

Figure 8.10. Simulation Run for DS # 1 Crash Event (LDW=OFF)

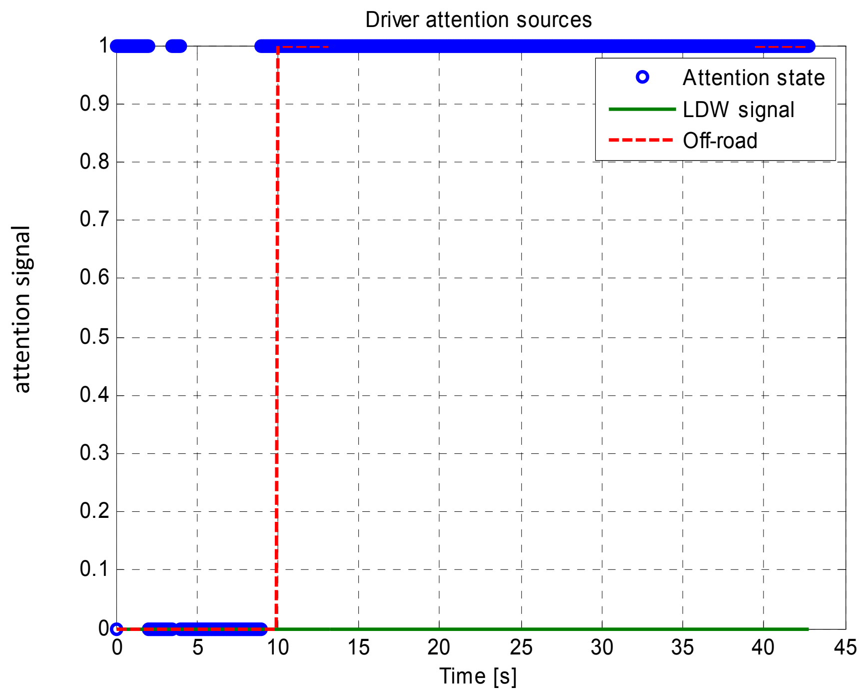

Figure 8.11. Driver Attention State for DS #1 Crash Event (LDW=OFF)

Figures 8.12-13 are the complementary set of figures for the same case with LDW=ON. It is clear from a comparison of Figures 8.11 and 8.13 that the driver attention switches back to the driving task earlier, enabling the vehicle to avoid the crash.

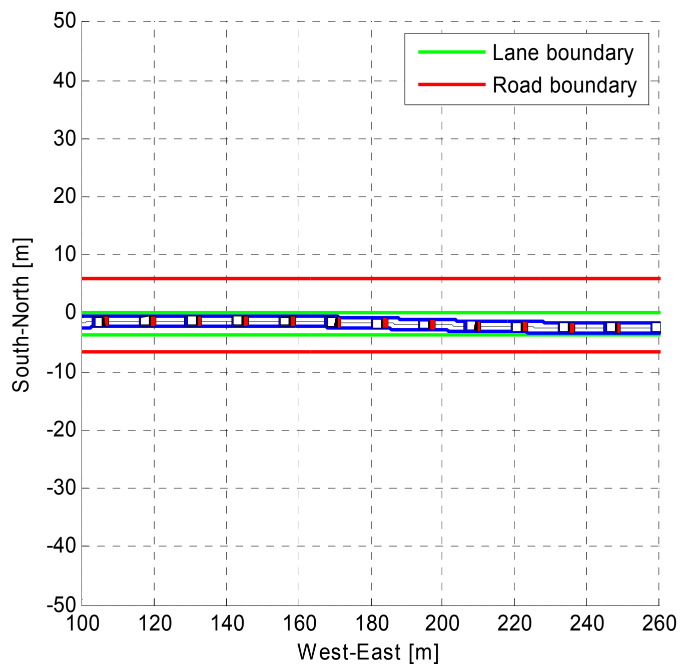

Figure 8.12. Simulation Run for DS #1 Crash Event (LDW=ON)

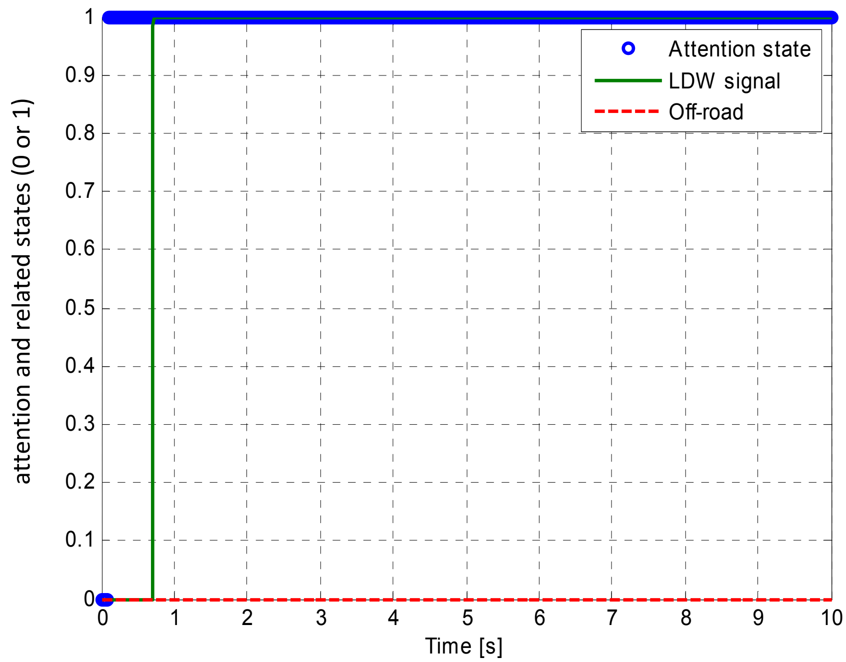

Figure 8.13. Driver Attention State for DS 1 Crash Event (LDW=ON)

8.3.2. CDS Crash for Driving Scenario 10 - Sample Comparison

Table 8.3. Example CDS Case for Driving Scenario 10

Case	2006-075-123
Researcher's narrative	Vehicle one (V1), a 1988 Acura Legend, was traveling in lane one of a three-lane, left-curve, level roadway. V1 traveled off the right side of the roadway and down a steep grass embankment. V1 rotated CW for approximately 13 meters. V1 then rolled three quarter-turns and and [sic] the left side struck a tree. V1 came to final rest on its left side near the point of impact with the tree facing SW. V1 was towed due to damage. V1 restrained driver was hospitalized with brain injury and lower back fracture. The unrestrained right front passenger was treated and released with minor contusion.
Additional information:	Driver distraction coded unknown. Travel speed 65 mph. Posted speed limit 65 mph.

Figure 8.14 shows a CDS file diagram of the crash. Note that the crash diagram and the researcher's narrative are descriptions of the crash and do not capture the trajectory or driver actions prior to the event. In the SIM computational model, the simulation will start prior to initiation of the conflict, so the initial maneuvers are not captured in the CDS narrative. As in the previous case for Driving Scenario 1, the intent is not to do crash reconstruction, but rather to

illustrate that the SIM approach captures a suitably broad spectrum of single-vehicle road departure crashes.

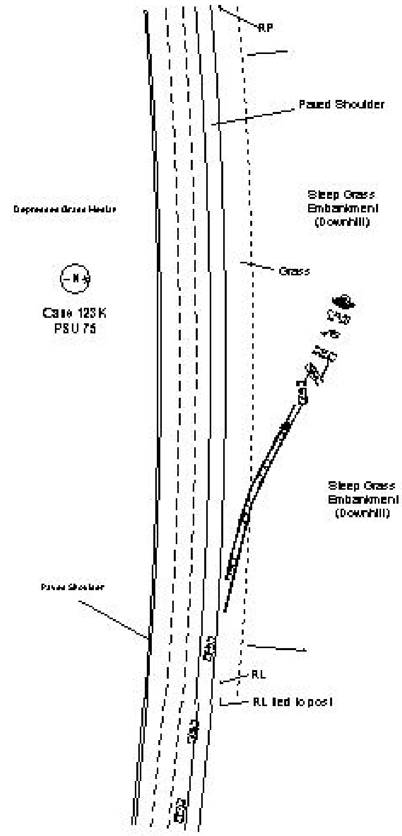

Figure 8.14. Crash Event Diagram for CDS Case 2006-075-123 (DS 10)

This case was compared with one of the batch runs from the SIM computational model. The initial conditions for the run were selected to be as close as possible to the researcher's description, road geometry and vehicle longitudinal speed. The run that appeared to match the CDS crash description the closest was: Run_65_S_10_E_48_R_3_9_C_5358381_Alt_1_LDW_0. The sample run has an ITTLC of 0.496 s^{-1} which corresponds to a vehicle veering to the right and the initial velocity of the vehicle was 75 mph. Figure 8.15 shows the road that was selected as one that was the closest in terms of approximating the crash site. A variation of this run with LDW ON was also evaluated.

Figures 8.16-8.19 show the results from the batch simulations for the case with LDW OFF and with LDW ON. Figures 8.16-8.17 show the simulation output for the vehicle trajectory with LDW OFF. Figure 8.16 shows the trajectory in road coordinates, while Figure 8.17 shows the attention switching by the driver as a function of time.

Figure 8.15. HPMS Road Segment Used for Simulation of CDS Case from DS 10

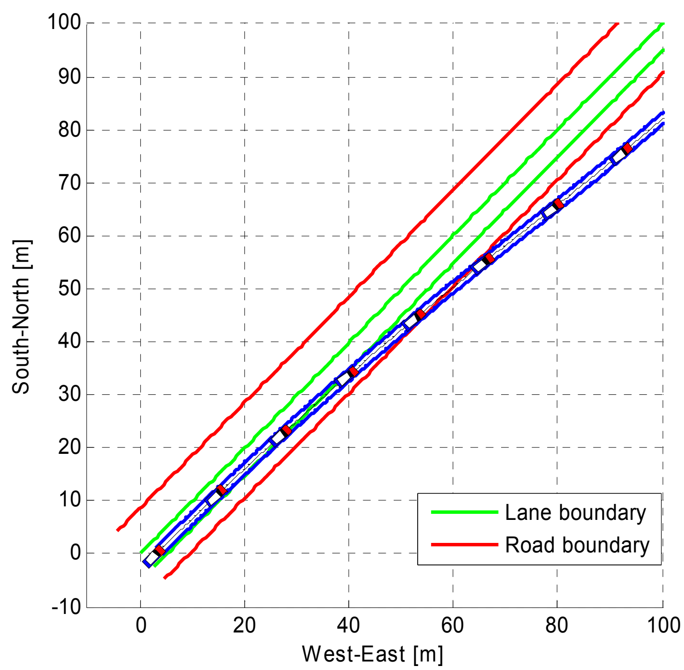

Figure 8.16. Simulation Run for DS #10 Crash Event (LDW=OFF)

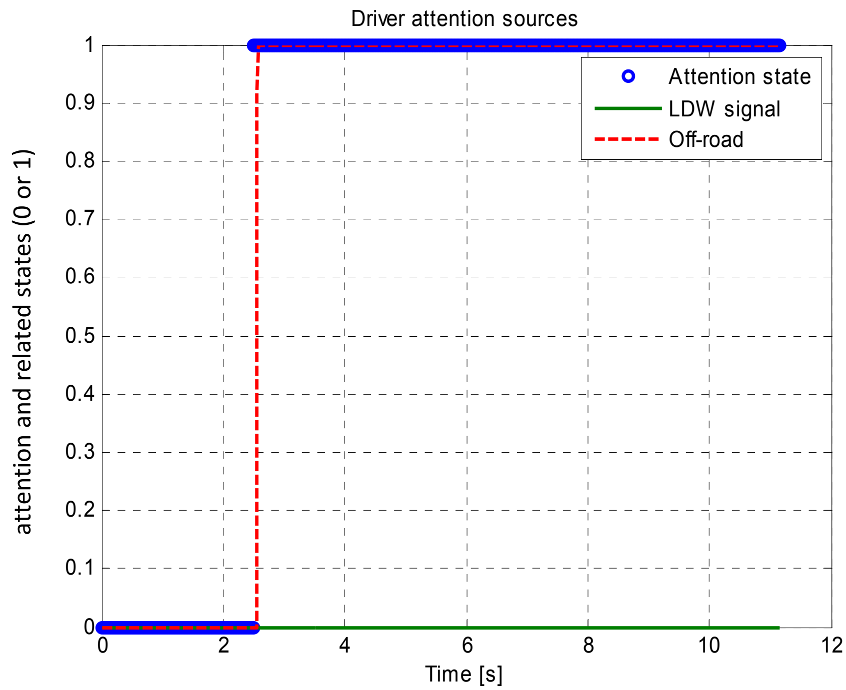

Figure 8.17. Driver Attention State for DS 10 Crash Event (LDW=OFF)

The simulation run indicates that the driver's attention shifted back to trying to control he vehicle too late in terms of the evolution of the crash dynamics. It is important to note that the simulation has several probabilistic and tunable parameters and random seeds which affect the initial driver attention to the driving task. In spite of this, it is quite clear that the SIM computational tool can successfully capture the spirit of the crash event and the vehicle trajectories at the coarse grain level are comparable.

Figures 8.18-8.19 show comparable results for the case with LDW=ON. Comparing Figures 8.16-8.17 with Figures 8.18-8.19, the SIM simulations predict that the driver attention would switch to the driving task earlier with the LDW warning and enable the vehicle to avoid a crash. Figure 8.18 shows that while the vehicle does leave the lane, it does not leave the road edge and the driver is able to steer the vehicle back within the lane boundaries.

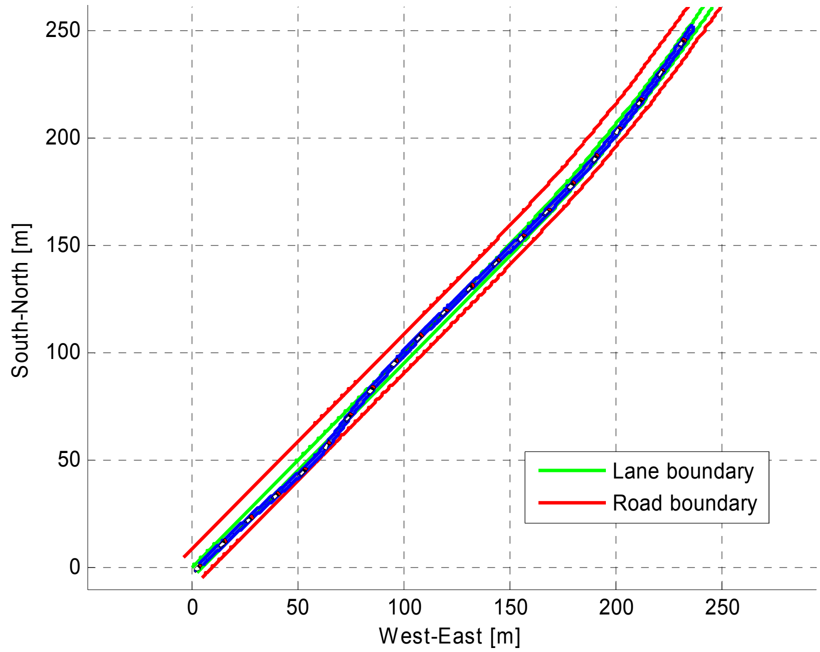

Figure 8.18. Simulation Run for DS 10 Crash Event (LDW=ON)

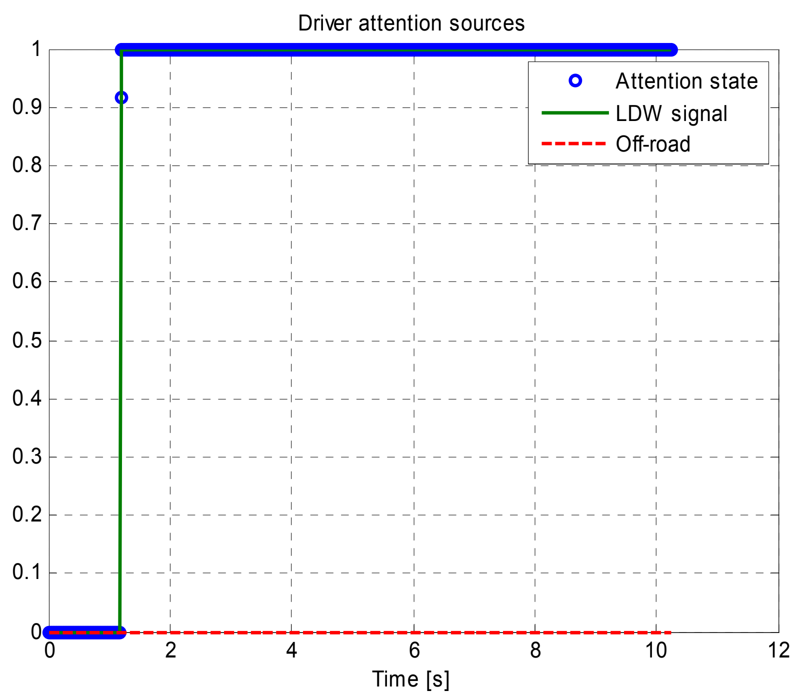

Figure 8.19. Driver Attention State for DS 10 Crash Event (LDW=ON)

9. Results of SIM Benefits Estimation

In the previous section we summarized the overall outcomes of the batch simulations in terms of the component metrics and presented some examples, with focus on the mechanisms at play and the effect of the LDW. In this section we present the results of the safety benefits analysis described in Section 2. The main difference is that we now include matching criteria to estimate scenario weights and crash metric scale factors. This is a fundamental step, taking the free parameters in the analysis and using them to create a best fit to the GES crash data. As we shall discuss, this process is not guaranteed to create a high degree of correlation between actual and virtual populations; if the simulation model does not capture the real crash mechanisms, or the GES codes do not capture the crashes of interest, then we can expect to see divergence between the populations. In this way the SIM benefits analysis of Section 2 incorporates intrinsic validation criteria.

Once the estimation of weights and scale factors has been implemented (Section 9.1) we are in a position to estimate crash reduction numbers resulting from the LDW technology, and also explore trends according to different driving scenarios. This is presented in Section 9.2, which provides the core set estimation results from safety benefits estimation based on batch simulations. The next two sections provide a number of variations on these results: the sensitivity of LDW benefits to driver reaction time when responding to an LDW alert – Section 9.3 – and the effects of system availability and driver compliance in Section 9.4. Further sensitivities are also considered in Section 9.4; the minimum operating speed for LDW, and the effect of any under-counting of fatigue in the GES crash data. In section 9.5 we consolidate the estimation results and provide approximate error bands for the estimates of LDW system effectiveness.

9.1. *Scenario Weights and the Virtual Crash Population*

Recall that in the analysis of virtual driving events there are four crash types arising from the simulations. In simulation these are based on trajectory deviations into the four zones of interest (paved shoulder, adjacent lane/same direction, adjacent lane/opposite direction, off highway) and correspond to the four components of the crash metric in Section 7.5. In term of GES crash data, the corresponding outcomes are set in Table_GES_13K (see Section 7.1) where a number of variables have been included to filter outcomes – see Table 9.1 below. Case weights were averaged over 5 years (2002-2006 inclusive) and are used for optimization. As well as the obvious scenario and Rural/Urban designations (for the preceding driving scenarios), the outcomes use the following criteria (see Table 9.1 for detail and references to the variables used):

- Metric Component 1: paved shoulder
  REL_RWY: 2, 7 or 8 (shoulder, parking lane or gore)
  EVENT1 is *not* 25 (*not* motor vehicle in transport)

- Metric Component 2: adjacent, same direction
  REL_RWY: 1 (on roadway)
  ACC_TYPE is *not* 50,52, 53 or 64 (*not* head-on or head-on sideswipe)
- Metric Component 3: adjacent, opposite direction
  REL_RWY: 1 (on roadway)
  ACC_TYPE is 50,52, 53 or 64 (head-on or head-on sideswipe)
- Metric Component 4: off highway
  REL_RWY: 3, 4, or 5 (median, roadside or outside trafficway)
  EVENT1 is *not* 25 (*not* motor vehicle in transport)

These outcome filters account for over 90% of all crashes in the designated scenarios (with the exception of one scenario: Scenario 24, urban region, where the proportion is 87%) with an overall mean of over 95% of ACAT relevant crashes. On this basis we are able to set up a "real" crash population for optimization and comparison against the virtual crash population.

Table 9.1. Variables Used to Filter Crash Outcomes

Variable Name (SAS code)	Description – for further details see the GES Analytical User's Manual (NHTSA, 2007a)	Page reference (NHTSA, 2007a)
ACC_TYPE	Accident Type (V23): Categorizes the pre-crash situation (via the relative position and motion of the vehicles immediately prior to the first harmful event)	87
REL_RWY	Relation to Roadway (A10): Indicates the location of the first harmful event.	37
EVENT1	First Harmful Event (A6NZ): Indicates the first property damaging or injury producing event in the crash.	27

To carry out the optimization, crash metric scale factors (k_1, k_2, k_3, k_4) are found by iterative optimization using the Matlab function *fminsearch* (Mathworks, 2009). In this process the Nelder-Mead Downhill Simplex method (Mathworks, 2009) is used to iteratively adjust the four scale factors to achieve best fit for the cost function

$$e = \frac{\sum_i \sum_j (n_{ji} - \overline{n}_{ji})^2}{\sum_i \sum_j n_{ji}^{\;2}} \tag{9.1}$$

which measures the deviation between the virtual crash numbers n_{ji} for outcome *j* and scenario *i*, and the corresponding GES number s $\overline{n}_{ji}$. Within each optimization step, once the k_i parameters are chosen, the crash metric function is fully defined, so the transition probabilities T_{ji} in equation (7.3) for may be evaluated. Then from equation (2.9) the normalized scenario weights $\overline{w}_i$ are found and hence virtual crash numbers

$$n_{ji} = T_{ji}\overline{w}_i N$$

are available. The overall optimization process is summarized in Figure 9.1.

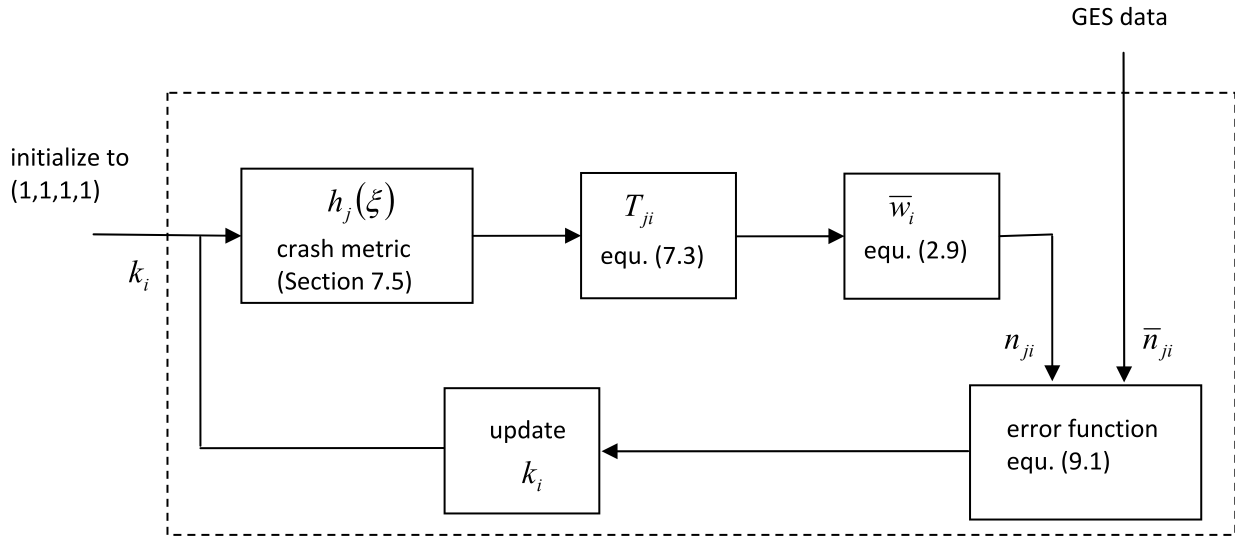

Figure 9.1. Optimization of Crash Metric Scale Factors and Scenario Weights

The optimization process is repeated for each of the 6 major road types (Type A Rural etc., as in Table 8.1) since we assume uniformity of crash metric scaling within these class types, but not between them. Note that since the overall optimization process includes a simple scale factor ambiguity (see Section 7.5) we set the maximum of the k_i to equal 1; in the above equations this is automatically linked to the overall scaling of the scenario weights $\overline{w}_i$.

The results of the optimization are now shown. Table 9.2 summarizes the fitted crash metric scale factors. According to this table, for paved shoulder (k_1) it is only the multi-lane highways (especially type A rural and type B urban) that appear to show a significant crash probability per unit distance travelled in the "wrong lane". Off highway (k_4) always ranks highly, and for type C (two lane) highways the off-highway component appears to dominate. However, care should be used in interpreting comparisons between k_4 and other modes, since k_4 is a scale factor for lateral distance while the others related to longitudinal distance.

Table 9.2. Crash Metric Scale Factors.

(*Type A: 4 lane divided, Type B: 4 lane undivided, Type C: 2 lane undivided*)

Major Road Type	Paved shoulder k_1	Adjacent lane same direction k_2	Adjacent lane opposite dirn. k_3	Off highway k_4
typeA_rural	0.871	1	0.988	0.679
typeA_urban	0.016	0.163	1	0.511
typeB_rural	0.089	0.128	0.127	1
typeB_urban	0.394	0.260	0.251	1
typeC_rural	0.004	0	0.024	1
typeC_urban	0.038	0	0.057	1

The crash numbers in the virtual population provide a simpler basis for the evaluation of the optimization results. These numbers include the effects of both the crash metric coefficients and the scenario weights, and provide a test of reasonableness by comparison with the GES population. The results are shown in Figure 9.2, where the light colored bars represent real-world data and the dark bars represent virtual data. Overall the trends in crash location are very similar, though there are certainly some differences, especially in the type A rural population (multi-lane divided highway – typically rural freeways and other limited access highways). Here the virtual population has too many first harmful events predicted on the paved shoulders, and not enough off the highway. This appears to show that, for this particular set of scenarios, simulations do not properly represent the real world pattern of trajectory and crash outcomes. While there are many potential sources of discrepancy, the most likely limitations are (a) over-simplified “2 lane” assumption, (b) no cross-slope or drop-off are modeled in unpaved area, and (c) approximate estimation of clear zone distances based on AASHTO design guidelines . All of these factors can be addressed, though to do so will require a depth of data collection or additional analysis that is beyond the scope of this initial study. Overall the match across major road types appears very reasonable, especially given the fact that these virtual population statistics are not based on any type of iteration and improvement, outside of the formal optimization of metric scale factors and scenario weights.

