

A Controller for Automated Drifting Along Complex Trajectories

Jonathan Y. Goh, T. Goel & J. Christian Gerdes

Department of Mechanical Engineering

Stanford University, California, USA

E-mail: jongoh@stanford.edu

Topics/Vehicle Dynamics and Chassis Control

The ability to operate beyond the stable handling limits is important for the overall safety and robustness of autonomous vehicles. To that end, this paper describes a controller framework for automated drifting along a complex trajectory. The controller is derived for a generic path, without assuming operation near an equilibrium point. This results in a physically insightful control law: the rotation rate of the vehicle's velocity vector is used to track the path, while yaw acceleration is used to stabilize sideslip. To accurately achieve these required state derivatives over a broad range of conditions, nonlinear model inversion is used in concert with low-level wheelspeed control. Experiments on a full-scale vehicle demonstrate excellent tracking of a trajectory with varying curvature, speed, and sideslip.

1 INTRODUCTION

Traditional vehicle control architectures typically assume independent lateral/longitudinal control, and stable sideslip dynamics. Exceeding a vehicle's handling limits, however, can lead to both strong input coupling and yaw/sideslip instability, rendering this simplified approach ineffective. Yet, drivers in professional 'drifting' events can achieve precise control over both sideslip and the vehicle's path, despite operating entirely outside the vehicle's stability limits. The development of controllers for automated drifting could extend the usable state-space beyond these limits, thereby helping ensure that the widest possible range of maneuvers is available to an autonomous vehicle.

Early examples in the literature demonstrated stabilization of the vehicle states about a single drift equilibrium, both in simulation by Velenis et al. [9], and experiment by Hindiyeh et al. [4]. Because the problem is underactuated with the standard set of inputs, steering and drive torque, tracking a path in addition to stabilizing sideslip is not straightforward. Some recent works have demonstrated this experimentally for simple, constant-curvature circles. The controller by Werling et al. [11] combined sideslip stabilization with tracking a vehicle course, while the formulation by Goh et al. [3] explicitly followed a path. However, due to restrictive assumptions in vehicle modelling or controller formulation, these approaches cannot be easily extended to more complex trajectories.

Throughout the literature, studies of drifting have used a large range of vehicle model fidelity. Two-state single-track models were used by Ono et al. [6] and Voser et al. [10] to discuss the unstable dynamics of drifting. Three-state single-track models with both forces as direct inputs [4] [3], and with explicit modelling of steering and throttle delay [11], have been used for experimentally-verified controller design. And a four wheel model with steady-state weight transfer and wheel-speed/differential dynamics was linearized by Velenis et al. [9] for a Linear Quadratic Regulator. Striking a balance between model accuracy and tractability of the equations of motion is still an open question.

Contrasting with these various approaches, this paper presents an automated drifting controller explicitly designed to handle complex trajectories. Sideslip and lateral error from a curvilinear path co-ordinate system are chosen as objectives. The desired controller dynamics are first derived without assuming a specific vehicle model, or that the vehicle is near an equilibrium point. The resulting control law, expressed in terms of vehicle state derivatives, is surprisingly simple and intuitive. It leverages



Figure 1: MARTY during an automated drifting test

the decoupling of sideslip and yaw dynamics that occurs when drifting: the rotation rate of the vehicle's velocity vector is used directly to track the path; then, by yawing the vehicle faster or slower than its velocity vector, we can simultaneously control the sideslip of the vehicle.

To realize this control law, a vehicle model is required to map these desired state derivatives to the inputs. Rather than relying on oversimplifying assumptions, nonlinear model inversion is used in conjunction with low-level wheelspeed control to achieve good accuracy over the broad range of conditions encountered in a complex trajectory. Experiments on the full-scale MARTY test vehicle (Figure 1) confirm the effectiveness of the controller on a trajectory with curvature that varies from $1/7$ to $1/20m$, speed that varies from 25 to 45km/h, and up to -40° of sideslip.

2 VEHICLE MODEL

In this section we first introduce the equations of motion for the force-based single-track model in a curvilinear co-ordinate system; tire force models are then described.

