Need for Speed Tokyo Drift: Hotwheels Edition (A Drifting Control Case Study)

Eric Wong and Frederick Chen

Abstract

Providing safety guarantees for automotive system controllers operating is an important and difficult problem across a wide range of industries. However, most approaches to automotive controllers don't have formal verification, ignore adverse conditions, or both. A particularly dangerous yet common adverse driving condition is the presence of rain or snow, resulting in a loss of traction with the ground. As an initial step in studying systems that operate after losing traction, we model and analyze the specific motion of drifting, the intentional loss of traction between the tires and the ground, and design controllers that choose a safe time to drift at a rate that reaches a desired degree of turn. Using Keymaera, a theorem prover for differential dynamic logic ($d\mathcal{L}$), we provide a formal proof of the safety guarantees afforded by our controllers of not drifting off the road and facing the desired direction.

1 Introduction

Safety guarantees form a critical component of cyber-physical systems, systems characterized by interleaved interactions between computational systems and and the physical world. These systems can range from large-scale systems like power grids to individual biomedical instruments, where a single error has the potential to shut down electric power for millions of homes or a faulty medical instrument can result in deathly surgeries followed by lawsuits. One of the most classical examples of these systems is an automotive system, which can consist of many sub-systems at various granularities. Cyber-physical systems can model specific parts of automotives, such as an automotive engine control unit for control of fuel injection, ignition timing, and vale retraction [Jensen et al., 2011], but also model general movements of the automotive to make safe and efficient driving decisions that to avoid collisions in various scenarios such as approaching road intersections, making left turns, and planning routes to avoid inevitable collisions [Loos and Platzer, 2011, Aréchiga et al., 2012, Chan et al., 2008]. Some of these systems, such as adaptive cruise control, are already present in the market.

While work has been done in designing controllers for various driving actions and techniques, there is an unanswered question: how do these controllers operate under sub-optimal conditions? Under hazardous conditions like rain or snow, an automotive can more easily lose traction with the ground resulting in a loss of control, and the direction of travel may no longer match the direction that the car is facing. Answering this question is critical not just for self-driving cars but for vehicles worldwide in order to maintain safety across different environments with varying climates and different road compositions, especially since people still cause collisions frequently on icy or snow-covered roads.

Furthermore, one of the major challenges to designing cyber-physical systems is that fundamental issues may have a low rate of occurrence, and consequently may not be discovered during

development with limited testing ability via simulations or empirical studies. It is practically impossible to test every possible configuration or environment for complex cyber-physical systems. While tests can give reasonable confidence in the safety of the system, there may be a missed edge case where the system is not safe. On the other hand, formal verification of systems is an area of growing interest, where safety and efficiency properties can be formally proved for models of cyber-physical systems. These types of guarantees are especially preferable over simulations and empirical studies, since a formal proof guarantees that the desired properties hold for a model over all states satisfying the initial conditions.

To this end, this paper takes an initial step in these directions and analyzes movement with a loss of traction in the case study of controlled drifting. Drifting is a popular racecar driving technique where drivers intentionally oversteer and lose traction to turn sharp corners at faster speeds. While this case is more benign than the case of poor weather conditions, it provides a fundamental first step in understanding and modeling movement with a loss of traction. This paper presents a model-based controller for a drifting vehicle, and proves various guarantees on the degree of turn while avoiding drifting off the road.

The structure of this paper is as follows: in the proceeding section we discuss related research concerning control of automobiles in various situations and the specific drifting action. Section 3 outlines formally the drifting problem and the desired safety and efficiency properties. Section 4 presents the drifting model, building up from the most basic version based on the challenges encountered, along with explanations on how each property was formally proved. Section 4.5 details some additional information requested in the report, including learned lessons, summary of deliverables, and work performed by each partner. Finally, Section 5 concludes the paper and outlines possible future work.

2 Related Work

Automated vehicle control systems that help the driver avoid accidents, or limit the damage in case of an accident, have become ubiquitous in modern cars. There is an increasing amount of work in this area attempting to test and verify the safety of automotive control software.

A large number of systems have been proposed and studied to varying degrees. One of the earliest works in verification of vehicles applied nonlinear simulations and road tests to verify that their proposed robust control system prevented skidding [Ackermann, 1997], while other researchers studied electronic stability control, a system that assists drivers in keeping on their intended path [Liebemann et al.]. Simulations remain common in later work, and are used to test and verify robustness of systems such as an integrated motion control system with various stability controls in electric vehicles and other systems [Kawashima et al., 2009, Hac and Bedner, 2007]. These simulations were carried out using software such as CarSim7.1.1 and MATLAB.

Meanwhile, drifting has been analyzed for small, remote-control cars [Ellefsen, 2012] and a controller for autonomous drifting was developed [Jakobsen, 2011] for model cars. This model was tested in a simulation environment as well, and while the results of their controller are largely in agreement with logged test data, the authors admit there are cases resulting in undesirable behavior.

