Autonomous vehicle control at the limits of handling

Krisada Kritayakirana*
and J. Christian Gerdes

Dynamic Design Laboratory,
Mechanical Engineering Department,
Stanford University,
California 94305, USA
E-mail: krisadak@alumni.stanford.edu
E-mail: gerdes@stanford.edu
*Corresponding author

Abstract: Racecar drivers have the ability to operate a vehicle at its friction limit without losing control. If autonomous vehicles or driver assistance systems had similar capabilities, many fatal accidents could be avoided. To advance this goal, an autonomous racing controller is designed to gain insights into vehicle control at the friction limits. A bicycle model and a '*g-g*' diagram are used to mimic racecar drivers' internal vehicle model. Lanekeeping steering feedback and wheel slip feedback controllers are used to imitate drivers making steering and throttle corrections according to the vehicle responses. Experimental results on a low friction surface demonstrate that the controller can robustly track a path while operating at the limits of tyre adhesion and provide insights for the future development of vehicle safety systems.

Keywords: autonomous vehicle; vehicle dynamics and control; limit handling; driver assistance systems; clothoid; traction circle; trail braking; lanekeeping; throttle on exit; '*g-g*' diagram; slip circle; bicycle model; speed tracking; racing; motorsport.

Biographical notes: K. Kritayakirana received the BEng degree from Chulalongkorn University, Thailand, in 2001 and the MPhil degree from University of Cambridge, England, in 2002. He is currently a PhD candidate at Stanford University, USA. His research interests focus on vehicle dynamics and control at the limits of handling. Prior to coming to Stanford University, he worked at Lotus Cars in England and led the Lotus 119c project, founded a startup focused on aftermarket car parts and was an Adjunct Professor at Chulalongkorn University. He is a Fulbright Science and Technology scholar as well as a Shell Centenary Scholarship alumnus.

J.C. Gerdes is an Associate Professor with the Department of Mechanical Engineering and, by courtesy, the Department of Aeronautics and

Astronautics at Stanford University, USA. He is also the Director of the Center for Automotive Research at Stanford (CARS). His research centres on feedback control techniques for mechanical systems with a specific focus on improving vehicle safety and efficiency. Prior to joining Stanford University, he was the Project Leader for virtual proving grounds development at Daimler-Benz Research and Technology, North America. He is a recipient of the Presidential Early Career Award for Scientists and Engineers for his work with driver assistance systems.

1 Introduction

Motorsport has consistently pushed the technology boundary and driven many developments in the automotive industry. Historically, race cars have provided the inspiration for many developments in production vehicles. With the emergence of autonomous vehicles, racecar drivers now provide similar inspiration for the design of vehicle safety systems. The ability to control an autonomous vehicle in any extreme condition is crucial for vehicle safety systems, which can greatly impact the public acceptance of autonomous vehicles and the automotive industry as a whole. Since racecar drivers have the ability to control a vehicle at its friction limits without losing control, they are natural models when designing controllers that mimic this behaviour.

Previous work in vehicle control at the limits of handling has been more mathematically focused, using methods that couple path generation (the racing line) and vehicle control into one optimisation problem. Various optimisation techniques such as optimal control (Gerdts et al., 2009; Velenis et al., 2008), Model Predictive Control (Timings and Cole, 2010) and genetic and evolutionary techniques (Muhlmeier and Muller, 2002) have been used to find an optimised racing line. In these works, since both path and control inputs are coupled, gaining a physical intuition from the results can be difficult. In addition, since the optimisation results are generally based on simulations, the robustness of the resulting controllers is not clear. The robustness of the controller is essential since friction variation and modeling errors inevitably exist on real vehicles.

The racecar driver's approach to the problem is somewhat different. Lopez (2001), Bentley (1998) and Taruffi (1959) explain how racecar drivers separate the tasks of finding a racing line and driving a vehicle at its limits. To understand vehicle limits, Rice (1973) introduced the concept of a friction limit circle on a '*g-g*' diagram. Racecar drivers use trail-braking and throttle-on-exit techniques to trace this friction limit circle, as described by Lopez (2001), Mitchell et al. (2004), Bentley (1998) and Velenis et al. (2008). Mimicking racecar drivers, the research presented in this paper separates the problem of autonomous racing into path planning and vehicle control. Thus, with any given desired path, the controller can calculate appropriate inputs that will drive the vehicle to its friction limits while tracking the path. This is similar to the plant inversion technique used by Mueller-Bessler et al. (2008) and Mueller-Bessler et al. (2008) in driving a double lane change manoeuvre. However, unlike their approach where the existing Electronic Stability Program (ESP) and Anti Lock Brake System (ABS) were used, the controller

presented in this paper only provides steering, throttle and brake inputs, similar to the inputs available to humans. It uses various vehicle states for feedback to ensure stability and tracking ability at the limits. To test the robustness of this controller, experiments are performed on a varying low friction surface with all the driving aids (ESP and ABS) switched off.

Moreover, in addition to designing a controller that drives an autonomous vehicle to its friction limits, the design process also considers the application of these algorithms to future driver assistance systems. With driver assistance systems in mind, the controller is constructed from individual control modules (i.e., feedforward steering controller, feedback steering controller, etc.) that serve specific tasks. Consequently, the attributes of each controller module and how they contribute to the robustness of the whole system can be understood. Furthermore, each control module can be used individually to assist a driver; a steering feedback module can, for instance, be used in conjunction with a driver to provide assistance when the driver is on the verge of losing control.

The remainder of this paper is organised as follows. Section 2 describes how the overall controller structure is divided into path generation and tracking a path at the limits. A 'g-g' diagram is used as a mathematical model to explain how racecar drivers understand their friction limits. Trail-braking and throttle-on-exit driving techniques, which make the vehicle's acceleration trace the friction limit circle on a 'g-g' diagram, are then discussed. Section 3 explains how a path is structured from clothoid curves, which offer a number of advantages due to their linear curvature. With a path defined, Section 4 describes how drivers' driving abilities are translated into feedforward and feedback controllers. To track a path at the limits, a quasi-static bicycle model and a 'g-g' diagram are used to calculate the feedforward steering, throttle and brake inputs. Imitating a racecar driver responding to tracking errors, a steering controller based on lanekeeping assistance is used for the steering feedback while a wheel slip controller is used for modulating brake and throttle. Section 5 describes the instrumentation of an Audi TTS used as a test vehicle and explains how feedforward and feedback commands are implemented. The experimental results are then discussed, explaining how the controller operates through a corner at the limits. The results also demonstrate how lanekeeping and wheel slip controllers provide stability and robustness to the system, even when the friction is over-estimated. Moreover, experimental results highlight the importance of coordinating steering and longitudinal inputs.

2 Overall architecture of the autonomous racing controller

A racecar driver's goal is to win a race, which means finishing the race with the fastest time. To be fast around a track, a racecar driver has to fully utilise tyre forces while tracking a racing line. To finish a race, a driver has to avoid losing control of the vehicle while driving at the limits. Furthermore, a great driver has to perform consistently, regardless of the changes in the environment or in the vehicle's characteristics (robustness). Thus, a racing controller has to robustly track a desired path, utilise tyre forces, and stabilise the vehicle.
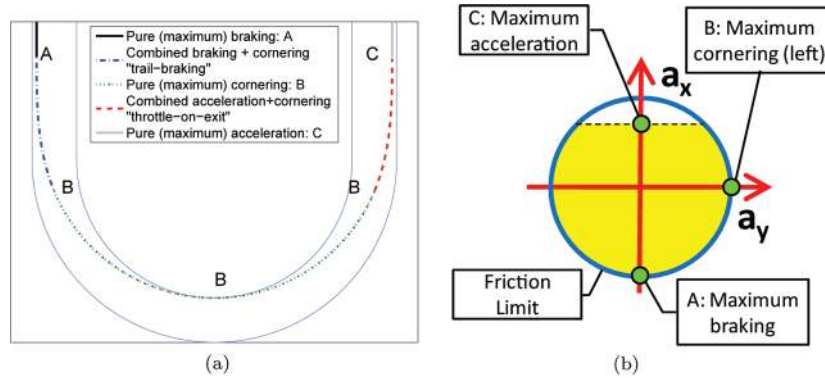
2.1 Understanding racecar drivers' behaviour

Racecar drivers can achieve the above objectives through their abilities to estimate friction and utilise friction information to control a vehicle at the limits. During practice, they formulate a racing line and through their understanding of vehicle limits, they coordinate their inputs to follow their racing line at the limits of tyre adhesion.