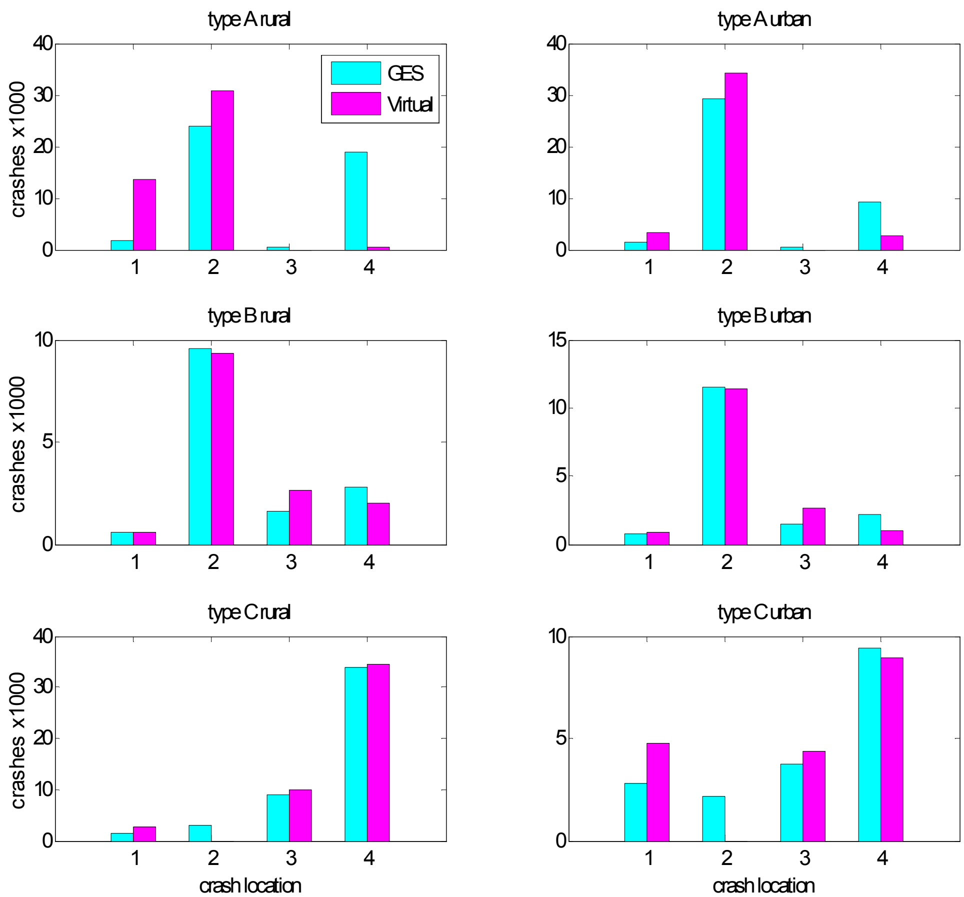

Figure 9.2. Comparison of GES and Estimated Crash Numbers According to Major Road Type.

A simple measure of the fit between the virtual and GES crash populations is provided by the following measure, derived from the cost function in equation (9.1):

$$P = \left[1 - \frac{\sum_{j,i} (n_{ji} - \bar{n}_{ji})^2}{\sum_{j,i} n_{ji}^{\;2}} \right] \times 100 \tag{9.2}$$

P represents the percentage of variation in the actual population explained by the virtual crash population, broken down by each individual scenario (there are 50 of these including the usual urban/rural split). For the data processed in this project we find *P* = 85, i.e. 85% of the variation in the population is explained. In making this comparison, recall there are 200 cells to fit (4 outcomes and 50 known scenarios) and a much smaller number of parameters being estimated: there are 6×3=18 independent *k* values in the crash metric (6 major road types and 3 free scale factors after setting $\max(k_i) = 1$), and 76 scenario weights, giving 94 parameters in all. The simulation model provides the remaining 100 plus degrees of freedom.

Other aspects of the virtual population are available for inspection to test overall reasonableness. In Figure 9.3 scenario weights are presented in an aggregated form via major road type. Here there is no specific real-world data for direct comparison and validation; the pre-conditions are modeled to be proportional to the driving scenario frequencies of the affected drivers, not necessarily to the whole driving population. Where there is more than 1 travel lane the numbers are mostly evenly distributed between the two lanes, which seems reasonable.

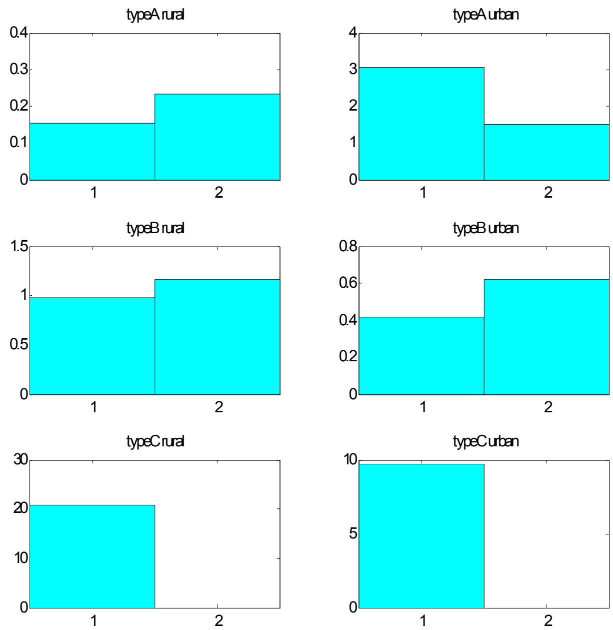

Figure 9.3. Scenario Weights by Road Type and Travel Lane (1=right lane, 2=left lane)

At this point we have defined a virtual crash population with properties that are similar to the actual crash population, in so far as crash numbers and crash types agree within some plausible degree of error. Provided we accept the hypothesis that the virtual crash population is a reasonable surrogate for the actual population we are in a position to develop an estimate of the potential safety benefits of the ACAT technology.

9.2. *LDW System Performance in the Virtual Crash Population*

First we consider overall estimated crash numbers by major road type – Figure 9.4. In every case the effect of LDW is to reduce the estimated crash numbers; this is certainly expected since the system triggers an earlier reaction from the driver or, if the LDW alert occurs too late to generate a useful response we assume its effect is at worst neutral. Perhaps the most striking result is the relative uniformity across road types and outcome types; numbers are projected to be reduced by around 50% in all cases. The corresponding crash numbers are shown in Table 9.3. Thus the initial raw estimate of overall "system effectiveness", i.e. the proportion of crashes reduced in total (see Section 2) turns out to be 47%. At this stage the raw estimates have not included several important effects, especially system availability and driver acceptance and compliance, which they need to be adjusted for. When these factors are considered later on in the section, more plausible estimates result.

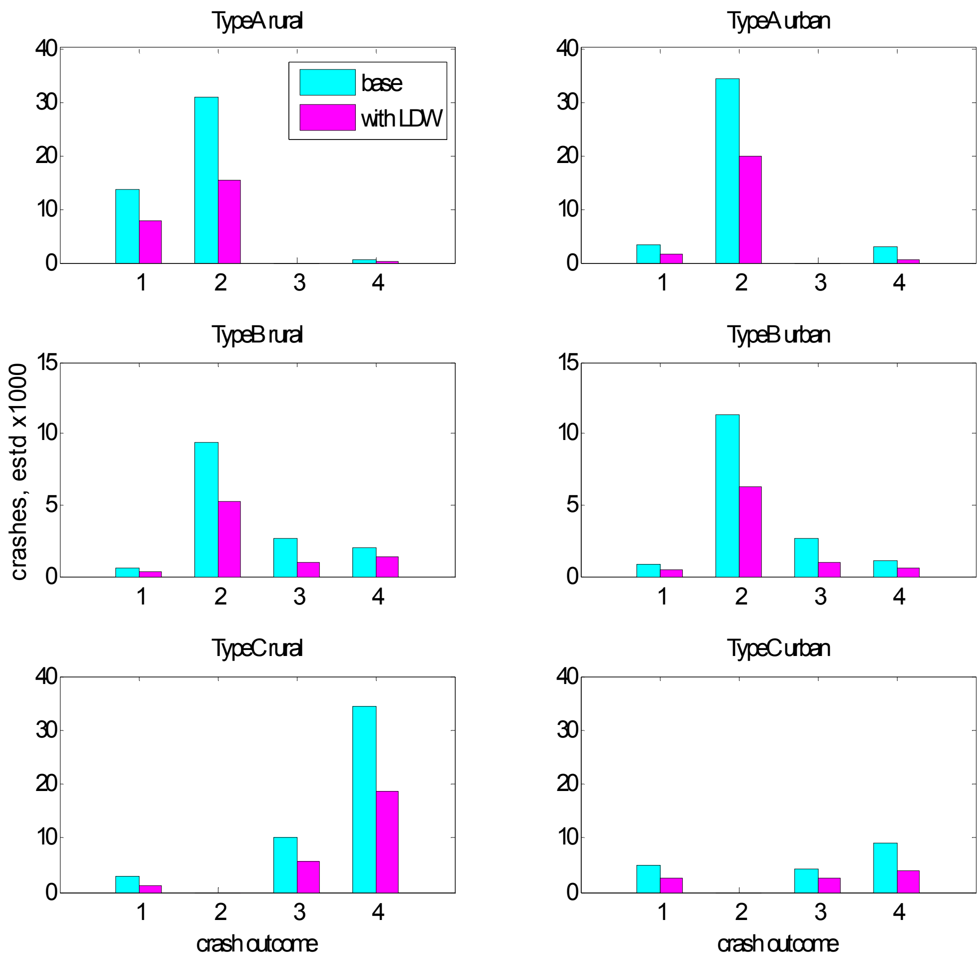

Figure 9.4. LDW Effect on Estimated Crash Numbers by Major Road Type (crash numbers in 1000's according to crash outcome: 1: paved shoulder, 2: adjacent lane – same direction, 3: adjacent lane – opposite direction, 4: off highway)

Table 9.3. Crash numbers (in thousands) aggregated by road type and outcome.
[*Type A: 4 lane divided, Type B: 4 lane undivided, Type C: 2 lane undivided*]

Estimated Crash Nos. (thousands)	Type A Rural	Type A Urban	Type B Rural	Type B Urban	Type C Rural	Type C Urban	Total
Base	44.9	40.3	14.5	15.9	47.2	18.1	181
With LDW	23.1	22.1	8.0	8.3	25.3	9.2	96
Difference	21.8	18.2	6.5	7.6	21.9	8.9	85

This baseline number is apparently quite large, indicating that in a perfect world the technology would eliminate nearly half of the crashes arising from drift out of lane; for the vehicle type being considered (automobiles – Table 4.1) this would amount to 85,000 crashes avoided annually. In the next section we consider the key real-world factors expected to reduce this effectiveness figure, but overall the picture seems both simple and positive.

The relative uniformity seen in Figure 9.4 does not necessarily mean that the LDW system is predicted to be equally effective under all circumstances. In Figure 9.5 we break down the effectiveness (relative reduction in crash numbers) according to the 25 major scenarios of Table 4.6, split by urban and rural roadways. There seems to be a broad trend towards greater effectiveness in rural conditions (only 8 scenarios have urban effectiveness higher than rural) and the "high scoring" scenarios 11, 15, 17 all have rural at the higher level. In fact, scenarios 11 and 15 both correspond to fatigued drivers, and this is not especially surprising given the analysis of Section 6. Figure 9.6 provides the same data, but this time in the form of crash numbers (estimated benefits by scenario identifier). Since the lower scenario identifiers correspond to higher crash numbers, the benefits scores tend to be higher there also.

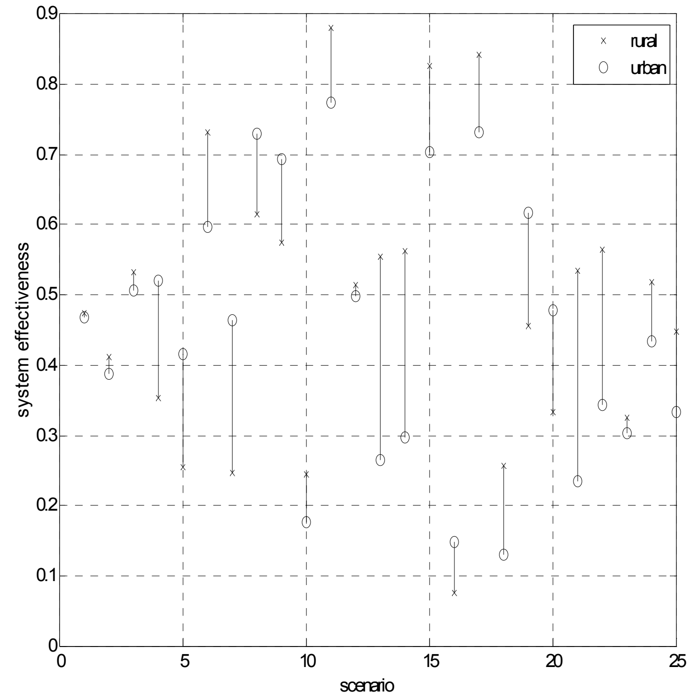

Figure 9.5. System effectiveness by scenario identification (Table 4.6) including urban/rural differences.

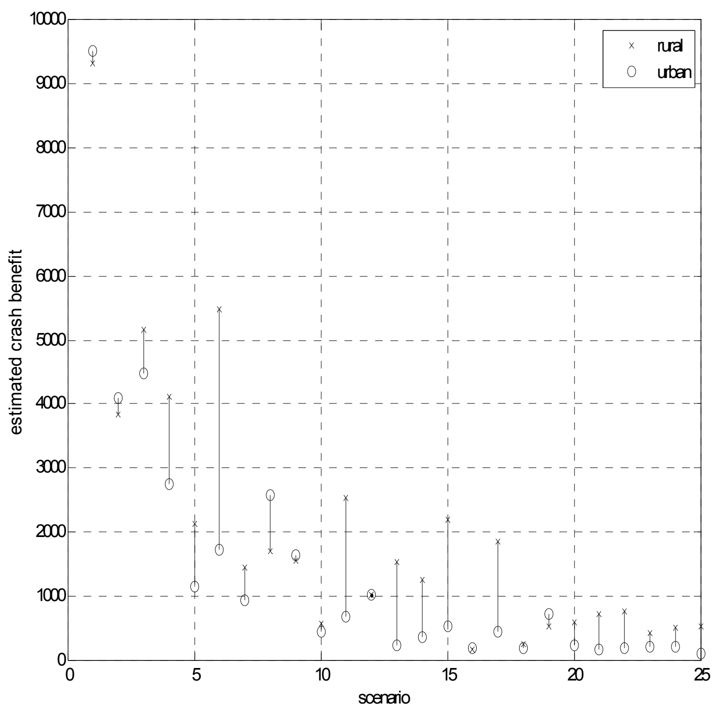

Figure 9.6. Estimated crash number benefits by scenario identification including urban/rural differences.

Again a possible trend towards greater benefits in rural areas seems to exist. To test this we plot effectiveness and benefits estimates according to urban and rural road types – Figure 9.7. Here it seems clear that overall the effectiveness is quite uniform between urban and rural cases, but that the greater numbers of rural crashes (particularly for Type C rural compared to Type C urban) concentrates the benefits in these cases; the apparent trend is therefore explained as a simple exposure effect.

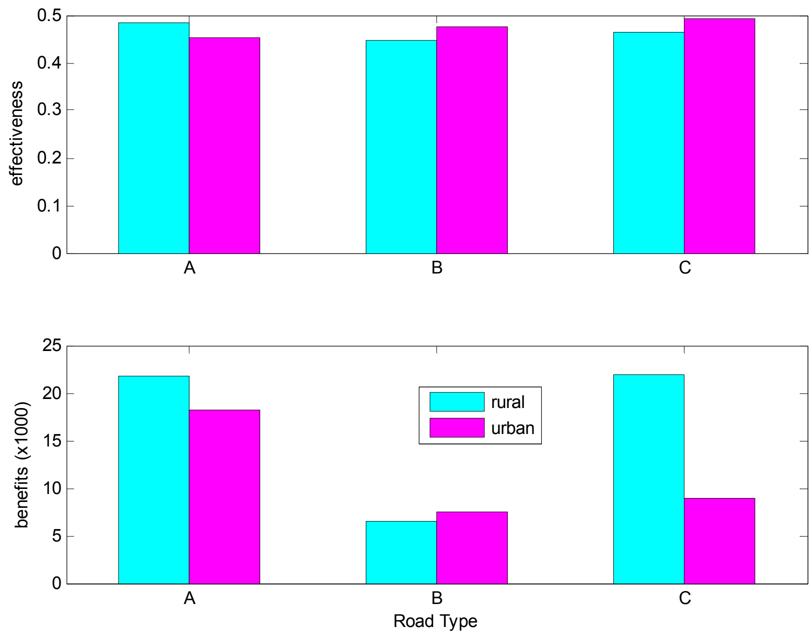

Figure 9.7. Effectiveness and Benefits by Major Road Type.

(*Type A: 4 lane divided, Type B: 4 lane undivided, Type C: 2 lane undivided*)

Other factors can be used to distinguish between the different scenarios in Table 4.6; the effects of a number of "adversity factors" (curved road, nighttime driving, fatigued driver and wet road) are presented in Figure 9.8. In these plots we use the label "1" to denote the presence of the condition labeled underneath (e.g. curve=1 denotes curved roads while 0 denotes non-curved, i.e. straight roads). Thus in the upper left plot we see curved road effectiveness represented at 0.3 compared to 0.5 straight roads. Overall we see the SIM analysis predicts system effectiveness is higher under the adverse conditions with the exception of *curved road*, where the LDW effectiveness is reduced. One can easily speculate on the reasons behind this. For the fatigued driver an early response is more critical to avoiding a crash, and the LDW alert is potentially more effective in these circumstances. But for curved roads the effect of drift will tend to translate into a significant lane departure more rapidly than for a straight road. Recall, lateral curvature is not presumed constant in this study; real road geometries linked to actual crashes are used for the simulations, so the time delay in responding to the LDW warning may be expected to degrade the overall LDW effectiveness in this case. The results in Figure 9.8 do not prove these assertions, but the results are certainly reasonable. It is feasible to formulate experimental designs (e.g. driving simulator or data mining from naturalistic driving databases) to provide independent verification of these suggested mechanisms, but again this is outside of the scope of the current study.

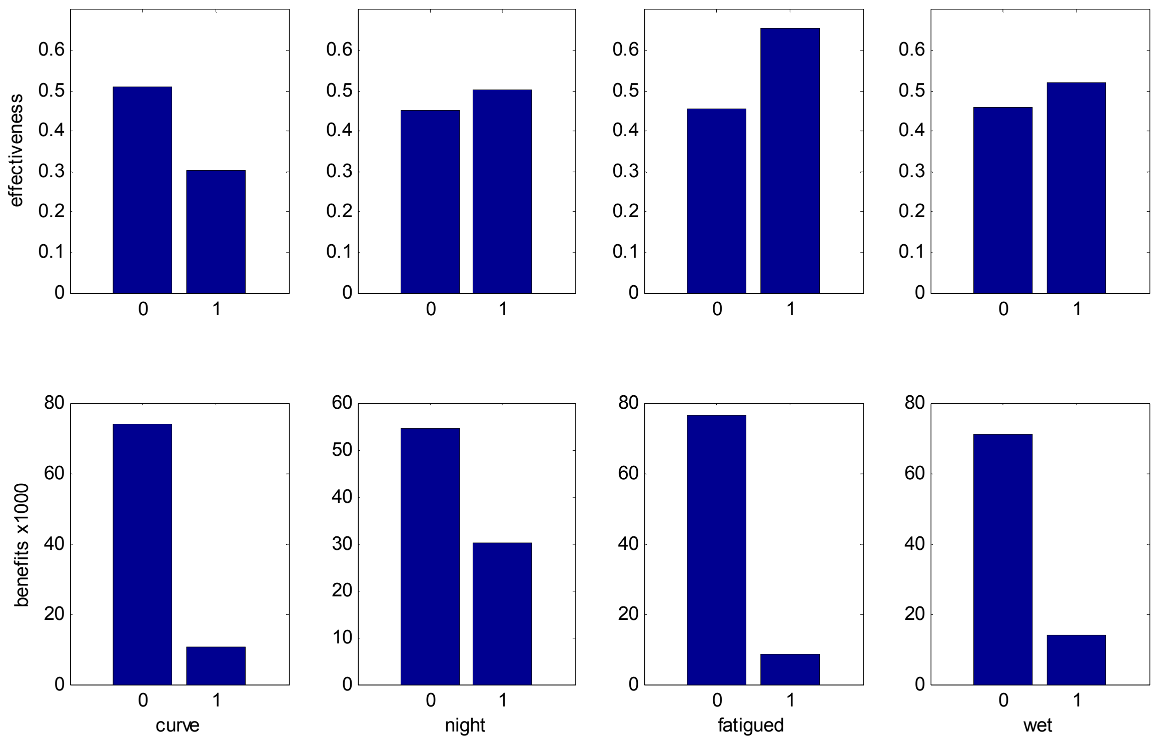

Figure 9.8. Estimates of Effectiveness and Benefits According to Adversity Factors

Finally in this section we consider the effect of speed on LDW benefits and effectiveness. Recall that speed was not used as an independent control variable for the driving scenarios that make up the virtual driving population. So, to determine speed effects we partition the driving events by initial speed and count crash numbers within 2.5 m/s (5.6 mph or 9 km/hr approx) speed bins. The results are shown in Figure 9.9-9.11. Figure 9.9 shows estimated crash numbers in the *without* and *with* populations for multi-lane divided highways (Type A, rural case) and there appears to be a trend towards reduced effectiveness at higher speeds. Figure 9.10 shows the same information, but now aggregated across all major road classes – and again there is a similar trend towards reduced effectiveness at the highest speeds. Figure 9.11 shows estimated benefits (in the form of numbers of crashes potentially prevented) broken down by 6 road types: Type A roads clearly show crashes reduced in a higher speed band than for Types B and C (undivided highways). Figure 9.11 also helps explain the bimodal nature of the distributions seen in Figure 9.10.

Note that in the simulations that generated the driving and crash populations we did not impose a specific lower speed filter for the initial conditions extracted from naturalistic driving data. We return to this point in Section 9.4.

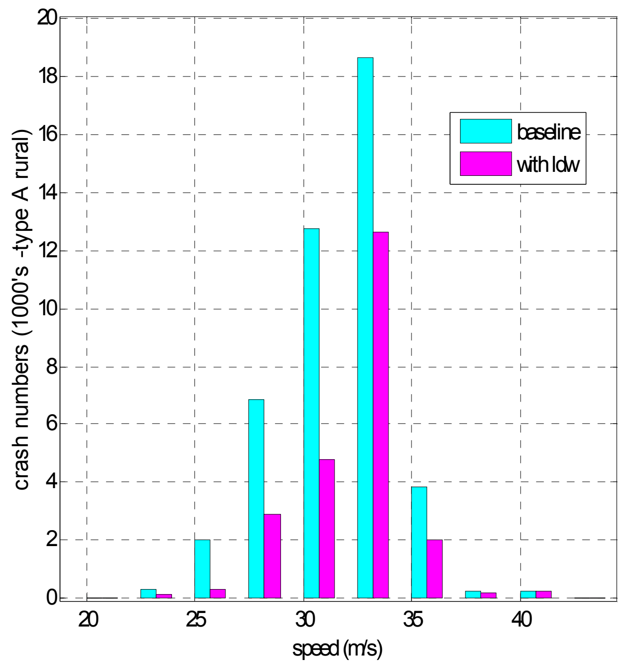

Figure 9.9. Crash Numbers in the Virtual Population (Type A roads, Rural)

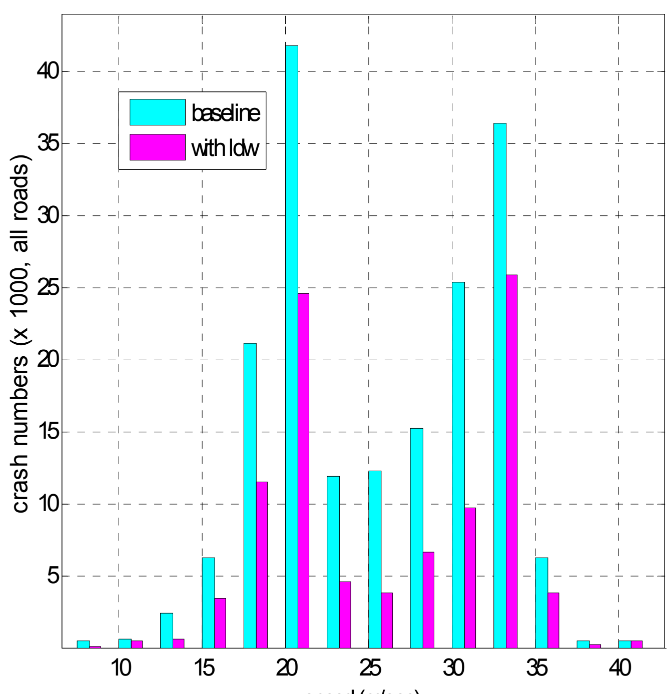

Figure 9.10. Crash Numbers in the Virtual Population (all roads)

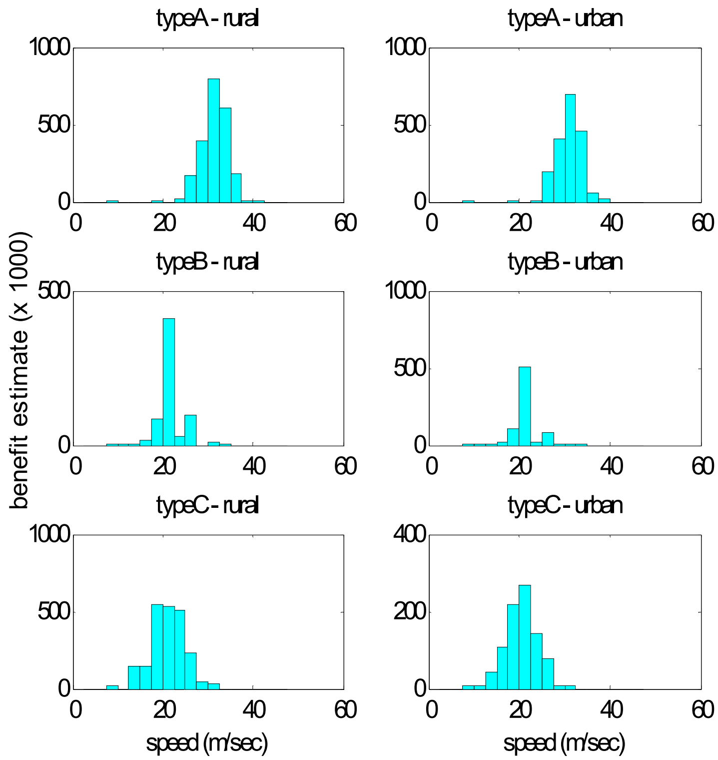

Figure 9.11. Estimated Benefits by Speed as a Function of Major Road Type

9.3. *Effect of Driver Response Time to LDW Warning*

In the results presented in Section 9.2 we have made use of LDW response delay times, randomly sampled within the range $0.4 \le T_{LDW} \le 0.8$ (and $0.3 \le T_{LDW} \le 0.4$ for the fatigued driver – see Table C.10 in Appendix C). These parameter ranges are potentially important to the overall system performance and safety benefits estimates, and the bounds were based on a relatively small set of driving simulator events (Appendix C). Hence, due to the central importance of T_{LDW} and the uncertainty in its distribution, in this section we consider the sensitivity to an overall adjustment to these upper and lower bounds. As mentioned in Table C.10 in (Appendix C) the effect of a uniform increase of 0.1s is to be considered. Hence, the following results are based on re-running the simulations where the LDW alert is issued, i.e. with sample range $0.5 \le T_{LDW} \le 0.9$ ($0.4 \le T_{LDW} \le 0.5$ in the fatigued case).

