Design and Development of an Optimal-Control-Based Framework for Trajectory Planning, Threat Assessment, and Semi-Autonomous Control of Passenger Vehicles in Hazard Avoidance Scenarios

Sterling J Anderson[†], Steven C. Peters[†], Tom E. Pilutti[‡], and Karl Iagnemma[†]*

†Department of Mechanical Engineering, Massachusetts Institute of Technology, 77 Massachusetts Ave. Cambridge, MA 02139, USA, E-mail: {ster | scpeters | kdi} @mit.edu

‡Ford Research Laboratories, Dearborn, MI 48124, USA, Email: tpilutti@ford.com

*Corresponding author

Abstract This paper describes the design of an optimal-control-based active safety framework that performs trajectory planning, threat assessment, and semi-autonomous control of passenger vehicles in hazard avoidance scenarios. This framework allows for multiple actuation modes, diverse trajectory-planning objectives, and varying levels of autonomy. A model predictive controller iteratively plans a best-case vehicle trajectory through a navigable corridor as a constrained optimal control problem. The framework then uses this trajectory to assess the threat posed to the vehicle and intervenes in proportion to this threat. This approach minimizes controller intervention while ensuring that the vehicle does not depart from a navigable corridor of travel. Simulation and experimental results are presented here to demonstrate the framework's ability to incorporate configurable intervention laws while sharing control with a human driver.

Introduction

Recent traffic safety reports from the National Highway Traffic and Safety Administration show that in 2007 alone, over 41,000 people were killed and another 2.5 million injured in motor vehicle accidents in the United States [1]. The long-standing presence of passive safety systems in motor vehicles, combined with the ever-increasing influence of active systems, has contributed to a decline in these

numbers from previous years. Still, the need for improved collision avoidance technologies remains significant.

Recent developments in onboard sensing, lane detection, obstacle recognition, and drive-by-wire capabilities have facilitated active safety systems that share steering and/or braking control with the driver [2,3]. These active safety systems operating with a "human in the loop" generally attempt to honor driver intentions, opposing them only when doing otherwise would lead to a collision or loss of control. Such modification of the driver's intended trajectory requires that these systems assess the threat posed to a vehicle in order determine when and how strongly to intervene. Such systems should honor safe driver inputs and maneuvers while intervening when necessary to correct or override those deemed unsafe.

Among existing proposals for semi-autonomous vehicle navigation, lane-keeping systems using audible warnings [4], haptic alerts [5], steering torque overlays [6], and various combinations of these have been developed with mixed results [7]. In a recent subproject of the European PReVENT consortium, a lane-keeping system was designed to prevent lane departure by perceiving the environment, making heuristic-based trajectory planning decisions based on perceived threat, and implementing warning mechanisms or a slight steering torque overlay when the vehicle drifts from the desired trajectory [8].

Many of the navigation systems developed in previous work address only one piece of the active safety problem. While some use planning algorithms such as rapidly-exploring random trees [3], evolutionary programming [9] or potential fields analysis [10] to plan a safe vehicle path, others simply begin with this path presumed [11]. The threat posed by a particular path is seldom assessed by the controller itself and is often only estimated by a simple threat metric such as lateral or longitudinal acceleration required to avoid a road hazard [12]. Finally, hazard avoidance is commonly performed using one or more actuation methods (steering, differential braking, etc.) without explicitly accounting for the effect of driver inputs on the vehicle trajectory [8]. Such controllers selectively replace (rather than assist) the driver in performing the driving task.

Yu addressed this problem in mobility aids for the elderly by designing an adaptive shared controller which allocates control authority between the human user and a controller in proportion to the user's performance [13]. These metrics and the associated intervention are designed to act on current and past user performance, however, and do not anticipate future states or performance. This reactive approach to semi-autonomy, while sufficient to control low-speed mobility aids, is not well suited for higher-speed applications with significant inertia effects and no pre-planned trajectory.

In this paper, a framework for passenger vehicle active safety is developed that performs vehicle trajectory planning, threat assessment, and hazard avoidance in a unified manner. This framework leverages the predictive and constraint-handling capabilities of Model Predictive Control (MPC) to plan trajectories through a pre-selected corridor, assess the threat this path poses to the vehicle, and regulate driver and controller inputs to maintain that threat below a given threshold.

The next section describes the semi-autonomous control framework and its associated trajectory prediction, control law, threat assessment, and intervention law. Simulation setup and results are then presented, followed by experimental setup and results, and the paper closes with general conclusions.

Framework Description

The framework described below leverages the predictive- and constraint-handling capabilities of MPC to perform trajectory planning, threat assessment, and hazard avoidance. First, an objective function is established to capture desirable performance characteristics of a safe or "optimal" vehicle path. Boundaries tracing the edges of the drivable road surface are assumed to have been derived from forward-looking sensor data and a higher-level corridor planner. These boundaries establish constraints on the vehicle's projected position. This constraint data, together with a model of the vehicle dynamics is then used to calculate an optimal sequence of inputs and the associated vehicle trajectory. The predicted trajectory is assumed to be a "best-case" scenario and used to establish the minimum threat posed to the vehicle given its current state and a series of best-case inputs. This threat is then used to calculate the necessary intervention required to prevent departure from the navigable region of travel and driver/controller inputs are scaled accordingly. Fig. 1 shows a block diagram of this system.