Research has also been done on the measurements of properties such as the yaw rate and vehicle slip angle, which are directly related to skidding, drifting, and rollovers. Daily and Bevly [2004] proposed the use of a global positioning system (GPS) to determine the vehicle slip angle regardless of the car model, and future work reduced the computational complexity for implementation in embedded hardware [Grip et al., 2009].

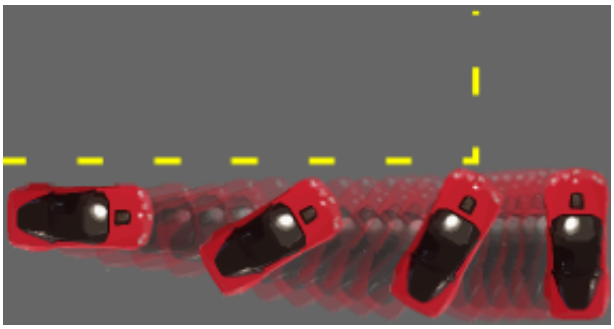

Figure 1: A graphic of a car drifting to a stop and turning 90 degrees

On the side of formal verification of vehicle control systems, a few areas have been studied. In particular, controllers made for adaptive cruise control and safe crossing of intersections have been proposed and with their safety properties formally verified [Aréchiga et al., 2012, Mitsch et al., 2012, Loos and Platzer, 2011].

Our work differs from the existing works due to the fact that we are modeling the drifting motion of a vehicle as opposed to preventing it. Our verification of safety also follows a formal, mechanized proof calculus, which verifies safety while avoiding the downfalls of simulations. To our knowledge, none of the existing works in formal verification consider a loss in traction, and are more concerned with collisions with other cars in various circumstances rather than the modeling more complicated motion. This combination of formal verification with tractionless motion, to the best of our knowledge, has not yet been addressed.

3 The Drifting Problem

The drifting problem can be defined as follows: Given a car and a road that turns to a new direction, can we create a controller that drifts the car in a direction consistent with the bend in the road?

This immediately brings up two desirable properties for a controller: first, the car should not drive off the road in order to remain safe, and second, after drifting, the car should end up pointing in the new direction consistent with the bend in the road, as this is the purpose of drifting. Since the case where the road bends right is entirely symmetrical to the case where the road bends left, we focus on the case where the road bends left.

Formally, let (x, y) be the coordinates of a car travelling parallel to the x axis approaching a turn in a road, with initial coordinates (x_0, y_0). Let (x_r, y_r) be the coordinates of the ending point of a straight segment of road, where $y_r = y_0$, and let (θ_l, θ_u) be the range of desired turning angles. Then, we can formulate safety as

$$y = y_0 \rightarrow x \leq x_r \tag{1}$$

which encapsulates the condition that if the car is still on the straight road, the car should not have gone past x_r in order to stay safe.

To formulate the second condition, note that if we represent the direction of the car as a unit vector (dx, dy), then the x-coordinate of the directional unit vector completely encapsulates the degree of turn via the relations $\cos(\theta_u) = dx_u$ and $\cos(\theta_l) = dx_l$. We can therefore reformulate the range of desired turn to be a unit vector with x-coordinate between (dx_l, dx_u), and formally express this as

$$\texttt{post drift} \rightarrow dx_l \leq dx \wedge dx \leq dx_u \tag{2}$$

where `post drift` is a condition that encapsulates the end of drifting. For our purposes, since our controller drifts to a stop, this condition is simply when the velocity is zero.

In addition, the car should not spin more than 180 degrees in order for the controller to be efficient. Therefore we also require that

$$dy \geq 0 \tag{3}$$

4 Drifting Control

The main decision for the drifting controller is to decide when and how fast should the car turn. For the most basic model of drifting, we assume a friction coefficient of F and an acceleration of A for the car, and we formally prove a controller that drifts between 90 to 180 degrees. In the second model, we impose the safety condition to not run off the road, and in the third model, we generalize this to to an arbitrary interval of turn under certain constraints. Lastly, we explore the idea of generalizing this to a top-down 2D perspective.

We make several assumptions in this study. To model the behavior of drifting, we assume the controller can control the angular velocity of the turn, and we evolve the unit direction vector around the unit circle with this angular velocity. While the vehicle itself may not be rotating around the center of the vehicle, when drifting, it is still rotating around some point relative to the vehicle, typically towards the front of the vehicle. Thus, we can represent the vehicle using this "drifting center" as the reference point, and choose a large enough buffer radius to encompass the vehicle. We also assume a fixed rate of angular velocity, because a driver can control how fast the car turns, turning the wheel sharply to roughly mimic the discontinuity in angular velocity. Since the direction vector is moving around the unit circle, we can use standard differential equations for circular motion to model movement around the unit circle.

We formally verify all presented models using KeYmaera, a hybrid systems verification tool for hybrid systems.