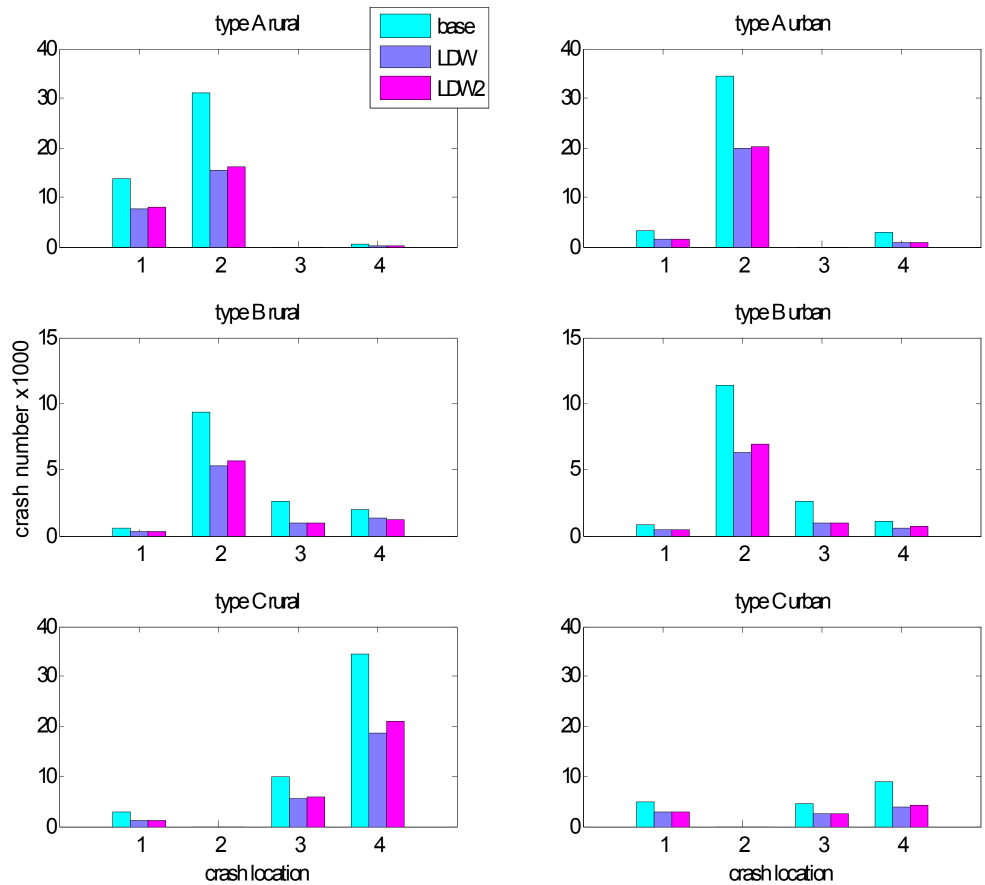

Figure 9.12 Effect of Delayed Driver Response to LDW Warning

(*LDW=standard range, LDW2= +0.1 second condition*)

Figure 9.12 shows results for estimated crash numbers with the initial and increased LDW response time. As expected, the additional delay reduces the estimated benefits but it is a relatively small change that is predicted. Table 9.4 gives the corresponding crash number estimates (c.f. Table 9.3). In this case the overall system effectiveness is reduced to 44% compared to the previous estimate of 47%. Thus the change in the benefits estimate due to a delayed driver response is quite small, certainly when compared to factors considered in the Section 9.4.

Table 9.4. Crash Number Sensitivity to LDW Response Time.
[*Type A: 4 lane divided, Type B: 4 lane undivided, Type C: 2 lane undivided*]

Estimated Crash Nos. (thousands)	Type A Rural	Type A Urban	Type B Rural	Type B Urban	Type C Rural	Type C Urban	Total
Base	44.9	40.3	14.5	15.9	47.2	18.1	181
LDW (std)	23.1	22.1	8.0	8.3	25.3	9.2	96
LDW (+0.1)	23.9	22.7	8.2	9.0	27.8	9.6	101.2
Increase	0.8	0.6	0.2	0.7	2.5	0.4	5.1

9.4. *Effect of System Availability and Other Influential Factors*

The LDW system relies on a camera to register lane markings and hence compute vehicle position relative to the lane. The system is only available when this data capture is successful, and hence the probability of the system being available depends on the quality and continuity of the lane markings, any contamination of the road surface (e.g. standing water or repair markings) and also the lighting conditions. For this project, availability data was obtained from a number of tests conducted in 2004, with test driving taking place in both Europe and the USA, and relevant results are shown in Table 9.5. Comparable data from the RDCW project (Leblanc et al, 2006) is shown in Table 9.6; the RDCW data also includes minor “surface” streets (i.e. not limited access highway) which we identify with our Type C designation (undivided two-lane highways). In Table 9.5, the combination of rain and nighttime driving is seen to dramatically reduce the availability of the system, and this was also the case for the system used in the RDCW project.

Table 9.5. Typical LDW availability (Volvo)

	Light condition	**Day**		**Night**	
	Weather	**Clear**	**Light rain**	**Clear**	**Light rain**
Road Type	**Interstate/Freeways/ Expressways**	91%	94%	95%	17%
	Arterials – Rural	86%	89%	94%	53%
	Arterials - Urban	86%	86%	84%	49%

Table 9.6 RDCW Availability data (approximate summary values from the RDCW final report)

		Rural	Urban
Road Type	**Interstate/Freeways/ Expressways**	80%	80%
	Major Surface	70%	40%
	Minor Surface	50%	35%

The data above depends on both the system used (Volvo or RDCW) and the environment (roadway and roadway maintenance, weather and lighting). Neither data set fully represents the population of interest, i.e. the Volvo LDW system when used on US roads in geographic locations that typify the crash population under consideration. Since the data are not specific enough to fully populate the scenario list of Table 4.6, we use simple assumptions to expand the availability data to the full scenario set (25 rural and 25 urban cases):

- Ignore any association with curved roads or fatigued drivers (i.e. we assume these factors are neutral)
- For Type A roads (typically freeway and major arterial) assume a baseline availability of 90%, based on Table 9.5
- For Type B roads (multi-lane, undivided, non-freeway) assume 85% in rural areas (Table 9.5) and 70% in urban areas (15% reduction, as seen in *Minor Surface*, and half of the 30% reduction seen for *Major Surface* in Table 9.6). This adjustment for urban areas is to take account of the clear trend seen in the RDCW study, while noting that the Volvo data in Table 9.5 shows little or no difference.
- For Type C roads (2 lane undivided highways) there is no appropriate data in Table 9.5, so we reference values from *Minor Surface* in Table 9.6: we assume 50% in rural areas and 40% in urban areas, again slightly diminishing the availability reduction seen in the RDCW data.
- Wet surface and nighttime driving combination: From Table 9.5, for Type A roads the condition reduces availability to 20%. The 25 scenarios do not include any case for Type B roads, so here no value is used. For Type C roads we have no specific data in the above tables, but assume the same 20% value. This is consistent with the rural (50%) Type C availability in Table 9.6 showing a reduction level of around 30% when there is rain at night (in Table 9.5, reductions of 30% or more are clearly seen)

From the above assumptions, we are able to fully populate the scenario set. Results are summarized in Table 9.7 and presented for each scenario in Figure 9.13. The effect on the safety benefits estimate is simple to implement. For each scenario (25 urban, 25 rural) the safety benefits estimate in terms of crash numbers is scaled by the availability factor. The resulting benefits estimate numbers are summarized in Table 9.8. The overall system effectiveness estimate (percentage reduction in crash numbers) in now 33%, compared to the unadjusted value of 47%.

Table 9.7. Availability Factors Applied in the SIM Benefits Analysis

		day		night	
		clear	wet	clear	wet
Type A	rural	90	90	90	20
	urban	90	90	90	20
Type B	rural	85	85	90	na
	urban	70	70	65	na
Type C	rural	50	50	50	20
	urban	40	40	40	20

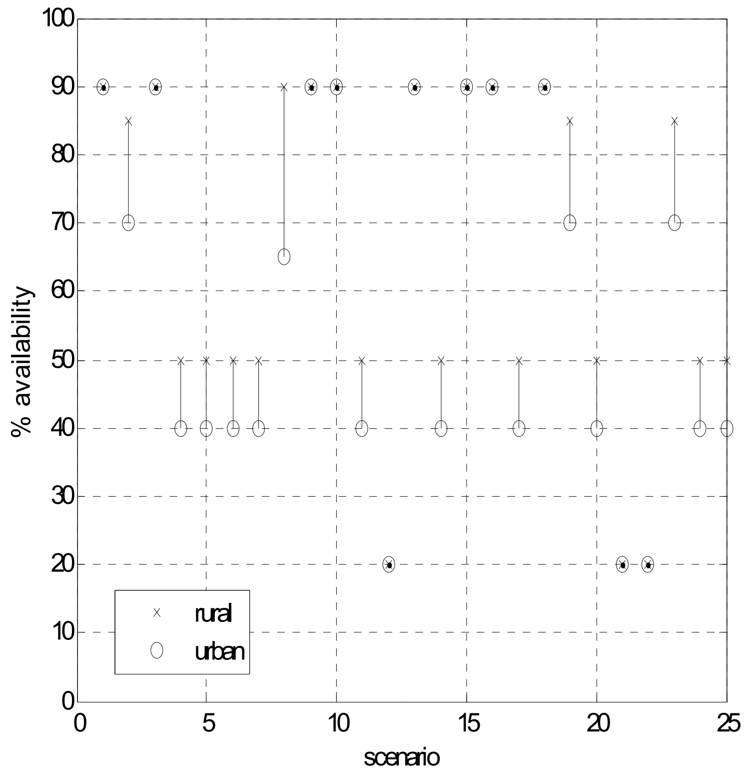

Figure 9.13. LDW System Availability by Scenario as Implemented in the SIM Benefits Estimation.

Table 9.8. Crash Number Estimates Including System Availability Effects.

(*Type A: 4 lane divided, Type B: 4 lane undivided, Type C: 2 lane undivided*)

Estimated Crash Nos. (thousands)	Type A Rural	Type A Urban	Type B Rural	Type B Urban	Type C Rural	Type C Urban	Total
Base	44.9	40.3	14.5	15.9	47.2	18.1	180.9
With LDW	26.0	24.6	8.9	10.7	36.7	14.6	121.6
Difference	18.9	15.7	5.6	5.2	10.5	3.5	59.3

Another factor closely related to the above is the reduction in availability due to the 40mph lower speed threshold – below this speed the system is disabled, so it is equivalent to the system being unavailable in this condition. Here, for convenience with the units used in the velocity distribution we assume a speed threshold of 17.5 m/s = 39.1 mph.

From the velocity distributions across all scenarios the speed-related availability is 92%. However the effect is highly scenario dependent, so for estimation purposes the full distribution of scenarios is used – Figure 9.14. In many cases the availability is 100%, as no cases were seen in the naturalistic driving data below the threshold speed. No specific speed filter was used in sampling the naturalistic driving data, but other filters such as high traffic density (Section 4.5) appear to have had the effect of excluding speeds below the threshold in many driving scenarios. On the other hand, in some cases the speed threshold has a significant effect, as in Scenario 5 urban region (2 lanes undivided on curves in daytime, dry road with no adverse weather) which has only 65% availability. The effect on benefits simply requires us to compound the availability percentages presented in Figures 9.13 and 9.13. The effect on crash number estimates is shown in Table 9.9, where system availability and speed threshold are both accounted for. It is seen that the reduction in benefits is largely limited to the Type C roads, and the effect is not that large; the overall system effectiveness is now 32%, compared to 33% when the potential reduction in the benefit estimate from the speed threshold was ignored.

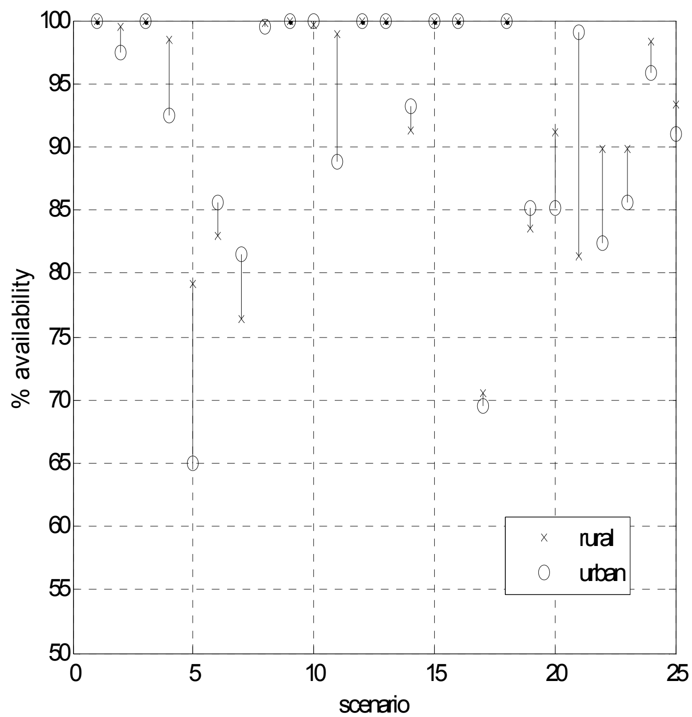

Figure 9.14. Estimated LDW System Availability due to Speed Threshold.

Table 9.9. Crash Number Estimates Including Speed Threshold and LDW Availability.
[*Type A: 4 lane divided, Type B: 4 lane undivided, Type C: 2 lane undivided*]

Estimated Crash Nos. (thousands)	Type A Rural	Type A Urban	Type B Rural	Type B Urban	Type C Rural	Type C Urban	Total
Base	44.9	40.3	14.5	15.9	47.2	18.1	180.9
With LDW	26.0	24.6	9.1	10.9	38.0	15.1	123.8
Difference	18.9	15.7	5.5	5.0	9.2	3.0	57.1

Overall, the combination of system availability and speed threshold for LDW operation reduces the estimate of crash benefits from 83,000 to 57,100; this amounts to a (scenario-weighted) mean availability of 67%.

We now consider a number of other factors likely to influence the overall benefits estimation. Unlike system availability there is insufficient objective data to fully quantify the effects of these additional factors, though where possible we indicate the broad magnitude of their influence.

Under-Reporting of Driver Fatigue

It is widely reported that the involvement of driver fatigue is underreported in crash data (NTSB, 1999). In the virtual population analyzed here, there are 4 major scenarios coded in this way (scenarios 11, 13, 15, and 20) accounting for 7% of all relevant crashes. According to Horne and Reyner (2001) the number in the UK is 15-20% on urban roads, and around 20% for freeways. So we may project that around twice the number of crashes really involves driver fatigue, i.e. around 14% instead of 7%. Given the increased effectiveness predicted for the fatigued driver scenarios (Figure 9.8: 65% effectiveness for the fatigued case, 45% for non-fatigued) the overall safety benefits would therefore be increased. The effect on an additional 7% of crashes involving a fatigued driver would be to add approximately +1.5% to the overall system effectiveness in the full population; since this improvement is then reduced by the system availability as discussed above, we conclude that adjusting for the effects of underreporting of fatigued drivers in crashes would be to add around 1% to the overall system effectiveness. Though important, the size of the effect is not great, especially given the considerable number of other uncertainties in the benefits estimation.

Driver Compliance

Driver acceptance of the technology and compliance with its alerts is one such area of uncertainty that should affect the prediction of safety benefits. If, for example, the driver is annoyed by repeated audible alerts, and does not address this by a positive adaptation, e.g. by increasing the use of the turn signal when exiting a lane, there is a distinct possibility that the driver will turn the system off, or simply ignore its output or be slow to respond. In each of these cases the effect is the same – reduced overall system availability and a reduction in the safety benefits. No appropriate data is available to the project team to evaluate the associated reduction in system availability, so no modification is made to the safety benefits estimates presented above. LeBlanc et al (2007) presented some relevant results relating to the RDCW data set. A sample of LDW alerts were analyzed from the RDCW field operational test; after removing 17% determined to be false alerts, and 33% of the sample associated with lane changes, it was found that drivers corrected their lane positions (meaning moved toward the center of the original travel lane within 6 seconds of the alert) 48% percent of the times. Thus, in this analysis, for half the time drivers were prepared to maintain the lateral position at which the alert was sounded (or drift even further from center) and didn't feel the need to correct.

From this, we might be tempted to estimate a 50% "compliance factor", but the issue is heavily complicated by context of the alert; for example a driver who has been distracted from the forward lane control for some time may well be more appropriately "primed" to respond, while a driver who is aware but unconcerned about the lane departure event (e.g. cutting inside the lane marking on the inside of a curve, allowing the vehicle to drift onto the paved shoulder while monitoring the zone for parked vehicles or other hazards) may well decide not to comply. We also note that compliance and response times are inter-connected; the distracted drivers in the VIRTTEX study (Section 6) often delayed their response in returning to steering control, and this again is a measure of reluctance to comply. Given the overall uncertainty on this complex driver behavior, the authors believe it appropriate to acknowledge its importance, but to consider a full analysis to be outside the scope of the present study. A similar conclusion will apply to our final "sensitivity factor", the exposure ratio.

Exposure Ratio

The exposure ratio was defined in Section 2 as the ratio between the expected frequencies (or probabilities) of a driving scenario occurring given that the ACAT system is enabled, compared to the baseline (*without*) population. To evaluate the exposure ratio would require a study of driver adaptation, for example to determine whether a driver is more likely to drive at night when fatigued given the system is supporting the driver. More subtle changes in driver behavior might well lead to changes in the within-scenario alpha parameter distributions, which is a more extended form of the concept behind the exposure ratio. In this case we might want to account for the different lane position control for a driver who is experienced in driving an equipped vehicle. This latter effect was well demonstrated in the RDCW field operational test (LeBlanc et al, 2006), for example in terms of a reduction in the standard deviation in lateral lane position. In our study, the effect would require a specific set of alpha parameter distributions for the population of adapted drivers. This is certainly a feasible data set to collect, but again is largely outside of our scope.

9.5. *Summary of Safety Benefits Estimation Results*

In this section we have presented a detailed analysis of the results based on the method of a simulated population of driving events and virtual crash outcomes. Starting with an experimental design of 200 simulations from each of 76 constituent driving scenarios (and hence over 15,000 simulations in the baseline set) a number of crash metric components were evaluated. The crash outcome results represented by these metrics were matched as far as possible to that of the mean annual crash population in the US averaged over a 5 year period. The quality of fit between actual and virtual populations was generally good, though for multi-lane rural divided highways there is a noticeable difference between crash number estimates when comparing paved shoulder and off-highway. While further work is needed to resolve this particular discrepancy, the overall fit was encouraging, especially given the large number of assumptions and simplifications adopted in the study.

Once the virtual-driving/virtual-crash population was created, safety benefits were estimated based on running repeat simulations for all cases where a lane departure took place; the exact same initial conditions and random number sequences were used to re-run the events with an LDW system represented in the simulation model, and hence a new set of crash numbers were obtained. Differences between the baseline and LDW-enabled populations then formed the basis for extensive analysis of the safety benefit estimates (crash numbers reduced) and the system effectiveness estimates (in terms of the relative reduction in crash numbers) provided by the methodology. All calculations were accomplished in Matlab via customized m-files, using input data from the detailed simulations, as well as data tables for averaged crash numbers by type and Michigan highway data associated with the crash types (referred to respectively as Table_GES_13K and Table_Mich_10K in Section 7.1). In spite of the many thousands of computations required, the only part that requires significant computation time was the initial calculation of the metric components, and even this only takes a few minutes on a Pentium PC. As mentioned previously, the major computational effort is due to the simulation of the virtual driving population.

The initial estimate of 47% crashes reduced (85,000 crashes annually for the target vehicle type and crash type) is only an initial estimate, and needs to be refined by taking into account real-world factors. The first of these is system availability; using available field data, and taking into account speed and sensor performance on a range of road types, the effectiveness is reduced to an overall revised estimate of 32% (57,000 crashes) which still represents something of an optimistic estimate of the system effectiveness. As mentioned above, the overall availability factor, including speed threshold effects, is 67%. We write this as $A_1 = 0.67$.

The next major factor is driver responsiveness (compliance and time response). Sensitivity to the driver's speed of response was estimated as an overall reduction of effectiveness of 3 percentage points relative to the 47% baseline for a 0.1 second overall delay time. To assess the overall magnitude of the effect, we assume the reduction in effectiveness is linear with respect to delay time. Treating this as an “equivalent availability” A_2 , we estimate

$$A_2 = 1 - 0.61 \times \Delta T$$

where ΔT is the additional delay in seconds. For example, substituting $\Delta T = 0.1$ in this equation yields the modified estimate for the crash number reduction: $A_2 \times 85{,}000 = 79{,}800$; this in agreement with the values stated in the final column of Table 9.4 (181,000-101,200 = 79,800).

The other factor identified as highly influential is that of driver compliance (and acceptance) – the willingness of the driver to leave the system switched on and respond appropriately when an alert is issued. We have no reason to believe the drivers will always respond appropriately to LDW warnings, or even leave the system switched on. While having no objective data to determine this factor, we can gain some insight on its relative importance by proposing a simple “50% rule” – that is we impose an error bound based on a range between neutral (no change from the default, in this case 100% compliance) and one that is 50% worse (i.e. 50% compliance). Again, an equivalent availability (or adjustment factor) is used, ranging between $A_3 = 1$ in the optimistic case to $A_3 = 0.5$ based on this simple rule.

For the driver time delay influence we similarly impose a range of no change (estimated mean delay parameter $T_{LDW} = 0.6$) to one that is 50% worse ($T_{LDW} = 0.9$, i.e. $\Delta T = 0.3$ s). In this case A_2 ranges between 1 in the optimistic case, to $A_2 = 1 - 0.61 \times 0.3 = 0.81$ based on the 50% rule. These error bounds are admittedly crude since the actual values are unknown, but they dominate over the effects of any statistical errors due to sampling variations. The overall range of estimates is shown in Figure 9.15, here shown as system effectiveness, and based on the equation

$$B = B_0 \times A_1 A_2 A_3$$

where B_0 = 83,000 is the basic initial estimate of crashes reduced and the factors A_i are the availabilities – actual or equivalent – respectively associated with system performance, driver response timing and driver compliance. The range of values is shown as cumulative moving to the right, so for example the final bar on the right indicates the overall estimation range due to system availability, reaction time and driver compliance; the final estimate for the effectiveness range is 13% - 31.5%, corresponding to an estimated number of crashes prevented annually in the range 24,000 – 57,000.

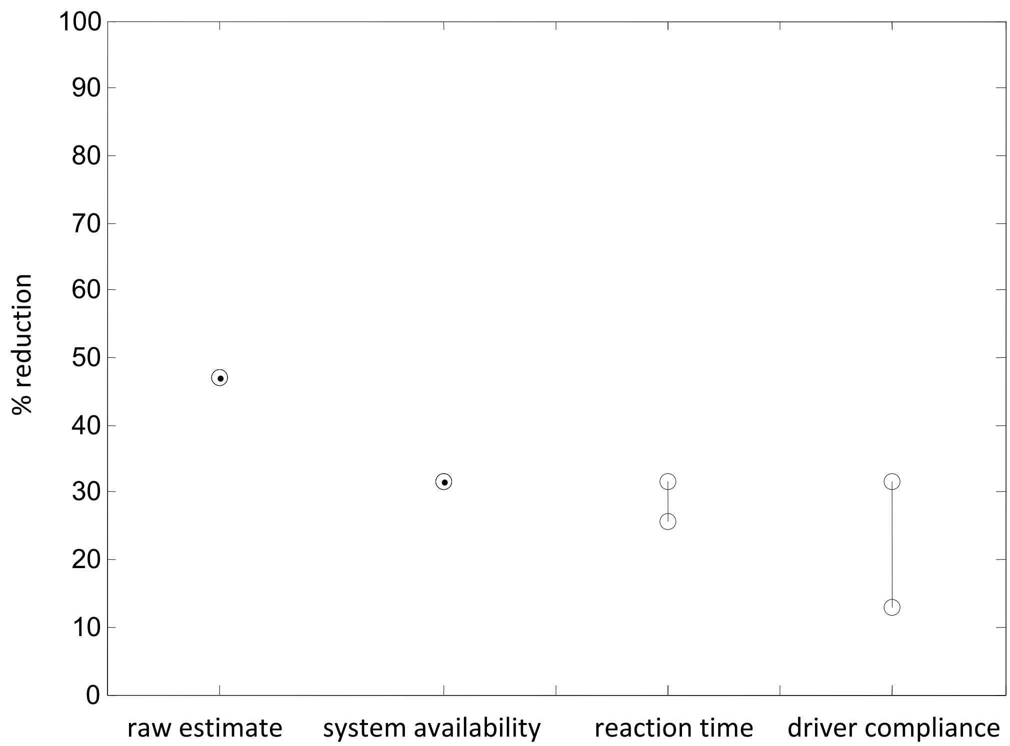

Figure 9.15. Error Estimates For Mean LDW System Effectiveness Based on Major Identified Factors.

It is worth comparing the results of the above analysis with that of the Volpe analysis of the RDCW FOT data (Wilson et al, 2007), which was based on the frequencies of LDW alerts relative to miles driven. Note that in our analysis we did not consider frequencies of lane departure events or warnings at all; we merely used the naturalistic driving data to initiate the simulations. Wilson et al, (2007) report an estimated 10% - 60% reduction in crashes, based on the observed reduction in conflicts. In their analysis, the LDW effectiveness in reducing crash numbers was evaluated based on two “pre-crash scenarios”, namely (i) going straight and departed road edge, and (ii) negotiating a curve and departed road edge. These were defined in terms of valid system alerts (LDW alerts deemed true positive after video review). Note that in the Volpe analysis, the scenario is defined at a later stage in the crash phase timing chart compared to the present ACAT analysis, i.e. during the conflict stage. From this FOT data analysis, the performance advantage is determined entirely by the exposure ratio – experimentally determined reductions in conflicts per unit distance traveled found to be statistically significant within a particular speed range. Interestingly, the speed range where benefits were most clearly found in the Volpe study (above 55mph, or 24.5m/s) corresponds with the dominant speed range of benefit for Type A roads (multi-lane, divided highways - see Figure 9.11). In the present analysis, lower speeds also show benefits, but this is not at variance with the results of Wilson et. al., because their analysis was limited by small sample sizes for the conflicts of interest, and was also conservative in assuming no crash reduction based on a prevention ratio – only conflict rates were used in their analysis. Overall then, methods differ but the results of the present study, while only provisional based on the large number of detailed assumptions, are generally consistent with that of the Volpe study. While the lower limits of around 10% are very similar, the approximate upper limit of around 30% effectiveness here is substantially lower than the 60% reported by Wilson et al, (2007). Because the present study considered a wide range of scenarios and analyzed the full response of the driver-vehicle system during a lane departure conflict, we believe the 30% value may represent a more plausible upper limit to the overall system effectiveness.

