From the Racetrack to the Road: Real-time Trajectory Replanning for Autonomous Driving

John K. Subosits and J. Christian Gerdes

***Abstract*—In emergency situations, autonomous vehicles will be forced to operate at their friction limits in order to avoid collisions. In these scenarios, coordinating the planning of the vehicle's path and speed gives the vehicle the best chance of avoiding an obstacle. Fast reaction time is also important in an emergency, but approaches to the trajectory planning problem based on nonlinear optimization are computationally expensive. This paper presents a new scheme that simultaneously modifies the desired path and speed profile for a vehicle in response to the appearance of an obstacle, significant tracking error, or other environmental change. By formulating the trajectory optimization problem as a quadratically constrained quadratic program, solution times of less than 20 milliseconds are possible even with a 10 second planning horizon. A simplified point mass model is used to describe the vehicle's motion, but the incorporation of longitudinal weight transfer and road topography mean that the vehicle's acceleration limits are modeled more accurately than in comparable approaches. Experimental data from on an autonomous vehicle in two scenarios demonstrate how the trajectory planner enables the vehicle to avoid an obstacle even when the obstacle appears suddenly and the vehicle is already operating near the friction limits.**

I. Introduction

With vehicle manufacturers and other companies pledging to deploy autonomous vehicles within the next few years, we seem to be on the cusp of realizing great improvements in road safety. However, autonomous vehicles will have to handle emergency situations caused by external factors such as other vehicles or wildlife. Curves, steep grades, and surface hazards such as dirt or gravel tracked onto the road only exacerbate the difficulty and pose new challenges for the planning capabilities of an autonomous vehicle [1]. While an emergency on public roads may require a brief use of the full potential of the vehicle to avoid a collision, sustained operation at the friction limits is a hallmark of auto racing. Racing, therefore, provides an avenue to develop control strategies which can make full use of the vehicle's capabilities and to test them in a controlled environment, so they may be used to maximize the potential safety benefits of autonomous vehicles.

This paper presents a real-time trajectory replanning algorithm derived from our work in offline racing line optimization. As described here, the scheme is intended to operate when the assumptions under which the nominal trajectory was generated are no longer valid. One possible situation is the detection of an obstacle that renders the original, desired trajectory no longer free of collision. Another use is for a vehicle facing a large change in the available friction that makes tracking the original trajectory impossible. The response can be either preemptive, for example in response to the detection of a wet or icy road ahead, or reactive, responding by modifying the desired trajectory in response to tracking error when the underlying path and speed tracking controllers having exhausted the tires' capability. Of course, if the original trajectory uses only a fraction of the available friction, even a large reduction in friction may not render the initial trajectory infeasible. At lower accelerations where the friction limits do not have to be considered, other models, such as the widely used kinematic single-track model described in [2] may be more suitable. However, it is shown by Polack et al. that this kinematic model is no longer valid when a maneuver requires more than half of the available tire friction [3]. The scheme presented here works in these more dynamic situations where the friction limitations are a significant constraint on the vehicle's motion.

The authors are with the Department of Mechanical Engineering, Stanford University, Stanford, CA, 94305.

Contact subosits@stanford.edu for further questions about this work.

A. Related Works

The problem of trajectory planning at the friction limits to obtain the minimum lap time and associated trajectory and control inputs for a given vehicle and racing circuit has been widely studied over the years. As shown by Perantoni and Limebeer, these approaches can be used to demonstrate the sensitivity of the optimal line to vehicle parameters [4] and the three dimensional geometry of the circuit [5]. Although Rucco et al. use a similarly complex double-track vehicle model with instantaneous weight transfer in [6], a comparative study of vehicle models by Berntorp et al. showed that model complexity has a small effect on the optimal trajectory [7]. While using nonlinear optimization to find the optimal trajectory for a full circuit is computationally expensive, Timings and Cole showed that convex optimization can be used to reduce the required computation time for these problems [8]. Kapania et al. achieved further reduction in computation time by using alternating convex subproblems to find approximately optimal trajectories [9]. This approach was experimentally validated by testing on an autonomous vehicle as was the receding horizon approach of Gerdts et al. [10]. Anderson and Ayalew showed that the performance of sequential receding horizon approaches could be improved by balancing the travel time over the horizon with the vehicle's final velocity [11]. While great strides have been made in reducing the computation time required, these remain offline trajectory planning methods.

While nominal trajectories for autonomous vehicles can be generated with the tools above, autonomous vehicles must be able to update or replan their trajectories in real-time to respond to the appearance of obstacles or other changes in the environment. Model predictive control has proven popular for

this task and can function as an integrated trajectory generation and tracking approach in some scenarios. In particular, rapid, simultaneous optimization of the driven path and speed via MPC has been used successfully in the control of miniature cars. An early example is given by Verschueren et al. in [12]. Work by Liniger et al. with a vehicle model that accounted for tire slip appeared in [13] and was similarly demonstrated through the racing of 1:43 scale RC cars. The transcription of the dynamics into spatial coordinates which has been widely used in lap time optimization was used in [14] and demonstrated in simulation. While these approaches work well with short planning horizons of approximately one second on model vehicles moving at lower speeds, full scale vehicles require longer planning horizons. In implementations on full size vehicles, real-time operation with longer planning horizons has been achieved by using MPC to compute the optimal steering inputs only after the longitudinal inputs are determined by another means. Liu et al. perform lane change maneuvers by first computing a velocity profile that will place the vehicle in a gap in traffic before performing the lane change with a steering only MPC [15]. In the approaches of Funke et al. [16] and Gutjahr et al. [17], the vehicle tracks a desired velocity profile while avoiding obstacles using steering only. However, being able to modify the vehicle speed would expand the number of situations in which obstacles can be avoided.

Simultaneous planning of a path and speed profile for obstacle avoidance in automated vehicles has been made tractable by only modeling the acceleration limits of the vehicle. In these approaches, the vehicle is modeled as a point mass with various forms of acceleration limits. Funke and Gerdes [18] showed that these acceleration limits could be used to choose an appropriate emergency lane change trajectory from a family of clothoid paths. Emergency lane changes were also studied by Singh and Nishihara [19] and Shiller and Sundar [20] who showed that combining braking and steering minimize the distance required to avoid an obstacle. The problem of minimizing the deviation from a circular reference path when the vehicle's speed is too high for the corner was addressed by Klomp et al. using optimal control theory [21]. Direct numerical methods have also been used for trajectory optimization as in the approaches of Ziegler et al. [22] and Falcone et al. [23]. However, use of nonlinear solvers limits the speed with which they can respond in an emergency. Altché et al. use a similar model, but instead of limiting the magnitude of the acceleration to a constant value, they choose limits that approximate those of a double-track model [24]. Since they work only with the performance limits of the vehicle, the planning approaches described above must be coupled with lower level controllers, such as the MPC controllers above or the feedforward-feedback controllers in [25] and [26].

B. Contributions

This paper presents an approach for autonomous vehicle control which can replan the vehicle's speed and path over a sufficiently long horizon, is suitable for real-time use in emergency obstacle avoidance scenarios, and captures the effects of road topography and individual axle friction limits on the vehicle's limits. The approach assumes the existence of a nominal trajectory, however crude, a low level controller capable of tracking the output of the replanning method, and a means of determining whether an obstacle should be passed to the left or right. It offers shorter computation times, a longer planning horizon, and a more accurate, physically-motivated model of the vehicle's acceleration limits compared to existing trajectory planning approaches for autonomous vehicles. Short computation times of less than 20 milliseconds are achieved even with a 10 second planning horizon by approximating the problem as a convex quadratically constrained quadratic program (QCQP) rather than a general nonlinear one. A simplified point mass model is used, but the vehicle's acceleration limits are modeled more accurately by including the effects of longitudinal weight transfer and, to the author's knowledge, the effects of road topography for the first time in online trajectory replanning. Finally, the effectiveness of the proposed approach is demonstrated experimentally through two aggressive obstacle avoidance maneuvers performed by a autonomous vehicle, including one at the limits of friction.

The next sections describe the form of QCQP used in this work and the reparameterization of a trajectory as a path and associated speed profile. This section is followed by the presentation of the dynamics of the point mass model in three dimensions, the derivation of the acceleration constraints, and the conversion of the dynamics and objective function into a spatial description. Efforts to minimize the approximation error introduced by the linearization of the dynamics and the discretization of the problem are reported in some detail. After describing the implementation of the planning algorithm, we demonstrate the successful real-time use of the algorithm in two scenarios: one inspired by racing and another that requires coordinated modification of speed and path to avoid an obstacle.