2.1 Equations of Motion

2.1.1 Path Tracking States and Dynamics

The vehicle, shown schematically in Figure 2, has three states: yaw rate r , velocity V and sideslip β . To incorporate path tracking, we introduce several additional states. The vehicle course, ϕ , is the direction of the vehicle's velocity vector measured to a fixed inertial frame χ . The dynamics of ϕ are simply:

$$\dot{\phi} = \dot{\beta} + r \quad (1)$$

The position of the vehicle is tracked using a curvilinear coordinate system relative to the reference trajectory. The lateral error

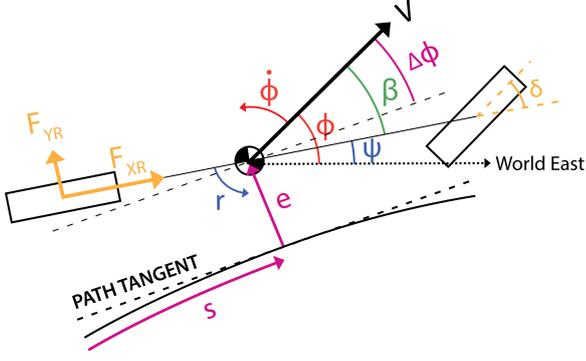


Figure 2: Three-state single-track model with reference path

e is the distance between the vehicle CG and the closest point on the reference trajectory; s is the distance along the path to this point.

The reference course ϕ_{ref} is the tangent angle of the reference path at s , measured to χ . The course error, $\Delta\phi = \phi - \phi_{ref}$, is thus the angular deviation between the vehicle course and the desired course.

The dynamics of e are:

$$\begin{aligned} \dot{e} &= V \sin \Delta\phi \\ &\approx V \Delta\phi \end{aligned} \quad (2)$$

$$\begin{aligned} \ddot{e} &\approx V \Delta\dot{\phi} + \dot{V} \Delta\phi \\ &\approx V \Delta\dot{\phi} \end{aligned} \quad (3)$$

Where we have used a few simplifying assumptions: $\sin \Delta\phi \approx \Delta\phi$ and $V \Delta\dot{\phi} \gg \dot{V} \Delta\phi$. This is a reasonable assumption as $\Delta\phi$ is driven small in closed-loop operation.

Finally, the dynamics of $\Delta\phi$ are:

$$\begin{aligned} \Delta\dot{\phi} &= \dot{\phi} - \dot{\phi}_{ref} \\ &= \dot{\phi} - \frac{\kappa_{ref} V \cos(\Delta\phi)}{1 - \kappa_{ref} e} \end{aligned} \quad (4)$$

where κ_{ref} is the curvature of the reference path at s .

2.1.2 Vehicle States and Dynamics

The nonlinear model inversion is conducted using the single-track ‘bicycle’ model as shown in Figure 2. The forces acting on the vehicle are the front lateral force F_{yf} , the rear lateral force F_{yr} , and the rear longitudinal force F_{xr} . The equations of motion are then:

$$\begin{aligned} \dot{\chi} &= \frac{aF_{yf} \cos(\delta) - bF_{yr}}{I_z} \\ \dot{\beta} &= \frac{F_{yf} \cos(\delta - \beta) + F_{yr} \cos(\beta) - F_{xr} \sin(\beta)}{mV} - r \\ \dot{V} &= \frac{-F_{yf} \sin(\delta - \beta) + F_{yr} \sin(\beta) + F_{xr} \cos(\beta)}{m} \end{aligned} \quad (5)$$

Where δ is the steering angle, a, b are the distance from the CG to the front and rear axles respectively, and m is the mass of the vehicle.

2.2 Front Tire Force Modelling

The front lateral force F_{yf} is modeled using the Fiala brush tire model [2], which was also used by the drifting controllers de-

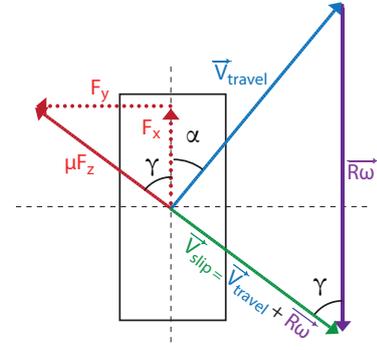


Figure 3: Velocity and force vectors for a sliding tire

scribed in [4] and [3]:

$$F_y = \begin{cases} -C_\alpha z + \frac{C_\alpha^2}{3\mu F_z} |z| - \frac{C_\alpha^3}{27\mu^2 F_z^2} z^3 & |z| < \alpha_{sl} \\ -\mu F_z \operatorname{sgn}(\alpha) & |z| \geq \alpha_{sl} \end{cases}$$

$$z = \tan \alpha$$

$$\alpha_{sl} = \arctan(3\mu F_z / C_\alpha) \quad (6)$$

Where F_z is the normal load on the tire, C_α is the cornering stiffness, α is the slip angle, and μ is the coefficient of friction.