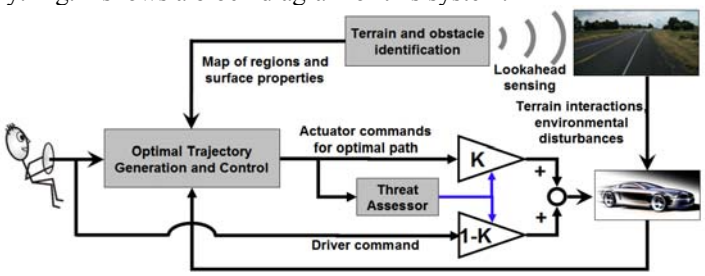

Fig. 1 Diagram of an active safety system

In this paper it is assumed that road lane data is available and that road hazards have been detected, located, and mapped into a 2-dimensional corridor of travel. Existing systems and previous work in onboard sensing and sensor fusion justify this as a reasonable assumption [14]. Radar, LIDAR, and vision-based lane-recognition systems [3,15], along with various sensor fusion approaches [16] have been proposed to provide the lane, position, and environmental information needed by this framework.

Additionally, where multiple corridor options exist (such as cases where the roadway branches or the vehicle must circumnavigate an obstacle in the center of the lane), it is assumed that a high-level path planner has selected a single corridor through which the vehicle should travel.

Vehicle Path Planning

The best-case (or baseline) path through a given region of the state space is established by a model predictive controller. As described in later sections, metrics from this predicted path will be used to assess threat.

Model Predictive Control is a finite-horizon optimal control scheme that iteratively minimizes a performance objective defined for a forward-simulated plant model subject to performance and input constraints. Stated another way, MPC uses a model of the plant to predict future vehicle state evolution and optimize a set of inputs such that this prediction satisfies constraints and minimizes a user-defined objective function. At each time step, t, the current plant state is sampled and a cost-minimizing control sequence spanning from time t to the end of a control horizon of n sampling intervals, $t+n\Delta t$, is computed subject to inequality constraints. The first element in this input sequence is implemented at the current time and the process is repeated at subsequent time steps. The basic MPC problem setup is described in [17].

The vehicle model used in MPC accounts for the kinematics of a 4-wheeled vehicle, along with its lateral and yaw dynamics. Vehicle states include the position of its center of gravity [x, y], its yaw angle ψ, yaw rate $\dot{\psi}$, and sideslip angle β, as illustrated in Fig. 2. Table 1 defines and quantifies this model's parameters.

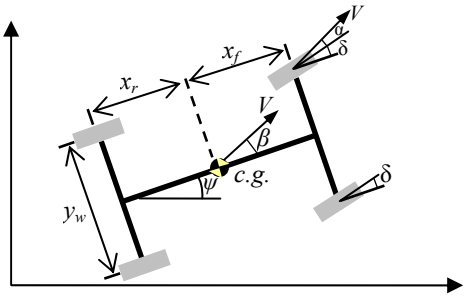

Fig. 2 Vehicle model used in MPC controller

Table 1 Vehicle model parameters

Symbol	Description	Value [units]
m	Total vehicle mass	2220 [kg]
I_{zz}	Yaw moment of inertia	3344 [kg·m^2]
x_f	C.g. distance to front wheels	1.43 [m]
x_r	C.g. distance to rear wheels	1.47 [m]
y_w	Track width	1.44 [m]
C_f	Front cornering stiffness	1433 [N/deg]
C_r	Rear cornering stiffness	1433 [N/deg]
μ	Surface friction coefficient	1

Tire compliance is included in the model by approximating lateral tire force (F_y) as the product of wheel cornering stiffness (C) and wheel sideslip (α or β for front or rear wheels respectively) as in

$$F_y = C\alpha\,. \tag{1}$$

Linearized about a constant speed and assuming small slip angles, the equations of motion for this model are (where δ represents the steering angle input)

$$\dot{x} = V \tag{2}$$

$$\dot{y} = V(\psi + \beta) \tag{3}$$

$$\dot{\beta} = \frac{-(C_r + C_f)}{mV}\beta + \left(\frac{(C_r x_r - C_f x_f)}{mV^2} - 1\right)\dot{\psi} + \frac{C_f}{mV}\delta \tag{4}$$

$$\ddot{\psi} = \frac{(C_r x_r - C_f x_f)}{I_{zz}}\beta - \frac{(C_r x_r^{\,2} + C_f x_f^{\,2})}{I_{zz}V}\dot{\psi} + \frac{C_f x_f}{I_{zz}}\delta \tag{5}$$

where C_f and C_r represent the cornering stiffness of the lumped front wheels and the lumped rear wheels, and x_f and x_r are the longitudinal distances from the c.g. of the front and rear wheels, respectively.

