SEMI-AUTONOMOUS STABILITY CONTROL AND HAZARD AVOIDANCE FOR MANNED AND UNMANNED GROUND VEHICLES

S. J. Anderson, S.C. Peters, and K. Iagnemma*
Department of Mechanical Engineering, Massachusetts Institute of Technology
Cambridge, MA 02139

J. Overholt
US Army Tank Automotive RDE Center (TARDEC)
Warren, MI 48397

1. ABSTRACT

This paper presents a method for trajectory planning, threat assessment, and semi-autonomous control of manned and unmanned ground vehicles. A model predictive controller iteratively replans a stability-optimal trajectory through the safe region of the environment while a threat assessor and semi-autonomous control law modulate driver and controller inputs to maintain stability, preserve controllability, and ensure that the vehicle avoids obstacles and hazardous areas. The efficacy of this approach in avoiding hazards while accounting for various types of human error, including errors caused by time delays, is demonstrated in simulation.

2. 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 (National Highway Traffic Safety Administration (NHTSA) 2008). The longstanding 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 opportunity for improved collision avoidance technologies remains significant.

The need for improved hazard avoidance extends to military operations as well, where accidents involving both manned and unmanned ground vehicles continue to represent a significant challenge to operational effectiveness. In manned ground vehicle operations, such as supply convoys, personnel transport, and reconnaissance missions, vehicle operators face many of the same challenges as those faced by civilian drivers, including path planning, lane-keeping, obstacle avoidance, and stability control. Additionally, operators of military vehicles often deal with significant terrain effects, low visibility conditions, and high situational uncertainty. Consequently, ground vehicle accidents currently represent the largest reported source of non-hostile deaths to U.S. soldiers during operation Iraqi Freedom (Defense Manpower Data Center, Statistical Information Analysis Division 2010).

Unmanned or "teleoperated" vehicles are increasingly being used in a variety of military functions, ranging from surveillance and reconnaissance to detecting and removing hazardous materials. While the advantages of teleoperation are compelling from tactical and human capital perspectives, the challenges associated with remotely operating a vehicle given current technology are daunting. Teleoperated vehicles are typically operated from a control station in which an operator monitors data transmitted from the vehicle and issues commands to the vehicle. Not only must the human operator cope with the challenges inherent to the manned driving task, but he/she must perform many of the same functions with a restricted field of view (FOV), limited depth perception, potentially disorienting camera viewpoints, and significant time delays. Telenavigating a ground vehicle under these conditions while monitoring the vehicle's health status, the status of the mission/tasks, and the condition of the environment leads to high failure rates. In a study of 10 field tests, UGV performance was shown to be relatively low, with mean time between failures ranging from 6 to 20 hours (Carlson and Murphy 2005). Given standard USAR and Department of Defense shifts of 12 and 20 hours, respectively, these results suggest that today's UGVs cannot be reliably depended upon to complete an entire shift.

Recent work on improving the reliability of teleoperated vehicles has focused largely on sensor processing and human interface design. Because of its substantial impact on wirelessly-teleoperated UGVs, much of this research has been devoted to reducing time delay and its perceived effects. Predictive displays using augmented reality, visual tracking, and image-based rendering or "virtualized reality" have been shown in many studies to improve UGV driving performance by 20-60% over standard (live video feed) teleoperation control (Kelly et al. 2009; Kim et al. 1998; Zhenyuan Deng and Jagersand 2003). Further performance improvements have been achieved using multimodal displays (Aleotti et al. 2002) and multimodal inputs (Weimer and Ganapathy 1989).

While advances in sensor processing and human interface design may improve the operator's distance estimation, spatial orientation, object detection, object identification, and situational awareness, and may even reduce the effects of sensing and control latency, reliable UGV operation remains fundamentally limited by the perception, judgment, and driving skill of the human operator. Just as maturation in manned ground vehicle design has resulted in driver error becoming the sole factor in 60% of automobile accidents and a contributing factor in 95% (Evans 1996) so it is

expected that maturation in UGV instrumentation and interface design will ultimately result in driver error being the limiting factor in the performance of teleoperated vehicles (Domme et al. 2007).

Further improvements in the safety and reliability of both manned and unmanned ground vehicles will come from control systems with enough autonomy to correct or override driver error.

Recent developments in onboard sensing, lane detection, obstacle recognition, and drive-by-wire capabilities have facilitated active safety systems that autonomously or semi-autonomously assist in the driving task (Weilkes et al. 2005). In addition to planning a route through the environment, these systems may also assess the threat posed to the vehicle by various hazards and ensure that the vehicle avoids them. In semi-autonomous systems, the planning, threat assessment, and control tasks are made more difficult by the inclusion of a human operator, who maintains at least partial control of the host vehicle. In these scenarios, the trajectory planner and semi-autonomous controller must allow for (and correct or reject if necessary) unanticipated inputs from the human operator.