10. Conclusions

The underlying purpose of the ACAT program has been to address gaps in current knowledge about the performance and likely effectiveness of new and emerging active safety technologies in reducing crash numbers. As described by NHTSA in the request for applications (NHTSA, 2006), an ACAT project should meet two goals. The first is to develop a formalized Safety Impact Methodology (SIM) tool to evaluate the ability of advanced crash avoidance technologies in full vehicle systems to address specific types of motor vehicle crashes. The second goal is to demonstrate how the results of objective tests can be used by the SIM to establish the safety impact of an actual driver assistance technology.

The VFU project team chose to address the first goal by developing and applying a SIM tool which at the core uses a detailed mechanism-based approach (continuous-time simulation) to represent the potential influence of driver assistance technologies. The basic SIM procedure starts by exploring real-world crash mechanisms using both statistical and in-depth analysis of reported crash events in order to understand contributory factors and event sequences in actual crashes. In this project, the crashes involved lane/road departures and included road departure crashes, head-on collisions, sideswipes, and other crash modes. This information is then used to develop a comprehensive set of Driving Scenarios (DS) which precede the crashes. The DS's are then further parameterized in all aspects necessary to represent them in software via a computational model. Representation in a computational model means that the four DVET components (driver, vehicle, countermeasure, road environment and technology) are represented by computational sub-models interacting in a virtual environment rather than by physical objects interacting in the real world.

Following the definition and parameterization of the full DS set, multiple cases are sampled and run in a Monte-Carlo simulation. The computational model time-steps from the starting point of each DS until the DS has run for a certain time interval (typically this is 10 - 20 seconds, though the exact value is dependent on how the simulation evolves). The driving scenarios precede any crashes but they are not pre-crash scenarios in the sense that a crash inevitably follows from DS development. Rather, crashes may or may not result from any given DS as it develops over time. Of course this applies both to real driving and simulations, but for efficiency sake there was a controlled process of over-sampling high risk initial conditions for the event simulations. The DS are thus "coarse-grain" in the sense that they cover a broad range of situations which include the ones that lead to crashes but also a number of situations where no crash occurs. Sub-sampling within the DS provides an essential improvement in resolution and detail in the resulting "virtual crash population". The point about efficiency is that crash risk is systematically increased in the virtual driving events, so that crashes are more likely (in the present study this just means that a significant lane departure is likely) while in real world driving a crash is much less likely.

In other words, rather than looking at single case accident reconstruction, the aim has been to generate an ensemble of crash/no-crash situations, and then study whether the crash avoidance technology under evaluation changes the overall proportions of crash/non-crash outcome for this ensemble. Running the simulation for all DS's therefore results in two distributions of virtual conflicts and crashes, one with the technology active and one without. These distributions have then been mapped to real-world crash types and frequencies using suitable crash metrics, as presented in Figure 10.1 (and previously as Figure 4.1)

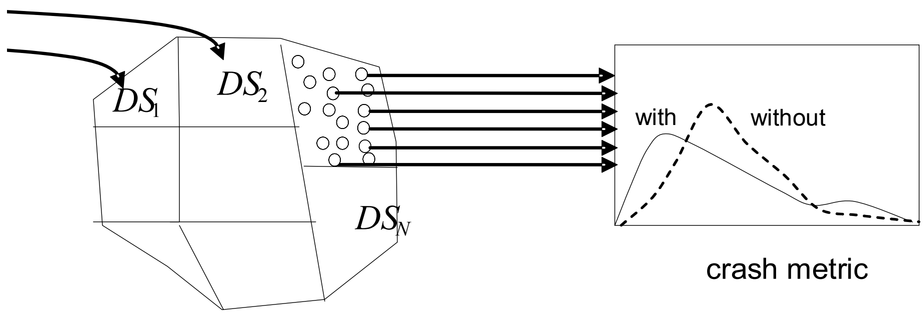

Figure 10.1. Driving Scenario Resolution for Monte Carlo Simulation With and Without the Specific Safety Technology Being Evaluated.

10.1. *Challenges in Data Collection*

In a general perspective, it can be argued that all approaches to the evaluation of active safety technologies face the same challenges. Regardless of where and in which form an evaluation takes place, the characteristics of the four components needed to perform the evaluation, i.e. Driver, Vehicle, Environment and Technology (DVET), must be represented in a way which is sufficiently similar to their respective counterparts in real world crashes.

The natural starting point when providing these DVET component descriptions is crash data from sources such as GES. However, while GES and CDS data play an essential part in defining the crash circumstances which the technologies under evaluation are meant to address, it has become clear during the project that crash data in itself contains limited or no detail on a number of the essential conditions and parameters which must be defined in order to create descriptions with sufficient information to set up a simulation-based evaluation. For example, while drivers may be coded as distracted or fatigued in the crash data, these factors are not described in sufficient detail to guide the setup of evaluation conditions in which an effect of the technology can be measured. For example, to evaluate how efficient a LDW technology is in redirecting a distracted driver's attention to the roadway, it is necessary to use test subjects who are distracted in a similar way to their real life counterparts. However, crash data says little about the duration or nature of the distractions cited as crash contributing factors.

To overcome some of the limitations of crash data, the project had to make use of a number of other data sources in order to create DVET component descriptions at a sufficiently detailed level to reach even a moderate level of confidence in the evaluation outcome. A substantial amount of work has therefore been devoted to retrieving and processing the structure and performance data needed for detailed DVET component descriptions. While this is not unique for the continuous time-simulation approach to technology evaluation, the importance of such data, as well as the proper tools for processing them, is very clear. Since all DVET components are virtually represented rather than physically, the veracity of the simulation outcome is solely dependent on how well the DVET component models capture the characteristics of their real life counterparts[9]. This means that a continuous time-simulation drives a very comprehensive data collection paradigm, which in this project has had the side effect of exposing some gaps in current data collection procedures. For example, in order to characterize normal driving behavior in terms of speed selection and lane positioning for the DS set in this project, it was necessary to have access to naturalistic driving data, as well as tools for extracting the portions of the driving data which represent the type of driving that precedes lane departure related crashes. Without access to those data and tools, the precision in applying the SIM-tool would have been severely reduced. Another example concerns the highway environmental model; since GES and CDS crashes are not geo-located, a detailed characterization of the road geometries at crash locations was not directly available. Imputed data was used instead, based on geo-located Michigan crashes. But the general point should be clear that detailed simulations require detailed environmental data, and future demands for such data will only increase.

The second goal for the ACAT projects was to demonstrate how the results of objective tests can be used by the SIM to establish the safety impact of a real driver assistance technology. As has been discussed above, in the continuous time-simulation approach used by the VFU-team, all four DVET components are represented virtually rather than physically. Objective testing has therefore taken on a somewhat different role compared to standard test track or driving simulator approaches. Rather than using results from objective testing directly to establish the safety impact of driver assistance technologies, the objective tests have instead provided an empirical basis for calibration and validation of the computational sub-models used to represent the DVET components. This has included objective tests in the form of detailed technical tests of the vehicle and the driver assistance technologies (principally on test tracks, but with highway driving to establish system availabilities) as well as objective tests designed to capture typical ranges of human performance where the driver is in the loop, with both track and driving simulator testing being used.

[9] To be clear, note that this is not to say that by physically representing components, one avoids the question of whether the DVET components used in evaluation are representative of their real life counterparts. For example, the extent to which driver behavior in a driving simulator or on a test track is the same as driver behavior in everyday driving outside a test situation remains to be fully explored.

10.2. *Component Model Fidelity*

An important discussion for continuous time-simulation approaches is how to verify DVET component model fidelity. This refers to how a model is implemented, i.e. the structures and processes used to recreate the conditions and behaviours of its real world counterpart. A possible criticism of any simulation approach is that if the component models it uses are simplifications of reality, the resulting simulation outcomes must be simplifications as well. Or put differently, one could argue that it's only when there is full bio-, mechanical-, neurological- or some other type of fidelity which mimics the properties of the physical component, that simulation results can be relied upon.

While this is true in some domains, in the context of active safety technology assessment, the primary concern is whether a certain technology influences the outcome of a set of relevant driving scenarios. Therefore, the primary type of fidelity to strive for is functional performance fidelity, i.e. to build component models which behave and interact like their real world counterparts, regardless of their internal design. Basically, if the component models used in an evaluation are able to re-create the real world initial states and behaviors of their real world counterparts for that DS set, then the simulations will indeed capture the influence of the technology on that DS set with some degree of certainty. The case for simplified system-level models of complex or interacting processes has been established over many decades in areas such as multi-body dynamics and control systems; in the engineering community it is well-established as a valid approach for system design iteration and verification.

If the key aspect of component modelling is functional performance, then any model which reproduces all relevant aspects of structure and behaviour displayed by its real world counterpart in the set of DS for which the technology is evaluated will suffice. One important implication of this is that complex component models are only necessary if a simpler model cannot generate the necessary functional performance. In this work it is not claimed that all models are in a fully mature and final form – further data analysis and development is likely to be beneficial in refining and validating aspects of component model behaviour, for example to verify driver steering and braking responses when an actual road departure happens in real-world driving.

10.3. *Simulations and Driver Behavior Variability*

A basic tenet of active safety technology evaluation is that the DVET components which have the largest influence on the evaluation outcome are most important to represent correctly. For evaluation of active safety technologies where a successful outcome is dependent on driver understanding and action (the driver is in the loop, either as part of the control system or with the possibility to override technology inputs), the most important component can arguably be said to be the driver. While the technologies are important too (without them there would be no opportunity to address pre-crash problems at all) the scenario outcomes to a very large extent depend on how the driver responds to the inputs from the technology.

This highlights a basic methodological challenge in active safety function evaluation, which is the very wide range of performance variability in human drivers. Due this high variability, one can rarely hope to devise a successful evaluation setup based on the performance of a single driver. Rather, the standard way of addressing this issue is to sample evaluation by drivers from within some definition of a standard driver range (e.g. 25-45 yrs old, minimum 5 years of driving experience, etc), and where that range is thought to capture "average" driver behavior. In this context, a simulation approach offers some very interesting opportunities. Assuming reliable and representative data on the range of driver performance parameters can be established, instead of restricting evaluation to a few drivers and their particular behaviors, a large-scale batch simulation allows for the exploration of the widest possible variations in driver behaviors in any given driving scenario.

Also, because drivers who participate in studies usually have to fulfill some minimum driving experience criteria to qualify as "average", they will have many hours of successful adaptation to the development of events in traffic behind them when they come to the test facilities. However, since the technologies under evaluation often are aimed at enhancing the driver's control during critical driving scenarios (such as inadvertently leaving the lane or being in an imminent collision phase), a primary evaluation challenge is to get drivers into those critical driving scenarios without them knowing or suspecting this beforehand, and consequently taking preventive action on their own to avoid the undesirable outcome. This can be very difficult in controlled experiments and, often, validity has to be sacrificed to some degree to accommodate other constraints. When the driver component is simulated however, this does not apply. A virtual driver does not remember a previously encountered scenario unless explicitly programmed to do so. The virtual driver's surprise is therefore genuine each time a critical driving scenario occurs, and thus neither the first nor the subsequent scenarios lead to driver adaptation. Thus while the particular challenge of driver modeling remains significant, especially when expanding the scope beyond simple lane-keeping, the future payoff may be worthwhile, in terms of the range and depth of future studies adopting a computational approach.

10.4. *Base Rates and Joint Distributions of Crash-relevant Parameters*

Separate from component fidelity, the benefits estimation process depends heavily on weighting a simulation outcome by the likelihood of that scenario occurring in the first place. Even a high-fidelity simulation with good rendering of an outcome can lead to imprecise estimates of safety benefits if the base rates of occurrence used in making those estimates are in error. In this project, the scenario-level weights have been estimated based on observed crash frequencies; therefore the scenario weights are subject to unknown errors induced by the practical limitations of the GES coding system and the supporting police-reports. Within the driving scenarios, distributions of driver behavior and vehicle kinematics have been derived from naturalistic driving, and the size of these samples is also a limiting factor, especially when the particular scenarios are not very common.

In the approach adopted in this project, simple inaccuracies associated with factor correlations have been largely removed, at least in principle, by using random sampling of multivariable databases for highway and driving data. However, geographical limitations on the data used may still influence the results – for example driving styles and highway design varies across the US in a way that has not been captured in the supporting databases, and this may also influence the scenario weights. Finally, relevant variations in driving style may show time variations as well as spatial variations. In particular, if driver adaptation to a safety system changes a driver's driving habits or risk-taking propensities, this can have an impact on the benefits, in a way that the estimation process did not account for.

10.5. *Reuse of Knowledge*

While the process of developing and validating DVET component models that can be used in simulation in one sense may be cumbersome (and nothing initially comes "for free") the simulation approach does have the advantage of being able to preserve the knowledge gained in a very efficient way. When a second simulation project starts up, the previous component models and the work which went into validation of those models is available from day one. A second project can therefore focus on improving models and validation rather than having to set out completely anew. True, evaluations where one or more of the components have a physical realization do not have to start from scratch each time either, but there always has to be some form of validation that the physical components used are of the right type and exhibit the right behaviors (i.e. verifying the baseline). For simulated components, that can be skipped unless there is a new scenario to be evaluated or the previously established baseline has been called into question, thus saving time and resources.

Also, in developing a SIM tool using a full simulation approach, one develops a tool where scalability becomes a question of computer time and speed, rather than of limits imposed by the size of testing facilities, availability of test vehicles equipped with the driver assistance technology, etc. Any need for increased testing can therefore quickly and cost effectively be addressed by comparatively small investments in computer time and speed, and there is much reduced risk of encountering "hard" limits, such as the current test facility reaching maximum capacity and further testing therefore requiring the building of new facilities.

10.6. ***Limitations and Future Potential***

This project has used diverse data sources and model components to perform event simulation and hence provide an estimation of safety benefits for a lane departure warning system based on objective data. The methodology developed is both comprehensive and general, even though the individual data sets and sub-models are specific to the systems studied. The research relies on objective data and scientific knowledge in a way that pushes the limits of what is available. For this reason, and also for reasons of time and scope, a number of simplifying assumptions and recognized limitations exist in the study, and these are highlighted in Appendix B. The authors hope that compiling this list will encourage future research and data gathering, so that limitations can be reduced and accuracy and confidence in the predictive estimation will be improved.

A simulation is by definition not reality, and the outcome of simulations in any project applying the SIM-tool developed here will therefore be approximations. Whether physical or computational simulations are used, results should be interpreted with some skepticism, since the real world is not directly represented. The results from objective testing (especially human performance results), cannot be taken at face value in any approach, since they may differ in important ways from the actions people perform when they encounter real emergencies in the real world. Moreover, the joint distributions of many performance shaping factors are not fully understood, and the sample of data used in the SIM may miss important aspects of the real-world crash population. This is of course not a unique problem for a continuous time-simulation approach, since it is true of all approaches to the evaluation of active safety technologies.

In formulating the analytical method, it was realized that lack of resolution within the driving scenarios can cause major unseen errors in the benefits estimation. In this project, naturalistic driving data has been used to fill in much of this detail. However some key variables, especially those relating to driver attention, have remained missing from these data, and so a number of simplifying assumptions have been made. Simplifying assumptions have also been made where limitations in time, resources and/or available knowledge has precluded full investigation or discovery. When reviewing the material in Appendix B, it is important to remember that such simplifications and assumptions are generally neither new nor unique to the simulation approach. They concern all approaches to active safety technology evaluation; it is just more typical that they are not explicitly identified and discussed.

It is particularly important to realize that the benefits estimates (numbers of crashes reduced) is not a prediction about the real world, but rather an estimate of change under hypothetical circumstances; if it were possible to go back in time and retro-fit passenger vehicles with this technology, then it is likely that a certain benefit would have accrued. To convert this kind of information into an actual benefit would require substantial further modeling, in terms of traffic and population trends, technology trends in other parts of the vehicle and traffic system, changes in highway design, changes in vehicle use etc., and especially in terms of the rate of penetration of the technology into the population.

The importance of understanding how drivers respond to inputs from the technology has become very apparent in the development of the SIM tool. The SIM tool developed in this project has mechanisms in place to assess the performance and likely effectiveness of both DAC and ELA, in addition to the analysis actually performed for LDW. However, for DAC and ELA, the objective testing necessary to validate the component models, and in particular the driver model in terms of responses to technology input, was found to be beyond the scope and resources of the project.

In summary, predicting the future benefit of an active safety technology is a difficult task. It depends on data that are sometimes unavailable, it requires assumptions that often cannot be easily verified, and generally presumes a rather simplistic view of the world. Predictions for systems which include human components are especially challenging because humans are very context-dependent creatures who exhibit creativity and change over time. SIM-tool predictions should therefore be treated as giving broad guidance into possible futures rather than as a precise ground truth. They provide insights into fruitful areas of research and development and indicate factors in active safety technology development and deployment that might be significant, but cannot be relied upon in their own right to predict the future.

On the positive side, simulation outcomes may be possible to calibrate against data on how the technology actually performs on the road when such data becomes available. Therefore, over time it will be possible to assess the accuracy and precision of the SIM-tool predictions, and based on this improve the tool. This is quite reminiscent of how evaluation is carried out today in the field of passive safety. While crash test dummies are not meant to be full physical replicas of humans, they can be said to be models which capture a number of important human physical characteristics related to injury outcome in particular crash modes for which the dummy is designed. The way forward in passive safety is often viewed as an iterative process where models (i.e. crash test dummies) are improved, and then calibrated against human body performance in real world crashes, which leads to improved models, etc. This process provides a good example for how work with active safety function evaluation by means of component models in simulation can be performed.

10.7. ***Research Results***

The results of this project are at two levels – development of a methodology for objective estimation of safety benefits based on data fusion and simulation, and the numerical results obtained from applying the methodology to a specific LDW system. Arguably it is the former that is of greatest significance at this point; the methodology has a wide and general applicability that goes far beyond the results of the individual estimation. The development and implementation process has highlighted a number of areas where either basic research knowledge or supporting data have significant gaps at present, and it is hoped that the research findings will stimulate fresh efforts to fill the gaps. On the other hand, the particular assessment of the particular LDW system does have inherent value: benefits are estimated to be in the range of 10%-30% of crashes reduced for the crash types and light vehicle involvements considered. Further refinement of this type of result is likely in the future. The results presented in Section 9 are both comprehensive (they apply to a large and complex population) and detailed: effectively all details are “known” in the virtual population, so for example anticipated trends over road type, speed, and driver factors can be determined. This combination of breadth and depth in the benefits estimation is also expected to support the design and development of active safety systems in the future.

In terms of methodology, the project outcome has been to define and implement an innovative and challenging program of systematic research. Though based on earlier work used by NHTSA, the VFU-ACAT research team is not aware of any other implementation of active safety benefits estimation that utilizes such a wide range of data sources and combines that information with detailed computer simulations of crash occurrence and avoidance. In particular it is worth noting the critical use of naturalistic data to refine detail in the driving scenario and to provide the necessary kinematics to initiate scenario simulation. While field operational test data has been used for this, its use is not specific to the system evaluated or even to the crash types considered; as mentioned above this provides for efficient re-use of the supporting data elements.

While the research output acknowledges uncertainty in a number of areas, the main sensitivity seems to be in the area of human acceptance and responsiveness – the driver needs to accept and understand system function and make appropriate and timely responses to warnings. It has not been possible to assess any long-term trends or adaptations in driver behavior that may affect this. The method also did not take account of possible negative unintended consequences of the system, though with more supporting data it is capable of doing this. Any modes of unintended consequences are not seen in crash data, so crash history and crash reconstruction are entirely inappropriate for analysis; but scenario weights may be derived from existing naturalistic driving data, provided any new modes can be referenced on frequencies seen in normal driving. Again, much of the power of the new methodology derives from the combined use of alpha and beta parameters whereby crash and naturalistic data are merged. The construction of a virtual population has also been advantageous is estimating beta weightings (crash-weighted frequencies) for “hidden” parameters, provided they are coupled strongly to known crash outcomes; in the research it has been possible to infer distributions of crash frequencies arising from the left and right travel lanes on multi-lane highways even though this information is not coded in GES.

Turning to the safety benefits estimation for the specific crash problem, around 15,000 simulations were made in setting up the underlying virtual crash population; by optimizing DS weights it was possible to produce a reasonable degree of fit to the actual (GES coded) crash population, albeit with some error in the locations of first harmful event predicted for off-highway crashes in multi-lane divided highways (see Section 9). The lack of perfect fit is hardly surprising in a first study of this type; however the chosen measure of fit (see equation 9.1) is high – 85% of variations are explained according to the analysis – the use of a virtual crash population is at least plausible, and there remains clear scope for improvement in the future, especially through the use of more detailed and geographically diverse highway data.

For the specific LDW technology, repeat simulations of the virtual crash population with the LDW system in place led to a “raw estimate” that approximately 50% of crashes can be prevented; this number assumes perfect compliance by the driver and perfect system availability. Specific predictions were also derived that the effectiveness (i.e. percentage crash reduction) in adverse conditions is higher than average, but with the exception of crashes on curved roads. It was also found that benefits are likely across a wide range of vehicle speeds. Such conclusions may be feasible in future studies or based on analyses of existing naturalistic data. Note however, that with few relevant conflicts detected directly in field operational tests, and with a lack of event-level comparisons between equipped and non-equipped vehicles, it is difficult to derive comparable trends in FOT analysis from the direct counting of alert rates.

Raw estimates were refined by including the effects of driver responsiveness and system availability. System availability alone reduces the mean estimated effectiveness to around 30%, and this seems to provide a realistic upper limit to the number of crashes that may be prevented. We emphasize that this assessment is based on crashes and dynamic performance involving light passenger vehicles, and any inference towards other vehicle types should be treated with caution. For the light vehicle crashes considered[10] this estimate is applied to a baseline set of 180,900 crashes annually in the USA that could be reduced to about 121,600 with LDW in place, so that around 59,300 crashes might be prevented (Table 9.8). These results appear not to be overly sensitive to the details of the driver response; in terms of reaction time a 50% increase in the mean driver reaction time caused the estimate to drop to around 50,000 prevented. However the influence of the drivers' willingness to use and be responsive to the system in the real world environment has the potential to have a major impact on this prediction and is largely unknown.

10.8. *Strengths and Weaknesses of the ACAT SIM Methodology*

The SIM benefits formulation presented in Section 2.3 expands the relatively standard formulation summarized in Section 2.2 by increasing the level of detail included in scenario definitions, and introducing a major analytical role for naturalistic driving data. Here we consider the potential strengths and weaknesses of the approach.

An important innovation of the methodology described above is the internal resolution of driving scenarios – it is not assumed that all scenarios are essentially the same. In particular it may be that multiple mechanisms and factors operate within the same scenario in terms of crash causation and avoidance. The importance of this is illustrated in the following hypothetical example. Suppose that a crash sequence of interest relates to a distracted driver who unintentionally drifts left out of lane on an undivided highway (although this is relevant to the research problem we are addressing, the numbers and suppositions in the following are purely for illustration). Suppose there are two types of distraction which we label D_0 and D_1 (this is hypothetical, but D_0 might be visual distraction without loss of situation awareness, while D_1 might also include loss of situation awareness). We assume D_0 is the most common form of distraction (occurs 99% of the time in the driving population, with D_1 only 1%) but that D_1 carries a much higher risk,

[10] annual total police-reported crashes for light passenger vehicle types averaged 6.1 million over the period 2002-2006

$P(C \mid S_i \cap D_1) = 1000 \times P(C \mid S_i \cap D_0)$. Further, if the safety system is highly effective, $\pi_i = 0.1$ (prevents 90% of crashes) in the D_0 case but ineffective in the D_1 case ($\pi_i = 0.9$ say) it turns out that failing to recognize the sub-scenario differences creates a large error in the benefits estimation. If the system is only evaluated in the common case D_0, the predicted effectiveness (fraction of crashes prevented assuming for simplicity that there is no change in exposure, $\varepsilon_i = 1$) is clearly

$$(1 - \pi_i \, \varepsilon_i) = 0.9$$

On the other hand, given the above assumptions, most crashes occur with D_1 distraction and the actual fraction of crashes prevented is therefore much lower. In detail, let N_0 represent the number of distraction events in the population and p_0 the probability of a crash resulting from a D_0 distraction event. The expected number of crashes in the baseline "without" population is

$$N = (0.99 N_0) \times p_0 + (0.01 N_0) \times 1000 p_0 = 10.99 N_0 p_0 .$$

The corresponding number "with" the safety system is

$$N' = 0.1 \times (0.99 N_0) \times p_0 + 0.9 \times (0.01 N_0) \times 1000 p_0 = 9.099 N_0 p_0$$

Hence the "true" effectiveness under these hypothetical conditions is

$$\frac{(N - N')}{N} = 0.172$$

which is clearly a lot less than the naïve estimate of 0.9. Of course this example is purely hypothetical, but it does highlight the potential for errors resulting in a safety benefits estimation process that does not fully resolve the underlying mechanisms of real-world crash events, especially those associated with driver state and driver behavior.

Real world crashes generally occur for complex reasons. In this study we are concerned with crashes occurring after an individual driver inadvertently allows the subject vehicle to exit the travel lane; if the vehicle had stayed in the lane it is assumed that no crash would have occurred, so in any particular case the fact of this error would need a deeper explanation, for example to include the actions and mental state of the driver, including potential discrepancies between the actual vehicle/traffic conditions and the drivers internal model of those conditions. Reconstruction of individual crash events containing such information is at the very least a challenging undertaking. In the approach described above, the reconstruction takes place instead at the population level; multiple driving scenarios are constructed and multiple simulations run from each particular case, with scenario weights optimized to match real world crash data and hence provide a candidate *virtual crash population* for technology testing.

Included in the virtual crash population is a simplified modeling approach that broadly represents fallible driving behavior, including gross effects of distraction and diminished vigilance, for example due to sleep deprivation. In both cases the driver "module" in the simulation is subject to errors in lane tracking, as well as delays in responding due to limitations in recognition and reaction processes. The model is inevitably a gross simplification of what really happens. In many branches of engineering such "system level" models are well accepted and routinely used, for example in vehicle engineering to explore the performance of stability control systems. In that case the fine detail of suspension bushings and tire mechanics are lumped into gross system parameters that may be identified from experiment or from more complex models. Once the high level system performance of the model has been calibrated and validated, it is used to explore the phenomena it was designed for, and the same general philosophy is adopted here. It seems clear that reliance on some form of model is necessary because of the critical role of the driver in our study area. And, with faithful crash reconstruction on a case-by-case basis probably out of reach of current data and knowledge, the "system model" approach has been adopted as the only one that appears both feasible and credible.

The current methodology is not without its limitations, and based on the example presented in Section 2.3, it is important that all major contributory mechanisms leading to crash are represented in the model. In this study, with emphasis on the overall methodological approach, it is to be expected that some elements of the analysis – including the fidelity of the model and the data used to support its calibration – can be improved in the future.