Constraint Setup and Objective Function Description

As mentioned above, this framework assumes that the environment has been delineated previously. The boundaries of the navigable road surface at each timestep are then described by the constraint vectors

$$\begin{aligned} \mathbf{y}^{y}{}_{\max}(k) &= \left[y^{y}{}_{\max}(k+1) \quad \cdots \quad y^{y}{}_{\max}(k+p)\right]^T \\ \mathbf{y}^{y}{}_{\min}(k) &= \left[y^{y}{}_{\min}(k+1) \quad \cdots \quad y^{y}{}_{\min}(k+p)\right]^T \end{aligned} \tag{6}$$

In (6), $\mathbf{y}^{y}{}_{\max}$ and $\mathbf{y}^{y}{}_{\min}$ represent the upper and lower limits on the vehicle lateral position (y) and must satisfy

$$\mathbf{y}^{y}{}_{\max} - \mathbf{y}^{y}{}_{\min} > 0 \tag{7}$$

in order for the constraint space to remain feasible.

By enforcing vehicle position constraints at the boundaries of the navigable region of the road surface (i.e. the lane edges on an unobstructed road), the controller forces the MPC-generated path to remain within the constraint-bounded corridor whenever dynamically feasible. Coupling this lateral position constraint with input constraints $\mathbf{u}_{\min/\max}$, input rate constraints $\Delta\mathbf{u}_{\min/\max}$, and vehicle dynamic considerations, the corridor delineated by $\mathbf{y}^{y}{}_{\max}$ and $\mathbf{y}^{y}{}_{\min}$ translates to a safe operating region within the state space.

The controller's projected path through this constraint-imposed region is shaped by the performance objectives established in the MPC cost function. While many options exist for characterizing desirable vehicle trajectories, here, the total sideslip angle at the front wheels (α) was chosen as the trajectory characteristic to be minimized in the objective function. This choice was motivated by the strong influence front wheel sideslip has on the controllability of front-wheel-steered vehicles since cornering friction begins to decrease above critical slip angles. In [18] it is shown that limiting tire slip angle to avoid this strongly nonlinear (and possibly unstable) region of the tire force curve can significantly enhance vehicle stability and performance. Further, the linearized tire compliance model described here does not account for this decrease, motivating the suppression of front wheel slip angles to reduce controller-plant model mismatch. Finally, trajectories that minimize wheel slip also tend to minimize lateral acceleration and yaw rates, leading to a safer and more comfortable ride.

The MPC objective function with weighting matrices $R_{(\cdot)}$ then takes the form

$$J_k = \sum_{i=k+1}^{k+p} \frac{1}{2}\alpha_i{}^{\mathrm{T}} R_\alpha \alpha_i + \sum_{i=k}^{k+p-1} \frac{1}{2}\delta_i{}^{\mathrm{T}} R_\delta \delta_i + \sum_{i=k}^{k+p-1} \frac{1}{2}\Delta\delta_i{}^{\mathrm{T}} R_{\Delta\delta} \Delta\delta_i + \frac{1}{2}\rho_\varepsilon \varepsilon^2 \tag{8}$$

where ε represents constraint violation and was included to soften select position constraints as $\mathbf{y}^{j}{}_{\min} - \boldsymbol{\varepsilon}\mathbf{V}^{j}{}_{\min} \le \mathbf{y}^{j} \le \mathbf{y}^{j}{}_{\max} + \boldsymbol{\varepsilon}\mathbf{V}^{j}{}_{\max}$.

In summary, the MPC controller uses vehicle position, input magnitude, and input rate constraints to satisfy safety requirements, while minimizing front wheel slip to maximize controllability.

Threat Assessment

The vehicle path calculated by the MPC controller is assumed to be the best-case or safest path through the environment. As such, key metrics from this prediction are used to assess instantaneous threat posed to the vehicle. By setting constraint violation weights (ρ_ε) significantly higher than the competing minimization weight (R_α) on front wheel sideslip, a hierarchy of objectives is created in order to force the optimal solutions to satisfy corridor constraints before minimizing front wheel

sideslip. When constraints are not active, as illustrated by the gray vehicle in Fig. 3, front wheel sideslip – and the corresponding controllability threat – is minimized. When the solution is constrained, predicted front wheel sideslip increases with the severity of the maneuver required to remain within the navigable corridor.

The dark vehicle in Fig. 3 illustrates how the MPC-predicted optimal vehicle trajectory might appear as the tire slip angles and corresponding threat increase in the presence of an active constraint. As predicted sideslip approaches tire-cornering-friction-imposed limits, the threat of leaving the navigable corridor increases.

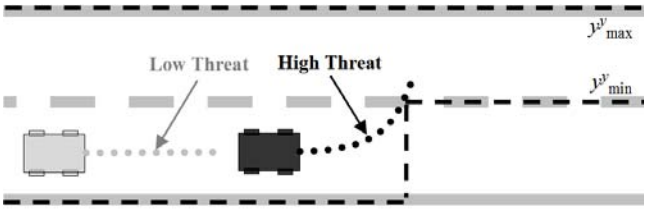

Fig. 3 Obstacle avoidance scenario showing MPC trajectory plans and corresponding threat

Various approaches are available to reduce the vector **α** to a scalar threat metric Φ. In this work,

$$\Phi(k) = \max\left(\alpha_{k+1} \quad \alpha_{k+2} \quad \cdots \quad \alpha_{k+p}\right) \tag{9}$$

was chosen for its good empirical performance when used to regulate controller intervention (described in the next section).

Hazard Avoidance

Given a best-case vehicle path through the environment and a corresponding threat, desired inputs from the driver and controller are blended and applied to the vehicle. This blending is performed based on the threat assessment: a low predicted threat causes more of the driver's input and less of the controller's input to be applied to the vehicle, while high threat allows controller input to dominate that of the driver. This "scaled intervention" may thereby allow for a smooth transition in control authority from driver to controller as threat increases.