2.3 Rear Tire Force Modelling

When drifting, the entire rear tire contact patch is in a fully sliding condition. Assuming an isotropic friction value, the lateral and longitudinal tire forces are then governed by the ‘friction circle’ relationship:

$$F_{yr} = \sqrt{(\mu F_z)^2 - F_{xr}^2} \quad (7)$$

This friction circle relationship has been used in [4], [3] to formulate controllers that map a desired rear lateral force F_{yr} to a rear longitudinal force F_{xr} , which is then treated directly as an input to the system by multiplying through the tire radius to yield a rear axle torque command.

This approach ignores the effects of wheelspeed dynamics, which are important when considering a more complex trajectory; F_{xr} is not a true input, but a result of slip between the tire and road. For a fully-saturated tire, the relationship between force and slip can be modelled in a way that is both simple and physically insightful: the magnitude of the force is μF_z , and its direction opposes the slip velocity vector V_{slip} of the contact patch. Defining this direction as the thrust angle γ , and referring to Figure 3, it can be immediately seen that changing the length of the $R\omega$ vector directly rotates γ . Geometrically, this relationship is:

$$\begin{aligned} \gamma &= \arctan(F_{yr} / F_{xr}) \\ &= \arctan\left(\frac{-V_{travel,y}}{-V_{travel,x} + R\omega}\right) \\ &= \arctan\left(\frac{-(V \sin \beta - br)}{-V \cos \beta + R\omega}\right) \end{aligned} \quad (8)$$

Where R is the radius of the tire, ω is the wheelspeed, and $V_{travel,x}$ and $V_{travel,y}$ refer to the longitudinal and lateral components of the vehicle velocity at the tire.

It is important to note that this relationship agrees with both the brush [7]¹ and simplified Pajacka magic formula tire models [1] at full saturation.

¹The brush model is described in many places, but the description of saturated tire dynamics is particularly comprehensive in this reference.

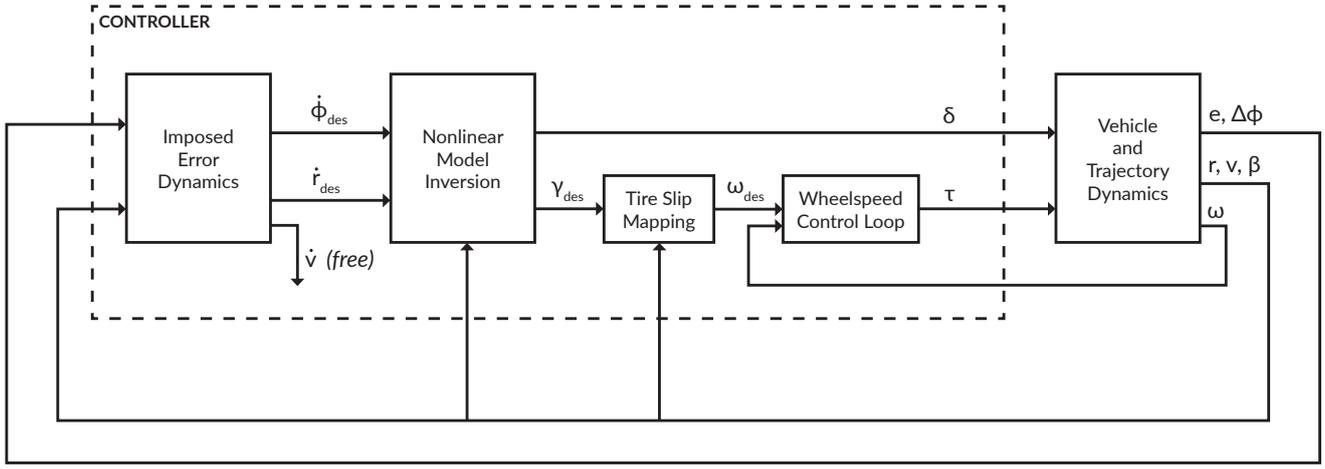


Figure 4: Block diagram of overall controller structure

The control input to the system, τ , finally appears in the model of ω dynamics:

$$I_w \dot{\omega} = \tau - R F_{xr} \quad (9)$$

Where I_w is the rotational inertia of the wheel-tire-drivetrain system, and τ is the applied torque.

3 CONTROLLER DEVELOPMENT

3.1 Overview

The task of the controller is to track the specified path and sideslip using the steering angle and rear drive torque. This motivates choosing the lateral error e and sideslip β as the control objectives.

The overall structure of the controller is shown in Figure 4. In the first part, the desired stable dynamics are imposed on e and β simultaneously, yielding a desired course rate $\dot{\phi}_{des}$ and synthetic yaw rate r_{syn} . Closing the loop around r_{syn} then yields a desired yaw acceleration, \dot{r}_{des} .