In the traditional approach to autonomous vehicle navigation, a path planner and controller are arranged as a tiered subsystem, with a planner designing a collision-free path and the controller seeking to track that path while rejecting disturbances. Common path planning approaches include rapidly-exploring random trees (Leonard et al. 2008), graph search methods (Vaidyanathan et al. 2001), potential fields analysis (Rossetter and Christian Gardes 2006), and neural optimization techniques (Dong Kwon Cho and Myung Jin Chung 1991). Control laws commonly employed in these systems include PID schemes (Cremean et al. 2006), linear-quadratic regulators (Bemporad et al. 2002), and nonlinear fuzzy controllers (Tsukagoshi and Wakaumi 1990).

With a human driver in the loop, tiered subsystem architectures that rely on a pre-planned path may be overly restrictive at best and inaccurate at worst. By seeking to limit the vehicle trajectory to a specific path, these approaches neither allow nor account for deviations from the nominal trajectory caused by human inputs or unanticipated hazard motion. Many existing semi-autonomous systems also seek to perform the hazard avoidance task without explicitly accounting for the effect of driver inputs on the vehicle trajectory (Netto et al. 2006). These systems generally estimate the threat posed by static or moving hazards with simple time-based, distance-based, or deceleration-based measures (Fuller 1981; Polychronopoulos et al. 2004; Engelman et al. 2006). While these metrics provide a useful estimate of threat posed by a given maneuver, they are poorly suited to consider multiple hazards, complex vehicle dynamics, actuator and controller limitations, or complicated environmental geometry with its attendant constraints.

In (Anderson et al. 2010a), a framework for semi-autonomous control of passenger vehicles is presented. This framework uses Model Predictive Control (MPC) to iteratively plan trajectories through a traversable corridor, assess the threat this trajectory poses to the vehicle, and regulate driver and controller inputs to prevent that threat from exceeding a given threshold. In the context of static hazards, this system's model-based threat assessment provides an efficient means of: 1) combining various roadway hazards such as lane boundaries and roadway obstacles into realistic spatial constraints and 2) combining these constraints with knowledge of the vehicle dynamics to predict the threat posed by those hazards given the current inputs of a human driver.

This paper demonstrates this framework's potential in military applications of manned and unmanned ground vehicles. The basic framework operation is first presented, followed by a discussion of advantages it offers in teleoperation scenarios. Simulations of controller performance in the presence of unsafe driver inputs and time delays are then presented, and the paper closes with general conclusions.

3. FRAMEWORK DESCRIPTION

The framework described in this paper leverages the predictive and constraint-handling capabilities of MPC to perform trajectory planning, threat assessment, and semi-autonomous 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 derived from (assumed) forward-looking sensor data and a higher-level corridor planner. These boundaries extrapolate the current state of road hazards (vehicles, pedestrians, etc.) to establish a traversable corridor constraining the vehicle's projected lateral 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 treated as 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 intervention required to prevent departure from the traversable corridor, and driver/controller inputs are scaled accordingly. Figure 1 shows a block diagram of this system.

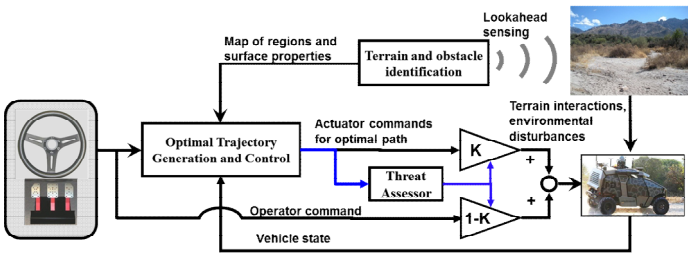

Figure 1. Diagram of an active safety system

3.1 Assumptions

In this paper it is assumed that road lane data is available and that the instantaneous position, velocity, and acceleration of road hazards have been measured or estimated by on-board sensors or vehicle-to-vehicle communication. Existing systems and previous

work in onboard sensing and sensor fusion justify this as a reasonable assumption. Radar, LIDAR, and vision-based lane-recognition systems (Leonard et al. 2008), along with various sensor fusion approaches (Zomotor and Franke 1997) have been proposed to provide the lane, hazard, and environmental information needed by this framework. Where multiple corridor options exist (such as cases where the roadway branches or the vehicle must avoid an obstacle in the center of the lane), it is assumed that a high-level planner has selected a single corridor through which the vehicle should travel.

3.2 Path Planning

The best-case (or baseline) path through the constrained corridor is predicted by an MPC controller. Model Predictive Control is a finite-horizon optimal control scheme that 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. Three important elements of the controller implemented in this paper – the plant model, objective function, and constraint setup – are described below.

3.3 Vehicle Dynamic Model

The vehicle model used by the controller consists of the linearized kinematics of a 4-wheeled vehicle along with its lateral (wheel slip) yaw, and roll dynamics. Vehicle states include the position of its center of gravity [x, y], the vehicle yaw angle ψ, yaw rate $\dot{\psi}$, sideslip angle β, roll angle φ, and roll rate $\dot{\phi}$ as illustrated in Figure 2. The input to the system is the front steer angle δ.

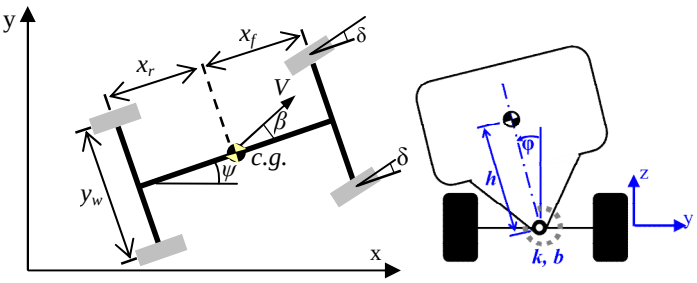

Figure 2. Vehicle model used in MPC controller

Table 1 defines and quantifies this model's parameters.