2.1.1 Cornering sequences

Figure 1(a) depicts how racecar drivers structure their racing line when driving through a turn (Lopez, 2001; Bentley, 1998; Taruffi, 1959). At the corner apex, they pick a constant radius arc (which corresponds to a constant steering angle) that has the largest radius, to maximise vehicle speed. In order to connect between straight and constant radius sections, a transition curve is needed during corner entry and corner exit. During these transitions, drivers have to increase or decrease their steering angle. Thus, in general, each cornering sequence consists of a corner entry where the curvature transitions from zero (straight) to a constant value, a mid corner where the turning radius is constant, and a corner exit where the curvature changes from a constant value back to zero. Once a racing line is established, racecar drivers have to understand their limits before they can provide correct inputs to the vehicle.

Figure 1 Driving at the limits: (a) racing line's structure and (b) introduction to a '*g-g*' diagram (see online version for colours)

Source: Adapted from Bentley (1998)

2.1.2 Understanding vehicle limits on a '*g-g*' diagram

Racecar drivers follow their racing line while ensuring that they maximise their tyre forces governed by the friction between the tyres and the track surface. They use a concept of a friction circle (or traction circle) to conceptualise these limits (Lopez, 2001; Bentley, 1998). To translate this concept into a mathematical model, a '*g-g*' diagram (Rice, 1973; Milliken and Milliken, 1995) is used.

The force an individual tyre can produce is limited by friction. For an isotropic tyre, these limits can be viewed as a friction circle on axes of lateral and longitudinal force. The friction forces from the four tyres define the vehicle's capability. While driving, the capability of each tyre is continuously changing due to longitudinal and lateral load transfer and aerodynamic downforce effects. These can be modelled to give a full sense of the vehicle's limits in different trim states, but for simplicity, the resulting diagram is often simplified to the same basic circular shape as the isotropic tyre. This is known as the friction circle or traction circle for the vehicle.

Force and acceleration are related through Newton's second law. Since the acceleration of a vehicle can be measured more easily than the forces, the ability of a driver to operate a vehicle at its limits is often evaluated using a friction limit circle on a '*g-g*' diagram, shown in Figure 1(b). Thus, the acceleration that the driver achieved at different points along the course can be compared to the maximum acceleration predicted by a simple '*g-g*' diagram or a more complicated representation incorporating load transfer.

In order to maximise tyre forces, racecar drivers try to operate the vehicle on the friction limit circle in Figure 1(b) at all times. They can maximise their braking without cornering and vice versa, or they can use a combination of lateral and longitudinal forces that generates the maximum tyre forces.

2.1.3 Driving at the limits

At this point, a racing line and the vehicle limits have been established. The remaining task is to modulate brake, throttle and steering inputs to track the racing line in Figure 1(a) while tracing the friction limit circle in Figure 1(b). To drive at the limits, racecar drivers apply pure braking during the last straight section before the corner entry at point A in Figure 1(b). Similarly, during pure cornering around the apex, they maximise their tyre forces in a lateral direction without using any longitudinal forces (point B). At point C, the driver applies the maximum longitudinal acceleration allowed by the engine's torque (the dashed line in Figure 1(b)). The challenge arises during the transition phases (corner entry and exit) where drivers have to coordinate their steering and longitudinal inputs by using trail-braking and throttle-on-exit techniques.

Trail-braking and throttle-on-exit

During corner entry, racecar drivers have to make a transition from the straight to constant radius section in Figure 1(a), while tracing the friction limit circle from point A to point B in Figure 1(b). They have to delicately balance their tyre forces between braking and steering to trace the fourth quadrant of the friction limit circle in Figure 1(b). This technique is called trail-braking (Lopez, 2001; Velenis et al., 2008) where drivers slowly decrease the amount of braking while increasing the amount of steering. Similarly, during corner exit, drivers use a throttle-on-exit technique where they slowly increase their throttle input while unwinding the steering wheel to trace the first quadrant of the friction limit circle (from point B to point C).

The tradeoff between the available cornering and braking forces during corner entry highlights the importance of corner entry speed. If drivers approach the corner entry too fast, tracking the original racing line becomes infeasible. Higher

vehicle speed requires higher cornering force, which consequently reduces the available braking force. Thus, as the corner becomes tighter, a vehicle traveling at an excessive speed will prematurely use all of the friction capability for cornering and not be able to track the racing line. If drivers choose to stay on the brake, they no longer have sufficient cornering force to remain on the path. Thus, no matter what they do, if the vehicle approaches the corner with excessive speed, it will deviate from the desired path. To avoid this problem, they have to hit the braking point properly in order to achieve the correct entry speed. If the drivers notice that they are approaching the corner too fast, they have to modify their path to create a 'recovery line' (Bentley, 1998; Klomp, 2010). Notice that this dilemma does not exist in the corner exit. Lifting the throttle reduces vehicle speed without sacrificing any cornering ability, because the curvature is decreasing.

Once the knowledge of driving at the limits is established, a racing controller can be designed to mimic these behaviors.

2.2 A controller architecture for racing

To achieve path tracking at the limits, the controller separates path generation from path tracking at the limits (Figure 2). The path is structured to mimic a racecar driver's racing line, and a friction limit circle on a 'g-g' diagram is used to describe the vehicle limits. In this paper, the controller uses a priori knowledge of friction and relies on the controller robustness to handle any variation of the track surface. This can easily be extended to include realtime friction estimation based on tyre slip or aligning moment (Hsu et al., 2010).

**Figure 2** Overall controller structure (see online version for colours)
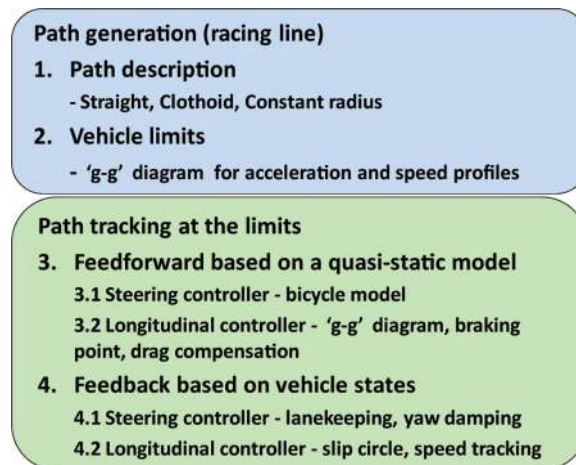
Similar to racecar drivers planning ahead, feedforward controllers are designed to calculate steering and longitudinal inputs that track the path while on the friction limit circle. While driving, feedback controllers imitate a driver's vehicle control ability, making adjustments based on vehicle responses and tracking errors. The next few sections discuss each of these aspects in detail.

3 Path description

To mimic racecar drivers' behaviour, the controller divides the cornering sequence into a corner entry, a constant radius and a corner exit, see Figure 3. Transition curves (corner entry and exit) are required in order to connect between straight and constant radius sections. Any abrupt change during these transitions is undesirable as it can upset vehicle balance. A clothoid curve, which has linear curvature (Figure 3(b)), is chosen to approximate this transition as it produces a linear steering input between straight and constant radius sections.

Figure 3 Clothoid map: showing (a) shape and (b) curvature (see online version for colours)
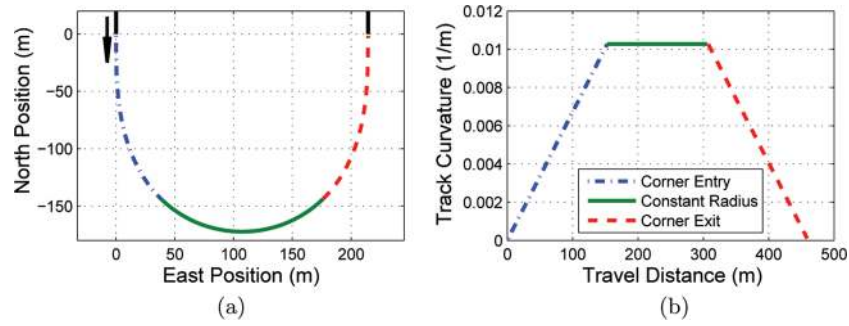
A clothoid path (also known as an Euler Spiral) is commonly used in highway road design (AASHTO, 2004). It is based on a Fresnel Integral (Jeffrey and Dai, 2008) and contains only one parameter c that describes the shape of the clothoid. This parameter c controls the rate that the curvature changes along the segment, i.e., the slope of the curvature in Figure 3(b). A clothoid can be expressed as:

$$x = C(\acute{s}) = \frac{1}{c} \int_0^{\acute{s}} \cos(q^2) \, dq$$
$$y = S(\acute{s}) = \frac{1}{c} \int_0^{\acute{s}} \sin(q^2) \, dq \tag{1}$$

where $\acute{s} = cs$ and s is the distance measured along the clothoid segment.