In the next step of the controller, the nonlinear vehicle dynamics model (5) is inverted to transform $\dot{\phi}_{des}$ and \dot{r}_{des} into a steering angle δ and desired thrust angle γ_{des} . The thrust angle is then mapped to a desired wheelspeed, ω_{des} . Finally, the loop is closed around ω_{des} to produce a torque command τ .

3.2 Imposed Error Dynamics

3.2.1 Path tracking

The rotation rate of the velocity vector, $\dot{\phi}$, is used directly to track the path. Desired second-order dynamics are imposed on the lateral error e :

$$\ddot{e}_{des} = V \Delta \dot{\phi}_{des} = -k_p e - k_d \dot{e} \quad (10)$$

This yields the desired value $\dot{\phi}_{des}$:

$$\begin{aligned} \dot{\phi}_{des} &= \Delta \dot{\phi}_{des} + \dot{\phi}_{ref} \\ &= \left(-\frac{k_p}{V} e - k_d \Delta \phi \right) + \dot{\phi}_{ref} \end{aligned} \quad (11)$$

3.2.2 Sideslip stabilization

Then, the yaw rate of the vehicle, relative to $\dot{\phi}_{des}$, is used to stabilize sideslip. Desired first-order dynamics are imposed on

the sideslip tracking error $e_\beta = \beta - \beta_{ref}$:

$$\dot{e}_\beta = -k_\beta e_\beta$$

$$\dot{\beta}_{des} = -k_\beta e_\beta + \dot{\beta}_{ref} \quad (12)$$

Similar to the approach in [4], [3], yaw rate is used as a synthetic input r_{syn} to yield the desired $\beta_{des} = \dot{\phi} - r$. Both of those approaches assumed a steady-state approximation for $\dot{\phi} \approx r_{eq}$. Here, however, because we are explicitly controlling $\dot{\phi}$, the desired $\dot{\phi}_{des}$ is used instead.

$$\dot{\beta} = \dot{\phi} - r$$

$$\begin{aligned} r_{syn} &= \dot{\phi}_{des} - \dot{\beta}_{des} \\ &= \dot{\phi}_{des} + k_\beta e_\beta - \dot{\beta}_{ref} \end{aligned} \quad (13)$$

First-order dynamics are then imposed on the yaw rate tracking error, $e_r = r - r_{syn}$, yielding an expression for the desired \dot{r}_{des} .

$$\dot{e}_r = -k_r e_r$$

$$\dot{r}_{des} = -k_r (r - r_{syn}) + \dot{r}_{syn} \quad (14)$$

Where \dot{r}_{syn} is the derivative of the synthetic yaw rate, and, based on our desired dynamics, can be approximated as:

$$\dot{r}_{syn} \approx (k_d^2 - k_p) \Delta \phi + e k_d k_p / V - k_\beta^2 e_\beta + \dot{r}_{ref} \quad (15)$$

Where \dot{r}_{ref} is the yaw acceleration of the reference trajectory.

3.3 Nonlinear Model Inversion

The pair $(\dot{\phi}_{des}, \dot{r}_{des})$ is inverted through the nonlinear single-track vehicle model (5) to yield the steering angle and rear longitudinal force/thrust angle commands.

A representative example of the surface of possible state derivatives is shown in Figure 5. When viewed on the \dot{r} vs. $\dot{\phi}$ axes (Figure 6), we see that this surface folds in on itself such that there is some region where there are two solutions for a given combination. This motivates splitting the surface into a larger ‘upper’ and smaller ‘lower’ surface; the splitting line is shown in red in Figures 5, 6 and 7.

Generally, only a small area of the $(\dot{\phi}, \dot{r})$ space is exclusive to the bottom surface. This motivates constraining the model inversion to the upper surface to make the mapping 1 : 1 and ensure

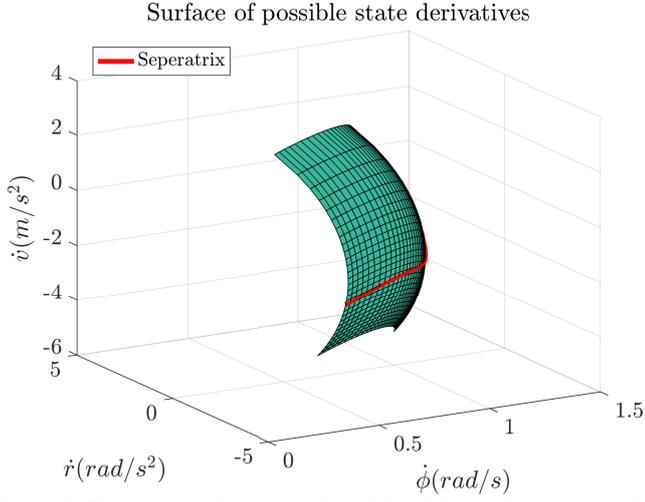


Figure 5: Representative example of feasible surface of possible state derivatives for the MARTY vehicle model at the state $r = 53.9^\circ$, $V = 9.35\text{m/s}$, $\beta = -40^\circ$