II. Overview of Method

The method for trajectory modification considered here closely approximates the finite-horizon optimal control problem by a convex QCQP which can be solved quickly on an autonomous vehicle. Specifically, we consider problems of the form

$$\begin{aligned} \underset{z}{\text{minimize}} \quad & \sum_{i=1}^{N} f_i^T z_i + \frac{1}{2} z_i^T H_i z_i \\ \text{subject to} \quad & C_{i-1} z_{i-1} + D_i z_i = c_i, \\ & \underline{z}_i \leq z_i \leq \overline{z}_i, \\ & A_i z_i \leq b_i, \\ & z_i^T Q_{i,k} z_i + L_{i,k}^T z_i \leq r_{i,k}, \ k = 1, ..., n_q \end{aligned}$$

where z_i is a vector combining the state, control, and slack variables at the ith discretization point. These variables are measured as deviations from a nominal trajectory which is determined ahead of time by another method such as those described in the introduction. The objective function in this work is a second-order approximation to the change in predicted travel time over a fixed-distance planning horizon. The discretization of the dynamics at fixed points in space instead

of time allows time to be treated as a state variable and minimized directly. While any objective of the given quadratic form could be used, the choice of the minimum time objective was dictated by the desire to use this method for racing and to push the car to its limits. The linear equality constraints encode the vehicle dynamics and enforce that the planned trajectory starts from the vehicle's current state and returns to the nominal path at the end of the planning horizon. The box constraints enforce that the planned path stays within the road edges, limit changes in acceleration to physically possible levels, and ensure that the speed at the end of the horizon is not too great for the upcoming section of path. Finally, the linear and quadratic inequality constraints are used to model engine power and tire friction limits, respectively. The optimization results can be combined with the original trajectory to give a useful, approximately minimum time trajectory that can be tracked with the feedforward-feedback controller architecture described in [25]. This method allows the car to avoid obstacles blocking the initial path, adjust to state offset from the nominal path, and address changes in the available friction.

III. Nominal Trajectory Description

A nominal trajectory for the vehicle can be conveniently decomposed into a path and a speed profile. In two dimensions, a path is a geometric object parameterized by its curvature $k(s)$ where s is the distance along the path. To be geometrically consistent, a path must satisfy the following integration constraints:

$$\psi(s) = \psi(0) + \int_0^s k(x)dx, \tag{1}$$

$$E(s) = E(0) - \int_0^s \sin(\psi(x))dx, \tag{2}$$

$$N(s) = N(0) + \int_0^s \cos(\psi(x))dx, \tag{3}$$

In the above, ψ, E, and N are the heading angle, east position, and north position coordinates, respectively, of the path at s. Heading is defined positive counterclockwise from north. The curvature of the path is required to be continuous so that the vehicle can track the path at speed. While any path meeting these requirements can be used, the nominal path in the experiments presented here is a racing line designed to complete the circuit in minimum time and created through an offline version of the algorithm presented in this paper, extended to plan a path for the entire track.

The other component of the trajectory is the speed profile expressed as the desired speed $V(s)$ along the path. The maximum speed is a function of the curvature of the path and the coefficient of friction between the road and the vehicle's tires. On a flat, level road the required centripetal acceleration is given by V^2k and is limited by the available friction between the tires and the road. To minimize time in racing, any grip that is not used for cornering should be used for accelerating or for braking for upcoming corners. The approach presented in [27] was used here to generate minimum time trajectories for the known path, but any speed profile that approximately satisfies the friction constraints could be used as part of the nominal trajectory.

While the above notation suffices to describe paths in a plane, roads are generally not flat and level, and it is necessary to consider the three-dimensional nature of the path to consistently obtain good performance near the limits of friction. In this work, the road is treated as locally planar but can move in pitch and roll. We introduce the following orthogonal unit vectors to describe a reference frame relative to the path at some value s: $\hat{p}_x$ tangent to and directed along the three dimensional path, $\hat{p}_y$ in the plane of the road and directed laterally towards the left edge, and $\hat{p}_z$ locally normal to the road surface and directed upwards. The orientation of this reference frame at s can be obtained from an inertial, Earth-fixed one by a series of right-handed elementary rotations. A rotation of $\psi(s)$ in yaw is followed by rotations of $\theta(s)$ in pitch and $\phi(s)$ in roll about the intermediate axes. Positive values of θ and ϕ correspond to the road sloping downhill and to the right. The orientation of the path along the road surface is still given by the heading angle $\psi(s)$, and s is now the distance along the path in three dimensions.

IV. System Dynamics and Inputs

Various levels of model fidelity have been used for racing line or vehicle trajectory optimization [4]- [11]. The simplest reasonable model for trajectory planning up to the limits, such as in racing applications, is a point mass with friction and engine power constraints. One key assumption of the point mass model is that the vehicle's yaw dynamics can be assumed to remain near steady-state. While transient dynamic models may seem far superior to such a quasi-static approach, actual comparisons of the two approaches show little difference in the optimal line [28] or the actual elapsed time [29] in racing applications where the accelerations are high. Furthermore, Casanova et al. showed that the yaw inertia of the vehicle has little effect on the optimal lap suggesting that the yaw dynamics are relatively insignificant in trajectory planning [30]. A low-level controller such as the one described by Kapania and Gerdes in [25] is required to stabilize the vehicle's yaw motion and generate feedforward steering commands for a desired path. Finally, it is assumed that lateral load transfer and its steady-state effect on the available grip of an axle [31] can be captured with sufficient accuracy by lumping the two tires together and choosing the friction coefficient for that axle accordingly. As pointed out by Liu et al. in the context of rollover avoidance with single-track vehicle models, simpler models may be used so long as knowledge of more complex dynamics are used to establish appropriate limits [32]. In addition to being supported by the literature, the assumptions above are further validated through the presented real world tests with a full size vehicle.

In the form considered here, the point mass model has 4 states and 2 control inputs. The lateral position of the center of mass of the car relative to the nearest point of the 3D path can be defined easily in the path tangent frame as $e\,\hat{p}_y$. The closest point on the path is a distance s along the path. The vehicle's velocity vector within the road plane has magnitude

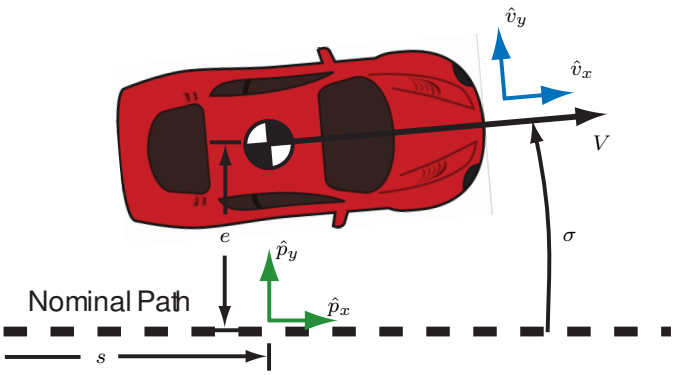

Fig. 1. Definition of reference frames and vehicle states relative to the nominal path.

V and orientation given by a rotation σ from $\hat{p}_x$ about $\hat{p}_z$. It is expedient to use the elementary rotation by σ to define another reference frame, the velocity frame, described by the unit vectors $\hat{v}_x$, $\hat{v}_y$, and $\hat{v}_z = \hat{p}_z$. The state vector x is given by $x = [s\ e\ V\ \sigma]^T$. Positive values of all four state variables are shown in Fig. 1. Other choices of variables are possible, but these simplify the expressions for the equations of motion and the constraints.