The curvature $k(s)$ of a track can simply be found from Gray (1997):

$$k(s) = \frac{1}{R(s)} = \frac{\dot{x}\ddot{y} - \dot{y}\ddot{x}}{(\dot{x}^2 + \dot{y}^2)^{\frac{3}{2}}} \tag{2}$$

where $R(s)$ is the radius of the clothoid curvature, $\dot{x} = \frac{dx}{ds}$, $\dot{y} = \frac{dy}{ds}$, $\ddot{x} = \frac{d^2x}{ds^2}$ and $\ddot{y} = \frac{d^2y}{ds^2}$. By substituting x and y from equation (1) into equation (2), the clothoid curvature is found to be:

$$k(s) = \frac{1}{R(s)} = 2c^2 s. \tag{3}$$

This linear relationship between curvature k and the distance traveled s produces a smooth feedforward steering that will be described in Section 4.1.2. The linear relationship between distance traveled and curvature greatly simplifies the feedforward longitudinal calculation in Section 4.1.1. In addition, a clothoid curve has many established methods for curve fitting within some geometry constraints (Walton and Meek, 2005; Baran et al., 2010).

4 Path tracking at the vehicle's limits

Both the steering and longitudinal controllers consist of feedforward and feedback components, which will be described in this section. First, a quasi-static model is used to predict feedforward longitudinal and steering inputs. A 'g-g' diagram is utilised in the feedforward longitudinal controller to ensure that the vehicle is operating at its limits, while a bicycle model is used to predict how much the controller should steer in order to track a desired path. To add stability and robustness, a lanekeeping steering feedback system (Rossetter et al., 2004) is used to minimise any errors caused by modeling errors or disturbances to the system. To minimise vehicle yaw oscillation when driving at the limits, a small yaw damping feedback is added into the controller. In addition, a slip circle that monitors wheel slip is used for the longitudinal feedback. This slip circle feedback minimises excessive tyre slip, enabling the steering system to work effectively and detecting if the vehicle is not at its limits. Furthermore, the controller monitors its speed as mentioned in Section 2.1.3. Thus, a speed feedback is used to ensure that a vehicle will not approach a corner with an excessive speed.

4.1 Feedforward based on a quasi-static model

4.1.1 Calculating desired longitudinal acceleration from a predefined path

The friction limit circle in Figure 1(b) is found from a ramp steer manoeuvre. Once the limits are known, the goal of the longitudinal feedforward algorithm is to compute throttle and brake inputs that fully utilise the tyre forces. The friction limit circle in Figure 1(b) can be translated into a simple mathematical relationship between a_x and a_y as follows:

$$(\mu g)^2 = a_x^2 + a_y^2 \tag{4}$$

where a_x is vehicle's longitudinal acceleration, a_y is vehicle's lateral acceleration and μ is an effective friction coefficient.

To estimate the feedforward longitudinal acceleration along the clothoid segment s, the calculation starts by substituting curvature k from equation (3) into a_y approximated by steady state cornering:

$$a_y(s) \approx \frac{U_x(s)^2}{R(s)} = U_x(s)^2 k(s) = U_x(s)^2 2c^2 s \tag{5}$$

where $U_x(s)$ is the desired speed along the traveled distance s and $R(s)$ is the radius of the curvature along the path s.

The curve in Figure 3 is divided into three sections and each has different calculation processes. The calculation starts by estimating the desired vehicle speed at the corner's apex and integrates along the path to find the desired speed along the path $U_x(s)$.

Constant radius: At mid corner, the vehicle is operating at maximum cornering (point B in Figure 1(b)), where $a_x = 0$ and $a_y = \mu g$. The speed at this point, can be found from using equations (4) and (5). As a consequence, $U_x(s) = \sqrt{\mu R(s)g}$, where R is the radius of the curvature at the apex.

Clothoid entry: The goal of the feedforward longitudinal control is to calculate the correct amount of $a_x(s)$ at each point along the segment s, so that when combined with the $a_y(s)$ generated from vehicle cornering, the vehicle's acceleration will trace counterclockwise along the fourth quadrant of the friction limit circle.

The overall concept of finding feedforward longitudinal acceleration $a_x(s)$ relies on the process of finding $U_x(s)$. Once $U_x(s)$ is known, $a_y(s)$ can be estimated from equation (5). Thus, by knowing the available friction μ and $a_y(s)$, $a_x(s)$ can be found from equation (4).

To find $U_x(s)$, the calculation starts from substituting $a_y(s)$ from equation (5) into equation (4), and rearranging to obtain:

$$a_x(s) = \sqrt{(\mu g)^2 - (2c^2 s U_x(s)^2)^2}. \tag{6}$$

From the differential equation of $a_x(s)$:

$$a_x(s) = \frac{\mathrm{d}U_x(s)}{dt} = \frac{\mathrm{d}U_x(s)}{ds}\frac{\mathrm{d}s}{dt} = \frac{\mathrm{d}U_x(s)}{ds}U_x(s) \tag{7}$$

where t is time. Substituting this definition of $a_x(s)$ into equation (6) obtains:

$$\frac{\mathrm{d}U_x(s)}{\mathrm{d}s} = \frac{1}{U_x(s)}\sqrt{(\mu g)^2 - (2c^2 s U_x(s)^2)^2}. \tag{8}$$

Equation (8) is then solved in realtime by using a backward integration. The integration starts from the end of the clothoid entry where $s = L^{\mathrm{cloth}}$ (point B in Figure 1(b)) and ends at the beginning of the clothoid where $s = 0$ and $k = 0$ (point A). The speed calculated from the constant radius section is used as the initial condition.

After $U_x(s)$ along the clothoid entry is found from equation (8), the desired feedforward longitudinal acceleration $a_x(s)$ along s can be found from equation (6).

Straight: Once the vehicle speed along the path $U_x(s)$ is found, a corner entry speed and the braking point on the preceding straight section can be determined. This braking point is an important parameter. Braking too early or too much will result in a slower corner entry speed (not fully utilising the tyre forces). Braking too late or not sufficiently will result in vehicle sliding or deviation from the intended path as mentioned in Section 2.1.3.

Thus, in order to achieve the correct corner entry speed, the controller has to start braking in the section before reaching the corner entry. If the section before corner entry is a straight section, the braking distance s_{braking} can be found from:

$$s_{\mathrm{braking}} = \frac{(U_x)^2 - (U_x^{\mathrm{entry}})^2}{2\mu g} \tag{9}$$

where U_x is the current vehicle speed and U_x^{entry} is the desired corner entry speed. In practice, the effect of the time delay in the braking system has to be compensated by the controller. Conservatism can also be added into equation (9), by slightly scaling μ down, to ensure sufficient braking distance is available to achieve the correct entry speed.

Clothoid exit: During corner exit, the controller transitions from full cornering (point B in Figure 1(b)) to full accelerating (point C). A throttle-on-exit technique is used to trace along the circle by modulating the correct amount of throttle and steering.

Similar to the clothoid entry, calculating the feedforward longitudinal acceleration $a_x(s)$ starts by estimating $U_x(s)$ at the constant radius section. Equation (8) has to be integrated from the beginning of the clothoid where $s = 0$, to the end of the clothoid where $s = L^{\mathrm{cloth}}$. Similar to the clothoid entry, the initial condition at $s = 0$ is found from the constant radius section.

Note that the estimated vehicle speed $U_x(s)$ could be overestimated as the powertrain may not have sufficient torque to accelerate the vehicle out from a corner at its friction limit, see Figure 1(b). From the controller standpoint, this will simply saturate the throttle command at wide open throttle, which is the desired behaviour.

Once $U_x(s)$ is known, equation (6) is used to find the desired $a_x(s)$ along the path. Thus, from Newton's second law, feedforward longitudinal force along the path $F_x^{\mathrm{feedforward}}(s)$ can be found from $F_x^{\mathrm{feedforward}}(s) = ma_x(s)$, where m is the vehicle mass.

4.1.2 Calculating steering angle from lateral acceleration and path information

Several methods can be used to determine the feedforward steering from the estimated lateral acceleration $a_y(s)$ in equation (5). It can be calculated from a complex vehicle model with four nonlinear tyres. Alternatively, empirical data generated from a ramp steer manoeuvre can create a vehicle specific steady-state feedforward steering lookup table. For simplicity in this initial study, a basic bicycle model, shown in Figure 4, using a linear tyre assumption is used. The effects of the tyre modeling errors are then handled by the steering feedback.