The SIM analysis, based on constructing a virtual crash population, is capable of predicting changes in crash numbers for that population as a result of installing the ACAT technology on the simulated vehicle. Assuming a reasonable degree of model performance, this is a major step forward in the analysis of active safety systems. However we should not equate this conditional predictive ability to call out the actual number of crashes that will be reduced in the real-world as the ACAT technology in introduced. Even if the model is assumed to be faithful to the way such crashes occur in the real world, and even if we assume the selected GES codes do accurately represent the crash population of interest, we should recognize other factors that limit our predictive ability:

- changes in demographics compared to the historical crash data referenced
- changes in traffic conditions and other vehicle technologies, as for example congestion increases, the mix of vehicle types changes and the proportion of vehicles equipped with electronic stability control increases
- changes in other driver behaviors, especially the increased use of hand-held mobile devices in the vehicle
- the nominal benefits estimation is based on *all* vehicles in the population being fitted with the technology under test.

More realistically, we regard the benefit *B* in equation (2.8) to be a measure of system benefit that has the units of crash numbers and represents a uniform metric to compare different technologies, provided *B* is estimated using a consistent methodology. This appears an exceptionally valuable metric to be able to derive, and an important step forward for forecasting and improving safety benefits estimation through advanced technologies. On the other hand, based on the above considerations, B is unlikely to represent the actual number of crashes that will be prevented if the specific ACAT technology is widely adopted, and this is quite apart from the fact that technology deployment is always a gradual process. For this reason we treat the benefit estimates with caution, and at this stage do not expand the analysis to include predictions of injuries, deaths, and economic costs reduced.

10.9. *Future Studies*

The simulation-based SIM tool developed by the VFU-ACAT team in this project, with its explicit description of how component sub-model calibration and validation influences benefit calculations, offers some further interesting thoughts on future studies. For example, suppose two states introduce different laws on some aspect of road maintenance which relates to driver assistance technologies, such as how often lane markings should be repainted. In this case, the sensitivity analysis on how technology availability influences the overall benefit calculation can be calibrated in the SIM tool by validating it against, for example, test drives of technology availability in these two states at some later time.

The SIM-tool developed here may be viewed as being complex compared to other approaches. However, the types of driving problems which this SIM tool was developed to address are complex problems shaped by multiple interactions between the DVET components. Almost by definition, a simpler approach would have been insufficient, as it would not address the actual real-world complexity of the problems. Moreover, in the future, the number of complex evaluation problems can be expected to increase. While most current active safety technologies today operate as stand-alone units in single vehicles, the future is likely to bring together a number of driver assistance technologies which operate in a coupled fashion, for example based on vehicle-to-vehicle communications. Evaluation of such driver assistance technologies will have to include multiple vehicles co-located in precise configurations for the relevant simulations of driving scenarios to take place. Conducting comprehensive and objective system evaluation and scenario development for these types of scenarios may well be too difficult to contemplate using any method other than via a full DVET simulation of the type presented in this research.

In the future, sampling from crash and roadway databases obtained from multiple states would be an advantage, because the roadway data would come directly from the same data that was used to develop the crash estimates. However, the purpose of the present work is to develop a methodology to estimate the safety impact of advanced technologies. Not all of the data needed for the methodology is directly available, so the Michigan data are used as an alternative, are developed to supply information of the type needed to allow the methodology to be demonstrated.

Analysis of Driver Alert Control (DAC) and Emergency Lane Assist (ELA) were not formally incorporated into the SIM benefits estimation (though see Appendix D for some preliminary analysis results). While there is no fundamental obstacle to achieving this in the future, both systems present challenges not present in the LDW system. The DAC system has potential safety benefits that rely on high level decisions made by real drivers balancing risks and goals in their normal lives. This cannot be modeled in any realistic detail, so some form of statistical or parametric model is essential, and any such model must be supported by actual field data. In Appendix D we show how part of the modeling can be performed by using detailed naturalistic driving data to emulate the real-world performance of the DAC system. However, it is clear that systems operating in a longer-term advisory or self-awareness mode are not as amenable to simulation-based analysis as those operating at the "sharp-end" of critical events such as lane departure, and the objective assessment of such systems remains a challenge for the future.

Compared to DAC, the implementation of ELA within the SIM benefits analysis is conceptually straightforward. However, there are several challenges to capturing the ELA functionality for full implementation in the SIM (Appendix D). ELA is a steering-based safety system providing a steering torque input to the vehicle, and its interaction with the driver is critical to the overall system performance – the driver, vehicle and technology elements of the DVET system come together at a single point, and the resulting complexity of interaction deserves a real depth of study, especially given that a wide range of naïve driver responses can be anticipated.

Overall, the results of this research provide a new methodology for the objective analysis of active safety technologies operating in the real-world. The approach may also help in closing the loop in the design and evaluation process for active safety systems. The underlying virtual crash population used is suitable for highlighting areas of high and low potential effectiveness of a given safety system, and can provide a new means, perhaps as a pre-production evaluation tool, to evaluate new safety concepts and make quantitative estimates of key sensitivities such as sensor performance and triggering thresholds.

References

Alm, H., and Nilsson, L. (1993), "Changes in Driver Behavior as a Function of Handsfree Mobile Phones – A Simulator Study," Accident Analysis and Prevention Journal, Elsevier Science, Vol. 26, No 4, pp 441-451, 1994.

AASHTO (American Association of State Highway and Transportation Officials), Task Force for Roadside Safety (2002), "Roadside Design Guide," 3rd edition, Washington D.C., AASHTO.

Aust, M.A., Gordon, T., Blower, D., Sardar, H., Isaksson-Hellman, I., Ivarsson, J., Jakobsson, L. (2009), "Requirements and Data Sources Needed for Validation of Component Properties and performance in Simulation Based Benefit Assessment of Driver Assistance Technologies," Paper No. 09-0438, Proceedings of the 21st International Technical Conference on the Enhanced Safety of Vehicles, Stuttgart, Germany, June 15-18, 2009.

Birk, W., Brännström, M. and Levin, D. (2006), "Method for determining a measure for evaluating the behaviour of a driver of a vehicle," Ford Global Technologies, LLC, European Patent Application EP 1 674 375 A1.

Blommer, M., Curry, R., Kozak, K., Greenberg, J., and Artz, B. (2006), "Implementation of Controlled Lane Departures And Analysis of Simulator Sickness for a Drowsy Driver Study," Proceedings of the 2006 Driving Simulation Conference, Paris, France.

Burgett, A., Srinivasan, G., and Ranganathan, R. (2008), "A Methodology for Estimating Potential Safety Benefits for Pre-Production Driver Assistance Systems," Sponsored by National Highway Traffic Safety Administration, Washington D.C., DOT HS 810 945.

Carter A. A., Burgett, A., Srinivasan, G., Ranganathan, R. (2009), "Safety Impact Methodology (SIM): Evaluation of Pre-production Systems," Proceedings of the 21st International Technical Conference on the Enhanced Safety of Vehicles, Stuttgart, Germany, June 15-18.

Dingus, T. and Klauer, S. (2008), „The relative risks of secondary task induced driver distraction," SAE Paper No. 2008-21-0001, Society of Automotive Engineers, Warrendale, PA.

Federal Highway Administration (FHWA) (2005), *Highway Performance Monitoring Field Manual*, chapter 1, page 1, May 2005.

Gillespie, T.D. (1992) "Fundamentals of Vehicle Dynamics", SAE Publications.

Gordon, T., Blankespoor, A., Barnes, M., Blower, D., Green, P., and Kostyniuk, L. (2009), "Yaw rate error: a dynamic measure of lane keeping control performance for the retrospective analysis of naturalistic driving data," Proceedings of the 21st International Technical Conference on the Enhanced Safety of Vehicles, Stuttgart, Germany, June 15-18.

Gordon, T., and Magnuski , N. (2006), "Modeling Normal Driving as a Collision Avoidance Process," Proceedings of the 8th International Symposium on Advanced Vehicle Control (AVEC), Taipei, Taiwan.

Green, M. (2000). "How long does it take to stop? Methodological analysis of driver perception-brake times," Transportation Human Factors **2**(3): 195-216.

Horne, J. and Reyner, L. (2001) "Sleep-related vehicle accidents: some guides for road safety policies," Transportation Research, F. 4, pp. 63-74.

Kozak, K., Pohl, J., Birk, W., Greenberg, J., Artz, B., Blommer, M., Cathey, L., and Curry, R. (2006), "Evaluation of lane departure warnings for drowsy drivers," Proceedings of the Human Factors and Ergonomics Society 50th Annual Meeting, Santa Monica, CA: Human Factors and Ergonomics Society, (CD-ROM).

LeBlanc, D., Sayer, J., Winkler, C., and Bogard, S. (2007): "Field test results of a road departure crash warning system: driver utilization and safety implications," Proc. 4th International Driving Symposium on Human Factors in Driver Assessment, Training, and Vehicle Design, Stevenson, Washington, U.S.A.

LeBlanc, D., Sayer, J., Winkler, C., Bogard, S., Devonshire, J., Mefford, M., Hagan, M., Bareket Z., Goodsell R., and Gordon T. (2006), "Road Departure Crash Warning System (RDCW) Field Operational Test Final Report," University of Michigan Transportation Research Institute (UMTRI),UMTRI reports No. 2006-9-1 and 2006-9-2.

Ljung, M., Blommer, M., Curry, R., Artz, B., Greenberg, J., Kochhar, D., Tijjerina, L., Fagerstrom, M., and Jakobsson, L. (2007), "The Influence of Study Design on Results in HMI Testing for Active Safety," Proceedings of the 20th International Technical Conference on the Enhanced Safety of Vehicles, Lyon, France.

MacAdam, C.C. (2003), "Understanding and Modeling the Human Driver," Vehicle System Dynamics, Vol. 40, pp 101-134.

NTSB (1999), National Transportation Safety Board, "Evaluation of U.S. Department of Transportation Efforts in the 1990s to Address Operator Fatigue", Safety Report NTSB/SR-99/01, Washington, DC.

The Mathworks (2009), Matlab/Simulink Users Manuals, Release 2009a.

Mechanical Simulation Corporation (2008), CarSim 7.1 Reference Manuals, May 2008.

NHTSA (2006), Cooperative Agreement for the Advanced Crash Avoidance Technologies Program (ACAT), U.S. DOT Announcement of a Request for Applications (RFA), CFDA Number: 20.600, Funding Opportunity Number: DTNH22-06-H-00011, June 2006.

NHTSA (2007a), National Automotive Sampling System (NASS) General Estimates System (GES) Analytical User's Manual 1988-2007, U.S. Department of Transportation.

NHTSA (2007b), National Automotive Sampling System (NASS) General Estimates System (GES) Coding and Editing Manual 2007, U.S. Department of Transportation.

Najm, W.G., Stearns, M.D.Howarth, H., Koopman, J., and Hitz, J. (2006), "Evaluation of an Automotive Rear-End Collision Avoidance System," DOT VNTSC-NHTSA-06-01, DOT HS 810 569.

Najm, W., Mironer, M., Koziol, J., Wang, J.S., and Knipling, R.R. (1995), "Evaluation of Target Vehicular Crashes and Potential ITS Countermeasures," DOT HS 808 263, Volpe National Transportation Systems Center.

Pack, R., Najm, W.G., and Koopman, J. (2006), "Exploratory Analysis of Pre-Crash Sensing Countermeasures," 2006 SAE World Congress, SAE Paper No. 2006-01-1438.

Sandin, J., and Ljung, M. (2007), "Understanding the Causation of Single-Vehicle Crashes: A Methodology for In-depth On-scene Multidisciplinary Case Studies," International Journal of Vehicle Safety, 2007.

Savkoor, A.R. (1990), "Tribology of Tyre Traction of Dry and Wet Roads," in Vehicle Tribology, Proceedings of the 17th Leeds-Lyon Symposium on Tribology, Leeds, UK, September 1990 (Elsevier, 1991).

Sayers M. W. and Riley S. M. (1996), "Modeling Assumptions for Realistic Multibody Simulations of the Yaw and Roll Behavior of Heavy Trucks," SAE Transactions, vol. 105, pp. 72-83.

Strang, G. (2003), Introduction to Linear Algebra, 3rd Edition, SIAM.

Wilson, B.H., Stearns M.D., Koopmann, J. and Yang C.Y.D. (2007), "Evaluation of a Road-Departure Crash Warning System," DOT HS 810 854, Volpe National Transportation Systems Center, National Highway Traffic Safety Administration.

Appendix A: Summary of Key Simulation Model Parameters

In this appendix we summarize the major parameters used in the simulation model for vehicle, environment and driver. Choice of these parameter values is described in Section 6, as are the parameters of the LDW system sub-model.

Table A1. Vehicle and Environment Models

Parameter Type	Description	Value
Vehicle geometry	front track (lateral distance between wheel centers)	1549 mm
	rear track	1549 mm
	tire width	215 mm
	wheelbase (distance between rear and front wheel centers)	2757mm
	steering ratio (ratio between steering wheel angle and road wheel angle)	16:1
Vehicle inertial properties	total mass	1567 kg
	yaw moment of inertia	2535 kgm2
	mass center height	406 mm
	mass center distance to front axle	962 mm
Road surface	coefficient of friction - dry road	1.0
	coefficient of friction - wet road	0.6
	coefficient of friction - off road	0.25
	lane marker width (used for animation only)	150 mm
	cross slope (null – road assumed flat)	0 deg
	vertical slope and curvature (null – flat road)	0 deg
	horizontal curvature (selected from table of 60 representative road locations, see Section 6)	variable

Table A2. Driver Model

Parameter Type	Description	Value
Manual Steering Control	τ : time constant associated with steering amplitude response (equation 6.5)	0.22-0.3 s
	T_c: time constant used in control output sub-model	0.02 s
	$\lvert\dot{\delta}\rvert_{\max}$: maximum steering rate achievable by the driver	10 rad/sec
Visual perception and information processing	$T_{\min}$: minimum preview time for scanning lane boundaries	1 s
	$T_{\max}$: maximum preview time for scanning lane boundaries	2.0-2.4 s
	T_{DA} : anticipation time used by driver to compensate for changing visual inputs	$T_{DA} = 2T_s + 0.22$

Parameter Type	**Description**	**Value**
	T_s: fundamental time step for information processing	0.05-0.1 (uniform distribution)
	$d_{\max}$: maximum available preview distance	120 m (daytime) 80 m (nighttime)
	$d_{\min}$: maximum available preview distance	2 m
	$dd_{\max}$: maximum relative distance change for association of existing and updated points in WM and VF respectively	10%
	$d\phi_{\max}$: maximum heading angle change for association of existing and updated points in WM and VF respectively	5 deg.
	age_max: maximum time boundary points may stay in working memory with being connected to existing visual image points	$10 \quad T_s$
	N_{WM} : Maximum number of records (rows) in the WM and VF data stores	40
	n_{pts} : number of boundary points captured from each lane marker	10
Response to LDW	T_{LDW} : internal delay time for LDW alert response	0.4-0.8 s (distracted) 0.3-0.4 s (fatigued)
Speed control	speed control gain	0.5
	reference speed	constant (at initial condition)
	AREF: reference acceleration for speed control	2.9 m/s^2
Attention switching	timing of visual attention (VA) to the road	
	Mean time for VA off	1.9 s
	Mean time for VA on	1.4 s
	Standard deviation for VA off	1.7 s
	Standard deviation for VA on	1.0 s
	STR_MAX: Steering rate threshold above which driver will become and remain attentive	15 rad/sec

Appendix B: Summary of Major Simplifying Assumptions and Limitations of the Estimation Process

This appendix presents a compiled summary of the major simplifying assumptions used in the research, including data limitations and opportunities for improvement via future research. These are grouped into five major categories: driving scenarios, modeling, batch simulations, analytical method, and finally the additional challenges presented for the driver alert control and emergency lane assist systems. Where appropriate, references to the corresponding report sections [§] are given.

B.1. *Setting up Driving Scenarios*

Driving scenarios (DS) are simplified representations of real-world driving conditions. They are configured based on a number of related considerations: the design intent of the system, the association with observed naturalistic driving and crash conditions, and the feasible scope of the overall vehicle-system-driver-environment model. Limitations relating to driving scenarios mainly result from uncertainty and misalignment of defining parameters and conditions relevant to these different perspectives. A simple illustrative example is the use of the turn signal. For design intent, the turn signal indicates a deliberate lane change, and the system is suppressed. The turn signal is directly available in the naturalistic data, but since many drivers change lanes without its use, it provides an imperfect measure of lane-change intention. In the GES crash data it is completely unknown, and finally in the event simulation, it can be set as a logical flag. However, without a complex decision sub-model that triggers a deliberate lane change, such a flag has limited applicability and has not been implemented (in the model we assume there is no intention to change lanes and also no turn signal use). In all four cases the status of the turn signal is different, yet in all cases it has some degree of relevance to the definition of the DS.

We now summarize the main limitations and deliberate simplifications associated with setting up driving scenarios and representing them in simulation:

- The outcome of DS simulation is either crash or no-crash; other outcomes, such as the vehicle drifting off the highway and coming to rest safely, are not considered [§7.1]
- Highway and environmental factors used in the definition of the DS are relatively coarse-grained, and factor correlations are provided only via random sampling of Michigan crashes, and so national geographic variations are not included [§4.4]
- The pre-crash maneuver variable P_CRASH1 in NASS GES has limited fidelity; uncertainty over the driver's intention (drift out of lane or active lane change) means that accurate counting of relevant crash numbers is currently impossible [§4.2.2]
- Pre-crash travel speed is uncertain in the GES data, so again "ACAT-relevant" crash number filtering is uncertain; note however that this problem is largely addressed via the naturalistic data analysis, where virtual crashes at speeds below the operational threshold of the system are excluded [§4.2.2]

- Crashes coded with associations to alcohol or drug use are excluded, even though it is possible that LDW might offer benefits in such cases [§4.3]
- The vehicle type was restricted in light passenger vehicles, even though other classes of vehicle, such as light truck, are also likely to benefit [§4.2.2]
- GES coding is based entirely on police accident reports, and there is significant variability and uncertainty in some of those data, e.g. coding of driver fatigue [§ 4.4, 9.4]
- To contain the number of scenarios included, some factors relating to lighting, weather and roadway were combined [§4.4]
- Initial conditions for DS event simulation were based on samples from naturalistic driving data – some restrictions on sample size were imposed, and data was obtained from a limited geographic region (mainly SE Michigan) [§4.5]
- It was not feasible to sample from fatigued driving as a population separate from normal or distracted driving [§7.2]

B.2. *Model-related Limitations and Simplifications*

Models are inherently simplified representations of the real world, but are fundamental to our ability to understand, simulate and predict. Even physical tests only simulate the real world processes of interest. To be useful, models have to provide adequate fidelity in terms of processes and mechanisms, but otherwise should be as simple as possible; increasing the detail in any model beyond its basic purpose is generally wasteful, especially in terms of the need to test, validate and provide supporting parametric data. In spite of the approach adopted in the project, there are large numbers of parameters and some influential model assumptions that have not been fully represented (e.g. driver emergency response and sensor performance). Most of the model-related simplifications have been made to complete the overall project within the available time and resources available.

We now consider the specific model-related limitations according to the sub-models for the vehicle, system, driver and environment.

Vehicle Model

- It was not feasible to simulate a wide range of light vehicles without greatly increasing the time needed for conducting batch simulations, so a vehicle model based on a single mid-sized sedan was used throughout – the simplifying assumption is that this single vehicle is sufficiently representative of the range of light passenger vehicles considered [§6.1]
- The CarSim model was not identical to the vehicle model used in the VIRTTEX driving simulator; although differences were small, some discrepancies in driver model calibration and validation may have been introduced in this way [§6.1]
- No external disturbances (due to road roughness, crosswinds etc.) were included in the CarSim model – the real vehicle may therefore have a greater tendency to drift in the lane compared to simulated drift events [§7.2]

- The simulation model did not include stability control; while this is the most appropriate choice for the baseline (unequipped) population, for the LDW-equipped vehicle it would be typical that stability control would also be implemented [§8.2]

ACAT System Model

- There is no realistic sensor/estimation model included, so whenever it is enabled the system accurately senses the location of the lane markings; in reality there may be some error in this process. (Note that system unavailability – including cases when the sensor cannot detect the lane markings at all – is included in the safety benefits analysis) [§6.2]
- Driver choice is not represented in enabling the LDW system; in reality the driver can decide to choose low or high sensitivities, or even turn the system off completely. While there was some simple consequence analysis of any unwillingness of the driver to switch the system on, it was not part of the formal objective analysis[11]. [§6.2]
- The active safety systems use proprietary algorithms and will therefore be approximated by anyone other than the OEM.

Driver Model

The driver is the most complex element of the overall system, and it is clear that the model used is highly simplified. High level decision making relating to distraction and fatigue are not within the scope of this model, and it seems that even the best feasible model would use statistical components to represent such behaviors. That said, the model does produce plausible responses and combines error-prone behavior with the ability to recover stable lane-keeping. Specific limitations are as follows.

- Driver distraction is represented as an idealized stationary random switching process; with additional research and expanded analysis of naturalistic data sources this aspect can be greatly improved, with the potential to include context-sensitive distraction behavior [§6.3]
- The modeled process of visual processing and recognition of lane markers is not sensitive to the several effects such as road scene complexity, lane marker quality or ambiguity in the lane boundaries [§6.3]
- There is no bio-mechanical sub-model to characterize the physical actions of the driver in terms of muscle force and limb motions etc. [§6.3]
- Although the model has the capability of including non-driving workload and cognitive distraction, this has not been included, and it is not clear whether these factors would have a significant effect on the benefits estimates [§6.3]
- The available range of lane boundary perception (assumed to be 80-120 m) has not been properly calibrated according to weather and lighting conditions [§7.4]
- The model assumes that the driver is experienced enough to have achieved a basic control adaptation, in the form of a best compromise between responsiveness and stability; thus no special consideration has been given to novice or teen drivers [§6.4.1]

[11] It is worth noting from recent customer clinic data (Volvo Cars, North America) that over 80% of drivers report having LDW "always on" or "usually on". The data is considered too limited to include in the formal analysis.

- Each block that performs information processing in the model is assumed to operate in "discrete time", i.e. it takes a fixed time to execute a single cycle of processing. While typically true of computer systems, this is a simplifying assumption for the biological information processing in the human brain [§6.3]
- The perceived "hardness" (h) of a boundary point is kept uniform in the model; more realistically h should be higher for a lane marker adjacent to a narrow gravel shoulder than for a wide paved shoulder [§6.3]
- Related to the previous point, the driver only responds to the boundaries of the initial travel lane, and not to other boundaries or threats (e.g. oncoming traffic) [§7.5]
- Parameter estimation and validation for the driver model is based on data from the VIRTTEX driving simulator; this process was somewhat limited due to the small number of subjects and the variable performance of naïve drivers in a complex simulator environment [§6.4-5]
- Although the model does generate large and rapid steering responses related to large excursions from the lane, no detailed calibration or model validation was available to characterize this type of startled or emergency response [§7.2]
- It is assumed that the driver control response is dominated by the steering intervention; an emergency braking response was not included in the model [§6.3.6]
- VIRTTEX-derived reaction times in the distracted driver model are dependent on the actual number-reading task used (since the numbers disappear from the display there is little opportunity for the driver to delay the task) ; this may or may not bias the reaction times relative to the naturalistic driving environment [§6.4.2]
- Driver compliance: it is assumed throughout that the driver (model) understands the nature of the LDW warning, and is motivated to respond to it. This assumption is key to estimating the overall system effectiveness. Currently there is little data to support this part of the safety benefits analysis and only simplified assumptions have been made [§9.4]
- Driver compliance and response times may be inter-connected (e.g. an annoying alert and the driver responding slowly) but analysis of this was considered to be outside the feasible scope of the study. [§9.4]
- The driver model does not capture driver motivations and adaptations to real driving contexts with real consequences. These motivations and adaptations can have substantial effects on the safety impacts of a given technology.
- Many assumptions were required to estimate availability. Actual availability will depend on specific characteristics of a given active safety system.

Environment Model

- The geometric road and off-highway parameters were derived from geo-location of Michigan crashes; while there is no assumption that Michigan is "typical" of the national experience in terms of weather severity or the physical condition of the roads, statistical correlations between the highway factors in the environmental model are indeed based on Michigan examples [§4.4]
- Only three friction levels were applied [§7.4]

- Road segments designated as “straight” in the GES coding were idealized as perfectly straight road segments in simulation [§7.1]
- In simulation the road surface does not contain details such as grade, roughness, crown and super-elevation [§7.1]
- The lack of geographic location data (or detailed highway characteristics) in the GES data meant that lateral curvature was based on a library of 60 “typical” curved roads [§7.1]
- Site specific clear-zone data was not available, so use was made of traffic and speed data for each location; clear-zone distance we then estimated from the MDOT Road Design Manual [§7.5]
- For both divided and undivided multi-lane highways, we simply assume two lanes in either direction [§7.1]
- The vertical geometry in the off-highway zone is also presumed flat; future research should represent the off-highway condition in greater detail as it strongly affects the ability of the vehicle to recover from a road departure. [§6.1]

B.3. *Batch simulations*

- To limit the number of simulations, attention was restricted to the most common 25 scenarios, accounting for 91% of relevant crash types [§7.1]
- For simulation efficiency, a stratified sampling scheme was devised based on a single “control variable” (inverse time to lane crossing, ITTLC); the sampling scheme has limited resolution however, using 10 bins of ITTLC values and based on an initial sample of 50 vehicle conditions per scenario (5 examples per bin) [§7.2]
- Due to lack of data, no separate ITTLC distributions could be constructed for fatigued drivers [§7.2]
- No corresponding sampling scheme was available for driver state (distraction or fatigue level); in future studies such a development could greatly reduce the uncertainty of the benefits analysis
- The batch simulations are to represent a single event as a potential conflict or crash, and the stopping conditions for any such event are somewhat arbitrary [§7.4]
- Differences in lane width between the naturalistic driving data and the simulated events had an unplanned effect of the ITTLC distribution for initializing batch simulations. It is assumed that this had only a minor influence on the safety benefits estimates. [§8.1]

B.4. *Safety Benefits Analysis*

There are a number of limitations arising from the application of the analytical method.
One important aspect not specifically regarded as a limitation is associated with the nature of the benefits estimate (numbers of crashes reduced). The estimate is not a prediction about the real world, but rather a measure of change under hypothetical circumstances; if it were possible to go back in time and retro-fit passenger vehicles with this technology, then it is likely that a certain benefit would have accrued. To convert this kind of information into an actual benefit prediction would require substantial further modeling, in terms of traffic and population trends, technology trends in other parts of the vehicle, changes in highway design, changes in vehicle use, etc., and especially in terms of the rate of penetration of the technology into the vehicle fleet. These important aspects have not formed part of the presented SIM methodology. Other restrictions, limitations and fruitful opportunities for further research noted in the report fall under two main categories: overall methodology and crash metric.