A. Kinematics and Dynamics

The velocity of the center of mass of the car relative to the inertial, Earth-fixed frame N is given by

$$^N\vec{v}^{cg} = \dot{s}\,\hat{p}_x + \dot{e}\,\hat{p}_y + {}^N\vec{\omega}^P \times (e\,\hat{p}_y) \quad (4)$$

$$= V\,\hat{v}_x + {}^N\vec{\omega}^P \times (e\,\hat{p}_y) \quad (5)$$

where $^N\vec{\omega}^P$ is the angular velocity of the path frame relative to N. The angular velocity of the path frame can be expressed in terms of the derivatives of the Euler angles as

$$^N\vec{\omega}^P = \omega_x\,\hat{p}_x + \omega_y\,\hat{p}_y + \omega_z\,\hat{p}_z \quad (6)$$

where $\omega_x = \dot{\phi} - \dot{\psi}\sin\theta$, $\omega_y = \dot{\theta}\cos\phi + \dot{\psi}\sin\phi\cos\theta$, and $\omega_z = \dot{\psi}\cos\theta\cos\phi - \dot{\theta}\sin\phi$. Note that the angular velocity of the velocity frame can then be expressed as

$$^N\vec{\omega}^V = (\omega_x\cos\sigma + \omega_y\sin\sigma)\,\hat{v}_x + (\omega_y\cos\sigma - \omega_x\sin\sigma)\,\hat{v}_y + (\omega_z + \dot{\sigma})\,\hat{v}_z. \quad (7)$$

The velocity can then be computed as

$$^N\vec{v}^{cg} = (\dot{s} - e(\dot{\psi}\cos\theta\cos\phi - \dot{\theta}\sin\phi))\,\hat{p}_x + \dot{e}\,\hat{p}_y + e(\dot{\phi} - \dot{\psi}\sin\theta)\,\hat{p}_z \quad (8)$$

or

$$^N\vec{v}^{cg} = V\cos\sigma\,\hat{p}_x + V\sin\sigma\,\hat{p}_y + e(\dot{\phi} - \dot{\psi}\sin\theta)\,\hat{p}_z. \quad (9)$$

The acceleration of the center of mass is then given by

$$^N\vec{a}^{cg} = {}^N\vec{\omega}^V \times {}^N\vec{v}^{cg} + \dot{V}\,\hat{v}_x + (\dot{e}(\dot{\phi} - \dot{\psi}\sin\theta) + e(\ddot{\phi} - \ddot{\psi}\sin\theta - \dot{\psi}\dot{\theta}\cos\theta))\,\hat{v}_z. \quad (10)$$

However, we would like to write the equations of motion solely in terms of the state variables and the path description. For example, the Euler angle rates can be written as

$$\dot{\psi} = \frac{d\psi}{ds}\dot{s} = \psi'\dot{s}. \quad (11)$$

The rate of progress along the path, $\dot{s}$, can then be written as

$$\dot{s} = \frac{V\cos\sigma}{1 - (\psi'\cos\theta\cos\phi - \theta'\sin\phi)e} \quad (12)$$

while for a path in two dimensions the simpler expression

$$\dot{s} = \frac{V\cos\sigma}{1 - k(s)e} \quad (13)$$

applies. This observation motivates a redefinition of k as

$$k = \psi'\cos\theta\cos\phi - \theta'\sin\phi \quad (14)$$

so that either the two or three dimensional cases, k describes how the path tangent rotates about the normal to the road surface with distance traveled along the path. We also introduce $i = \phi' - \psi'\sin\theta$, the local twisting of the road surface, and $j = \theta'\cos\phi + \psi'\sin\phi\cos\theta$, the component of the path's total curvature in the plane of the road. Meanwhile, the rate of change of lateral position is simply

$$\dot{e} = V\sin\sigma. \quad (15)$$

In this model, the only forces acting on the vehicle are gravity, wind resistance, and the tire forces. The point mass model has two control inputs, the total lateral force, F_y, and the total longitudinal force, F_x, generated by the vehicle's tires. From (10) and Newton's second law, the rate of change of the track angle offset σ can be computed as

$$\dot{\sigma} = \frac{F_y + m\vec{g}\cdot\hat{v}_y}{mV} - \dot{s}k + \frac{\dot{s}^2(i^2\cos\sigma + ij\sin\sigma)}{V}e \quad (16)$$

where m is the vehicle mass and $\vec{g}$ is the vector describing the gravitational acceleration. Longitudinal dynamics are given by

$$\dot{V} = \frac{F_x - \frac{1}{2}\rho A C_d V^2 + m\vec{g}\cdot\hat{v}_x}{m} - \dot{s}^2(ij\cos\sigma - i^2\sin\sigma)e \quad (17)$$

including a standard model for aerodynamic drag. Some intuition can be gained by examining the simplified equations for planar roads,

$$\dot{\sigma} = \frac{F_y}{mV} - \dot{s}k(s). \quad (18)$$

and

$$m\dot{V} = F_x - \frac{1}{2}\rho A C_d V^2. \quad (19)$$

Finally, the expressions can be simplified slightly by introducing $a_x = \frac{F_x}{m}$ and $a_y = \frac{F_y}{m}$ and working solely in terms of accelerations.

B. Constraints

For a minimum time optimal control problem, the limitations on the control inputs have a major impact on the optimal trajectory. For a vehicle, the two major constraints are the limited friction between the tires and the road and the maximum torque that the engine can produce. The longitudinal force is subject to an upper limit imposed by the vehicle's limited engine power P:

$$F_x \leq \frac{P}{V}. \tag{20}$$

This limit comes from modeling the engine as producing constant power independent of vehicle speed. This approximation is rough but reasonable because the vehicle transmission is used to keep the engine as near peak power output as possible when racing. The friction limit of a single tire is typically expressed by a friction circle,

$$\sqrt{F_x^2 + F_y^2} \leq \mu F_z \tag{21}$$

where μ is the coefficient of friction between the tire and road and F_z is the normal load on the tire. The following presentation uses the simpler, equivalent expression,

$$\sqrt{a_x^2 + a_y^2} \leq \mu a_z \tag{22}$$

where (21) is normalized by the vehicle mass.

1) Influence of Road Topography on Normal Loads: While road topography has a minor direct effect on the evolution of the vehicle states, it can have a dramatic effect on the tire normal loads. From kinematics, the acceleration of the point mass normal to the road surface is

$${}^N\vec{a}^{cg} \cdot \hat{p}_z = \dot{e}\dot{s}i + e\ddot{s}i + e\dot{s}^2 i' - V\dot{s}(j\cos\sigma - i\sin\sigma) \tag{23}$$

where the chain rule has been used to express $\dot{i}$ as $\dot{s}i'$. Therefore the tire constraints can be updated using

$$a_z = -\vec{g}\cdot\hat{p}_z + \dot{e}\dot{s}i + e\ddot{s}i + e\dot{s}^2 i' - V\dot{s}(j\cos\sigma - i\sin\sigma). \tag{24}$$

The first term means that a_z generally has a value near the constant of gravitational acceleration, 9.81 m/s^2. However, the term $V\dot{s}j\cos\sigma$ can be extremely significant. This normal or "vertical" curvature, j, describes how the pitching down of the road surface at hill crests causes a reduction in normal load. Conversely, normal load increases where the road is concave upwards. Variations of normal load of greater than 25% with a corresponding change in the acceleration potential of the vehicle have been observed at speeds as low as 15 m/s on various circuits due to this effect.

2) Longitudinal Weight Transfer: Typically the friction circle model is meant for a single tire, and would only be an approximation for the whole vehicle. While modern chassis control systems such as torque vectoring, Traction Control Systems, Electronic Brake-bias Variation, Anti-lock Brake Systems, and Electronic Stability Control improve the quality of such an approximation by dividing the tire forces to more closely match the force potential at each wheel, the fidelity with which the model predicts the vehicle's acceleration limits can be improved by considering the friction limits of each axle independently and modeling longitudinal weight transfer. By balancing the pitching moments about the center of mass and assuming aerodynamic drag acts at the center of mass, the normal load on the front axle can be calculated as

$$F_{zf} = \frac{b}{L}F_z - \frac{h_{cg}}{L}F_x \tag{25}$$

where L is the vehicle's wheelbase, b is the horizontal distance from the c.g. to the rear axle, and h_{cg} is the height of the c.g. above the ground. The equivalent expression for the rear tires is given by

$$F_{zr} = \frac{a}{L}F_z + \frac{h_{cg}}{L}F_x \tag{26}$$

where $a = L - b$ is the horizontal distance from the center of mass to the front axle. For a vehicle traveling on a flat, level roadway, F_z is just the mass of the vehicle multiplied by the gravitational acceleration, but road topography and speed have an effect in the general case. Two simplifying assumptions are used in the derivation of the above expressions. The first is that suspension motion is ignored and only steady state weight transfer is modeled since suspension dynamics are much faster than the dynamics governing path tracking. Also, the fraction of the pitching moment that keeps the chassis aligned with the surface of the road is neglected since this effect is very small for larger road features such as hills.