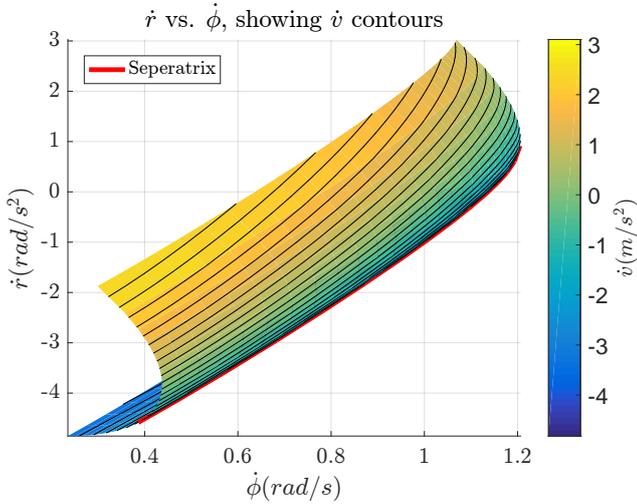


Figure 6: \dot{r} vs. $\dot{\phi}$, with \dot{v} contours

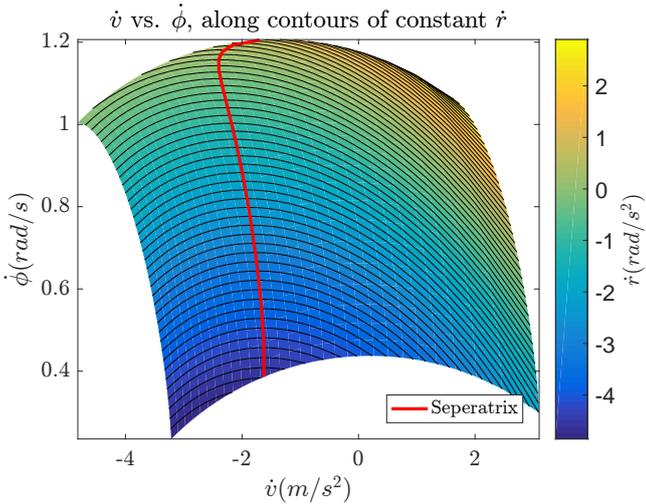


Figure 7: $\dot{\phi}$ vs. \dot{v} along contours of constant \dot{r}

continuous actuator commands in a simple manner. It is additionally worth noting that the upper surface generally contains the combinations that have $\dot{v} = 0$, $\dot{r} = 0$, as shown in Figure 7, and is at least a sufficient space of possible state derivatives for tracking a quasi-equilibrium trajectory.

It is possible for the commanded $(\dot{\phi}_{des}, \dot{r}_{des})$ pair to fall outside the feasible actuation space. For this initial presentation and testing of the controller structure, we elect to handle these cases in a naive way. The value of $\dot{\phi}_{des}$ computed in (11) is saturated

to fall within the feasible values of $\dot{\phi}$, before being used in (13) to compute \dot{r}_{des} . The resultant $(\dot{\phi}_{des}, \dot{r}_{des})$ pair is then simply projected into the feasible space along the $\dot{r} = \dot{r}_{des}$ line. This strategy roughly prioritizes the unstable yaw/sideslip dynamics over the lateral error dynamics.

3.4 Wheelspeed Control

The nonlinear model inversion step yields a rear longitudinal force/thrust angle command, which can be mapped through the force-slip relationship (8) to a desired wheelspeed; this ω_{des} is in turn achieved with acceptable bandwidth using a simple form inspired by Dynamic Surface Control [8]:

$$\begin{aligned} \tau_{command} &= -k_\omega I_\omega (\omega - \omega_{des,f}) + I_\omega \dot{\omega}_{des,f} + R F_{xrdes} \\ \dot{\omega}_{des,f} &= -(1/t_\omega) (\omega_{des,f} - \omega_{des}) \end{aligned} \quad (16)$$

Where $\omega_{des,f}$ is the desired wheelspeed after a first order filter with time constant t_ω , and we have used the desired rear longitudinal force F_{xrdes} to compute the feedforward torque.