Denoting the current driver input by u_{dr} and the current controller input by u_{MPC}, the blended input seen by the vehicle, u_v, is defined as

$$u_v = K(\Phi)u_{MPC} + (1 - K(\Phi))u_{dr} \tag{10}$$

The intervention function K is used to translate predicted vehicle threat (Φ) into a scalar blending gain. This function is bounded by 0 and 1 and may be linear, piecewise-linear, or nonlinear. Linear and piecewise-linear forms of this function may be described by

$$K(\Phi) = \begin{cases} 0 & 0 \le \Phi \le \Phi_{eng} \\ \dfrac{\Phi_{aut} - \Phi}{\Phi_{aut} - \Phi_{eng}} & \Phi_{eng} \le \Phi \le \Phi_{aut} \\ 1 & \Phi \ge \Phi_{aut} \end{cases} \tag{11}$$

In (11), the shape of K is described by the threat level at which the semi-autonomous controller engages (Φ_{eng}) and the level at which it is given full control authority and effectively acts as an autonomous controller (Φ_{aut}).

Using predicted threat (Φ) as calculated in (9) with an appropriate cost function formulation of the form (8) ensures that 1) the threat metric regulating controller intervention is minimized in the path plan (and associated control calculation) and 2) the controller maintains full control authority when constraints are binding.

Increasing Φ_{eng} widens the "low threat" band in which the driver's inputs are unaffected by the controller. While this provides greater driver freedom for low-threat situations, this freedom comes at the cost of increasing the rate of controller intervention when Φ_{eng} is exceeded. This increased rate of intervention may adversely affect driver experience, as discussed in the results below.

Increasing the value of Φ_{aut}, on the other hand, delays complete controller intervention until more severe maneuvers are predicted. The friction-limited bounds on the linear region of the tire force curve (1) suggest a natural upper limit of $\Phi \le 5$ degrees on surfaces with a friction coefficient of 1.0 in order to ensure that by the time the predicted maneuver required to remain within the safe region of the state space reaches this level of severity, the controller has full control authority and can – unless unforeseen constraints dictate otherwise – guide the vehicle to safety.

Simulation Setup

Controller performance was simulated using a vehicle plant model provided by researchers at Ford. This model describes longitudinal and lateral tire forces with the semi-empirical Pacejka tire model where the longitudinal and cornering forces are assumed to depend on the normal force, tire slip angle, surface friction, and longitudinal slip [19].

The vehicle model described by (2-5), with the parameters given in Table 1 was used in the receding horizon controller. Controller parameters are defined and quantified in Table 2 and vehicle velocity was chosen as 14 meters per second.

Table 2 Controller parameters

Symbol	Description	Value [units]
p	Prediction horizon	40
n	Control horizon	20
$R_y^{(\alpha)}$	Weight on front wheel slip	0.2657
R_u	Weight on steering input	0.01
$R_{\Delta u}$	Weight on steering input rate (Δ per Δt)	0.01
$u_{min/max}$	Steering input constraints	± 10 [deg]
$\Delta u_{min/max}$	Steering input rate (per Δt) constraints	± .75 [deg] (15 deg/s)
$y^y_{min/max}$	Lateral position constraints	Scenario-dependent
ρ_ε	Weight on constraint violation	1×10^5
$[\Phi_{eng}\ \Phi_{aut}]$	Thresholds for controller intervention	{[0 3],[1 3],[0 4],[2 4]}°
V	Variable constraint relaxation on vehicle position	[1.25, ···, 1.25, 0.01]

Simulation Results

Simulation results were obtained for various maneuvers, driver inputs, objective function configurations, and intervention laws. Results below are shown for double lane change maneuvers with a driver steer input $\delta_{driver} = 0$. This driver behavior was intended to simulate a drowsy or otherwise inattentive driver.

Semi-autonomous controllers using varying threat thresholds for controller engagement (Φ_{eng}) and full autonomy (Φ_{aut}) successfully satisfied safety constraints while allowing significant driver control whenever possible. In the figure legends below, sideslip thresholds Φ_{eng} and Φ_{aut} (in units of degrees) are labeled as $[\Phi_{eng}\ \Phi_{aut}]$.

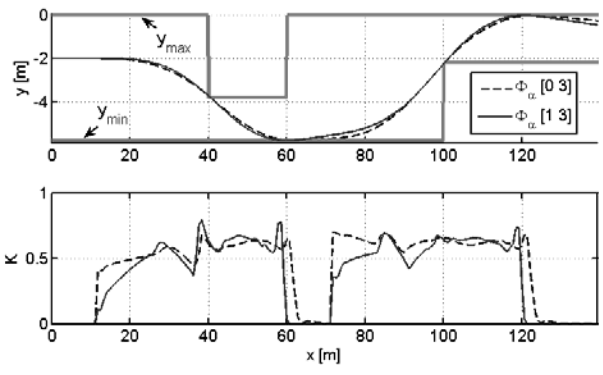

Fig. 4 Simulation results showing the effect of intervention thresholds ($[\Phi_{eng}\ \Phi_{aut}] = [0\ 3], [1\ 3]$) on semi-autonomous corridor-tracking performance