While computation of the normal load is straightforward, there are several possible ways to model the division of longitudinal and lateral forces between the two tires, and this choice can have major implications for the vehicle's acceleration limits. An imbalance between the front and rear lateral forces produces a yawing moment. In racing, however, yawing moment is generally considered undesirable since it reduces the amount of grip that can be used for lateral and longitudinal acceleration. Since the point mass model does not consider yaw dynamics explicitly, we choose to proportion the lateral force so that there is no yawing moment; an alternative implementation would constrain the the vehicle sideslip angle β to remain constant ($\dot{\beta} = 0$) without explicitly modeling the yaw dynamics. However, assuming the yaw moment can be neglected results in a simpler form for the tire friction constraints.

In contrast, the model for longitudinal force must be tailored to the mechanical design of the vehicle. Handling both front and rear wheel drive vehicles is an obvious example. However, in some cases, the longitudinal force distribution is not known exactly or may vary rapidly in response to wheel slip. These situations can occur in AWD vehicles with active center differentials and in vehicles with ABS. Besides these examples, it would be unusual for the distribution of engine and brake torque in a vehicle to the be same. We use a model that addresses this difficulty for AWD vehicles while avoiding making the acceleration constraints overly conservative. In this work, longitudinal forces are nominally distributed according to the static normal load on each axle, but a new variable, da_x, is introduced to model a changing distribution of longitudinal forces between the two axles. While introducing this additional variable slightly complicates the resulting optimization problem, it significantly improves the validity of the model for vehicles with variable force distribution. Under the above

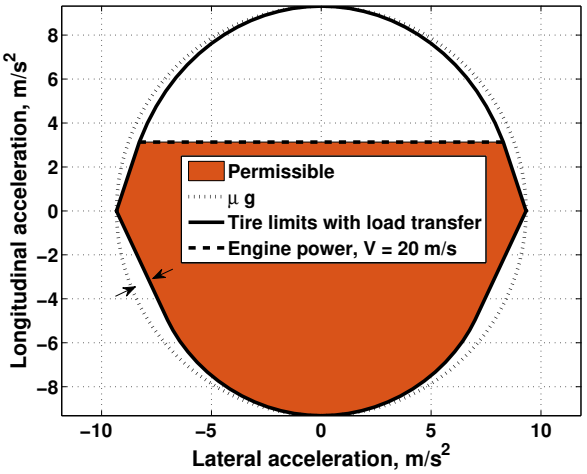

Fig. 2. Illustration of control limits with $\mu = 0.95$ and the engine power limit evaluated at a vehicle speed of 20 m/s. The arrows indicate a difference in capability of 0.88 m/s^2 or almost 10% when longitudinal weight transfer is modeled as compared to the circle of radius μg.

assumptions, the friction circle constraint for the front axle is given by

$$\left(\frac{b}{L}a_x - da_x\right)^2 + \left(\frac{b}{L}a_y\right)^2 \leq \left(\mu(\frac{b}{L}a_z - \frac{h_{cg}}{L}a_x)\right)^2. \quad (27)$$

Similarly, for the rear axle we have

$$\left(\frac{a}{L}a_x + da_x\right)^2 + \left(\frac{a}{L}a_y\right)^2 \leq \left(\mu(\frac{a}{L}a_z + \frac{h_{cg}}{L}a_x)\right)^2. \quad (28)$$

While these expressions are designed to model a vehicle with variable drive and brake force allocation, FWD or RWD vehicles can be easily modeled by constraining the total longitudinal force from the undriven axle to be non-positive.

When the engine power limit is combined with both the tire friction constraints, the possible vehicle accelerations are as shown in Fig. 2. When the vehicle is accelerating, the front tires limit the total acceleration, and the grip of the rear tires limits the vehicle while braking and turning. While modeling longitudinal weight transfer does not radically change the acceleration limits, it does reduce the amount of braking or acceleration that is possible at high lateral loads. This change reduces the risk that the vehicle experiences significant limit oversteer while braking and cornering or understeer while accelerating out of a corner. In addition, while friction is assumed to be isotropic here, the friction circle can easily be generalized to a friction ellipse. Such a change allows the model to capture more phenomena such as limit understeer at constant speed.

C. Transformation of Dynamics

In order to shift time from the independent to an optimization variable, we can rewrite the ODE's governing the vehicle's motion in terms of distance along the path. Instead of writing the equations of motion as

$$\dot{x}(t) = F(x(t),\ u(t)) \quad (29)$$

with $x(t) = [s(t)\ e\ V\ \sigma]^T$ and $u(t) = [a_x\ a_y]^T$, we wish to express the state as $x(s) = [t(s)\ e\ V\ \sigma]^T$ which is unambiguous so long as $\dot{s}$ is always positive. The dynamics can be transformed into ordinary differential equations in terms of s by dividing through by $\dot{s}$ since by the chain rule,

$$\frac{\dot{x}(t)}{\dot{s}} = \frac{dx}{dt}\frac{dt}{ds} = \frac{dx}{ds} = x' \quad (30)$$

As examples, (15) becomes

$$e(s)' = (1 - k(s)e(s))\tan\sigma(s) \quad (31)$$

and, importantly but somewhat trivially,

$$t(s)' = \frac{1}{\dot{s}} = \frac{1 - k(s)e(s)}{V(s)\cos\sigma(s)}. \quad (32)$$

The change of the independent variable from time to distance is not just useful in a racing context where we might wish to explicitly minimize time. It also allows the precise and straightforward encoding of obstacle avoidance constraints for stationary obstacles: For a given value of s, there is a fixed upper and lower limit on e before a part of the vehicle must be outside of the lane or off the road. If time is used as the independent variable, obstacle avoidance constraints are coupled constraints in the two position coordinates. Furthermore, the form of the constraints on these variables may not be convex, making the problem more difficult to solve quickly.

V. Minimum Time Objective Function

The problem of racing line optimization is typically posed as finding a trajectory that completes a lap in minimum time, minimizing

$$t_f = \int_0^{t_f} dt = \int_0^l \frac{ds}{\dot{s}(s)}, \quad (33)$$

where we have made a change of variables using the chain rule. Note that l is the distance along the path over which time is measured and that $s = l$ at t_f. For a finite horizon approach, the distance l is the length of the planning horizon. A sequence of control inputs must be found to minimize the above expression subject to the limitations on the controls and the dynamics described previously.

VI. Linearization and Discretization

To transform the problem into a form that can be solved quickly, the dynamics are first linearized about a nominal trajectory. While not required to produce a useful modified trajectory, the nearer this nominal trajectory is to being feasible and time-optimal the better since less approximation error will then be introduced by the linearization. The method of discretization is chosen to limit the resulting error as much as possible.

A. Linearization

To include the vehicle dynamics as equality constraints in a quadratic program (QP), we place the equations of motion in the following linear form:

$$x(s)' = A(s)x(s) + B(s)u(s), \quad (34)$$

where $x(s) = [\Delta t \ e \ \Delta V \ \sigma]^T$, $u(s) = [\Delta a_x \ \Delta da_y]^T$ and Δ denotes a perturbation from the nominal value at that point on the path. The dynamics matrix and controls matrices are given by

$$A(s) = \begin{bmatrix} 0 & \frac{-k}{V} & \frac{-1}{V^2} & 0 \\ 0 & 0 & 0 & 1 \\ 0 & \frac{-k\sum F_x}{mV} - Vij & \frac{-\sum F_x}{mV^2} - \frac{\rho AC_d}{m} & \frac{\vec{g}\cdot\hat{p}_y}{V} \\ 0 & \frac{-k(a_y+\vec{g}\cdot\hat{p}_y)}{V^2} + i^2 & \frac{-2(a_y+\vec{g}\cdot\hat{p}_y)}{V^3} & -\frac{\vec{g}\cdot\hat{p}_x}{V^2} \end{bmatrix} \tag{35}$$

and

$$B(s) = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ \frac{1}{V} & 0 \\ 0 & \frac{1}{V^2} \end{bmatrix} \tag{36}$$

respectively, where

$$\frac{\sum F_x}{m} = a_x - \frac{\frac{1}{2}\rho AC_d V^2}{m} + \vec{g}\cdot\hat{p}_x. \tag{37}$$

Since it has been assumed up to this point that the vehicle is tracking the path, the nominal values of e and σ are zero everywhere. The dynamics are in now in a form which can easily be discretized.