Table 1. Vehicle Model Parameters

Symbol	Description	Value [units]
m	Total vehicle mass	2450 [kg]
I_{zz}	Yaw moment of inertia	3053 [kg•m^2]
x_f	C.g. distance to front wheels	1.13 [m]
x_r	C.g. distance to rear wheels	1.43 [m]
y_w	Track width	1.56 [m]
C_f	Front cornering stiffness	1640 [N/deg]
C_r	Rear cornering stiffness	1140 [N/deg]
μ	Surface friction coefficient	1
m_s	Chassis sprung mass	1880 [kg]
I_{xx}	Roll moment of inertia	834 [kg•m^2]
h	Sprung c.g. height	0.34 [m]
k_f , k_r	Front & rear axle roll stiffness	30 x 10^3 [N•m/rad]
b_f , b_r	Front & rear axle roll damping	1600 [N•m•s/rad]

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) as in

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

and illustrated in Figure 2.

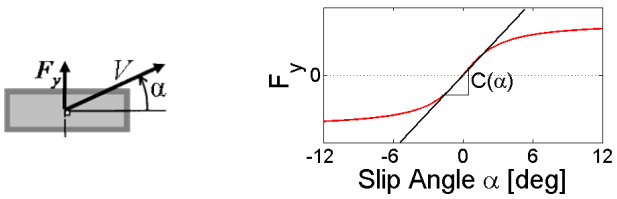

Figure 3. Tire compliance model illustrating velocity *V*, slip angle *α*, and lateral force F_y

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

$$\begin{aligned}
\dot{x} &= V \\
\dot{y} &= V(\psi + \beta) \\
\dot{\beta} &= \frac{-D(C_f + C_r)}{mV}\beta + \left(\frac{D(C_r x_r - C_f x_f)}{mV^2} - 1\right)\dot{\psi} + \\
&\quad \frac{m_s g h^2 - h(k_f + k_r)}{VI_{xx}}\varphi - \frac{h(b_f + b_r)}{VI_{xx}}\dot{\phi} + \frac{DC_f}{mV} \\
\ddot{\psi} &= \frac{(C_r x_r - C_f x_f)}{I_{zz}}\beta - \frac{(C_r x_r^2 + C_f x_f^2)}{I_{zz}V}\dot{\psi} + \frac{C_f x_f}{I_{zz}}\delta \\
\ddot{\phi} &= \frac{-(C_f + C_r)h}{I_{xx}}\beta + \frac{(C_r x_r - C_f x_f)h}{VI_{xx}}\dot{\psi} + \frac{(m_s g h - k_f - k_r)}{I_{xx}}\varphi + \frac{C_f h}{I_{xx}}\delta
\end{aligned} \tag{2}$$

where C_f and C_r represent the (linearized) cornering stiffness of the lumped front wheels and the lumped rear wheels, x_f and x_r are the longitudinal distance from the c.g. of the front and rear wheels, respectively, and $D = 1 + m_s h^2/I_{zz}$. As described in (Anderson et al.

2010a) and briefly discussed below, small slip angles are maintained by configuring the MPC objective function and semi-autonomous intervention law to keep wheel slip within the roughly linear range of the tire force curve. The constant speed assumption used in this linearization reflects the common driver intention in highway driving scenarios.

3.4 Constraint Setup

Assuming that the environment has been delineated previously (see assumptions above), the boundaries of the traversable road surface ($y_{min}(x)$ and $y_{max}(x)$) are sampled over the prediction horizon to generate the constraint vectors

$$\begin{aligned} \mathbf{y}^{y}{}_{\max}(k) &= \left[y^{y}{}_{\max}(k+1) \quad \mathsf{L} \quad y^{y}{}_{\max}(k+p) \right]^T \\ \mathbf{y}^{y}{}_{\min}(k) &= \left[y^{y}{}_{\min}(k+1) \quad \mathsf{L} \quad y^{y}{}_{\min}(k+p) \right]^T \end{aligned} \tag{3}$$

Moving hazards are considered in the autonomous control problem by estimating their future position based on their current position, velocity, and acceleration and excluding predicted collision states from the constraint-bounded corridor (Anderson et al. 2010b).

Constraints on vehicle position are constructed at each sampling instant to form a convex (in y) corridor from the outline of each hazard's anticipated position and depth at time t_c. Figure 4 illustrates what a snapshot aof this time-varying constraint placement might look like to the controller.