By using vehicle kinematics and assuming small angles, the feedforward steering is found from:

$$\delta_{\mathrm{feedforward}} = \frac{Lr}{U_x} - \alpha_f + \alpha_r \tag{10}$$

where δ is the steering angle, L is the vehicle length in Figure 4, r is the vehicle yaw rate, α_f and α_r are the front and rear slip angles respectively. These slip angles can be found from equations of motion, using a linear tyre model.

The equations of motion are derived from a bicycle model, assuming the vehicle is cornering at a steady state.

$$F_y^f + F_y^r = ma_y = m\frac{U_x^2}{R(s)}$$
$$aF_y^f - bF_y^r = 0 \tag{11}$$

Figure 4 Bicycle model
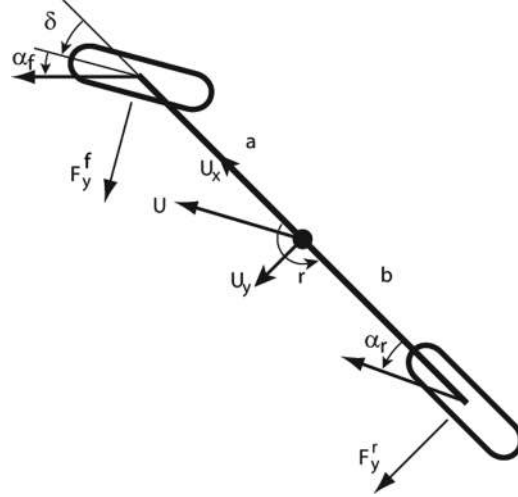
where a and b are the distances from the front and the rear axle to the vehicle's centre of gravity, F_y^f and F_y^r are the front and the rear lateral forces, and a_y is the vehicle's lateral acceleration.

To derive the feedforward steering, a linear tyre assumption is used:

$$F_y = -C\alpha \tag{12}$$

where C is the lateral axle cornering stiffness.

The distribution between front axle normal load W_f and rear axle normal load W_r can be found from:

$$W_f = \frac{b}{L}mg$$
$$W_r = \frac{a}{L}mg \tag{13}$$

substituting equation (12) and (13) into equation (11) yields the following relations:

$$-C_f\alpha_f = \frac{W_f}{g}\frac{U_x^2}{R(s)}$$
$$-C_r\alpha_r = \frac{W_r}{g}\frac{U_x^2}{R(s)}. \tag{14}$$

The slip angles from the above equations are then used in equation (10), resulting in the following expression for the feedforward steering:

$$\delta_{\text{feedforward}} = \left(L + \frac{KU_x^2}{g}\right)\frac{1}{R(s)} \tag{15}$$

where K is the vehicle understeer gradient:

$$K = \frac{W_f}{C_f} - \frac{W_r}{C_r}. \tag{16}$$

By using a clothoid curve, the curvature along the path $(1/R(s))$ in equation (15) can easily be found from equation (3).

With the feedforward steering controller designed above, the system will have some tracking error to be compensated by the feedback controller due to the linear tyre assumption. Figure 5 depicts an example of variations in the feedforward steering between using a linear tyre model and a brush tyre Fiala model (Pacejka, 2002). It is possible to incorporate a nonlinear tyre model into equation (15) to improve the accuracy of the feedforward steering, if desired.

Figure 5 Example of feedforward steering from a handling diagram, using ramp steer $U_x = 15\,\text{m/s}$ (see online version for colours)
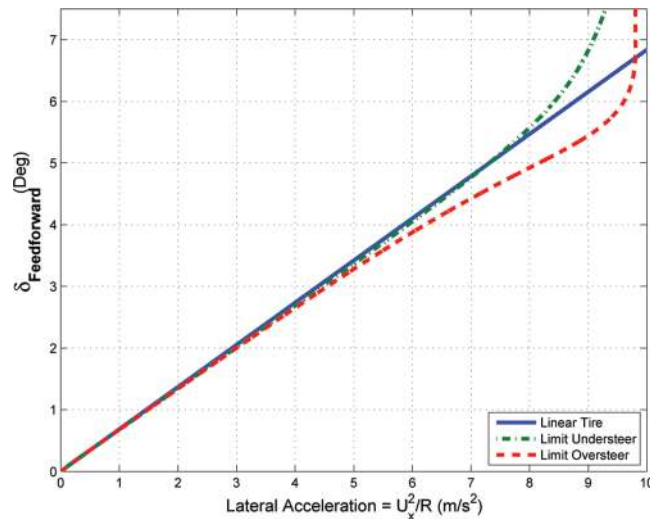
4.2   Feedback controller based on different elements

The feedback controller imitates a driver trimming the steering, throttle, and brake inputs according to the responses of the vehicle. The lanekeeping steering feedback makes adjustments to the steering to minimise tracking error and provide stability to the system, while the slip circle longitudinal feedback controller modulates the throttle and brake to regulate wheel slip. In addition, a speed feedback is used to track a desired speed and ensure that the vehicle will not overshoot the corner entry speed. To minimise yaw motion of the vehicle at the limits, yaw damping is added into the system.

4.2.1   Lanekeeping

The controller is built around a potential field lanekeeping system (Rossetter et al., 2004; Switkes et al., 2006), which provides stability, robustness and path tracking

up to the limits of tyre adhesion (Talvala et al., 2011). The lanekeeping feedback mitigates tracking errors due to possible disturbances from the testing environment or from modeling errors caused by the linear tyre assumption used in Section 4.1.2.

The lanekeeping feedback produces an additional steering angle, which is proportional to the lookahead error e_{la}, shown in Figure 6(a).

$$\delta_{\text{control}} = -\frac{2K_p}{C_f} e_{la} \tag{17}$$

where K_p is the lanekeeping potential field gain.

Figure 6 Lanekeeping system: (a) definition of heading error ($\Delta\psi$) and lookahead error (e_{la}) and (b) lanekeeping on a curved path
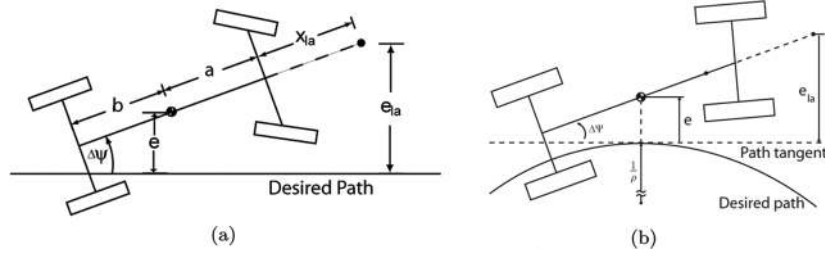
Note that the lookahead error e_{la} in Figure 6(a) is a combination of both lateral error e and heading error $\Delta\psi$:

$$e_{la} = e + (a + x_{la})\sin\Delta\psi \tag{18}$$

substituting e_{la} in equation (18) into equation (17) gives:

$$\delta_{\text{control}} = -\frac{2K_p}{C_f}(e + (a + x_{la})\sin(\Delta\psi)). \tag{19}$$

The lanekeeping potential field gain K_p is chosen to ensure stability as described in Rossetter et al. (2004) and Talvala et al. (2011), but the lookahead distance x_{la} is tuned for sustained driving at the handling limits. As x_{la} is increased, the lanekeeping system becomes more sensitive to the vehicle's heading error $\Delta\psi$ (see equation (19)). This produces a quick response to vehicle yaw motion, which improves yaw stability. However, if x_{la} is too high, the lanekeeping steering command will become too sensitive to vehicle yaw motion and could induce yaw oscillation in the system (Talvala et al., 2011). Thus, choosing x_{la} is a tradeoff between yaw stability and yaw oscillation during tyre saturation.

It is important to point out that the error signal e_{la} used in this lanekeeping system is based on the projection to the closest point on the desired path, as shown in Figure 6(b). As a consequence, the lanekeeping feedback does not depend on the future path information and does not have any inherent conflict with the feedforward steering. Such conflicts may occur in lanekeeping systems that use feedback based on future path information, such as one proposed by Tseng et al. (2005).

In addition, the steering controller limits the amount of steering to avoid exceeding a maximum lateral slip, α_{lim}, on the front axle. Since no additional lateral tyre force is generated after the lateral slip exceeds the peak slip, α^{peak}, where the maximum lateral force is generated, there is no significant benefit from additional steering when the lateral slip $\alpha > \alpha^{\text{peak}}$. In practice, α_{lim} is chosen to be slightly higher than α^{peak} to ensure that the front axle could generate maximum force and to enable the longitudinal feedback controller to detect any excessive tyre slip.