Because the MARTY test vehicle has independently powered rear wheels that are mechanically decoupled, for experimental testing we have to use a separate control loop for each wheel. Setting $\gamma_{des,RL} = \gamma_{des,RR} = \gamma_{des}$ yields the target rear-left and rear-right wheelspeeds $\omega_{des,RL} = \omega_{des} - dr/(2R)$ and $\omega_{des,RR} = \omega_{des} + dr/(2R)$, where d is the vehicle track width. The feedforward rear longitudinal force for each wheel is computed by making a rough steady-state approximation for weight transfer:

$$\begin{aligned} \Delta F_{zrest} &= P_r h m r V \cos(\beta) / d \\ F_{xrdes,RL} &= (0.5 - \Delta F_{zrest} / F_{zr}) F_{xrdes} \\ F_{xrdes,RR} &= (0.5 + \Delta F_{zrest} / F_{zr}) F_{xrdes} \end{aligned} \quad (17)$$

Where P_r is an empirically guided guess at the proportion of lateral weight transfer that occurs over the rear axle, and h is the CG height.

4 EXPERIMENTAL VALIDATION

Experimental validation confirms the effectiveness of the controller for a complex trajectory with curvature between $1/7$ and $1/20\text{m}$, speeds between 25 and 45km/h , and reference sideslip of up to -40° .

4.1 Test Procedure

The experiments are conducted on the MARTY research testbed, depicted in Figure 1. MARTY is a modified 1981 DMC DeLorean that is equipped with steer-by-wire and an electric drivetrain that can drive the rear-left and rear-right wheels independently. Vehicle state information is provided at 250Hz by an integrated RTK-GPS/IMU. The controller is implemented on a real-time computer that also runs at 250Hz .

A simple entry and exit clothoid are appended to the start and end of the desired trajectory, and a basic path-tracking controller, similar to that presented in [5], is used to guide the vehicle on this trajectory. The drifting controller is activated at $s = 57\text{m}$, and de-activated at $s = 463\text{m}$. The controller and vehicle parameters are shown in Table 1.

4.2 Trajectory Generation

As suggested by [3], the trajectory is generated as a sequence of unstable drifting equilibrium points. First, a desired curvature κ_{ref} and sideslip β_{ref} profile as a function of path distance s is selected. Then, for each point on the trajectory, the equations of