Methodology

- It is assumed that the nature of each driving scenario is unaffected by the presence of the safety technology, and that the technology does not introduce any new driving scenarios [§2.2]
- The analysis did not include any effects of driver adaptation to the safety technology, e.g., increased inattention to the driving task because the driver expects the LDW system to warn when necessary [§2.2]
- Crashes involving intoxicated, suicidal, and deliberately aggressive drivers were excluded, to the extent possible, from the analysis, and the performance of drivers in such states were not modeled
- Except for the possibility of driver alert control (see below) a unit exposure ratio is assumed, i.e. the presence of the technology does not affect the frequency at which the various driving scenarios occur [§9.4]
- We assume in any scenario that the α and β parameters are statistically independent. [§2.3.2]

Crash Metric

- Crash numbers for the virtual population are based on a simplified measure of crash risk; further work is needed to assess and improve the validity of this metric [§7.5]
- Related to the previous point, we do not represent the detailed locations of potential collision objects, or assess the density and type of off-highway obstructions
- As part of the crash metric, a somewhat arbitrary "boundary layer" of 30 cm minimum intrusion is included; this parameter also needs further assessment and refinement in the future [§7.5]
- Crash metric parameters (k) are estimated from crash outcome frequencies, and no corresponding exposure measures are used. Thus k values are only estimated relative to each other within each of the 6 major road categories. In the future it should be possible to

use exposure data and naturalistic data to validate and improve estimates, and also estimate absolute values of crash risk [§9.1]

- The SIM analysis predicts system effectiveness is higher under adverse conditions with the exception of *curved road* conditions – this suggests future research is needed based on naturalistic driving data to confirm and/or better understand these apparent trends. [§9.2]

B.5. *Driver Alert Control and Emergency Lane Assist*

As mentioned in Section 10, some progress was made with assessing the methodological challenges of these systems, and future work is certainly feasible to fill in the gaps and attempt an estimation of potential safety benefits of these types of safety technology. See also Appendix D below. Here we note the main outstanding research questions.

Driver Alert Control [Appendix D]

- A field study is needed that provides an independent measure of vigilance and fatigue, for example via real-world measurement of eye closure and gaze direction.
- The willingness and availability of opportunities for a driver to respond to DAC warnings needs to be evaluated.
- The effectiveness of different strategies (rest break, short sleep, caffeine, engaging in conversation, etc.) needs to be evaluated in the context of the objective analytical framework; the persistence of any improvement in vigilance is an important aspect of any such evaluation.

Emergency Lane Assist [Appendix D]

- Detailed studies are needed to characterize the interaction between the driver and the vehicle system when ELA provides a steering torque intervention (e.g. to give information on the extent to which drivers oppose such a disturbance). Preliminary testing carried out in the VIRTTEX driving simulator with non-naïve drivers demonstrated that this interaction has significant effects on vehicle trajectories and hence crash risk.
- In a system of this type, the performance of the sensor system is vitally important, as both lane markings and the relative position and velocity of the principal other vehicle is required. Track-based and real-world testing is needed to characterize this performance, and the results need to be included in a sub-model in the SIM tool.
- In practice, for a complete benefits assessment of ELA, it is necessary to fully define the position and motion of the principal other vehicle in the simulations, and to provide corresponding scenario weights based on real world conditions (e.g. how likely is it that a POV will become a detectable threat during a lane excursion). While several approaches are available to achieve this, it was not possible to explore this within the scope of the current project.

Appendix C: Parameter Estimation and Validation for the Driver Model

This appendix provides supporting detail for the discussion in Section 6.4, and the values given in Table 6.2. Parameter estimation and validation for the driver model is mostly but not entirely based on data from the VIRTTEX driving simulator. We consider how the values or ranges for each of the parameters in Table 6.2 were chosen

As mentioned in Section 6.4, T_c cannot be determined independently of others, and a reasonable "initial set value" was chosen. Only if T_c were improperly large is it possible that T_c could not be reduced sufficiently to compensate. Another effect of large T_c would be to reduce the maximum frequency of sharp steering changes as the filtering effect of the driver's mechanical output becomes more pronounced (i.e. filter bandwidth is reduced). Choosing T_c to be a small value (20ms) renders the limiting speed of driver motor response to be substantially faster than the other processes in the driver model, so its effect on the DVET system response is unimportant.

Before we consider VIRTTEX data to set and validate parameters for the driver model we emphasize that the source data from VIRTTEX is highly variable due to the influence of the naïve subject driver, and parameter fitting is neither a precise nor unique process. The expectation here is that after some amount of tuning, reasonable and representative parameter values will be found for the overall simulation model to provide outcomes which are similar to those measured in the VIRTTEX driving simulator.

VIRTTEX data includes video files synchronized via time stamps with digital data. Figure C.1 shows a sample image from the video. From the video it is possible to estimate the driver's direction of gaze and provide a rough estimate of overall reaction time. However, following the approach of Kozak, et al. (2006), reaction time (RT) is more reliably estimated from steering response initiation, i.e. the time between the LDW warning being issued and a defined level of the steering response.

Figure C.1. Video Image from the VIRTTEX Data

In the VIRTTEX study, to create true positives for a distracted driver, a extraneous “yaw deviation” was applied to the simulated vehicle – a small yaw angle deviation was applied, except the motion control system was deliberately not excited; hence drivers experienced no associated lateral acceleration that might have otherwise have alerted them to the disturbance. This method was found effective, as most drivers believed the lane departure was presented as being natural and not due to an external intervention. Table C.1 lists the main data channels available from the VIRTTEX simulator tests.

Table C.1. VIRTTEX Data channels used to validate the driver model.

Name	Units	Description
current_scenario	--	1: Distraction only 2: Distraction with yaw deviation 11: LDW(left), false positive, no distraction 12: LDW(left), false positive, distraction 13: LDW(right), false positive, no distraction 14: LDW(right), false positive, distraction
distraction_state	--	> 0: Voice-prompt for distraction task < 0: Negative of number in distraction number-reading task (i.e., Numbers displayed to driver are positive. Six numbers, each displayed for 0.5 seconds)
euler_yaw		Vehicle yaw. Need to subtract 90 from euler_yaw to align 0 yaw with vehicle driving straight forward at beginning of the drive. (euler_yaw – 90) increasing: CW rotation (view from top of vehicle)

Name	Units	Description

lane_pos	Ft	Position of vehicle in lane.
		Lane widths are 12.5 ft.
		Define:
		lw = 12.5; // lane width
		laneExcursion = lane_pos - (lw/2);
		
		laneExcursion is deviation to right (> 0) or left (< 0) of vehicle from center of lane.

lat_accel		Lateral acceleration of the vehicle
ldw_type	--	0: No LDW Non-zero: Different types of LDWs
ldw_warn	--	0 => 1 transition : activate LDW corresponding to ldw_type 1: ldw for deviation to the left 2: ldw for deviation to the right
steering	Degrees	Steering wheel angle
Subject	--	Subject number
T	Seconds	Simulation time
Tapa	Percent	Throttle Angle
Vmph	Mph	Vehicle speed
Yawdev	--	0: No deviation 1: yaw deviation (ramp to 2-degree deviation over 3 seconds. Deviation goes to right or left depending on whether vehicle is closer to right or left lane line)

Parameter estimation involves several basic tests of the model followed by matching with the VIRTTEX events. The tests of the model relate to the control loops represented in Figure C.2. Here the driver model is once again broken down into information processing and control application blocks (compare Figure 6.18), however here it is convenient to expand the domain of information processing to include basic image processing (BIP) and also compute the single critical yaw rate (see above). This allows us to define a fast "inner control loop" that deals with vehicle control relative to a yaw error feedback, and an "outer control loop" that also includes BIP/Information processing. In fact both control loops also include the vehicle yaw response, so when considering the overall response times we should include some measure T_v of the vehicle yaw response time (see below for how this is defined and estimated). Based on the timing information given in Section 6.3.7, we estimate the overall delay time in the full ("outer") control loop that includes *BIP* and *information processing* to be

$$T_{outer} = 2T_s + 2T_c + T_v$$

Again this is expected to be a minimum value, depending on other conditions in the simulation. On the other hand, the inner control loop does not make use of the housekeeping process (*HK*), so in this case we expect a reduced time:

$$T_{inner} = T_s + 2T_c + T_v$$

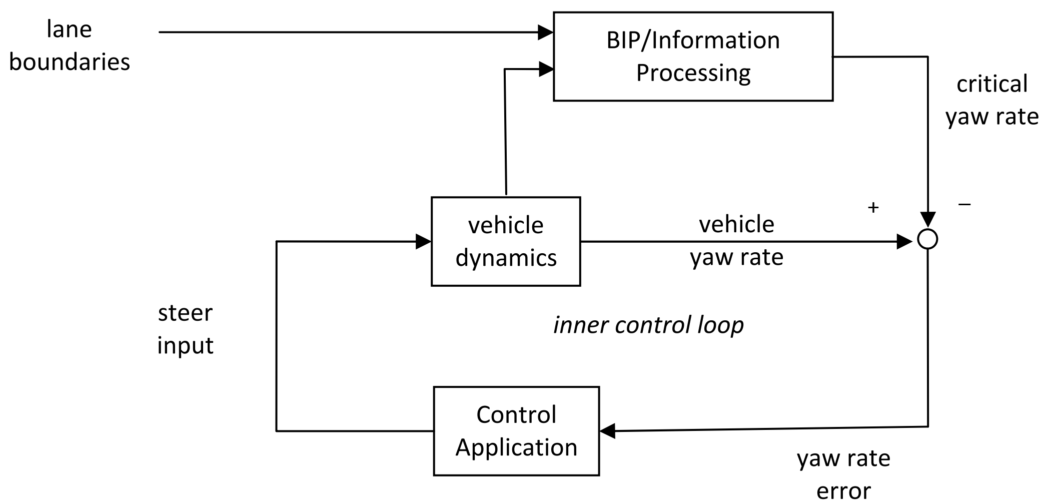

Figure C.2. Structure of the Driver Model Control Loops Used in Parameter Estimation

With this background preparation, the following three tests are applied to the vehicle and/or driver model to provide reference data in the parameter estimation process:

Test 1: Open-Loop Step-Steer

This is a simple test of the vehicle model to determine an estimate for T_v. A step input is applied to the steering and the yaw velocity $r(t)$ is recorded. The following time constant

$$T_v = \frac{r_{ss}}{\max(\dot{r})} \tag{C.1}$$

is a measure of the initial response time of the vehicle alone, and will be referred to as the *rise time*; here r_{ss} is the steady-state yaw rate following the steer input and $\max(\dot{r})$ is the peak yaw acceleration during the event.

Test 2: Inner Loop Yaw Response Test

Starting from a steady-state condition, an artificial step change in critical yaw rate (in Figure C.2 a step change is applied to the signal shown as output from BIP/Information processing) is fed into the control part of the driver model and the corresponding rise time T_r is found. In fact the BIP/Information block is not active in this test; rather, a test signal is used to represent changes that may arise in the critical yaw rate. T_r is defined via equation (C.1) in the same way as T_v above (i.e. the steady-state change in yaw rate divided by peak yaw acceleration). In this test the response is modulated by the driver in as he/she controls the vehicle's yaw motion, and the response is naturally slower than for the vehicle alone; thus we expect $T_r > T_v$.

Test 3: Outer Loop Yaw Response Test
In this test an artificial yaw deviation is introduced into the model, much as in the VIRTTEX simulator tests. Initially driver attention is switched off; after the attention signal is switched back on, the response is assessed for initial transient and also further steady-state values. To enable direct comparison with the VIRTTEX data, initial conditions were directly obtained from that source, and the initial attention switch was adjusted manually to directly match the VIRTTEX response.

Here the VIRTTEX data was used to estimate an overall reaction time (RT): this is determined directly from the steer response in VIRTTEX, and confirmed by video review (see below). The overall reaction time is used to estimate the LDW response time, once other parameters have been estimated.

The critical model parameters are summarized in Table C.2, while a full set of driver model parameters is given in Appendix B. In this table, $T_{\min}$ and $T_{\max}$ define the minimum and maximum preview times for the lane boundary. The driver anticipation time T_{DA} is used to compensate for driver delay during tracking of the road boundary – equation (6.6). T_{DA} is therefore set equal to the outer loop delay time to achieve this:

$$T_{DA} = T_{outer} = 2T_s + 2T_c + T_v \tag{C.2}$$

Table C.2. Summary of Key Driver Model Parameters and Their Initialization

Parameter	Initial set value [sec]	Characteristics/comments
$T_{\min}$	1	Influences stability in a transient lane change; must exceed T_{DA}
$T_{\max}$	3	Increasing $T_{\max}$ produces more frequent and larger steering interventions during regular lane keeping (star point moves further from the vehicle)
τ	0.35	τ^{-1} is related to steering control action via equations (6.3) and (6.5). To preserve stability it must be increased as other delays increase, thus reducing the steering gain.
T_{DA}	see equation (C.2)	Note subject to independent variation
T_s	0.05	Affects overall delay in both control loops
T_c	0.02	Affects inner loop delay, but is confounded with T_c in the overall control loop – fixed at the initial set value
T_{LDW}	no initial set value - see below	Reaction time for a distracted driver to begin visual data capture after an LDW alert

The “initial set values” shown were based on initial simulation tests used to confirm that the model operated in a stable and reasonable manner; the tests and iterations presented below were then used to calibrate the key parameters in a more systematic fashion.

Results of Test 1: Open Loop Step Steer Test

A number of open loop step-steer tests of the simulated vehicle were conducted, keeping the resulting lateral acceleration below $2ms^{-2}$. In this low acceleration regime the response timing is that of a linear dynamic system and independent of amplitude. According to the results in Table C.3 the resulting rise time T_v in the yaw response is reasonably constant across a wide speed range; here we show the steady state yaw rate, r_{ss}, and corresponding T_v for five initial speeds, U_0. Based on these results a constant value T_v=0.18 was estimated.

Table C.3. Rise time of yaw rate response to unit step steering wheel input

U_0 [m/s]	r_{ss} [rad/s]	T_v [s]
10	0.0552	0.1632
15	0.0643	0.1836
20	0.0675	0.1824
25	0.0673	0.1788
30	0.0653	0.1716

Results of Test 2: Inner Loop Yaw Response Test

In this test we use relative yaw stability as a criterion for choosing the steering gain parameter τ . A value for T_s must be selected to conduct the test; since it is unknown at this point we allow values to span a broad range, [0.05, 0.1, 0.15, 0.2] seconds. The choice of T_s is left open until after Test 3 is completed. For each value of T_s we carry out the following steps:

a) Choose conservative values for the preview time, in the range 2-3 seconds; i.e. set $T_{\min}$=2s, $T_{\max}$=3s. This is to avoid any outright instability that can arise with shorter preview times and un-tuned steering gain. (Note however, once the steering gain parameter has been tuned, a shorter minimum preview time $T_{\min}$ = 1 s is permitted).
b) Obtain responses to a step yaw rate input for a range of values of τ (0.1 - 0.5s in steps of 0.02s).
c) Compute the settling time for each τ, and find the best value τ^*, associated with minimum settling time.

In Figures C.3 and C.4 we show typical results, here for the case $T_s = 0.05$. The rise times and settling time were computed for the range of τ values. We see the rise time increases uniformly as τ increases (this is expected: the overall response slows down as the steering gain reduces) but the settling time, and hence the relative stability of the inner control loop, has a minimum value.

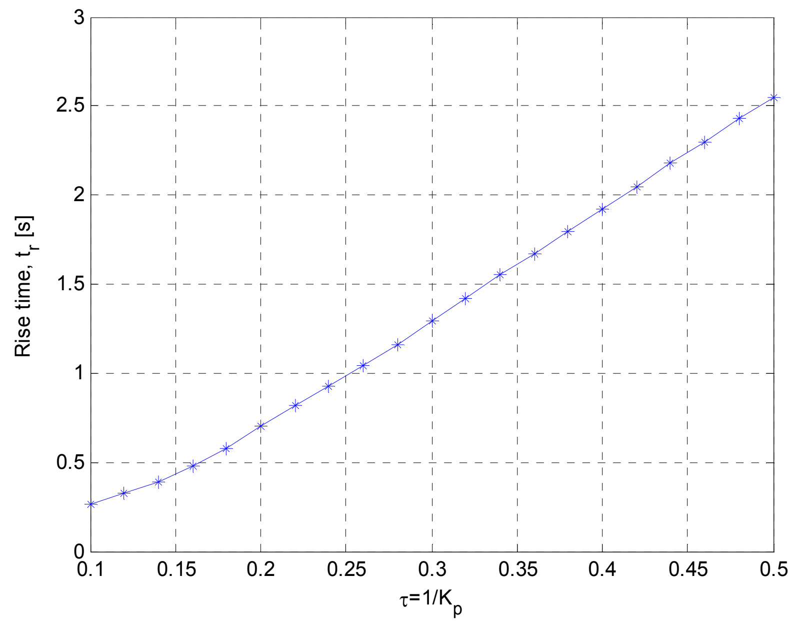

Figure C.3. Rise Time for Different Control Gains ($T_s = 0.05$)

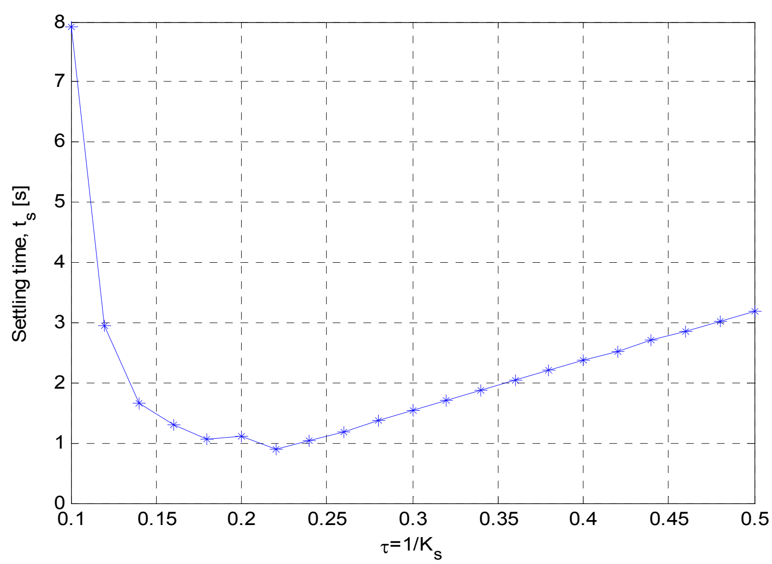

Figure C.4. Settling Time for Different Control Gains ($T_s = 0.05$)

This provides an optimum value $\tau^* = 0.22\sec$ for the chosen value of T_s. Table C.4 summarizes all the results, together with the rise times. Note that, as mentioned above, the inner loop response is slower than for the vehicle-only result (Test 1), and also the speed of response of the coupled vehicle-driver system becomes slower as the basic driver processing time step is increased. The tuning shown here is analogous to a driver adapting to a particular vehicle so that the best compromise between responsiveness and stability is achieved; a basic assumption of the model is that the driver is experienced enough to have achieved this basic control adaptation. As required by stability, drivers with slower inherent response are expected to drive with a correspondingly higher value of τ , in other words with a lower gain in the steering response.

Table C.4. Summary of Test 2 Results

T_s [s]	τ^* [s]	Settling time [s]	Rise time T_r [s]
0.05	0.22	0.9	0.82
0.1	0.3	1.19	1.09
0.15	0.36	1.36	1.25
0.2	0.44	1.67	1.53

Results of Test 3: Outer Loop Yaw Response Tests

Using the steering gain time constant τ^* found from Test 2, T_{DA} is computed by equation (C.2). The value of $T_{\min}$ is then set back to its initial setting of 1 second, providing a wide range of boundary points between minimum and maximum preview for the driver model to respond to during lane keeping. In fact the star point (Section 6.3.1) rarely approaches the minimum preview point, only when there is a significant drift out of lane does the nearest point becomes the most critical. The value of $T_{\min}$ does therefore affect the transient driver response during a yaw deviation event, but the value of T_s also affects this. Without more extensive data it was decided to fix $T_{\min}$ at the nominal 1 second value and adapt T_s to capture the timing and amplitude of the transient steering corrections.

On the other hand $T_{\max}$ does have a specific effect on the steering control responses – for larger values of $T_{\max}$ the amplitude of steering corrections in steady-state lane keeping is reduced as $T_{\max}$ increases. Based on this condition, root-mean-square values of steer angle were matched to the driving data by adjusting $T_{\max}$ in the simulation. This is done for each of the values of T_s and the results are given in Table C.5. Again, as expected, the required maximum preview time increases as driver model information processing time increases.

Table C.5. Driver Anticipation (T_{DA}) and Maximum Preview ($T_{\max}$) Times Estimated from Test 3.

T_s [s]	T_{DA} [s]	T_{max} [s]
0.05	0.32	2.0
0.10	0.42	2.4
0.15	0.52	2.6
0.20	0.62	3.0

At this point, the transient responses seen in the VIRTTEX data are used in the calibration, comparing results with simulations obtained using different T_s values. Not surprisingly the main differences relate to the speed and amplitude of response. The following four figures show sample results. Given the considerable variability between drivers and events, a precise match is beyond expectation, so we are interested in establishing a plausible match in terms of amplitude, decay and frequency of steering oscillations.

In each of Figures C.5-C.8 the upper plot gives steering wheel angle and the lower plot is steering wheel rate. In the upper plot, to improve visual comparison, a dashed line has been fitted to the VIRTTEX steering response, indicating the overall amplitude decay of the real steering oscillations. Except for $T_s = 0.20$ the amplitude decay also fits the driver model quite well; in the final case the steering response clearly does not decay fast enough. The overall best match is for $T_s = 0.05$, having similar amplitudes and frequencies in both plots. By comparison, in Figure C.7 the frequency is too low, and the simulated peak steering angles are too small. Figure C.6 gives the second best match; steer angles are similar in magnitude, response is slower and steering velocities lower than for the sample data, but the differences do not seem unreasonable. From this we propose to use a range $0.05 \le T_s \le 0.1$ for the basic driver model sample time.

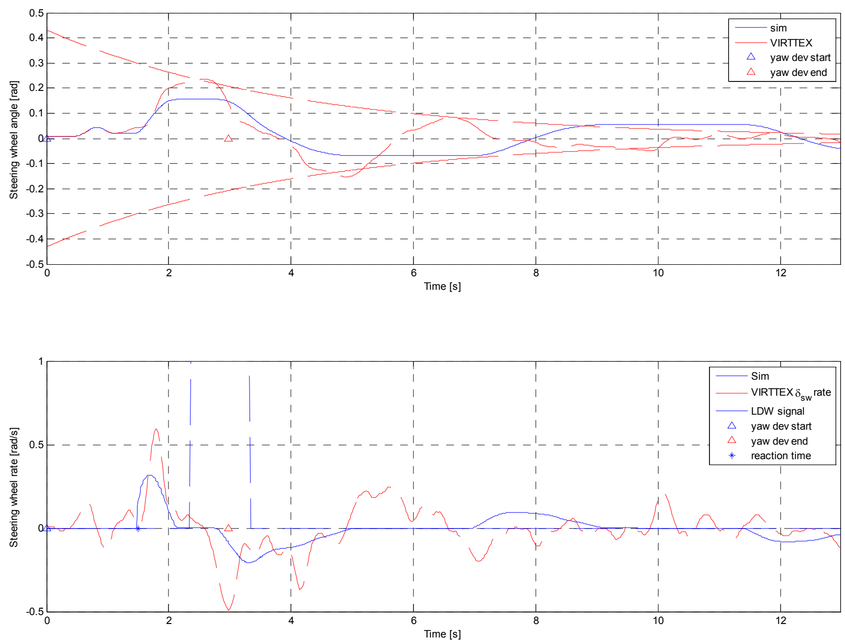

Figure C.5. Steering Responses (T_s=0.05s).

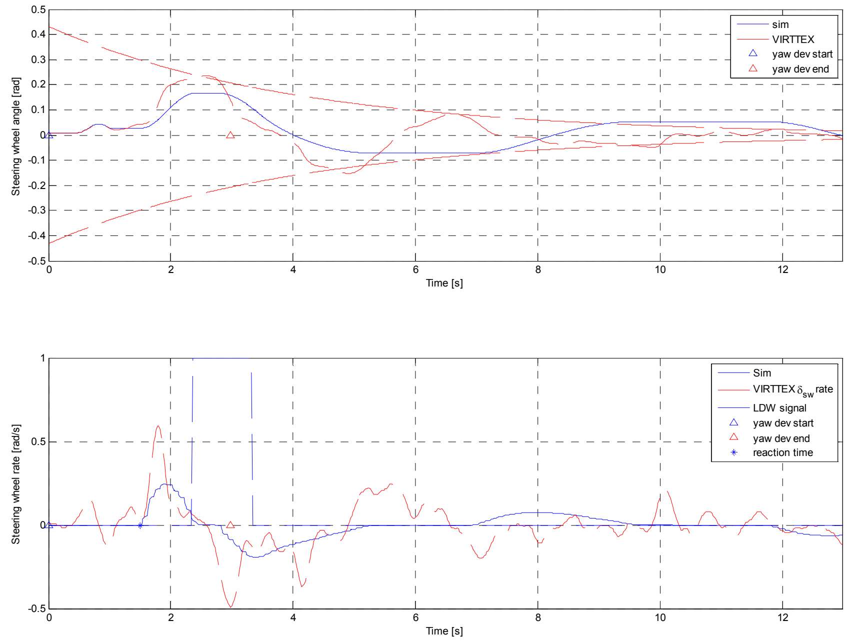

Figure C.6. Steering Responses (T_s=0.1s)

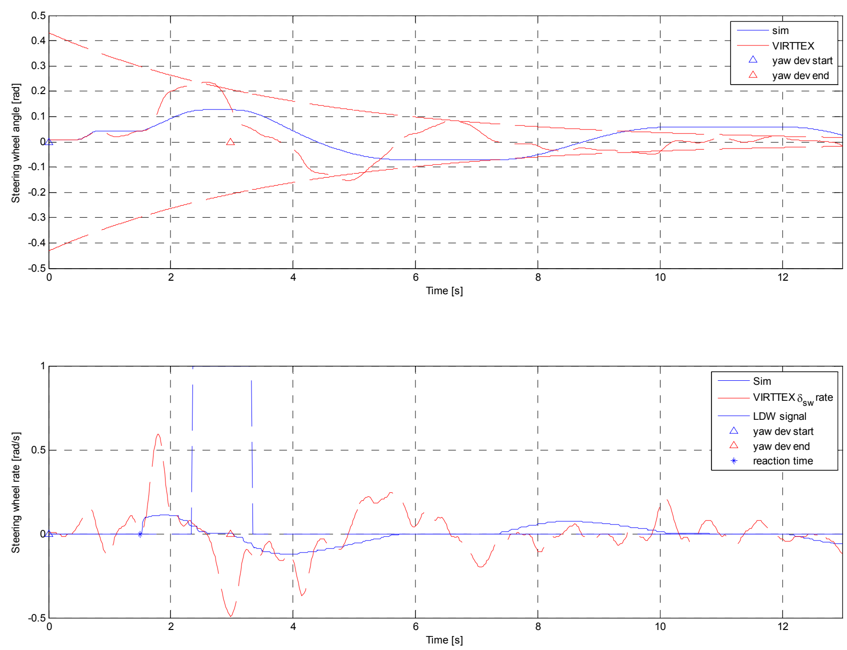

Figure C.7. Steering Responses (T_s=0.15s)

.