Since our goal is to minimize the value of Δt at the end of the planning horizon and we are formulating the problem as a (QC)QP we can also make use of the second-order term in the Taylor series approximation to $\Delta t'$;

$$\nabla_x^2 \Delta t' = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & \frac{k}{V^2} & 0 \\ 0 & \frac{k}{V^2} & \frac{2}{V^3} & 0 \\ 0 & 0 & 0 & \frac{1}{V} \end{bmatrix}. \tag{38}$$

While this matrix is not positive semi-definite, as is required to rapidly find a unique solution to the trajectory planning problem, this issue will be addressed after the equations have been discretized.

B. Discretization

The linearized equations of motion are used to generate linear discrete equations describing the state transitions from one point in space to the next. The discretization points are evenly spaced in time along the nominal trajectory which has the advantage of giving finer spatial resolution in the corners than on straightaways. While many methods exist to form discrete equations from linear continuous time ones, an approach from the theory of linear time varying systems was used in this work. The linearized state at a given point in space is given by

$$x(s) = \Phi(s, s_0)x_0 + \int_{s_0}^{s} \Phi(s,\ \tau)B(\tau)u(\tau)d\tau \tag{39}$$

where the state transition matrix $\Phi(s,\ s_0)$ is given by the Peano-Baker series,

$$\Phi(s,\ s_0) = I + \int_{s_0}^{s} A(s_1)ds_1 + ... \tag{40}$$

In this work, the higher-order terms were neglected since the above equation was used only to calculate the transition matrices between the discretization points of the nominal trajectory which are much closer together than the discretization points used for trajectory replanning. Trapezoidal integration with step sizes corresponding to the discretization interval of the original path was used here and throughout the discretization process. The state transition matrices are then computed over the longer discretization intervals used in the trajectory optimization problem via the following property of the state transition matrix,

$$\Phi(s,\ s_0) = \Phi(s,\ \tau)\Phi(\tau,\ s_0), \tag{41}$$

so many discrete integration steps are combined to produce a single state transition matrix.

The calculation of the effect of the control inputs is slightly more involved. Since the control inputs are ultimately implemented via a first-order hold over track distance, the discretized equations of motion are placed in the form,

$$x_i = A_{i-1}x_{i-1} + B_{i-1}u_{i-1} + B_i u_i \tag{42}$$

where $A_{i-1} = \Phi(s_i,\ s_{i-1})$. Between the discretization points, the value of the control input is given by

$$u(\tau) = \frac{s-\tau}{s-s_0}u(s_0) + \frac{\tau - s_0}{s-s_0}u(s) \tag{43}$$

The above equation can be substituted into (39) and integrated to obtain the form in (42).

The Hessian of the objective function also needs to be expressed in terms of the optimization variables. Since the dependence of the state between the discretization points on the state and controls at those points is known, the quadratic costs for the states and controls at the discretization points can then be computed via integration:

$$H(s_i) = \int_{s_{i-1}}^{s_i} x(\tau)' \left(\nabla_x^2 \Delta t'\right) x(\tau) d\tau \tag{44}$$

While the matrix in (38) is sparse, the matrix for the discretized problem has many more nonzero entries and depends additionally on the controls and how they change. The resulting discrete matrix is not necessarily positive definite as required by the convex solver, so it is approximated by performing a symmetric eigenvalue decomposition, setting the non-positive eigenvalues to a small positive number, and performing the multiplication to obtain a positive definite approximation to the original matrix.

It is mildly computationally expensive to compute this accurate discretization, but this only has to be done once for a given nominal trajectory. Considering this discretization as part of the nominal trajectory generation reveals it to be only a small additional cost. Also, several quantities appear repeatedly in the various calculations, so they need only be computed once. While simpler discretization approaches exist, the one used here has the advantage of using all available data from the nominal trajectory.

VII. Implementation details

Both the engine power limit and the tire friction limits are represented as convex versions of the second-order Taylor series approximations around the nominal trajectory. To second-order, the engine power limit is simply an affine inequality.

The rate of change of the total acceleration in both the lateral and longitudinal directions is also limited to account for the maximum rate of response of the vehicle and its actuators.

A. Slack variables

In an attempt to ensure that a solution always exists to the replanning problem, the friction circle constraint at each stage is implemented with a slack variable. The friction coefficient μ is replaced with $\mu + \nu$ in (27) and (28). A large penalty on the slack variable, $100,000 \sum_i \nu_i^2$, is placed in the objective function, weighting an increase in the friction coefficient by 0.01 as heavily as the loss of 10 seconds. The friction coefficient μ is also difficult to measure accurately, particularly when the car is operating far from the friction limits. In light of this fact, the numerical value of μ used in the optimization problem can be viewed as limiting the total acceleration to a desired level. Setting this value higher than the true tire-road friction coefficient will result in replanned trajectories that the vehicle cannot track while setting it lower results in perfectly usable trajectories, albeit with a probable increase in the value of the cost function. Note that when the slack variable is forced positive, it is usually only at the beginning of the planning horizon which encourages the car to use all the friction currently available.

B. Terminal constraints and cost

To ensure that the transition between the replanned and nominal trajectory at the end of the planning horizon is seamless, we require $e = 0$, $\sigma = 0$, $\dot{\sigma} = 0$, and $V \leq V_{max}(s_N)$ at the end of the trajectory. To maintain convexity, an affine approximation of $\dot{\sigma} = 0$ is used. The generation of the initial speed profile provides the maximum speed the car can be traveling at any point on the path while able to track the future parts of the trajectory. To mitigate the limitations of a finite horizon, the effect of a speed difference from the nominal trajectory at the end of the planning horizon is estimated and incorporated into the cost function as a linear term. Approximating the cost-to-go in this fashion allows differences in the problem depending on where the horizon ends to be approximately quantified. For example, when driving onto a long straightaway, the time penalty for a reduction in speed is high while the maximum permissible speed is far above the nominal trajectory. The situation is reversed if the horizon ends on a part of the track where the car is slowing for an upcoming corner. The optimal computed trajectory is therefore both safe because it uses information gleaned from the generation of the initial trajectory and nearer optimal than would otherwise be possible even with such a long planning horizon because of the chosen terminal cost.

Choosing the constraints at the start of the planning horizon also takes some care. The plan needs to start from the current state of the car and the controls must be continuous. While matching lateral accelerations or curvatures may seem like an obvious choice, matching steering angle is actually key when performing dynamic maneuvers. Since replanning from the current state of the vehicle means that any tracking error vanishes, large changes in steering commands are possible. The exact means for handling this issue will depend on the architecture of the underlying path-tracking controller.

Fig. 3. Autonomous Audi TTS used for vehicle control research at the limits of friction

VIII. Experimental Validation

Two experiments with an autonomous vehicle were performed to demonstrate the advantages of this replanning approach in obstacle avoidance scenarios. The tests were performed at the Thunderhill West circuit in Willows, CA, using the 2009 Audi TTS shown in Fig. 3 which has been configured to allow fully autonomous operation. Guidance is provided by a DGPS/IMU system capable of centimeter level position accuracy. Control of steering, throttle, braking, and shifting is performed by a low-level control computer operating at 200 Hz and communicating with the vehicle via multiple CAN buses. The trajectory replanning operation is performed by a second computer (Intel Core i7-3610QE CPU @ 2.30GHz, 8GB RAM, standard Linux kernel) which receives vehicle state information via a UDP packet from the low-level control computer. The QCQP solver is auto-generated code developed with the FORCES Pro framework [33] and typically takes between 10 and 12 milliseconds to reach a solution. In the worst case observed in testing, the modified trajectory is received by the low-level control computer via UDP within 20 milliseconds of sending the initial state information. This time is in addition to any required by the vehicle's perception system to perceive and locate the obstacle. The fact that a new trajectory can be decided upon and put into practice in such a short amount of time means trajectory planning algorithms need not limit the reaction time of autonomous vehicles in emergencies.