4.1 Basic Linear Model

In the basic model, we assume the car is already in motion and is immediately deciding how much to turn. The goal of this controller was to prove that it is possible to guarantee a turn of between 90 and 180 degrees. These bounds correspond to the resulting dx component of the unit direction vector being in the interval $(-1, 0)$, so $dx_l = -1, dx_u = 0$.

The explanation for this model is that the controller can decide the rate of turn, and then friction is the only decelerating force since we assume the car has lost traction with the ground. The dynamic part evolves the car in the same direction it started with taking into account the new force of friction, and evolves the direction vector according to the rate of turn.

Note that turning the car corresponds to moving (dx, dy) around the unit circle, and so we can reduce the problem of of deciding how much to drift to deciding the velocity at which (dx, dy) goes around the unit circle, such that after drifting, $-1 \leq dx \leq 0$. The lower bound on dx can easily be obtained by under-approximating the distance travelled on the semi-circle with the diameter of the circle, since choosing a velocity that travels at most diameter of the circle certainly won't pass the opposite side of the circle.

The main challenge here is to pick a velocity that provably satisfies the upper bound on dx, which corresponds to a lower bound on the amount of turn. To overcome this challenge, we note that $dx = \cos(av \cdot t)$ and $dy = \sin(av \cdot t)$ and that the Taylor series approximations for these functions provide alternating upper and lower bounds of dx, dy. More specifically, using up to the

Algorithm 1 Basic Linear Model (BLM)

$$
\begin{aligned}
BLM \equiv & ctrl; dyn && (4)\\
ctrl \equiv & av := 2F/vel; && (5)\\
& t := 0; && (6)\\
& acc := -F && (7)\\
dyn \equiv & (pos' = vel, vel' = acc, t' = 1 && (8)\\
& dx' = av \cdot ddx, dy' = av \cdot ddy, && (9)\\
& ddx' = -ddy \cdot av, ddy' = ddx \cdot av \ \&\ vel \geq 0) && (10)
\end{aligned}
$$

4th order terms, we have:

$$
\begin{aligned}
& dx \leq 1\\
& dy \leq (av \cdot t)\\
& dx \geq 1 - \frac{(av \cdot t)^2}{2}\\
& dy \geq (av \cdot t) - \frac{(av \cdot t)^3}{6}\\
& dx \leq 1 - \frac{(av \cdot t)^2}{2} + \frac{(av \cdot t)^4}{24}
\end{aligned}
\tag{11}
$$

These equations provide both upper and lower bounds for both dx and dy. To use these in our proof, we would like these bounds to be differential invariants. This results in the following lemma:

Lemma 1. *The equations in equation (11) can be proven as differential invariants for the Basic Linear Model.*

Proof. The inital equation $dx \leq 1$ follows from (dx, dy) being a point on the unit circle. Furthermore, $ddx = -dy$ and $ddy = dx$ can also be proven using the standard differential invariant proof techniques on the unit circle. The remaining bounds are proven inductively as differential invariants using the (DI) proof rule. Since the step is the same in each one, we present a brief proof of one such step as an example. Define:

$$\alpha \equiv dx \leq 1 - \frac{(av \cdot t)^2}{2} + \frac{(av \cdot t)^4}{24}$$

and

$$\beta \equiv dy \geq (av \cdot t) - \frac{(av \cdot t)^3}{6}$$

Furthermore, let $\gamma \equiv (dx, dy, ddx, ddy)$ and let $\theta \equiv (av \cdot ddx, av \cdot ddy, -ddy \cdot av, ddx \cdot av)$ so that

$$(\gamma' = \theta) \equiv (dx' = av \cdot ddx, dy' = av \cdot ddy, ddx' = -ddy \cdot av, ddy' = ddx \cdot av)$$

are the differential equations that evolve (dx, dy) in circular motion with angular velocity av.

$$
\dfrac{
\dfrac{
\dfrac{
\dfrac{
\dfrac{
\dfrac{
\dfrac{*}{\beta, ddx = -dy, ddy = dx \vdash -dy \leq -av \cdot t + \frac{av^3 \cdot t^3}{6}}(QE)
}{\beta, ddx = -dy, ddy = dx \vdash ddx \leq -av \cdot t + \frac{av^3 \cdot t^3}{6}}(ApplyEqn)
}{\beta, ddx = -dy, ddy = dx \vdash av \cdot ddx \leq -av^2 \cdot t + \frac{av^4 \cdot t^3}{6}}(Arithmetic)
}{\beta \wedge ddx = -dy \wedge ddy = dx \vdash av \cdot ddx \leq -av^2 \cdot t + \frac{av^4 \cdot t^3}{6}}(\wedge_l, \wedge_l)
}{\beta \wedge ddx = -dy \wedge ddy = dx \vdash \left(dx' \leq -av^2 \cdot t + \frac{av^4 \cdot t^3}{6}\right)^{\theta}_{\gamma'}}
}{\beta \wedge ddx = -dy \wedge ddy = dx \vdash (\alpha)'^{\theta}_{\gamma'}}(DI)
}{\alpha \vdash [(\gamma \ \& \ \beta \wedge ddx = -dy \wedge ddy = dx)]\alpha}
$$

