Annual Reviews in Control 35 (2011) 137–148

Contents lists available at ScienceDirect

Annual Reviews in Control

journal homepage: www.elsevier.com/locate/arcontrol

Pushing the limits: From lanekeeping to autonomous racing

Kirstin L.R. Talvala *, Krisada Kritayakirana, J. Christian Gerdes

Dynamic Design Laboratory, Department of Mechanical Engineering, Stanford University, Building 550, Room 102, 416 Escondido Mall, Stanford, CA 94305, USA

ARTICLE INFO

Article history:
Received 29 November 2010
Accepted 18 February 2011

Keywords:
Autonomous racing
Driver assistance system
Lanekeeping
Limit handling
Lyapunov stability
Vehicle dynamics

ABSTRACT

The success of Electronic Stability Control (ESC) has demonstrated the potential life-saving benefits of vehicle control systems. Lanekeeping presents an obvious next step in vehicle control, but the performance of such systems must be guaranteed before lanekeeping can be viewed as a safety feature. This paper demonstrates that simple lookahead control schemes for lanekeeping are provably robust even at the limits of tire adhesion. By responding to the heading error relative to the desired path, these schemes provide the countersteer behavior necessary to compensate for rear tire saturation and stabilize the vehicle. Using a Lyapunov-based analysis, vehicle stability can be proven even with a highly saturated tire. Taking this a step further by developing a desired path based on the racing line, this lookahead controller can be coupled with longitudinal control based on path position and wheel slip to create an autonomous racecar. The performance of this algorithm shows the potential for lanekeeping control that can truly assist even the best drivers.

1. Introduction

With the introduction of Electronic Stability Control (ESC) systems in automobiles, control engineering has begun to make a huge impact on vehicle safety. Studies of ESC effectiveness have concluded that the system reduces the risk of a fatal crash by more than 40%. In the US, researchers estimate that as many as 10,000 fatal crashes could be avoided each year if all vehicles had this technology (IIHS, 2006). As a result, the US National Highway Traffic Safety Administration has mandated ESC for all new vehicles beginning in the 2012 model year. Similar legislation also exists in Japan, the EU and Australia.

Yet even with this success, automobile accidents remain one of the leading causes of death worldwide. To build upon the safety improvements of ESC, control systems must move beyond the stability of the vehicle itself and incorporate knowledge of the surrounding environment. In particular, future systems must recognize and respond to road boundaries, other vehicles and pedestrians. Among these, fixed lane boundaries represent a logical starting point. Keeping the vehicle in the lane is in many ways the simplest control incorporating the environment, yet one that has significant impact, potentially saving over 11,000 lives a year in the US alone (NHTSA, 2009).

While such lanekeeping systems do exist in production, they are currently marketed as convenience systems and, by design, turn off in critical situations where they are arguably the most needed. For lanekeeping to truly be part of vehicle safety, the robustness of these systems to varying conditions must be guaranteed. In particular, the operation of a lanekeeping system at the very limits of handling (when the tire forces saturate at the limits of friction) must be clearly established.

Several analytical frameworks for lanekeeping systems have emerged in the research community. Reichardt and Schick (1994) view the vehicle as an electron, assign a risk potential to points in the environment and then use an electric field analogy to determine how these risks repel the vehicle. Hennessey, Shankwitz, and Donath (1995) discuss the idea of a virtual bumper, based on the idea of impedance control, where the lane boundaries and other vehicles apply virtual forces to the vehicle that push the vehicle away from danger. Inspired by these approaches, Rossetter and Gerdes (2006) interpret hazards as potential fields and apply an additional guidance force on the vehicle based on the gradient of these potentials. Brandt, Sattel, and Wallaschek (2007) use the idea of potential fields to assign hazards to road edges and obstacles. A virtual string of elastic bands protrudes from the front of the vehicle and stretches according to the defined hazards, defining a collision-free trajectory for the vehicle. Minoiu Enache, Mammar, Netto, and Lusetti (2010) phrase the problem as a hybrid system, switching control authority between the driver and controller, and use Lyapunov methods to design a control law to avoid lane departures. A decision algorithm is employed by Eidehall, Pohl, Gustafsson, and Ekmark (2007) to assess the risk of the current situation, and only activates the system if a lane departure is detected and determined to be dangerous. Other systems are designed around the idea of a virtual driver that either works cooperatively with the human driver (such as in Leelavansuk, Yoshida,

---

* Corresponding author.
E-mail addresses: kirstin.talvala@stanford.edu (K.L.R. Talvala), krisadak@stanford.edu (K. Kritayakirana), gerdes@stanford.edu (J.C. Gerdes).

1367-5788/$ - see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.arcontrol.2011.03.009

& Nagai (2003)) or functions as a steering robot (as in Tseng et al. (2005)).

This paper examines stability and robustness issues for a simple lookahead control scheme. While the development parallels the potential field approach of Rossetter and Gerdes (2006), the connections between lanekeeping and stability that arise generalize to any system that looks ahead on the road using a combination of lateral and heading errors. Using a linear vehicle model, it becomes clear that this system not only works to keep the vehicle in the lane but also provides additional yaw stability. Such conclusions can be extended to nonlinear tire models by accounting for the change in the force–slip relationship as the vehicle reaches the friction limits. It becomes clear that the lanekeeping system always works to assist in vehicle stabilization. These conclusions are borne out by experiments that demonstrate the ability of a simple lanekeeping system to control the vehicle on an uneven gravel surface at the very limits of handling.