The following parameters were used in the experiments. The discretization points are spaced by 1/3 seconds along the nominal trajectory allowing $N = 30$ points to cover a planning horizon of approximately 10 seconds. The changes in the controls were bounded by $|\dot{a}_y| \leq 19\,\mathrm{m/s^3}$ and $-25\,\mathrm{m/s^3} \leq \dot{a}_x \leq 15\,\mathrm{m/s^3}$. A buffer of 0.5 m was added to the road edge constraints to accommodate tracking error on the test vehicle. If the position of the obstacle could only be measured approximately by the vehicle's sensors, this buffer would have to be expanded. The parameters used to model the vehicle are given in Table I.

TABLE I
VEHICLE MODEL PARAMETERS

Parameter	Value	Units
a	1.015	m
b	1.453	m
h_{cg}	0.5	m
m	1659	kg
P	120	kW
$\frac{1}{2}\rho C_D A$	0.499	$\frac{\text{Ns}^2}{\text{m}^2}$
μ	varies	none

The experimental tests begin with the vehicle autonomously tracking the nominal trajectory. When the vehicle reaches a fixed point along the path, the replanning operation is triggered. Updating the constraints on the lateral position e, forces the vehicle to modify the trajectory so that it passes to the right of the obstacle. The vehicle then tracks the modified trajectory until the test is ended by the people supervising the vehicle. Footage from both scenarios is available at https://stanford.box.com/s/hzczr66rq4kkat830yz920s98dbkmopz. Prior to beginning these tests, the exact location of the obstacle is given to the vehicle software rather than relying on the vehicle's sensors. The position measurements of the extreme points of the obstacle in the global coordinate frame are transformed into local s and e position coordinates using the same map matching procedure that is used to compute the test vehicle's position relative to the nominal path. This map matching procedure finds the closest point on the nominal path to a given position using a global then a local search. The constraints on the vehicle's lateral position can then be updated for appropriate values of s based on the computed s and e coordinates of the obstacle corners, taking care to account for the width and length of the test vehicle. When the replanning horizon includes an obstacle, the updated bounds on the lateral position cause the vehicle to plan a path around the side of the obstacle chosen. Testing in this manner removes possible sources of variation when testing the trajectory replanning algorithm.

A. Racing Scenario

The first experimental scenario is drawn from racing where a car may need to avoid a stopped car while minimizing lost time. In this scenario, the car is rounding a hairpin turn when, suddenly, a car-sized obstacle is blocking the intended path. An illustration of the situation is shown in Fig. 4 and a view of the obstacle from the perspective of the autonomous vehicle is shown in Fig. 5. Since the intended path is a racing line that uses the full width of the track, the car must tighten the latter part of the corner. The friction coefficient for both the nominal and replanned trajectories is 0.92, very near the true grip limits of the car.

Despite the need to deviate by more than a full car width from the nominal path, the required changes to the control inputs are quite modest as shown in the middle plot of Fig. 6. Slightly more lateral acceleration is used, particularly at the end of the right hand corner, to avoid the obstacle. No longitudinal forces are applied so that the available lateral grip

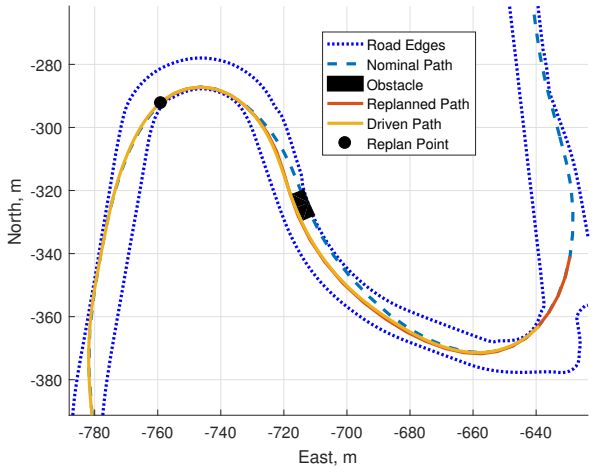

Fig. 4. Illustration of nominal and replanned paths showing both the road edges and the obstacle location for the racing scenario

Fig. 5. Picture from test with obstacle location marked by mock vehicle

is maximized. Full acceleration is still used, but later in the corner, resulting a similar planned speed profile out of the corner as shown in Fig. 7. However, the vehicle is not able to track the desired speed despite operating at full throttle. The vehicle continues to accelerate as it passes the obstacle. Note that the replanned path passes immediately adjacent to the obstacle as shown at the bottom of Fig. 6 to minimize the extra distance traveled. Despite operating very near the limits of tire friction, the vehicle successfully avoided the obstacle in each experimental trial. Once the car is even with the obstacle, the slight bend to the left is tightened to set the vehicle up nearly parallel to the road edge to permit it to brake for the last corner. The radius of the final left hand corner is reduced, allowing the car to brake later and enter the corner with more speed, as shown in Fig. 7, recouping some lost time.

Since this scenario is inspired by racing, minimizing lost time is important. The vehicle does not change its intended speed much during this test, but it does not need to. As shown in the top plot of Fig. 6, the predicted loss of time over the horizon is small. Because of the speed tracking error over much of the horizon, the actual loss of time displayed in Fig.

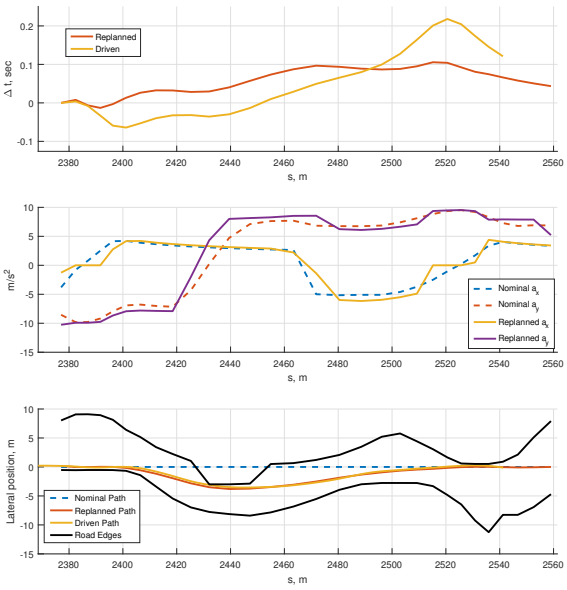

Fig. 6. Illustration of nominal and replanned state and input trajectories

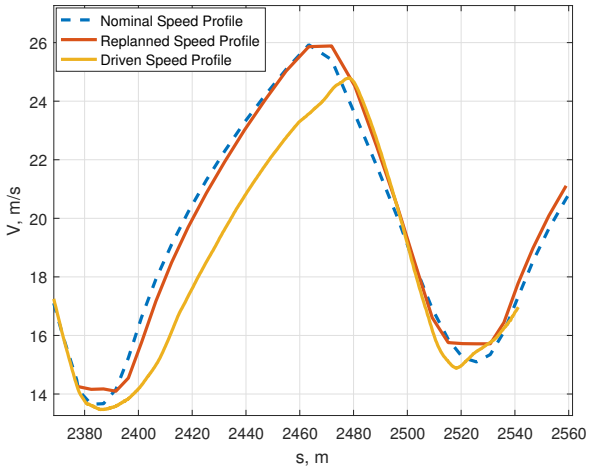

Fig. 7. Speed profiles for the racing scenario

6 is computed relative to that of the vehicle attempting to track the nominal trajectory. With this correction, the prediction of lost time is quite good. By modeling the connection between speed and path, the vehicle is able to understand that the path can be changed as much as needed without significantly reducing the target speed or giving up much time.