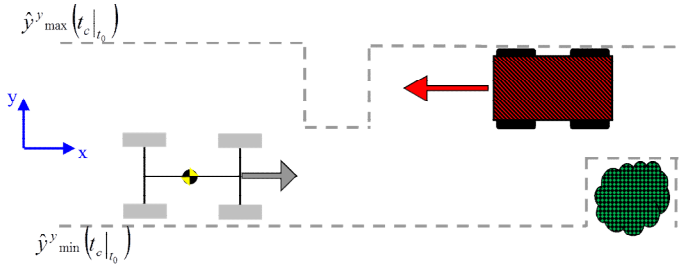

Figure 4. Illustration of constraint placement ($\mathbf{y}^{y}_{\max}$, $\mathbf{y}^{y}_{\min}$) for moving hazards

By enforcing vehicle position constraints at the boundaries of the traversable environment[1], 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 $\mathbf{\Delta u}_{min/max}$, and vehicle dynamic considerations, the traversable operating corridor delineated by $\boldsymbol{y}^{y}{}_{\max}$ and $\boldsymbol{y}^{y}{}_{\min}$ translates to a safe operating region within the state space.

[1] Position constraints may be applied to the vehicle profile or offset and applied to its center of gravity. In this paper, the later approach was used. Results shown here illustrate the host vehicle c.g. traveling through the offset corridor.

3.5 Objective Function

Inside the constraint-bounded corridor, various vehicle outputs such as load transfer, yaw rates, or lateral acceleration, may be minimized to improve vehicle performance, stability, and controllability. In this work, front wheel sideslip ($\alpha = (x_f / V)\psi + \beta - \delta$) is chosen for its influence on the controllability of front-wheel-steered vehicles. As illustrated in Figure 3, cornering friction begins to decrease above critical slip angles. These critical angles are well-known and provide a direct mapping from environmental conditions to vehicle handling limitations. The linearized tire compliance model's failure to account for this decrease further motivates the suppression of front wheel slip angles to reduce controller-plant model mismatch. In (Falcone et al. 2007) 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. Finally, trajectories that minimize wheel slip also tend to minimize lateral acceleration and yaw rates, leading to a safer and more comfortable ride.

Describing the discretized vehicle plant model by

$$\mathbf{x}_{k+1} = \mathbf{A}\mathbf{x}_k + \mathbf{B}_u\mathbf{u}_k + \mathbf{B}_v\mathbf{v}_k \tag{4}$$

$$\mathbf{y}_k = \mathbf{C}\mathbf{x}_k + \mathbf{D}_v\mathbf{v}_k \tag{5}$$

with **x, y, u,** and **v** representing states, outputs, inputs, and disturbances respectively, a quadratic objective function over a prediction horizon of p sampling intervals is defined as

$$J_k = \sum_{i=k+1}^{k+p} \frac{1}{2}\mathbf{y}_i{}^{\mathrm{T}}\mathbf{R}_y\mathbf{y}_i + \sum_{i=k}^{k+p-1} \frac{1}{2}\mathbf{u}_i{}^{\mathrm{T}}\mathbf{R}_u\mathbf{u}_i + \sum_{i=k}^{k+p-1} \frac{1}{2}\Delta\mathbf{u}_i{}^{\mathrm{T}}\mathbf{R}_{\Delta u}\Delta\mathbf{u}_i + \frac{1}{2}\rho_\varepsilon\varepsilon^2 \tag{6}$$

where $\mathbf{R_y}$, $\mathbf{R_u}$, and $\mathbf{R}_{\Delta\mathbf{u}}$ represent diagonal weighting matrices penalizing deviations from $\mathbf{y}_i = \mathbf{u}_i = \Delta\mathbf{u}_i = 0$, ρ_ε represents the penalty on constraint violations, n denotes the number of free control moves, and ε represents the maximum constraint violation over the prediction horizon p. Inequality constraints on vehicle position (**y**), inputs (**u**), and input rates (**Δu**) are then defined as:

$$\begin{aligned} \mathbf{y}^{j}{}_{\min}(i) - \varepsilon\mathbf{V}^{j}{}_{\min}(i) \le \mathbf{y}^{j}(k+i+1|k) \le \mathbf{y}^{j}{}_{\max}(i) + \varepsilon\mathbf{V}^{j}{}_{\max}(i) \\ \mathbf{u}^{j}{}_{\min}(i) \le \mathbf{u}^{j}(k+i+1|k) \le \mathbf{u}^{j}{}_{\max}(i) \\ \Delta\mathbf{u}^{j}{}_{\min}(i) \le \Delta\mathbf{u}^{j}(k+i+1|k) \le \Delta\mathbf{u}^{j}{}_{\max}(i) \\ i = 0,\ldots,\ p-1 \\ \varepsilon \ge 0 \end{aligned} \tag{7}$$

where the vector $\Delta\mathbf{u}$ represents the change in input from one sampling instant to the next, the superscript "(●)j " represents the j^{th} component of a vector, k represents the current time, and the notation $(\bullet)j(k+i|k)$ denotes the value predicted for time $k+i$ based on the information available at time k. The vector **V** allows for variable constraint softening over the prediction horizon, p, when ε is included in the objective function. Table 2 defines and quantifies this controller parameters used in this paper.