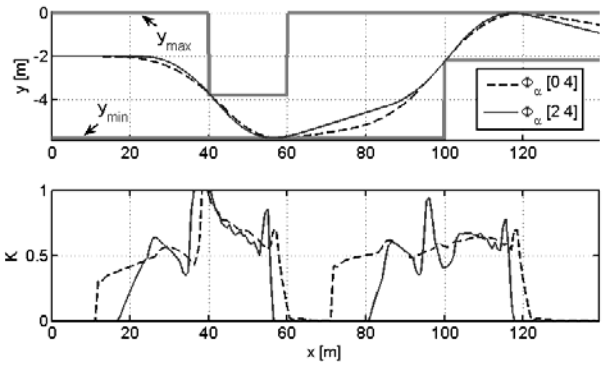

Fig. 5 Simulation results showing the effect of different intervention thresholds ([Φ_{eng} Φ_{aut}] = [0 4], [2 4]) on semi-autonomous corridor-tracking performance

As Fig. 4 and Fig. 5 illustrate, increasing Φ_{eng} delays controller intervention K at the cost of more rapid increases and more frequent saturation of the control authority allotment. This late intervention, while allowing the human driver greater autonomy far away from constraints/hazards, often requires a similar average control authority allotment as it must rapidly and forcefully regain control of the vehicle if the driver does not make the correction on their own. For example, increasing Φ_{eng} from 0 to 1 deg as shown in Fig. 4 ultimately decreased the average intervention K over the entire maneuver by only 12 %. Similar results were observed over the entire range of interest in Φ_{eng} and Φ_{aut} ($0 \leq \Phi_{eng} \leq 2$ and $2.5 \leq \Phi_{aut} \leq 5$), with average intervention K varying by less than 0.19. These results suggest that over the course of some maneuvers, this framework tends to average out controller intervention for various Φ_{eng} and Φ_{aut} settings, allowing for considerable driver preference tuning without dramatically changing average K.

Experimental Setup

Experimental testing was performed using a test vehicle and three human drivers. Driver and actuator steering inputs were coupled via an Active Front Steer (AFS) system. An inertial and GPS navigation system was used to measure vehicle position, sideslip, yaw angle, and yaw rate while a 1 GHz dSPACE processor ran controller code and interfaced with steering actuators.

Three common scenarios, including lane keeping, hazard avoidance, and multile hazard avoidance were used to analyze system performance. In each scenario, obstacles, hazards, and driver targets were represented to the driver by cones and lane markings and to the controller by a constrained corridor (with onboard sensing and constraint mapping assumed to have been performed previously by "virtual sensors" and high-level planners respectively).

Lane keeping experiments tested the threat assessment and intervention characteristics of the controller when the driver maneuvered inside and outside of the given lane. Six pairs of cones were set up along ~200 meters of the 3.35-meter-

wide lane to guide the driver's intended path. As shown in Fig. 6 (not to scale), the second and third sets of cones required the driver to steer the vehicle to the edge of the safe lane while the final two targets required that he attempt to leave the lane.

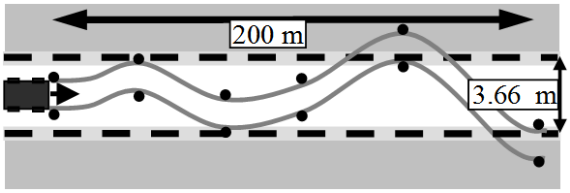

Fig. 6 Lane keeping test setup showing circles where cones were placed to guide the human driver's (unassisted) path. Lane boundaries delineated by dashed lines were represented as constraints y^{y}_{min} and y^{y}_{max} to the controller

Hazard avoidance tests required that the vehicle avoid an obstacle in the current lane of travel. In these tests, the vehicle was driven at a constant velocity in the center of a lane with the driver holding the steering wheel at $\delta = 0$ as if drowsy or inattentive. A row of cones blocked the vehicle's lane of travel, requiring the controller to: 1) plan a stable lane change maneuver around them, 2) assess the threat posed by that maneuver, and 3) intervene as necessary to avoid the hazard. Fig. 7 illustrates this test setup.

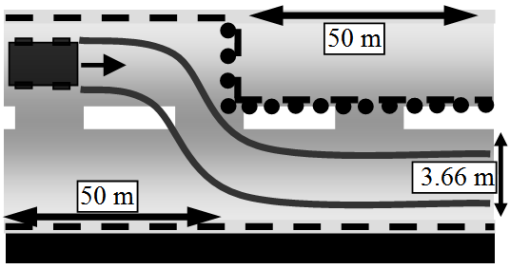

Fig. 7 Hazard avoidance test setup showing hazard cone placement (large circles) and lane boundaries (dashed) enforced by the controller

Multiple hazard avoidance experiments tested the controller's ability to navigate more complex road/hazard setups that required maneuvers with appreciable load transfer. In these tests (illustrated in Fig. 8), both lanes of travel were blocked at different locations, forcing the vehicle to change lanes to avoid the first hazard, then change lanes again to avoid the second as in a double-lane-change maneuver.