B. Obstacle Avoidance Scenario

While the previous example was drawn from racing, this obstacle avoidance scenario begins with the vehicle traveling at freeway speed while cornering fairly aggressively. As the vehicle crests a hill, an obstacle appears slightly more than 100 m ahead of the car. The vehicle must avoid the obstacle blocking the original planned path while still negotiating the

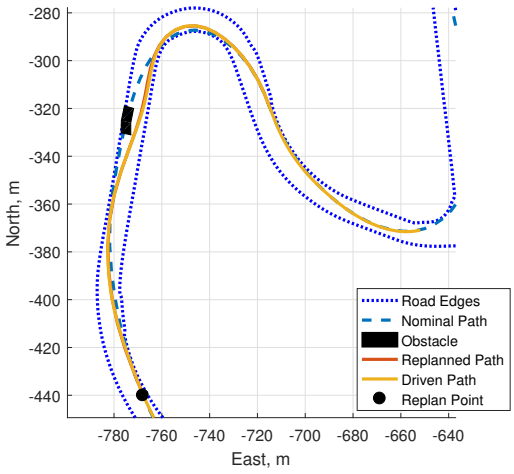

Fig. 8. Illustration of nominal and replanned paths in the obstacle avoidance scenario

hairpin corner just beyond the obstacle as shown in Fig. 8. The placement of the obstacle makes entering the hairpin challenging because the obstacle is blocking the path that would allow the planned high speed entry into the corner. The road topography also complicates the situation. Tracking the nominal trajectory, the vehicle experiences normal accelerations ranging from 8 $\mathrm{m/s^2}$ at the crest of the hill to 11.8 $\mathrm{m/s^2}$ at the bottom. On the replanned trajectory, the changes are slightly smaller at 8.4 to 11.3 $\mathrm{m/s^2}$ because the vehicle is traveling more slowly over the crest of the hill and through the dip at the bottom. This change in normal load still has a significant effect on the accelerations the car can generate while avoiding the obstacle. For safety, the value of the friction coefficient was set at 0.7 in this test, a value that matches the limits of the test vehicle on wet asphalt.

Since the car has to complete the second half of a high speed corner, avoid the obstacle, and negotiate the low speed corner in quick succession, the control inputs are complex and rapidly varying. As shown in the middle plot of Fig. 9, the vehicle initially responds by reducing lateral acceleration to almost zero, thereby straightening the path to allow the vehicle to brake aggressively and quickly reduce speed. The vehicle then returns to maximum cornering, decreasing the radius of the initial path, using the full width of the road despite the initial brake application as shown in the bottom plot of Fig. 9. This reduced radius allows the vehicle to avoid the obstacle as well. As the vehicle passes the obstacle, the vehicle begins braking for the hairpin corner while turning left to approach alongside the inside road edge. The lower speed in the hairpin corner allows the vehicle to turn through a greater angle in the same amount of road and even slightly exceed the original exit speed. The planned and driven velocity profiles for the sequence of maneuvers are shown in Fig. 10. The modified path deviates significantly from the nominal one until very near the end of the planning horizon demonstrating the advantage of using a long planning horizon that reaches all the way through the hairpin corner.

The vehicle must operate at the imposed acceleration limits

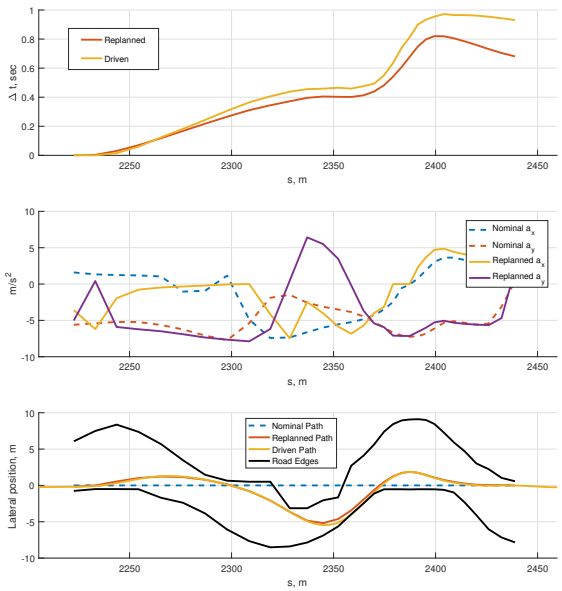

Fig. 9. Change in elapsed time, control inputs, and lateral position during the obstacle avoidance scenario

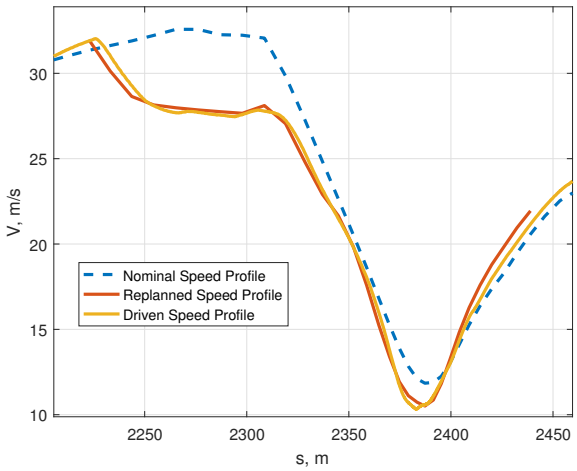

Fig. 10. Avoiding the obstacle requires an initial reduction in speed before slowing earlier for the now tighter corner.

for most of the planning horizon. Fig. 11 makes this necessity obvious. Both the nominal and replanned trajectory require operating at the edge of the friction circle, but the replanned trajectory requires making larger, more rapid moves around it. The non-zero values of the first seven slack variables in the optimization problem also support this conclusion. While successfully handling this situation requires a slight relaxation of the friction circle constraints, the slack variable exceeds 0.01 at only two horizon points with a maximum of 0.015. In practice, it is difficult to estimate the acceleration limits of the vehicle to within this accuracy, particularly for road that has yet to be traversed, and it is preferable for the solver to return a trajectory that the car might have slight difficulty tracking

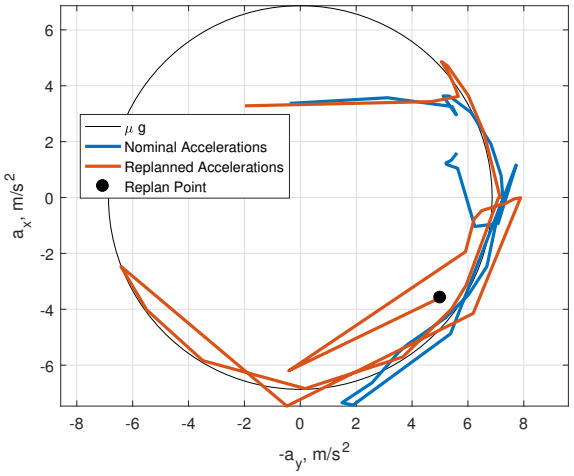

Fig. 11. Comparison of nominal and replanned control inputs via a G-G diagram. The circle is only for comparison. The actual constraints vary over the horizon and are state-dependent.

than fail completely.

Since the trajectory of the vehicle is so highly constrained in this scenario, the loss of time is significant. The plot at the top of Fig. 9 shows the linear approximation for time from the optimization problem. The minimum time loss is 0.68 seconds according to the model while the car actually takes 0.93 additional seconds relative to the desired nominal trajectory. This discrepancy is easily explained by the fact that the car generally tracks slightly slower than the desired speed throughout this test. The predictive power of the model is quite good considering most of the divergence occurs at the exit of the hairpin where the car is already operating at full throttle.

IX. Conclusion

A method for rapid replanning of autonomous vehicle trajectories over a long horizon which more accurately captures the significant effects of road topography and the vehicle friction limits has been presented. The rapidness with which the autonomous car can respond to new obstacles, changes in the predicted level of grip, and significant tracking errors gives it a superior chance of avoiding a collision in an emergency compared to replanning methods that take longer to return an answer. Experiments on a full-size autonomous vehicle have validated the effectiveness of the method for controlling cars in practice. The first test scenario emphasized the connection of the replanning algorithm to racing while the second test demonstrated its performance in a real world obstacle avoidance scenario. This planning method, with the cost function modified to better suit the task of everyday driving, offers future autonomous cars the capability to respond rapidly to emergency scenarios, using the entire physical capability of the vehicle if necessary.

Acknowledgment

The authors would like to thank Vincent Laurense and Nathan Spielberg for assisting with the experiments. Testing was possible with the help of Thunderhill Raceway Park.

REFERENCES

[1] L. Figueiredo, I. Jesus, J. T. Machado, J. R. Ferreira, and J. M. De Carvalho, "Towards the development of intelligent transportation systems," in *Intelligent Transportation Systems, 2001. Proceedings. 2001 IEEE*. IEEE, 2001, pp. 1206–1211.

[2] B. Paden, M. Čáp, S. Z. Yong, D. Yershov, and E. Frazzoli, "A survey of motion planning and control techniques for self-driving urban vehicles," *IEEE Transactions on Intelligent Vehicles*, vol. 1, no. 1, pp. 33–55, 2016.

[3] P. Polack, F. Altché, B. d'Andréa Novel, and A. de La Fortelle, "The kinematic bicycle model: A consistent model for planning feasible trajectories for autonomous vehicles?" in *Intelligent Vehicles Symposium (IV), 2017 IEEE*. IEEE, 2017, pp. 812–818.