4.2.2 Yaw damping

In addition to the lanekeeping system, a small amount of yaw damping is added into the steering feedback controller to minimise yaw oscillation. At the limits of handling, the yaw dynamics become more oscillatory since the inherent damping due to the tyre forces decrease. Yaw oscillation is undesirable, as it creates unnecessary lateral tyre slip. This reduces the available longitudinal tyre force that can be used for accelerating the vehicle along the path.

To reduce yaw motion of the vehicle, a steering feedback term δ_{damping} based on $\Delta\dot{\psi} = \mathrm{d}\Delta\psi/\mathrm{d}t$ is added into the lanekeeping equation (19):

$$\delta_{\text{damping}} = -k_{\Delta\dot{\psi}}\Delta\dot{\psi} \tag{20}$$

where $k_{\Delta\dot{\psi}}$ is the yaw damping gain.

Rather than differentiate $\Delta\psi$ to obtain $\Delta\dot{\psi}$, which could amplify high frequency noise, $\Delta\dot{\psi}$ is calculated from the measured vehicle states:

$$\begin{aligned}\Delta\dot{\psi} &= \dot{\psi}_{CG} - \dot{\psi}_r \\ &= r - \frac{U_x}{R(s)}(\cos\Delta\psi - \tan\beta\sin\Delta\psi)\end{aligned} \tag{21}$$

where ψ_{CG} is the heading angle of the vehicle relative to an earth frame, ψ_r is the heading angle of the path relatives to earth frame and β is the vehicle sideslip.

The final steering angle δ is therefore:

$$\delta = \delta_{\text{feedforward}} + \delta_{\text{control}} + \delta_{\text{damping}}. \tag{22}$$

4.2.3 Slip circle

A longitudinal feedback controller based on a slip circle fulfils two purposes. First, it provides a longitudinal input that controls tyre slip. Controlling rear tyre slip enhances vehicle stability by avoiding rear axle saturation, while controlling front axle slip ensures that sufficient control authority is provided to the steering controller. For instance, when a vehicle is understeering (the front tyre is saturated), additional steering no longer provides additional lateral tyre force to the vehicle. Thus, the slip circle controller reduces the amount of longitudinal force to increase lateral force available to the steering. Secondly, the slip circle controller ensures that the tyres are operating at their limits. If the slip circle controller detects that none of the tyres is operating at their limits, it can command additional longitudinal force to utilise the friction.

Concept Before introducing the idea of combined slip, individual lateral and longitudinal tyre curves are shown in Figure 7(a). A lateral tyre curve demonstrates how lateral tyre slip α generates lateral tyre force. The initial slope of the lateral tyre curve is the lateral cornering stiffness C used in equation (12). The maximum lateral tyre force is achieved when the lateral slip $\alpha = \alpha^{\text{peak}}$, and is equal to the maximum force that a tyre could generate.

Figure 7 Relationship between tyre curves and a slip circle: (a) lateral and longitudinal tyre curves and (b) a slip circle (see online version for colours)
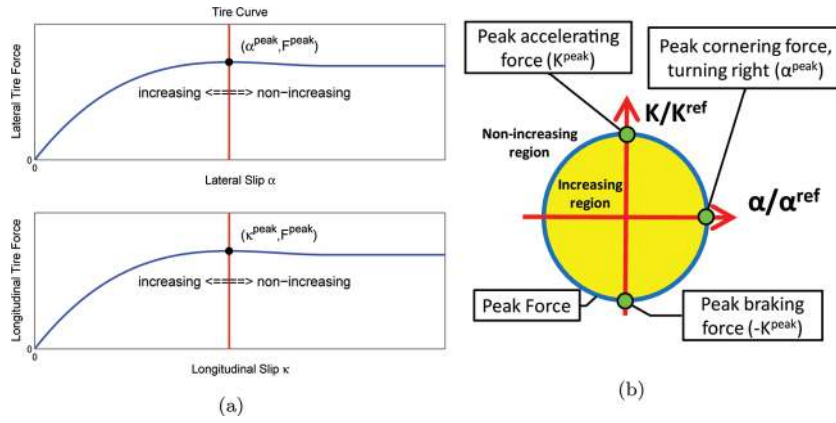
Similarly, a longitudinal tyre curve shows the relationship between the longitudinal tyre force and the longitudinal slip κ, which is defined as:

$$\kappa = \frac{R_e\omega - V}{V} \tag{23}$$

where ω is the angular velocity of a wheel, R_e is the effective radius of the tyre and V is the wheel centre's velocity along the tyre's longitudinal axis. When only κ is present, the maximum longitudinal force occurs when $\kappa = \kappa^{\text{peak}}$, which corresponds to the peak braking or accelerating.

A slip circle explains the state of the combined longitudinal and lateral tyre slip, which can be used for estimating the available longitudinal and lateral tyre force as described in Schuring et al. (1996). Alternatively, a combined slip can be used as a feedback state for the longitudinal controller. This slip represents how force is distributed between lateral and longitudinal force. By normalising lateral slip by $\alpha^{\text{ref}} = \alpha^{\text{peak}}$ and longitudinal slip by $\kappa^{\text{ref}} = \kappa^{\text{peak}}$, any point on a unit circle ($\overline{\alpha}^2 + \overline{\kappa}^2 = 1$) in Figure 7(b) produces peak tyre force, where normalised quantities $\overline{\alpha}$ and $\overline{\kappa}$ can be defined as:

$$\overline{\alpha} = \frac{\alpha}{\alpha^{\text{ref}}}$$
$$\overline{\kappa} = \frac{\kappa}{\kappa^{\text{ref}}}. \tag{24}$$

For instance, $|\overline{\alpha}| = 1$ corresponds to a tyre generating peak cornering force while $\overline{\kappa} = -1$ represents a tyre generating maximum braking.

Note that the unit circle in Figure 7(b) can be used to define the tyre force boundary between increasing and non-increasing tyre force regions. Inside this unit circle, the tyre force does not reach its limit and increasing slip increases tyre force. However, outside of this unit circle, increasing slip no longer generates additional force. In theory, a racecar driver will try to stay on these unit circles defined for each tyre, to obtain maximum tyre forces. In practice, due to weight transfer, suspension geometries, etc., it is difficult to be on the unit circle of every tyre at once. To gain a fundamental understanding of how this slip circle concept works, this paper simplifies the concept into front and rear slip circles. The left and right κ of each axle are averaged to create a slip circle of each axle.

In addition, it should be pointed out that there is a fundamental difference between the 'g-g' diagram in Figure 1(b) and the slip circle in Figure 7(b). A 'g-g' diagram represents the vehicle's combined acceleration limits. It is useful to show how close the vehicle is to the limits but does not show when a tyre passes its limit. In contrast, a slip circle can detect if a tyre is sliding beyond its peak force (outside of a slip circle) or 'gripping' (inside of a slip circle) based on the state of the slip relative to the unit circle. Thus, a slip circle is used as a feedback state while a 'g-g' diagram is used for calculating desired acceleration and speed profiles.

Implementing the slip circle controller

A racecar driver has the ability to control the amount of tyre slip to ensure vehicle stability and to operate at the peak tyre force. Similarly, the objective of the longitudinal slip circle controller is to track the edge of the unit circle in Figure 7(b). The controller should minimise excessive slip that could cause understeer or oversteer and ensure that the tyres are operating close to their peak forces.

To reach the full limits of the vehicle, one must choose α^{peak} and κ^{peak} to normalise the unit circle in equation (24). In practice, different reference values may be used depending upon the desired vehicle behaviour. Since rear axle saturation could cause vehicle instability, different levels of conservatism can be chosen to balance between operating at the peak force and minimising the chances of rear wheel saturation when applying throttle. For instance, α_f^{ref} is chosen from the peak tyre force using a ramp steer manoeuvre while α_r^{ref} is detuned from the α^{peak} to minimise rear axle sliding. Similarly, κ_f^{ref} and κ_r^{ref} in Table 1 are found experimentally from the peak tyre force in a longitudinal tyre curve.
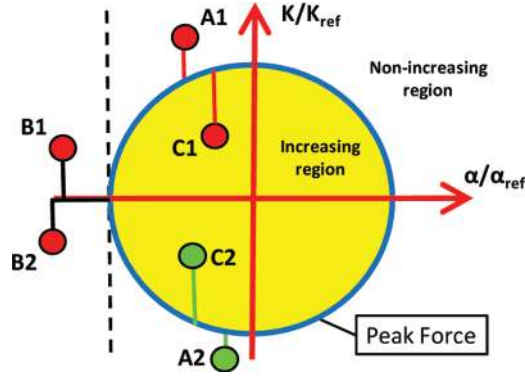
To explain how the longitudinal slip circle controller operates in various scenarios, different tyre slip conditions shown in Figure 8 are used.