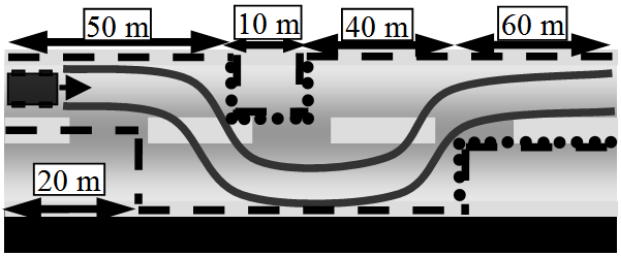

Fig. 8 Multiple hazard avoidance test setup showing hazard cone placement (circles) and lane boundaries (dashed)

Hazard avoidance and multiple hazard avoidance tests were conducted using two different types of driver inputs. Drowsy, inattentive, or otherwise impaired drivers were represented by a constant driver steer input of zero degrees. In these tests, the unassisted driver's path formed a straight line directly through the obstacle(s). To represent active driver steer inputs, the drivers were asked in separate tests to steer either around or into the obstacles. The urgency of these driver steer events was varied – sometimes avoiding the obstacle(s) with a smooth input and other times, steering at the last minute. Controller parameters were chosen as defined in Table 2. Experiments were conducted at 5, 10, and 14 meters per second.

Experimental Results

The semi-autonomous framework proved capable of keeping the vehicle within the navigable corridor for each of the maneuvers tested, with three different human drivers, and using multiple intervention laws. Additionally, the 50 ms sample time proved sufficient for control calculations.

Fig. 9 shows the results of lane-keeping tests.

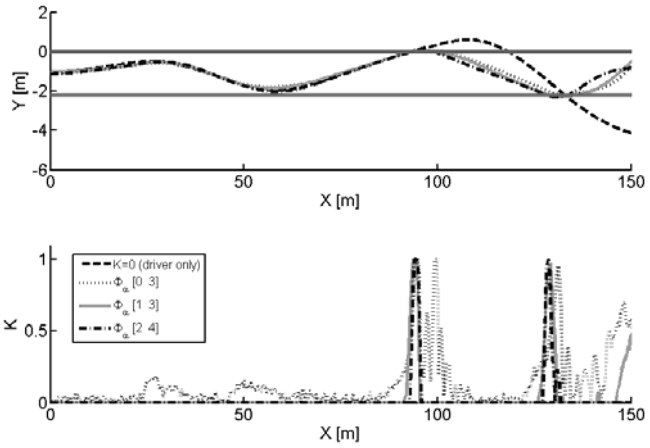

Fig. 9 Results of lane keeping tests with no controller action (dashed), and semi-autonomous controller intervention (dotted, solid, and dash-dot)

The dashed black line in Fig. 9 represents the vehicle trajectory under complete driver control ($K = 0$), and is shown here and in subsequent plots as a reference for the trajectory the driver would have followed had the semi-autonomous controller not engaged. Note that for multiple engagement (Φ_{eng}) and autonomous (Φ_{aut}) intervention thresholds, the semi-autonomous controller successfully kept the driver within the navigable corridor while allowing him significant control authority as long as he remained inside the navigable corridor (x~-20m to x~70m). Only when the vehicle was about to depart from the corridor did the controller intervene. In each case, this intervention was early and large enough to arrest departure. The

slight intervention observed between x~10m and x~50m illustrates the framework's response to a predicted trajectory that required appreciable sideslip in order to remain within the lane. When the driver corrected the vehicle heading, *K* returned to approximately zero. Also note that the inclusion of a low-threat "deadband" ([1 3] and [2 4]) removed much of the noise seen for the experiment without a deadband ([0 3]).

Fig. 10 shows the results of multiple hazard avoidance experiments.

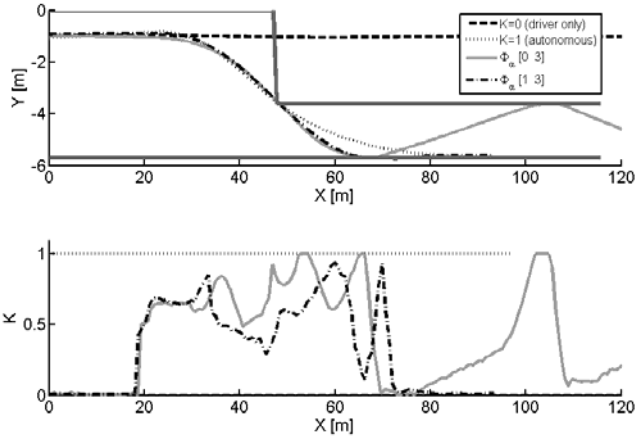

Fig. 10 Results of hazard avoidance tests with no controller action (dashed), fully-autonomous control (dotted), and semi-autonomous control (solid and dash-dot)

Similar to the lane keeping results shown in Fig. 9, controller intervention in hazard avoidance tests was sufficient to maintain the driver's natural/unassisted trajectory for as long as possible before taking control. When the framework did intervene, it allocated enough control authority to the controller to avert corridor departure or loss of control. Note that the trajectory oscillation observed in the [0 3] semi-autonomous experiment was a result of an overcorrection on the part of the controller at x~65m. The vehicle trajectory proceeded to rebound from of $y^{y}_{\max}$ because the driver's input remained at $\delta_{driver} = 0$. Were the driver more attentive as a result of the first intervention incident, the low levels of *K* directly following the initial rebound would have allowed him significant control authority to correct and straighten out the vehicle.