[4] G. Perantoni and D. J. Limebeer, "Optimal control for a formula one car with variable parameters," *Vehicle System Dynamics*, vol. 52, no. 5, pp. 653–678, 2014.

[5] D. Limebeer and G. Perantoni, "Optimal control of a formula one car on a three-dimensional trackpart 2: Optimal control," *Journal of Dynamic Systems, Measurement, and Control*, vol. 137, no. 5, p. 051019, 2015.

[6] A. Rucco, G. Notarstefano, and J. Hauser, "An efficient minimum-time trajectory generation strategy for two-track car vehicles," *IEEE Transactions on Control Systems Technology*, vol. 23, no. 4, pp. 1505–1519, 2015.

[7] K. Berntorp, B. Olofsson, K. Lundahl, and L. Nielsen, "Models and methodology for optimal trajectory generation in safety-critical road–vehicle manoeuvres," *Vehicle System Dynamics*, vol. 52, no. 10, pp. 1304–1332, 2014.

[8] J. P. Timings and D. J. Cole, "Minimum maneuver time calculation using convex optimization," *Journal of Dynamic Systems, Measurement, and Control*, vol. 135, no. 3, p. 031015, 2013.

[9] N. R. Kapania, J. Subosits, and J. C. Gerdes, "A sequential two-step algorithm for fast generation of vehicle racing trajectories," *Journal of Dynamic Systems, Measurement, and Control*, vol. 138, no. 9, p. 091005, 2016.

[10] M. Gerdts, S. Karrenberg, B. Müller-Beßler, and G. Stock, "Generating locally optimal trajectories for an automatically driven car," *Optimization and Engineering*, vol. 10, no. 4, pp. 439–463, 2009.

[11] J. R. Anderson and B. Ayalew, "Modelling minimum-time manoeuvering with global optimisation of local receding horizon control," *Vehicle System Dynamics*, pp. 1–24, 2018.

[12] R. Verschueren, S. De Bruyne, M. Zanon, J. V. Frasch, and M. Diehl, "Towards time-optimal race car driving using nonlinear mpc in real-time," in *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*. IEEE, 2014, pp. 2505–2510.

[13] A. Liniger, A. Domahidi, and M. Morari, "Optimization-based autonomous racing of 1: 43 scale rc cars," *Optimal Control Applications and Methods*, vol. 36, no. 5, pp. 628–647, 2015.

[14] R. Verschueren, M. Zanon, R. Quirynen, and M. Diehl, "Time-optimal race car driving using an online exact hessian based nonlinear mpc algorithm," in *Control Conference (ECC), 2016 European*. IEEE, 2016, pp. 141–147.

[15] K. Liu, J. Gong, A. Kurt, H. Chen, and U. Ozguner, "Dynamic modeling and control of high-speed automated vehicles for lane change maneuver," *IEEE Transactions on Intelligent Vehicles*, vol. 3, no. 3, pp. 329–339, 2018.

[16] J. Funke, M. Brown, S. M. Erlien, and J. C. Gerdes, "Collision avoidance and stabilization for autonomous vehicles in emergency scenarios," *IEEE Transactions on Control Systems Technology*, vol. 25, no. 4, pp. 1204–1216, 2017.

[17] B. Gutjahr, L. Gröll, and M. Werling, "Lateral vehicle trajectory optimization using constrained linear time-varying mpc," *IEEE Transactions on Intelligent Transportation Systems*, 2016.

[18] J. Funke and J. C. Gerdes, "Simple clothoid lane change trajectories for automated vehicles incorporating friction constraints," *Journal of Dynamic Systems, Measurement, and Control*, vol. 138, no. 2, p. 021002, 2016.

[19] A. S. P. Singh and O. Nishihara, "Nondimensionalized univariate equation characterizing optimal state feedback control for collision avoidance," *IEEE Transactions on Intelligent Transportation Systems*, no. 99, pp. 1–16, 2018.

[20] Z. Shiller and S. Sundar, "Emergency lane-change maneuvers of autonomous vehicles," *Journal of dynamic systems, measurement, and control*, vol. 120, no. 1, pp. 37–44, 1998.

[21] M. Klomp, M. Lidberg, and T. J. Gordon, "On optimal recovery from terminal understeer," *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, vol. 228, no. 4, pp. 412–425, 2014.

[22] J. Ziegler, P. Bender, T. Dang, and C. Stiller, "Trajectory planning for Bertha - a local, continuous method," in *Intelligent Vehicles Symposium Proceedings, 2014 IEEE*. IEEE, 2014, pp. 450–457.

[23] P. Falcone, F. Borrelli, H. E. Tseng, J. Asgari, and D. Hrovat, "A hierarchical model predictive control framework for autonomous ground vehicles," in *American Control Conference, 2008*. IEEE, 2008, pp. 3719–3724.

[24] F. Altché, P. Polack, and A. de La Fortelle, "High-speed trajectory planning for autonomous vehicles using a simple dynamic model," in *Intelligent Transportation Systems (ITSC), 2017 IEEE 20th International Conference on*. IEEE, 2017, pp. 1–7.

[25] N. R. Kapania and J. C. Gerdes, "Design of a feedback-feedforward steering controller for accurate path tracking and stability at the limits of handling," *Vehicle System Dynamics*, vol. 53, no. 12, pp. 1687–1704, 2015.

[26] K. Kritayakirana and J. C. Gerdes, "Autonomous vehicle control at the limits of handling," *International Journal of Vehicle Autonomous Systems*, vol. 10, no. 4, pp. 271–296, 2012.

[27] J. Subosits and J. C. Gerdes, "Autonomous vehicle control for emergency maneuvers: The effect of topography," in *American Control Conference (ACC), 2015*. IEEE, 2015, pp. 1405–1410.

[28] D. Brayshaw and M. Harrison, "A quasi steady state approach to race car lap simulation in order to understand the effects of racing line and centre of gravity location," *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, vol. 219, no. 6, pp. 725–739, 2005.

[29] B. Siegler, A. Deakin, and D. Crolla, "Lap time simulation: Comparison of steady state, quasi-static and transient racing car cornering strategies," SAE Technical Paper, Tech. Rep., 2000.

[30] D. Casanova, R. Sharp, and P. Symonds, "Minimum time manoeuvring: The significance of yaw inertia," *Vehicle system dynamics*, vol. 34, no. 2, pp. 77–115, 2000.

[31] P. Haney, *The racing & high-performance tire: using the tires to tune for grip and balance*. McFarland, 2003, vol. 351.

[32] J. Liu, P. Jayakumar, J. L. Overholt, J. L. Stein, and T. Ersal, "The role of model fidelity in model predictive control based hazard avoidance in unmanned ground vehicles using lidar sensors," in *ASME 2013 Dynamic Systems and Control Conference*. American Society of Mechanical Engineers, 2013, pp. V003T46A005–V003T46A005.

[33] A. Domahidi and J. Jerez, "FORCES Professional," embotech GmbH (http://embotech.com/FORCES-Pro), Jul. 2014.

John K. Subosits received the B.S.E. degree in mechanical and aerospace engineering from Princeton University, Princeton, NJ, USA in 2013 and the M.S. degree in mechanical engineering from Stanford University, Stanford, CA, USA in 2015. He is currently pursuing the Ph.D. degree in mechanical engineering with Stanford University, Stanford, CA, USA. His current research interests include optimal control and trajectory planning for vehicles operating at the limits of handling.

J. Christian Gerdes received the Ph.D. degree from the University of California at Berkeley, Berkeley, CA, USA, in 1996.He is currently a Professor in mechanical engineering with Stanford University, Stanford, CA, USA, and the Director of the Center for Automotive Research at Stanford, Stanford University. He is a Co-Founder of Peloton Technology, Mountain View, CA, USA. His laboratory studies how cars move, how humans drive cars, and how to design future cars that work cooperatively with the driver or drive themselves. When not teaching on campus, he can often be found at the racetrack with students, instrumenting historic race cars, or trying out their latest prototypes for the future.

Prof. Gerdes and his team have been recognized with several awards, including the Presidential Early Career Award for Scientists and Engineers, the Ralph Teetor Award from SAE International, and the Rudolf Kalman Award from the American Society of Mechanical Engineers.