Note that after applying the differential invariant to both sides and applying $ddx = -dy$, we get $\beta \vdash \beta$, the previous Taylor series bound. This is the proof structure that allows the above equations to be used as differential invariants. Applying the same steps starting from the lower order terms and building up completes the proof. □

With these differential invariants, the reasoning behind the choice of $av = 2F/vel$ is clear. Since F is the friction coefficient, we set $acc := -F$, and the car will stop after vel_0/F time, where vel_0 is the value of vel before the differential evolution. Then, when the car stops, we can plug in $t = vel_0/F$ and $av = 2F/vel_0$ into the Taylor series bounds to get

$$
\begin{aligned}
dy &\leq 2 \\
dx &\geq 1 - \frac{2^2}{2} = -1 \\
dy &\geq 2 - \frac{2^3}{6} = \frac{2}{3} \\
dx &\leq 1 - \frac{2^2}{2} + \frac{2^4}{24} = -\frac{1}{3}
\end{aligned}
\tag{12}
$$

It immediately follows that we have the desired conditions $dy \geq 0, -1 \leq dx, dx \leq 0$. Therefore the Basic Linear Model satisfies the drifting conditions outlined in Equation (2) and Equation (3).

4.2 Linear Model with an Obstacle

In the previous model, there was no concept of safety from staying on the road. The goal of the second model is to avoid drifting off the road, and we place a wall at the x coordinate x_r, in line with the condition presented in Equation (1).

The model must be modified to allow for acceleration or movement prior to drifting, since the car could start far away from the obstacle. We thus add another differential equation to our model for this purpose, with a time-triggered controller that decides when to start drifting. Since the edge of the road and the location of the car are represented by a point, we can over approximate their sizes with some radius and use these radii to form a buffer variable, which is the distance that guarantees safety between the car and the end of the road. The controller here decides whether or not to accelerate, based on whether acceleration can allow the car to safely drift to a stop before before reaching the end of the road.

The proof technique used for this model is the same as the previous, except that we must prove that $\frac{vel^2}{2F} + buffer < ox - x$ is a loop invariant in order to prove that the car controller is safe. This is sufficient for safety because $\frac{vel^2}{2F}$ is the stopping distance for a deceleration of F, which is

Algorithm 2 Linear Model with an Obstacle (LMO)

$$LMO \equiv (ctrl_1; dyn)^*; ctrl_2; dyn \tag{13}$$

$$ctrl_1 \equiv ?(vel \cdot T + \frac{1}{2} \cdot A \cdot T^2 + \frac{(A \cdot T + vel)^2}{2F} + buffer); \tag{14}$$

$$t := 0; \tag{15}$$

$$av := 0; \tag{16}$$

$$acc := A \tag{17}$$

$$ctrl_2 \equiv av := 2F/vel; \tag{18}$$

$$t := 0; \tag{19}$$

$$acc := -F \tag{20}$$

$$dyn_2 \equiv (pos' = vel, vel' = acc, t' = 1 \tag{21}$$

$$dx' = av \cdot ddx, dy' = av \cdot ddy, \tag{22}$$

$$ddx' = -ddy \cdot av, ddy' = ddx \cdot av \ \&\ vel \geq 0) \tag{23}$$

the stopping distance for a car the drifts with a force of friction F. With this invariant, the rest of the proof proceeds in a similar fashion to the previous model.

4.3 Arbitrary Turn Model

While the previous model is excellent for city-layouts with right-angle turns in the road, what if the road turns less than 90 degrees, or requires a more exact turn? In this section, we generalize the previous model to turning between an arbitrary range of acceptable turn. Formally, we relax the previous assumptions on the values of (dx_l, dx_u) from $(-1, 0)$ to an arbitrary interval, and we verify that subject to a constraint on the size of this interval, the controller can choose a safe angular velocity.

The major challenge here was determining how to pick an angular velocity for a general interval while maintaining provability. We approached this problem by noting that if the lower bound on dx is greater than dx_l, and if the upper bound on dx is less than dx_u, then dx is certainly within the range (dx_l, dx_u). This gives rise to the following equations:

$$\begin{aligned} 1 - \frac{(av \cdot t)^2}{2} &\geq dx_l \\ 1 - \frac{(av \cdot t)^2}{2} + \frac{(av \cdot t)^4}{24} &\leq dx_u \end{aligned} \tag{35}$$

The hope is that if these bounds are satisfied, then for some choice of av, we would like $dx_l \leq dx \leq dx_u$. However, in general it may not always be the case that such an av exists. Therefore, we prove the following lemma:

Lemma 2. *Suppose* $1 + 4 \cdot dx_l + dx_l^2 \leq 6 \cdot dx_u$. *Then, there is a choice of* av *such that Equation (35) is satisfied.*