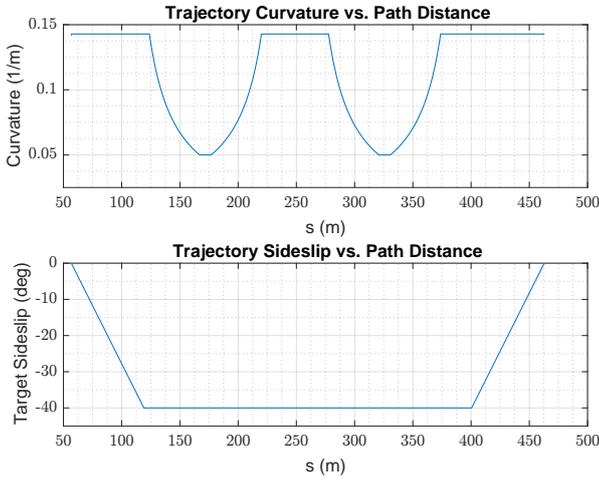


Figure 8: Selected curvature and sideslip profile vs. distance along path

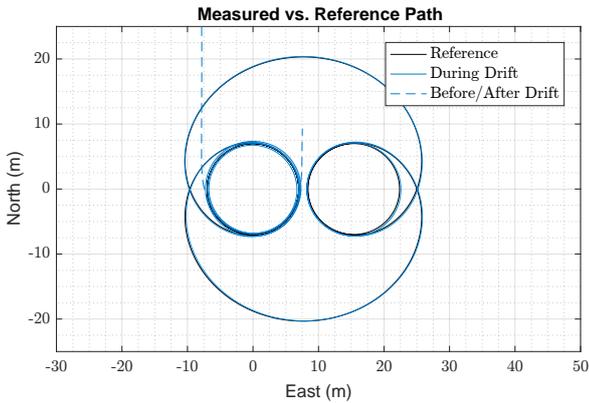


Figure 9: Measured vs. reference path for experimental run

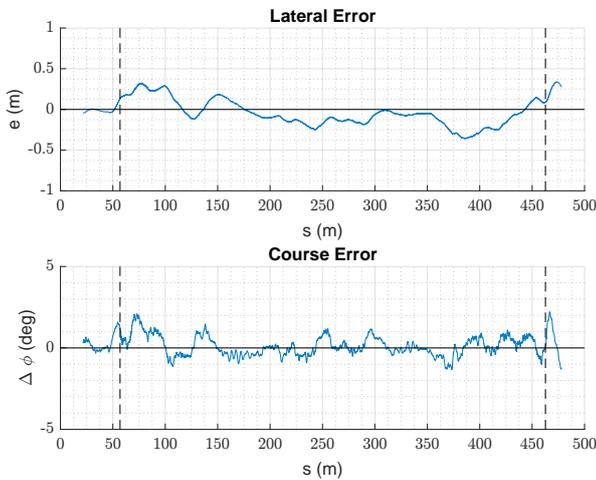


Figure 10: Path tracking performance during the experimental run

motion (5) are solved given the conditions $\beta = \beta_{ref}$, $\dot{\phi} = \kappa_{ref}V$, $\dot{V} = \dot{r} = 0$, yielding the reference trajectory values of $\dot{\phi}_{ref}$, V_{ref} . The reference sideslip rate $\dot{\beta}_{ref}$ is approximated as $\frac{d\beta}{ds}V_{ref}$.

The selected profile (Figure 8) has curvature between $1/7$ and $1/20m$, with reference sideslip of up to -40° , and results in a trajectory with equilibrium speed between 25 to $45km/h$.

4.3 Results

The controller tracks the reference trajectory well, as seen in the excellent lateral error and sideslip tracking performance.

The measured path of the vehicle is shown in Figure 9, and adheres very closely to the reference path. This qualitative observation is confirmed by the lateral error plotted in Figure 10: the

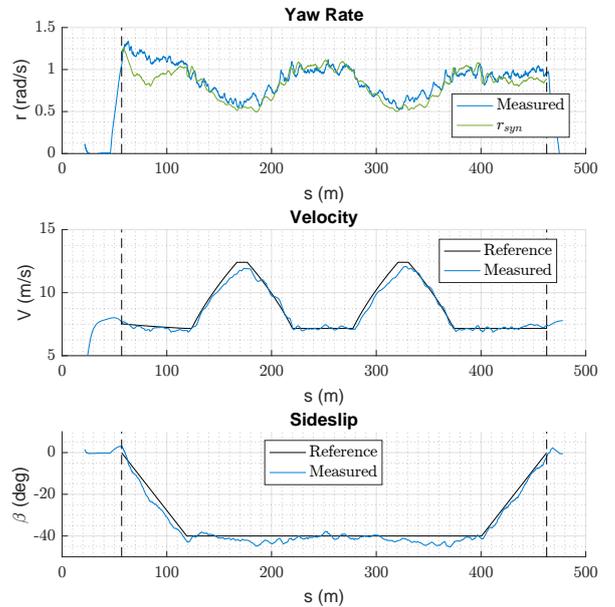


Figure 11: States vs. path distance during the experimental run

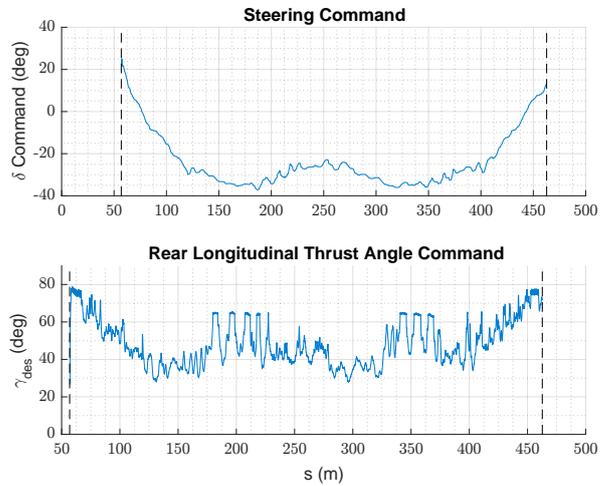


Figure 12: Steering and thrust angle inputs during the experimental run

RMS lateral error is $0.18m$, and the largest deviation is $-0.36m$.

This high level of path tracking performance occurs while the vehicle is operating near the large target sideslip of -40° . Indeed, the measured vehicle states are plotted in Figure 11, and show good tracking of the reference sideslip: the RMS sideslip tracking error is 2.4° , and the largest deviation is just -6.1° . The closed-loop velocity stays quite close to the trajectory values, indicating both that the velocity dynamics are stable, and that the model used for planning is fairly accurate. The measured yaw rate tracks the synthetic yaw rate quite well, except during the initial transition to a drifting condition.

Finally, it is worth noting that this trajectory exercises the non-linear model inversion over a broad range of states, resulting in a wide range of inputs – this is most clearly exemplified by the $\sim 65^\circ$ of steering angle the controller sweeps through during the test, as shown in Figure 6.