Table 2. Controller Parameters

Symbol	Description	Value [units]
$[\Phi_{eng}\ \Phi_{aut}]$	Semi-autonomous intervention thresholds	[0 3] deg
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 ($\Delta/\Delta t$)	0.01
$u_{min/max}$	Constraint on steering input	$\pm$10 deg
Δu_{min}	Constraint on steering input rate ($/\Delta t$)	$\pm$ 0.75 deg
Δu_{max}		(15 deg/s)
y^y_{min} y^y_{max}	Lateral position constraints	Scenario - dependent
ρ_ε	Weight on constraint violation	1 x 10^5
$\mathbf{V}$	Variable constraint relaxation on vehicle position	$\mathbf{V}(1...p\text{-}1)$ =1.25 $\mathbf{V}(p)$ = 0.01
Δt	Controller timestep (update rate)	50 msec

3.6 Threat Assessment and Controller Intervention

Similar to the controller described in (Anderson et al. 2010a) the MPC controller used here constrains vehicle position, input magnitude, and input rates to satisfy safety requirements, while minimizing front wheel slip to maximize controllability and minimize plant-controller model mismatch. At each controller timestep, the predicted front wheel sideslip (α) is converted to a scalar threat metric Φ by

$$\Phi(k) = \max\left(\alpha_{k+1} \quad \alpha_{k+2} \quad \mathrm{L} \quad \alpha_{k+p}\right). \tag{8}$$

This threat assessment is then used in a piecewise-linear intervention function $K(\Phi, u_{driver}, u_{MPC}) \in [0 \quad 1]$ to blend driver and controller inputs as

$$u_{vehicle} = K(\Phi)u_{MPC} + (1 - K(\Phi))u_{driver}. \tag{9}$$

The intervention function K used in (9) is parameterized by threshold threat values Φ_{eng} and Φ_{aut}. While predicted threat remains below Φ_{eng}, the driver maintains full control. As predicted threat increases, so does the proportion of control allocated to the controller. In the simulations and experiments shown below, K was calculated as

$$K = \begin{cases} 0 & 0 \le \Phi \le \Phi_{eng} \\ \dfrac{\Phi_{aut} - \Phi}{\Phi_{aut} - \Phi_{eng}} & \Phi_{eng} \le \Phi \le \Phi_{aut} \\ 1 & \Phi \ge \Phi_{aut} \end{cases} \tag{10}$$

Adjusting intervention thresholds Φ_{eng} and Φ_{aut} can significantly alter the "feel" – though not the effectiveness of – the semi-autonomous controller (Anderson et al. 2010a). 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, it may also increase the rate of controller intervention if Φ_{eng} is exceeded. 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 $\Phi_{aut} \le 5$ 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.

3.7 Key Advantages

This corridor-based approach to path planning, threat assessment, and semi-autonomous vehicle control offers some key advantages to military applications of both manned and unmanned ground vehicles. By planning an optimal state trajectory through a constraint corridor rather than rigidly adhering to a preplanned path, this approach minimizes controller intervention while ensuring that the vehicle does not depart from the safe or navigable region of the environment. It also provides a unified architecture into which various vehicle models, actuation modes, trajectory-planning objectives, driver preferences, and levels of autonomy can be seamlessly integrated without changing the underlying controller structure.

The ability to rapidly adapt to various vehicles, operators, and environmental conditions makes this framework particularly well suited for manned ground vehicles. Unmanned and teleoperated vehicles may also benefit from its ability to ensure safe vehicle operation in presence of unsafe, uninformed, or latency-delayed operator inputs, or in the absence of operator inputs altogether (as when communications are lost). In the absence of control inputs, or when safety-critical scenarios require corrective action, the controller seamlessly transitions from semi-autonomous (human-in-the-loop) to fully autonomous control as necessary to avoid hazards, returning control to the human operator as the threat is reduced or communications are restored. This shared control approach exploits the human operator's complex perception and decision-making skills, while simultaneously reducing or eliminating the occurrence of collisions and loss of control caused by operator error.

4. SIMULATION

4.1 Setup

Controller performance was simulated using a nonlinear ADAMS® model of a generic light truck featuring a double wishbone suspension, passive roll stabilizers, and rack and pinion steering. Tire forces were approximated using a Pacejka tire model, which describes longitudinal and cornering forces as a function of normal force, tire slip angle, surface friction, and longitudinal slip. The vehicle model described by (2), with the parameters given in Table 1 was used in the MPC controller.

A pure pursuit driver model was used to simulate operator inputs as the operator seeks to track a predefined path through the center of a safe corridor. This model implements proportional

feedback on the path tracking error, with the main tuning parameter being the lookahead distance L. Driver steering inputs δ were calculated as

$$\delta = \left(y_{des}(t) - y(t)\right)\frac{2\left(x_f + x_r\right)}{L^2}\sin(\Theta) \qquad (1)$$

where Θ is illustrated in Figure 5.

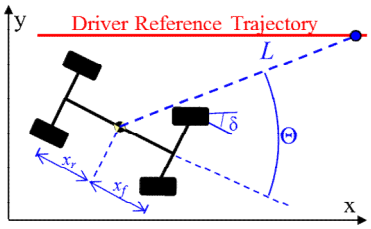

Figure 5. Driver model parameter illustration

Various driver inputs were tested. In the first, the driver failed to steer around a hazard. Such a scenario can occur when a vehicle operator fails to notice an impending hazard or when communications and controls are impaired. The second type of driver input simulated a poor driver. In these simulations, the lookahead distance L was set to 10m, creating relatively large gains on the operator's feedback control. This class was chosen to mimic the impaired driver in manned vehicles and replicate the overactuation observed in the teleoperation of systems with unpredictable latency (Chen et al. 2007). The third type of driver input mimicked a skilled driver. Using a lookahead distance of L = 14m, and in the absence of system latency, this driver model formulation successfully tracked a double-lane-change reference trajectory. The final type of driver input simulated a skilled driver (L = 14m) in the presence of system latency. In these simulations, time delays of 100 – 200 ms were introduced to simulate the effect of feedback and control delays on a remote human operator. Note that because the semi-autonomous controller runs on the vehicle, its state feedback and control inputs are not directly affected by time delays.