Algorithm 3 Abitrary Turn Model (ATM)

$$ATM \equiv (ctrl_1; dyn)^*; ctrl_2; dyn \tag{24}$$

$$ctrl_1 \equiv ?(vel \cdot T + \frac{1}{2} \cdot A \cdot T^2 + \frac{(A \cdot T + vel)^2}{2F} + buffer); \tag{25}$$

$$t := 0; \tag{26}$$

$$av := 0; \tag{27}$$

$$acc := A \tag{28}$$

$$ctrl_2 \equiv av := (2 \cdot (1 - xlow))^{1/2} \cdot (F/vel); \tag{29}$$

$$t := 0; \tag{30}$$

$$acc := -F \tag{31}$$

$$dyn_2 \equiv (pos' = vel, vel' = acc, t' = 1 \tag{32}$$

$$dx' = av \cdot ddx, dy' = av \cdot ddy, \tag{33}$$

$$ddx' = -ddy \cdot av, ddy' = ddx \cdot av \ \&\ vel \geq 0) \tag{34}$$

Proof. Solving the first inequality in Equation (35) results in:

$$\begin{aligned} 1 - \frac{(av \cdot t)^2}{2} &\geq dx_l \\ 2(1 - dx_l) &\geq (av \cdot t)^2 \\ \frac{1}{t^2}(2(1 - dx_l)) &\geq av^2 \end{aligned}$$

Solving the second inequality in Equation (35) and applying the quadratic formula with respect to $(av \cdot t)^2$, we get

$$\begin{aligned} 1 - \frac{(av \cdot t)^2}{2} + \frac{(av \cdot t)^4}{24} &\leq dx_u \\ (1 - dx_u) - \frac{(av \cdot t)^2}{2} + \frac{(av \cdot t)^4}{24} &\leq 0 \\ \frac{\frac{1}{2} - (\frac{1}{2^2} - 4(1 - dx_u)/24)^{\frac{1}{2}}}{2/24} \leq (av \cdot t)^2 &\leq \frac{\frac{1}{2} + (\frac{1}{2^2} - 4(1 - dx_u)/24)^{\frac{1}{2}}}{2/24} \\ 6 - 12\left(\frac{1}{12} + \frac{dx_u}{6}\right)^{\frac{1}{2}} \leq (av \cdot t)^2 &\leq 6 + 12\left(\frac{1}{12} + \frac{dx_u}{6}\right)^{\frac{1}{2}} \\ \frac{1}{t^2}(6 - (12 + 24 \cdot dx_u)^{1/2}) \leq av^2 &\leq \frac{1}{t^2}(6 + (12 + 24 \cdot dx_u)^{1/2}) \end{aligned}$$

Thus, so long as av^2 satisfies all three of the resulting inequalities, then it satisfies Equation (35). Note further that since $dx_l \leq dx_u$, we have

$$2(1 - dx_l) \leq 6 + (12 + 24 \cdot dx_l)^{1/2} \leq 6 + (12 + 24 \cdot dx_u)^{1/2}$$

and thus we can reduce these three inequalities to the following condition:

$$\frac{1}{t^2}(6 - (12 + 24 \cdot dx_u)^{1/2}) \leq av^2 \leq \frac{1}{t^2}(2(1 - dx_l)) \tag{36}$$

This interval will be non empty if and only if:

$$\begin{aligned} 6 - (12 + 24 \cdot dx_u)^{1/2} &\leq 2(1 - dx_l) \\ (12 + 24 \cdot dx_u)^{1/2} &\geq 4 + 2dx_l \\ 12 + 24 \cdot dx_u &\geq 16 + 16dx_l + 4dx_l^2 \\ 6 \cdot dx_u &\geq 1 + 4dx_l + dx_l^2 \end{aligned}$$

which concludes the proof. □

Thus, by requiring that the target interval satisfies this inequality, there must exists an av that satisfies Equation 35. Looking at Equation 36, we can simply choose $av = \frac{1}{t}(2(1 - dx_l))^{\frac{1}{2}}$, where $t = \frac{vel}{F}$ as before.

The formal verification for this controller proceeds in a similar fashion to the previous models. After proving the Taylor series bounds and substituting in the value for av, and using the precondition $1 + 4dx_l + dx_l^2 \leq 6dx_u$, the following is provable via quantifier elimination:

$$\begin{aligned} 1 - \frac{2 - 2dx_l}{2} \leq dx &\leq 1 - \frac{(2 - 2dx_l)}{2} + \frac{(2 - 2dx_l)^2}{24} \\ dx_l \leq dx &\leq dx_l + \frac{4 - 8dx_l + 4dx_l^2}{24} \\ dx_l \leq dx &\leq \frac{1 + 4dx_l + dx_l^2}{6} \\ dx_l \leq dx &\leq dx_u \end{aligned}$$

and thus the controller is pointing in the desired range when it stops.