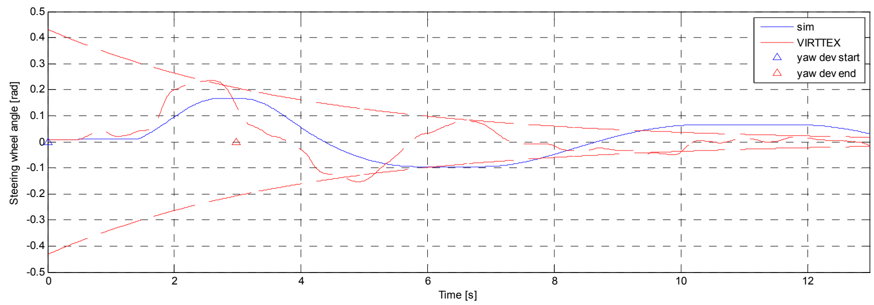

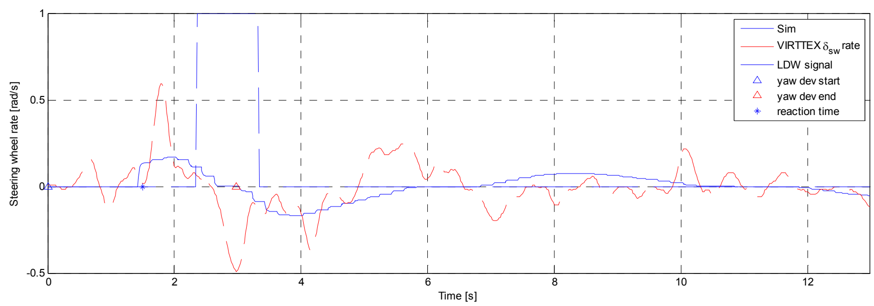

Figure C.8. Steering Responses (T_s=0.20s).

Estimation of Drivers' Reaction Time to LDW Audible Alerts

Continuing with the data from VIRTTEX Study 2 (distracted driver) the key parameters in Table C.2 are now all estimated except for the LDW response time T_{LDW}. We note that the range (0.05-0.10s) chosen for T_s is based on limited data and so the precise limits could certainly be refined in future studies. However, values much smaller than 0.05s are excluded since unrealistically short overall reaction times are then possible, while values much larger than 0.10s give unrealistically slow steering response to yaw deviation events.

We now consider the estimation of T_{LDW} based on a comparison of overall reaction times between the model and the VIRTTEX driving data. Reaction time is a function of many sub-processes within the driver model (and within the real human driver), and since it is not a fundamental component quantity set by a single parameter, its value depends on making a functional definition. Starting with the data available from VIRTTEX we define two particular reaction times.

The first is a video-based reaction time (RT_V). From recorded data we know the onset of LDW warning t_0. From video it is possible to estimate the time t_1 when the head first returns to looking forward after the number reading task. Then

$$RT_V = t_1 - t_0 \tag{C.3}$$

Estimating the timing of the head turn t_1 is not precise and in some cases the head turning action is not obvious to see; however the results of RT_V estimation is useful as a test of reasonableness for a more objective measure based on steering response.

Steering-based reaction time RT_S was estimated from steering wheel angle similar to a method reported by Kozak et al (2006). In that paper the *RT* was found by detecting a zero crossing of steering wheel acceleration, $\ddot{\delta}_{sw}$. While this has the advantage in detecting the earliest possible onset of steering motion, it was found to be very sensitive to random variations in the steering, at least without applying additional filters to the steering measurement. The preference was to use a more stable signature, even if that is somewhat delayed. We define t_2 as the time at which steering wheel rate, $\dot{\delta}_{sw}$ reaches a maximum in the direction of the corrective steer. Then

$$RT_S = t_2 - t_0 \tag{C.4}$$

and we add the condition that t_2 should precede the point of maximum lane deviation. Compared to (6.11) or the initial zero crossing of steering acceleration this does introduce additional biomechanical response delay; but since we also have a representation of this in the model (via the steering control gain equation (6.3) and in the control action delay parameter T_c) a direct comparison between model and experiment is available, and hence estimation of T_{LDW} is not affected by this additional delay.

Table C.6 summarizes results from the 16 events analyzed. LDW alert onset time t_0 is always available but the times of response t_1 and t_2 are not always so sharply defined. From published literature (Macadam, 2003) under ideal conditions and for a fully anticipated event, reactions times can be as quick as around 140ms for tactile and auditory triggers, and around 180ms for visual stimuli. Thus in reviewing the VIRTTEX data a reaction time less than 100ms was taken to indicate an event was not a direct reaction but rather a coincidence– for example when the head simply turned as part of the multi-tasking activity. (This was not particularly unusual, as the driver is switching attention between number reading and lane following; indeed, in video review, some drivers employ a tactic of making multiple head movements during the number reading task). Similarly, for the objective steering data, a clear steering response may not exist, so again for some events no reliable result is available. Entries marked (*) under t_1 and t_2 in Table C.6 were deemed unreliable and were not used in reaction time estimation. This includes some cases where the steering correction happened at the same time the driver moved his head forward, suggesting that he or she may be using peripheral vision to monitor the yaw deviation, and in any event the response was not clean enough to have confidence in it. The table shows RT_S computed for 16 yaw events - 4 drivers and 4 yaw deviation events for each driver (driver # and yaw deviation # are specified by the first two-digit number and the second number in filename, respectively).

Table C.6. Estimated reaction times for VIRTTEX distracted drivers. Numbers marked (*) indicate values deemed unreliable due to factors described above.

File name	t_0	t_1 (head turn)	t_2 (steer response)	RT_V	RT_S	$RT_S - RT_V$
sub02_1.mat	492.44	492.65	*493.04	0.21	-	-
sub02_2.mat	691.00	691.90	692.36	0.9	1.36	0.46
sub02_3.mat	1164.42	1165.30	1165.44	0.88	1.02	0.14
sub02_4.mat	1391.22	*1391.65	*1391.92	-	-	-
sub14_1.mat	500.34	* video timer missing	501.10	-	0.76	-
sub14_2.mat	710.70	711.50	711.82	0.8	1.12	0.32
sub14_3.mat	1179.76	*1179.90	*1179.76	-	-	-
sub14_4.mat	1414.28	1415.00	1415.20	0.72	0.92	0.20
sub31_1.mat	513.60	514.40	514.66	0.80	1.06	0.26
sub31_2.mat	694.30	*695.00	*694.30	-	-	-
sub31_3.mat	1140.62	-	*1140.62	-	-	-
sub31_4.mat	1348.74	*1349.90	1349.44	-	0.70	-
sub46_1.mat	493.16	-	*495.40	-	-	-
sub46_2.mat	706.06	*706.90	*706.06	-	-	-
sub46_3.mat	1129.92	*1130.60	*1129.92	-	-	-
sub46_4.mat	1339.12	*1339.60	*1339.96	-	-	-
Mean				0.72	0.99	0.28

After screening for unreliable events, the VIRTTEX data set used is very small, so we should certainly treat the mean values in the table as indicative rather than definitive – further testing and analysis was however beyond the scope of the present project. However some analysis of additional events was possible using the same experimental data; in some cases yaw deviations occurred when number reading was not associated with a yaw deviation. The results were consistent with those shown in Table C.6, providing general support for the numbers presented and also indicating that the artificial yaw deviation did not systematically affect the driver responses.

Further confidence in the results follows from the fairly constant difference between head and steering times seen in the table. Also, the key variable to be used in estimating T_{LDW} is the mean value of RT_S, i.e. 0.99s. This is in line with many results reported in the literature (Alm and Nilsson 1993; Green, 2000) and indicates an appropriate target for matching the simulation model.

Now we consider corresponding model-based reaction times. First we define a basic *system reaction time*

$$RT_0 = T_{LDW} + 2T_s + 2T_c \tag{C.5}$$

based on the minimum driver delay given in equation (6.8). Here we use the subscript 0 to denote a reaction time based on the simulation model. We note that reaction time for information processing may be longer than RT_0 depending on which parts of the driver model are active and also on when the LDW alert is issued relative to the *HK* processing cycle.

From simulation we have access to the steering-based reaction time (RT_{S0}) using the same peak steering rate condition as above. We expect $RT_{S0} > RT_0$ not only because of information processing, but also because of the extra time taken to generate the peak steering velocity. The differences are not constant, but based on a large number of test simulations with T_s sampled uniformly in the range (0.05, 0.1) and the assumed value T_c =0.02 (Table C.2) we obtained the mean result

$$RT_{S0} - RT_0 \approx 0.2 \tag{C.6}$$

This now gives a basis to estimate a mean value for T_{LDW} . From the above two equations

$$RT_{S0} \approx RT_0 + 0.2 = T_{LDW} + 2T_s + 2T_c + 0.2 \tag{C.7}$$

Using the mean value $T_s = 0.075$, and identifying the steering-based reaction time with the mean value in VIRTTEX ($RT_{S0} = RT_{S0} = 0.99$) we obtain

$$T_{LDW} \approx 0.6 \tag{C.8}$$

as the mean value for the LDW response delay in the driver model.

We now need to define a range of variation for T_{LDW} to be used in the batch simulations. To achieve this we compare the dispersion in the modeled RT_{S0} reaction times with the dispersion in RT_S obtained from VIRTTEX. Variations in RT_{S0} are obtained by random sampling T_{LDW} in a uniform range centered at 0.6. The simulations were to be based on Driving Scenario 1 (see Table 4.6) which has similar road and driving conditions to the VIRTTEX study. After a small number of iterations, the range (0.4, 0.8) was found to give plausible results, based on the informal criteria that the simulated dispersion should be wider than the valid VIRTTEX events in Table C.6, but not so large that unrealistically short reaction times emerge. Figure C.9 shows the basis of this informal validation, with results of 40 simulation runs from - together with the corresponding VIRTTEX data.

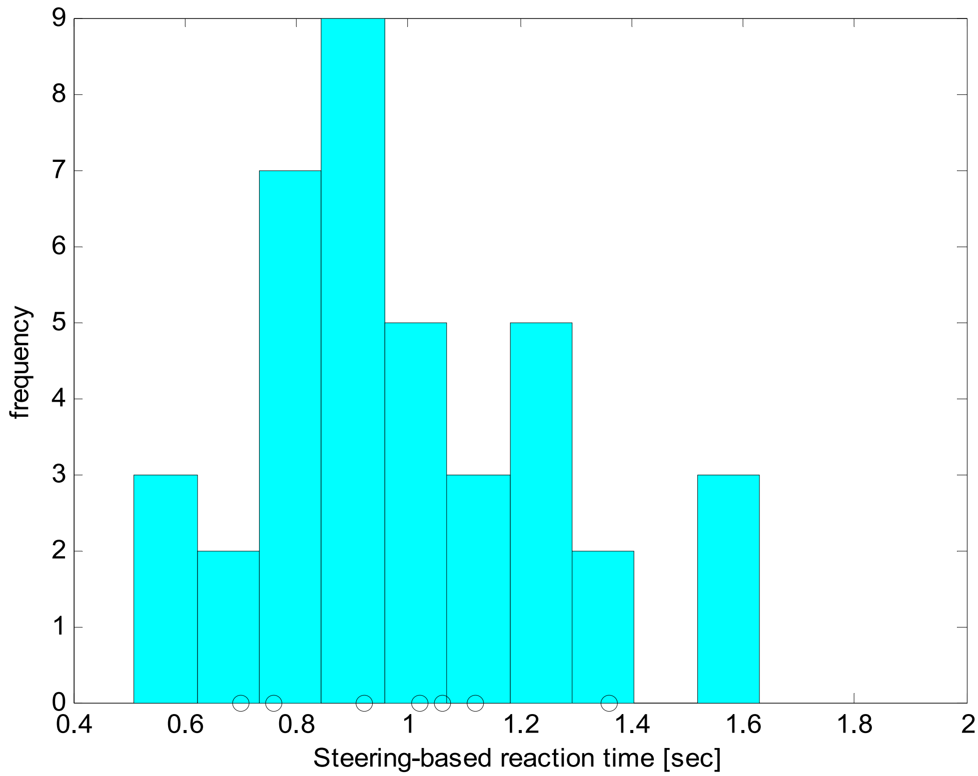

Figure C.9. Reaction Times RT_{S0} of 40 Simulations (histogram) and VIRTTEX Distracted-Driver Reaction Times RT_S (circles).

Driver Model Parameters for Fatigued Driving Scenarios

For the fatigued driver, again VIRTTEX test data with naïve subjects is used, and again the supporting event data was not very extensive. The approach taken was to hypothesize several plausible trends for the fatigued driver relative to the distracted drivers in the previous section. The VIRTTEX data would then be used test these trends as far as possible and then establish consequent changes in the driver model parameters. Three trends were initially proposed:

(a) Increase in the basic driver sampling rate T_s to provoke a slower response timing
(b) Reduction in the situational awareness of the fatigued driver after becoming alert
(c) Increase or otherwise modify to the LDW response time based on the reaction timing via peak steering velocity.

Using the same approach as above for the distracted driver, the experimental data of the sleep deprived driver VIRTTEX study (Study 1 of Section 5.1.1.) was used. Trip data was now based on only two subjects, one female driver and one male. Each trip contains 10 yaw deviation events with 5 different types of LDW warning (or "intervention type", namely, 0) no signal, 1) HUD + steering torque, 2) rumble strip sound + steering torque, 3) steering wheel vibration + steering torque, and 4) steering torque only. Each LDW warning type was used twice in a trip. The level of reduced vigilance was represented by the average percentage eyelid closure (PERCLOS, defined as the percentage of time that the pupils of the eyes are 80-100% occluded over a moving one minute time frame).

Not all events were appropriate for analysis; valid events were expected to meet the following criteria:

1) Average PERCLOS $\geq$ 8%.
2) RT is not too short (as before, to exclude coincident, non-triggered, responses).
3) An LDW alert of some kind was issued (i.e. exclude IT_0 events where the warning was suppressed).
4) There is no large steering action just before the LDW warning is issued.

Invalid events are marked (*) in the results, Table C.7, leaving only 8 events for analysis. The first thing we notice is that the mean reaction time RT_S is substantially reduced from the distracted case (here $mean(RT_S)$ =0.60s compared to 0.99s previously). It appears that the sleep deprived driver reacts more quickly to the LDW alert, and this is supported also by the mean of the peak steering rate: here the maximum value of $|\dot{\delta}|$ is 1.40 rad/s (for valid events) compared to 0.78 rad/s for the previous case of distracted drivers.

Table C.7. Reaction Times from Yaw Deviation Events (VIRTTEX Sleep Deprived Driver Study).

ID of Yaw Deviation Event	t0(LDW on)	t2(max_strwhl_rate)	max_strwhl_rate	RT(=t2-t0)	avgPERCLOS[%]
f_1_IT_3_yawdev_2.mat*	2227.70	2227.96	1.10	-	3.50
f_2_IT_4_yawdev_1.mat	2583.78	2584.14	0.94	0.36	8.44
f_3_IT_0_yawdev_2.mat*	3294.64	3294.64	0.38	-	13.36
f_4_IT_2_yawdev_2.mat*	3834.38	3834.96	0.64	-	10.81
f_5_IT_1_yawdev_2.mat*	4251.54	4252.22	1.01	-	6.89
f_6_IT_3_yawdev_2.mat	6042.88	6043.64	1.77	0.76	32.47
f_7_IT_1_yawdev_1.mat	6426.86	6427.36	0.62	0.50	9.94
f_8_IT_0_yawdev_2.mat*	6960.74	6961.50	0.48	-	9.36
f_9_IT_4_yawdev_1.mat	7319.34	7319.60	0.50	0.26	34.56
f_10_IT_2_yawdev_2.mat	7735.72	7736.40	3.09	0.68	33.50
m_1_IT_0_yawdev_1.mat*	2162.66	2162.66	0.34	-	5.00
m_2_IT_3_yawdev_2.mat*	2519.84	2520.00	0.85	-	5.53
m_3_IT_1_yawdev_1.mat	3249.18	3249.72	1.13	0.54	32.86
m_4_IT_4_yawdev_2.mat*	3779.38	3779.40	1.07	-	48.94
m_5_IT_2_yawdev_1.mat	4179.34	4180.26	0.93	0.92	41.67
m_6_IT_0_yawdev_2.mat*	5893.66	5894.40	0.92	-	45.67
m_7_IT_4_yawdev_1.mat	6442.88	6443.68	2.25	0.80	50.67
m_8_IT_3_yawdev_2.mat*	6994.42	6995.52	1.57	-	48.08
m_9_IT_3_yawdev_1.mat*	6994.42	6994.80	0.29	-	48.08
m_10_IT_1_yawdev_1.mat*	7323.04	7323.50	1.96	-	50.81
Mean			1.09 (all) 1.40 (valid)	0.60	27 (all) 31 (valid)

Note: ID includes female (f) or male (m) designation and intervention type (IT_0 etc.). Events marked (*) were invalid based on the above stated criteria.

Given what seems to be a counter-intuitive result, some further analysis was conducted on the effectiveness of the corrections made. It might be supposed that the sleep deprived driver responds more quickly but less effectively. To test this we compared the “settling time” of the recovery to stable lane keeping. Settling time is defined here as the elapse time between the start of the LDW alert and the time when the steer angle drops within a tolerance of $\pm 0.1\,\mathrm{rad}$ ($\pm 5.7\,\mathrm{deg}$) and the vehicle remains within the lane boundaries (both conditions to hold for a minimum of 2 seconds from the settling point). Results are shown in Tables C.8 and C.9 below. The settling times are smaller for the sleep deprived driver, even though the time out of lane is slightly longer. As with peak steering velocities, the peak steering angle is larger for the sleep deprived driver, again indicating a higher amplitude of intervention.

Table C.8. Settling Times (VIRTTEX Study 1, Sleep deprived driver, “valid” events)

Filename	lane excursion duration [s]	max strwhl angle [rad]	settling time [s]
f_2_IT_4_yawdev_1.mat	1.64	0.56	4.34
f_6_IT_3_yawdev_2.mat	5.70	0.65	5.70
f_7_IT_1_yawdev_1.mat	1.98	0.37	3.82
f_9_IT_4_yawdev_1.mat	3.82	0.27	5.06
f_10_IT_2_yawdev_2.mat	3.26	0.81	3.82
m_3_IT_1_yawdev_1.mat	1.96	0.49	2.68
m_5_IT_2_yawdev_1.mat	2.28	0.43	4.08
m_7_IT_4_yawdev_1.mat	2.92	0.51	3.54
Mean	3.13	0.50	4.10

Table C.9. Settling Times (VIRTTEX Study 2, Distracted driver, “valid” events)

Filename	lane excursion duration [s]	max strwhl angle [rad]	settling time [s]
sub02_2.mat	2.76	0.21	9.48
sub02_3.mat	2.44	0.39	6.74
sub14_1.mat	2.96	0.30	4.24
sub14_2.mat	2.32	0.39	5.00
sub14_4.mat	2.18	0.40	4.18
sub31_1.mat	2.70	0.41	6.28
sub31_4.mat	2.36	0.47	3.40
Mean	2.53	0.37	5.62

Returning to our initial anticipated trends listed at the beginning of this section, under (a) there is no evidence to reduce T_s, and while there may be some small indication to reduce it, with limited data to re-calibrate its range we retain the previous range (0.05-0.10) seconds.

Under (b) we have no objective data to support any loss of situation awareness, but intuitively a fatigued or sleep deprived driver who is drifting out of lane without any specific distraction activity is less likely to be fully tracking the lane boundaries. In this case, the driver model provides a simple means to represent this proposed degradation in tracking awareness. As the attention level switches to zero in the model, the working memory buffer is cleared, so that the blocks object recognition and classification (ORC) and threat assessment (TA) in Figure 6.19 must be utilized to allow effective tracking to resume.

Turning to (c), the LDW reaction time, it is clear that from the VIRTTEX data this appears to be reduced relative to the distracted case. From the discussion in Section 6.4 it seems reasonable to accept the hypothesis that T_{LDW} may be reduced for the sleep deprived driver.

The mean and standard deviation of RT_S of valid events in Table C.7 are 0.63 and 0.23 respectively. After a number of iterations with the simulation model (including the effect of clearing the working memory buffer during a loss of attention) it was found that selecting T_{LDW} in the range (0.3, 0.4) gives comparable results of 0.63 (mean) and 0.18 (standard deviation). These were based on 66 simulated lane excursions from Scenario 11 (see Table 4.6 – this is comparable to the condition in the VIRTTEX study) where the LDW signal was issued and the driver directly responded. The distributions are shown in Figure C.10. In the figure, the measured 0.26 sec result appears very quick, and may be an outlier due to coincidence rather than a stimulus-response reaction. Overall the distribution seems plausible, including the events that show response times over 1 second, comparable to the distracted driver responses.

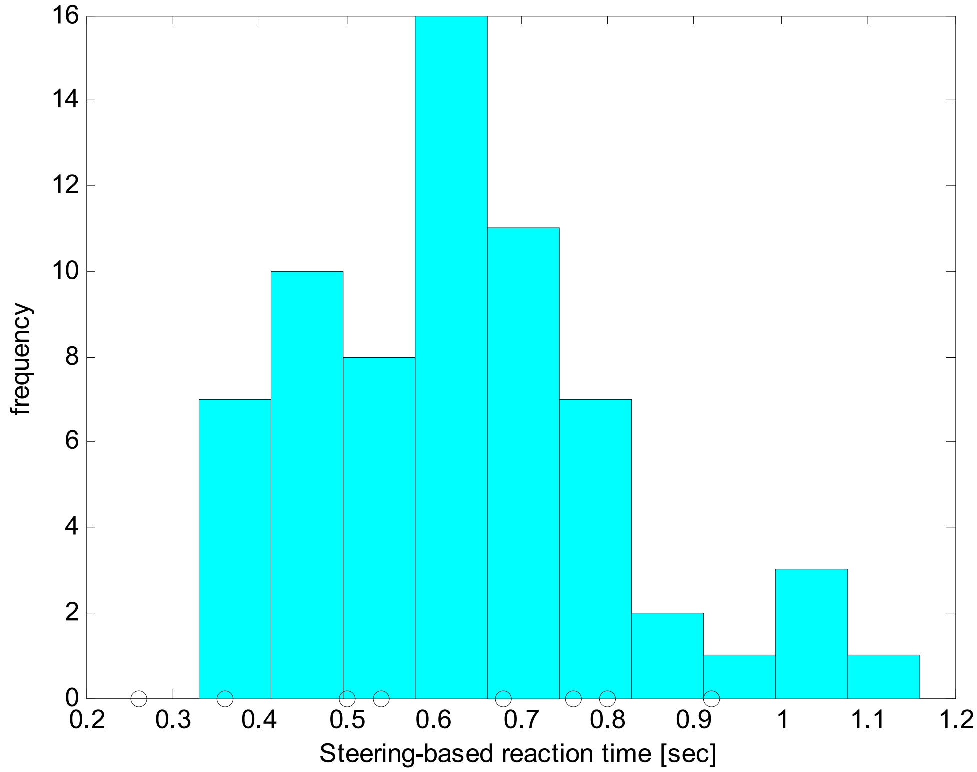

Figure C.10. Reaction Times RT_{S0} of 40 Simulations (histogram) and VIRTTEX Sleep Deprived Driver Reaction Times RT_S (circles).

Appendix D: Driver Alert Control and Emergency Lane Assist

In this section we present some partial analysis and discussion for the Driver Alert Control (DAC) and Emergency Lane Assist (ELA) systems, even though the formal SIM analysis is not applied to these systems. In the context of the SIM methodology presented for LDW it is highly instructive to consider the challenges presented by these very different ACAT technologies. Although we consider the same basic crash types and hence the same baseline driving population, the two systems considered here operate quite differently from LDW. DAC is an advisory system, and operates at a different point in the crash phase timing diagram, while ELA is a control system that does not depend on driver response timing, but has the potential for more complex interactions with the driver.

D.1. Driver Alert Control

For Driver Alert Control (DAC) we consider factors that influence benefits estimation in a somewhat remote fashion, i.e. not via microscopic dynamic simulations, but via influences on the weights in the benefits equation. As we shall describe, DAC is based on an estimated "vigilance state" of the driver and may be relevant to both distracted driver states and diminished driver vigilance states, for example due to fatigue. For simplicity, we refer to the detection and warning of a "drowsy" driver state and assume that the benefits are limited to a potential reduction in crash likelihood for scenarios involving a fatigued driver (as coded by GES). In reality (and if shown during a future in-depth study of DAC or similar systems) the benefits may be wider.

Thus DAC potentially reduces the scenario weights associated with fatigued driving, and hence reduces estimated crash numbers via the exposure ratio. In this section we develop a candidate method for estimation of potential DAC benefits using naturalistic driving data and a functional model of DAC. The timing of DAC is broadly illustrated in Figure 1.1 (Section 1); it operates remotely from any particular conflict, so we do not simulate its performance in the same way as LDW. Rather, we simulate its performance based on the broad range of naturalistic driving data in the RDCW database and estimate potential benefits by the degree to which it could reduce the frequency with which certain driving scenarios occur. In this case, the effect of DAC is modeled as an exposure ratio for conflicts where the driver is coded as fatigued.

The extent of exposure reduction for fatigued driver scenarios depends on the extent to which drivers:

a) receive timely warning from DAC,
b) respond appropriately (take a rest break or adopt another appropriate refresh strategy), and
c) maintain vigilance during subsequent driving.

There is a potential fourth factor at play – the frequency of false alarms experienced during normal driving. The first three factors are to be explicitly represented in the benefits estimation for DAC, while the fourth factor is less direct, but is expected to have a strong influence on both (a) and (b); as noted in Section 9 in the context of LDW, nuisance alerts will increase the likelihood of the driver switching the system off (reducing frequencies in (a)) or simply ignoring the system's advice (reducing frequencies in (b)). False positives, if perceived as a nuisance by the majority of drivers, clearly have the potential to render the system less effective.

Estimating DAC benefits depends on the above three factors, and it is not a simple matter to obtain reliable estimates; all thee depend on complex real-world factors, and a full analysis would require field testing that goes beyond the scope of this project. However some progress can be made and a methodology is proposed, even if the supporting data are not fully available. The approach is based on using naturalistic driving data as a surrogate for the driving population to estimate the effect of (a). The effects of (b) and (c) are combined into a simple "post alert model": given an alert we assume a probability P_{com} that the driver "complies" with the alert, i.e. responds in some appropriate manner; a second factor *Q* determines the time duration for which increased vigilance is maintained, assuming there was initial compliance. The meaning of *Q*, as well as vigilance level *V*, is derived from the DAC algorithm itself.