4.2 Results

Simulation results were obtained for various maneuvers, driver inputs, objective function configurations, and intervention laws. Figure 6 compares the vehicle trajectory when the operator (traveling at a constant 20 m/s) does not take appropriate evasive action to avoid a hazard to the trajectory obtained with the semi-autonomous controller in the loop. Note that the controller is able to semi-autonomously avoid the hazard without ever taking more than 50% of the available control authority (K) from the human driver. At any time, had the human regained the ability to control the vehicle s/he would have had significant control authority to modify the trajectory of the vehicle.

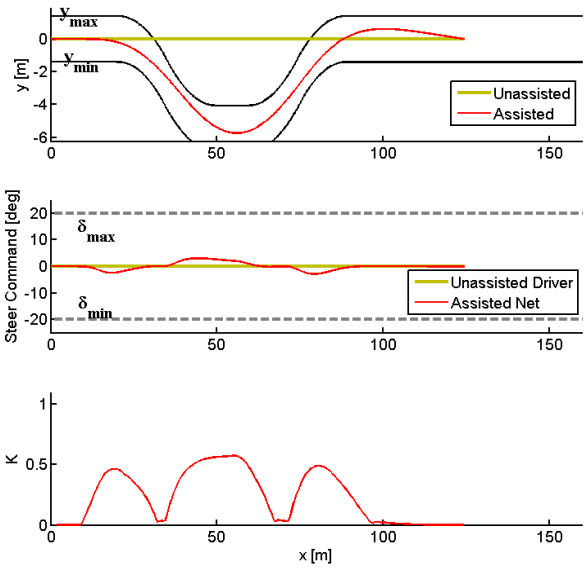

Figure 6. Simulation results demonstrating the effect of semi-autonomous assistance on the vehicle trajectory

In the presence of a closed-loop "pure-pursuit" driver control, the moderating effect of semi-autonomous intervention on both the vehicle trajectory and driver input becomes apparent. Figure 7 shows one such simulation in which the semi-autonomous controller improved the driver's ability to track a desired path. In this simulation, the pure pursuit driver model was designed with a short lookahead (L = 10 m). At 20 m/s, this corresponds to an effective lookahead horizon of 0.5 seconds, which lead to large steering gains and consequent difficulty in tracking the desired trajectory without losing control of the vehicle.

In this scenario, including the semi-autonomous controller in the control loop reduces the magnitude of the driver's inputs. Whereas a short lookahead distance and its attendant high steering gains (1) caused the unassisted driver to oversteer and lose control of the vehicle, the assisted driver was more moderate in its steer commands and thus maintained control of the vehicle. Moreover, allocating less than 50% of the available control authority to the MPC controller was sufficient to keep the vehicle inside the corridor and within 0.4 meters of the desired trajectory. The combined effect of both inputs (driver and controller) is a vehicle trajectory that more closely tracks the driver's desired trajectory than either the pure pursuit controller or corridor-based MPC controller would have done on its own.

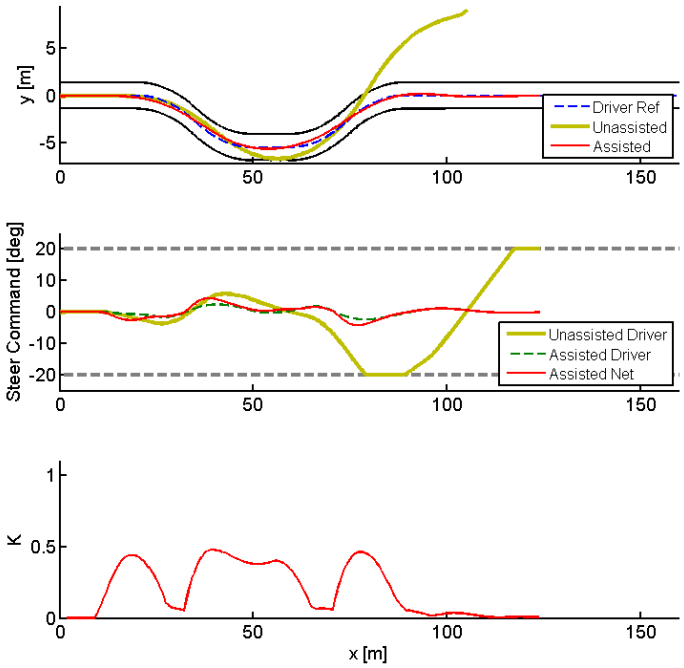

Figure 7. Effect of semi-autonomous intervention on the inputs and performance of an unskilled driver

It is important to note that improved path tracking is not entirely the result of the MPC controller's actions; the controller seeks only to keep the vehicle within a navigable corridor, asserting only as much control authority as the predicted (and stability-related) threat warrants. The "assisted" (i.e. semi-autonomous) trajectory shown in Figure 7 tracks the desired trajectory specifically due to a control authority allocation that allows the driver significant freedom to track a desired path while allowing the MPC controller just enough control authority to stabilize and keep the vehicle within a navigable corridor.