An extremely useful proof technique here was to abbreviate our choice for av to speed up quantifier elimination while adding the Taylor series approximations as differential invariants. By abbreviating av and using the fact that $av > 0$, we can hide the actual definition of av and avoid having to do reasoning over square roots, until we absolutely have to use the definition of av at the end of the proof. Without hiding the square roots, quantifier elimination tended to take extremely long amounts of time.

One would hope that by increasing the order of the Taylor series approximations, we can reduce the constraint on the size of (dx_l, dx_u). Theoretically, this makes sense, as higher order approximations will cause the upper and lower bounds to get closer to the true value and reduce the size of the interval; however it is unfortunately not practical. The approach outlined above requires solving the Taylor series bound for $(av \cdot t)^2$. As the order of the approximation increases, the order of the polynomial to be solved also increases, and there is no general closed form formula for polynomials of degree 5, so we cannot get a general closed form for a Taylor series approximation of order 8 or higher. As a result, unfortunately this particular approach does not generalize to arbitrary precision, and it is still an open question as to whether this interval can be reduced further with other means.

In practice, the constraint on the size of the interval of turn is small enough that the controller is still useful. As can be seen in Figure 4.3, the minimum range of (dx_l, dx_u) (the vertical distance between the two lines) is at most $1/6$ for $dx_l \in (0, 1)$, which is the target range of virtually all drifting manoeuvres. However, a small interval along an axis near -1 or 1 could translate to a large interval of turn in radians. Transforming the coordinates to radians, however, we see the same behavior: the interval for (θ_l, θ_u) is reasonably small within the drifting range of up to 90 degrees (at most 0.17 radians, or 9.78 degrees), and the minimum range has the same pattern as dx goes

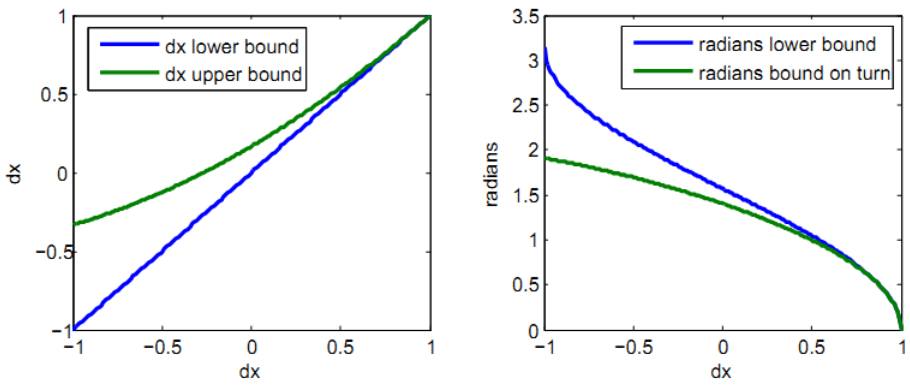

Figure 2: On the left is a plot of dx_l with the minimum possible value for dx_u plotted against dx_l as given in the constraint defined in Lemma 2. On the right is a plot of the upper and lower bounds in radians plotted against dx_l.

.

to -1. If further accuracy is still desired, note that it is technically possible to go up to the 8th order Taylor series approximation using the exact same proof strategy, since we can solve the roots of polynomials of degree 4.

4.4 Proposed Top-Down 2D Model

Previous sections assumed the initial direction of the vehicle is the x-axis of the coordinate plane. While this is acceptable for making decisions local to that particular vehicle, it is an interesting problem to model the controller from a top-down perspective in a 2D plane to make it possible to manage many vehicles at once, from the perspective of an overwatching controller.

Model 4 represents our proposed 2D model $d\mathcal{L}$: it first repeatedly executes a control decision followed by dynamics representing the linear movement. Then, it executes another control followed by dynamics representing the drifting motion cf. (37). While it is safe to accelerate (38), the time is set to zero, the acceleration was set to A, and the angular velocity is set to zero since we are still travelling straight and are not turning cf. (39)-(41). The vehicle is safe when the car can drift to a stop at least $buffer$ distance away in both x and y position from the obstacle. Thus, it is safe to accelerate when the vehicle is safe after accelerating for another time interval cf. (42),(43). After the car is done accelerating, it enters into the drifting phase. The drift control decision (44)-(46) is identical to the decision in Model 1 followed by the dynamic evolution. The basic setup of our proposed 2D model is that we have the car travelling in an initial random direction with respect to the original frame denoted by (dx_{orig}, dy_{orig}) and we denote the direction with respect to the car's axis as (dx_{car}, dy_{car}). Finally the continuous dynamics of the car are defined in (47)-(49).

In our proposed 2D model, we assume that the car is already in motion and heading towards a random direction with an obstacle in front of it, since this is the only case where drifting may be unsafe. The goal of this controller was to prove that our controller guarantees that it turns away from the obstacle after drifting, and stops before the car hits the obstacle when starting in an arbitrary initial direction. Note that we continue to calculate the change in direction due to drifting from the perspective of the car, so that the previous Taylor series approximations are still relevant.