Front or all tyres are saturated: When the front axle slip is outside of the unit circle, the slip circle algorithm focuses on controlling the front axle slip. If $|\bar{\alpha}| \leq 1$ (points A1 and A2 in Figure 8), the feedback longitudinal force is proportional to the distance between the slip and the edge of the unit circle (along the $\bar{\kappa}$-axis). If $|\bar{\alpha}| > 1$ (points B1 and B2), the feedback longitudinal force is proportional to the the distance along the $\bar{\alpha}$ and $\bar{\kappa}$ axes.

$$F_x^{\mathrm{feedback}} = \begin{cases} K_\kappa \Delta\bar{\kappa} + K_\alpha \Delta\bar{\alpha}, & \bar{\kappa} \leq 0 \\ -K_\kappa \Delta\bar{\kappa} - K_\alpha \Delta\bar{\alpha}, & \bar{\kappa} > 0 \end{cases} \qquad (25)$$

Table 1 Controller gains and parameters for Bonneville Salt Flats

Steering controller			
Lanekeeping	K_p	3500	N/m
Lookahead distance	x_{la}	20	m
Yaw damping	$k_{\Delta\dot\psi}$	.1	s
Longitudinal controller			
Front axle lateral slip	α_f^{ref}	7	deg
Rear axle lateral slip	α_r^{ref}	5	deg
Front axle longitudinal slip	$\kappa_{\mathrm{front}}^{\mathrm{ref}}$	.1	
Rear axle longitudinal slip	$\kappa_{\mathrm{rear}}^{\mathrm{ref}}$	.1	
Kappa slip	K_κ	3000	N
Alpha slip	K_α	2000	N
Speed	K_{speed}	2000	$N \cdot \mathrm{s/m}$

Figure 8 Different control scenarios on a slip circle (see online version for colours)

where

$$\Delta\bar\kappa = \begin{cases} |\bar\kappa| - \sqrt{1 - \bar\alpha^2}, & \bar\alpha \le 1 \\ |\bar\kappa|, & \bar\alpha > 1 \end{cases}$$

$$\Delta\bar\alpha = \begin{cases} 0, & \bar\alpha \le 1 \\ |\bar\alpha| - 1, & \bar\alpha > 1 \end{cases}$$

K_κ and K_α are the gains of the front longitudinal and lateral slip respectively.

Note that when the front axle slip is at point A2 or B2 in Figure 8, the slip circle commands positive feedback force. This is similar to a racecar driver reducing the amount of braking when the front axle is locking up. On the contrary, when the front slip is at point A1 or B1, the feedback controller will reduce the amount of throttle, mimicking a racecar driver letting off the throttle to reduce wheel spin.

Notice that with this longitudinal controller, there is no direct control over lateral slip α. Pulling the front slip back into a unit circle becomes a challenge when $|\bar\alpha| > 1$. Changing the longitudinal controller input could only have a secondary effect on vehicle speed, which then influences the lateral slip α. Only the front steering input has direct control over front lateral slip α_f.

Only rear tyres are saturated: When only the rear axle is sliding (rear slip is outside of the unit circle), equation (25) is used with parameters from the rear axle.

None of the tyres are saturated: When none of the tyres is operating at their peak forces (point C1 or C2 in Figure 8), a racecar driver can adjust the inputs to fully utilise the tyre forces. This may be emulated by allowing the controller to increase longitudinal force:

$$F_x^{\text{feedback}} = K_{\text{noSlip}} \Delta \kappa_{\text{min}}$$

$$\Delta \kappa_{\text{min}} \quad = \min \begin{cases} \sqrt{1 - (\bar{\alpha}_f)^2} - |\bar{\kappa}_f| \\ \sqrt{1 - (\bar{\alpha}_r)^2} - |\bar{\kappa}_r| \end{cases} \tag{26}$$

where K_{noSlip} is the gain when every axle is in the increasing region. Because this behaviour can conflict with setting appropriate corner entry speed, the no-slip gain is inactive during braking on a straight section.

4.2.4 Tracking of desired speed profile

As mentioned in Section 2.1.3, vehicle speed plays an important role in vehicle control, especially during corner entry. Although the vehicle speed is indirectly related to the longitudinal acceleration, the lateral acceleration directly relates to the square of the vehicle speed. By using a steady state assumption to approximate a_y in equation (5), equation (4) also depends on vehicle speed U_x. Thus, in order to control the vehicle's lateral acceleration a_y, controlling the vehicle speed U_x becomes critical.

In order to track a desired speed profile, a proportional speed feedback is used:

$$F_x^{\text{speed Feedback}} = K_{\text{speed}}(U_x(s) - U_x) \tag{27}$$

where K_{speed} is the speed proportional feedback controller gain. Desired vehicle acceleration $a_x(s)$ and desired vehicle speed $U_x(s)$ are calculated from Section 4.1.1. The speed tracking is active in every segment except during the constant radius section, where the tyre slip circle controller and feedforward drag compensation F_x^{drag} govern the longitudinal force. F_x^{drag} is added to compensate for drag from rolling resistance, aerodynamic force, inclination and the longitudinal component from the front axle turning.

The total longitudinal force F_x is the combination of the feedforward longitudinal force calculated from a 'g-g' diagram in Section 4.1.1, the drag compensation, the slip circle feedback in Section 4.2.3 and the speed feedback. Thus, the longitudinal force command F_x is therefore:

$$\begin{aligned} F_x = {} & F_x^{\text{feedforward}} + F_x^{\text{drag}} \\ & + F_x^{\text{feedback}} + F_x^{\text{speed Feedback}}. \end{aligned} \tag{28}$$

If the total longitudinal force is positive, the force is converted to throttle position via an engine map. On the other hand, if the total longitudinal force is negative, the force is converted to either throttle position that will produce engine brake using an engine map or a commanded brake pressure using a lookup table.

5 Experimental results

This research uses an Audi TTS (four-wheel drive, Figure 9) with production electronic power steering motor, active brake booster and throttle by-wire that are modified to receive commands from the controller. A lefthanded oval track at the Bonneville Salt Flats in Utah is used for testing. This surface has an approximate friction coefficient of 0.5. The vehicle is equipped with a Differential Global Positioning System (DGPS) and Inertial Navigation System (INS), from which vehicle position and other states can be obtained. The map matching algorithm (Rossetter et al. (2004)) utilises the information from the DGPS/INS system to find the vehicle's path tracking errors used in equation (19). The sampling rate of the controller is 200 Hz.

Figure 9 Autonomous Audi TTS at Bonneville Salt Flats (see online version for colours)


5.1 Overall controller performance

Figure 10 provides an overview of how the controller performs on an oval track.

Straight: $t = 67.6$–71.1 s: In this section, the vehicle transitions from full throttle to full braking in Figure 10(b), when the distance to the corner entry satisfies equation (9) at $t = 68.5$ s. During braking, the longitudinal slip κ in Figure 10(d) is negative in order to generate braking force and the longitudinal acceleration (Figure 10(g)) is approximately equal to -0.5 g.

Corner entry: $t = 71.1$–75.1 s: During this transient phase, the controller trail-brakes by gradually releasing the brake (see Figure 10(b)) while the feedforward steering increases the steering angle (see Figure 10(a)). Although the lateral slip α increases and the longitudinal slip κ decreases in Figure 10(c) and (d), the magnitude of the combined slip (slip norm $|\text{slip}| = \sqrt{\bar{\alpha}^2 + \bar{\kappa}^2}$) in Figure 10(e) is approximately constant.