Fig. 10 also shows the results of an autonomous experiment in which the controller was given full control authority (*K*=1). Notice that for the given driver input ($\delta_{driver} = 0$), the vehicle path under semi-autonomous control closely resembles the "best case" path achieved using autonomous control while exerting only an average intervention level (*K*) of 0.34.

Fig. 11 compares a semi-autonomous multi-hazard-avoidance maneuver to an autonomous maneuver (*K*=1).

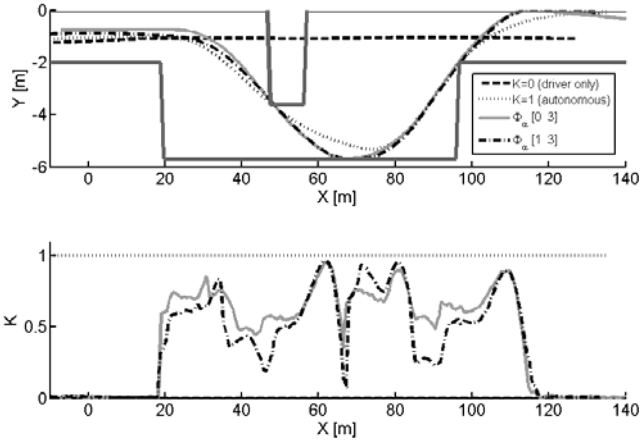

Fig. 11 Multiple hazard avoidance tests showing the similarity between vehicle trajectories under semi-autonomous (solid and dash-dot) and autonomous (dotted) control

Notice that both semi-autonomous controller configurations delayed intervention until the driver's inputs put the vehicle at risk of leaving the navigable road surface. When the framework did intervene, it allocated enough control authority to the controller to avert safe lane departure or loss of control. Even with average controller intervention K_{ave}=0.44, the ultimate vehicle trajectory using the semi-autonomous controller very closely resembles the "best case" trajectory taken by the autonomous controller. This arises from the selective nature of the semi-autonomous system – it intervenes only when necessary, then relinquishes control to the driver once predicted threat to the vehicle has been reduced.

Fig. 12 shows experiments in which the driver was instructed to swerve at the last minute to avoid hazards.

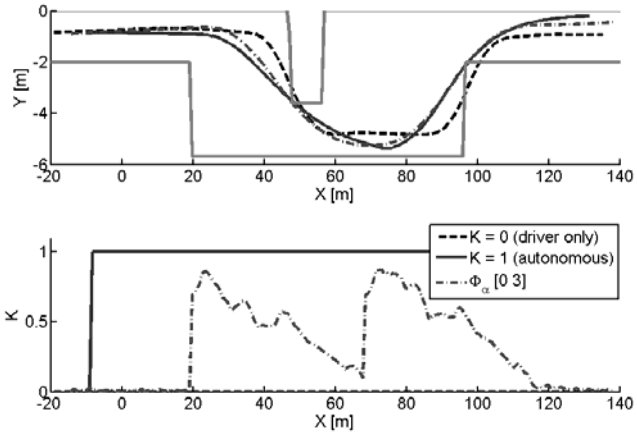

Fig. 12 Multiple hazard avoidance tests showing the vehicle trajectory with an unassisted driver input (dashed) and autonomous controller (solid), and semi-autonomous controller (dash-dot). In each case, the driver swerved to avoid hazards

As Fig. 12 shows, intervention by the semi-autonomous controller slightly preceded an otherwise-late driver reaction. The combined effect of both inputs was then sufficient to avoid both road hazards.

In each of experimental results shown above, the shared-adaptive controller behaves as a stable closed-loop system. While this was also true of all of the other simulated and experimental results conducted to date, no rigorous stability proof is presented in this paper.

Conclusions

This paper presented an optimal-control-based framework that performs trajectory planning, threat assessment, and semi-autonomous control of passenger vehicles in hazard avoidance. This framework has been shown in simulation and experiment to satisfy position, input, and dynamic vehicle constraints using multiple threat metrics and intervention laws. Additionally, this framework has been shown to provide significant autonomy to a human driver, intervening only as necessary to keep the vehicle under control and within the navigable roadway corridor. Simulation and experimental results have shown this control framework to be stable even in the presence of system-inherent time delays, though a rigorous stability proof is a topic of current investigation.

Finally, while human factors have not been studied in depth here, it is expected that with additional investigation, a best-case, or average driver-preferred intervention law may be described and intervention settings tuned accordingly. Further work is needed before this research is road-ready.

Acknowledgments The authors would like to thank Eric Tseng, Matt Rupp, Mitch McConnell, Len Johnson, Steve Hermann, Kyle Carey, Reid Steigler, Tim Zwicky, Roger Trombley, Chris Wallis, and Jeff Rupp, all of Ford Motor Co. for their assistance in conducting the abovementioned experiments.

References

[1] National Highway Traffic Safety Administration (NHTSA), *2007 Traffic Safety Annual Assessment - Highlights*, NHTSA National Center for Statistics and Analysis, 2008.

[2] M. Weilkes, L. Burkle, T. Rentschler, and M. Scherl, "Future vehicle guidance assistance - combined longitudinal and lateral control," *Automatisierungstechnik*, vol. 42, Jan. 2005, pp. 4-10.