5 CONCLUSIONS

This paper develops a controller for autonomous drifting along a general complex trajectory. The controller is first formulated in terms of the vehicle's state derivatives without referring to a specific vehicle model. The rotation of the vehicle's velocity vector is used to track the lateral error in a curvilinear co-ordinate system. Then, the vehicle's yaw rate, relative to this velocity vector rotation rate, is used to regulate sideslip.

To map these required state derivatives to the actuator inputs,

Parameter	Symbol	Value
Mass	m	1700kg
Rotational inertia	I_z	2385kgm ²
CG to front axle	a	1.392m
CG to rear axle	b	1.008m
CG height	h	0.45m
Track width	d	1.6m
Rear weight transfer proportion	P_r	0.75
Tire radius	R	0.33m
Single wheel + drivetrain inertia	I_w	3kgm ²
Max steering angle	δ_{max}	38°
Yaw rate gain	k_r	6
Sideslip gain	k_β	2
Path tracking proportional gain	k_p	2
Path tracking damping gain	k_d	2.8

Table 1: Vehicle and controller parameters

steering and drive torque, nonlinear model inversion is used in concert with a simple wheelspeed control loop. Experimental validation on the MARTY full-scale test vehicle demonstrates good tracking of both sideslip and lateral error for a trajectory with curvature between $1/7$ and $1/20m$, speeds between 25 and 45km/h, and sideslip of up to -40° .

This successful demonstration of path-tracking with sideslip stabilization for a complex drifting trajectory points the way for future autonomous systems to be able to operate outside the open-loop stability limits if the situation requires it. Further work could explore trajectory planning and execution beyond the ‘quasi-equilibrium’ approximation, in order to handle much faster changes in the vehicle states.

6 ACKNOWLEDGEMENTS

The authors wish to thank the Revs Program at Stanford, Bridgestone Corporation, Brembo Corporation, and Renovo Motors for supporting this work, as well as our Dynamic Design Lab colleagues for assisting during testing at Thunderhill Raceway Park.

REFERENCES

- [1] Egbert Bakker, Lars Nyborg, and Hans B Pacejka. *Tyre modelling for use in vehicle dynamics studies*. Tech. rep. SAE Technical Paper, 1987.
- [2] E Fiala. “Seitenkrafte am rollenden Luftreifen”. In: *ZVD I 96* (1954), Nr-29.
- [3] Jonathan Y Goh and J Christian Gerdes. “Simultaneous stabilization and tracking of basic automobile drifting trajectories”. In: *Intelligent Vehicles Symposium (IV), 2016 IEEE*. IEEE. 2016, pp. 597–602.
- [4] Rami Y. Hindiyeh and J. Christian Gerdes. “A Controller Framework for Autonomous Drifting: Design, Stability, and Experimental Validation”. In: *Journal of Dynamic Systems, Measurement, and Control* 136.5 (July 2014), p. 051015.
- [5] Krisada Kritayakirana and J. Christian Gerdes. “Autonomous Vehicle Control at the Limits of Handling”. In: *International Journal of Vehicle Autonomous Systems* 10.4 (2012), pp. 271–296.
- [6] E. Ono, S. Hosoe, H.D. Tuan, and S. Doi. “Bifurcation in Vehicle Dynamics and Robust Front Wheel Steering Control”. In: *Control Systems Technology, IEEE Transactions on* 6.3 (May 1998), pp. 412–420.
- [7] Jacob Svendenius. “Tire Modeling and Friction Estimation”. PhD thesis. Lund University, 2007.
- [8] D Swaroop, J Karl Hedrick, Patrick P Yip, and J Christian Gerdes. “Dynamic surface control for a class of nonlinear systems”. In: *IEEE transactions on automatic control* 45.10 (2000), pp. 1893–1899.

- [9] Efsthathios Velenis, Diomidis Katzourakis, Emilio Frazzoli, Panagiotis Tsiotras, and Riender Happee. “Steady-state drifting stabilization of RWD vehicles”. In: *Control Engineering Practice* 19.11 (2011), pp. 1363–1376.
- [10] Christoph Voser, Rami Y. Hindiyeh, and J. Christian Gerdes. “Analysis and Control of High Sideslip Manoeuvres”. In: *Vehicle System Dynamics* 48.sup1 (2010), pp. 317–336.
- [11] Moritz Werling, Philipp Reinisch, and Lutz Gröll. “Robust power-slide control for a production vehicle”. In: *International Journal of Vehicle Autonomous Systems* 13.1 (2015), pp. 27–42.