In Figure 10(a), notice that the lanekeeping steering feedback commands additional steering throughout the corner. This is expected, as a linear tyre model is used in

Figure 10   Controller commands (a)–(b) and vehicle states (c)–(h), with zero no-slip gain. Feedback force in Figure (b) only shows slip circle controller command (see online version for colours)
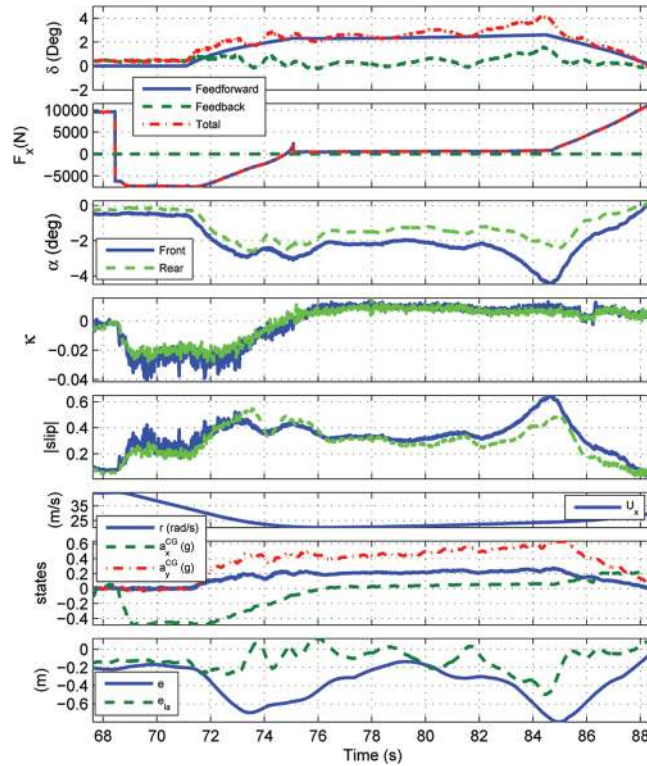
the feedforward steering calculation. Since the experimental vehicle has an inherent limit understeer characteristic, additional steering is required when cornering at high lateral acceleration as in Figure 5.

Constant radius: $t = 75.1$–84.5 s: At this point in the manoeuvre, all of the feedforward commands (Figure 10(a) and (b)) are relatively constant. The lanekeeping steering feedback in Figure 10(a) makes small corrections to minimise the lookahead error e_{la} in Figure 10(h). Note that the no-slip gain K_{noSlip} is set to zero in Figures 10–12 to investigate the effectiveness of the feedforward longitudinal controller. Since the vehicle is inside a unit slip circle and the no-slip gain is off, the slip circle feedback commands zero in Figure 10(b). The slight speed increase in Figure 10(f) is a result of overcompensation for vehicle drag.

Corner exit: $t = 84.5$–88.5 s: The controller applies throttle-on-exit. It progressively increases the throttle while unwinding the steering in Figure 10(a). A small dip in the longitudinal slip κ around $t = 86$ s in Figure 10(d) indicates clutch disengagement from the Direct-Shift Gearbox (DSG). This consequently creates a small drop in the longitudinal acceleration in Figure 10(g).

While the vehicle is driving at its predicted friction limits, it still maintains good tracking performance. The lookahead error e_{la} in Figure 10(h), which is the state

that the lanekeeping tries to track, remains within $0.5\,\text{m}$. The lateral error e, which is a by-product of controlling e_{la}, stays within $0.8\,\text{m}$ throughout the corner.

Figure 11 shows the effectiveness of the feedforward longitudinal controller when the no-slip gain is inactive. The goal of the feedforward longitudinal controller is to design a longitudinal force command (brake and throttle input) that will make the vehicle's acceleration trace the friction limit circle ($\mu g = 0.5\,\text{g}$) on a 'g-g' diagram. Figure 11 shows that the vehicle's acceleration traces the friction limit circle very well, up to the engine power limit. The dip in the clothoid exit around $a_y = 4$ and $a_x = 1\,\text{m/s}^2$ corresponds to the gearshift mentioned previously.

Figure 11 'g-g' diagram from Audi TTS, with zero no-slip gain (see online version for colours)
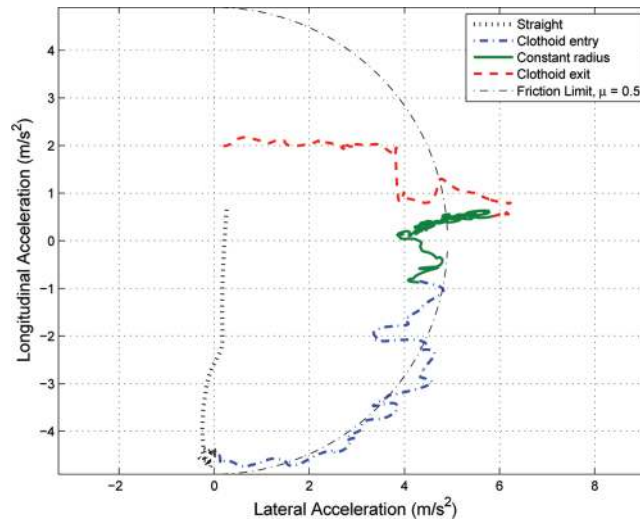

5.2 Performance with over-estimated friction

This section shows how the same controller with the same parameters performs on a different corner of the oval. Figure 12 demonstrates the robustness of the controller and how it reacts when the actual friction coefficient is lower than the predicted value ($\mu = 0.5$). Although the vehicle oversteers during 33.8–35.2 s and understeers during 44–45 s (from $|\text{slip}| > 1$), the vehicle remains stable as described below.

During the corner entry from 32.2–36.1 s, the vehicle is trail-braking into the corner. The rear slip starts to grow due to weight transfer, which reduces the normal force on the rear axle. Since the predicted friction value is too large, the slip norm of the rear axle in Figure 12(e) grows outside of the slip circle ($|\text{slip}| > 1$). The longitudinal slip circle feedback controller reduces the amount of braking force (see Figure 12(b)) to allow the rear axle to regain grip. At the same time, the lanekeeping steering controller countersteers and ensures stability of the system. The saturation of the rear tyres causes the lateral error e and lookahead error e_{la} to grow as shown in Figure 12(h), but the vehicle remains stable.

Figure 12   Controller commands (a)–(b) and vehicle states (c)–(h) when the actual friction is less than the predicted value, with zero no-slip gain. Feedback force in Figure (b) only shows slip circle controller command (see online version for colours)
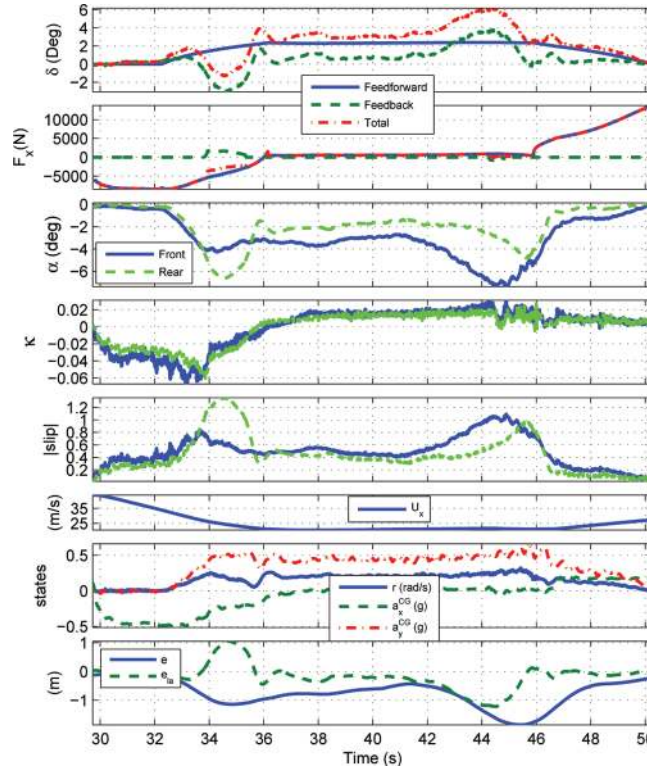
During the constant radius turn from 36.1–45.9 s, the vehicle speed slowly increases as shown in Figure 12(f). This is due to the longitudinal controller overcompensating for drag force. Although the increase in speed is minimal, the vehicle is already close to its limits and thus this excessive speed is sufficient to cause the vehicle to understeer (front slip norm in Figure 12(e) > 1). This results in increasing tracking error in both e and e_{la}, as demonstrated in Figure 12(h). Since the front axle is already saturated, any additional steering does not improve tracking performance. Thus, in this scenario, the controller reduces the amount of throttle (see Figure 12(b)) to minimise the front slip, and hence reduces the understeer and improves the tracking performance.

Note that modulating the longitudinal input only has direct control over longitudinal slip κ, i.e., along the $\bar{\kappa}$-axis in Figure 8. When heavy understeer occurs ($\bar{\alpha} > 1$), controlling longitudinal slip is not sufficient to bring the wheel slip back into a unit circle. A steering input is required to move the lateral slip α back into the unit circle. Work is in progress to coordinate steering and throttle inputs to ensure that wheel slip can be moved back into the unit circle from both the longitudinal and lateral directions.