With the change in direction relative to the original direction of the car, we can calculate the new direction with respect to the x, y coordinate axis. First note that $cy^2 = a^2(dx^2 + dy^2) = a^2$

Algorithm 4 Proposed 2D Model (2DM)

$$2DM \equiv (ctrl_{linear}; dyn)^*; ctrl_{drift}; dyn; \tag{37}$$
$$ctrl_{linear} \equiv ?safe_{acc} \tag{38}$$
$$av := 0; \tag{39}$$
$$t := 0; \tag{40}$$
$$acc := A \tag{41}$$
$$safe_{acc} \equiv ?((vel \cdot T + \frac{AT^2}{2} + \frac{(vel + AT)^2}{2F} + buffer)^2 < (ox - x)^2 \tag{42}$$
$$\wedge (vel \cdot T + \frac{AT^2}{2} + \frac{(vel + AT)^2}{2F} + buffer)^2 < (oy - y)^2); \tag{43}$$
$$ctrl_{drift} \equiv av := 2F/vel; \tag{44}$$
$$t := 0; \tag{45}$$
$$acc := -F \tag{46}$$
$$dyn \equiv (x' = vel \cdot dx_{orig}, y' = vel \cdot dy_{orig}, vel' = acc, t' = 1 \tag{47}$$
$$dx'_{car} = av \cdot ddx, dy'_{car} = av \cdot ddy, \tag{48}$$
$$ddx' = -ddy \cdot av, ddy' = ddx \cdot av \ \&\ vel \geq 0 \wedge t \leq T) \tag{49}$$

and $cx^2 = b^2(dx^2 + dy^2) = b^2$. Then, we can calculate the new direction after drifting with respect to the original x, y coordinate axis with the following formula

$$\langle cx \cdot dx - cy \cdot dy, cy \cdot dx + cx \cdot dy \rangle \tag{50}$$

where (dx, dy) is the initial direction of the car, and (cx, cy) is the change in direction relative to the car's perspective (see Figure 3).

In order to prove the safety of our model, we need to guarantee that the car turns and faces away from the obstacle and stops before the it hits the obstacle. This requirement can be captured by the following condition:

$$safety \equiv (ox - x)^2 > buffer^2 \tag{51}$$
$$\wedge (oy - y)^2 > buffer^2 \tag{52}$$
$$\wedge (vel = 0 \rightarrow \tag{53}$$
$$(ox - x) \cdot (dx_{car} dx_{orig} - dy_{car} dy_{orig}) \tag{54}$$
$$+ (oy - y) \cdot (dy_{car} dx_{orig} + dx_{car} dy_{orig}) < 0) \tag{55}$$

Using Equation (50), the resulting direction of the car after doing linear algebra turns out to be

$$\langle dx_{car} dx_{orig} - dy_{car} dy_{orig}, dy_{car} dx_{orig} + dx_{car} dy_{orig} \rangle$$

We intended to prove this condition by following a similar procedure in proving Model 2 for (51) and (52) while using Model 1's approximations along with using the dot product to determine whether the actual direction vector is facing away from the obstacle. If the dot product of the

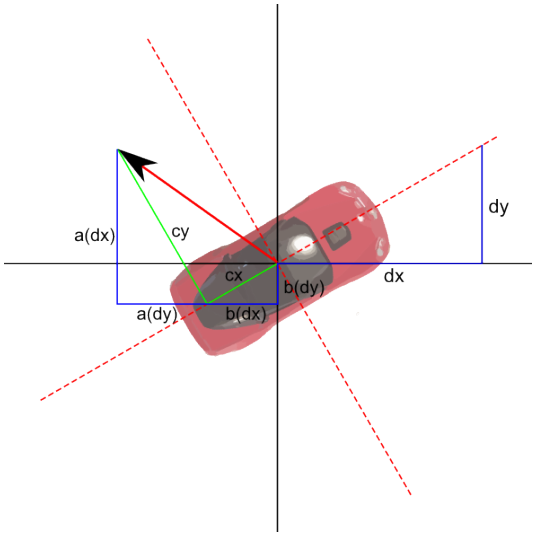

Figure 3: An outline of the variables used for the 2D controller. Note that (dx, dy) is the initial direction vector, and (cx, cy) is the drifting direction relative to the axis aligned with the perspective of the car. a, b are constants, that show that the blue triangles all have the same proportions.

direction the car is facing and the direction of the obstacle is less than 0, then the cosine of the angle between the two vectors must be negative, so the angle between the two vector must be more than 90 degrees. If the angle between the direction of travel and the direction to the obstacle is more than 90 degrees, then the car is moving away from the obstacle (53).