[3] J. Leonard, J. How, S. Teller, M. Berger, S. Campbell, G. Fiore, L. Fletcher, E. Frazzoli, A. Huang, S. Karaman, O. Koch, Y. Kuwata, D. Moore, E. Olson, S. Peters, J. Teo, R. Truax, M. Walter, D. Barrett, A. Epstein, K. Maheloni, K. Moyer, T. Jones, R. Buckley, M. Antone, R. Galejs, S. Krishnamurthy, and J. Williams, "A perception-driven autonomous urban vehicle," *Journal of Field Robotics*, vol. 25, 2008, pp. 727-774.

[4] J. Jansson, "Collision avoidance theory with application to automotive collision mitigation," Doctoral Dissertation, Linkoping University, 2005.

[5] J. Pohl, W. Birk, and L. Westervall, "A driver-distraction-based lane-keeping assistance system," *Proceedings of the Institution of Mechanical Engineers. Part I: Journal of Systems and Control Engineering*, vol. 221, 2007, pp. 541-552.

[6] R. Mobus and Z. Zomotor, "Constrained optimal control for lateral vehicle guidance," *Proceedings of the 2005 IEEE Intelligent Vehicles Symposium, 6-8 June 2005*, Piscataway, NJ, USA: IEEE, 2005, pp. 429-34.

[7] A. Alleyne, "A comparison of alternative obstacle avoidance strategies for vehicle control," *Vehicle System Dynamics*, vol. 27, Jun. 1997, pp. 371-92.

[8] M. Netto, J. Blosseville, B. Lusetti, and S. Mammar, "A new robust control system with optimized use of the lane detection data for vehicle full lateral control under strong curvatures," *ITSC 2006: 2006 IEEE Intelligent Transportation Systems Conference, Sep 17-20 2006*, Piscataway, NJ 08855-1331, United States: Institute of Electrical and Electronics Engineers Inc., 2006, pp. 1382-1387.

[9] R. Vaidyanathan, C. Hocaoglu, T.S. Prince, and R.D. Quinn, "Evolutionary path planning for autonomous air vehicles using multiresolution path representation," *2001 IEEE/RSJ International Conference on Intelligent Robots and Systems, Oct 29-Nov 3 2001*, Institute of Electrical and Electronics Engineers Inc., 2001, pp. 69-76.

[10] E.J. Rossetter and J. Christian Gardes, "Lyapunov based performance guarantees for the potential field lane-keeping assistance system," *Journal of Dynamic Systems, Measurement and Control, Transactions of the ASME*, vol. 128, 2006, pp. 510-522.

[11] P. Falcone, M. Tufo, F. Borrelli, J. Asgarit, and H. Tseng, "A linear time varying model predictive control approach to the integrated vehicle dynamics control problem in autonomous systems," *46th IEEE Conference on Decision and Control 2007, CDC, Dec 12-14 2007*, Piscataway, NJ 08855-1331, United States: Institute of Electrical and Electronics Engineers Inc., 2008, pp. 2980-2985.

[12] G. Engelman, J. Ekmark, L. Tellis, M.N. Tarabishy, G.M. Joh, R.A. Trombley, and R.E. Williams, "Threat level identification and quantifying system," U.S. Patent US 7034668 B2, April 25, 2006.

[13] H. Yu, M. Spenko, and S. Dubowsky, "An adaptive shared control system for an intelligent mobility aid for the elderly," *Autonomous Robots*, vol. 15, 2003, pp. 53-66.

[14] L. Guo, J. Wang, and K. Li, "Lane keeping system based on THASV-II platform," *IEEE International Conference on Vehicular Electronics and Safety, ICVES 2006, Dec 13-15 2006*, Piscataway, NJ 08855-1331, United States: Institute of Electrical and Electronics Engineers Computer Society, 2006, pp. 305-308.

[15] J.R. McBride, J.C. Ivan, D.S. Rhode, J.D. Rupp, M.Y. Rupp, J.D. Higgins, D.D. Turner, and R.M. Eustice, "A perspective on emerging automotive safety applications, derived from lessons learned through participation in the DARPA Grand Challenges," *Journal of Field Robotics*, vol. 25, 2008, pp. 808-840.

[16] Z. Zomotor and U. Franke, "Sensor fusion for improved vision based lane recognition and object tracking with range-finders," *Proceedings of the 1997 IEEE Conference on Intelligent Transportation Systems, ITSC, Nov 9-12 1997*, Piscataway, NJ, USA: IEEE, 1997, pp. 595-600.

[17] C. Garcia, D. Prett, and M. Morari, "Model predictive control: theory and practice-a survey," *Automatica*, vol. 25, May. 1989, pp. 335-48.

[18] P. Falcone, F. Borrelli, J. Asgari, H.E. Tseng, and D. Hrovat, "Predictive active steering control for autonomous vehicle systems," *IEEE Transactions on Control Systems Technology*, vol. 15, 2007, pp. 566-580.

[19] T. Keviczky, P. Falcone, F. Borrelli, J. Asgari, and D. Hrovat, "Predictive control approach to autonomous vehicle steering," *2006 American Control Conference, Jun 14-16 2006*, Piscataway, NJ 08855-1331, United States: Institute of Electrical and Electronics Engineers Inc., 2006, pp. 4670-4675.