At the corner exit during 45.9–50.2 s, since the vehicle has recovered from understeering in the constant radius section, the tracking performance improves as both e and e_{la} reduce to zero in Figure 12(h).

5.3 *Effectiveness of yaw damping*

Figure 13 shows how the controller performs without steering yaw damping. Without steering yaw damping, the system dynamics include only the intrinsic damping of the tyres, which decreases as the tyres approach their limits. This in turn creates an oscillation in the vehicle states, especially yaw motion. This unnecessary yaw motion is undesirable since it requires additional lateral tyre force, which is not necessarily available when the vehicle is operating at its limits. After adding yaw damping (in Figures 10 and 12), the oscillation in the vehicle states is significantly reduced, demonstrating the effectiveness of injecting yaw damping to compensate for tyre saturation.

Figure 13 Oscillation in vehicle states when no yaw damping is added into the system (see online version for colours)
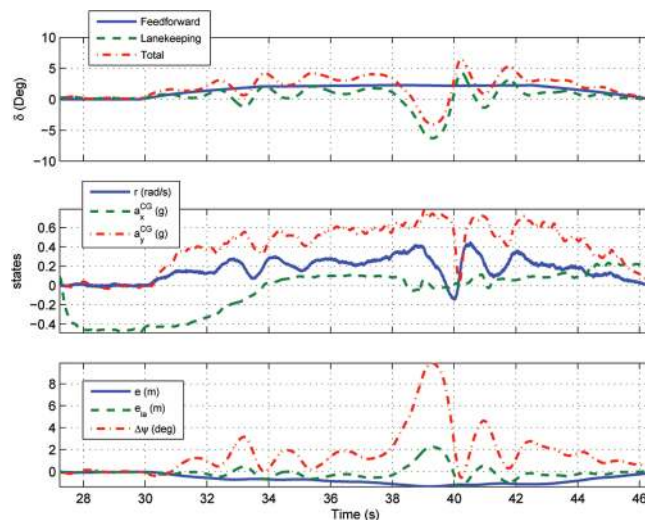
6 Conclusion

Understanding how racecar drivers control a vehicle at its limits provides many insights into designing an autonomous racing controller. The desired path is structured so that it imitates a racecar driver's cornering sequence while a clothoid path is used for corner entry and exit sections. The controller consists of individual controller modules which cover different aspects of driving, yet work concurrently to control the vehicle at its limits. A linear bicycle model and a 'g-g' diagram are used to imitate racecar drivers' mental model (pre-planning), while lanekeeping steering and a slip circle feedback are utilised to mimic drivers making corrections

to minimise tracking errors. The experimental results demonstrate the controller's ability to robustly track a desired path with a lookahead error e_{la} of less than $0.5\,\mathrm{m}$ while tracing a friction limit circle on a variable salt surface. Moreover, the lanekeeping steering provides stability to the vehicle when the controller overestimates the friction value, by providing countersteering when the vehicle oversteers. When the vehicle understeers, the slip circle controller regulates the front slip, allowing the front tyres to regain grip.

The results also highlight several areas that will be addressed in future research. Realtime friction estimation can further improve controller performance and minimise tracking error caused by overestimating friction. Furthermore, work is in progress to coordinate steering and longitudinal inputs to improve wheel slip control. Tighter path tracking is also possible with more accurate feedforward steering that takes the tyre nonlinearities into account.

The racing controller's abilities to track a path at the limits while maintaining minimal tracking error shows promise for future vehicle safety systems. This controller can be applied to drive an autonomous vehicle while ensuring stability and tracking ability even in extreme conditions. Alternatively, each controller module can be adapted to create driver assistance systems that work in conjunction with the driver, assisting the driver during emergency manoeuvres.

References

AASHTO (2004) *AASHTO Green Book: A Policy on Geometric Design of Highways and Streets*, 5th ed., American Association of State Highway and Transportation Officials, Washington D.C., USA.

Baran, I., Lehtinen, J. and Popovi, J. (2010) 'Sketching clothoid splines using shortest paths', *Computer Graphics Forum*, Vol. 29, No. 2, pp.655–664.

Bentley, R. (1998) *Speed Secrets: Professional Race Driving Techniques*, 1st ed., Number ISBN-10: 0760305188, Motorbooks.

Gerdts, M., Karrenberg, S., Mueller-Bessler, B. and Stock, G. (2009) 'Generating locally optimal trajectories for an automatically driven car', *Optimization and Engineering*, Vol. 10, No. 4, pp.439–463.

Gray, A. (1997) *Modern Differential Geometry of Curves and Surfaces with Mathematics*, 2nd ed., CRC Press, Boca Raton.

Hsu, Y-H.J., Laws, S.M. and Gerdes, J.C. (2010) 'Estimation of tire slip angle and friction limits using steering torque', *IEEE Transactions of Control Systems Technology*, Vol. 18, No. 4, pp.896–907.

Jeffrey, A. and Dai, H-H. (2008) *Handbook of Mathematical Formulas and Integrals*, 4th ed., Academic Press, Massachusetts, USA.

Klomp, M. (2010) *Longitudinal Force Distribution and Road Vehicle Handling*, PhD Thesis, Department of Applied Mechanics, Chalmers University of Technology, Gothenburg, Sweden.

Lopez, C. (2001) *Going Faster Mastering the Art of Race Driving*, Robert Bentley, Inc., 1734 Massachusetts Avenue, Cambridge, MA 02138-1804, USA.

Milliken, W.F. and Milliken, D.L. (1995) *Race Car Vehicle Dynamics*, SAE International, Pennsylvania, USA, pp.345–359.

Mitchell, W.C., Schroer, R. and Grisez, D.B. (2004) *Driving the Traction Circle*, SAE (2004-01-3545).

Mueller-Bessler, B., Henze, R. and Kucukay, F. (2008) 'Reproducible transverse dynamics vehicle evaluation in the double lane change', *ATZ Worldwide*, Vol. 110, ATZ 0412008, pp.44–49.

Mueller-Bessler, B., Stock, G. and Hoffmann, J. (2008) 'Customer oriented safety and handling evaluation via adjusted driver model using real vehicle', *ATZ Online*, F2008-12-014, pp.1–8.

Muhlmeier, M. and Muller, N. (2002) *Optimization of the Driving Line on a Race Track*, SAE Technical Paper Series (2002-01-3339), pp.1–16.

Pacejka, H.B. (2002) *Tire and Vehicle Dynamics*, Society of Automotive Engineers, Inc., 400 Commonwealth Dr. Warrendale, PA 15096-0001.

Rice, R.S. (1973) 'Measuring car-driver interaction with the g-g diagram', *Society of Automotive Engineers*, Warrendale, PA 730018, pp.1–19.

Rossetter, E.J., Switkes, J.P. and Gerdes, J.C. (2004) 'Experimental validation of the potential field lanekeeping system', *International Journal of Automotive Technology*, Vol. 5, No. 2, pp.95–108.

Schuring, D.J., Pelz, W. and Pottinger, M.G. (1996) 'A model for combined tire cornering and braking forces', SAE Technical Paper 960180, pp.61–73.

Switkes, J.P., Rossetter, E.J., Coe, I.A. and Gerdes, J.C. (2006) 'Handwheel force feedback for lanekeeping assistance: combined dynamics and stability', *Journal of Dynamic Systems, Measurement, and Control*, Vol. 128, No. 3, pp.532–542.

Talvala, K.L.R., Kritayakirana, K. and Gerdes, J.C. (2011) 'Pushing the limits: from lanekeeping to autonomous racing', *Annual Reviews in Control*, Vol. 35, No. 1, pp.137–148.

Taruffi, P. (1959) *The Technique of Motor Racing*, Robert Bentley, Inc.

Timings, J.P. and Cole, D.J. (2010) 'Minimum manoeuvre time of a nonlinear vehicle at constant forward speed using convex optimisation', *Proc 10th International Symposium on Advanced Vehicle Control*, AVEC 2010, August 2010, Loughborough, pp.1–6.

Tseng, H.E., Asgari, J., Hrovat, D., van der Jagt, P., Cherry, A. and Neads, S. (2005) 'Evasive manoeuvres with a steering robot', *Vehicle System Dynamics*, Vol. 43, No. 3, pp.199–216.

Velenis, E., Tsiotras, P. and Lu, J. (2008) 'Optimality properties and driver input parameterization for trail-braking cornering', *European Journal of Control*, Vol. 14, No. 4, pp.308–320.

Walton, D. and Meek, D. (2005) 'A controlled clothoid spline', *Computers & Graphics*, Vol. 29, pp.353–363.