The challenges we faced while developing this model were due to multiple revisions and simplifications of the model. We initially had a 2D model that repeatedly executed the acceleration phase and drift phase where the acceleration phase was not travelling in a straight line. That initial model was too hard to prove due to the need to come up with complex invariants that held throughout the execution. Then we tried removing the repetition but the model still proved to be too hard due to the differential equation involving in a nonlinear motion. Finally we developed this final proposed 2D model where we can reduce our problem down to a linear motion and reuse similar proof concepts from previous models.

4.5 Additional discussion

Due to the difference in direction of travel by the vehicle and direction of the car, we've learned that modeling drifting can be quite complicated. We also learned proving lower bounds on circular motion is a significantly harder task that proving upper bounds, and that the structure of a Taylor series for trigonometric functions almost exactly matches the structure of differential invariant proof rules. When trying to create the 2D model, these Taylor series approximations don't do well if they aren't centered around 0, which is why we had to measure turn from the perspective of the vehicle instead.

We were unable to prove the 2D bot, due to a lack of time. Another unreached goal was a model that accelerates before it comes to a complete stop. This is because such a model requires tracking the residual velocity of the vehicle and combining it with the new velocity from acceleration in a different direction, or combining the residual velocity with more additional residual velocity if drifting is done repeatedly, which is a non-trivial task.

The project deliverables are the four model `.key` files and proofs for all models except for the

2D model. Frederick Chen created the 2D Model and wrote the corresponding section along with the literature review and the abstract. Eric Wong created and proved the Basic Linear Model, Linear Model with an Obstacle, and the Arbitrary Turn Model, and wrote the remaining sections.

5 Conclusion

While automotive control research is a growing field, little work has been done in modeling situations under adverse conditions. Motivated by the particular situation where vehicles lose traction with the ground, possibly due to factors such as snow or ice, it is important to ensure that vehicles make reliable decisions outside of optimal conditions. As a first step in this direction, we presented controller for a cyber-physical system of a vehicle that purposefully loses traction with the ground and drifts to a stop, facing a new target direction. We build from a simple model, culminating in a formal verification and proof of the desired safety and efficiency conditions: namely, our controller guarantees that the car does not drift off the road, and it also ensures the car faces within some arbitrary interval of turn, subtract to a minor constraint on the size of the interval. Finally, we propose a generalized controller that would work for an arbitrary starting direction that may not be parallel to an axis, which is useful when managing multiple cars at once.

Formal verification of vehicle systems under adverse conditions such as loss of traction is non-existent, as most systems either ignore or attempt to avoid such conditions. As a result, there is much space for future work, including planning for unexpected losses of traction, acceleration after loss of traction, drifting on roads with less discrete curvatures, and avoiding collisions with other vehicles while drifting.

References

Jeff C Jensen, Danica H Chang, and Edward A Lee. A model-based design methodology for cyber-physical systems. In *Wireless Communications and Mobile Computing Conference (IWCMC), 2011 7th International*, pages 1666–1671. IEEE, 2011.

Sarah M Loos and André Platzer. Safe intersections: At the crossing of hybrid systems and verification. In *Intelligent Transportation Systems (ITSC), 2011 14th International IEEE Conference on*, pages 1181–1186. IEEE, 2011.

Nikos Aréchiga, Sarah M Loos, André Platzer, and Bruce H Krogh. Using theorem provers to guarantee closed-loop system properties. In *American Control Conference (ACC), 2012*, pages 3573–3580. IEEE, 2012.

Nicholas Chan, James Kuffner, and Matthew Zucker. Improved motion planning speed and safety using regions of inevitable collision. In *17th CISM-IFToMM symposium on robot design, dynamics, and control*, pages 103–114, 2008.

Juergen Ackermann. Robust control prevents car skidding. *Control Systems, IEEE*, 17(3):23–31, 1997.

EK Liebemann, K Meder, J Schuh, and G Nenninger. Safety and performance enhancement: The bosch electronic stability control (esp).

Kiyotaka Kawashima, Toshiyuki Uchida, and Yoichi Hori. Rolling stability control based on electronic stability program for in-wheel-motor electric vehicle. *system*, 2:4, 2009.

Aleksander Hac and Edward Bedner. Robustness of side slip estimation and control algorithms for vehicle chassis control. In *Proceedings of ESV Conference*, 2007.

Stian Ellefsen. Analysis of drifting for a remotely controlled car. 2012.

Jakob Lieng Jakobsen. Autonomous drifting of a 1: 5 scale model car. 2011.

Robert Daily and David M Bevly. The use of gps for vehicle stability control systems. *Industrial Electronics, IEEE Transactions on*, 51(2):270–277, 2004.

Havard Fjaer Grip, Lars Imsland, Tor A Johansen, Jens C Kalkkuhl, and Avshalom Suissa. Vehicle sideslip estimation. *Control Systems, IEEE*, 29(5):36–52, 2009.

Stefan Mitsch, Sarah M Loos, and André Platzer. Towards formal verification of freeway traffic control. In *Cyber-Physical Systems (ICCPS), 2012 IEEE/ACM Third International Conference on*, pages 171–180. IEEE, 2012.