When latency as little as 100ms was introduced between a well-tuned (ie. stable in the absence of time delays) driver model (L = 14m) and the vehicle, the unassisted driver was unable to maintain control of the vehicle while negotiating around a hazard. This instability observed in the presence of time delays as short as a few hundred milliseconds is consistent with experimental observations (Chen et al. 2007). With the semi-autonomous controller in the loop (operating on the vehicle itself), the vehicle successfully negotiates the turn and again causes the human operator to moderate his/her inputs and thereby maintain control and better track his/her desired trajectory.

Figure 8 compares the performance of 1) the (well-tuned) driver model without time delays, 2) the same driver model in the presence of a 200ms time delay, and 3) the semi-autonomously-assisted driver model in the presence of a 200ms delay. Note that, as observed in local or undelayed remote vehicle operation, the semi-autonomous controller has no trouble keeping the vehicle under control while ceding the majority of the available control authority to the human operator. This is unsurprising given the driver-agnostic nature of the MPC-based controller. Because the controller seeks only to keep the vehicle within a constraint-bounded envelope of operation, its threat assessment and intervention function treat any error – human-caused, latency-caused, or otherwise – the same way; if it poses an imminent risk of causing the vehicle to leave the save region of operation, it will intervene as necessary to correct the vehicle's current (and predicted future) trajectory.

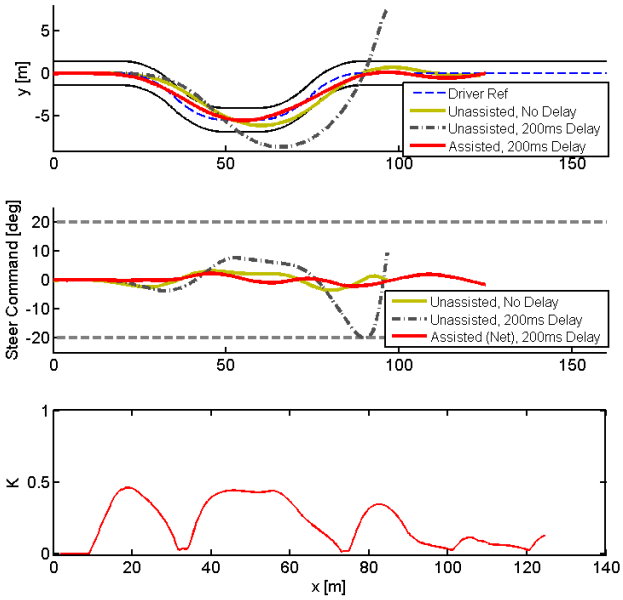

Figure 8. Simulation results illustrating the effect of a 200ms time delay on the driver model and demonstrating the ability of the semi-autonomous framework to assist the driver

5. CONCLUSION

This paper described a semi-autonomous hazard avoidance framework that performs trajectory planning, threat assessment, and shared-adaptive control in both manned and unmanned ground vehicles. This method was shown in simulation to efficiently avoid collisions with environmental hazards while satisfying position, input, and dynamic vehicle constraints in the presence of unsafe, delayed, or absent human inputs. Additionally, the framework was shown to provide significant autonomy to a human driver during low threat situations, intervening only as necessary to keep the vehicle within a safe corridor. Finally, while the simulation results shown here closely correlate with experimental results on a semi-autonomously-controlled passenger vehicle (Anderson et al. 2010a), further research remains to be conducted to understand how latencies and nuances common to teleoperation affect the performance of both the human operator and the vehicle.

REFERENCES

Aleotti, J., S. Bottazzi, S. Caselli, and M. Reggiani, 2002: A Multimodal User Interface for Remote Object Exploration in Teleoperation

Systems. (Accessed August 26, 2010).

Anderson, S. J., S. C. Peters, T. E. Pilutti, and K. Iagnemma, 2010a: An Optimal-Control-Based Framework for Trajectory Planning, Threat Assessment, and Semi-Autonomous Control of Passenger Vehicles in Hazard Avoidance Scenarios. *IJVAS*, **(in press)**.

Anderson, S. J., S. C. Peters, T. E. Pilutti, H. E. Tseng, and K. D. Iagnemma, 2010b: Semi-autonomous avoidance of moving hazards for passenger vehicles. *Proceedings of the 2010 ASME Dynamic Systems and Controls Conference*, Vol. (in press) of, Cambridge, MA.

Bemporad, A., M. Morari, V. Dua, and E. Pistikopoulos, 2002: The explicit linear quadratic regulator for constrained systems. *Automatica*, **38**, 3-20.

Carlson, J., and R. Murphy, 2005: How UGVs physically fail in the field. *Robotics, IEEE Transactions on*, **21**, 423-437.