The DAC algorithm provides a measure of lane keeping performance as derived from the on-board lane tracker, using vehicle path estimation and lane-keeping error prediction (Birk, 2006). When the performance of lane-keeping control is reduced (variations in lane position and path direction increase), an internally stored vigilance score *V* is similarly reduced. If lane-keeping performance improves, the *V* score is increased. *V* takes values on an integer scale of 1 to 5, and is initially set to 5 (interpreted as high driver alertness); a DAC alert is issued if *V* falls to 1 (interpreted as a fatigued driver). The DAC algorithm was applied to the naturalistic driving data from the RDCW database, via a Simulink model, and Figure D.1 shows a typical time history of *V*. Until about 2400 seconds into the trip, the DA algorithm evaluated the driver's vigilance to be between 5 (most alert) and 3 (somewhat vigilant). The resulting warning signal was generated at about 2400 seconds as shown in Figure D.2 – i.e. a single right-side alert was issued at the time *V* dropped to 1.

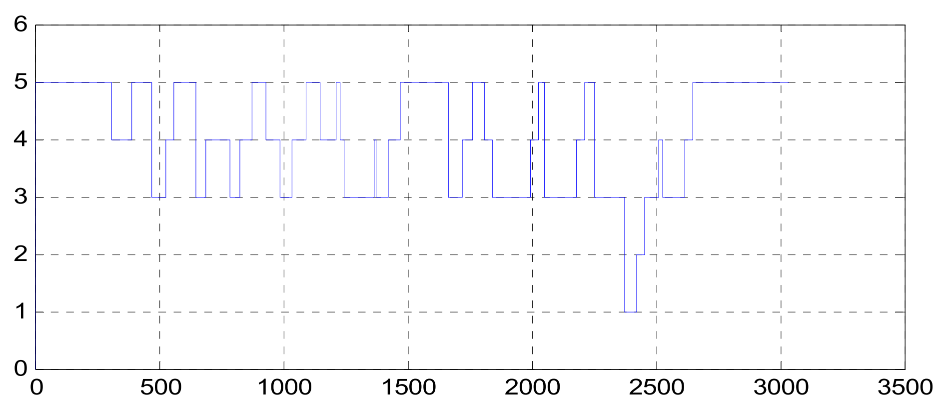

Figure D.1. DAC Vigilance-level (*V* vs. trip time (sec))

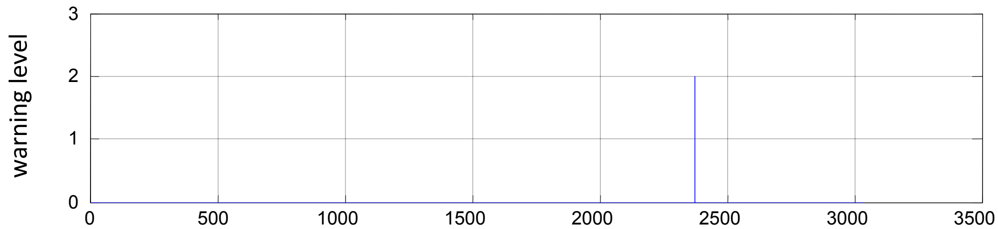

Figure D.2. Time History of DAC Warning Signal (0 = off, 1 = left side warning, 2 = right side warning)

Up to the point where the warning is shown, it is safe to assume that if DAC had been installed in that particular vehicle during that particular trip – its performance and output would be appropriately represented by the above simulation output. However, if we assume driver compliance, the vigilance level would be expected to improve after the alert, so some form of *post-alert model* is needed for how this might improve. Note that any apparent improvement seen in Figure D.1 after the alert time was due to other factors, since DAC was not fitted to the RDCW vehicle that provided the analysis data, and hence no such alert was given in the real driving situation.

Figures D.1 and D.2 show vigilance estimation in the DAC system, as well as the timing of the single warning issued. To estimate the effect of safety benefits we need to estimate the proportion of "drowsy driving time" that is reduced because of the DAC warnings, and hence – as well as a post-alert model – we need an estimate of when the driver is actually fatigued. Recall that in this very preliminary analysis, vigilance score is an assumed metric for a particular driver state that we refer to as "drowsy"; at present it is not directly correlated with other measures of drowsiness, or with reports of driver fatigue found in real world crash data.

Figure D.3 illustrates this approach: the naturalistic data (no warnings issued) provides an estimate of when the driver is actually drowsy (the green bar labeled "No-DA Drowsy") while the post-alert model defines an interval of time ("Mask") during which the compliant driver recovers vigilance. In the following we define a simple but illustrative post-alert model that defines the masking interval in terms of the alertness at the start of the trip. As part of the simplified approach, it is further assumed that there is no delay between the first DAC warning and the driver achieving a heightened state of alertness (i.e. we assume no significant delay in the recovery and hence the mask is initiated at the DAC Alert time)

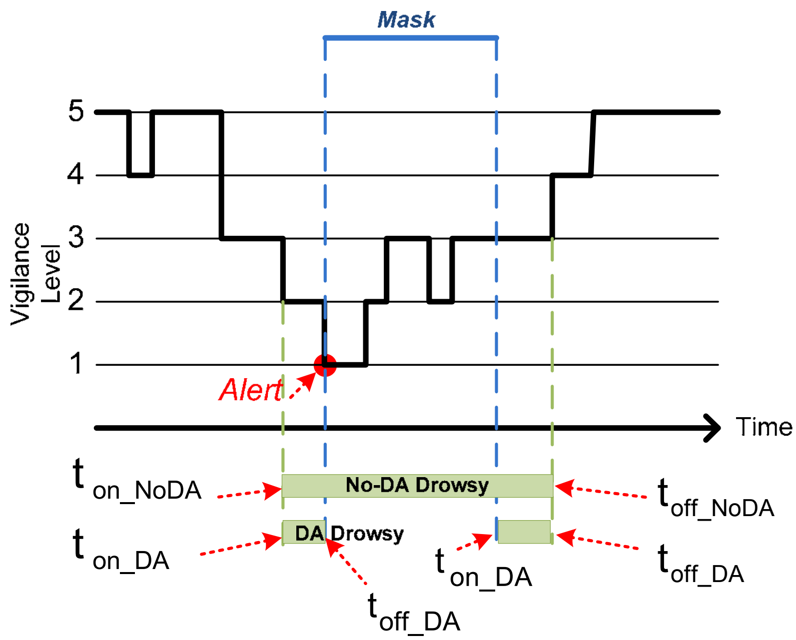

Figure D.3. DAC Exposure Effect: Estimated Reduction in "Drowsy Driving Time"

In Figure D.3 we use the DAC vigilance estimator to define the onset of a drowsy driving interval (t_{on}) when *V* falls to 2, and the end of that interval when *V* increases again to *V*=4. Thus vigilance levels of 4 or 5 are taken to indicate the non-fatigued state, levels 1 and 2 indicate a drowsy state, and *V*=3 is considered a transition state that is not directly used. Ideally some independent indicator of the drowsy state would be used here, typically one that involves eye closure percentage (*perclos*) as in Section 6. Without this, and for the purpose of discussion and establishing some benchmark performance metrics, we use the vigilance measure as described.

According to our analysis approach the exposure ratio ε_i for fatigued scenarios (Scenarios 11, 13, 15, and 20 in Table 4.6) is estimated via the following equations. The distance traveled "while drowsy" in the non-DA case is

$$D_{NoDA} = \sum_{NoDA} \bar{v}\left(t_{off} - t_{on}\right) \tag{D.1}$$

where $\bar{v}$ is the mean vehicle speed during identified drowsy driving episode in the naturalistic data. The corresponding distance in the "masked" data, for when DAC is active and the driver is fully compliant, is given by

$$D_{comDA} = \sum_{DA} \bar{v}\left(t_{off} - t_{on}\right) \tag{D.2}$$

More generally, assuming a compliance probability $0 \le P_{com} \le 1$, the mean distance traveled while fatigued is intermediate between the above values:

$$D_{DA} = P_{com} D_{comDA} + (1 - P_{com}) D_{NoDA} \tag{D.3}$$

The exposure ratio is then estimated as the ratio $D_{DA} \div D_{NoDA}$, i.e.

$$\varepsilon_i = (1 - P_{com}) + P_{com} \frac{\sum_{DA} \bar{v}\left(t_{off} - t_{on}\right)}{\sum_{NoDA} \bar{v}\left(t_{off} - t_{on}\right)} \tag{D.4}$$

The RDCW naturalistic driving data used in this analysis could have been broken down according road type and hence provide different exposure ratios indicated for each scenario. However, given the relatively small amount of data available, plus the lack of maturity and refinement of this method, maintaining simplicity via aggregation of factors is preferred; hence a single estimated exposure ratio is derived across all relevant driving scenarios in the naturalistic data.

Having defined the strategy for estimating benefits, the details of the masking algorithm in the post-alert model must be determined. The scientific basis for the masking algorithm is slender at present: we need to know how long a driver maintains vigilance after a rest break or other refresh strategy (e.g., engage in conversation to increase attentional arousal). Given the complexity of real-world conditions that influence the driver state (availability of rest areas, perceived urgency of the trip, appreciation of the risks on drowsy driving, presence of passengers, physiological response to the rest activity, etc.) this is not directly known. One possibility is to assume a fixed time over which the individual is likely to be recovered, e.g. to assume that a 30 minute respite will typically be achieved; this would be constant across all individuals. Another possibility (also somewhat arbitrary, but variable among individuals and individual trips) is to start with a time parameter, T_{V5}, which is the time the driver's vigilance state remained initially at *V*=5, and use this as a benchmark for the individual trip. The duration of the masking period is then set to be proportional to the benchmark:

$$t_{mask} = QT_{V5}. \tag{D.5}$$

Here *Q* is a dimensionless scale factor relating the ability of the driver to regain high vigilance after a DAC warning, based on vigilance at the start of the trip. In reality *Q* will be a random variable, but its distribution and dependencies are currently unknown, so for the analysis here we simplify by assuming a constant value. This approach is based on the idea that a driver who is more fatigued at the start of the trip (and would be expected to have a reduced T_{V5} time) is provided a smaller masking time than a driver who is initially more vigilant. Even when we assume a constant value for *Q* in equation (D.5), the parameter is completely unknown at this time, and for simplicity we shall assume *Q*=1 as a reference value (other values were also considered – see below).

The above post-alert model is based entirely on the driver response, albeit in a very simplified manner. The DAC system itself also has some post alert behavior that is worth mentioning. If the driver stops and exits the vehicle, the DAC algorithm will reset itself to *V*=5 (this happens if the ignition is switched off and the door is opened and closed). On the other hand, if there is no such interruption, the system will automatically increment its value to *V*=3 within 2 minutes of the alert, and, during this short interval, the system is not responsive to lane-keeping performance. This system effect is short-lived and unlikely to strongly influence subsequent warnings or long-term behavior of V. Because implementation of these conditions in the post-alert phase is unlikely to influence the value obtained for ε_i, we simply ignore this in our analysis.

The data analysis is based on a subset of trips in the naturalistic driving data chosen with

- trip length of at least 30 miles, and
- lane tracking system available on both sides of the lane for at least 70% of the distance traveled.

The second condition means that the lane markings, which are crucial for DAC to function, were sufficiently visible for at least that percentage distance. These conditions resulted in 440 trips being selected for analysis, yielding the results below.

No specific predictions are made for the safety benefits deriving from DAC. However, to demonstrate the potential scope of the method, we note that under the assumption Q=1, and based on the masking data analysis, the exposure ratio for intervals of fatigued driving is estimated as a function of compliance probability according to the following equation:

$$\varepsilon_i = (1 - P_{com}) + 0.327 P_{com} \tag{D.6}$$

For a fully compliant driver with Q=1, the analysis suggests an exposure ratio of 33%, and the crash numbers associated with Scenarios 11, 13, 15, and 20 would be reduced by 67%. If instead of Q=1 we use smaller or larger values (Q=0.5 or 1.5), repeating the analysis leads to the following modified equations for the exposure ratio

$$\varepsilon_i = \begin{cases} (1 - P_{com}) + 0.310 P_{com} & (\text{if } Q = 0.5) \\ (1 - P_{com}) + 0.395 P_{com} & (\text{if } Q = 1.5) \end{cases} \tag{D.7}$$

These results are intended to indicate the order of magnitude of the effect on exposure ratio, since at present there are insufficient data to validate or develop further the simple masking interval method; for now the exposure ratio estimate equations are only provisional, and the values of Q and P_{com} mentioned are purely illustrative.

In the above analysis we have implied the simplifying assumption that only crashes coded as fatigued would benefit from reduced exposure. It is also quite possible that a larger number of baseline crash numbers would be affected: because of underreporting of driver fatigue, because scenarios with a distracted driver would benefit, or even because of long-term improvements of awareness of reduced vigilance and its risks. At present no objective data is available to quantify all of these effects.

D.2. Emergency Lane Assist

A description of the ELA system and its design intent were given in Sections 1.3 and 5.3. The ELA is a steering-based intervention system that essentially applies a steering torque input to the vehicle to obtain a desired heading and trajectory. Objective testing involved experimental testing at Volvo test facilities with an XC90, as well as some limited driving with Ford's VIRTTEX driving simulator; in the latter tests, team members acted as subjects for a set of pre-defined driving scenarios. Both sets of tests were intended to map the basic system performance; no naïve test subjects were used in the short ELA study, so there was no attempt made to characterize relevant driver-system interaction in a safety critical situation.

Based on the data from the above testing, a generic model of the ELA functionality has been developed. The model is able to provide a reasonably accurate representation of the ELA functionality corresponding to this limited test data and could be used for event simulation similar to that used for LDW. The focus of the remainder of the section is the ELA model and can be viewed as an extension of the model development work reported in Section 6.

Functionally, the ELA system is activated when the LDW system is triggered in the vehicle, which would occur when the subject vehicle is leaving its current lane. The system then a) identifies the likely threat based on the likelihood of collision if the subject vehicle maintained its current path – which in its simplest form could be an oncoming vehicle in the adjacent lane, and b) evaluates a desired correction path to avoid a collision.

The desired path is based on design metrics that combine the subject vehicle's current and projected kinematic states to ensure a safe return to its original lane without causing yaw instability. The path is calculated so that the vehicle is safely in its original lane for a certain amount of time (safety window) before the threat vehicle passes the spatial point where a collision would have occurred if the subject vehicle did not have a course correction. Once the ELA system is activated, a desired steering wheel angle is calculated based on the desired path. The system uses a feedback controller to track the desired steering signal in the curved section of the trajectory (Figure D.4). Once the desired path changes from curved to straight, the system uses a feedback controller with position error feedback to complete the desired maneuver. Figure D.4 shows a schematic of the ELA operation as implemented in the SIM model.

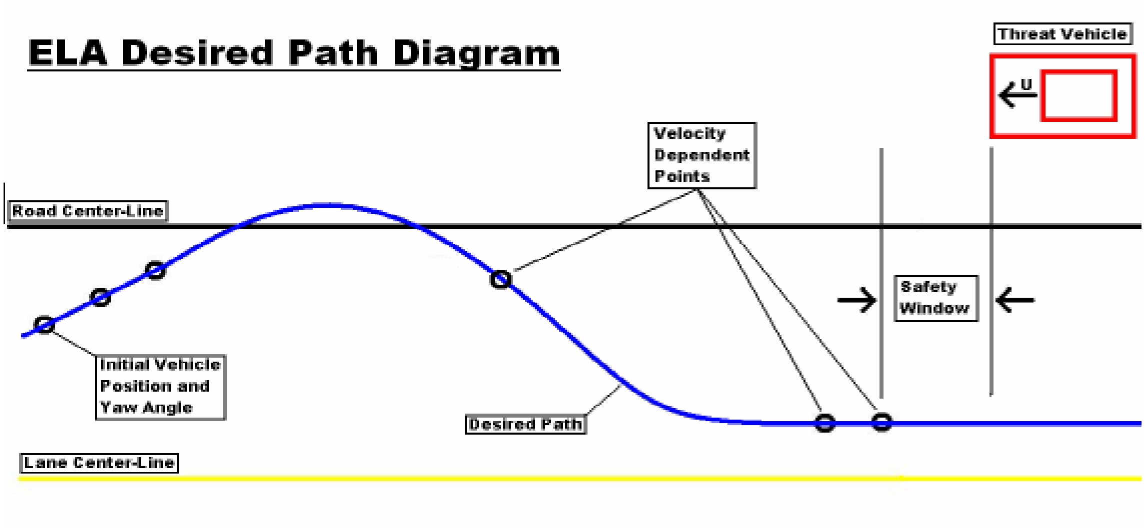

Figure D.4. Schematic of ELA Operation

In terms of the model development within Simulink, the ELA function was added to the existing LDW model described in Section 6. The ELA system is simulated in Simulink, and to allow both system and driver torques to be simultaneously input, a simple steering system model was also created in Simulink. This portion of the model required detailed tire force outputs from the underlying CarSim model so that the steering system model experiences representative dynamics in the coupled simulation. Figure D.5 shows the system model with the ELA function included instead of LDW (compare Figure 6.1). The system adds a steering torque correction designed to steer the vehicle back into the initial lane. The model was implemented in Simulink and coupled to the CarSim model. To be consistent, the driver model was adapted slightly so the driver input was also in the form of torque input; in this way it is possible for both driver and ELA to apply input simultaneously – a feasible, though generally undesirable operating condition.

Figures D.6-D.8 show sample simulation results for a straight road segment that corresponds to Driving Scenario 1. In these simulations, it is assumed that the driver does not interfere with the ELA operation and is largely passive for the duration of the ELA activation. In Figure D.6, the subject vehicle is drifting to the right, into the path of an oncoming vehicle. The subject vehicle's longitudinal velocity at the start of the maneuver is 64.4 mph. For the simulation, the driver model is turned off to avoid confounding the effect of driver input with the ELA. Therefore the correction to the subject vehicle's trajectory is due to the ELA steering torque. Figure D.7 shows the desired and output steering wheel angles. The desired angle is computed based on the system functionality discussed earlier. Figure D.8 shows the desired and resulting vehicle path in relation to the lane boundaries. The model is broadly representative of the dynamic operation of the Volvo ELA system and provides a feasible starting point for parameter tuning, validation and simulation for safety benefits estimation.

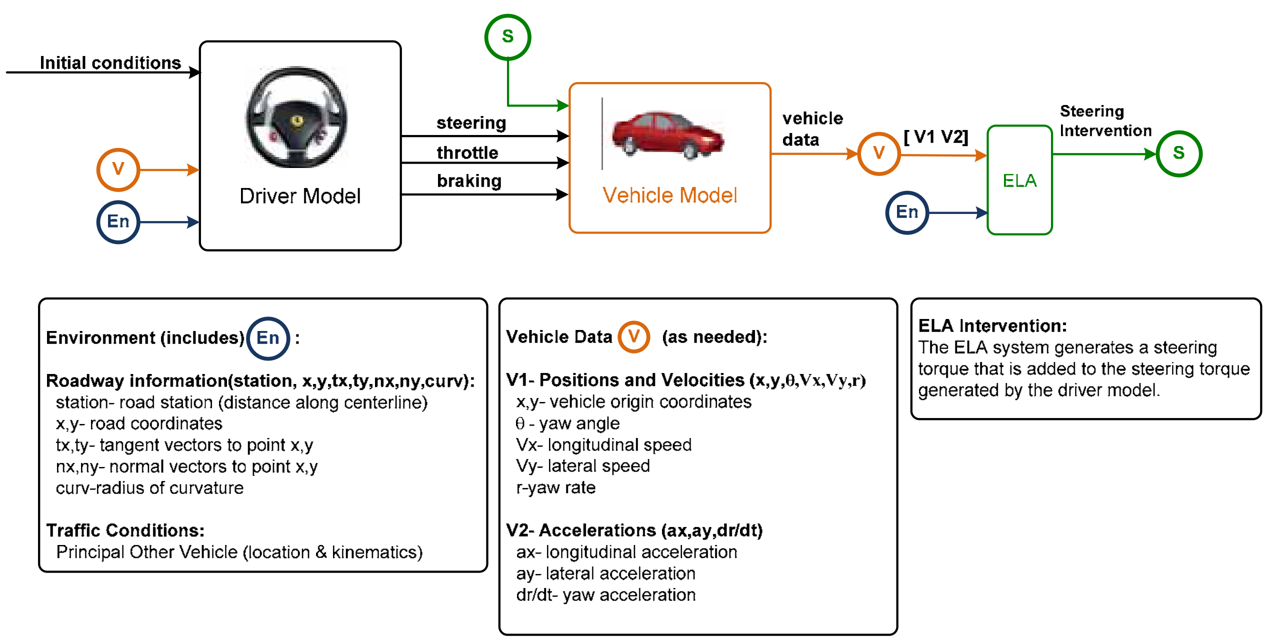

Figure D.5. ELA Functionality Added to the DVET Model

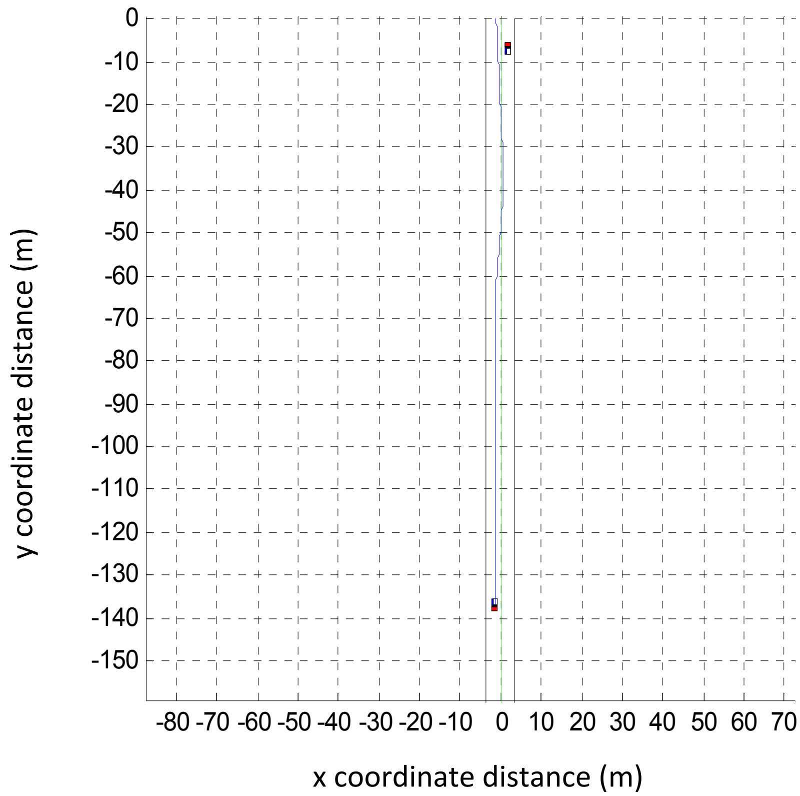

Figure D.6. Simulation of ELA Intervention to Avoid Oncoming Vehicle (Driving Scenario 1)

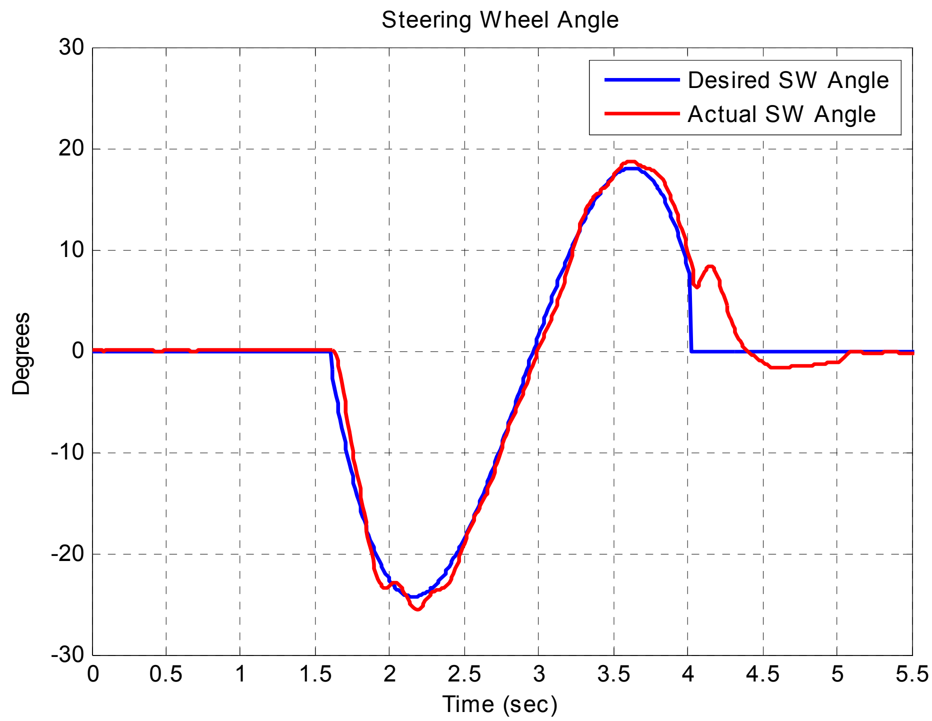

Figure D.7. Actual and Desired Steering Wheel Angle for ELA Activation

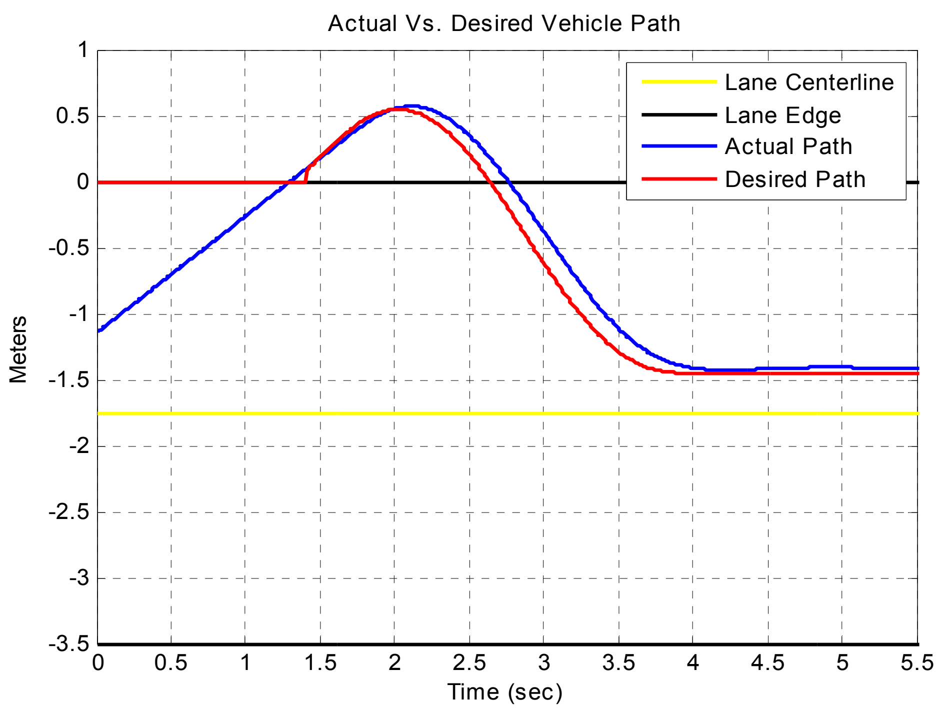

Figure D.8. Actual and Desired Vehicle Path for ELA Activation

DOT HS 811 405
October 2010

U.S. Department
of Transportation

National Highway
Traffic Safety
Administration

NHTSA
www.nhtsa.gov