Going one step further, a Lyapunov analysis is used to establish stability for various degrees of tire saturation. These results show that a very large handling envelope, extending even past the peak force generation capabilities of the tires, can be mathematically guaranteed. Thus very simple lanekeeping algorithms provide stability even during extreme vehicle handling.

The robustness is so great, in fact, that the simple lanekeeping controller can be used as the basic feedback mechanism for an autonomous racecar, capable of driving at the limits on dirt or paved surfaces. This application not only demonstrates the inherent robustness of the algorithm but also how a system designed to assist an ordinary driver on the highway can also assist even the best driver in extreme driving. The last part of the paper demonstrates how lanekeeping can be extended with feedforward steering, brake and throttle commands and a simple slip controller to form an autonomous racecar. Only the addition of a small amount of yaw damping is necessary beyond the basic system. Experimental results demonstrate how the feedback properties of lanekeeping prevent spins due to rear tire saturation and consistently steer back to the desired path when the front tires saturate.

2. Lanekeeping control system

The lanekeeping controller used in this paper is based on the idea of a virtual potential field, shown conceptually in Fig. 1, that seeks to control the vehicle's lateral error, e, and heading error, ψ, as measured relative to the desired path's tangent and defined in Fig. 2. It is intended to be used as a driver assistance system, working in conjunction with a driver. A discussion on this particular lanekeeping system can be found in Rossetter and Gerdes (2006) and details on the experimental implementation and validation can be found in Rossetter, Switkes, and Gerdes (2004).

For this work, it is assumed that the desired path is known and that measurements of the lateral and heading errors (e and ψ) are available. The desired path can, for example, be defined from Global Positioning System (GPS) data, as is the case for the experiments discussed in Sections 3 and 6. It can also be constructed from vision and/or ranging sensor data as is done, for example, in Eidehall et al. (2007). Measurements of e and ψ can similarly be obtained from GPS data (integrated with an Inertial Navigation System (INS) as in Sections 3 and 6) or from various vision systems (McCall & Trivedi, 2006).

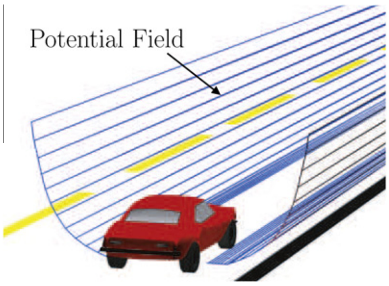

Fig. 1. Visualization of potential field.

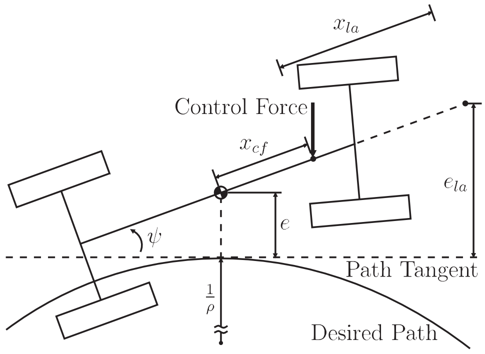

Fig. 2. Definition of lateral error, e, and heading error, ψ, as measured relative to the desired path's tangent. The instantaneous road curvature is denoted ρ.

The potential field, V_c, is defined as:

$$V_c = k_{LK} e_{la}^2 = k_{LK}(e + (x_{cf} + x_{la})\sin\psi)^2 \quad (1)$$

where the lookahead term, x_{la}, is included for stability. The potential field gain, k_{LK}, can be chosen to control, for example, the maximum lane deviation. Using this potential field, the control force, F_c, is defined as:

$$F_c = -\partial V_c / \partial e_{la} = -2k_{LK} e_{la} = -2k_{LK}(e + (x_{cf} + x_{la})\sin\psi) \quad (2)$$

In other words, the force applied to the vehicle is simply a linear gain multiplied by a projected error that nudges the vehicle back towards the center of the lane. Alternatively, this controller can be viewed as separate gains on the lateral error, e, and the heading error, ψ. The heading error component functions as a torsional spring to align the vehicle's direction of travel with the tangent to the desired path in the lane, while the lateral error component functions as a linear spring to align the vehicle's lateral position with that of the desired path.

The control force is applied to the vehicle at a distance x_{cf} in front of the center of gravity (CG) of the vehicle, as shown in Fig. 2. If the front steering is used to apply this force, the distance x_{cf} is set to be the distance to the front axle, a. If other actuators are available, the distance x_{cf} can be moved. The lookahead distance, x_{la}, can be chosen to guarantee that the linear system is stable (Rossetter & Gerdes, 2006) and can be used to fine tune the closed loop vehicle dynamics.

2.1. Vehicle model

The base model used for the vehicle dynamics in this paper is the "bicycle model," Fig. 3, in which the vehicle is treated as a mass with inertia moving in the plane. Assuming a constant longitudinal velocity, U_x, the equations of motion are:

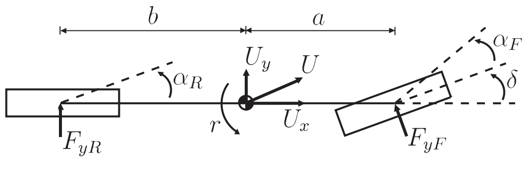

Fig. 3. Planar vehicle model.