Chen, J., E. Haas, and M. Barnes, 2007: Human Performance Issues and User Interface Design for Teleoperated Robots. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, **37**, 1231-1245.

Cremean, L. B., T. B. Foote, J. H. Gillula, G. H. Hines, D. Kogan, K. L. Kriechbaum, J. C. Lamb, J. Leibs, L. Lindzey, C. E. Rasmussen, A. D. Stewart, J. W. Burdick, and R. M. Murray, 2006: Alice: An information-rich autonomous vehicle for high-speed desert navigation. *Journal of Field Robotics*, **23**, 777-810.

Defense Manpower Data Center, Statistical Information Analysis Division, 2010: *Military Casualty Summary: Global War on Terrorism*. http://siadapp.dmdc.osd.mil/personnel/CASUALTY/gwot_reason.pdf (Accessed August 24, 2010).

Domme, S., P. Hwang, P. Kim, M. Wang, and B. Sperling, 2007: A value based approach to determining top hazards in army ground vehicle operations. *2006 IEEE Systems and Information Engineering Design Symposium, SIEDS'06, April 28, 2006 - April 28, 2006*, Charlottesville, VA, United states, Inst. of Elec. and Elec. Eng. Computer Society, 124-129 (Accessed August 26, 2010).

Dong Kwon Cho, and Myung Jin Chung, 1991: Intelligent motion control strategy for a mobile robot in the presence of moving obstacles. *Proceedings IROS '91. IEEE/RSJ International Workshop on Intelligent Robots and Systems '91. Intelligence for Mechanical Systems (Cat. No.91TH0375-6), 3-5 Nov. 1991*, New York, NY, USA, IEEE, 541-6 (Accessed March 13, 2010).

Engelman, G., J. Ekmark, L. Tellis, M. N. Tarabishy, G. M. Joh, R. A. Trombley, and R. E. Williams, 2006: Threat level identification and quantifying system.

Evans, L., 1996: The dominant role of driver behavior in traffic safety. *Am J Public Health*, **86**, 784-786.

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

Fuller, R., 1981: Determinants of time headway adopted by truck drivers. *Ergonomics*, **24**, 463-474.

Kelly, A., E. Capstick, D. Huber, H. Herman, P. Rander, and R. Warner, 2009: Real-Time Photorealistic Virtualized Reality Interface for Mobile Robot Control. *Proceedings of the 14th International Symposium on Robotics Research*, The 14th International Symposium on Robotics Research, Lucerne, Switzerland.

Kim, D., S. Kim, J. Kim, C. Lee, and J. Park, 1998: Active operator guidance using virtual environment in teleoperation. *Proceedings of the 1998 IEEE/RSJ International Conference on Intelligent Robots and Systems, October 13, 1998 - October 17, 1998*, Vol. 2 of, Victoria, Canada, IEEE, 1084-1089.

Leonard, J., J. How, S. Teller, et al., 2008: A perception-driven autonomous urban vehicle. *Journal of Field Robotics*, **25**, 727-774.

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

Netto, M., J. Blosseville, B. Lusetti, and S. Mammar, 2006: 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, Institute of Electrical and Electronics Engineers Inc., 1382-1387.

Polychronopoulos, A., M. Tsogas, A. Amditis, U. Scheunert, L. Andreone, and F. Tango, 2004: Dynamic situation and threat assessment for collision warning systems: The EUCLIDE approach. *Proceedings of the 2004 IEEE Intelligent Vehicles Symposium, June 14,2004 - June 17,2004*, Parma, Italy, Institute of Electrical and Electronics Engineers Inc., 636-641.

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

Tsukagoshi, T., and H. Wakaumi, 1990: Highly-reliable semi-autonomous vehicle control on lattice lane. *Proceedings of the IROS '90. IEEE International Workshop on Intelligent Robots and Systems '90. 3-6 July 1990*, New York, NY, USA, IEEE, 731-8.

Vaidyanathan, R., C. Hocaoglu, T. S. Prince, and R. D. Quinn, 2001: Evolutionary path planning for autonomous air vehicles using multiresolution path representation. *Proceedings of the 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems, Oct 29-Nov 3 2001*, Vol. 1 of, Maui, HI, United States, Institute of Electrical and Electronics Engineers Inc., 69-76.

Weilkes, M., L. Burkle, T. Rentschler, and M. Scherl, 2005: Future vehicle guidance assistance - combined longitudinal and lateral control. *Automatisierungstechnik*, **42**, 4-10.

Weimer, D., and S. Ganapathy, 1989: A synthetic visual environment with hand gesturing and voice input. *Conference on Human Factors in Computing Systems (CHI 89), 30 April-4 May 1989, SIGCHI Bull. (USA)*, SIGCHI Bulletin, USA, 235-40.

Zhenyuan Deng, and M. Jagersand, 2003: Predictive display system for tele-manipulation using image-based modeling and rendering. *Proceedings of the 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems, 27-31 Oct. 2003*, Vol. vol.3 of, Piscataway, NJ, USA, IEEE, 2797-802.

Zomotor, Z., and U. Franke, 1997: 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, IEEE Conference on Intelligent Transportation Systems, Proceedings, ITSC*, Piscataway, NJ, USA, IEEE, 595-600.